METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

CROSS-REFERENCE TO RELATED APPLICATIONS

[01] This application is a continuation-in-part of and claims priority to U.S. Patent Application Serial No. 13/657,010, filed October 22, 2012, and entitled "METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK," the disclosure of which is incorporated by reference herein in its entirety.

BACKGROUND

[02] The TCP/IP network protocols (*e.g.*, the Transmission Control Protocol (TCP) and the Internet Protocol (IP)) were designed to build large, resilient, reliable, and robust networks. Such protocols, however, were not originally designed with security in mind. Subsequent developments have extended such protocols to provide for secure communication between peers (*e.g.*, Internet Protocol Security (IPsec)), but the networks themselves remain vulnerable to attack (*e.g.*, Distributed Denial of Service (DDoS) attacks).

[03] Most existing approaches to protecting such networks are reactive rather than proactive. While reactive approaches may identify the source of an attack and assist in subsequent mitigation efforts, in most instances, the attack will have already been successfully launched.

[04] Proactive solutions, however, have often been deemed untenable due to an inability to scale to larger networks. A significant challenge associated with building a scalable proactive solution is the need to filter substantially all network traffic at a high resolution. In a large network, where traffic volumes may be enormous, the time required to provide high resolution filtering has traditionally been thought to render a proactive solution infeasible.

SUMMARY

[05] The following presents a simplified summary in order to provide a basic understanding of some aspects of the disclosure. It is neither intended to identify key or critical elements of the disclosure nor to delineate the scope of the disclosure. The following summary merely presents some concepts in a simplified form as a prelude to the description below.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1

[06] Aspects of this disclosure relate to protecting a secured network. In some embodiments, one or more packet security gateways are associated with a security policy management server. At each of the packet security gateways, a dynamic security policy may be received from the security policy management server, packets associated with a network protected by the packet security gateway may be received, and at least one of multiple packet transformation functions specified by the dynamic security policy may be performed on the packets. Performing the at least one of multiple packet transformation functions specified by the dynamic security policy on the packets may include performing at least one packet transformation function other than forwarding or dropping the packets.

[07] In some embodiments, two or more of the packet security gateways may be configured in series such that packets forwarded from a first of the packet security gateways are received by a second of the packet security gateways. In some embodiments, the dynamic security policy may include two rules requiring sequential execution. A first of the packet security gateways may perform a packet transformation function specified by one of the rules on the packets and a second of the packet security gateways may subsequently perform a packet transformation function specified by the other of the rules on packets received from the first packet security gateway.

[08] In some embodiments, the dynamic security policy may include a rule specifying a set of network addresses for which associated packets should be dropped and a rule specifying that all packets associated with network addresses outside the set should be forwarded. Additionally or alternatively, the dynamic security policy may include a rule specifying a set of network addresses for which associated packets should be forwarded and a rule specifying that all packets associated with network addresses outside the set should be dropped. In some embodiments, the security policy management server may receive information associated with one or more Voice over Internet Protocol (VoIP) sessions and the set of network addresses for which associated packets should be forwarded may be created or altered utilizing the information associated with the one or more VoIP sessions.

[09] In some embodiments, the packet security gateways may receive three or more dynamic security policies from the security policy management server. A first of the dynamic

security policies may specify a first set of network addresses for which packets should be forwarded. A second of the dynamic security policies may be received after the first and may specify a second set of network addresses, which includes more network addresses than the first set, for which packets should be forwarded. A third of the dynamic security policies may be received after the second and may specify a third set of network addresses, which includes more network addresses than the second set, for which packets should be forwarded.

[10] In some embodiments, the dynamic security policy may include two rules that each specify a set of network addresses. The dynamic security policy may specify that packets associated with the first set of network addresses should be placed in a first forwarding queue and packets associated with the second set of network addresses should be placed in a second forwarding queue. The first forwarding queue may have a different queueing policy, for example, a higher forwarding rate, than the second forwarding queue.

[11] In some embodiments, the dynamic security policy may include a rule specifying a set of network addresses and an additional parameter. The packet transformation function specified by the dynamic security policy may include routing packets that fall within the specified set and match the additional parameter to a network address different from a destination network address specified by the packets. In some embodiments, the additional parameter may be a Session Initiation Protocol (SIP) Uniform Resource Identifier (URI). The network address different from the destination network address may correspond to a device configured to copy information contained within the packets and forward the packets to the destination network address specified by the packets.

[12] In some embodiments, the packet transformation function may forward the packets into the network protected by the packet security gateway. In some embodiments, the packet transformation function may forward the packets out of the network protected by the packet security gateway. In some embodiments, the packet transformation function may forward the one or more packets to an IPsec stack having an IPsec security association corresponding to the packets. In some embodiments, the packet transformation function may drop the packets.

[13] In some embodiments, the dynamic security policy may include multiple rules. One of the rules may specify the packet transformation function. In some embodiments, one of the rules may specify a five-tuple of values selected from packet header information. The five-tuple may specify one or more protocol types, one or more IP source addresses, one or more source ports, one or more IP destination addresses, and one or more destination ports. In some embodiments, one of the rules may specify a Differentiated Service Code Point (DSCP) that maps to a DSCP field in an IP header of one of the packets.

[14] In some embodiments, one of the packet security gateways may operate in a network layer transparent manner. For example, the packet security gateway may send and receive traffic at a link layer using an interface that is not addressed at the network layer and simultaneously perform the packet transformation function at the network layer. Additionally or alternatively, the packet security gateway may include a management interface having a network layer address. Access to the management interface may be secured at the application level.

[15] In some embodiments, the dynamic security policy may include a rule generated based, at least in part, on a list of known network addresses associated with malicious network traffic. In some embodiments, the list of known network addresses associated with malicious network traffic may be received from a subscription service that aggregates information associated with malicious network traffic.

[16] In some embodiments, the packets associated with the network protected by the packet security gateway may originate within the network protected by the packet security gateway and may be destined for a network distinct from the network protected by the packet security gateway. Additionally or alternatively, the packets associated with the network protected by the packet security gateway may originate within a network distinct from the network protected by the packet security gateway and may be destined for a host within the network protected by the packet security gateway.

[17] In some embodiments, one of the packet security gateways may be located at each boundary between a protected network associated with the security policy management server and an unprotected network.

[18] Other details and features will be described in the sections that follow.

BRIEF DESCRIPTION OF THE DRAWINGS

[19] The present disclosure is pointed out with particularity in the appended claims. Features of the disclosure will become more apparent upon a review of this disclosure in its entirety, including the drawing figures provided herewith.

[20] Some features herein are illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings, in which like reference numerals refer to similar elements.

[21] FIG. 1 illustrates an exemplary network environment in which one or more aspects of the disclosure may be implemented.

[22] FIG. 2 illustrates an exemplary packet security gateway.

[23] FIG. 3 illustrates an exemplary dynamic security policy.

[24] FIG. 4 illustrates an exemplary configuration of multiple packet security gateways in series.

[25] FIG. 5 illustrates an exemplary security policy management server.

[26] FIG. 6 illustrates an exemplary network environment for implementing a monitoring service.

[27] FIG. 7 illustrates an exemplary network environment that includes a secured network having multiple boundaries with unsecured networks.

[28] FIG. 8 illustrates an exemplary network environment that includes multiple distinct secured networks.

[29] FIG. 9 illustrates an exemplary secure LAN environment.

[30] FIG. 10 illustrates an exemplary method for protecting a secured network.

[31] FIG. 11 illustrates an exemplary method for protecting a secured network that includes identifying packets based on application-layer packet-header information.

[32] FIG. 12 illustrates an exemplary method for protecting a secured network that includes performing a packet digest logging function on packets that correspond to specified packet-identification criteria.

[33] FIG. 13 illustrates an exemplary method for protecting a secured network that includes receiving security updates from multiple sources.

DETAILED DESCRIPTION

[34] In the following description of various illustrative embodiments, reference is made to the accompanying drawings, which form a part hereof, and in which is shown, by way of illustration, various embodiments in which aspects of the disclosure may be practiced. It is to be understood that other embodiments may be utilized, and structural and functional modifications may be made, without departing from the scope of the present disclosure.

[35] Various connections between elements are discussed in the following description. These connections are general and, unless specified otherwise, may be direct or indirect, wired or wireless. In this respect, the specification is not intended to be limiting.

[36] FIG. 1 illustrates an exemplary network environment in which one or more aspects of the disclosure may be implemented. Referring to FIG. 1, network environment **100** may include networks A – E **102**, **104**, **106**, **108**, and **110**. One or more networks within network environment **100** may be a Local Area Network (LAN) or a Wide Area Network (WAN). Such a LAN or WAN may be associated, for example, with an organization (*e.g.*, a company, university, enterprise, or government agency). For example, networks A – D **102**, **104**, **106**, and **108** may be LANs, any combination of which may be associated with one or more organizations. One or more networks within network environment **100** may interface with one or more other networks within network environment **100**. For example, network environment **100** may include a WAN that interfaces one or more LANs within network environment **100** or network environment **100** may include one or more Internet Service Providers (ISPs) that interface one or more LANs or WANs within network environment

100 via the Internet. For example, network E **110** may comprise the Internet and may interface networks A – D **102**, **104**, **106**, and **108**.

[37] As used herein, a packet security gateway includes any computing device configured to receive packets and perform a packet transformation function on the packets. Optionally, a packet security gateway may further be configured to perform one or more additional functions as described herein. As used herein, a security policy management server includes any computing device configured to communicate a dynamic security policy to a packet security gateway. Optionally, a security policy management server may further be configured to perform one or more additional functions as described herein. As used herein, a dynamic security policy includes any rule, message, instruction, file, data structure, or the like that specifies criteria corresponding to one or more packets and identifies a packet transformation function to be performed on packets corresponding to the specified criteria. Optionally, a dynamic security policy may further specify one or more additional parameters as described herein.

[38] Network environment **100** may include one or more packet security gateways and one or more security policy management servers. For example, network environment **100** may include packet security gateways **112**, **114**, **116**, and **118**, and security policy management server **120**. One or more security policy management servers may be associated with a protected network. For example, networks A – D **102**, **104**, **106**, and **108** may each be distinct LANs associated with a common organization and may each form part of a protected network associated with security policy management server **120**. Many network protocols route packets dynamically, and thus the path a given packet may take cannot be readily predicted. Accordingly it may be advantageous to locate a packet security gateway at each boundary between a protected network and an unprotected network. For example, packet security gateway **112** may be located at the boundary between network A **102** and network E **110**. Similarly, packet security gateway **114** may be located at the boundary between network B **104** and network E **110**; packet security gateway **116** may be located at the boundary between network C **106** and network E **110**; and packet security gateway **118** may be located at the boundary between network D **108** and network E **110**. As will be described in greater detail below, each of one or more packet security gateways associated

with a security policy management server may be configured to receive a dynamic security policy from the security policy management server, receive packets associated with a network protected by the packet security gateway, and perform a packet transformation function specified by the dynamic security policy on the packets. For example, each of packet security gateways **112**, **114**, **116**, and **118** may be configured to receive a dynamic security policy from security policy management server **120**. Each of packet security gateways **112**, **114**, **116**, and **118** may also be configured to receive packets respectively associated with networks A – D **102**, **104**, **106**, and **108**. Each of packet security gateways **112**, **114**, **116**, and **118** may further be configured to perform a packet transformation function specified by the dynamic security policy received from security policy management server **120** on the packets respectively associated with networks A – D **102**, **104**, **106**, and **108**.

[39] FIG. 2 illustrates an exemplary packet security gateway according to one or more aspects of the disclosure. Referring to FIG. 2, as indicated above, packet security gateway **112** may be located at network boundary **200** between network A **102** and network E **110**. Packet security gateway **112** may include processor **202**, memory **204**, network interfaces **206** and **208**, packet filter **214**, and management interface **222**. Processor **202**, memory **204**, network interfaces **206** and **208**, packet filter **214**, and management interface **222** may be interconnected via data bus **210**. Network interface **206** may connect packet security gateway **112** to network E **110**. Similarly, network interface **208** may connect packet security gateway **112** to network A **102**. Memory **204** may include one or more program modules that when executed by processor **202**, configure packet security gateway **112** to perform various functions as described herein.

[40] Packet security gateway **112** may be configured to receive a dynamic security policy from security policy management server **120**. For example, packet security gateway **112** may receive dynamic security policy **212** from security policy management server **120** via management interface **222** (*i.e.*, out-of-band signaling) or network interface **206** (*i.e.*, in-band signaling). Packet security gateway **112** may include one or more packet filters or packet discriminators, or logic for implementing one or more packet filters or packet discriminators. For example, packet security gateway **112** may include packet filter **214**,

which may be configured to examine information associated with packets received by packet security gateway **112** and forward the packets to one or more packet transformation functions based on the examined information. For example, packet filter **214** may examine information associated with packets received by packet security gateway **112** (*e.g.*, packets received from network E **110** via management interface **222** or network interface **206**) and forward the packets to one or more of packet transformation functions 1 – N **216**, **218**, and **220** based on the examined information.

[41] As will be described in greater detail below, dynamic security policy **212** may include one or more rules and the configuration of packet filter **214** may be based on one or more of the rules included in dynamic security policy **212**. For example, dynamic security policy **212** may include one or more rules specifying that packets having specified information should be forwarded to packet transformation function **216**, while all other packets should be forwarded to packet transformation function **218**. Packet transformation functions 1 – N **216**, **218**, and **220** may be configured to perform one or more functions on packets they receive from packet filter **214**. For example, packet transformation functions 1 – N **216**, **218**, and **220** may be configured to forward packets received from packet filter **214** into network A **102**, forward packets received from packet filter **214** to an IPsec stack having an IPsec security association corresponding to the packets, or drop packets received from packet filter **214**. In some embodiments, one or more of packet transformation functions 1 – N **216**, **218**, and **220** may be configured to drop packets by sending the packets to a local "infinite sink" (*e.g.*, the /dev/null device file in a UNIX/LINUX system).

[42] In some embodiments, packet security gateway **112** may be configured in a network layer transparent manner. For example, packet security gateway **112** may be configured to utilize one or more of network interfaces **206** and **208** to send and receive traffic at the link layer. One or more of network interfaces **206** and **208**, however, may not be addressed at the network layer. Because packet filter **214** and packet transformation functions 1 – N **216**, **218**, and **220** operate at the network layer, PSG **112** may still perform packet transformation functions at the network layer. By operating in a network layer transparent manner, packet security gateway **112** may insulate itself from network attacks (*e.g.*, DDoS attacks) launched at the network layer because attack packets cannot be routed to the network interfaces **206**

and **208**. In some embodiments, packet security gateway **112** may include management interface **222**. Management interface **222** may be addressed at the network level in order to provide packet security gateway **112** with network level addressability. Access to management interface **222** may be secured, for example, at the application level by using a service such as SSH, or secured at the transport level using, *e.g.*, TLS, or secured at the network level by attaching it to a network with a separate address space and routing policy from network A **102** and network E **110**, or secured at the link level, *e.g.*, using the IEEE 802.1X framework, etc.

[43] The flows illustrated by Figure 2 are merely exemplary and show packets that originate within a network distinct from network A **102** and are destined for a host within network A **102** in order to simplify the illustration. Packet security gateway **112** may be configured to receive and filter packets that originate within a network other than network A **102** (*e.g.*, networks B – E **104**, **106**, **108**, or **110**) and are destined for a host within network A **102**, as well as packets that originate within network A **102** destined for a network distinct from network A **102** (*e.g.*, network B – D **104**, **106**, **108**, or **110**). That is, packet security gateway **112** may be configured to filter and perform one or more packet transformation functions on packets flowing in either direction and may thus be utilized, for example, to both protect network A **102** from malicious network traffic and to prevent malicious network traffic from leaving network A **102**.

[44] FIG. 3 illustrates an exemplary dynamic security policy in accordance with one or more embodiments. Referring to FIG. 3, dynamic security policy **300** may include rules 1 – 5 **302**, **304**, **306**, **308**, and **310**. Each rule may specify criteria and one or more packet transformation functions that should be performed for packets associated with the specified criteria. The specified criteria may take the form of a five-tuple of values selected from packet header information, specifying a protocol type of the data section of the IP packet (*e.g.*, TCP, UDP, ICMP, or any other protocol), one or more source IP addresses, one or more source port values, one or more destination IP addresses, and one or more destination ports. For example, rule 1 **302** may specify that IP packets containing TCP packets, originating from a source IP address that begins with 140, having any source port, destined for an IP address that begins with 130, and destined for port 20 should have an accept packet

transformation function (*e.g.*, the identity function) performed on them. Similarly, rule 2 **304** may specify that IP packets containing TCP packets, originating from a source IP address that begins with 140, having any source port, destined for any IP address, and destined for port 80 should have an accept packet transformation function performed on them; rule 3 **306** may specify that IP packets containing TCP packets, originating from a source IP address that begins with 150, having any source port, destined for any IP address that begins with 120, and destined for port 90 should have an accept packet transformation function performed on them; rule 4 **308** may specify that IP packets containing UDP packets, originating from a source IP address that begins with 150, having any source port, destined for any IP address, and destined for port 3030 should have an accept packet transformation function performed on them; and rule 5 **310** may specify that IP packets containing any data, originating from any source IP address, having any source port, destined for any IP address, and destined for any port should have a deny packet transformation function performed on them. One or more rules included in dynamic security policy **300** may be specified in IP version 4 or IP version 6. In some embodiments, the specified criteria may comprise values selected from application-layer packet-header information. For example, application-layer hypertext transfer protocol (HTTP) packets, which transfer information between instances of web client applications (*e.g.*, web browsers) and web server applications (*e.g.*, named web sites such as www.examplewebsite.com), may contain headers that specify methods and/or one or more uniform resource identifiers (URIs) for transferring resources (*e.g.*, files) between client and servers. For example, the HTTP GET method may transfer URI-labeled resources (*e.g.*, hypertext markup language (HTML) files) from a server to a client. Similarly, the HTTP PUT method may transfer resources from a client to a server. A rule may specify that IP packets containing HTTP packets with a GET method and transferring a resource with a particular URI should have an accept packet transformation function performed on them. Another rule may specify that IP packets containing HTTP packets with a PUT method and transferring a resource to a server with a particular URI should have a deny packet transformation function performed on them. Rules may specify criteria comprising values selected from five-tuple header information and/or values selected from application-layer header information.

[45] As will be described in greater detail below, dynamic security policy **300** may include one or more rules that specify a packet transformation function other than forwarding (accepting or allowing) or dropping (denying) a packet. For example, rule 3 **306** may specify that IP packets containing TCP packets, originating from a source IP address that begins with 150, having any source port, destined for any IP address that begins with 120, and destined for port 90 should not only have an accept packet transformation function performed on them, but should also be routed to a monitoring device.

[46] One or more rules within dynamic security policy **300** may be required to execute in a specific order. For example, it may be required that rule 5 **310** be executed last. Because rule 5 **310** specifies that any packet should have a deny packet transformation function performed on it, if it were executed before a rule specifying an accept packet transformation function (*e.g.*, one or more of rules 1 – 4 **302**, **304**, **306**, or **308**), no packets matching the criteria specified by the rule specifying the accept packet transformation function would pass through a packet security gateway implementing dynamic security policy **300**. Similarly, two or more rules within dynamic security policy **300** may specify overlapping criteria and different packet transformation functions. In such cases, the order-of-application of the rules may determine which rule is applied to a packet that would match the two or more rules. Such rules may be merged together or otherwise transformed into a different set of rules without overlapping criteria, which may produce the same result as the original set of rules, when applied to any packet.

[47] A dynamic security policy may utilize the combination of one or more rules to create policies for governing packets within a network environment or effectuating one or more services within a network environment. For example, a dynamic security policy may include one or more rules, the combination of which may effectuate a blocklist service within a network environment. A dynamic security policy that effectuates a blocklist service within a network environment may include one or more rules specifying criteria (*e.g.*, a set of network addresses) for which associated packets should be blocked, dropped, or denied, and at least one rule specifying that all packets outside the specified block sets should be forwarded, accepted, or allowed. Such a dynamic security policy may be constructed by including one or more rules specifying criteria (*e.g.*, a set of network

addresses) for which associated packets should be dropped, and a wildcard rule, designated to be executed last, and specifying that all packets should be allowed. One or more dynamic security policies that effectuate a blocklist service may be utilized to implement one or more Virtual Private Networks (VPNs).

[48] A dynamic security policy may also include one or more rules, the combination of which may effectuate an allowlist service within a network environment. A dynamic security policy that effectuates an allowlist service within a network environment may include one or more rules specifying criteria (*e.g.*, a set of network addresses) for which associated packets should be forwarded, allowed, or accepted, and at least one rule specifying that all packets outside the specified allow sets should be blocked, denied, or dropped. Such a dynamic security policy may be constructed by including one or more rules specifying criteria (*e.g.*, a set of network addresses) for which associated packets should be forwarded, and a wildcard rule, designated to be executed last, and specifying that all packets should be blocked. For example, dynamic security policy **300** includes rules 1 – 4 **302**, **304**, **306**, and **308**, each of which specifies a set of network addresses for which packets should be allowed, and rule 5 **310** which specifies that all packets should be dropped. Thus, if rules 1 – 5 **302**, **304**, **306**, **308**, and **310** are executed in order, dynamic security policy **300** will effectuate an allowlist service.

[49] A dynamic security policy may also include one or more rules, the combination of which may effectuate a VoIP firewall service within a network environment. As will be discussed in greater detail below, a security policy management server may receive information associated with VoIP sessions. For example, a security policy management server may receive information associated with VoIP sessions from one or more softswitches (*e.g.*, H.323 softswitches, SIP IP Multimedia Subsystem (IMS) softswitches) or session border controllers when a VoIP session is initialized or set up. In order to allow packets associated with such a VoIP session within a network protected by one or more packet security gateways associated with the security policy management server, the security policy management server may utilize the received information associated with the VoIP sessions to construct one or more rules for allowing the packets associated with the VoIP session. When the VoIP session is terminated or torn down, the softswitch or session border

controller may notify the security policy management server, which may create or alter one or more rules to reflect the termination of the VoIP session (*e.g.*, to deny future packets which may match criteria previously associated with the VoIP session).

[50] A dynamic security policy may also include one or more rules or rule sets, the combination of which may effectuate a phased restoration service within a network environment. Such a phased restoration service may be used in the event of a network attack (*e.g.*, a DDoS attack). When an attack occurs a network may be overwhelmed with network traffic and be unable to route all or any of the traffic. In the event of such an attack, it may be beneficial to utilize a dynamic security policy which effectuates a phased restoration service. Such a dynamic security policy may include one or more rules or rule sets configured for execution in time-shifted phases. Each of the rules or rule sets may specify progressively larger sets of network addresses. For example, a dynamic security policy may include three rules or rule sets which may be configured for execution in time-shifted phases. A first of the rules or rule sets may specify a relatively small set of network addresses for which packets should be forwarded (*e.g.*, network addresses corresponding to mission critical network devices). A second of the rules or rule sets may specify a relatively larger set of network addresses for which packets should be forwarded (*e.g.*, network addresses corresponding to trusted network devices). A third of the rules or rule sets may specify an even larger set of network addresses for which packets should be forwarded (*e.g.*, network addresses corresponding to all network devices that would be allowed under ordinary circumstances). The dynamic security policy may specify that the rules or rule sets should be implemented in time-shifted phases. That is, the dynamic security policy may specify that the first rule or rule set should be executed first, and that the second rule or rule set should be executed at a time after the time at which the first rule or rule set is executed, and the third rule or rule set should be executed at a time after the time at which the second rule or rule set is executed. Such a dynamic security policy may assist a network in recovering from an attack, by allowing the network to isolate itself from the attack or recover in a controlled manner.

[51] A dynamic security policy may also include one or more rules, the combination of which may effectuate an enqueueing service within a network environment. A dynamic security policy that effectuates an enqueueing service may include one or more rules that specify sets

of network addresses and packet transformation functions that queue packets in one or more queues corresponding to the sets. These queues may then be serviced at varying rates. For example, a dynamic security policy may include two rules, each of which specify a set of network addresses. A first of the rules may specify that packets corresponding to its specified set should be queued in a first forwarding queue. A second of the rules may specify that packets corresponding to its specified set should be queued in a second forwarding queue. The first forwarding queue may be serviced at a higher forwarding rate than the second forwarding queue. Such an enqueueing service may be utilized during or following a network attack, or generally to provide prioritized service to critical network devices (*e.g.*, when network resources are strained). In some embodiments, one or more rules contained within a dynamic security policy may include an arbitrary selector which may correspond to one or more parameters or fields associated with a packet. For example, a dynamic security policy rule may include a Differentiated Service Code Point (DSCP) selector that corresponds to a DSCP field in an IP header. Thus, two packets having different values within the specified DSCP field may correspond to two distinct rules within a dynamic security policy and have different packet transformation functions performed on them. For example, two otherwise identical packets having different values within the specified DSCP field may be queued in two different forwarding queues that have different forwarding rates, and may thus receive differentiated service.

[52] A dynamic security policy may also include one or more rules, the combination of which may effectuate a multi-dimensional routing service or a multi-dimensional switching service within a network environment. For example, in some embodiments, a dynamic security policy may include one or more rules that specify a set of network addresses and an additional parameter. Such rules may further specify a packet transformation function configured to route packets within the specified set of network addresses that match the additional parameter to a network address distinct from the packets' respective destination network addresses. For example, the packet transformation function may be configured to encapsulate such packets (*e.g.*, as described by Internet Engineering Task Force (IETF) Request For Comment (RFC) 2003) with an IP header specifying a network address different from their respective destination addresses. The packets may then be routed to the network address specified by the encapsulating IP header, which may correspond to a

network device configured to utilize such packets or data contained within them, strip the IP header from the packets, and forward the packets to their respective destination addresses. In some embodiments, the packet transformation function may be configured to alter or modify the destination address of the packets, which may then be routed to the altered or modified destination address. Additionally or alternatively, the packet transformation function may be configured to assign such packets to a particular Layer-2 VLAN (*e.g.*, as described by IEEE 802.1Q). The packets may then be switched to another device on the same VLAN, which may or may not be on the IP-layer path that the packet would have taken if it were routed according to the packet's destination IP address instead of being switched through the VLAN.

[53] As will be described in greater detail below, in some embodiments a dynamic security policy may include one or more rules, the combination of which may effectuate an implementation of a multi-dimensional routing service for performing a monitoring service within a network environment. For example, a dynamic security policy may include one or more rules that specify a set of network addresses (*e.g.*, a set of network addresses from which a call that is to be monitored is expected to originate within) and an additional parameter (*e.g.*, a SIP URI corresponding to a caller to be monitored). As indicated above, such rules may further specify a packet transformation function configured to route or switch packets within the specified set of network addresses that match the additional parameter (*e.g.*, the SIP URI) to a network address corresponding to a monitoring device. The network address corresponding to the monitoring device may be different from the packets' destination network address (*e.g.*, an address corresponding to the called party or a softswitch associated with the called party). For example, the packet transformation function may be configured to encapsulate the packets with an IP header specifying the network address corresponding to the monitoring device. The packets may then be routed (or rerouted) to the monitoring device, which may be configured to copy the packets or data contained within them (*e.g.*, for subsequent review by a law enforcement or national security authority), strip the IP header from them, and then forward the packets to their destination address (*e.g.*, the address corresponding to the called party or softswitch associated with the called party).

[54] In some embodiments, a dynamic security policy may include one or more rules, the combination of which may effectuate an implementation of an informational service for performing a network communications awareness service, a network security awareness service, and/or a network threat awareness service (*e.g.*, for a particular network environment). For example, a dynamic security policy may include one or more rules that specify criteria such as one or more network addresses, protocol types, method types, and/or directions (*e.g.*, inbound, outbound, or the like) that are indicative of packet communications that are of interest to an organization that operates the secured network. Such rules may further specify a packet transformation function that, when applied to a packet that matches such a rule, produces a digest version, or log, of the packet. This packet log (or digest) may contain selected packet information and/or system information (*e.g.*, associated network addresses, ports, protocol types, URIs, arrival times, packet sizes, directions, interface names, media access control (MAC) addresses, matching rule IDs, metadata associated with matching rules, enforced policy names, or the like). The associated packet security gateway may store and/or forward packet logs using a logging system based on a standard (*e.g.*, syslog, or the like). Awareness application servers may read packet logs and perform various transformations on the logs to produce awareness information, which may be accessed by client applications.

[55] As indicated above, a significant challenge associated with building a scalable proactive solution for protecting a secured network, is the need to filter substantially all network traffic at a high resolution. Filtering traffic at a high resolution often requires the use of many rules. In a large network, where traffic volumes may be enormous, the time required to provide high resolution filtering (*e.g.*, the time required to apply a large number of rules to a large volume of traffic) has traditionally been thought to render proactive network protection solutions infeasible. This concern may be particularly acute in network environments that utilize low-latency applications (*e.g.*, VoIP).

[56] Recent advances in packet filtering technology have reduced the time required to apply large rule sets to network traffic. For example, U.S. Patent Application Publication Nos. 2006/0195896 and 2006/0248580 to Fulp et al., and U.S. Patent Application Publication No.

2011/0055916 to Ahn, describe advanced packet filtering technologies, and are each incorporated by reference herein in their entireties.

[57] One approach to providing high resolution filtering, while reducing the number of rules applied to network traffic, may be utilized when a dynamic security policy is combinatorially complete. For example, a dynamic security policy may be configured to allow bi-directional communication between a set of N internal hosts $\{I_1, I_2, \ldots, I_N\}$ within a protected network and a set of M external hosts $\{E_1, E_2, \ldots, E_M\}$ outside the protected network. To enable communications between the internal hosts and the external hosts, the dynamic security policy may be constructed to include a set of rules containing each possible combination of internal hosts and external hosts (*e.g.*,$\{\{I_1, E_1\}, \{I_1, E_2\}, \ldots \{I_1, E_M\}, \{I_2, E_1\}, \{I_2, E_2\}, \ldots \{I_2, E_M\}, \ldots, \{I_N, E_1\}, \{I_N, E_2\}, \ldots \{I_N, E_M\}\}$), each of the rules being associated with an allow packet transformation function. Such a dynamic security policy would have N*M rules for allowing communication between the internal hosts and the external hosts that originate from one of the internal hosts and are destined for one of the external hosts, and an additional N*M rules for allowing communications between the internal hosts and the external hosts that originate from one of the external hosts and are destined for one of the internal hosts. An equivalent result may be achieved, however, by constructing two smaller dynamic security policies: a first dynamic security policy that includes rules specifying the N internal hosts (*e.g.*, $\{\{I_1\}, \{I_2\}, \ldots, \{I_N\}\}$), each rule being associated with an accept packet transformation function; and a second dynamic security policy that includes rules specifying the M external hosts (*e.g.*, $\{\{E_1\}, \{E_2\}, \ldots, \{E_M\}\}$), each rule being associated with an accept packet transformation function. Such a construct of dynamic security policies may be implemented using a system of packet security gateways configured in series.

[58] FIG. 4 illustrates an exemplary configuration of multiple packet security gateways connected in series. Referring to FIG. 4, packet security gateway **112** may include one or more packet security gateways configured in series. For example, packet security gateway **112** may include packet security gateways 1 – N **400**, **402**, and **404**. Packet security gateways 1 – N **400**, **402**, and **404** may be configured so that packets forwarded by packet security gateway 1 **400** are received by packet security gateway 2 **402**, and packets

forwarded by packet security gateway 2 **402** are received by the next packet security gateway in the series, all the way through packet security gateway N **404**. Each of packet security gateways 1 – N **400**, **402**, and **404** may include a packet filter, similar to packet filter **214** described above with respect to FIG. 2, and one or more packet transformation functions, similar to packet transformation functions 1 – N **216**, **218**, and **220** described above with respect to FIG. 2. Packet security gateways 1 – N **400**, **402**, and **404** may be utilized to implement a construct of dynamic security policies similar to that described above.

[59] For example, packet security gateway 1 **400** may be configured to implement P_1, which may include rules specifying M external hosts (*e.g.*, $\{\{E_1\}, \{E_2\}, \ldots, \{E_M\}\}$), each rule being associated with an accept packet transformation function. Packet security gateway 2 **402** may be configured to implement P_2, which may include rules specifying N internal hosts (*e.g.*, $\{\{I_1\}, \{I_2\}, \ldots, \{I_N\}\}$), each rule being associated with an accept packet transformation function. A packet received by packet security gateway **112** may be initially received via packet security gateway 1 **400**'s network interface. Packet security gateway 1 **400** may apply one or more of the rules in P_1 to the received packet until the packet matches criteria specified by a rule in P_1, at which point packet security gateway 1 **400** may perform a packet transformation function specified by the rule on the packet. For example, a packet may be received by packet security gateway **112** that originates from external host E_5 (*e.g.*, a host within network E **110**) and is destined for internal host I_7 (*e.g.*, a host within network A **102**). Packet security gateway 1 **400** may apply one or more of the rules in P_1 (*e.g.*, $\{\{E_1\}, \{E_2\}, \ldots, \{E_M\}\}$) to the received packet and the received packet may match the criteria specified by one of the rules in P_1 (*e.g.*, $\{\{E_5\}$). The rule may specify that an accept packet transformation function should be performed, and packet security gateway 1 **400** may utilize one or more of its packet transformation functions to perform the accept packet transformation function on the packet and forward the packet to packet security gateway 2 **402**. Packet security gateway 2 **402** may apply one or more of the rules in P_2 (*e.g.*, $\{\{I_1\}, \{I_2\}, \ldots, \{I_N\}\}$) to the packet and the packet may match the criteria specified by one of the rules in P_2 (*e.g.*, $\{\{I_7\}$). The rule may specify that an accept packet transformation function should be performed, and packet security gateway 2 **402** may utilize one or more of its

packet transformation functions to perform the accept packet transformation function on the packet and forward the packet to network A **102**.

[60] It will be appreciated that utilizing multiple packet security gateways in series to implement dynamic security policy constructs may increase performance and decrease memory resource requirements. For example, in the described scenario packet security gateway 1 **400** may have only been required to compare the packet to five rules and packet security gateway 2 **402** may have only been required to compare the packet to seven rules. In a worst case scenario, packet security gateway 1 **400** may have only been required to compare the packet to M rules and packet security gateway 2 **402** may have only been required to compare the packet to N rules. Moreover, the series configuration may enable packet security gateway 1 **400** to begin implementing P_1 with respect to a subsequently received packet, while packet security gateway 2 **402** simultaneously implements P_2 with respect to the packet forwarded by packet security gateway 1 **400**. Furthermore, the memory requirements for this scenario with packet security gateways in series may be comparable to M+N, whereas originally the combinatorially complete set of rules contained in a single packet security gateway may have required memory comparable to N*M.

[61] FIG. 5 illustrates an exemplary security policy management server. Referring to FIG. 5, security policy management server **120** may include processor **500**, memory **502**, and network interface **504**. One or more of processor **500**, memory **502**, and network interface **504** may be interconnected via data bus **506**. Network interface **504** may interface security policy management server **120** with network E **110**. Memory **502** may include one or more program modules that when executed by processor **500**, configure security policy management server **120** to perform functions described herein. It will be appreciated that as used herein the term "server" designates one or more computing devices configured to perform one or more functions described herein. The term "server" should not be construed to imply that a client/server relationship (*e.g.*, a relationship in which a request is received from a client and then serviced by a server) necessarily exists.

[62] Security policy management server **120** may be configured to communicate one or more dynamic security policies to one or more packet security gateways within network

environment **100**. For example, security policy management server **120** may communicate one or more dynamic security policies stored in memory **502** to one or more of packet security gateways **112**, **114**, **116**, and **118**. For example, security policy management server **120** may be configured to communicate one or more dynamic security policies to one or more of packet security gateways **112**, **114**, **116**, and **118** on a periodic basis, under specified network conditions, whenever security policy management server **120** receives a new dynamic security policy, whenever a dynamic security policy stored on security policy management server **120** is changed or altered, or in response to a request from one or more of packet security gateways **112**, **114**, **116**, and **118**.

[63] Security policy management server **120** may also be configured to provide one or more administrators associated with security policy management server **120** with management interface **510**. For example, security policy management server **120** may be configured to provide one or more administrators with a Graphical User Interface (GUI) or Command Line Interface (CLI). An administrator of security policy management server **120** may utilize security policy management server **120**'s management interface **510** to configure security policy management server **120**. For example, an administrator may configure security policy management server **120** in order to associate security policy management server **120** with one or more of packet security gateways **112**, **114**, **116**, and **118**. An administrator of security policy management server **120** may also utilize security policy management server **120**'s management interface **510** to construct one or more dynamic security policies or to load one or more dynamic security policies into security policy management server **120**'s memory **502**. For example, an administrator associated with security policy management server **120** may manually construct one or more dynamic security policies offline and then utilize security policy management server **120**'s management interface **510** to load such dynamic security policies into security policy management server **120**'s memory **502**.

[64] In some embodiments, security policy management server **120** may be configured to add, remove, or alter one or more dynamic security policies stored in memory **502** based on information received from one or more devices within network environment **100**. For example, security policy management server **120**'s memory **502** may include a dynamic

security policy having one or more rules that specify a list of network addresses known to be associated with malicious network traffic. Security policy management server **120** may be configured to automatically create or alter one or more of such rules as new network addresses associated with malicious network traffic are determined. For example, security policy management server **120** may receive updates (*e.g.* as part of a subscription) from malicious host tracker service **508**. Malicious host tracker service **508** may aggregate information associated with malicious network traffic and updates received from malicious host tracker service **508** may include one or more network addresses that have been determined to be associated with malicious network traffic. Security policy management server **120** may be configured to create or alter one or more rules included within a dynamic security policy associated with malicious host tracker service **508** to block traffic associated with the network addresses received from malicious host tracker service **508**. Additionally or alternatively, as indicated above, security policy management server **120** may be configured to create or alter one or more dynamic security policies, or one or more rules included in one or more dynamic security policies, to account for VoIP sessions being initiated or terminated by a network device within network environment **100**.

[65] In some embodiments, security policy management server **120** may be configured to add, remove, and/or alter one or more dynamic security policies stored in memory **502** based on information received from two or more devices within network environment **100**. For example, security policy management server **120** may receive updates (*e.g.*, as part of a subscription) from malicious host tracker service **508** and/or one or more other services. The updates from the two or more services may be correlated (*e.g.*, by security policy management server **120**). For example, a network address received from host tracker service **508** may be a duplicate of a network address received from another (*e.g.*, different) service (*e.g.*, a range of network addresses received from host tracker service **508** may overlap with a range of network addresses received from another service). Security policy management server **120** may combine the rules associated with these correlated updates from two or more services within one or more dynamic security policies. Security policy management server **120** may be configured to reduce the size of and/or to de-correlate dynamic security policies (*e.g.*, because the performance of a packet security gateway may be dependent on the size of and/or correlations within dynamic security policies). For

example, duplicate network addresses may be removed, and/or overlapping ranges of network addresses may be combined into one or more new network address ranges.

[66] As indicated above, a dynamic security policy may include one or more rules, the combination of which may effectuate an implementation of a multi-dimensional routing service for performing a monitoring service within a network environment. FIG. 6 illustrates an exemplary network environment for implementing a monitoring service in accordance with one or more embodiments. Referring to FIG. 6, a user of network environment **100** (*e.g.*, a law enforcement or national security authority) may desire to obtain a copy of packets associated with one or more VoIP sessions (*e.g.*, sessions associated with SIP URI exampleuser@exampledomain.com) within network environment **100**. Because many SIP-signaled services are designed to address sessions dynamically, it may not be possible to determine, prior to a session being set up, a particular network address and port from which packets should be copied. Moreover, due to privacy concerns, regulators may require that only packets associated with the specified VoIP sessions (*e.g.*, sessions associated with SIP URI exampleuser@exampledomain.com) be copied.

[67] For example, a user associated with SIP URI exampleuser@exampledomain.com may utilize User Equipment (UE) **600** within network A **102** to place a VoIP call to a user utilizing UE **602** within network B **104**. SIP switch **604** may be utilized by an operator of network A **102** for switching SIP signals within network A **102**. Similarly, SIP switch **606** may be utilized by an operator of network B **104** for switching SIP signals within network B **104**. One or more of SIP switches **604** and **606** may include an analysis application configured to monitor SIP signals and publish SIP messages associated with specified users to one or more subscribers. For example, the operator of network A **102** may have installed analysis application **610** on SIP switch **604** (*e.g.*, accessed via a SIP IMS Service Control (ISC) interface associated with SIP switch **604**) and configured analysis application **610** to search for and publish SIP messages associated with SIP URI exampleuser@exampledomain.com to security policy management server **120**. Similarly, the operator of network B **104** may have installed analysis application **612** on SIP switch **606** and configured analysis application **612** to publish SIP messages associated with SIP URI exampleuser@exampledomain.com to security policy management server **120**.

[68] When the user associated with SIP URI exampleuser@exampledomain.com utilizes UE **600** to place a VoIP call to the user utilizing UE **602**, analysis application **610** may detect one or more SIP signaling messages associated with the call (*e.g.*, SIP signaling messages for setting up the call) and publish the messages to security policy management server **120**. Security policy management server **120** may extract one or more network addresses and port numbers from the SIP signaling messages (*e.g.*, a network address and port number utilized by UE **600** for placing the VoIP call to UE **602**). Security policy management server **120** may utilize the extracted network addresses and port numbers to create a new dynamic security policy or alter one or more rules within an existing dynamic security policy. For example, security policy management server **120** may construct a new dynamic security policy that includes a rule specifying one of the extracted network addresses and port numbers, as well as a packet transformation function configured to route associated packets to monitoring device **608**. Security policy management server **120** may communicate the new or modified dynamic security policy to packet security gateway **112**.

[69] When packets associated with the VoIP call between UE **600** and UE **602** are received by packet security gateway **112**, packet filter **214** may identify the packets as matching the criteria specified by the dynamic security policy received from security policy management server **120** (*e.g.*, packets addressed to or from the extracted address and port number) and may perform the packet transformation function configured to route the packets to monitoring device **608**. For example, the packet transformation function configured to route the packets to monitoring device **608** may be packet transformation function 2 **218**. When packet transformation function 2 **218** receives the packets from packet filter **214**, it may encapsulate them with an IP header having an address corresponding to monitoring device **608** and may then forward them to network E **110**. Once forwarded, the packets may be routed based on the address specified by the encapsulating header, and may thus be communicated to monitoring device **608**. When the packets are received by monitoring device **608**, monitoring device **608** may copy the packets or data contained within them, and strip the encapsulating header from them. Monitoring device **608** may then forward the packets, without the encapsulating header, to network E **110**. Network E **110** may receive the packets forwarded by monitoring device **608** and may route them based on their destination address (*e.g.*, to UE **602**).

[70] In some embodiments, packet security gateway **112** may be configured to perform multiple packet transformation functions on the packets associated with the VoIP call between UEs **600** and **602**. For example, packet filter **214** may identify the packets as matching the criteria specified by the dynamic security policy received from security policy management server **120** and may forward the packets to packet transformation functions 1 **216** and 2 **218**. Packet transformation function 1 **216** may be configured to forward the packets to their destination address (*e.g.*, to UE **602**) and packet transformation function 2 **218** may be configured to encapsulate the packets (or a copy of the packets) with an IP header having an address corresponding to monitoring device **608** and then forward the encapsulated packets to network E **110**. Once forwarded, the encapsulated packets may be routed based on the address specified by the encapsulating header, and may thus be communicated to monitoring device **608**, which may store the packets or data contained within them for subsequent review or analysis (*e.g.*, by a law enforcement or national security authority). In such embodiments, it may not be necessary for monitoring device **608** to strip the encapsulating header from the packets or route them based on their destination address (*e.g.*, to UE **602**) because packet transformation function 1 **216** may have already forwarded the packets to their destination address (*e.g.*, to UE **602**).

[71] It will be appreciated that SIP switch **604**'s analysis application **610** may similarly detect SIP signaling associated with the termination of the VoIP call between UE **600** and UE **602** and may publish the SIP messages to security policy management server **120**. Security policy management server **120** may utilize one or more network addresses and port numbers within the messages to construct a new dynamic security policy or modify one or more rules within an existing dynamic security policy and communicate the new or modified dynamic security policy to packet security gateway **112** in order to ensure that future packets associated with the network address and port number but not associated with SIP URI exampleuser@exampledomain.com are not routed to monitoring device **608**. Security policy management server **120** may communicate any dynamic security policy constructed or modified based on SIP messages to any of multiple packet security gateways (*e.g.*, packet security gateways **114** and **116**) within network environment **100** in order to ensure that all packets associated with the VoIP call between UE **600** and UE **602** are forwarded to monitoring device **608**.

[72] FIG. 7 illustrates an exemplary network environment that includes a secured network having multiple boundaries with unsecured networks in which one or more embodiments may be implemented. Network environment **700** may include networks A – C **702**, **704**, and **706**. Networks A **702** and B **704** may be a LAN or WAN associated with an organization (*e.g.*, a company, university, enterprise, or government agency). One or more networks within network environment **700** may interface with one or more other networks within network environment **700**. For example, the organizations associated with networks A **702** and B **704** may subscribe to an ISP to provide interconnectivity between their respective networks or allow public access to their respective networks (*e.g.*, via the Internet). Each of networks A **702** and B **704** may be connected to network C **706**, which may be the ISP's network. The ISP may desire to offer an interconnection service between networks A **702** and B **704**, but may also want to enforce one or more dynamic security policies with respect to traffic traversing network C **706**.Accordingly, one or more packet security gateways may be located at each boundary between network A **702** and network C **706**, and each boundary between network B **704** and network C **706**. For example, packet security gateway **708** and packet security gateway **710** may be respectively located at first and second boundaries between networks A **702** and C **706**. Similarly, packet security gateways **712** and **714** may be respectively located at first and second boundaries between networks B **704** and C **706**. Each of packet security gateways **708**, **710**, **712**, and **714** may be associated with security policy management server **716**.

[73] Security policy management server **716** may maintain one or more dynamic security policies configured for protecting network C **706**, and may be managed by the ISP associated with network C **706**. Security policy management server **716** may ensure that each of packet security gateways **708**, **710**, **712**, and **714** protect each of their respective boundaries with network C **706** in a uniform manner. For example, security policy management server **716** may be configured to communicate one or more dynamic security policies it maintains to each of packet security gateways **708**, **710**, **712**, and **714** on a periodic basis, in response to being directed to by a network operator associated with network environment **700**, in response to detected network conditions (*e.g.*, an attack or high resource utilization), or in response to a request from one or more of packet security gateways **708**, **710**, **712**, or **714**.

[74] In some embodiments, security policy management server **716** may be configured to communicate different dynamic security policies to one or more of packet security gateways **708**, **710**, **712**, and **714** based on, for example, their respective locations within network environment **700**. For example, security policy management server **716** may be configured to implement one or more anti-spoofing techniques (*e.g.*, ingress filtering or Best Current Practice (BCP) 38, as described by Internet Engineering Task Force (IETF) Request For Comment (RFC) 2827) with respect to network environment **700**. Effective implementation of such techniques may require that a dynamic security policy be based on the location at which it is being implemented. For example, a dynamic security policy that implements ingress filtering may comprise one or more rules that filter based on a packet's source address, identifying packets having source addresses that could not possibly have originated from a network downstream of the ingress filtering point (*e.g.*, packets having spoofed source addresses). Such rules may vary depending on the boundary point for which they are implemented (*e.g.*, a packet for one boundary may be properly identified as spoofed, yet a packet having the same source address may be legitimate traffic at a different boundary point). Accordingly, security policy management server **716** may be configured to communicate different dynamic security policies to one or more of packet security gateways **708**, **710**, **712**, and **714** based on their respective locations within network environment **700**. For example, security policy management server **716** may communicate a dynamic security policy to packet security gateways **708** and **710** that includes one or more rules for performing ingress filtering for network A **702** (*e.g.*, for identifying packets having source addresses that could not have originated within network A **702**) and a different dynamic security policy to packet security gateways **712** and **714** that includes one or more rules for performing ingress filtering for network B **704** (*e.g.*, for identifying packets having source addresses that could not have originated within network B **704**).

[75] It will be appreciated that by maintaining uniform dynamic security policies at each boundary between networks A **702** and C **706**, as well as at each boundary between networks B **704** and C **706**, security policy management server **716** and packet security gateways **708**, **710**, **712**, and **714** may aid the ISP associated with network C **706** in protecting network C **706** from network attacks.

[76] FIG. 8 illustrates an exemplary network environment that includes multiple distinct secured networks in which one or more embodiments may be implemented. Referring to FIG. 8, network environment **800** may include networks A **802**, B **804**, and C **806**. Each of networks A **802** and B **804** may interface with network C **806** at multiple boundaries within network environment **800**. Packet security gateways **808** and **810** may be respectively located at first and second boundaries between networks A **802** and C **806**. Similarly, packet security gateways **812** and **814** may be respectively located at first and second boundaries between networks B **804** and C **806**.

[77] Network A **802** and B **804** may both be associated with a common organization (*e.g.*, a company, university, enterprise, or government agency), or may each be associated with a distinct organization. In the former case, the common organization may desire to utilize one or more dynamic security policies with respect to network A **802** and one or more different dynamic security policies with respect to network B **804**. In the latter case, an organization associated with network A **802** may desire to utilize one or more dynamic security policies with respect to network A **802** and a different organization associated with network B **804** may desire to utilize one or more different dynamic security policies with respect to network B **804**. Network environment **800** may include security policy management servers A **816** and B **818**. Security policy management server A **816** may be associated with network A **802** and may maintain one or more dynamic security policies configured for protecting network A **802**. Similarly, security policy management server B **818** may be associated with network B **804** and may maintain one or more dynamic security policies configured for protecting network B **804**.

[78] Packet security gateways **808** and **810** may be associated with security policy management server A **816**. Similarly, packet security gateways **812** and **814** may be associated with security policy management server B **818**. Security policy management server A **816** may ensure that packet security gateways **808** and **810** protect each of their respective boundaries with network C **806** in a uniform manner. For example, security policy management server A **816** may be configured to communicate one or more dynamic security policies it maintains to packet security gateways **808** and **810** on a periodic basis, in response to being directed to by a network operator associated with network A **802**, in response to detected

network conditions (*e.g.*, an attack or high resource utilization), or in response to a request from packet security gateway **808** or **810**. Similarly, security policy management server B **818** may ensure that packet security gateways **812** and **814** protect each of their respective boundaries with network C **806** in a uniform manner. For example, security policy management server B **818** may be configured to communicate one or more dynamic security policies it maintains to packet security gateways **812** and **814** on a periodic basis, in response to being directed to by a network operator associated with network B **804**, in response to detected network conditions (*e.g.*, an attack or high resource utilization), or in response to a request from packet security gateway **812** or **814**. By utilizing distinct security policy management servers (*e.g.*, security policy management servers A **816** and B **818**), one or more operators associated with distinct networks (*e.g.*, networks A **802** and B **804**) may maintain uniform dynamic security policies at each boundary of their respective networks, while simultaneously enabling different dynamic security policies to be maintained for each network. Similarly, by utilizing distinct security policy management servers (*e.g.*, security policy management servers A **816** and B **818**), one or more operators associated with a single organization that desires to maintain distinct networks (*e.g.*, networks A **802** and B **804**) may maintain uniform dynamic security policies at each boundary of their distinct networks, while simultaneously enabling different dynamic security policies to be maintained for each network.

[79] FIG. 9 illustrates an exemplary secure LAN environment protected in accordance with one or more aspects of the disclosure. Referring to FIG. 9, network environment **900** may be a LAN, including hosts A **902**, B **904**, and C **906**. It may also include LAN switch **908**. LAN switch **908** may be configured to switch network traffic (*e.g.*, packets) between one or more of hosts A **902**, B **904**, and C **906**. For example, LAN switch **908** may include a switching matrix configured to switch packets received from one or more of hosts A **902**, B **904**, and C **906** to one or more of hosts A **902**, B **904**, and C **906**. LAN switch **908** may be associated with packet security gateway **910**, and network environment **900** may include security policy management server **912**.

[80] In some embodiments, packet security gateway **910** may be embedded within LAN switch **908**. Alternatively, packet security gateway **910** may be a device distinct from LAN switch

908, and LAN switch **908** may be configured to route network traffic through packet security gateway **910** (*e.g.*, by modifying LAN switch **908**'s switching matrix). Packet security gateway **910** may be configured to receive one or more dynamic security policies from security policy management server **912**. The dynamic security policies received from security policy management server **912** may include one or more rules specifying criteria associated with one or more of hosts A **902**, B **904**, and C **906**, and may further specify one or more packet transformation functions to be performed on packets matching the specified criteria. Packet security gateway **910** may identify packets matching one or more of the criteria specified by the rules and may perform the associated packet transformation functions on the identified packets. By utilizing packet security gateway **910** within network environment **900**, an operator of network environment **900** may be able to protect network environment **900** from network attacks, as well as implement one or more services (*e.g.*, blocklist service, allowlist service, VoIP firewall service, phased restoration service, enqueueing service, multi-dimensional routing service, or monitoring service) within network environment **900**. Network environment **900** may include multiple LAN switches with embedded or associated packet security gateways, each of the packet security gateways configured to receive one or more dynamic security policies from security policy management server **912**.

[81] FIG. 10 illustrates an exemplary method for protecting a secured network in accordance with one or more embodiments. The steps may be performed at each of one or more packet security gateways associated with a security policy management server. For example, each of packet security gateways **112**, **114**, **116**, and **118** may be associated with security policy management server **120**, and the steps may be performed at each of packet security gateways **112**, **114**, **116**, and **118**. At step **1000**, a dynamic security policy is received from the security policy management server. For example, packet security gateway **112** may receive dynamic security policy **300** from security policy management server **120**. At step **1002**, packets associated with a network protected by each respective packet security gateway are received. For example, packet security gateway **112** may receive UDP packets from a device within network E **110** having an address that begins with 150 and that are destined for port 3030 of a device within network A **102**. At step **1004**, a packet transformation function specified by the dynamic security policy is performed on the

packets. For example, rule **308** of dynamic security policy **300** may specify that packets using the UDP protocol, coming from a source address that begins with 150, having any source port, destined for any address, and destined for port 3030 should have an accept packet transformation function performed on them, packet filter **214** may identify the UDP packets received from the device within network E **110** as matching the criteria specified by rule **308**, packet transformation function 1 **216** may be configured to forward packets, and packet security gateway **112** may utilize packet transformation function 1 **216** to perform the accept packet transformation function specified by rule **308** on the UDP packets received from the device within network E **110**.

[82] FIG. 11 illustrates an exemplary method for protecting a secured network that includes identifying packets based on application-layer packet-header information. Referring to FIG. 11, at step **1102**, a plurality of packet security gateways associated with a security policy management server may receive a dynamic security policy that includes at least one rule specifying application-layer packet-header information and a packet transformation function to be performed on packets comprising the application-layer packet-header information. For example, packet security gateways **112**, **114**, **116**, and **118** may be associated with security policy management server **120** and may receive a dynamic security policy that includes one or more rules specifying application-layer packet-header information (*e.g.*, information that identifies one or more HTTP packets) and a packet transformation function to be performed on packets comprising the application-layer packet-header information (*e.g.*, an accept or deny packet transformation function) from security policy management server **120**. At step **1104**, a packet security gateway, of the plurality of packet security gateways associated with the security policy management server, may receive packets associated with a network protected by the packet security gateway. For example, packet security gateway **112** may receive packets associated with network A **102** (*e.g.*, from a computing device located in network A **102**, network B **104**, network C **106**, network D **108**, and/or network E **110**).

[83] At step **1106**, the packet security gateway may identify, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets comprising the application-layer packet-header

information. For example, in some embodiments, the rule(s) specifying the application-layer packet-header information may identify one or more HTTP packets (*e.g.*, one or more HTTP packets comprising an HTTP GET method call and/or an HTTP PUT method call), and packet security gateway **112** may identify, from amongst the packets associated with network A **112**, one or more packets comprising the application-layer packet-header information. In some embodiments, identifying the packets comprising the application-layer packet-header information may include the packet security gateway determining that the packets are among the HTTP packets identified by the rule(s) (*e.g.*, the HTTP GET method call and/or the HTTP PUT method call may specify one or more URIs, and packet security gateway **112** may determine that the packet(s) originated from and/or are destined for a network address corresponding to the URI(s)). At step **1108**, the packet security gateway may perform, on a packet-by-packet basis, the packet transformation function on the identified packet(s). For example, if the packet transformation function is an accept packet transformation function, packet security gateway **112** may forward the packet(s) toward their respective destinations. Similarly, if the packet transformation function is a deny packet transformation function, packet security gateway **112** may drop the packet(s).

[84] FIG. 12 illustrates an exemplary method for protecting a secured network that includes performing a packet digest logging function on packets that correspond to specified packet-identification criteria. Referring to FIG. 12, at step **1202**, each of a plurality of packet security gateways associated with a security policy management server may receive a dynamic security policy that includes at least one rule specifying packet-identification criteria and a packet transformation function comprising a packet digest logging function to be performed on packets corresponding to the packet-identification criteria from the security policy management server. For example, packet security gateways **112**, **114**, **116**, and **118** may be associated with security policy management server **120** and may receive a dynamic security policy that includes one or more rules specifying packet-identification criteria (*e.g.*, application-layer packet-header information and/or a five-tuple specifying one or more transport-layer protocols, a range of source addresses, a range of source ports, a range of destination addresses, and a range of destination ports) and a packet transformation function comprising a packet digest

logging function to be performed on packets corresponding to the packet-identification criteria from security policy management server **120**. At step **1204**, a packet security gateway, of the plurality of packet security gateways associated with the security policy management server, may receive packets associated with a network protected by the packet security gateway. For example, packet security gateway **112** may receive packets associated with network A **102** (*e.g.*, from a computing device located in network A **102**, network B **104**, network C **106**, network D **108**, and/or network E **110**).

[85] At step **1206**, the packet security gateway may identify from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets corresponding to the packet-identification criteria. For example, packet security gateway **112** may identify from amongst the packets associated with network A **102** one or more packets corresponding to the packet-identification criteria (*e.g.*, by determining that the packet(s) comprise the specified application-layer packet-header information and/or correspond to at least one of the specified transport-layer protocols, have a source address within the specified range of source addresses, have a source port within the specified range of source ports, have a destination address within the specified range of destination addresses, and/or have a destination port within the specified range of destination ports). At step **1208**, the packet security gateway may perform, on a packet-by-packet basis, the packet digest logging function on each of the packets corresponding to the packet-identification criteria. For example, packet security gateway **112** may perform the packet digest logging function on each of the identified packets.

[86] In some embodiments, performing the packet digest logging function may include identifying a subset of information specified by the packet digest logging function (*e.g.*, a portion of data from the packet, a source address of the packet, a source port of the packet, a destination address of the packet, a destination port of the packet, a transport-protocol type of the packet, a uniform resource identifier (URI) from the packet, an arrival time of the packet, a size of the packet, a flow direction of the packet, an identifier of an interface of the packet security gateway that received the packet, and/or one or more media access control (MAC) addresses associated with the packet) for each of the

identified packets, and generating a record comprising the subset of information for each of the identified packets. In some embodiments, packet security gateway **112** may be configured to temporarily store data comprising the subset of information for each of the identified packets, and to utilize the stored data to generate a message comprising the subset of information (or a portion thereof) for each of the identified packets. In such embodiments, packet security gateway **112** may communicate the message to a different computing device (*e.g.*, security policy management server **120**). In some embodiments, the subset of information (or a portion thereof) may be reformatted (*e.g.*, by packet security gateway **112** and/or security policy management server **120**) in accordance with a logging system standard (*e.g.*, syslog).

[87] FIG. 13 illustrates an exemplary method for protecting a secured network that includes receiving security updates from multiple sources. Referring to FIG. 13, at step **1302**, a security policy management server may receive a security update comprising a set of network addresses from a computing device. For example, security policy management server **120** may receive a security update comprising a set of network addresses *(e.g.*, {network address A, network address B, and network address C}) from malicious host tracker service **508** (or a computing device associated therewith). At step **1304**, the security policy management server may update one or more rules stored in its memory to include the set of network addresses. For example, security policy management server **120** may update one or more rules of dynamic security policy **300** (*e.g.*, a rule specifying a deny packet transformation function) to include the set of network addresses (*e.g.*, network addresses A, B, and C). At step **1306**, the security policy management server may receive a security update comprising a different set of network addresses from a different computing device. For example, security policy management server **120** may receive a security update comprising a different set of network addresses (*e.g.*, {network address B, network address C, network address D, and network address E}) from a different computing device (*e.g.*, a different computing device associated with malicious host tracker service **508** or a computing device associated with a malicious host tracker service different from malicious host tracker service **508**).

[88] At step **1308**, the security policy management server may determine that the different set of network addresses includes at least a portion of network addresses included in the set of network addresses, and, at step **1310**, may identify the at least a portion of network addresses included in the set of network addresses. For example, security policy management server **120** may determine that the different set of network address (*e.g.*, {network address B, network address C, network address D, and network address E}) includes at least a portion of network addresses (*e.g.*, network addresses B and C) included in the set of network addresses (*e.g.*, {network address A, network address B, and network address C}), and may identify the at least a portion of network addresses included in the set of network addresses (*e.g.*, network addresses B and C). At step **1312**, the security policy management server may identify at least one of the rules stored in its memory that specifies a range of network addresses comprising the at least a portion of network addresses. For example, security policy management server **120** may identify the rule(s) of dynamic security policy **300** updated in step **1304** (*e.g.*, the rule(s) updated to include network addresses A, B, and C). At step **1314**, the security policy management server may update the identified rule(s) to include one or more other network addresses included in the different set of network addresses. For example, security policy management server **120** may update the identified rules of dynamic security policy **300** to include one or more other network addresses included in the different set of network addresses (*e.g.*, network addresses D and E).

[89] In some embodiments, the security policy management server may identify two or more rules that each specify a range of network addresses comprising the at least a portion of network addresses. For example, security policy management server **120** may identify two rules of dynamic security policy **300** that each include the at least a portion of network addresses (*e.g.*, a rule that includes network addresses A, B, C, and F, and a rule that includes network addresses A, B, C, and G). In such embodiments, the security policy management server may combine the two or more rules into a rule that specifies a range of network addresses that includes network addresses specified by each of the two or more rules and the one or more other network addresses included in the different set of network addresses. For example, security policy management server **120** may combine the two rules of dynamic security policy **300** that each include the at least a portion of network addresses

(*e.g.*, the rule that includes network addresses A, B, C, and F, and the rule that includes network addresses A, B, C, and G) into a rule that specifies a range of network addresses that includes network addresses specified by each of the two identified rules and the one or more other network addresses included in the different set of network addresses (*e.g.*, a range that includes network addresses A, B, C, D, E, F, and G).

[90] The functions and steps described herein may be embodied in computer-usable data or computer-executable instructions, such as in one or more program modules, executed by one or more computers or other devices to perform one or more functions described herein. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types when executed by one or more processors in a computer or other data processing device. The computer-executable instructions may be stored on a computer-readable medium such as a hard disk, optical disk, removable storage media, solid state memory, RAM, etc. As will be appreciated, the functionality of the program modules may be combined or distributed as desired in various embodiments. In addition, the functionality may be embodied in whole or in part in firmware or hardware equivalents, such as integrated circuits, application-specific integrated circuits (ASICs), field programmable gate arrays (FPGA), and the like. Particular data structures may be used to more effectively implement one or more aspects of the disclosure, and such data structures are contemplated to be within the scope of computer executable instructions and computer-usable data described herein.

[91] Although not required, one of ordinary skill in the art will appreciate that various aspects described herein may be embodied as a method, an apparatus, or as one or more computer-readable media storing computer-executable instructions. Accordingly, those aspects may take the form of an entirely hardware embodiment, an entirely software embodiment, an entirely firmware embodiment, or an embodiment combining software, hardware, and firmware aspects in any combination.

[92] As described herein, the various methods and acts may be operative across one or more computing servers and one or more networks. The functionality may be distributed in

any manner, or may be located in a single computing device (*e.g.*, a server, a client computer, etc.).

[93] Aspects of the disclosure have been described in terms of illustrative embodiments thereof. Numerous other embodiments, modifications, and variations within the scope and spirit of the appended claims will occur to persons of ordinary skill in the art from a review of this disclosure. For example, one of ordinary skill in the art will appreciate that the steps illustrated in the illustrative figures may be performed in other than the recited order, and that one or more steps illustrated may be optional.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 37

What is claimed is:

1. A method, comprising:

receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy that includes at least one rule specifying application-layer packet-header information and a packet transformation function to be performed on packets comprising the application-layer packet-header information;

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway;

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets comprising the application-layer packet-header information; and

performing, by the packet security gateway and on a packet-by-packet basis, the packet transformation function on each of the one or more packets comprising the application-layer packet-header information.

2. The method of claim 1, wherein the at least one rule indicates that packets having the application-layer packet-header information should have an accept packet transformation function performed on them, and wherein performing the packet transformation function on each of the one or more packets comprising the application-layer packet-header information comprises the packet security gateway forwarding each of the one or more packets comprising the application-layer packet-header information toward its respective destination.

3. The method of claim 1, wherein the at least one rule indicates that packets having the application-layer packet-header information should have a deny packet transformation function performed on them, and wherein performing the packet transformation function on each of the one or more packets comprising the application-layer packet-header information comprises the packet security gateway dropping each of the one or more packets comprising the application-layer packet-header information.

4. The method of claim 1, wherein the application-layer packet-header information identifies one or more hypertext transfer protocol (HTTP) packets, and wherein identifying the one or more packets comprising the application-layer packet-header information comprises determining by the packet security gateway that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets.

5. The method of claim 4, wherein the one or more HTTP packets comprise an HTTP GET method call, and wherein determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets comprises determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP GET method call.

6. The method of claim 5, wherein the HTTP GET method call specifies a uniform resource identifier (URI), and wherein determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP GET method call comprises determining that the one or more packets comprising the application-layer packet-header information originated from or are destined for a network address corresponding to the URI.

7. The method of claim 4, wherein the one or more HTTP packets comprise an HTTP PUT method call, and wherein determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets comprises determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP PUT method call.

8. The method of claim 7, wherein the HTTP PUT method call specifies a uniform resource identifier (URI), and wherein determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP PUT method call comprises determining that the one or more packets comprising the application-layer packet-header information originated from or are destined for a network address corresponding to the URI.

9. The method of claim 1, wherein the at least one rule comprises a five-tuple specifying one or more transport-layer protocols, a range of source addresses, a range of source ports, a range of destination addresses, and a range of destination ports, and wherein identifying the one or more packets comprising the application-layer packet-header information comprises determining by the packet security gateway that each of the one or more packets comprising the application-layer packet-header information corresponds to at least one of the one or more transport-layer protocols, has a source address within the range of source addresses, a source port within the range of source ports, a destination address within the range of destination addresses, and a destination port within the range of destination ports.

10. A method, comprising:

receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy that includes at least one rule specifying packet-identification criteria and a packet transformation function comprising a packet digest logging function to be performed on packets corresponding to the packet-identification criteria;

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway;

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets corresponding to the packet-identification criteria; and

performing, by the packet security gateway and on a packet-by-packet basis, the packet digest logging function on each of the one or more packets corresponding to the packet-identification criteria.

11. The method of claim 10, wherein performing the packet digest logging function on each of the one or more packets corresponding to the packet-identification criteria comprises:

identifying a subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria; and

generating, for each of the one or more packets corresponding to the packet-identification criteria, a record comprising the subset of information specified by the packet digest logging function.

12. The method of claim 11, wherein the subset of information comprises at least one of a portion of data from the packet, a source address of the packet, a source port of the packet, a destination address of the packet, a destination port of the packet, a transport-protocol type of the packet, a uniform resource identifier (URI) from the packet, an arrival time of the packet, a size of the packet, a flow direction of the packet, an identifier of an interface of the packet security gateway that received the packet, or one or more media access control (MAC) addresses associated with the packet.

13. The method of claim 11, comprising:

temporarily storing data in a memory of the packet security gateway, the data comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria; and

utilizing, by the packet security gateway, the data to generate a message comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria.

14. The method of claim 13 comprising communicating, by the packet security gateway and to the security policy management server, the message comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria.

15. The method of claim 11, comprising reformatting, for each of the one or more packets corresponding to the packet-identification criteria, the subset of information specified by the packet digest logging function in accordance with a logging system standard.

16. The method of claim 15, wherein reformatting the subset of information specified by the packet digest logging function in accordance with the logging system standard comprises

reformatting the subset of information specified by the packet digest logging function in accordance with a syslog logging system standard.

17. The method of claim 10, wherein the packet-identification criteria comprises a five-tuple specifying one or more transport-layer protocols, a range of source addresses, a range of source ports, a range of destination addresses, and a range of destination ports, and wherein identifying the one or more packets corresponding to the packet-identification criteria comprises determining by the packet security gateway that each of the one or more packets corresponding to the packet-identification criteria corresponds to at least one of the one or more transport-layer protocols, has a source address within the range of source addresses, a source port within the range of source ports, a destination address within the range of destination addresses, and a destination port within the range of destination ports.

18. The method of claim 10, wherein the packet-identification criteria comprises application-layer packet-header information, and wherein identifying the one or more packets corresponding to the packet-identification criteria comprises determining by the packet security gateway that each of the one or more packets corresponding to the packet-identification criteria comprises at least a portion of the application-layer packet-header information.

19. A method, comprising:

receiving, by a security policy management server and from a computing device, a security update comprising a set of network addresses;

updating, by the security policy management server, one or more rules stored in a memory of the security policy management server to include the set of network addresses;

receiving, by the security policy management server and from a different computing device, a security update comprising a different set of network addresses;

determining, by the security policy management server, that the different set of network addresses includes at least a portion of network addresses included in the set of network addresses; and

responsive to determining that the different set of network addresses includes the at least a portion of network addresses included in the set of network addresses:

identifying, by the security policy management server, the at least a portion of network addresses included in the set of network addresses;

identifying, by the security policy management server, at least one of the one or more rules stored in the memory of the security policy management server that specifies a range of network addresses comprising the at least a portion of network addresses; and

updating, by the security policy management server, the at least one of the one or more rules to include one or more other network addresses included in the different set of network addresses.

20. The method of claim 19, comprising:

identifying at least two rules, of the one or more rules stored in the memory of the security policy management server, that each specify a range of network addresses comprising the at least a portion of network addresses; and

combining the at least two of the one or more rules into a rule that specifies a range of network addresses that includes network addresses specified by each of the at least two rules and the one or more other network addresses included in the different set of network addresses.

ABSTRACT

Methods and systems for protecting a secured network are presented. For example, one or more packet security gateways may be associated with a security policy management server. At each packet security gateway, a dynamic security policy may be received from the security policy management server, packets associated with a network protected by the packet security gateway may be received, and at least one of multiple packet transformation functions specified by the dynamic security policy may be performed on the packets.

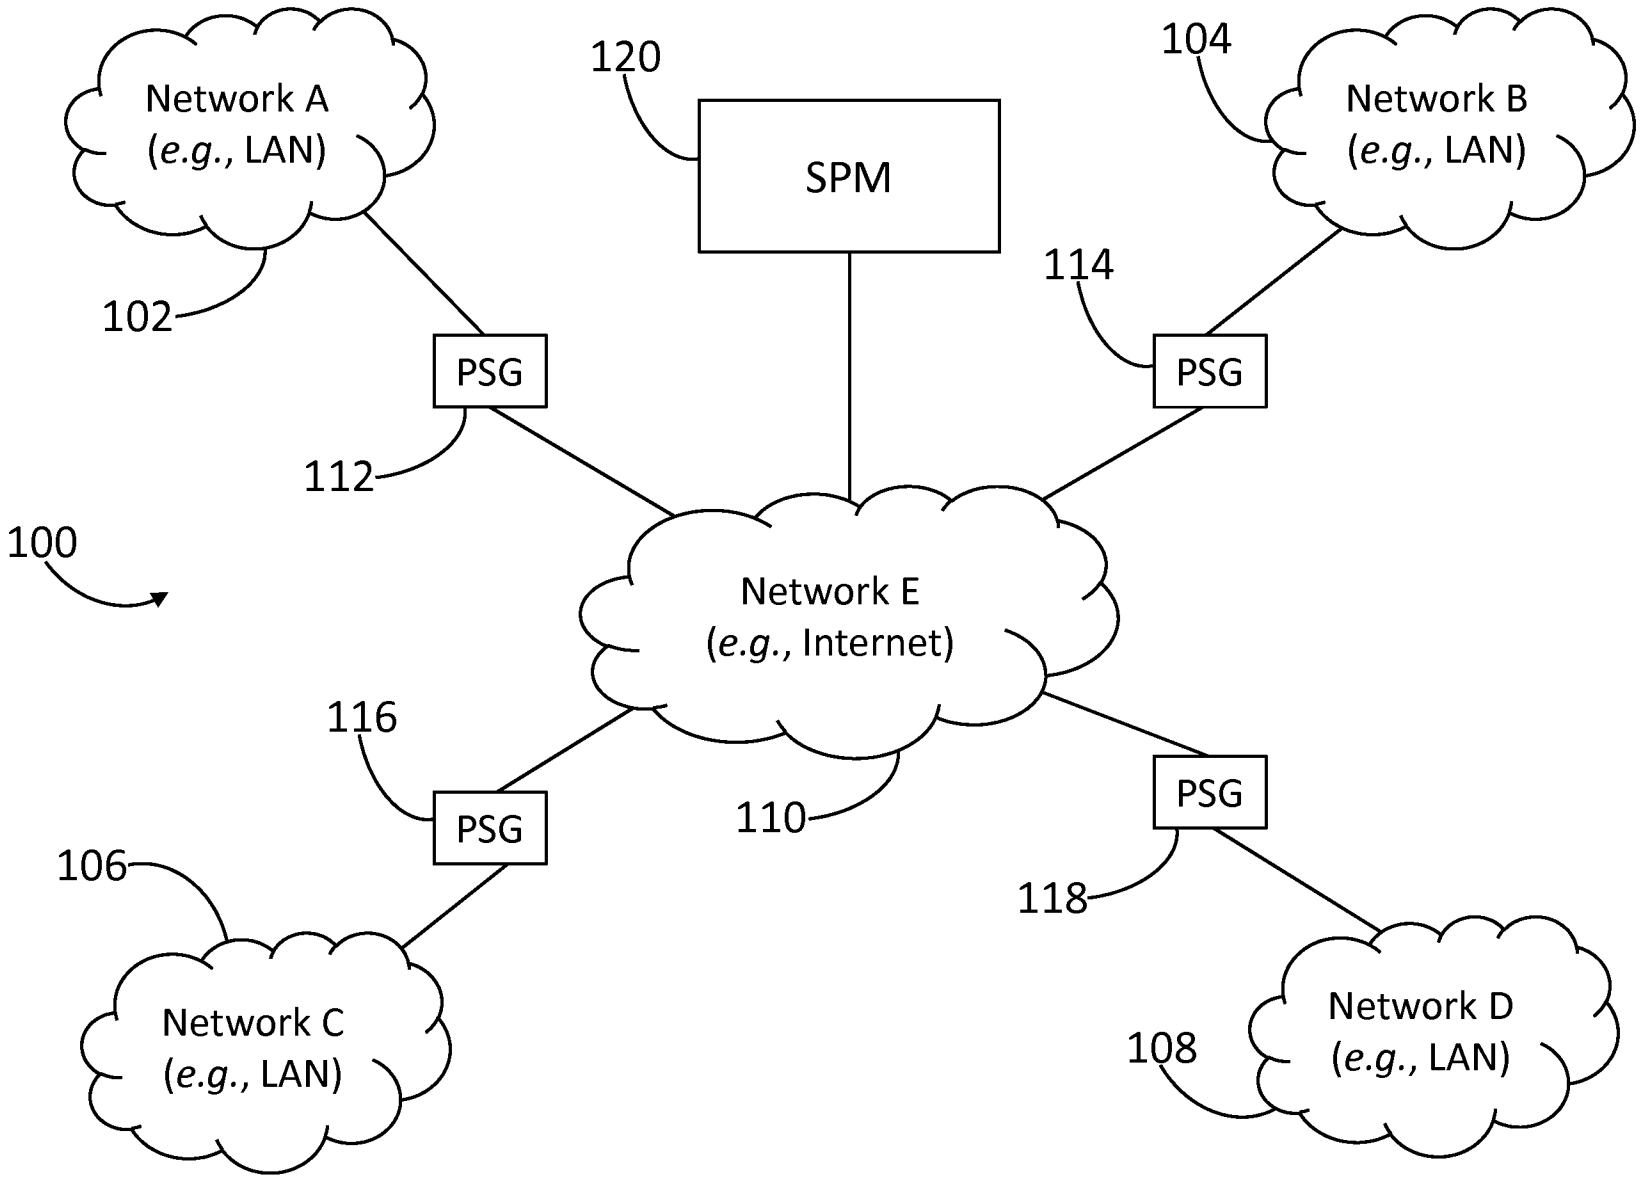

FIG. 1

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 45

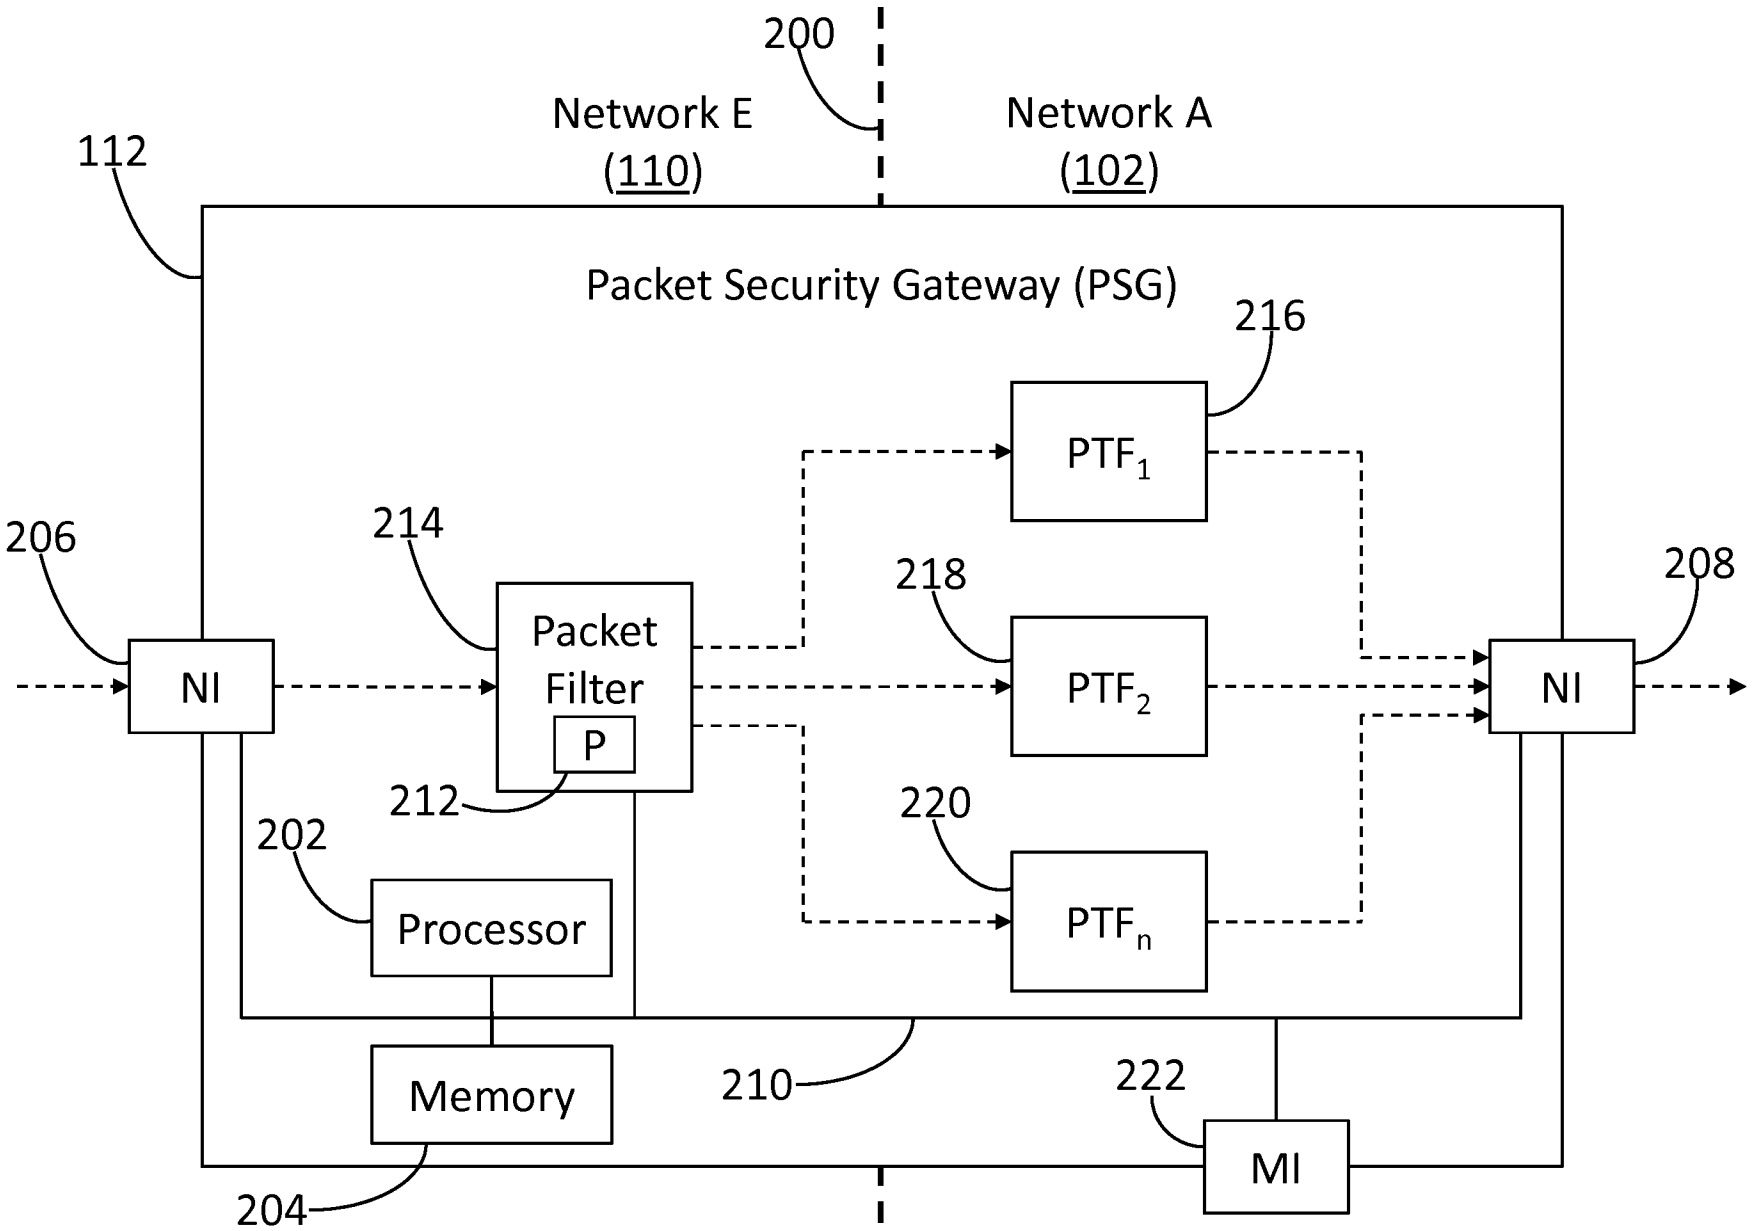

FIG. 2

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 46

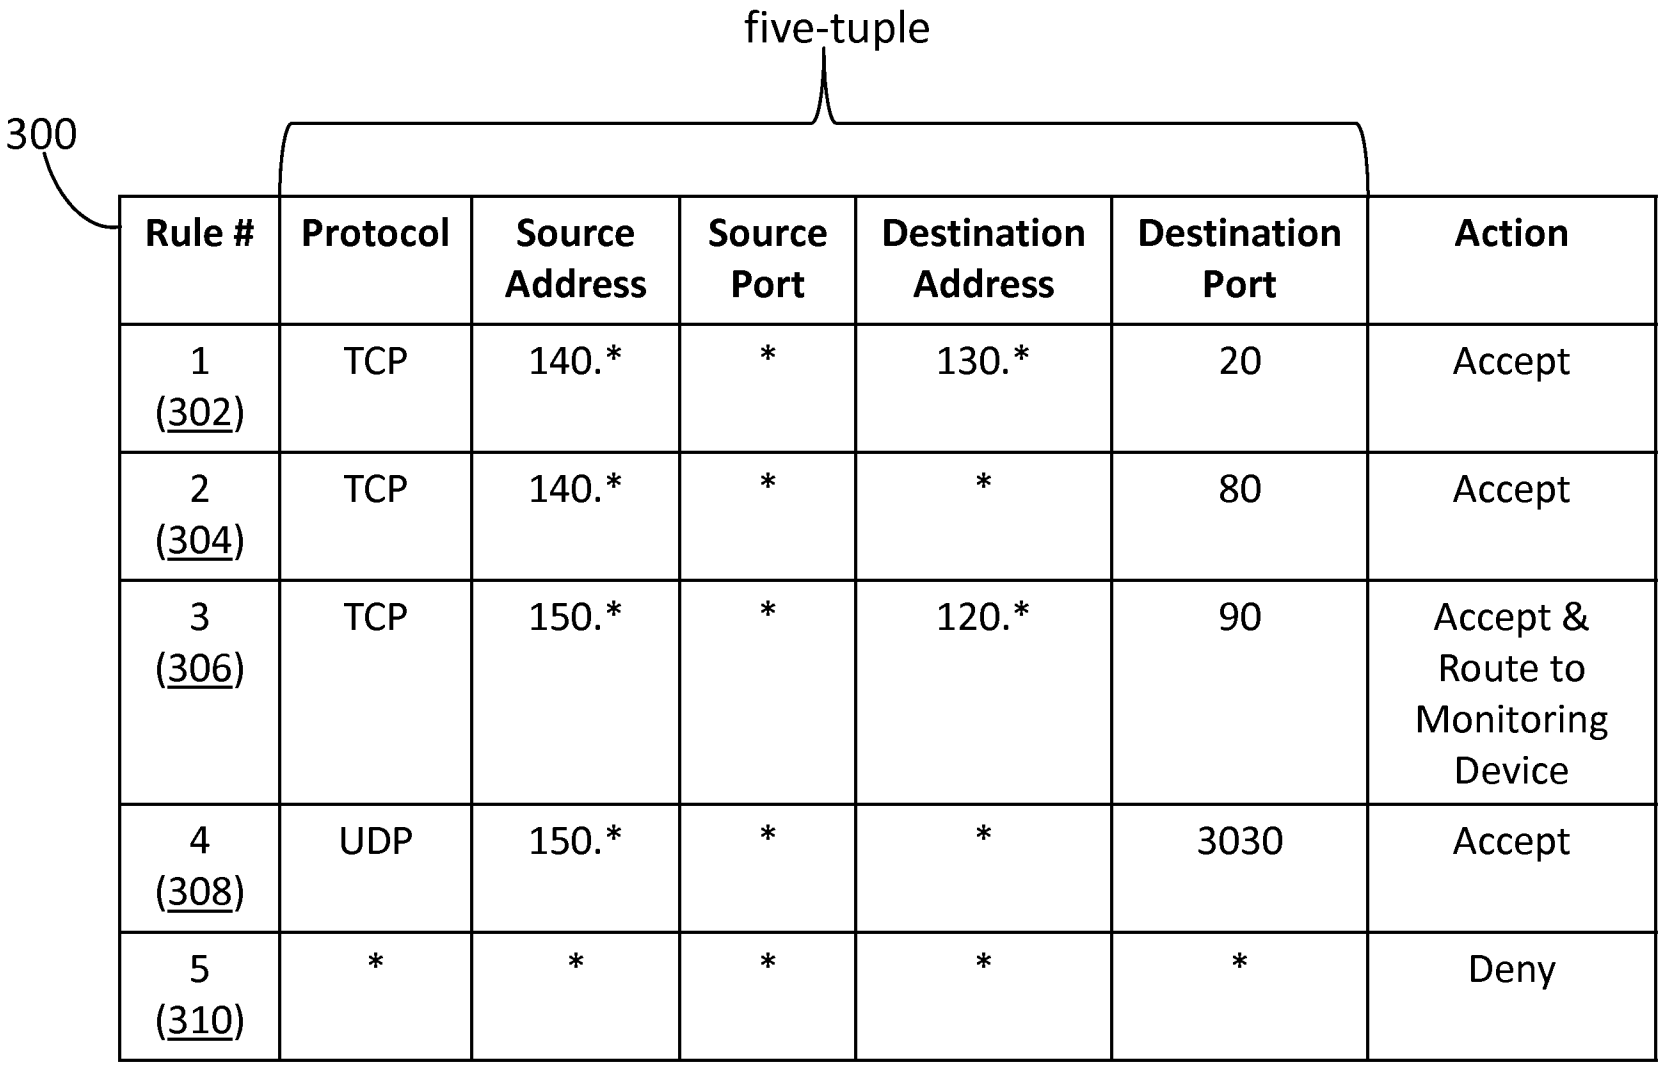

Rule #	Protocol	Source Address	Source Port	Destination Address	Destination Port	Action
1 (302)	TCP	140.*	*	130.*	20	Accept
2 (304)	TCP	140.*	*	*	80	Accept
3 (306)	TCP	150.*	*	120.*	90	Accept & Route to Monitoring Device
4 (308)	UDP	150.*	*	*	3030	Accept
5 (310)	*	*	*	*	*	Deny

FIG. 3

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 47

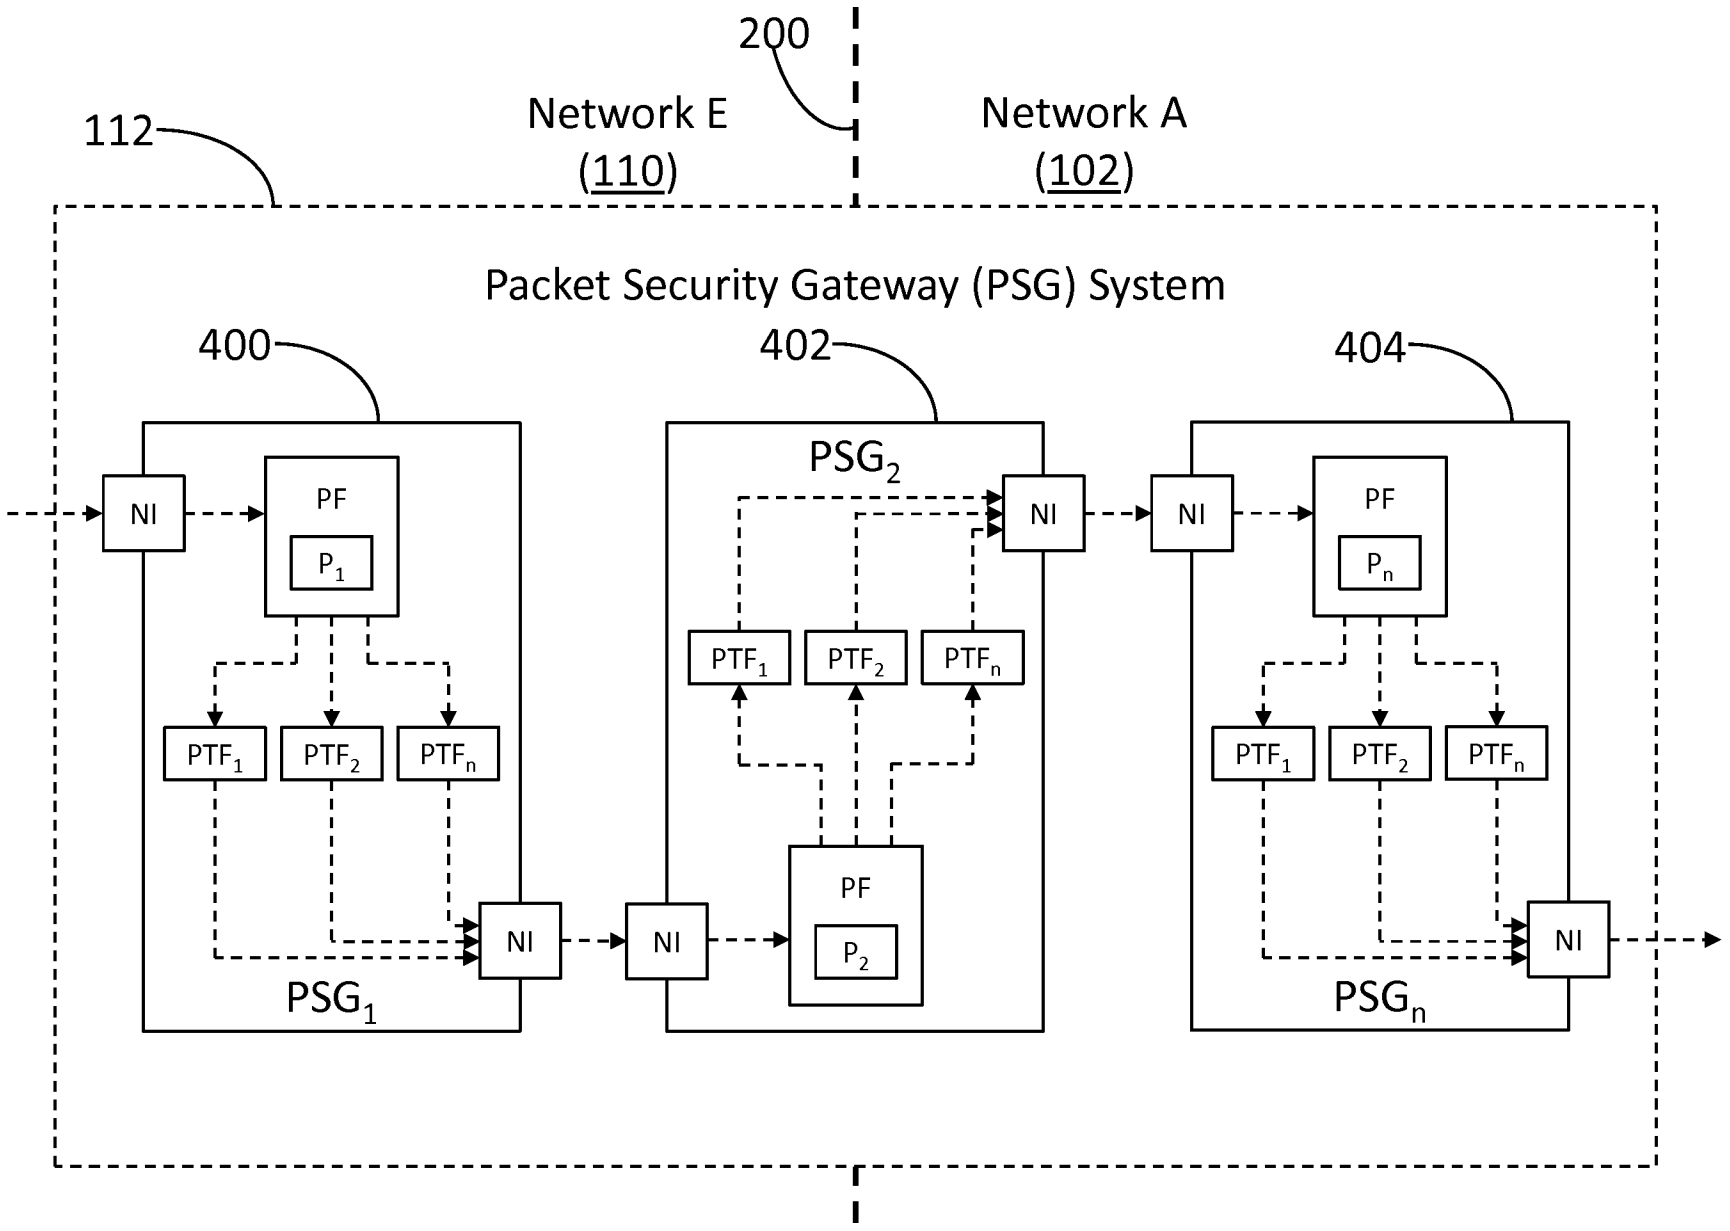

FIG. 4

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 48

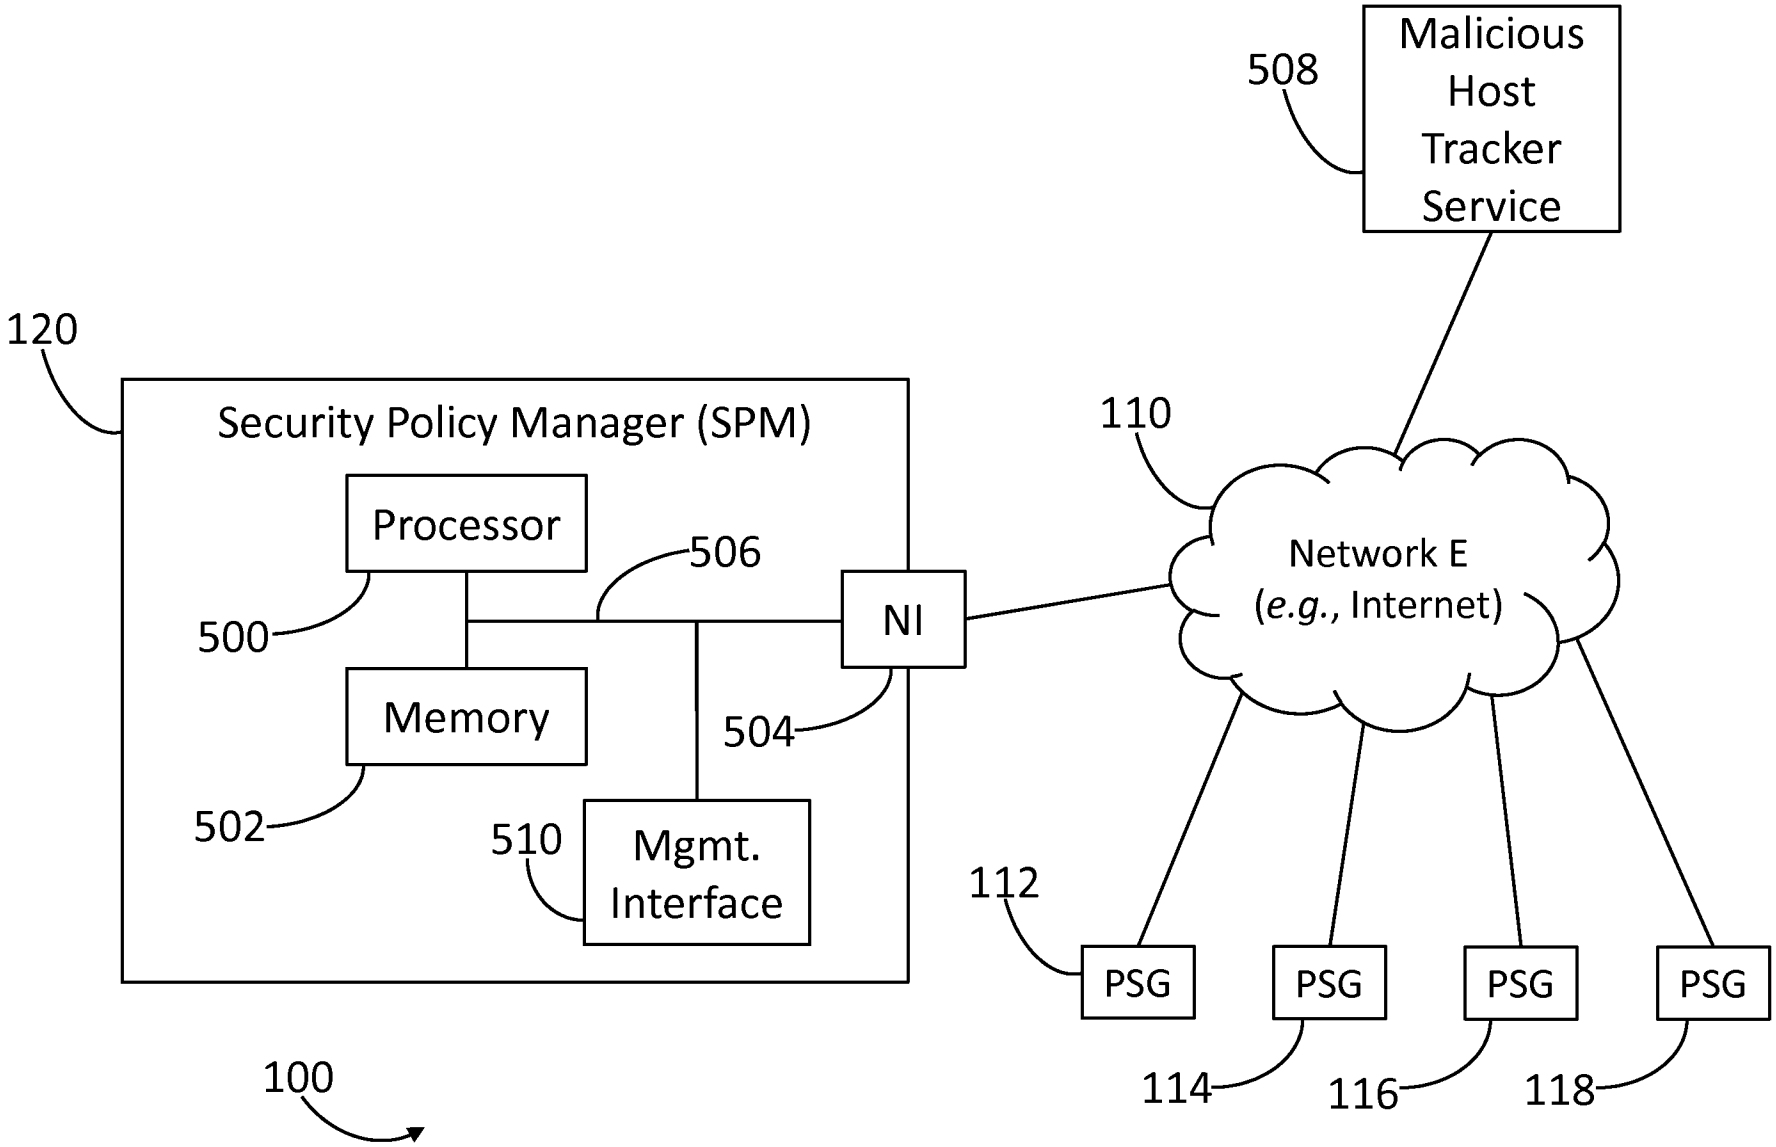

FIG. 5

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 49

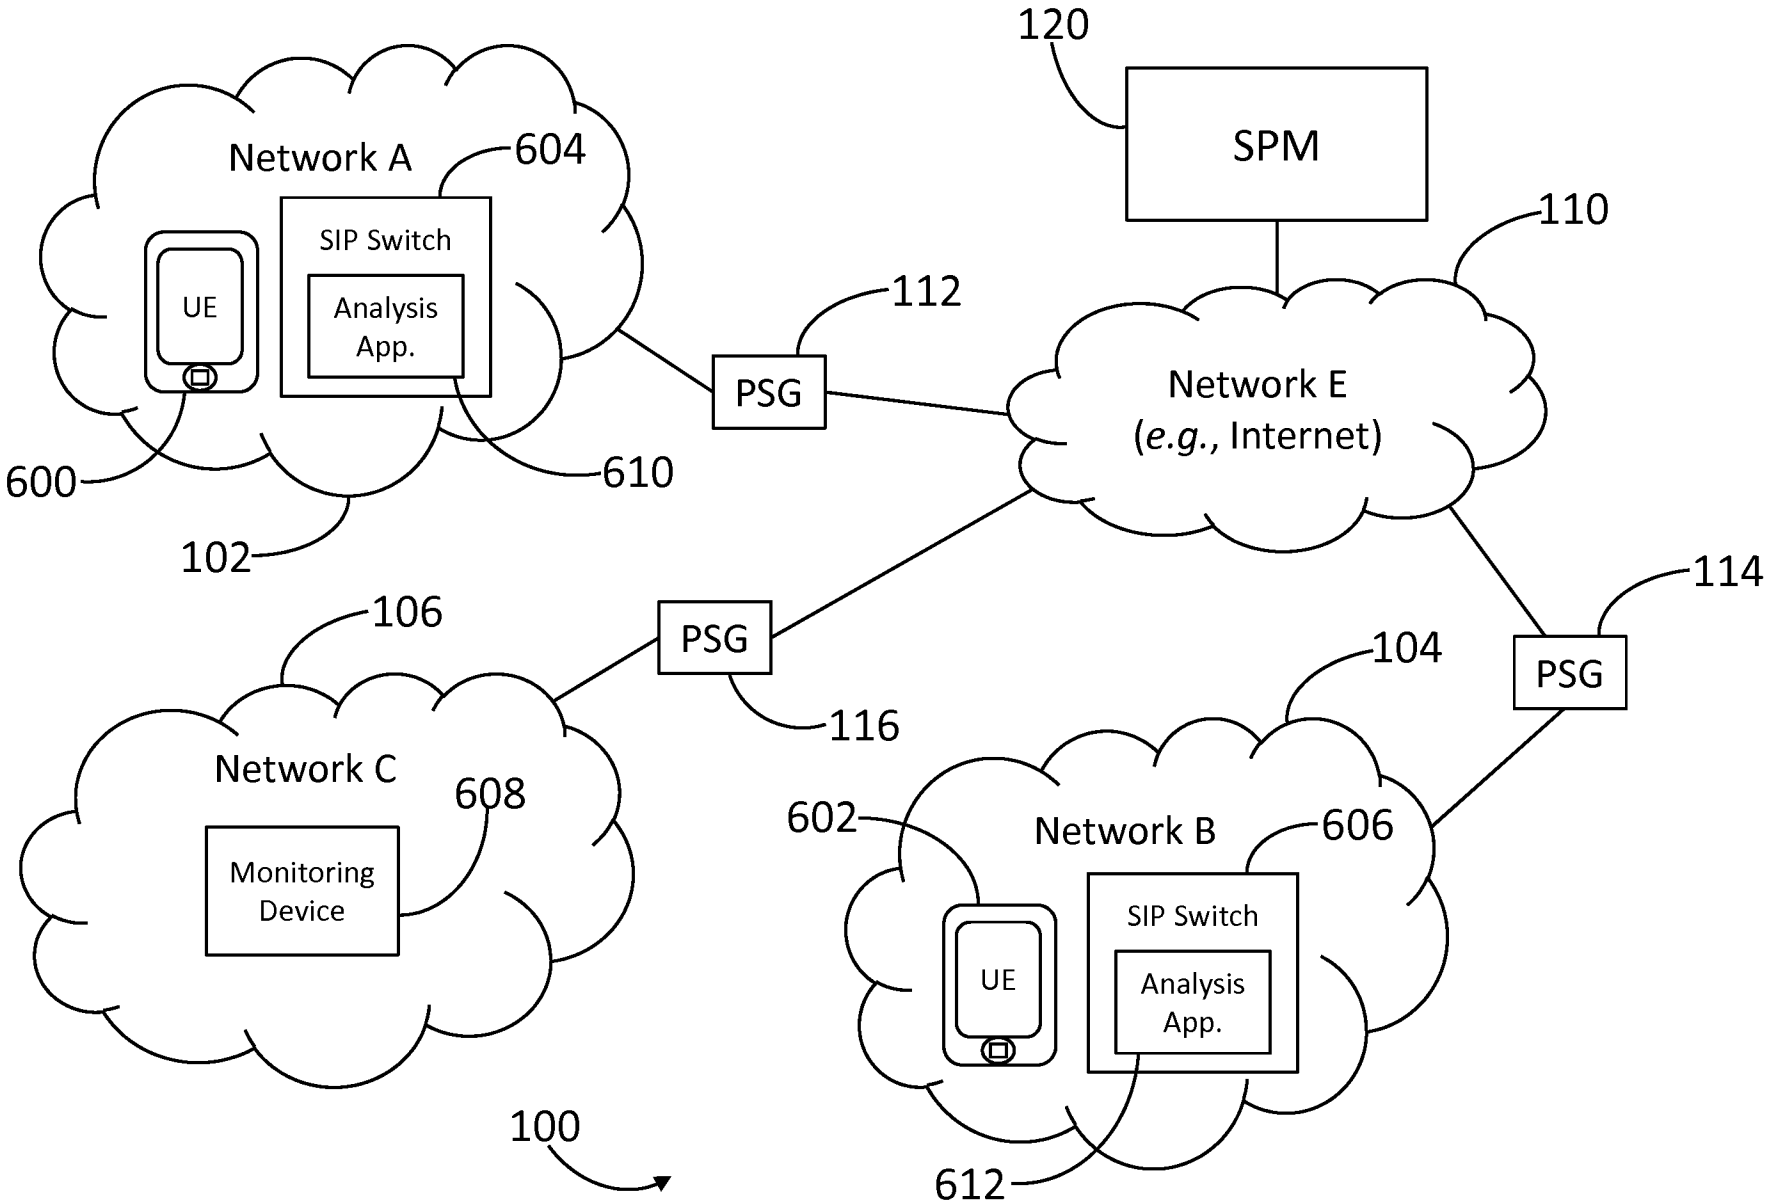

FIG. 6

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 50

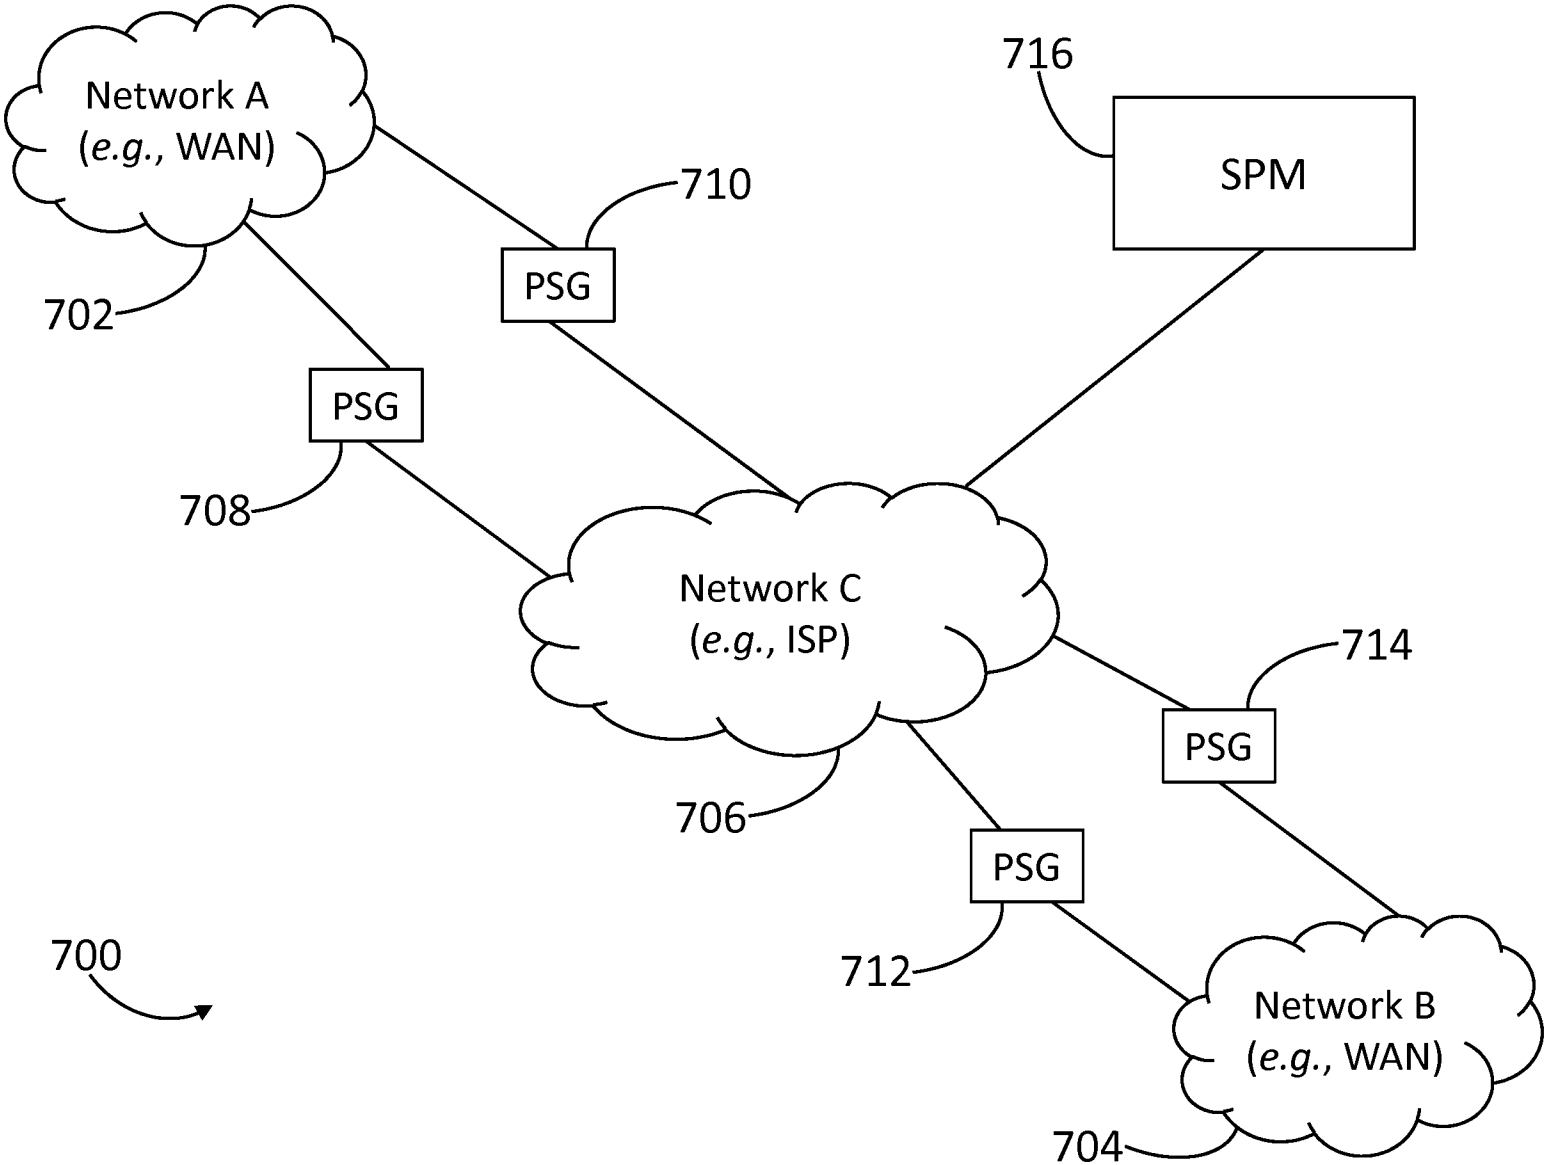

FIG. 7

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 51

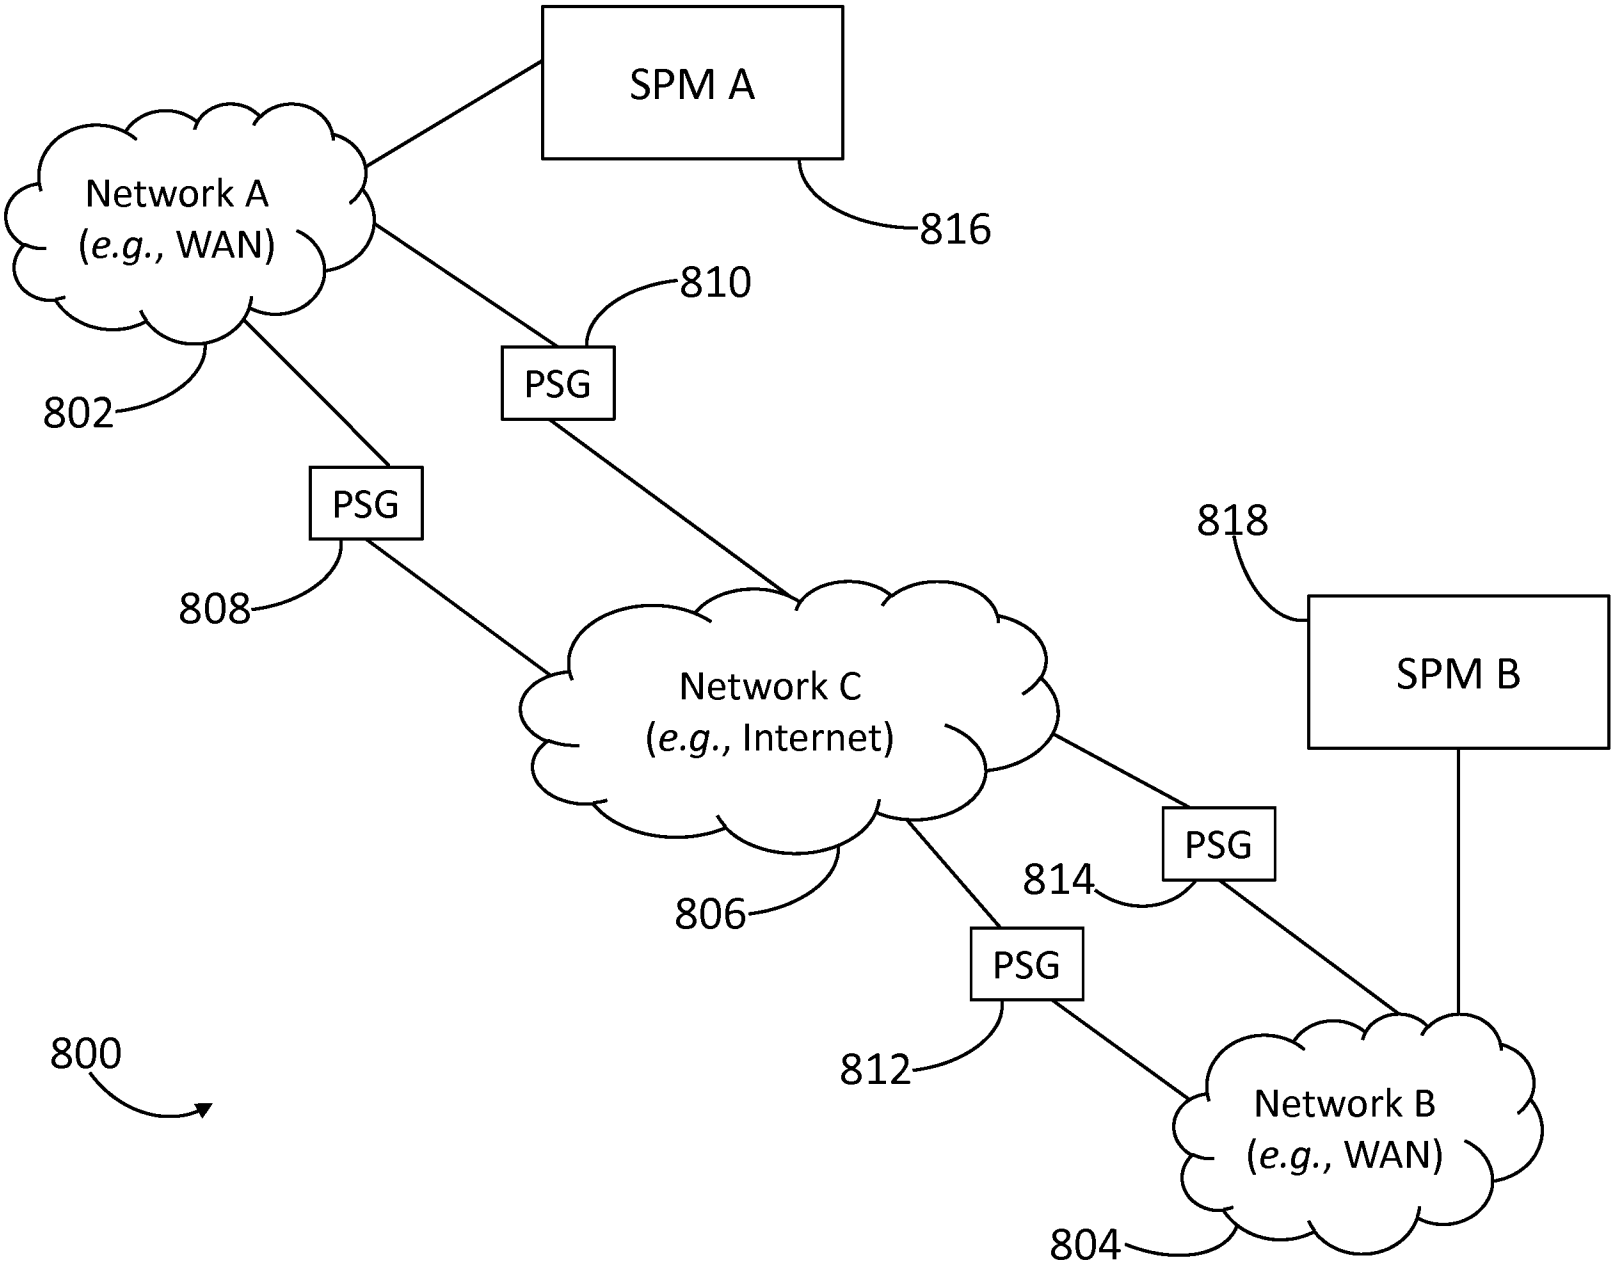

FIG. 8

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 52

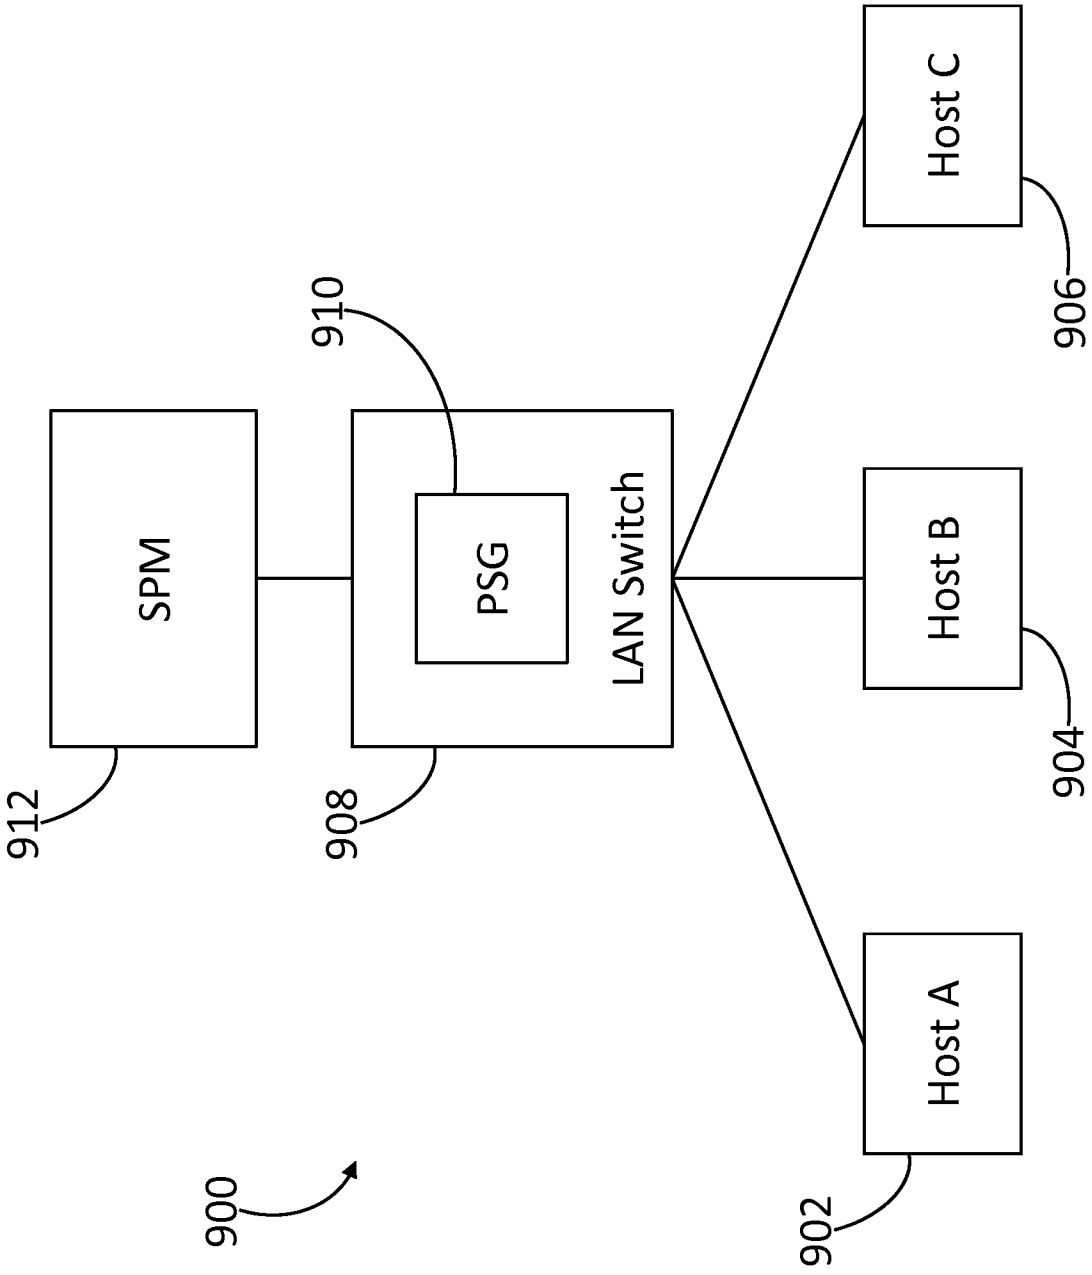
900
902
Host A
904
Host B
906
Host C
908
LAN Switch
910
PSG
912
SPM

FIG. 9

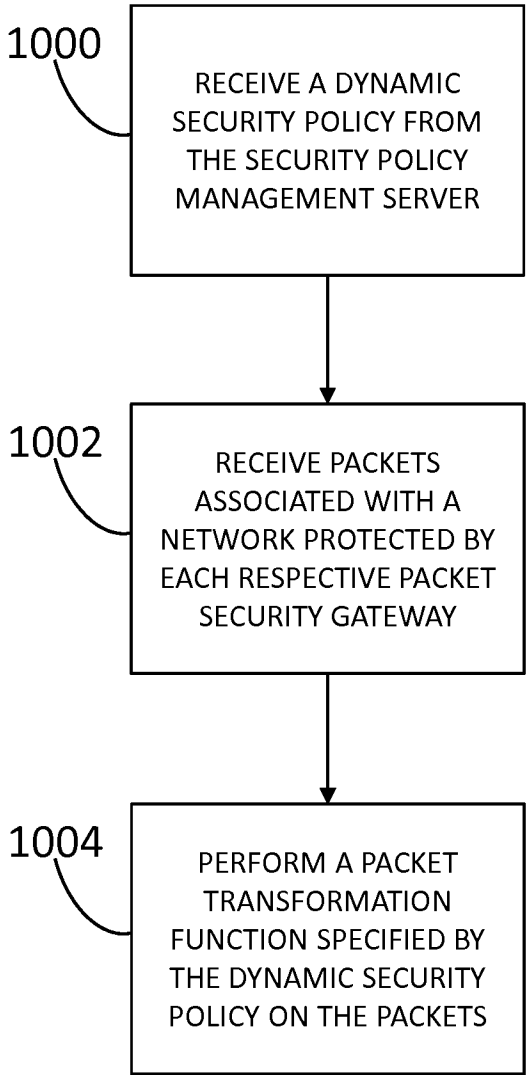

FIG. 10

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 54

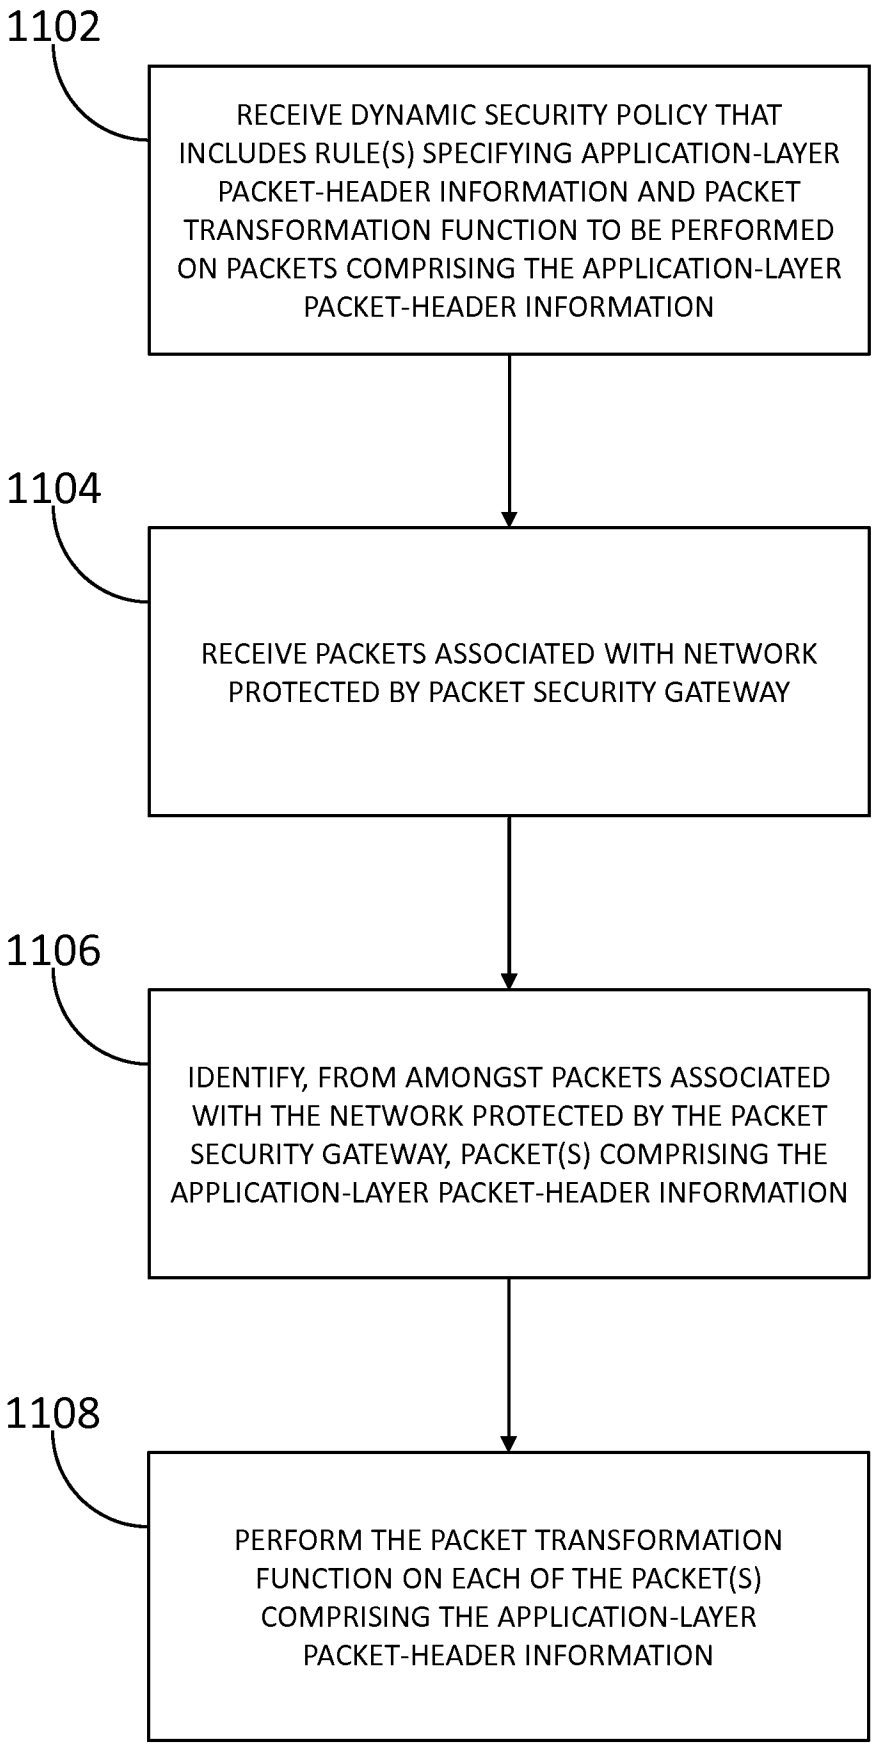
1102
RECEIVE DYNAMIC SECURITY POLICY THAT INCLUDES RULE(S) SPECIFYING APPLICATION-LAYER PACKET-HEADER INFORMATION AND PACKET TRANSFORMATION FUNCTION TO BE PERFORMED ON PACKETS COMPRISING THE APPLICATION-LAYER PACKET-HEADER INFORMATION
1104
RECEIVE PACKETS ASSOCIATED WITH NETWORK PROTECTED BY PACKET SECURITY GATEWAY
1106
IDENTIFY, FROM AMONGST PACKETS ASSOCIATED WITH THE NETWORK PROTECTED BY THE PACKET SECURITY GATEWAY, PACKET(S) COMPRISING THE APPLICATION-LAYER PACKET-HEADER INFORMATION
1108
PERFORM THE PACKET TRANSFORMATION FUNCTION ON EACH OF THE PACKET(S) COMPRISING THE APPLICATION-LAYER PACKET-HEADER INFORMATION

FIG. 11

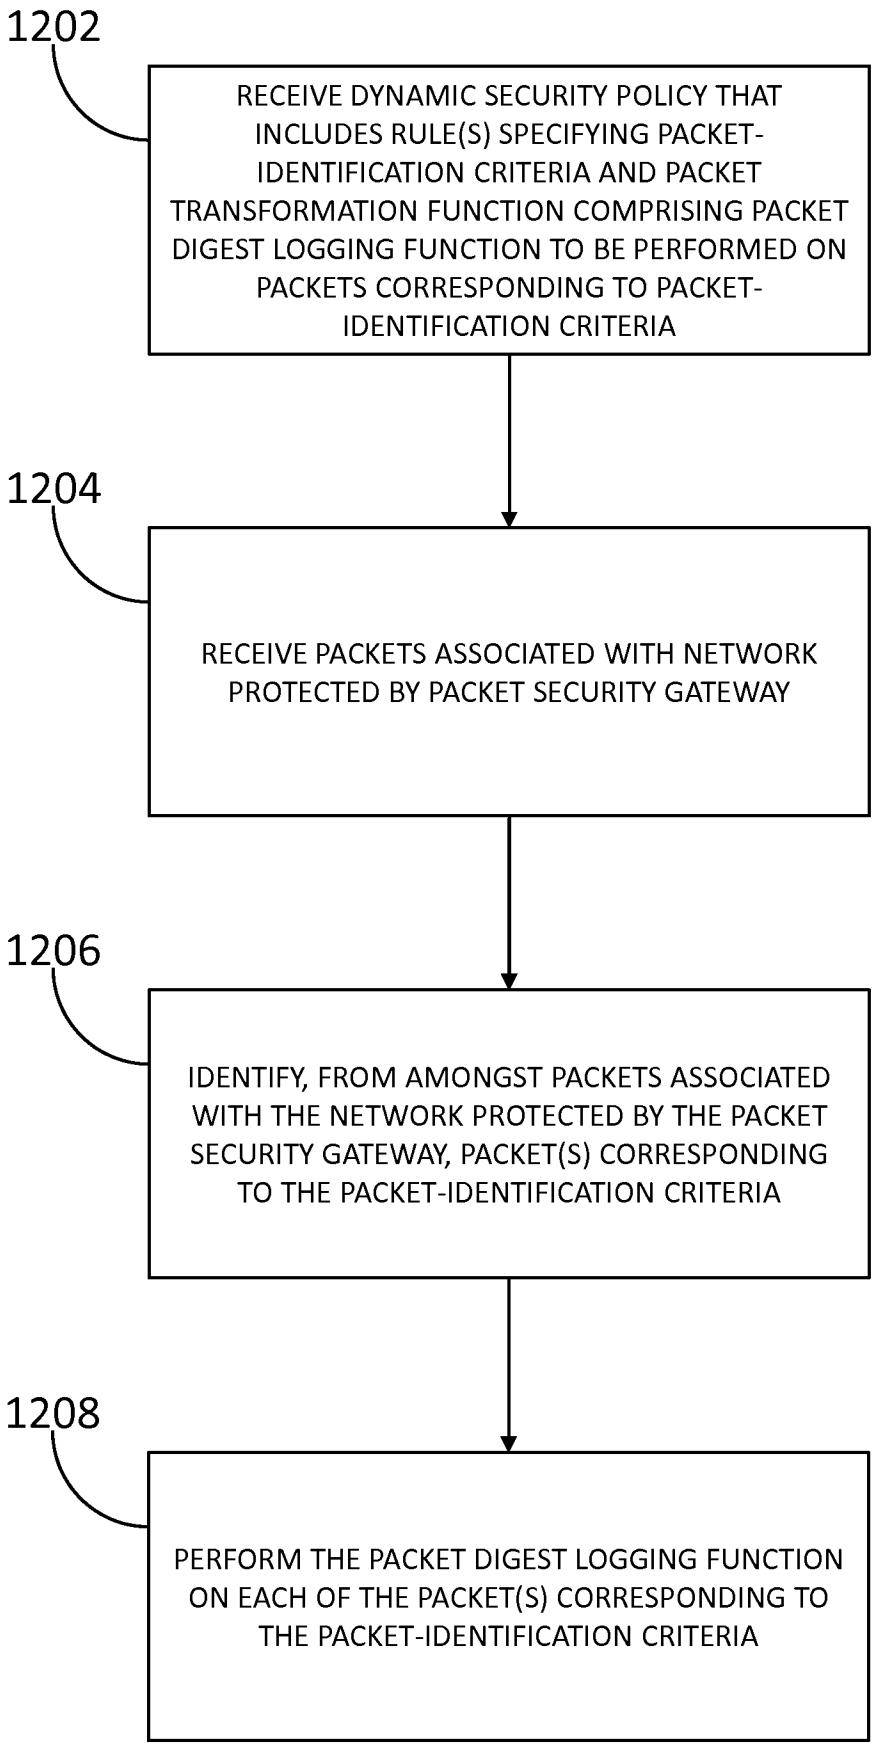
1202
RECEIVE DYNAMIC SECURITY POLICY THAT INCLUDES RULE(S) SPECIFYING PACKET-IDENTIFICATION CRITERIA AND PACKET TRANSFORMATION FUNCTION COMPRISING PACKET DIGEST LOGGING FUNCTION TO BE PERFORMED ON PACKETS CORRESPONDING TO PACKET-IDENTIFICATION CRITERIA
1204
RECEIVE PACKETS ASSOCIATED WITH NETWORK PROTECTED BY PACKET SECURITY GATEWAY
1206
IDENTIFY, FROM AMONGST PACKETS ASSOCIATED WITH THE NETWORK PROTECTED BY THE PACKET SECURITY GATEWAY, PACKET(S) CORRESPONDING TO THE PACKET-IDENTIFICATION CRITERIA
1208
PERFORM THE PACKET DIGEST LOGGING FUNCTION ON EACH OF THE PACKET(S) CORRESPONDING TO THE PACKET-IDENTIFICATION CRITERIA

FIG. 12

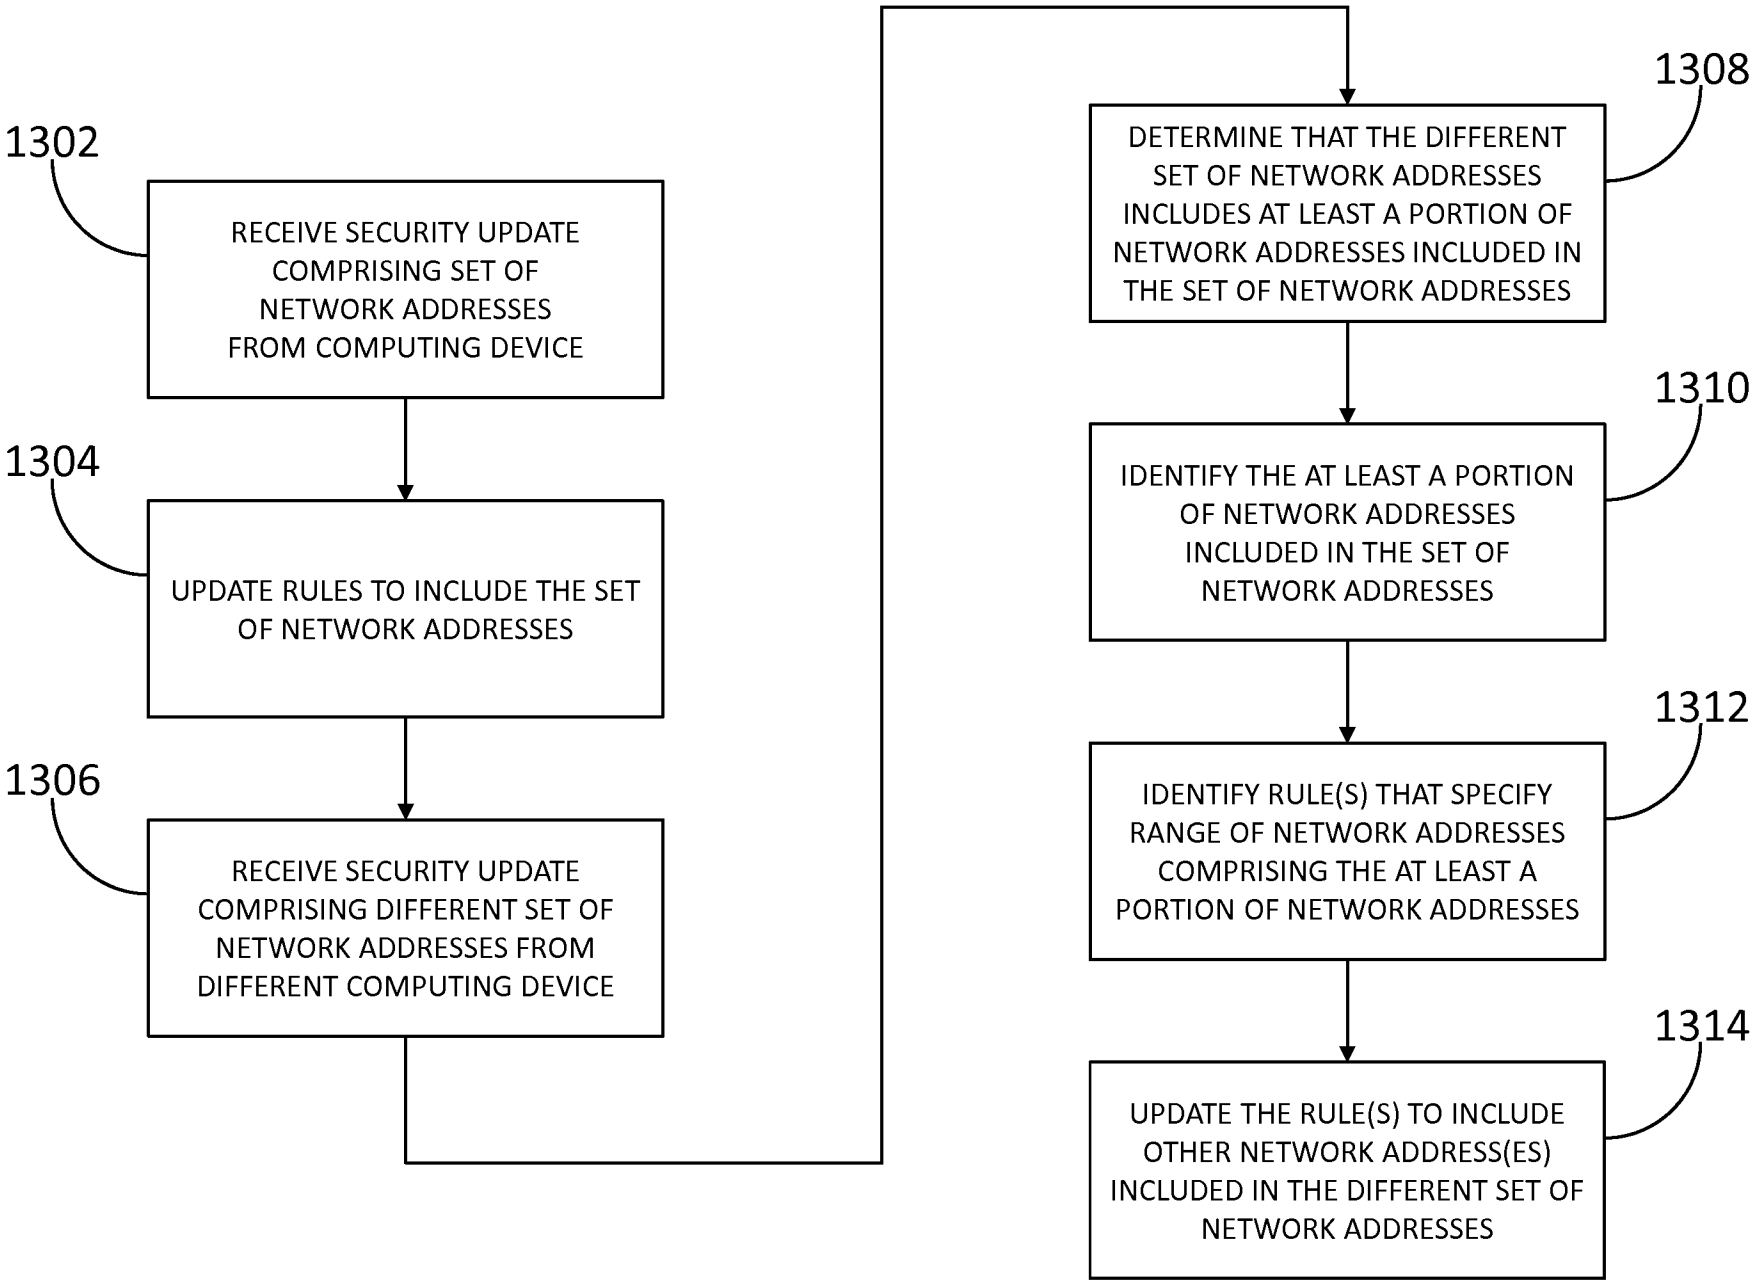

FIG. 13

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 57

PTO/AIA/01 (06-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

As the below named inventor, I hereby declare that:

This declaration is directed to:

- [■] The attached application, or
- [ ] United States application or PCT international application number _____

filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Steven Rogers Date (Optional): 04/16/14

Signature: _____

Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/01 (08-12)
Approved for use through 01/31/2013. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

As the below named inventor, I hereby declare that:

This declaration is directed to:

- [■] The attached application, or
- [ ] United States application or PCT international application number _____
  filed on _____

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Sean Moore Date (Optional): 11-APR-2014

Signature: Sean Moore

Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/01 (06-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

As the below named inventor, I hereby declare that:

This declaration is directed to:

- [x] The attached application, or
- [ ] United States application or PCT international application number _____

  filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: David K. Ahn Date (Optional): 4-16-14

Signature: _____

Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/01 (06-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

As the below named inventor, I hereby declare that:

This declaration is directed to:

- [■] The attached application, or
- [ ] United States application or PCT international application number _____ filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Peter P. Geremia Date (Optional): 4/15/14

Signature: /Peter P. Geremia/

Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Steven		Rogers	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	Leesburg	State/Province	VA	Country of Residence i	US

Mailing Address of Inventor:

Address 1	c/o Centripetal Networks, Inc.
Address 2	11720 Sunrise Valley Drive

City	Reston	State/Province	VA
Postal Code	20191	Country i	US

Inventor 2 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Sean		Moore	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	Hollis	State/Province	NH	Country of Residence i	US

Mailing Address of Inventor:

Address 1	c/o Centripetal Networks, Inc.
Address 2	11720 Sunrise Valley Drive

City	Reston	State/Province	VA
Postal Code	20191	Country i	US

Inventor 3 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	David	K.	Ahn	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

City	Winston-Salem	State/Province	NC	Country of Residence i	US

Mailing Address of Inventor:

Address 1	c/o Centripetal Networks, Inc.		
Address 2	11720 Sunrise Valley Drive		
City	Reston	State/Province	VA
Postal Code	20191	Country i	US

Inventor 4 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Peter	P.	Geremia	

Residence Information (Select One) (•) US Residency ( ) Non US Residency ( ) Active US Military Service

City	Portsmouth	State/Province	NH	Country of Residence i	US

Mailing Address of Inventor:

Address 1	c/o Centripetal Networks, Inc.		
Address 2	11720 Sunrise Valley Drive		
City	Reston	State/Province	VA
Postal Code	20191	Country i	US

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. Add

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

☐ **An Address is being provided for the correspondence Information of this application.**

Customer Number	22907		
Email Address		Add Email	Remove Email

Application Information:

Title of the Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK		
Attorney Docket Number	007742.00029	Small Entity Status Claimed	☒
Application Type	Nonprovisional		
Subject Matter	Utility		
Total Number of Drawing Sheets (if any)	13	Suggested Figure for Publication (if any)	1

Filing By Reference :

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

Only complete this section when filing an application by reference under 35 U.S.C. 111(c) and 37 CFR 1.57(a). Do not complete this section if application papers including a specification and any drawings are being filed. Any domestic benefit or foreign priority information must be provided in the appropriate section(s) below (i.e., "Domestic Benefit/National Stage Information" and "Foreign Priority Information").

For the purposes of a filing date under 37 CFR 1.53(b), the description and any drawings of the present application are replaced by this reference to the previously filed application, subject to conditions and requirements of 37 CFR 1.57(a).

Application number of the previously filed application	Filing date (YYYY-MM-DD)	Intellectual Property Authority or Country

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)

☐ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number	22907		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) or indicate National Stage entry from a PCT application. Providing this information in the application data sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.
When referring to the current application, please leave the application number blank.

Prior Application Status	Pending		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
	Continuation in part of	13657010	2012-10-22

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button. Add

Foreign Priority Information:

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

This section allows for the applicant to claim priority to a foreign application. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55(d). When priority is claimed to a foreign application that is eligible for retrieval under the priority document exchange program (PDX) [i] the information will be used by the Office to automatically attempt retrieval pursuant to 37 CFR 1.55(h)(1) and (2). Under the PDX program, applicant bears the ultimate responsibility for ensuring that a copy of the foreign application is received by the Office from the participating foreign intellectual property office, or a certified copy of the foreign priority application is filed, within the time period specified in 37 CFR 1.55(g)(1).

Remove

Application Number	Country [i]	Filing Date (YYYY-MM-DD)	Access Code [i] (if applicable)

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button. Add

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications

☒ This application (1) claims priority to or the benefit of an application filed before March 16, 2013 and (2) also contains, or contained at any time, a claim to a claimed invention that has an effective filing date on or after March 16, 2013.
NOTE: By providing this statement under 37 CFR 1.55 or 1.78, this application, with a filing date on or after March 16, 2013, will be examined under the first inventor to file provisions of the AIA.

Authorization to Permit Access:

☒ Authorization to Permit Access to the Instant Application by the Participating Offices

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

If checked, the undersigned hereby grants the USPTO authority to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the World Intellectual Property Office (WIPO), and any other intellectual property offices in which a foreign application claiming priority to the instant patent application is filed access to the instant patent application. See 37 CFR 1.14(c) and (h). This box should not be checked if the applicant does not wish the EPO, JPO, KIPO, WIPO, or other intellectual property office in which a foreign application claiming priority to the instant patent application is filed to have access to the instant patent application.

In accordance with 37 CFR 1.14(h)(3), access will be provided to a copy of the instant patent application with respect to: 1) the instant patent application-as-filed; 2) any foreign application to which the instant patent application claims priority under 35 U.S.C. 119(a)-(d) if a copy of the foreign application that satisfies the certified copy requirement of 37 CFR 1.55 has been filed in the instant patent application; and 3) any U.S. application-as-filed from which benefit is sought in the instant patent application.

In accordance with 37 CFR 1.14(c), access may be provided to information concerning the date of filing this Authorization.

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Applicant 1 Remove

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

Clear

(•) Assignee	( ) Legal Representative under 35 U.S.C. 117	( ) Joint Inventor
( ) Person to whom the inventor is obligated to assign.	( ) Person who shows sufficient proprietary interest	

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor :

If the Applicant is an Organization check here. [X]

Organization Name	Centripetal Networks, Inc.		
Mailing Address Information:			
Address 1	11720 Sunrise Valley Drive		
Address 2			
City	Reston	**State/Province**	VA
Country [i]	US	Postal Code	20191
Phone Number		Fax Number	

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

Email Address	

Additional Applicant Data may be generated within this form by selecting the Add button. Add

Assignee Information including Non-Applicant Assignee Information:

Providing assignment information in this section does not subsitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Assignee 1

Complete this section if assignee information, including non-applicant assignee information, is desired to be included on the patent application publication . An assignee-applicant identified in the "Applicant Information" section will appear on the patent application publication as an applicant. For an assignee-applicant, complete this section only if identification as an assignee is also desired on the patent application publication.

Remove

If the Assignee or Non-Applicant Assignee is an Organization check here. ☐

Prefix	**Given Name**	Middle Name	**Family Name**	Suffix

Mailing Address Information For Assignee including Non-Applicant Assignee:

Address 1			
Address 2			
City		**State/Province**	
Country i		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Assignee or Non-Applicant Assignee Data may be generated within this form by selecting the Add button. Add

Signature:

Remove

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications

Signature	/William E. Wooten/			Date (YYYY-MM-DD)	2014-04-16
First Name	William	Last Name	Wooten	Registration Number	60049

Additional Signature may be generated within this form by selecting the Add button. Add

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Patent Application Fee Transmittal

Application Number:	
Filing Date:	
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Filer:	William Eugene Wooten
Attorney Docket Number:	007742.00029

Filed as Small Entity

Utility under 35 USC 111(a) Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Utility filing Fee (Electronic filing)	4011	1	70	70
Utility Search Fee	2111	1	300	300
Utility Examination Fee	2311	1	360	360
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		**730**

Electronic Acknowledgement Receipt

EFS ID:	18775695
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	16-APR-2014
Filing Date:	
Time Stamp:	11:19:50
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Deposit Account
Payment was successfully received in RAM	$730
RAM confirmation Number	8887
Deposit Account	190733
Authorized User	

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

Charge any Additional Fees required under 37 C.F.R. Section 1.16 (National application filing, search, and examination fees)

Charge any Additional Fees required under 37 C.F.R. Section 1.17 (Patent application and reexamination processing fees)

Charge any Additional Fees required under 37 C.F.R. Section 1.19 (Document supply fees)

Charge any Additional Fees required under 37 C.F.R. Section 1.20 (Post Issuance fees)

Charge any Additional Fees required under 37 C.F.R. Section 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Specification	Spec.pdf	240274 3c3264178c85a3a7574c20a23d062ae9b1f6ac8d	no	44
Warnings:					
Information:					
2	Drawings-only black and white line drawings	Figs.pdf	203926 bd0950ccf94abda924edc7418f2bfd804b87e972	no	13
Warnings:					
Information:					
3	Oath or Declaration filed	Decs.pdf	3197951 6ac2d08f8b4c7eb67d3f861cb6274b3666ac6ba8	no	4
Warnings:					
Information:					
4	Application Data Sheet	ADS.pdf	1561477 f0475a3e696024b7963ff72017d20487daa1246b	no	8
Warnings:					
Information:					
5	Fee Worksheet (SB06)	fee-info.pdf	32951 493d365f0b548df19c0dfca8bb70740a4f9f304a	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	5236579		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	TBA
Examiner Name	TBA
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	6226372		2001-05-01	Beebe et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	US200601958		2006-08-31	Fulp et al.	
	2	US200602485		2006-11-02	Fulp et al.	
	3	20110055916		2011-03-03	Ahn	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	1677484	EP	A2	2006-07-05	Microsoft Corp [us]		☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	TBA
	Examiner Name	TBA
	Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	John Kindervag; "Build Security Into Your Network's DNA: The Zero Trust Network Architecture", Forrester Research Inc.; Novermber 5, 2010, pps: 1-26.	☐
	2	Palo Alto Networks; "Designing A Zero Trust Network With Next-Generation Firewalls"; pps. 1-10; last viewed on October 21, 2012.	☐
	3	Copy of International Search Report for PCT/US2013/057502, dated November 7, 2013	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	TBA
	Examiner Name	TBA
	Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2014-04-16
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Espacenet

Bibliographic data: EP1677484 (A2) — 2006-07-05

Method and system for distributing security policies

No documents available for this priority number.

Inventor(s): PALEOLOGU EMANUEL [US]; YOUNGBLUT ERIC E [US]; IVANOV MAXIM A [US]; NAGAMPALLI NARASIMHA RAO S S [US]; SHETH SACHIN C [US]; KOTI SHIRISH R [US]; LIN YUN [US] + (PALEOLOGU, EMANUEL, ; YOUNGBLUT, ERIC E, ; IVANOV, MAXIM A, ; NAGAMPALLI, NARASIMHA RAO, ; SHETH, SACHIN C, ; KOTI, SHIRISH R, ; LIN, YUN, ; NAGAMPALLI, NARASIMHA RAO S.S)

Applicant(s): MICROSOFT CORP [US] + (MICROSOFT CORPORATION)

Classification: **- international:** ***H04L29/06***
- cooperative: **H04L63/0218**; **H04L63/0263**; **H04L63/20**

Application number: EP20050110420 20051107

Priority number(s): US20040993688 20041119

Also published as: EP1677484 (A3) US2006129808 (A1) US7509493 (B2) US2009172774 (A1) US7831826 (B2) more

Abstract of EP1677484 (A2)

A method and system for distributing and enforcing security policies is provided. A firewall agent executing at a host computer system that is to be protected receives security policies for the enforcement engines responsible for enforcing the security policies on the host computer system. A security policy has rules that each provide a condition and action to be performed when the condition is satisfied. A rule also has a rule type that is used by the distribution system to identify the security components that are responsible for enforcing the rules. To distribute the security policies that have been received at a host computer system, the firewall agent identifies to which enforcement engine a rule applies based in part on rule type. The firewall agent then distributes the rule to the identified enforcement engine, which then enforces the rule.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 677 484 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
05.07.2006 Bulletin 2006/27

(51) Int Cl.:
H04L 29/06 *(2006.01)*

(21) Application number: **05110420.6**

(22) Date of filing: **07.11.2005**

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IS IT LI LT LU LV MC NL PL PT RO SE SI SK TR
Designated Extension States:
AL BA HR MK YU

(30) Priority: **19.11.2004 US 993688**

(71) Applicant: **MICROSOFT CORPORATION**
Redmond, Washington 98052-6399 (US)

(72) Inventors:
- **Paleologu, Emanuel,**
  cO Microsoft Corporation
  Redmond,
  WA 98052 (US)
- **Youngblut, Eric E.,**
  c/o Microsoft Corporation
  Redmond,
  WA 98052 (US)
- **Ivanov, Maxim A.,**
  c/o Microsoft Corporation
  Redmond,
  WA 98052 (US)
- **Nagampalli, Narasimha Rao S.S.,**
  c/o Microsoft Corp
  Redmond,
  WA 98052 (US)
- **Sheth, Sachin C.,**
  c/o Microsoft Corporation
  Redmond,
  WA 98052 (US)
- **Koti, Shirish R.,**
  c/o Microsoft Corporation
  Redmond,
  WA 98052 (US)
- **Lin, Yun,**
  c/o Microsoft Corporation
  Redmond,
  WA 98052 (US)

(74) Representative: **Schuster, Thomas**
Grünecker, Kinkeldey,
Stockmair & Schwanhäusser
Anwaltssozietät
Maximilianstrasse 58
D-80538 München (DE)

(54) Method and system for distributing security policies

(57) A method and system for distributing and enforcing security policies is provided. A firewall agent executing at a host computer system that is to be protected receives security policies for the enforcement engines responsible for enforcing the security policies on the host computer system. A security policy has rules that each provide a condition and action to be performed when the condition is satisfied. A rule also has a rule type that is used by the distribution system to identify the security components that are responsible for enforcing the rules. To distribute the security policies that have been received at a host computer system, the firewall agent identifies to which enforcement engine a rule applies based in part on rule type. The firewall agent then distributes the rule to the identified enforcement engine, which then enforces the rule.

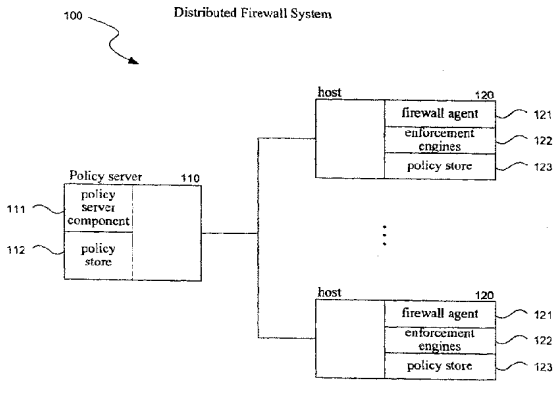

FIG. 1

EP 1 677 484 A2

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 81

Description

TECHNICAL FIELD

[0001] The described technology relates generally to systems that prevent the exploitation of vulnerabilities of computer systems.

BACKGROUND

[0002] Although the Internet has had great successes in facilitating communications between computer systems and enabling electronic commerce, the computer systems connected to the Internet have been under almost constant attack by hackers seeking to disrupt their operation. Many of the attacks seek to exploit vulnerabilities of software systems including application programs or other computer programs executing on those computer systems. Developers of software systems and administrators of computer systems of an enterprise go to great effort and expense to identify and remove vulnerabilities. Because of the complexity of software systems, however, it is virtually impossible to identify and remove all vulnerabilities before software systems are released. After a software system is released, developers can become aware of vulnerabilities in various ways. A party with no malicious intent may identify a vulnerability and may secretly notify the developer so the vulnerability can be removed before a hacker identifies and exploits it. If a hacker identifies a vulnerability first, the developer may not learn of the vulnerability until it is exploited—sometimes with disastrous consequences.

[0003] Regardless of how a developer finds out about a vulnerability, the developer typically develops and distributes to system administrators "patches" or updates to the software system that remove the vulnerability. If the vulnerability has not yet been exploited (e.g., might not be known to hackers), then a developer can design, implement, test, and distribute a patch in a disciplined way. If the vulnerability has already been widely exposed, then the developer may rush to distribute a patch without the same care that is used under normal circumstances. When patches are distributed to the administrators of the computer systems, they are responsible for scheduling and installing the patches to remove the vulnerabilities.

[0004] Unfortunately, administrators often delay the installation of patches to remove vulnerabilities for various reasons. When a patch is installed, the software system and possibly the computer system on which it is executing may need to be shut down and restarted. If the vulnerability is in a software system that is critical to the success of an organization, then the administrator needs to analyze the tradeoffs of keeping the software system up and running with its associated risk of being attacked and of shutting down a critical resource of the enterprise to install the patch. Some administrators may delay the installation of the patch because they fear that, because of a hasty distribution, it might not be properly tested and have unintended side effects. If the patch has an unintended side effect, then the software system, the computer system, or some other software component that is impacted by the patch may be shut down by the patch itself. Administrators need to factor in the possibility of an unintended side effect when deciding whether to install a patch. These administrators may delay installing a patch until experience by others indicates that there are no serious unintended side effects.

[0005] Intrusion detection systems have been developed that can be used to identify whether an attempt is being made to exploit a known vulnerability that has not yet been patched. These intrusion detection systems can be used to prevent exploitations of newly discovered vulnerabilities for which patches have not yet been developed or installed. These intrusion detection systems may define a "signature" for each way a vulnerability can be exploited. For example, if a vulnerability can be exploited by sending a certain type of message with a certain attribute, then the signature for that exploitation would specify that type and attribute. When a security enforcement event occurs, such as the receipt of a message, the intrusion detection system checks its signatures to determine whether any match the security enforcement event. If so, the intrusion detection system may take action to prevent the exploitation, such as dropping the message.

[0006] Signatures for newly discovered exploitations of vulnerabilities can be created in different ways. Developers of intrusion detection systems may create and distribute new signatures when they become aware of new exploitations. An administrator can then install the new signatures to prevent the exploitation. A developer may not, however, provide signatures for all known exploitations. For example, the vulnerability may be in a special-purpose application program that the developer does not support. To prevent exploitation of such a vulnerability, intrusion detection systems may allow administrators to create their own signatures.

[0007] A set of one or more signatures may be considered a security policy. Developers of intrusion detection systems may provide various security policies. For example, a developer may provide one security policy that defines signatures of vulnerabilities of an operating system and many other security policies that are specific to an application or a class of applications. Similarly, an administrator may define a security policy that is specific to custom applications used by the enterprise.

[0008] Because intrusions can occur at various points within an operating system or an application, intrusion detection systems have been developed to detect and prevent exploitation of vulnerabilities at each of these points. For example, an intrusion detection system may be developed to prevent exploitation of vulnerabilities that can be detected at the network layer of a communication protocol, and another intrusion detection system may be developed to prevent exploitation of vulnerabilities that can be detected when a file is accessed. Each

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 82

developer of an intrusion detection system can provide their own implementation and security policies that need to be distributed to each computer system to be protected. In addition, any updates to the implementation or security policies also need to be distributed to the computer systems to be protected.

[0009] The distribution and updating of security policies for an enterprise with a large number of computer systems can be a complex and time-consuming task. Each security policy needs to be provided to each computer system and then to the appropriate intrusion detection system that is installed on the computer system. Since the developers of intrusion detection systems typically develop their systems independent of each other, each intrusion detection system may have a proprietary mechanism for defining and distributing security policies. A system administrator of an enterprise may need to be aware of each of these mechanisms to use the intrusion detection system effectively.

[0010] It would be desirable to have a mechanism that would provide a uniform way to distribute security policies received at a computer system to components responsible for enforcing the security policies.

SUMMARY

[0011] A method and system for distributing and enforcing security policies is provided. A distributed firewall system includes a policy server component, firewall agents, and enforcement engines. A firewall agent executing at a host computer system that is to be protected receives from the policy server component security policies for the enforcement engines responsible for enforcing the security policies on the host computer system. A security policy has rules that each provide a condition and action to be performed when the condition is satisfied. A rule also has a rule type (also referred to as a security type) that is used by the firewall agent to identify the enforcement engines that are responsible for enforcing the rules. The firewall agent may distribute the security policies to enforcement engines executing in user mode and in kernel mode. The enforcement engines may provide a layered approach to security enforcement. To distribute the security policies that have been received at a host computer system, the firewall agent identifies to which enforcement engine a rule applies based in part on rule type. The firewall agent then distributes the rule to the identified enforcement engine, which then enforces the rule.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] Figure 1 is a block diagram that illustrates the distributed firewall system in one embodiment.

[0013] Figure 2 is a block diagram that illustrates the components of the portions of the distributed firewall system executing on a host computer system in one embodiment.

[0014] Figure 3 is a flow diagram that illustrates the processing of the distribute policy component of the firewall agent in one embodiment.

[0015] Figure 4 is a flow diagram that illustrates the processing of the flow manager component in one embodiment.

[0016] Figure 5 is a flow diagram that illustrates the processing of an enforcement engine for a specific layer in one embodiment.

DETAILED DESCRIPTION

[0017] A method and system for distributing and enforcing security policies is provided. A distributed firewall system includes a policy server component, firewall agents, and enforcement engines. A firewall agent executing at a host computer system that is to be protected receives from a policy server security policies for the enforcement engines responsible for enforcing the security policies on the host computer system. A security policy has an execution scope and associated rules. An execution scope defines the attributes of a process or application to which the associated rules are to be applied. A rule provides a condition and action to be performed when the condition is satisfied. A rule also has a rule type (also referred to as a security type) that is used by the firewall agent to identify the enforcement engines that are responsible for enforcing the rules. The firewall agent may distribute the security policies to enforcement engines executing in user mode and in kernel mode. The enforcement engines may provide a layered approach to a distributed firewall. For example, enforcement engines may be installed on a host computer system to prevent intrusions at various layers of the ISO 7-layer reference model. The network layer, the transport layer, and the session layer may each have an enforcement engine that is developed to detect attempted intrusions at that layer. To distribute the security policies that have been received at a host computer system, the firewall agent identifies to which enforcement engine a rule applies based in part on rule type. The firewall agent then distributes the rule to the identified enforcement engine, which then enforces the rule. Upon receiving a rule, an enforcement engine may distribute the rule to its subcomponents as appropriate. For example, the firewall agent may distribute all rules destined for kernel-mode enforcement engines to a single kernel-mode security component that then distributes the rules to the enforcement engines. In this way, the distributed firewall system provides a single mechanism for distributing rules of security policies to the enforcement engines providing different layers of protection for a host computer system.

[0018] In one embodiment, a security policy includes rules that specify conditions, actions, and optionally exceptions. For example, a rule may indicate that an application may not receive messages on any network port other than port 80. The condition of the rule may be satisfied when a message is received on a network port

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 83

other than port 80, and the action may result in the message being discarded. As another example, a rule may indicate that when an application attempts to send a message on network port 80, authorization is to be solicited from a user of the computing system on which the rule is enforced. The rules may be categorized into rule or security types based on their behavior. For example, rules with a network security type may be directed to security enforcement for network traffic. Each security type may have its own security component for enforcing rules of that security type. Languages for specifying security policies is described in U.S. Patent Application No. 10/882,438, entitled "Languages for Expressing Security Policies" and filed on July 1, 2004, which is hereby incorporated by reference.

[0019] The condition, action, and exception of the rules may be specified as expressions. A rule may be semantically expressed as "IF conditions THEN actions EXCEPT exceptions." Conditions of a rule are expressions of circumstances under which security enforcement actions of the rule are to be performed. An action is an expression of activity to be performed when the condition is satisfied. A rule may have multiple actions. An exception is an expression of when the actions may not be performed even though the condition is satisfied. A condition may be either static or dynamic. A static condition is one which refers, for example, to a hard-coded list of files. A dynamic condition is one which, for example, performs a query to determine a list of files matching a provided criterion. Security enforcement actions may include allowing a request that caused the security enforcement event, denying the request, soliciting input from a user, notifying the user, and so on. Other rule constructs are also possible. As an example, an "else" construct could be added to perform alternate actions when a condition is false.

[0020] Figure 1 is a block diagram that illustrates the distributed firewall system in one embodiment. An enterprise may have a network system 100 that includes a policy server computer system 110 and host computer systems 120. A system administrator uses a policy server component 111 of the policy server computer system to prepare and distribute security policies stored in a security policy store 112 to the host computer systems. Each host computer system includes a firewall agent 121, enforcement engines 122, and a policy store 123. The firewall agent provides a central mechanism through which a host computer system receives and distributes security policies to the enforcement engines of that host computer system. Upon receiving a security policy from the policy server computer system, the firewall agent of a host computer system stores the security policy in the policy store. The firewall agent then processes the security policies on a rule-by-rule basis and provides the rules to the appropriate enforcement engine. The firewall agent may combine various security policies, order the rules of the security policies, and dynamically notify the enforcement engines of changes to security policies or system configurations that will affect the enforcement of the security policies. Such combining, ordering, and notifying is described in U.S. Patent Application No. 10/966,800, entitled "Method and System for Merging Security Policies" and filed on October 14, 2004, which is hereby incorporated by reference. The enforcement engines then enforce the rules by applying them to network events. The enforcement engines may provide various type of security enforcement. An enforcement engine may detect an attempt to perform a certain behavior and block the attempt. For example, an application that is infected with a virus may attempt to delete certain files of the operating system. A security engine executing in the same process space as the application may detect and block the attempt to delete the file. A security engine for performing behavior blocking is described in U.S. Patent Application No. 10/832,798, entitled "A Method and System for Enforcing a Security Policy via a Security Virtual Machine" and filed on April 27, 2004, which is hereby incorporated by reference. The enforcement engines may be adapted to apply rules to network events of various protocols, such as IP, ICMP, TCP, FTP, DNS, HTTP, RPC, and so on. Techniques preventing exploitation of vulnerabilities are described in U.S. Patent Application No. 10/955,963, entitled "Method and System for Filtering Communications to Prevent Exploitation of a Software Vulnerability" and filed on September 30, 2004, which is hereby incorporated by reference.

[0021] Figure 2 is a block diagram that illustrates the components of the distributed firewall system executing on a host computer system in one embodiment. The host computer system includes firewall agent components 220, kernel-mode components 230, and application user-mode components 240. The firewall agent components include a receive policy component 221, a parse policy component 222, and a distribute policy component 223. The firewall agent components also include a policy store 224 and a configuration manager 225. The receive policy component receives security policies provided by the policy server computer system and stores the policies in the policy store. The parse policy component retrieves the policies from the policy store and converts them to a lower-level language for processing by the enforcement engines. The distribute policy component receives the parsed security policies and distributes the rules of the security policies to the appropriate enforcement engines. The configuration manager provides configuration information for use in determining which rules should be distributed to which enforcement engines based on the current configuration of the host computer system.

[0022] The kernel-mode components include enforcement engines 231, flow manager components 233, and protocol hooks 234. The components may be layered from layer 0 to layer N, which may correspond to various layers of the ISO 7-layer reference model. Each layer may have an associated protocol hook, flow manager, and enforcement engine that is specific to the layer, parses network events specific to that layer, and enforces the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 84

rules provided by the firewall agent to that layer. The rules for each layer are stored in a rule store 232 that may be specific for that layer. The flow manager components are responsible for intercepting network events and invoking the corresponding enforcement engine to enforce the rules. Each flow manager component receives network events from the corresponding protocol hook that is responsible for detecting network events at the corresponding layer. The application user-mode components may be provided as a dynamic link library that links into an application for providing layered security in a manner that is similar in overall architecture to the kernel-mode components.

[0023] The computing device on which the distribution system is implemented may include a central processing unit, memory, input devices (e.g., keyboard and pointing devices), output devices (e.g., display devices), and storage devices (e.g., disk drives). The memory and storage devices are computer-readable media that may contain instructions that implement the distribution system. In addition, the data structures and message structures may be stored or transmitted via a data transmission medium, such as a signal on a communications link. Various communication links may be used, such as the Internet, a local area network, a wide area network, a point-to-point dial-up connection, a cell phone network, and so on.

[0024] The distribution system may be implemented in various operating environments that include personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, programmable consumer electronics, digital cameras, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and so on. The host computer systems may be cell phones, personal digital assistants, smart phones, personal computers, programmable consumer electronics, digital cameras, and so on.

[0025] The distribution system may be described in the general context of computer-executable instructions, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, objects, components, data structures, and so on that perform particular tasks or implement particular abstract data types. Typically, the functionality of the program modules may be combined or distributed as desired in various embodiments.

[0026] Figure 3 is a flow diagram that illustrates the processing of the distribute policy component of the firewall agent in one embodiment. Upon receiving policies from the parse policy component, the component identifies the enforcement engines to which the rules apply and distributes the rules to the identified enforcement engines. In blocks 301-307, the component loops selecting the security policies. In block 301, the component selects the next security policy. In decision block 302, if all the security policies have already been selected, then the component completes, else the component continues at block 303. In blocks 303-307, the component loops processing each rule within the selected security policy. In block 303, the component selects the next rule of the selected security policy. In decision block 304, if all the rules of the selected security policy have already been selected, then the component loops to block 301 to select the next security policy, else the component continues at block 305. In block 305, the component identifies the rule type (or security type) of the selected security rule. In block 306, the component identifies the enforcement engine that is responsible for enforcing the rules of the identified rule type. In block 307, the component provides the selected rule to the identified component and then loops to block 303 to select the next rule of the selected security policy.

[0027] Figure 4 is a flow diagram that illustrates the processing of the flow manager component in one embodiment. The flow manager component is passed a network event and invokes the appropriate enforcement engine for processing the network event. The flow manager component may be implemented as a separate flow manager for each layer that processes network events only for that layer or may be implemented as a single flow manager component that process network events for all layers. In this illustrated embodiment, the flow manager component processes all network events for all layers. In decision blocks 401-403, the component identifies the layer to which a network event applies. The component then invokes the associated enforcement engine for that layer in blocks 411-413. The component then completes. The flow manager may maintain a separate flow for each network connection. The enforcement engines and flow managers may be specific to the layer and the protocol implemented by the layer. For example, when a connection uses the Transmission Control Protocol ("TCP") or the User Datagram Protocol ("UDP") at the packet layer, the flow manager invokes the enforcement engine that is appropriate for the protocol. The invoked enforcement engines may instantiate data structures for tracking state information from one invocation to the next. Each application that is protected by user-mode enforcement engine may also have a flow manager with a separate flow for each connection.

[0028] Figure 5 is a flow diagram that illustrates the processing of an enforcement engine for a specific layer in one embodiment. The enforcement engine, which is a component of the distributed firewall system, is provided a network event, identifies the rules to apply to the event, determines whether the conditions of the rules are satisfied, and if so, performs the action associated with the rule. In blocks 501-504, the component loops selecting each rule associated with the layer of the enforcement engine. In block 501, the component selects the next rule. In decision block 502, if all the rules have already been selected, then the component completes, else the component continues at block 503. In decision block 503, if the condition of the rule is satisfied, then the component continues at block 504, else the component loops to block

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 85

501 to select the next rule. In block 504, the component performs the action associated with the rule and then loops to block 501 to select the next rule. Each enforcement engine may be adapted to process the network events and rules that are appropriate to its layer. For example, an enforcement engine for the network layer may parse network packets and enforce rules related to the network packets. An enforcement engine for the application layer may detect unauthorized attempts to access or modify files relating to that application.

[0029] From the foregoing, it will be appreciated that specific embodiments of the distributed firewall system have been described herein for purposes of illustration, but that various modifications may be made without deviating from the spirit and scope of the invention. Accordingly, the invention is not limited except as by the appended claims.

Claims

1. A method in a computer system for distributing rules of security policies to enforcement engines for enforcing the security policies, the method comprising:

   providing at the computer system enforcement engines that implement different layers of security enforcement;
   receiving and storing at the computer system security policies having rules, each rule having a rule type;
   under control of a firewall agent executing on the computer system,
   retrieving the stored security policies; and
   for rules of a retrieved security policy,
   identifying an enforcement engine to which a rule applies based on the rule type of the rule; and
   providing the rule to the identified enforcement engine; and under control of the enforcement engines, enforcing the rules provided to the enforcement engine by the firewall agent,
   wherein the firewall agent provides a mechanism for distributing the rules to multiple enforcement engines of the computer system.

2. The method of claim 1 wherein the security policies are received from policy computer system that distributes the security policies for multiple computer systems.

3. The method of claim 1 wherein the enforcement engines provide a layered firewall.

4. The method of claim 1 wherein the enforcement engines include a behavior blocking security component.

5. The method of claim 1 wherein an enforcement engine includes subcomponents and wherein upon being provided with a rule, the enforcement engine provides the rule to a subcomponent.

6. The method of claim 1 wherein at least one enforcement engine executes in kernel mode and at least one enforcement engine executes in user mode.

7. The method of claim 1 wherein the firewall agent includes a user-mode subcomponent and a kernel-mode subcomponent, and the user-mode subcomponent distributes rules to the kernel-mode subcomponent for providing the rules to enforcement engines that execute in kernel mode.

8. The method of claim 1 wherein a flow manager component intercepts network events at the various layers and invokes an enforcement engine associated with a layer to enforce the rules for that layer.

9. The method of claim 8 wherein a single flow manager component handles network events from multiple layers,

10. The method of claim 8 wherein a flow manager component handles network events from only one layer.

11. The method of claim 1 wherein enforcement engines provide for security at various layers of a network protocol stack.

12. The method of claim 11 wherein layers of different enforcement engines are used for different connections.

13. A computer system architecture for distributing security policies, the architecture comprising:

   a policy server computer system for establishing security policies, the security policies having rules with rule types, and for providing the security policies to host computer systems for enforcement of the security policies at the host computer systems; and
   multiple host computer systems that include:

      enforcement engines for enforcing a type of security at the host computer system; and
      a firewall agent that receives the security policies from the policy server computer system, identifies the enforcement engine to which the rules apply based on rule type and enforcement engine type, and provides the rules to the identified enforcement engines for enforcement of the security policies.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 86

14. The computer system architecture of claim 13 wherein the enforcement engines are layered so that when one enforcement engine determines that the rules provided to it do not to apply to a network event, another enforcement engine determines whether the rules provided to it apply to the network event.

15. The computer system architecture of claim 14 wherein the layers include a protocol interception hook and an enforcement engine.

16. The computer system architecture of claim 13 wherein the firewall agent distributes rules to enforcement engines executing in user mode.

17. The computer system architecture of claim 13 wherein the firewall agent distributes rules to enforcement engines executing in kernel mode.

18. The computer system architecture of claim 13 wherein the enforcement engines are layered and including a flow manager component that handles network events at each layer and invokes an enforcement engine of the layer to enforce the rules of that layer on the network event.

19. The computer system architecture of claim 18 wherein the flow manager component handles network events for multiple layers.

20. The computer system architecture of claim 18 wherein the flow manager component handles network events for a single layer.

21. A host computer system for distributing rules of security policies to enforcement engines for enforcing the security policies of the host computer system, comprising:

multiple enforcement engines that implement different layers of firewall security enforcement at the host computer system by receiving and enforcing rules of security policies;
a component that receives at the host computer system security policies having rules, each rule having a rule type; and
a component that identifies enforcement engines to which a rule applies based on the rule type of the rule and that provides the rule to the identified security component;
wherein a mechanism is provided for distributing the rules to the layered enforcement engines of the host computer system.

22. The computer system of claim 21 wherein the security policies are received from a policy computer system that distributes the security policies for multiple host computer systems.

23. The computer system of claim 21 wherein the enforcement engine includes a behavior blocking security component.

24. The computer system of claim 21 wherein an enforcement engine includes subcomponents and wherein upon being provided with a rule, the enforcement engine provides the rule to a subcomponent.

25. The computer system of claim 21 wherein at least one enforcement engine executes in kernel mode and at least one enforcement engine executes in user mode.

26. The computer system of claim 21 wherein the component that provides the rules to the enforcement engines includes a user-mode subcomponent and a kernel-mode subcomponent, and the user-mode subcomponent distributes rules to the kernel-mode subcomponent for providing the rules to enforcement engine that execute in kernel mode.

27. The computer system of claim 21 including a flow manager component that intercepts network events at the various layers and invokes an enforcement engine associated with a layer to enforce the rules for that layer.

28. The computer system of claim 27 wherein a single flow manager component handles network events from multiple layers.

29. The computer system of claim 27 wherein a flow manager component handles network events from only one layer.

30. A computer-readable medium containing instructions for controlling a host computer system to distribute security policies to enforcement engines for enforcing the security policies, by a method comprising:

receiving and storing at the host computer system security policies having rules, each rule having a rule type;
under control of a firewall agent executing on the host computer system,
retrieving the stored security policies; and
for rules of a retrieved security policy,
identifying an enforcement engine to which a rule applies based on the rule type of the rule; and
providing the rule to the identified enforcement engine; and under control of the enforcement engines, enforcing the rules provided to the enforcement engines by the distribution component.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 87

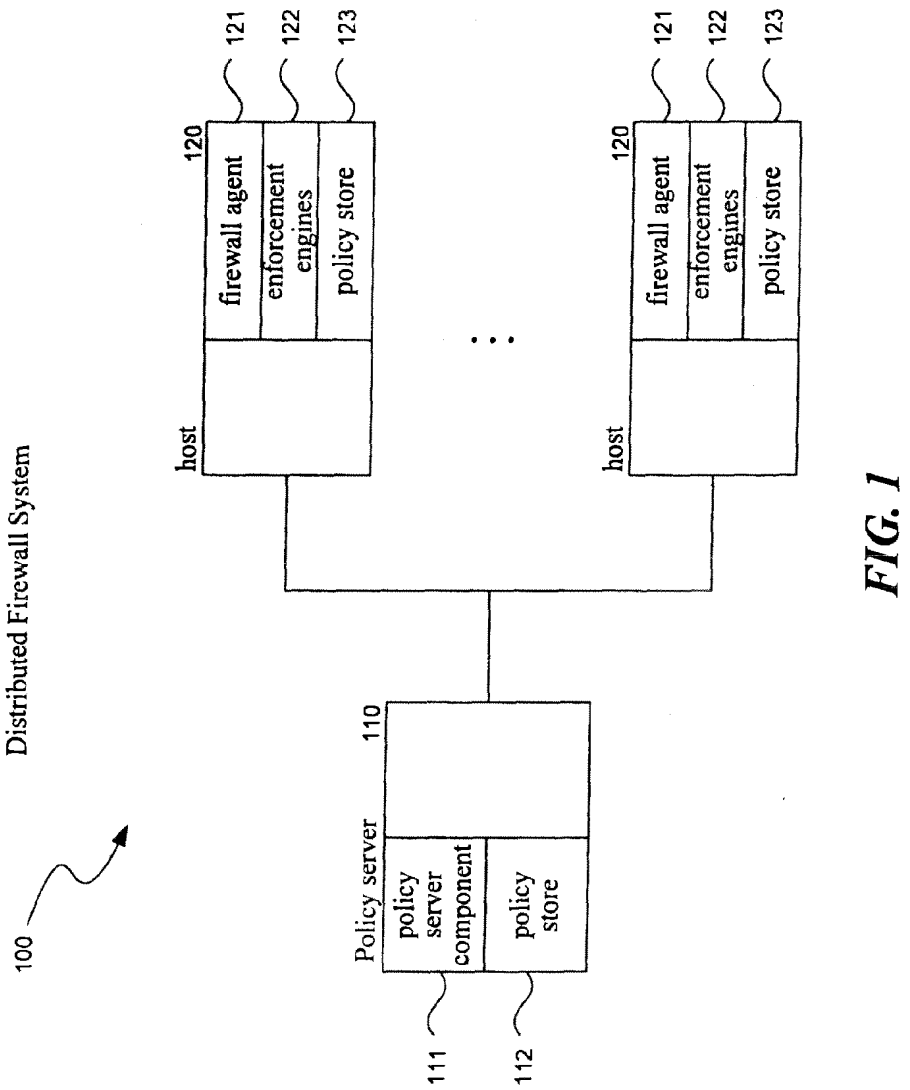

FIG. 1

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 88

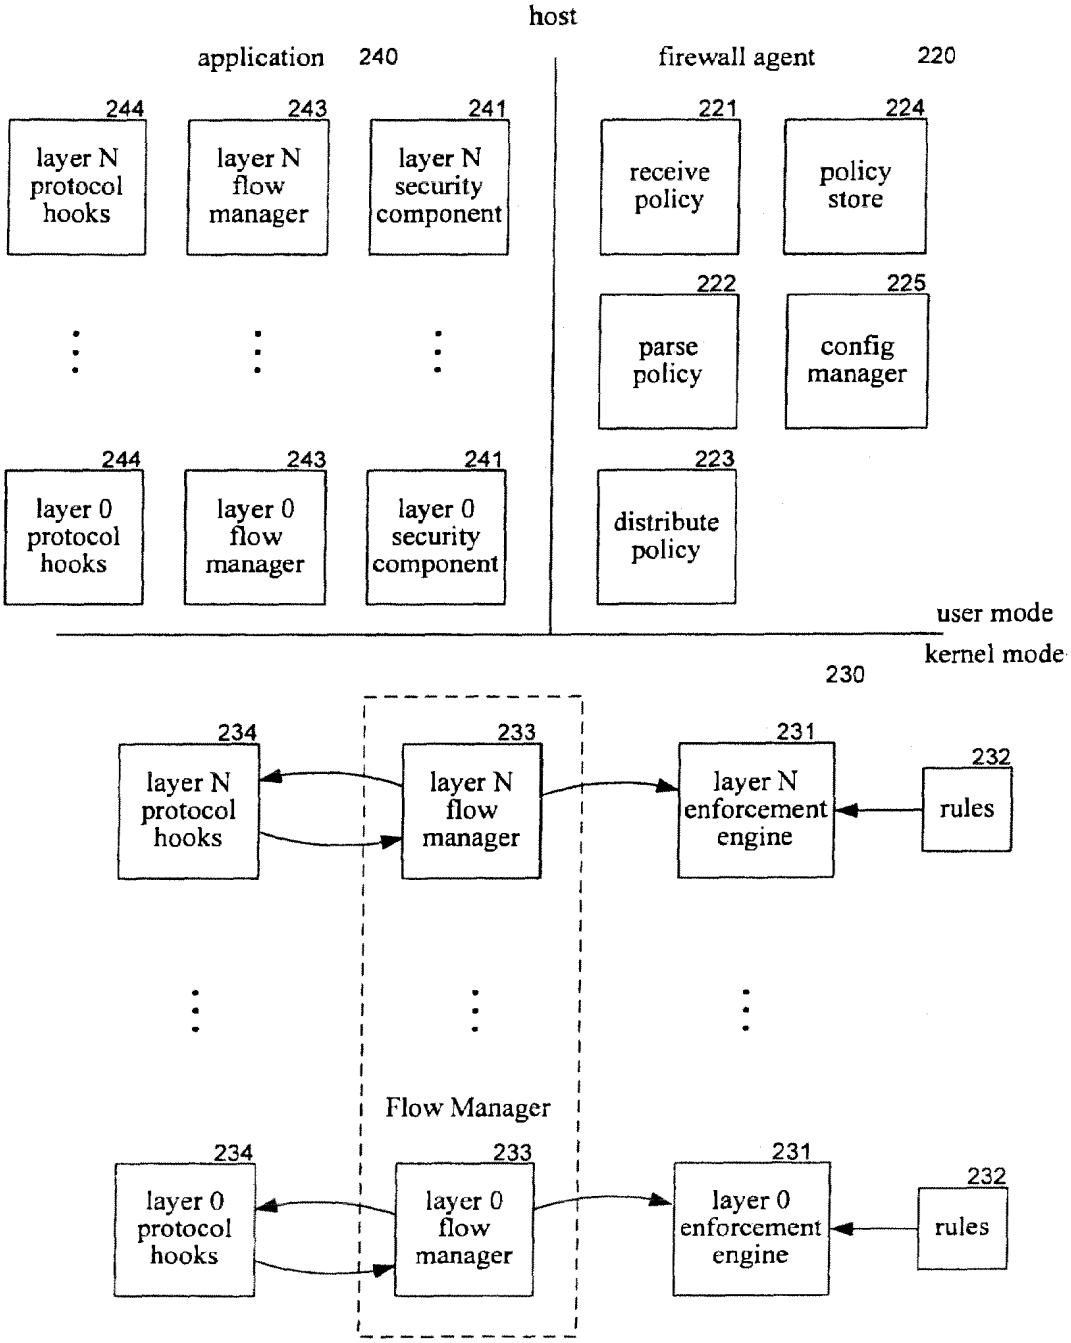

FIG. 2

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 89

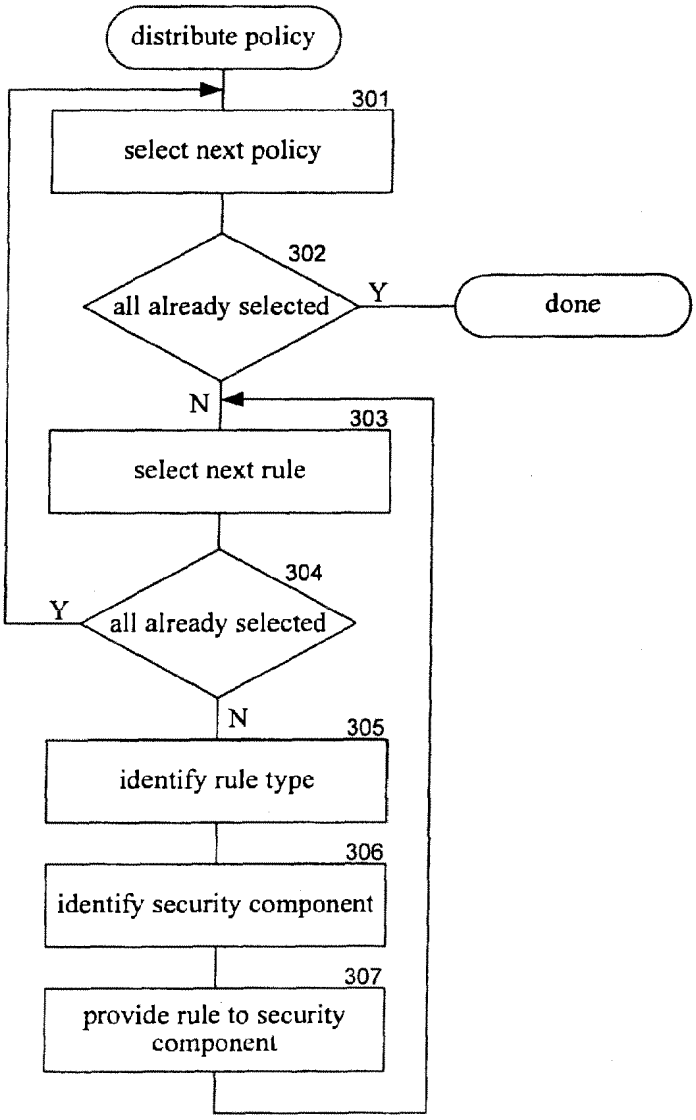

FIG. 3

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 90

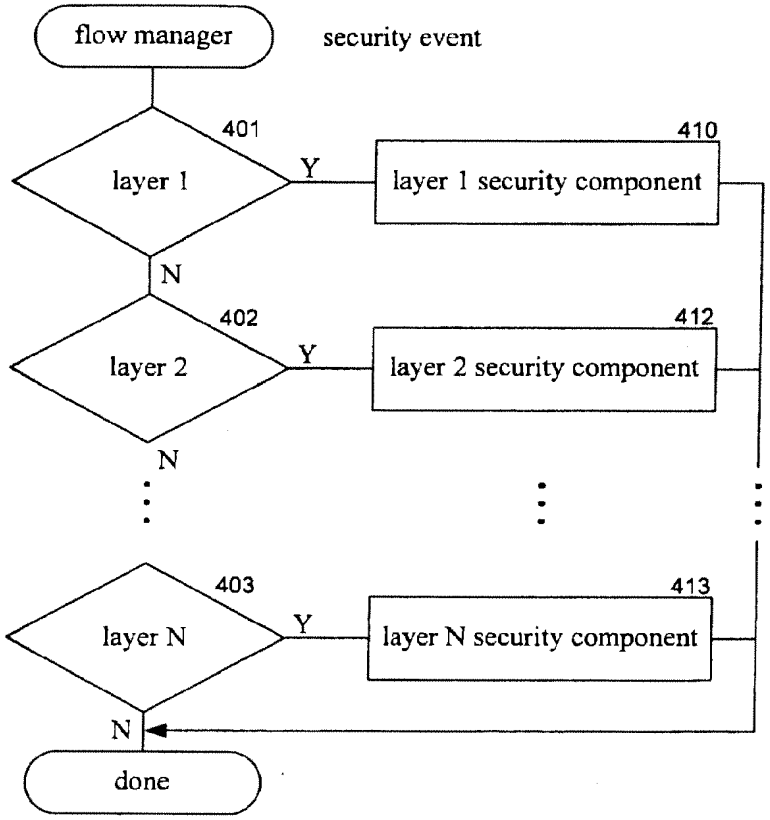

FIG. 4

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 91

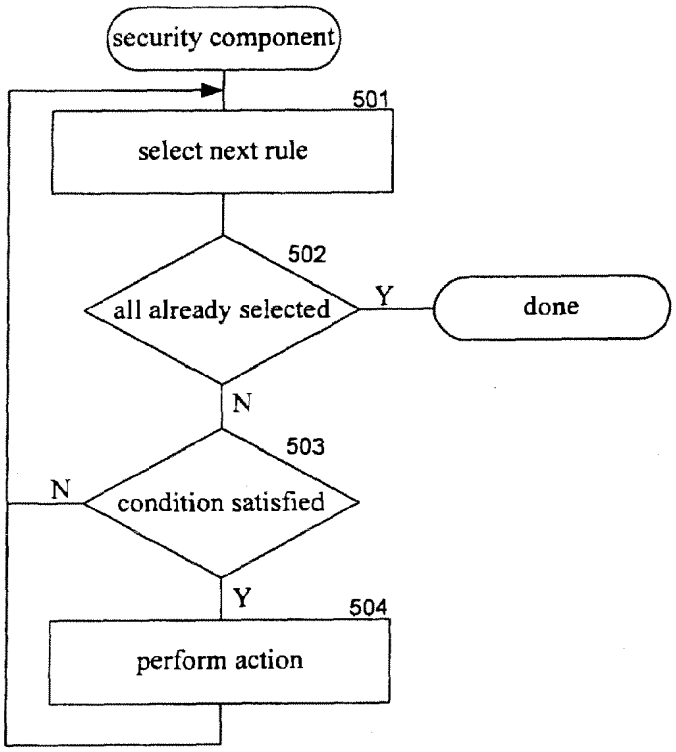

FIG. 5

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 92

Electronic Acknowledgement Receipt

EFS ID:	18782948
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten/Rocko Bund
Filer Authorized By:	William Eugene Wooten
Attorney Docket Number:	007742.00029
Receipt Date:	16-APR-2014
Filing Date:	
Time Stamp:	16:51:09
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Non Patent Literature	NPL-Kindervag_16712444.PDF	2836187 7320c5df93a3c25807d980d1ff9302ac28250a18	no	27

Warnings:	
Information:	

2	Non Patent Literature	NPL-PaloAltoNetworks_16712439.PDF	2240313 1a612b6fe32ae1fdcce919ae95025d31865f5c76	no	11
Warnings:					
Information:					
3	Non Patent Literature	ISR_to_Cite_in_IDS_16477204.PDF	879049 0c802d5b79a6d7cc80dc33c3685743441eceec44	no	8
Warnings:					
Information:					
4	Information Disclosure Statement (IDS) Form (SB08)	Information_Disclosure_Statement_Fillable_PDF_16712455.PDF	612510 0691b66628f0f730dba2ebcadfdaed15a578d979	no	4
Warnings:					
Information:					
5	Foreign Reference	EP1677484A2_16712298.PDF	773874 47dd5a5c781f0fd2a220829489e9b07dca656db2	no	14
Warnings:					
Information:					
		Total Files Size (in bytes):	7341933		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

PCT

NOTIFICATION OF TRANSMITTAL OF THE INTERNATIONAL SEARCH REPORT AND THE WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

To:
Wright, Bradley
BANNER & WITCOFF, LTD.
1100 13th Street, N.W. Suite 1200
Washington, DC 20005-4051
ETATS-UNIS D'AMERIQUE

Date of mailing *(day/month/year)*	7 November 2013 (07-11-2013)
Applicant's or agent's file reference 007742.00023	**FOR FURTHER ACTION** See paragraphs 1 and 4 below
International application No. PCT/US2013/057502	International filing date *(day/month/year)* 30 August 2013 (30-08-2013)
Applicant CENTRIPETAL NETWORKS, INC.	

1. [X] The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

   Filing of amendments and statement under Article 19:
   The applicant is entitled, if he so wishes, to amend the claims of the International Application (see Rule 46):

   When? The time limit for filing such amendments is normally two months from the date of transmittal of the International Search Report.

   Where? Directly to the International Bureau of WIPO, 34 chemin des Colombettes 1211 Geneva 20, Switzerland, Fascimile No.: (41-22) 338.82.70

   For more detailed instructions, see *PCT Applicant's Guide*, International Phase, paragraphs 9.004 - 9.011.

2. [ ] The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.

3. [ ] **With regard to any protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

   [ ] the protest together with the decision thereon has been transmitted to the International Bureau together with any request to forward the texts of both the protest and the decision thereon to the designated Offices.

   [ ] no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. **Reminders**

   The applicant may submit comments on an informal basis on the written opinion of the International Searching Authority to the International Bureau. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established. Following the expiration of 30 months from the priority date, these comments will also be made available to the public.

   Shortly after the expiration of **18 months** from the priority date, the international application will be published by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau before completion of the technical preparations for international publication (Rules 90*bis*.1 and 90*bis*.3).

   Within **19 months** from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase **until 30 months** from the priority date (in some Offices even later); otherwise, the applicant must, **within 20 months** from the priority date, perform the prescribed acts for entry into the national phase before those designated Offices.

   In respect of other designated Offices, the time limit of **30 months** (or later) will apply even if no demand is filed within 19 months.

   For details about the applicable time limits, Office by Office, see www.wipo.int/pct/en/texts/time_limits.html and the *PCT Applicant's Guide*, National Chapters.

Name and mailing address of the International Searching Authority European Patent Office, P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk Tel. (+31-70) 340-2040 Fax: (+31-70) 340-3016	Authorized officer VINCI, Rossana Tel: +31 (0)70 340-2812

NOV 12 2013

Form PCT/ISA/220 (July 2010)

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference	FOR FURTHER ACTION	see Form PCT/ISA/220 as well as, where applicable, item 5 below.
007742.00023		
International application No.	International filing date *(day/month/year)*	(Earliest) Priority Date *(day/month/year)*
PCT/US2013/057502	30 August 2013 (30-08-2013)	22 October 2012 (22-10-2012)

Applicant

CENTRIPETAL NETWORKS, INC.

This international search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 3 sheets.

[X] It is also accompanied by a copy of each prior art document cited in this report.

1. **Basis of the report**

a. With regard to the **language,** the international search was carried out on the basis of:

[X] the international application in the language in which it was filed

[ ] a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))

b. [ ] This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6*bis*(a)).

c. [ ] With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.

2. [ ] **Certain claims were found unsearchable** (See Box No. II)

3. [ ] **Unity of invention is lacking** (see Box No III)

4. With regard to the **title,**

[X] the text is approved as submitted by the applicant

[ ] the text has been established by this Authority to read as follows:

5. With regard to the **abstract,**

[X] the text is approved as submitted by the applicant

[ ] the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority

6. With regard to the **drawings,**

a. the figure of the **drawings** to be published with the abstract is Figure No. 1

[ ] as suggested by the applicant

[X] as selected by this Authority, because the applicant failed to suggest a figure

[ ] as selected by this Authority, because this figure better characterizes the invention

b. [ ] none of the figures is to be published with the abstract

NOV 12 2013
CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 96

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2013/057502

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L29/06
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 677 484 A2 (MICROSOFT CORP [US]) 5 July 2006 (2006-07-05) paragraphs [0011], [0017] - [0019]	1-20
X	US 6 226 372 B1 (BEEBE TODD [US] ET AL) 1 May 2001 (2001-05-01) column 3, line 38 - column 4, line 32; claims 1-11	1-20

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
30 October 2013	07/11/2013

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Veen, Gerardus

1

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2013/057502

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 1677484 A2	05-07-2006	CN 1777179 A	24-05-2006
		EP 1677484 A2	05-07-2006
		JP 4914052 B2	11-04-2012
		JP 2006146891 A	08-06-2006
		KR 20060056231 A	24-05-2006
		US 2006129808 A1	15-06-2006
		US 2009172774 A1	02-07-2009
US 6226372 B1	01-05-2001	AU 1950301 A	18-06-2001
		US 6226372 B1	01-05-2001
		US 2001014150 A1	16-08-2001
		WO 0143343 A1	14-06-2001

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2013/057502	30.08.2013	22.10.2012

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

NOV 12 2013

Name and mailing address of the ISA:	Date of completion of this opinion	Authorized Officer
European Patent Office P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk - Pays Bas Tel. +31 70 340 - 2040 Fax: +31 70 340 - 3016	see form PCT/ISA/210	Veen, Gerardus Telephone No. +31 70 340-3811

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed

   ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43bis.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing filed or furnished:

   a. (means)

      ☐ on paper

      ☐ in electronic form

   b. (time)

      ☐ in the international application as filed

      ☐ together with the international application in electronic form

      ☐ subsequently to this Authority for the purposes of search

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	
	No:	Claims	1-20
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-20
Industrial applicability (IA)	Yes:	Claims	1-20
	No:	Claims	

2. Citations and explanations

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1 The following documents are referred to:

D1 EP 1 677 484 A2 (MICROSOFT CORP [US]) 5 July 2006 (2006-07-05)

D2 US 6 226 372 B1 (BEEBE TODD [US] ET AL) 1 May 2001 (2001-05-01)

2 The present application does not meet the criteria of Article 33(2) PCT, because the subject-matter of claims 1-20 is not new.

2.1 In the words of independent claim 1, D1 discloses:

'A method (column 4, line 13), comprising:

at each of one or more packet security gateways (c.4, l.16, 'enforcement engine') associated with a security policy management server (c.4, l.15, 'policy server'):

receiving a dynamic security policy from the security policy management server (c.4, l.16-19);

receiving packets associated with a network protected by each respective packet security gateway; and

performing, on a packet by packet basis, at least one of multiple packet transformation functions specified by the dynamic security policy on the packets (c.4, l.19-20, 'enforcing the security policies'), wherein performing the at least one of multiple packet transformation functions specified by the dynamic security policy on the packets comprises performing at least one packet transformation function other than forwarding or dropping the packets (c.5, l.30-36)'.

Claim 1 is thus not new (Art. 33(2) PCT).

2.2 The same reasoning applies, mutatis mutandis, to independent claims 19 and 20 which are, therefore, also not new.

2.3 Dependent claims 2-18 do not contain any features which, in combination with the features of the claims to which they refer, meet the requirements of the PCT in respect of novelty and inventive step, see D1, D2 and the corresponding citations in the search report.

Doc Code: PA..
Document Description: Power of Attorney

PTO/AIA/82A (07-13)
Approved for use through 11/30/2014. OMB 0651-0051
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMITTAL FOR POWER OF ATTORNEY TO ONE OR MORE REGISTERED PRACTITIONERS

NOTE: This form is to be submitted with the Power of Attorney by Applicant form (PTO/AIA/82B) to identify the application to which the Power of Attorney is directed, in accordance with 37 CFR 1.5, unless the application number and filing date are identified in the Power of Attorney by Applicant form. If neither form PTO/AIA/82A nor form PTO/AIA82B identifies the application to which the Power of Attorney is directed, the Power of Attorney will not be recognized in the application.

Application Number	14/253,992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Title	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
Art Unit	2431
Examiner Name	TBA
Attorney Docket Number	007742.00029

SIGNATURE of Applicant or Patent Practitioner

Signature	/William E. Wooten/	Date (Optional)	2014-04-18
Name	William E. Wooten	Registration Number	60,049
Title (if Applicant is a juristic entity)			
Applicant Name (if Applicant is a juristic entity)			

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications. If more than one applicant, use multiple forms.

☑ *Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.131, 1.32, and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Doc Code: PA..

Document Description: Power of Attorney

PTO/AIA/82B(07-12)
Approved for use through 11/30/2014. OMB 0651-0035
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

POWER OF ATTORNEY BY APPLICANT

I hereby revoke all previous powers of attorney given in the application identified in the attached transmittal letter.

☒ I hereby appoint Practitioner(s) associated with the following Customer Number as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A or equivalent): 22907

OR

☐ I hereby appoint Practitioner(s) named below as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A or equivalent):

Name	Registration Number	Name	Registration Number

Please recognize or change the correspondence address for the application identified in the attached transmittal letter to:

☒ The address associated with the above-mentioned Customer Number.

OR

☐ The address associated with Customer Number:

OR

☐ Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

I am the Applicant:

☐ Inventor or Joint Inventor

☐ Legal Representative of a Deceased or Legally Incapacitated Inventor

☒ Assignee or Person to Whom the Inventor is Under an Obligation to Assign

☐ Person Who Otherwise Shows Sufficient Proprietary Interest (*e.g.*, a petition under 37 CFR 1.46(b)(2) was granted in the application or is concurrently being filed with this document)

SIGNATURE of Applicant for Patent

Signature		Date	July 19 2013
Name	Steven Rogers	Telephone	(571) 252-5082
Title and Company	Chief Executive Officer, Centripetal Networks, Inc.		

NOTE: Signature - This form must be signed by the applicant in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. Submit multiple forms for more than one signature, see below *.

☒ *Total of 1 forms are submitted.

This collection of information is required by 37 CFR 1.31, 1.32 and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Electronic Acknowledgement Receipt

EFS ID:	18800409
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	18-APR-2014
Filing Date:	
Time Stamp:	10:42:18
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Power of Attorney	POA.pdf	912728 cb9385638145bdf8ca203d401d5db7cd224d78c9	no	2

Warnings:	
Information:	

Total Files Size (in bytes):	912728

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re the Application of:		Atty. Docket No.:	007742.00029
Steven ROGERS *et al.*			
Serial No.:	14/253,992	Group Art Unit:	2431
Filed:	April 16, 2014	Examiner:	TBA
For:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	Confirmation No.:	5163

PRELIMINARY AMENDMENT

Mail Stop Amendment
Commissioner for Patents
Randolph Building
401 Dulany Street
Alexandria, VA 22314

Dear Sir:

This paper is a preliminary amendment.

Please enter the following:

Amendments to the Specification begin on page 2 of this paper.

Remarks begin at page 3 of this paper.

If any fees are required or if an overpayment is made, the Commissioner is authorized to debit or credit Deposit Account No. 19-0733, accordingly.

AMENDMENTS TO THE SPECIFICATION

Please delete the section heading captioned "CROSS-REFERENCE TO RELATED APPLICATIONS" and paragraph **[01]** from the specification as illustrated below.

~~CROSS-REFERENCE TO RELATED APPLICATIONS~~

~~This application is a continuation-in-part of and claims priority to U.S. Patent Application Serial No. 13/657,010, filed October 22, 2012, and entitled "METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK," the disclosure of which is incorporated by reference herein in its entirety.~~

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 108

REMARKS

By the instant preliminary amendment, the specification is amended to remove a section describing related applications and priority claims. Applicant submits an amended application data sheet (ADS) reflecting the same herewith. Applicant respectfully requests entry, consideration of the application, and allowance of all pending claims.

If any fees are required or an overpayment is made, the Commissioner is authorized to debit or credit Deposit Account No. 19-0733, accordingly.

Respectfully submitted,

BANNER & WITCOFF, LTD.

Date: April 22, 2014

By: /William E. Wooten/
William E. Wooten
Registration No. 60,049

Banner & Witcoff, Ltd.
1100 13th Street NW, Suite 1200
Washington, DC 20005-4051
Telephone: (202) 824-3000
Facsimile: (202) 824-3001

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 109

PTO/AIA/14 (12-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Steven		Rogers	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	Leesburg	State/Province	VA	Country of Residence i	US

Mailing Address of Inventor:

Address 1	c/o Centripetal Networks, Inc.		
Address 2	11720 Sunrise Valley Drive		
City	Reston	State/Province	VA
Postal Code	20191	Country i	US

Inventor 2 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Sean		Moore	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	Hollis	State/Province	NH	Country of Residence i	US

Mailing Address of Inventor:

Address 1	c/o Centripetal Networks, Inc.		
Address 2	11720 Sunrise Valley Drive		
City	Reston	State/Province	VA
Postal Code	20191	Country i	US

Inventor 3 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	David	K.	Ahn	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

City	Winston-Salem	State/Province	NC	Country of Residence i	US

Mailing Address of Inventor:

Address 1	c/o Centripetal Networks, Inc.		
Address 2	11720 Sunrise Valley Drive		
City	Reston	State/Province	VA
Postal Code	20191	Country i	US

Inventor 4 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Peter	P.	Geremia	

Residence Information (Select One) (•) US Residency ( ) Non US Residency ( ) Active US Military Service

City	Portsmouth	State/Province	NH	Country of Residence i	US

Mailing Address of Inventor:

Address 1	c/o Centripetal Networks, Inc.		
Address 2	11720 Sunrise Valley Drive		
City	Reston	State/Province	VA
Postal Code	20191	Country i	US

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. Add

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

☐ **An Address is being provided for the correspondence Information of this application.**

Customer Number	22907		
Email Address		Add Email	Remove Email

Application Information:

Title of the Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK		
Attorney Docket Number	007742.00029	Small Entity Status Claimed	☒
Application Type	Nonprovisional		
Subject Matter	Utility		
Total Number of Drawing Sheets (if any)	13	Suggested Figure for Publication (if any)	1

Filing By Reference :

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

Only complete this section when filing an application by reference under 35 U.S.C. 111(c) and 37 CFR 1.57(a). Do not complete this section if application papers including a specification and any drawings are being filed. Any domestic benefit or foreign priority information must be provided in the appropriate section(s) below (i.e., "Domestic Benefit/National Stage Information" and "Foreign Priority Information").

For the purposes of a filing date under 37 CFR 1.53(b), the description and any drawings of the present application are replaced by this reference to the previously filed application, subject to conditions and requirements of 37 CFR 1.57(a).

Application number of the previously filed application	Filing date (YYYY-MM-DD)	Intellectual Property Authority or Country

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)

☐ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number	22907		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) or indicate National Stage entry from a PCT application. Providing this information in the application data sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.
When referring to the current application, please leave the application number blank.

Prior Application Status	~~Pending~~		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
	~~Continuation-in-part of~~	~~13657010~~	~~2012-10-22~~

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button. Add

Foreign Priority Information:

/William E. Wooten/
William E. Wooten, Reg. No. 60,048, 2014-04-23

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

This section allows for the applicant to claim priority to a foreign application. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55(d). When priority is claimed to a foreign application that is eligible for retrieval under the priority document exchange program (PDX) the information will be used by the Office to automatically attempt retrieval pursuant to 37 CFR 1.55(h)(1) and (2). Under the PDX program, applicant bears the ultimate responsibility for ensuring that a copy of the foreign application is received by the Office from the participating foreign intellectual property office, or a certified copy of the foreign priority application is filed, within the time period specified in 37 CFR 1.55(g)(1).

Remove

Application Number	Country	Filing Date (YYYY-MM-DD)	Access Code (if applicable)

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button. Add

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications

This application (1) claims priority to or the benefit of an application filed before March 16, 2013 and (2) also contains, or contained at any time, a claim to a claimed invention that has an effective filing date on or after March 16, 2013.

N/A NOTE: By providing this statement under 37 CFR 1.55 or 1.78, this application, with a filing date on or after March 16, 2013, will be examined under the first inventor to file provisions of the AIA.

/William E. Wooten/
William E. Wooten, Reg. No. 60,049, 2014-04-22

Authorization to Permit Access:

☒ Authorization to Permit Access to the Instant Application by the Participating Offices

PTO/AIA/14 (12-13)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

If checked, the undersigned hereby grants the USPTO authority to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the World Intellectual Property Office (WIPO), and any other intellectual property offices in which a foreign application claiming priority to the instant patent application is filed access to the instant patent application. See 37 CFR 1.14(c) and (h). This box should not be checked if the applicant does not wish the EPO, JPO, KIPO, WIPO, or other intellectual property office in which a foreign application claiming priority to the instant patent application is filed to have access to the instant patent application.

In accordance with 37 CFR 1.14(h)(3), access will be provided to a copy of the instant patent application with respect to: 1) the instant patent application-as-filed; 2) any foreign application to which the instant patent application claims priority under 35 U.S.C. 119(a)-(d) if a copy of the foreign application that satisfies the certified copy requirement of 37 CFR 1.55 has been filed in the instant patent application; and 3) any U.S. application-as-filed from which benefit is sought in the instant patent application.

In accordance with 37 CFR 1.14(c), access may be provided to information concerning the date of filing this Authorization.

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Applicant 1 Remove

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

Clear

(•) Assignee	( ) Legal Representative under 35 U.S.C. 117	( ) Joint Inventor
( ) Person to whom the inventor is obligated to assign.	( ) Person who shows sufficient proprietary interest	

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor :

If the Applicant is an Organization check here. [X]

Organization Name	Centripetal Networks, Inc.

Mailing Address Information:

Address 1	11720 Sunrise Valley Drive		
Address 2			
City	Reston	**State/Province**	VA
Country [i]	US	Postal Code	20191
Phone Number		Fax Number	

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

Email Address	

Additional Applicant Data may be generated within this form by selecting the Add button. Add

Assignee Information including Non-Applicant Assignee Information:

Providing assignment information in this section does not subsitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Assignee 1

Complete this section if assignee information, including non-applicant assignee information, is desired to be included on the patent application publication . An assignee-applicant identified in the "Applicant Information" section will appear on the patent application publication as an applicant. For an assignee-applicant, complete this section only if identification as an assignee is also desired on the patent application publication.

Remove

If the Assignee or Non-Applicant Assignee is an Organization check here. ☐

Prefix	**Given Name**	Middle Name	**Family Name**	Suffix

Mailing Address Information For Assignee including Non-Applicant Assignee:

Address 1			
Address 2			
City		**State/Province**	
Country i		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Assignee or Non-Applicant Assignee Data may be generated within this form by selecting the Add button. Add

Signature:

Remove

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications

Signature	/William E. Wooten/			Date (YYYY-MM-DD)	2014-04-16
First Name	William	Last Name	Wooten	Registration Number	60049

Additional Signature may be generated within this form by selecting the Add button. Add

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	007742.00029
	Application Number	
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	18821718
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	22-APR-2014
Filing Date:	
Time Stamp:	08:57:19
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Preliminary Amendment	PrelimAmend.pdf	105723 54a0b444102dac2b24a57913be24c189738 50606	no	3

Warnings:	
Information:	

2	Application Data Sheet	SuppADS.pdf	2353666 c2e8450d8c20610a66c3964f29802430f9a1e1c1	no	8

Warnings:	
Information:	
This is not an USPTO supplied ADS fillable form	
Total Files Size (in bytes):	2459389

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 14/253,992	Filing Date 04/16/2014	☐ To be Mailed

ENTITY: ☐ LARGE ☒ SMALL ☐ MICRO

APPLICATION AS FILED – PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	

APPLICATION AS AMENDED – PART II

AMENDMENT **04/22/2014**

	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
Total (37 CFR 1.16(i))	* 20	Minus	** 20	= 0	X $40 =	0
Independent (37 CFR 1.16(h))	* 3	Minus	***3	= 0	X $210 =	0
☐ Application Size Fee (37 CFR 1.16(s))						
☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
					TOTAL ADD'L FEE	0

AMENDMENT

	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =	
Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =	
☐ Application Size Fee (37 CFR 1.16(s))						
☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
					TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

LIE
/NICOLE LAWRENCE/

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**
If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number: 14/253,992

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY FEE($)
BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	70		N/A	
SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A	300		N/A	
EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	360		N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	20 minus 20 =	*	x 40 =	0.00	OR		
INDEPENDENT CLAIMS (37 CFR 1.16(h))	3 minus 3 =	*	x 210 =	0.00			
APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			0.00			
MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				0.00			
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	730		TOTAL	

APPLICATION AS AMENDED - PART II

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
AMENDMENT A	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
AMENDMENT B	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest found in the appropriate box in column 1.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING or 371(c) DATE	GRP ART UNIT	FIL FEE REC'D	ATTY.DOCKET.NO	TOT CLAIMS	IND CLAIMS
14/253,992	04/16/2014	2431	730	007742.00029	20	3

CONFIRMATION NO. 5163

22907
BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051

FILING RECEIPT

OC000000068211094

Date Mailed: 05/02/2014

Receipt is acknowledged of this non-provisional patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF APPLICANT, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection. Please verify the accuracy of the data presented on this receipt. **If an error is noted on this Filing Receipt, please submit a written request for a Filing Receipt Correction. Please provide a copy of this Filing Receipt with the changes noted thereon. If you received a "Notice to File Missing Parts" for this application, please submit any corrections to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections**

Inventor(s)

Steven Rogers, Leesburg, VA;
Sean Moore, Hollis, NH;
David K. Ahn, Winston-Salem, NC;
Peter P. Geremia, Portsmouth, NH;

Applicant(s)

Centripetal Networks, Inc., Reston, VA

Assignment For Published Patent Application

Centripetal Networks, Inc., Reston, VA

Power of Attorney: The patent practitioners associated with Customer Number 22907

Domestic Applications for which benefit is claimed - None.
A proper domestic benefit claim must be provided in an Application Data Sheet in order to constitute a claim for domestic benefit. See 37 CFR 1.76 and 1.78.

Foreign Applications for which priority is claimed (You may be eligible to benefit from the **Patent Prosecution Highway** program at the USPTO. Please see http://www.uspto.gov for more information.) - None.
Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

Permission to Access - A proper **Authorization to Permit Access to Application by Participating Offices** (PTO/SB/39 or its equivalent) has been received by the USPTO.

If Required, Foreign Filing License Granted: 05/02/2014

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 14/253,992**

Projected Publication Date: 10/22/2015

Non-Publication Request: No

Early Publication Request: No
**** SMALL ENTITY ****
Title

METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

Preliminary Class

726

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at http://www.uspto.gov/web/offices/pac/doc/general/index.html.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, http://www.stopfakes.gov. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4258).

LICENSE FOR FOREIGN FILING UNDER

Title 35, United States Code, Section 184

Title 37, Code of Federal Regulations, 5.11 & 5.15

GRANTED

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign AssetsControl, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop technology, manufacture products, deliver services, and grow your business, visit http://www.SelectUSA.gov or call +1-202-482-6800.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431
Examiner Name	Not Yet Assigned
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20030123456	A1	2003-07-03	Denz et al.	
	2	20070083924	A1	2007-04-12	Lu	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2431
	Examiner Name	Not Yet Assigned
	Attorney Docket Number	007742.00029

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	International Search Report off International Application No. PCT/US2014/023286, dated June 24, 2014.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2431
	Examiner Name	Not Yet Assigned
	Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2014-06-25
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	19409080
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten/Rocko Bund
Filer Authorized By:	William Eugene Wooten
Attorney Docket Number:	007742.00029
Receipt Date:	25-JUN-2014
Filing Date:	16-APR-2014
Time Stamp:	15:14:45
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Non Patent Literature	NPL-ISR_dated_24-6-14774230_16806852.PDF	519073 8f3879c91d1dd61cd1a7d73ef98e5330996 4ad85	no	9

Warnings:	
Information:	

2	Information Disclosure Statement (IDS) Form (SB08)	Information_Disclosure_Statement_Fillable_PDF_16807037.PDF	612171 146dc2593df60d2fa62c68f55c7240844ef95355	no	4

Warnings:	
Information:	
Total Files Size (in bytes):	1131244

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431
Examiner Name	Not Yet Assigned
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20060195896		2006-08-31	Fulp et al.	
	2	20060248580		2006-11-02	Fulp et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	1313290	EP	A1	2003-05-21	Stonesoft Corp		☐
	2	2005046145	WO	A1	2005-05-19	Ericsson Telefon Ab L M		☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2431
	Examiner Name	Not Yet Assigned
	Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button			Add
NON-PATENT LITERATURE DOCUMENTS			Remove
Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	ISR off International Application No. PCT/US2013/072566, dated March 24, 2014.	☐
If you wish to add additional non-patent literature document citation information please click the Add button			Add

EXAMINER SIGNATURE			
Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431
Examiner Name	Not Yet Assigned
Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2014-06-27
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Espacenet

Bibliographic data: EP1313290 (A1) — 2003-05-21

A personal firewall with location dependent functionality

No documents available for this priority number.

Inventor(s): SYVAENNE TUOMO [FI] ± (SYVAENNE, TUOMO)

Applicant(s): STONESOFT CORP [FI] ± (STONESOFT CORPORATION)

Classification: **- international:** ***H04L29/06;*** (IPC1-7): H04L29/06
- cooperative: **H04L63/0227**; **H04L63/0263**; **H04L63/107**

Application number: EP20020102581 20021114

Priority number (s): US20010988356 20011119

Also published as: EP1313290 (B1) US2003097590 (A1) US7325248 (B2) DE60216218 (T2) AT346446 (T)

Abstract of EP1313290 (A1)

A computer device is provided with a local security mechanism, a personal firewall, for protecting the computer device from attacks from a foreign network, in addition to or instead of a firewall in the internal network which protects the computer when connected to a home network. The personal firewall is provided with different sets of security rules for the home network and foreign networks. The personal firewall is arranged to detect its current location, i.e. determine to which network it is connected to at each particular moment. The personal firewall activates one of the given sets of security rules according to the detected current location of the computer device, i.e. the personal firewall automatically uses the security rules predefined for the network to which the computer device is connected at each particular moment.; Upon detecting a change in the location, the personal firewall immediately adapts to use security rules predefined for the new location.

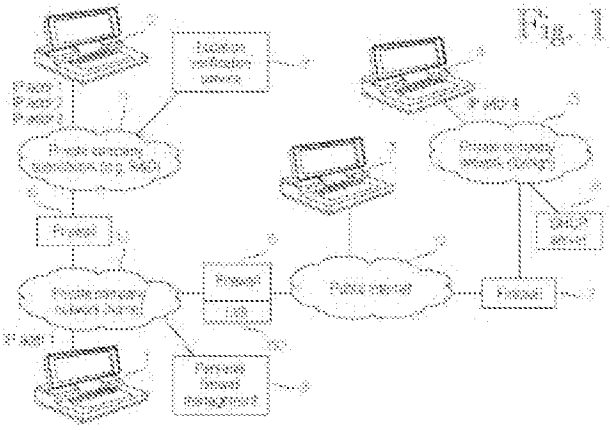

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 313 290 A1

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
21.05.2003 Bulletin 2003/21

(51) Int Cl.[7]: **H04L 29/06**

(21) Application number: **02102581.2**

(22) Date of filing: **14.11.2002**

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR IE IT LI LU MC NL PT SE SK TR
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **19.11.2001 US 988356**

(71) Applicant: **Stonesoft Corporation**
00210 Helsinki (FI)

(72) Inventor: **Syvänne, Tuomo**
Fin-01450, Vantaa (FI)

(74) Representative: **Äkräs, Tapio Juhani**
Kolster Oy Ab,P.O. Box 148,
Iso Roobertinkatu 23
00121 Helsinki (FI)

(54) A personal firewall with location dependent functionality

(57) A computer device is provided with a local security mechanism, a personal firewall, for protecting the computer device from attacks from a foreign network, in addition to or instead of a firewall in the internal network which protects the computer when connected to a home network. The personal firewall is provided with different sets of security rules for the home network and foreign networks. The personal firewall is arranged to detect its current location, i.e. determine to which network it is connected to at each particular moment. The personal firewall activates one of the given sets of security rules according to the detected current location of the computer device, i.e. the personal firewall automatically uses the security rules predefined for the network to which the computer device is connected at each particular moment. Upon detecting a change in the location, the personal firewall immediately adapts to use security rules predefined for the new location.

Fig. 1

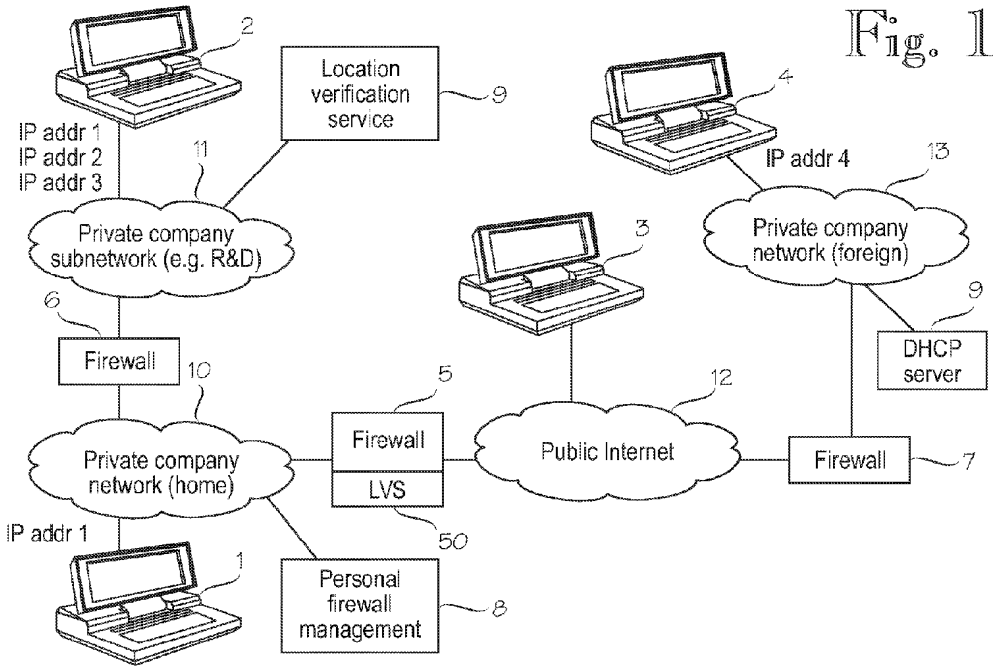

EP 1 313 290 A1

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 137

Description

Field of the Invention

[0001] The present invention relates to network security and, more particularly, to personal firewalls.

Background of the Invention

[0002] Traditionally, a firewall is considered as a set of components forming a gateway between two or more networks. Thus, a firewall has been a gateway which operates at the same time as a connector and a separator between the networks in a sense that the firewall keeps track of the traffic that passes through it from one network to another and restricts connections and packets that are defined as unwanted by the administrator of the system. Physically a firewall is a machine with appropriate software to do the tasks assigned to it. It can be a router, a personal computer (PC), or any other device that can be used for such purposes. Although firewalls are mostly used to connect Local Area Networks (LANs), i.e. internal networks, to the Internet and to protect against attackers or undesired traffic in general, they may also be used to separate and connect different segments of internal network for security purposes. The advantages of having a firewall are numerous. A firewall secures the network and can be used as a tool for monitoring the traffic especially from the outside to the inside of the network guarded by a firewall. Because all traffic intended for the internal network must pass through the firewall, most of the network security actions and policies can be concentrated in this particular point. This is of course a cost and administrative advantage.

[0003] Nowadays, laptop computers and other portable computer devices are widely used. While outside the internal network, the laptop cannot make benefit of the protection provided by the conventional "gateway-type" firewall. Therefore, approaches to improve security of a client located in a foreign network (a public network or an internal network of a foreign organisation) have been proposed. These approaches are based on protecting the laptop itself by means of a local security mechanism, called a personal firewall herein, installed in the laptop (in addition to or instead of a firewall in an internal network, which protects the computers connected to the internal network). The personal firewall may be implemented as software installed in the computer device, or as a separate electronic device connected to the computer device.

[0004] European patent application EP 0 952 715 discloses a firewall security device connected to an external communication port of a computer device. The incoming communications stream to the computer device from e.g. public networks is passed through the firewall security device. The firewall device applies standard security measures, thereby protecting the computer device.

[0005] There is a particular need for such protection by means of a personal firewall if the laptop is allowed to have a remote access, e.g. make a VPN (Virtual Private Network) connection to company network while being connected to a foreign network. In order to improve security of the VPN connections, one prior art solution is to enforce a protection level of a laptop, when a VPN tunnel to a company network is created. This means for example that, during a VPN connection, the IP address forwarding is not allowed, or that any connection attempts to the laptop are denied.

[0006] Clearly this is not enough, since the laptop must be protected as soon as it is connected to a foreign network, not only during a VPN connection. The laptops are often used by non-technical people, which increases the risk of overlooking security aspects. Laptops contain sensitive material, such as customer emails. If a laptop is unprotected, when connected to a foreign network, even for a short period of time, there is a risk of getting infected by a hostile application. Such application can be activated later, when the laptop is connected to an internal network and offer inside help for attacks.

[0007] Thus, there is a need to protect the laptop by means of a personal firewall always when the laptop is connected to a foreign network. However, when the laptop is connected to a company internal network, such personal firewall may unduly prevent some essential traffic. For example, the personal firewall should allow use of a laptop at home (internal) network and access to all services, such as disk-share. In a home network even non-IP protocols are sometimes used. Therefore, it is not feasible to have a personal firewall running at all times, at least not with the same configuration, since the protection needs in an internal network are different from those in a foreign network.

[0008] Some of the current solutions allow changing the set of rules used in the personal firewall, that is, they allow the user of the laptop to use different rule sets when connected to the internal network and when connected to a foreign network. However this is a manual operation. Since manual action is required, there is a high risk that operation is not done. Risk is even higher if the end user does not fully understand the need of a firewall.

Summary of the Invention

[0009] An object of the invention is to improve the security and flexibility of a personal firewall. The object is achieved by the invention disclosed in the attached independent claims. Preferred embodiments are disclosed in the attached dependent claims.

[0010] A computer device which can be connected to a home network (such as an internal network of a company or other organisation where the user is employed) and to a foreign network (such as a public network or an internal network of a foreign organisation) is provided with a local security mechanism, called a personal fire-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 138

wall herein, for protecting the computer device from attacks from a foreign network, in addition to or instead of a firewall in the internal network which protects the computer when connected to the internal network. The personal firewall is provided with different sets of security rules, at least one set of rules for the home network and at least one set of rules for foreign networks. In its simplest form, the set of rules for the home network contains no restrictions for the communication or use of service in the home network. The personal firewall is arranged to detect its current location, i.e. to determine the network to which it is connected at each particular moment. The personal firewall activates one of the given sets of security rules according to the detected current location of the computer device, i.e. the personal firewall automatically uses the security rules predefined for the network to which the computer device is connected at each particular moment. Upon detecting a change in the location, the personal firewall immediately adapts to use security rules predefined for the new location. A benefit of the invention is that the protection of a personal firewall is always enabled at the correct level, depending on the current location. On the other hand, when the computer device is located in the home network, a lower level of protection, or no protection at all, can be automatically provided by the personal firewall, so that the communication and services are not unduly restricted in the home network. Thus, the automated location-dependent management of different sets of rules offers optimal protection in different networks, while not unduly restricting operation in the home network.

[0011] The current location of the computer device is preferably determined on the basis of a currently used IP address of the computer device. This is based on the common practice that a computer device has a different IP address, either a fixed address or a dynamic address, in different networks. The IP address can thereby be utilized for identifying the current network and the location of the computer device.

[0012] However, there are situations where the IP address fails to indicate current location of the computer device. Therefore, in an embodiment of the invention, the current location determined on the basis of the current IP address of the computer device is verified by carrying out an additional location verification procedure with a predetermined network element. In a still further embodiment of the invention, availability of said predetermined network element related to the current IP address is checked. The predetermined network element is such that it responses only if the computer device is located in the network in which it is assumed to be on the basis of the current IP address. If the predetermined network element responses and identifies itself properly, the current location determined based on the current IP address is considered to be verified. Otherwise the computer device determines that the current IP address fails to indicate current location of the computer device. The additional verification process makes it even possible to automatically create a secured tunnel, such as a VPN tunnel to a home network even if the computer device uses the same IP address in the current location as in the internal (home) network. The present invention offers benefits even with stand alone personal firewalls wherein the security rules can be defined locally by the user, although the use of these rules is automated and location-dependent. However, more advantages are achieved when the basic invention is used with a central management of personal firewalls.

[0013] According to an aspect of the invention, security rules are defined, updated and distributed centrally by a centralized rule-based server. Especially the updating of the rules is challenging, because the rule updates must be applied as soon as possible, and therefore the process of updating rules in the personal firewalls must be automated. Updating of rules by push method from the centralized rule base server is not a sufficient option in this case. Use of DHCP (Dynamic Host Configuration Protocol), frequent travelling and the fact that at times the laptop may not be connected to any network makes it next to impossible for the centralized management to initiate contacts with the personal firewalls in the computer devices, because there is no way for the centralized management to know the IP address the computer device is using at a given moment. Therefore, according to an aspect of the invention, the personal firewall is configured to periodically query the availability of updated security rules from the centralized management. The queries should only be made, while the computer device is located in the home network, or optionally, when the computer device has a remote access (e.g. VPN connection) to the home network while being located in a foreign network. In other words, also the updating process is dependent on the current location of the computer device in a similar manner as the selection of the active rules, and similar methods can be utilized for determining the current location.

[0014] According to another aspect of the invention, log files containing information of a status and usage of resources of the computer device are handled in a centralized management location. This enables personnel aware of security aspects to verify whether there have been any attacks against the computer device or not. To that end, the personal firewall sends the log files to the central management, such as to a centralized log server, when the computer device is located in the home network. However, when the computer device is disconnected from the home network, the log files are collected and stored locally in the firewall. In order to enable central handling of the log files, the personal firewall transfers the collected log files to the central log server when such is available. This is performed automatically, whenever the computer device is located in, or optionally, connected to the home network. Again, the handling of the log files in the personal firewall is automated and location dependent in a way similar to the selection of active rules, and similar methods can be used for de-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 139

termining the current location of the computer device.

[0015] The present invention allows use of a computer device in a home (internal) network and access to all services, such as disc-share, and even use of non-IP protocols, which are often denied in foreign networks.

Brief description of the drawings

[0016] Preferred embodiments of the invention will now be described with reference to the attached drawings, in which

Figure 1 is a schematic block diagram of an exemplary network configuration where the present invention can be applied;
Figure 2 shows an exemplary protocol stack of a computer device containing a personal firewall according to the present invention,
Figure 3 illustrates exemplary selection rules and a security rule basis, and the association therebetween,
Figure 4 is a flow diagram illustrating a location-dependent rule base selection according to an embodiment of the invention;
Figures 5 and 6 are flow diagrams illustrating the location verification procedure according to an embodiment of the invention;
Figures 7 and 8 are flow diagrams illustrating the rule base updating according to an embodiment of the invention; and
Figure 9 is a flow diagram illustrating the handling of log files according to an embodiment of the invention.

Preferred embodiments of the invention

[0017] The present invention can be applied in personal firewalls in any computer device which can be moved and connected to different networks. Typically such devices are portable computer devices, such as laptop computers, PDAs, communicators, smart phones, intelligent telecommunication devices, etc. In the following illustrative embodiments of the invention, a laptop computer is used as an example of suitable computer devices.

[0018] Figure 1 shows a schematic block diagram of an exemplary network configuration. The configuration is shown only to facilitate the understanding and description of the present invention. The present invention is not intended to be restricted to any particular network configuration. Further, in order to improve clarity, only network elements which are somehow involved with the present invention are shown in Figure 1.

[0019] As illustrated in Figure 1, private local networks 10 and 13 are coupled to a public network, such as the Internet 12 via firewalls 5 and 7, respectively. Naturally, the coupling between the private networks and the public Internet 12 may include also routers and Internet service providers (ISPs not shown in Figure 1). As is well known in the art, private networks 10 and 13 may be, for example, company networks, such as local area networks (LANs) which connect users and resources, such as workstations, servers, printers and the like of the company. A private internal network may also consist of several sub-networks separated by internal firewalls. In the exemplary network configuration shown in Figure 1, the private company sub-network 11 is connected via a firewall 6 to the private local network 10. Such a sub-network 11 may be, for example, a dedicated network for a specific department of the organisation, such as the research and development (R&D) department which must have a restrictive access and higher protection level compared with other part of the company network. Sub-networks of the company, such as the local networks of organisation headquarters and branch offices may be interconnected by secure connections, such as virtual private network (VPN).

[0020] As already described above, the firewalls 5, 6 and 7 are gateways which operate at the same time as connectors and separators between the networks in a sense that the firewall keeps track of the traffic that passes through it from one network to another and restricts connections and packets that are defined as unwanted by the administrator of the system. Physically a firewall is a machine with appropriate software to perform the task assigned to it. It can be a router, a personal computer (PC), or whatever that can be used for such purposes.

[0021] However, the firewalls between the networks, or the implementations thereof, are not relevant to the present invention.

[0022] The present invention relates to protecting of the computer device, e.g. laptop itself, by means of a local security mechanism, called a personal firewall herein, installed on the laptop in addition to or instead of a firewall in a private network. The personal firewall may be implemented as software installed and run in the computer device, which is a preferred embodiment, or as a separate electronic device connected to the computer device. In Figure 1, the laptops 1, 2, 3 and 4 illustrate laptops provided with a personal firewall.

[0023] Figure 2 illustrates the basic principle of a personal firewall installed in a laptop. Physical and network layers 200 refer to all protocols and physical connections required for transferring protocol data units (PDUs) of the upper layers. The upper layers 200 include applications and any transmission protocols employed, such as Internet protocol (IP) transmission control protocol (TCP), NetPEUI, IPX, etc. Basically the personal firewall protection layer 201 operates in a manner analogue to a firewall between networks. More particularly, the personal firewall protection layer 201 operates at the same time as a connector and a separator between the underlying layers and the upper layer in a sense that the personal firewall keeps track of the traffic that passes through it from underlying layers to the upper layers, and

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 140

vice versa, and restricts connections and packets that are defined as unwanted according to the security rules used. The personal firewall protection layer 201 is implemented or controlled by a personal firewall application 203 run in the laptop. In a preferred embodiment of the invention, the personal firewall application 203 carries out the location detection and the location-dependent functions described below, such as the selection of the active rule base according to the current location of the laptop. However, it should be appreciated that the present invention is not intended to be restricted to any specific practical implementation of the personal firewall.

[0024] In accordance with the principles of the present invention, the personal firewall has different sets of rules for the home network (such as the private company network 10) and foreign network, such as the public Internet 12, or the foreign private network 13, or a network of another department of the company. It is not relevant to the present invention what kind of security rules are applied, but some examples are given in Figure 3. For example, a rule base 301 for the foreign company network may list as allowed connections of protocols: hypertext transfer protocol (http), secured http (Https), domain name service (DNS), single message transfer protocol (SMTP) and a VPN connection with IPsec. In the preferred embodiment of the invention these rules are exclusive, in other words, other protocols and connections are denied and blocked by the personal firewall. For a default network, which may be the public Internet 12, the rule base 302 is similar to the rule base 301, except that the SMTP protocol is no longer allowed. For the home network 10, a rule base 300 is defined. The allowed protocols include, in addition to the http, https and the SMTP, also other transmission protocols, such as NetBEUI and IPX. The rule base 300 also allows a disc-share for predefined servers using NetBIOS. Other protocols and connections are denied. It is also possible that the rule base 300 allows all protocols and connections in the home network. Since the home network is protected by a company firewall, the use of a personal firewall in the home network may be regarded as unnecessary. However, the company firewall gives protection only against attacks from the outside of the home network, and the use of a personal firewall protection may be necessary for protecting against attacks from within the home network.

[0025] The different rule bases could be activated manually by a user. However, according to the basic principle of the present invention, the personal firewall automatically selects and activates the proper rule base according to the current location of the laptop.

[0026] Figure 4 is a flow diagram illustrating the selection of the rule base according to one embodiment of the present invention. The simplest way to determine the current location of the laptop is to do it on the basis of the currently used IP address only. This is possible in the cases where the laptop has a different IP address, either a fixed or a dynamic address, in different networks. As is well known in the art, a part of the IP address identifies the network, and can thus be used for detecting the current network of the laptop. The personal firewall may, for example, contain information on the IP address space of home network, and optionally, foreign networks, or a list of addresses available for the laptop in the home network.

[0027] When the current IP address of the laptop matches to a given address space or a list of addresses of the home network 10, for example, it can be assumed that the laptop is located in the home network 10 and the rule base 300 of the home network 10 is used. Thus, the current IP address is used as a selection rule for activating the rule base 300. However, there is some uncertainty in determining the location based on the current IP address only, and some approaches to overcome this problem are described with reference to further embodiments of the invention below.

[0028] Referring again to the generic flow diagram shown in Figure 4, the current IP address of the laptop is firstly determined in the step 401. The current IP address may be obtained simply by asking for it from the operating system of the laptop by means of using IP configuration routine. The current location of the laptop is monitored constantly, and therefore the personal firewall may be configured to periodically query the current IP address from the operating system. More preferably, the operating system of the laptop may be configured to inform any changes in the IP address to the personal firewall, and therefore a need for query the IP address time-to-time can be avoided. The step 401 may also include verification of the location determined based on the IP address by a verification procedure described below. In step 402, the personal firewall compares the current IP address with the current IP address stored in the personal firewall. If the IP address has not changed, the present active rule base can be maintained. However, if the IP address has changed, the personal firewall checks whether the new IP address matches to any IP address space or IP address belonging to one of the networks on the selection rule list in the personal firewall (step 403). If the new IP address does not belong to any of the networks on the selection rule list, the personal firewall considers the current network an unidentified network, and a default rule base 302 is selected (step 404). If the network cannot be identified and the default rule base must be used, it is normally assumed that the laptop is in a potentially hostile environment, most likely in the public Internet 12. Therefore, the default rule base is typically defined to provide the maximum protection needed. If the new IP address belongs to one of the networks defined on the selection rule list in step 403, it means that the network has been identified and a rule base linked to the identified network (or the corresponding selection rule) is selected and activated (step 405). In the simplest implementation, the selection rules include only the home network of the laptop and the cor-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 141

responding IP address space or list of addresses . If the current IP address belongs to the home network, the rule base 300 of the home network 10 is used. Otherwise the rule base 310 for foreign network or the default rule base 302 is used. In a more complicated implementation, there are selection rules (i.e. IP addresses and associated rule bases) also for at least one foreign network and/or different segments of the home network 10.

[0029] In the examples described above there are two or more rule bases which are enabled or disabled on the basis of the current location of the laptop. However, there are also alternative ways to implement different rule bases. One alternative is to provide only one rule base in which the rules are enabled and disabled in different combinations on the basis of the current location of the laptop.

[0030] As noted above, there are situations where the location (the current network) determined on the basis of the current IP address is uncertain, i.e. the IP address fails to indicate the current location of the laptop. If the IP address does not match the current network, use of the Internet protocol (IP) to attack against the laptop is not likely, and one may reason that in that case a personal firewall does not need to be used. However, there is still a possibility that there is an attack using other protocols, such NetBEUI or IPX. By detecting the situation where the IP address of the laptop is not an IP address of the current network, it is possible to block such protocols while in foreign networks. Further, NAT (network address translation) and private IP addresses are frequently used. This means that the same IP address is in use in several networks. In that case it is not enough to trust IP address information only when determining the location of the laptop. It is even possible that while being connected to a hostile network, the DHCP (dynamic host configuration protocol) gives familiar IP address to make it easier to attack the laptop. Basically, the DHCP enables individual computers on a network to connect to a DHCP server, such as the server 9 in Figure 1, and be assigned a dynamic IP address of the current network.

[0031] Thus, according to an aspect of the invention, in addition to the detection of location based on the current IP address described above, a further location verification procedure is carried out with a predetermined network element, which is preferably reachable only from the location to be verified. More generally, the network element is selected in such a way that it responses to the verification request only if the request originates from the location (i.e. the network) to be verified. Preferably, the specific network element is provided with a location verification service supporting the verification according to the invention. The verification procedure requires that the verification method be specified for the personal firewall, preferably at the same time as the different locations are specified. In other words, the methods to verify the location are specified for the personal firewall in the initial configuration, for example. It is also possible that the verification methods are updated or changed by means of the updating procedure described below, in a manner similar to other security rules.

[0032] A generic location verification procedure according to one embodiment of the invention is described with reference to Figures 5 and 6. Firstly, the current location is determined on the basis of the currently used IP address as described above (step 501). Next, the personal firewall selects a predetermined network element that should be available for verification from the determined current location (step 502). Then the personal firewall sends to the selected network element a request to send a response with some data proving the identity of the network element (step 503). Referring now to Figure 6, the verifying network element receives the verification request (step 601) from a personal firewall. Then, either always in response to the verification request or only if predetermined requirements are met, the network element sends the response with the required identity data to the personal firewall (step 602). Referring again to Figure 5, the personal firewall waits for a response (step 504), and if no response is received (preferably within a predetermined period of time), the location determined on the basis of the current IP address is rejected, and the location is determined to be unknown (step 505). In that case, a default location and an associated rule base, such as the rule base 302, can be used.

[0033] However, if the response is received from the network element in the step 504, the personal firewall verifies the identity of the network element on the basis of the received identity data, e.g. by comparing the received identity data with identity data stored in the personal firewall (step 506). If the verification of the identity is unsuccessful (step 507), the procedure proceeds to the step 505 described above. However, if the verification of the identity of the network element is successful, also the location of the laptop determined on the basis of the current IP address has been successfully verified and can be accepted.

[0034] Additionally, it is possible that one IP address is included in more than one selection rule in the personal firewall. In that case, if the verification of the laptop being located in a first network indicated by the current IP address fails, it is checked if the laptop is located in a second network indicated by the current IP address. There are various ways to implement the generic location verification procedure described above. The simplest way to implement the location verification service is to probe some known (known to the personal firewall) element. For example, it is possible to ask the MAC address of the known network element located in the home network and having a known IP address. The network element returns the MAC address in response, and if the AMC address is the one that it is assumed to be (e.g. matches with a MAC address stored in the personal firewall), it is verified that the laptop is located in the home network. If the MAC address is not the correct one, the laptop is determined to be outside the internal

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 142

network.

[0035] One possible implementation is that the location verification service is implemented in such network element in an internal network which can be reached only from inside the internal network. For example, the firewall protecting the internal network, such as firewall 5 in Figure 1, may be configured to discard all connections from outside the internal network to this network element. If the personal firewall is able to reach this specific network element, it is verified that the personal firewall resides inside the internal network. Otherwise the personal firewall resides outside the internal network. There may be a plurality of such location verification services in a plurality of internal networks or sub-networks, and if the personal firewall cannot reach any of them, then it is assumed to be outside this plurality of internal networks or sub-networks.

[0036] The location verification service may be incorporated into the network firewall, such as the location verification service 50 in the firewall 5 in Figure 1. For example, verification requests only from the direction of the internal network (such as home network 10) may be responded by the location verification service 50. This implementation is rather straightforward, since the present firewalls are readily capable of detecting from which direction, that is, from which interface, a data packet is coming from.

[0037] In any case it is preferred that the personal firewall communicates with the location verification service by using some cryptographically strong method, such as public key encryption. For example SSL can be used. The certainty of the location verification can be further improved by setting the TTL (time-to-live) field in the location verification request to a relatively low value, so that the request is capable of reaching only a nearby location verification service. The TTL value is decremented each time the verification request passes through a router connecting different networks or network segments. If the TTL value is set to, for example, a zero value the verification request is not able to pass through a router to a different network or a network segment.

[0038] The use of additional location verification makes it even possible to automatically create a VPN tunnel to the home network even if the current location (a foreign network) is using the same IP address as in the internal network.

[0039] All the embodiments described above are effective both in standalone computers and in centrally managed computers. The central management of personal firewalls enables a uniform protection level in all computer devices using the private network. One feature of the central management is that preferably all of the personal firewalls have essentially similar security rules. It should also be possible to update these. It is preferable that rule updates are applied in the personal firewalls as soon as possible after they have been made in the central management. Because it is not sufficient to rely on the manual updating by the user, the process of updating the rules must be automated. However, distribution of the updated rules by a push transmission from the central management is not a sufficient option in a case where the personal firewalls can move from one network to another. Therefore, according to an aspect of the invention, the personal firewalls are arranged to periodically query the availability of updated rules from the central management. An updating procedure according to one embodiment of the invention is described with reference to Figures 7 and 8.

[0040] Firstly, a personal firewall measures a predetermined updating period, which can be any period of time, preferably one day or a few days (step 701). In step 702, the personal firewall checks whether the updating period has expired, and if not, the procedure returns to the step 701. When the updating period has expired, the personal firewall checks whether the current location of the laptop is in the home network (step 703) or in another subnetwork of the same company. The location determination is preferably based on the methods described above. If the current location is in the home network, the process proceeds directly to the step 705. However, if the current location is not in the home network, the personal firewall waits for the laptop to return to or establish a connection (e.g. VPN) to the home network (step 704), before proceeding to the step 705. In the step 705, the personal firewall sends a rule update query to the central management, such as the personal firewall management server 8 in Figure 1. Referring now to Figure 8, the personal firewall management 8 receives the rule update query from the personal firewall (step 801) and sends updated rules, if there are any, to the personal firewall (step 802). Referring again to Figure 7, the personal firewall checks whether the reply received from the personal firewall management 8 contains rule updatings (step 706), and if not, the process returns to the step 701 to measure the next updating period. However, if rule updatings have been received, the personal firewall updates the relevant rule bases stored in the laptop (step 707).

[0041] It is also preferable that the logs relating to the communication transactions of the laptop are handled in a central location. Since the laptops are frequently disconnected from the home network, logs must be collected locally. In order to enable central handling, the logs must be transferred to a central log server, such as the personal firewall management 8, when such is available. This should take place automatically when the computer device provided with a personal firewall is connected to the home network. Figure 9 illustrates log handling according to one embodiment of the invention.

[0042] Firstly, the personal firewall creates a log file each time the laptop is involved with a communication transaction, such as an Internet session (step 91). Then the personal firewall determines the current location of the laptop, preferably based on the location determining methods described above (steps 92 and 93). If the lo-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 143

cation of the laptop is in the home network or in another subnetwork of the same company, the personal firewall sends the log file to the central log server 8 immediately (step 94). However, if the current location of the laptop is not in the home network, the process proceeds to the step 95, where the log file is stored locally. Similarly, a number of log files is collected locally while the laptop is disconnected from the home network. When the personal firewall next time detects that the computer device is relocated in the home network, it sends the collected log files to the central log server 8. Optionally, the personal firewall may also send the collected log files to the personal firewall management 8 when the laptop has established a (e.g. VPN) connection to the home network.

[0043] It is apparent for those skilled in the art that the illustrative embodiments described are only examples and that various modifications can be made within the scope and spirit of the invention as defined in the appended claims.

Claims

1. A method of controlling a personal firewall in a client computer, said method comprising

providing said personal firewall with at least one set of security rules to be used when said client computer is connected to a home network of said client computer, and at least one set of security rules to be used when said client computer is connected to foreign networks,

monitoring the current location of said client computer based on an Internet Protocol (IP) address currently used by said client computer, and

automatically selecting one of said sets of security rules by said personal firewall according to said current location of said client computer.

2. A method according to claim 1, wherein

said step of providing comprises providing said client computer with said set of security rules in form of at least two rule bases,

said step of selecting comprises enabling one of said rule bases at a time according to the current location of said client computer.

3. A method according to claim 1, wherein

said step of providing comprises providing said client computer with said set of security rules in form of one rule base, and

said step of selecting comprises enabling and disabling rules in said one rule base in different combinations according to the current location of said client computer.

4. A method according to claim 1, wherein said step of monitoring comprises

storing in said client computer an IP address space of said home network,

comparing the current IP address of said client computer with said IP address space, and

if the current IP address of said client computer matches said IP address space, determining said personal firewall to be located in said home network.

5. A method according to any one of claims 1, 2 or 3, wherein said step of monitoring comprises

storing in said client computer a list of IP addresses of said home network,

comparing the current IP address of said client computer with said list of IP addresses, and

if the current IP address of said client computer matches one of said addresses on said list, determining said client computer to be located in said home network.

6. A method according to any one of claims 1, 2 or 3, comprising

verifying the current location determined on the basis of the current IP address of said client computer by carrying out a location verification procedure with a predetermined network element.

7. A method according to claim 6, wherein said step of verifying comprises

checking availability of said predetermined network element related to the current IP address, said predetermined network element responding only if said client computer is located in the network in which it is assumed to be on the basis of the current IP address,

verifying the current location determined based on said current IP address, if said predetermined network element responses with a specific identity data.

8. A method according to claim 7, wherein said specific identity data is a Media Access Control (MAC) address of said predetermined network element.

9. A method of managing a personal firewall in a client computer, comprising

storing in said personal firewall at least one set of security rules to be used when said client computer is connected to a home network of said client computer, and at least one set of security rules to be used when said client computer is connected to foreign networks,

storing updated sets of security rules, if any, in a centralized unit in said home network of said client computer,

configuring said personal firewall to periodically query the availability of said updated sets of security rules from said centralized rule base server

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 144

when being located in said home network, or when having a remote access to said home network while being located in a foreign network, and

loading said updated sets of security rules from said centralized rule base server to said personal firewall in response to said query, if such updated sets of security rules are available.

10. A method according to claim 9, comprising

monitoring the current location of said client computer based on an Internet Protocol (IP) address currently used by said client computer, and

automatically activating said periodical query, when the current location of said client computer is in said home network.

11. A method according to claim 9, comprising

monitoring the current location of said client computer based on an Internet Protocol (IP) address currently used by said client computer,

sending log files to a centralized log server from said personal firewall, when the current location of said client computer is in said home network, said log files containing information on communication transactions in said client computer,

collecting log files locally at said personal firewall, when the current location of said client computer is not in said home network, and

transferring said locally collected log files from said personal firewall to said centralized log server, when said client computer is connected to said home network.

12. A computer terminal, comprising

a personal firewall provided with at least one set of security rules to be used when said computer terminal is connected to a home network of said computer terminal, and at least one set of security rules to be used when said computer terminal is connected to foreign networks,

said personal firewall having a mechanism monitoring the current location of said computer terminal based on an Internet Protocol (IP) address currently used by said computer terminal, and

said personal firewall having a mechanism automatically selecting one of said sets of security rules by said personal firewall according to said current location of said computer terminal.

13. A computer terminal according to claim 12, further comprising

a mechanism verifying the current location determined on the basis of the current IP address of said computer terminal by carrying out a location verification procedure with a predetermined network element.

14. A computer-readable medium, containing a computer software which, when executed in a computer device, causes the computer device to provide a personal firewall routine comprising

storing at least one set of security rules to be used when said computer device is connected to a home network of said computer device, and at least one set of security rules to be used when said computer device is connected to foreign networks,

monitoring the current location of said computer device based on an Internet Protocol (IP) address currently used by said computer device, and

automatically selecting one of said sets of security rules by said personal firewall according to said current location of said computer device.

15. A computer-readable medium according to claim 14, said personal firewall routine further comprising

verifying the current location determined on the basis of the current IP address of said client computer by carrying out a location verification procedure with a predetermined network element

16. A computer device, comprising a personal firewall routine configured to store security rules for a home network of said computer device, and for foreign networks,

periodically query the availability of updated security rules from a centralized rule base server in said home network of said computer device when said computer device is located in said home network, and

download said updated security rules from said centralized rule base server, if such updated security rules are available in said centralized rule base server.

17. A computer device according to claim 16, said personal firewall routine being further configured to

activate said periodical query also when said computer device has a remote access to said home network while being located in a foreign network.

18. A computer device, comprising a personal firewall routine configured to

send log files to a centralized log server, when a current location of said computer device is in a home network of said computer device, said log files containing information on communication transactions in said computer device,

collect log files locally in said computer device, when the current location of said client device is not in said home network,

transfer said locally collected log files to said centralized log server, when said computer device is reconnected to said home network.

19. A computer-readable medium, containing computer software which, when executed in a computer de-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 145

vice, causes the computer device to provide a personal firewall routine comprising

storing security rules for a home network of said computer device, and for foreign networks,

periodically querying the availability of updated security rules from a centralized rule base server in said home network of said computer device when said computer device is located in said home network, and

downloading said updated security rules from said centralized rule base server, if such updated security rules are available in said centralized rule base server.

20. A computer-readable medium, containing computer software which, when executed in a computer device, causes the computer device to provide a personal firewall routine comprising

sending log files to a centralized log server, when a current location of said computer device is in a home network of said computer device, said log files containing information on communication transactions in said computer device,

collecting log files locally in said computer device, when the current location of said client device is not in said home network,

transferring said locally collected log files to said centralized log server, when said computer device is reconnected to said home network.

21. A computer program comprising program code means for performing all the steps of any one of claims 1 to 11 when the program is run on a computer.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 146

Fig. 1

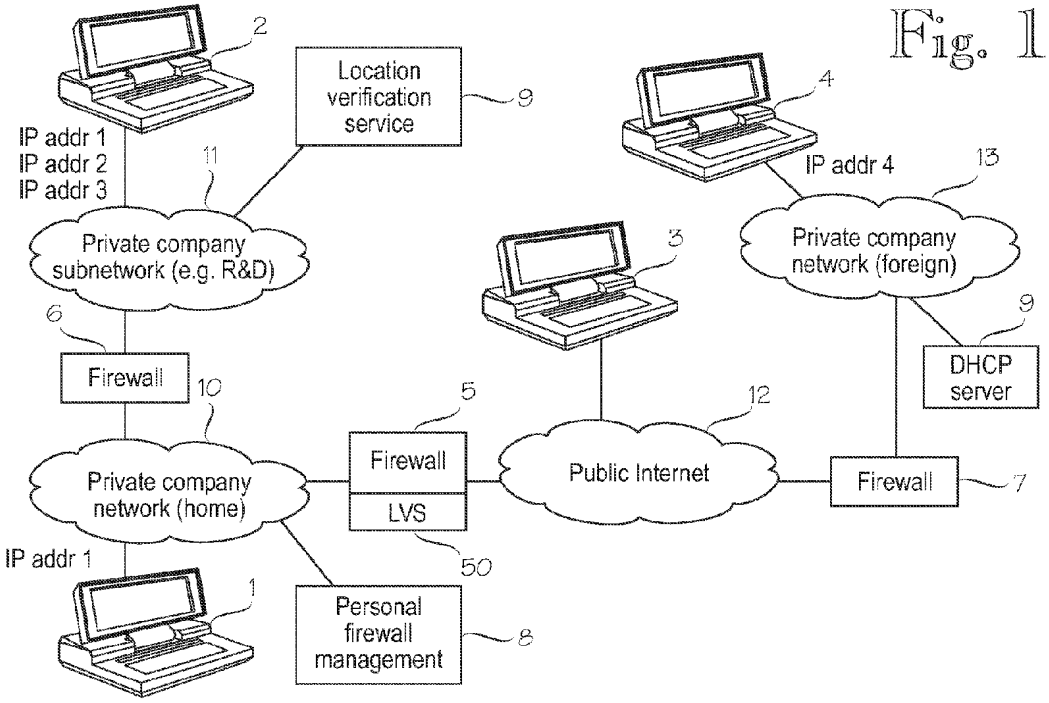

Fig. 2

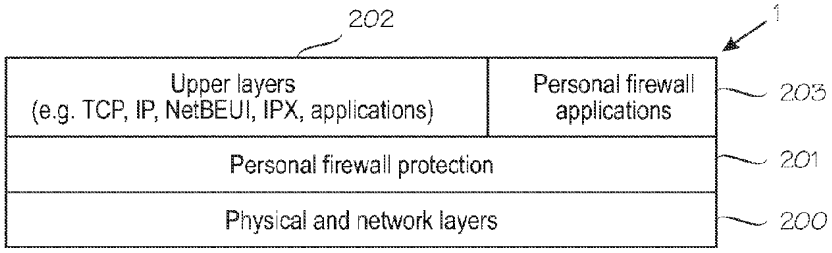

Fig. 3

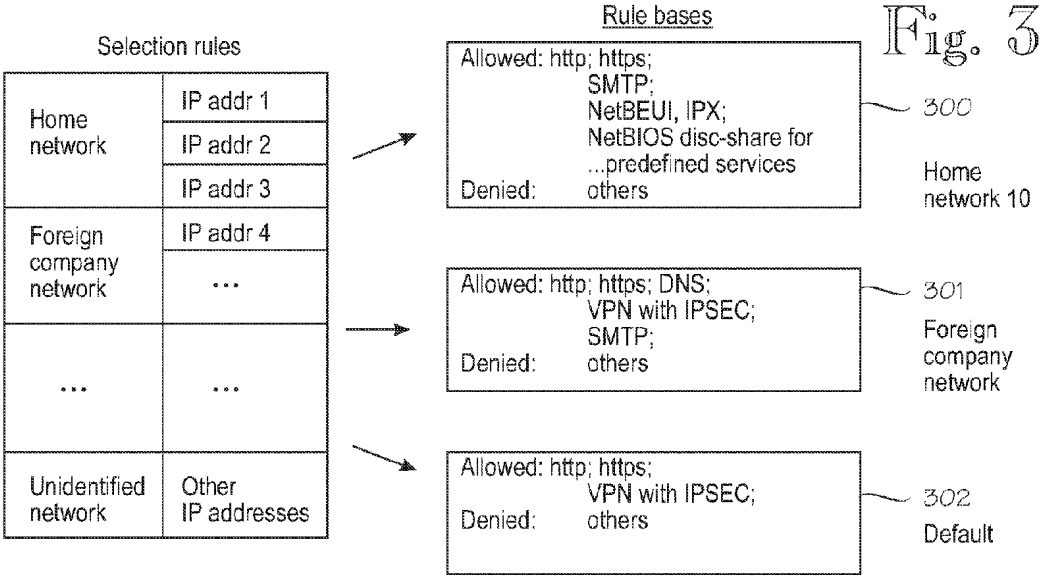

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 147

Fig. 4

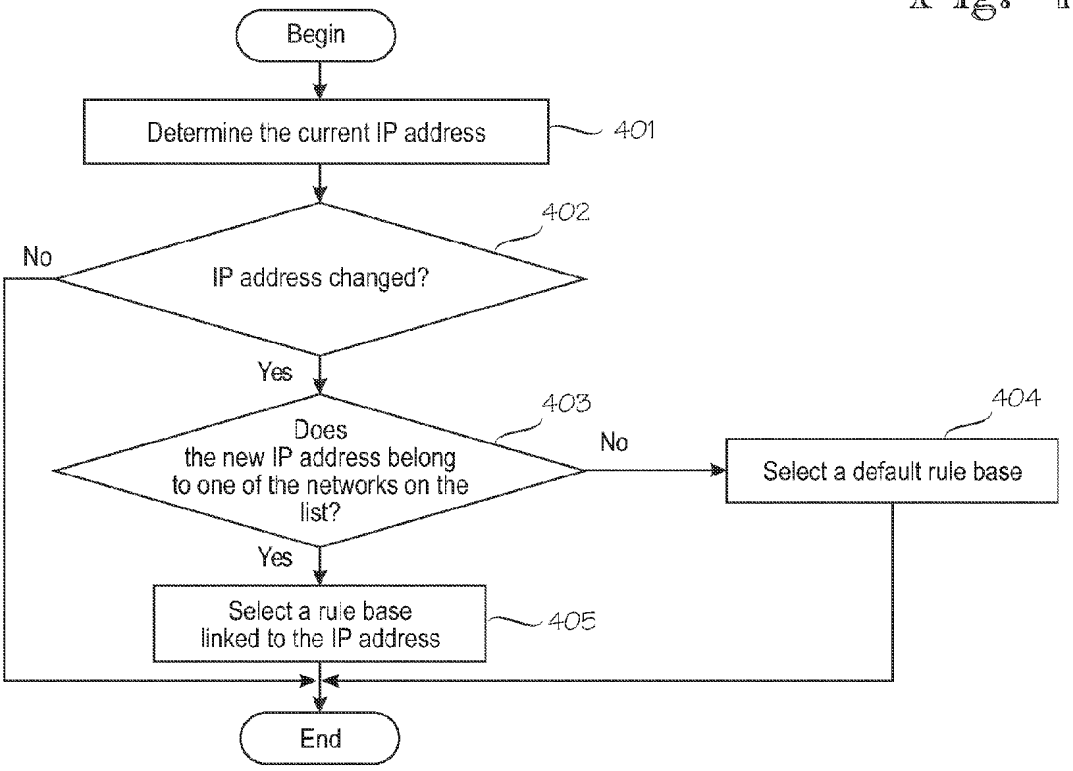

Fig. 9

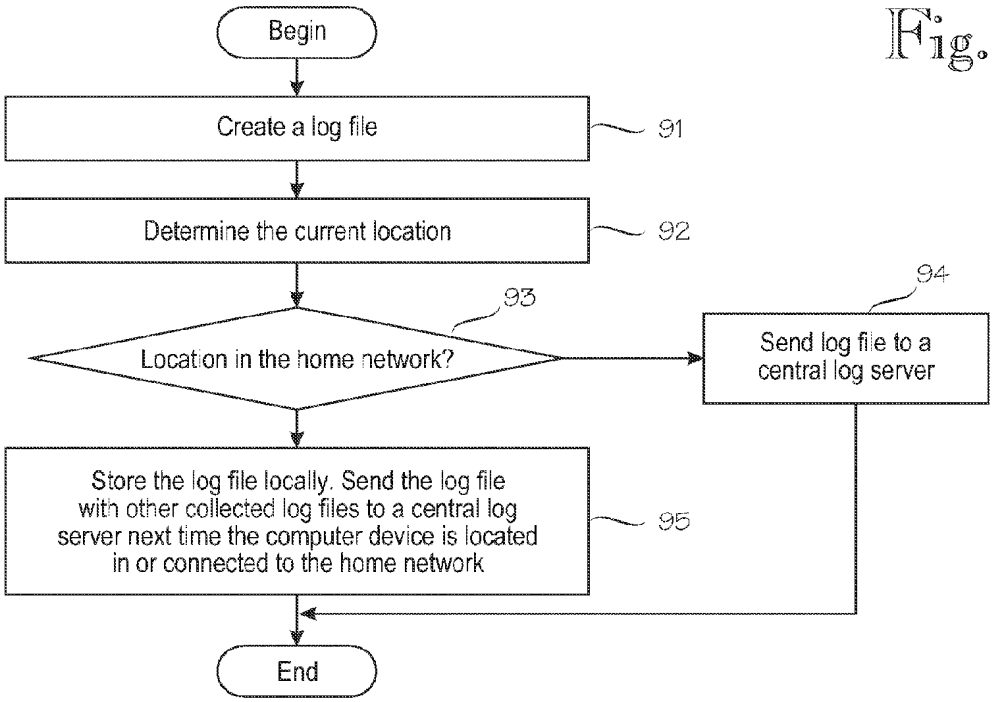

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 148

Fig. 5

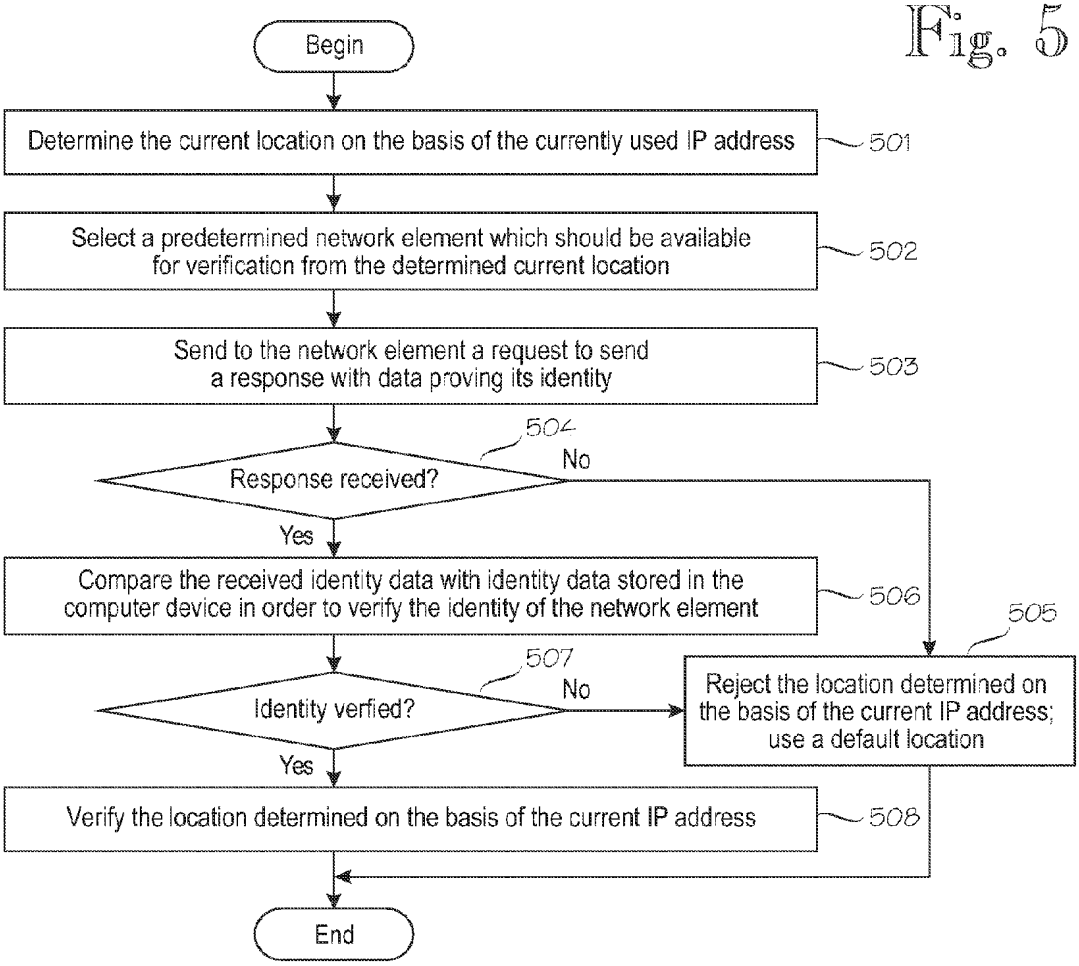

Fig. 6

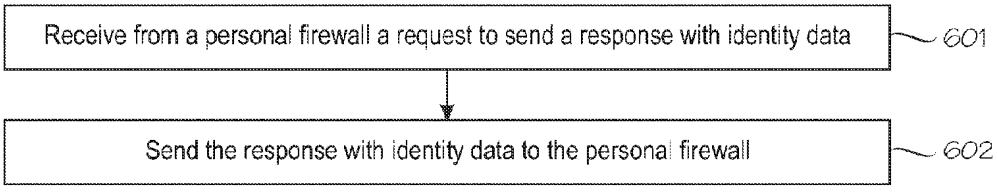

Fig. 8

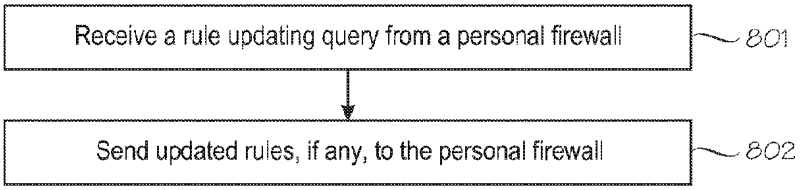

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 149

Fig. 7

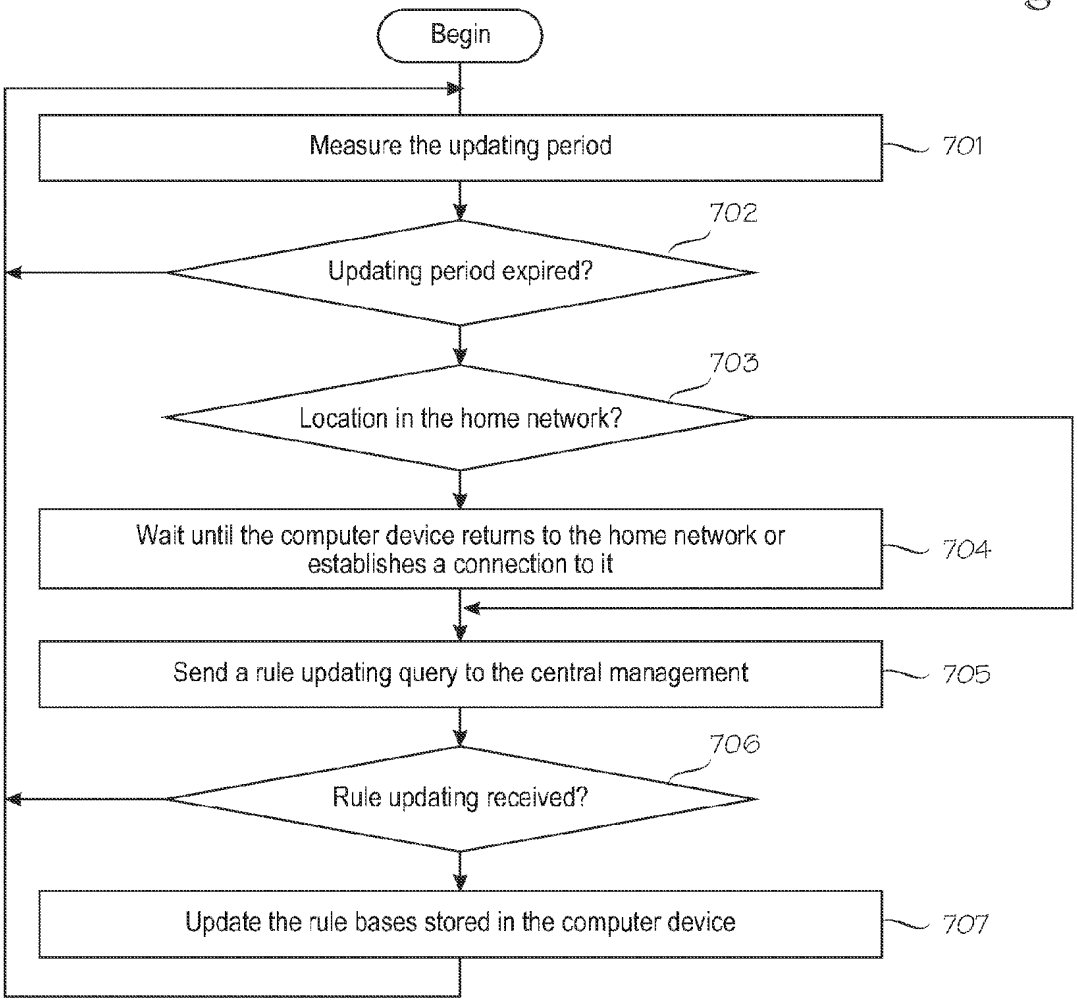

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 150

European Patent Office

EUROPEAN SEARCH REPORT

Application Number
EP 02 10 2581

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
A	EP 0 854 621 A (AT & T CORP) 22 July 1998 (1998-07-22) * page 4, line 28 - line 47 *	1-21	H04L29/06
A	US 6 158 008 A (OWENS LESLIE DALE ET AL) 5 December 2000 (2000-12-05) * claim 1 *	1-21	

TECHNICAL FIELDS SEARCHED (Int.Cl.7)

H04L

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
THE HAGUE	14 March 2003	Veen, G

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P04C01)

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 151

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO. EP 02 10 2581

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

14-03-2003

Patent document cited in search report		Publication date	Patent family member(s)		Publication date
EP 0854621	A	22-07-1998	US	6233686 B1	15-05-2001
			CA	2226814 A1	17-07-1998
			EP	0854621 A1	22-07-1998
			JP	10229418 A	25-08-1998
US 6158008	A	05-12-2000	WO	9921339 A1	29-04-1999

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 152

Espacenet

Bibliographic data: WO2005046145 (A1) — 2005-05-19

ADAPTABLE NETWORK BRIDGE

No documents available for this priority number.

Inventor(s): HALEN JOACIM ± (HALEN, JOACIM)

Applicant(s): ERICSSON TELEFON AB L M [SE] ± (TELEFONAKTIEBOLAGET LM ERICSSON (PUBL))

Classification: **- international:** ***H04L12/24; H04L12/46; H04L12/56;*** (IPC1-7): H04L12/46; H04L12/56
- cooperative: **H04L12/4625**; **H04L41/0893**; **H04L41/0896**; **H04L45/04**; **H04L45/56**; **H04L45/66**; **H04L49/351**

Application number: WO2004SE01611 20041105

Priority number (s): US20030517737P 20031106 ; US20040977724 20041029

Also published as: US2005111434 (A1) US7672318 (B2) JP2007534219 (A) JP4323523 (B2) EP1683313 (A1) more

Abstract of WO2005046145 (A1)

An adaptable network bridge such as an Ethernet bridge (10) includes a Rule Engine (16) that processes data frames according to a set of switching rules. Each switching rule identifies a type of frame to which the rule applies, and how the type of frame is to be processed. The bridge also includes means (18) for dynamically adding, subtracting, or modifying the switching rules in response to changes in the bridge's state information, or changes in context information. If the bridge includes a plurality of input/output ports (14, 15), the Rule Engine (16) may include a plurality of sets of switching rules, with each set of switching rules being applicable to a different port in the bridge. In this way, different network providers can share the same equipment, but have different network behavior.

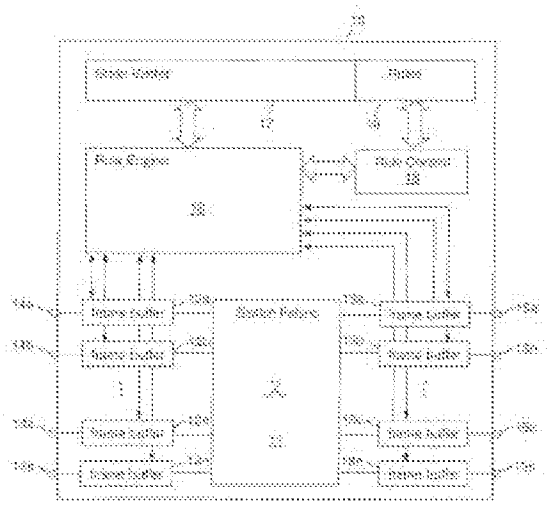

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
19 May 2005 (19.05.2005)

PCT

(10) International Publication Number
WO 2005/046145 A1

(51) International Patent Classification[7]: H04L 12/56, 12/46

(21) International Application Number:
PCT/SE2004/001611

(22) International Filing Date:
5 November 2004 (05.11.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/517,737 6 November 2003 (06.11.2003) US
10/977,724 29 October 2004 (29.10.2004) US

(71) Applicant: TELEFONAKTIEBOLAGET LM ERICSSON (publ) [SE/SE]; S-164 83 Stockholm (SE).

(72) Inventor: HALEN, Joacim; Anhaltsvägen 8, S-191 40 Sollentuna (SE).

(74) Agent: KARIN BOESTAD et al.; Ericsson AB, Patent Unit Core Networks Kista, S-164 80 Stockholm (SE).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*
— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: ADAPTABLE NETWORK BRIDGE

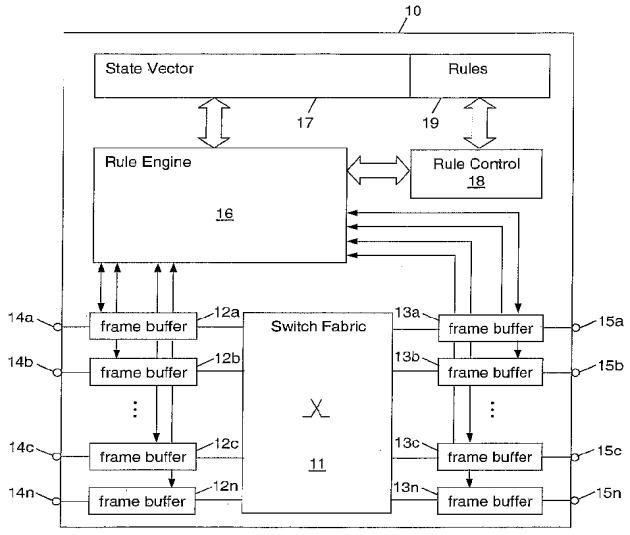

(57) Abstract: An adaptable network bridge such as an Ethernet bridge (10) includes a Rule Engine (16) that processes data frames according to a set of switching rules. Each switching rule identifies a type of frame to which the rule applies, and how the type of frame is to be processed. The bridge also includes means (18) for dynamically adding, subtracting, or modifying the switching rules in response to changes in the bridge's state information, or changes in context information. If the bridge includes a plurality of input/output ports (14, 15), the Rule Engine (16) may include a plurality of sets of switching rules, with each set of switching rules being applicable to a different port in the bridge. In this way, different network providers can share the same equipment, but have different network behavior.

WO 2005/046145 A1

ADAPTABLE NETWORK BRIDGE

CROSS-REFERENCE TO RELATED APPLICATIONS

This nonprovisional application claims priority under 35 U.S.C. § 119(e) on U.S. Provisional Application No. 60/517,737 entitled, *Programmable Bridge*, filed November 6, 2003, the disclosure of which is incorporated herein in its entirety.

BACKGROUND OF THE INVENTION

Technical Field of the Invention

The present invention relates generally to digital communication systems. More particularly, and not by way of limitation, the invention is directed to an adaptable network bridge having a Rule Engine that simplifies adapting the bridge to dynamically changing conditions.

Description of Related Art

Ethernet is a layer-2 packet-based transmission protocol that is primarily used in local area networks (LANs). Ethernet is the common name for the Institute of Electrical and Electronics Engineers (IEEE) 802.3 industry specification. Data is transmitted in Ethernet frames. Each frame includes a preamble with 64 bits utilized for synchronization. A Start of Frame Delimiter (SFD), a destination address, a source address, and a length/type identifier follow the preamble. Media Access Control (MAC) client data, together with a Packet Assembler/Disassembler (PAD) may vary in length from 46 to 1500 bytes (octets). A Frame Check Sequence (FCS) adds four more octets. Higher-level protocols reside in the content part of the frame. The frame size is counted from the destination address to the FCS, inclusive, and thus may vary between 64 and 1518 octets, not including a Virtual Local Area Network (VLAN) tag, which adds 4 octets.

Today most packet-based networks are built on Ethernet technology, as it is a cheap and reliable technology with good performance. Ethernet has been used in office networks for a long time and is now used to build Wide Area Networks (WANs) as well as small networks in homes. The communications

industry is now expanding the use of Ethernet to what is called "Public Ethernet" where millions of users are connected in the same network. Unfortunately, Ethernet does not scale to such large networks due to flooding of broadcasted messages. Some solutions have been proposed, but these solutions require modifications to the existing standards if they are to be implemented on standard equipment.

Ethernet bridging is a way to connect networks together to form a larger network. The standard for bridging is ANSI/IEEE 802.1d, which is hereby incorporated by reference herein. A bridge is a way to connect two separate network segments together in a protocol-independent way. Packets are forwarded based on Ethernet address, rather than IP address. Since forwarding is done at Layer 2, all protocols can go transparently through a bridge.

There are basically two approaches to adding new functionality to Ethernet bridges (and other Ethernet equipment as well): standardization thru IEEE, or proprietary solutions, which make the equipment non-standard. Adding new functionality and behavior to IEEE 802.1d MAC bridges, better known as Ethernet switches (or other Ethernet equipment), is a difficult and time-consuming process without any guarantees of success. Each new function must be proposed to and scrutinized by an IEEE standardization group (an existing working group or a new group has to be accepted and established). If one is lucky, the proposed function is accepted and added to the standard. However, since each new function requires a change in the standard, there is reluctance to add new functions. If the proposal is rejected, one is limited to a proprietary solution on non-standard equipment. The members of the standardization groups come from different vendors and other organizations, and although they are supposed to act in the interest of IEEE and the entire community, they may also act on agendas from their own organization/company that might be in conflict with the proposed functionality, thereby turning it down. It should also be noted that if a new function is added to the standard, it is then up to the different vendors to add the standardized functionality. A feature may be standardized but not made mandatory.

The proprietary approach also has its drawbacks. For example,

programmable network processors have been available for some time, and bridges may be purchased off the shelf containing programmable network processors. Proprietary solutions can be implemented on these. However, these network processors are vendor specific, and the solution will only work on that network processor or processor family. The solution will not be portable, and it will not be possible to run it on other vendors' equipment. Additionally, programming a network processor is done at a very low level, basically in assembly language, and the programming model is hard, requiring special skill. It should also be noted that a proprietary solution will make the equipment non-standard, and customers are more reluctant to buy non-standard equipment. There might also be inter-operational problems with other equipment.

SUMMARY OF THE INVENTION

The present invention is an adaptable network bridge having a Rule Engine that simplifies the adaptation of the bridge to dynamically changing conditions related, for example, to network conditions and user context. Simplicity is achieved by specifying how data frames should be processed and switched in terms of switching rules. Flexibility is achieved by enabling dynamic loading and updating of rules in the bridge via a protocol. In addition, one bridge may generate or modify rules for loading at another, neighboring, bridge. In a preferred embodiment, the bridge is an Ethernet bridge.

Thus, in one aspect, the invention is directed to an adaptable network bridge that includes a Rule Engine that processes data frames according to a set of switching rules. Each switching rule identifies a type of frame to which the rule applies, and how the type of frame is to be processed. The bridge also includes means for dynamically modifying the switching rules. In particular, the rules may be dynamically loaded and updated, for example, in response to changes in the bridge's state information, or changes in context information. In another embodiment, the bridge includes a plurality of input/output ports, and the Rule Engine includes a plurality of sets of switching rules, wherein each set of switching rules is applicable to a different port in the bridge.

In one aspect the rules are forced on the switch by external units such as a

Network Management System (NMS).

In a further aspect, a switch generates or modifies rules that are forced on other switches.

In yet another aspect, the invention is directed to a method of processing and switching Ethernet frames in an adaptable Ethernet bridge. The method includes the steps of defining a set of switching rules for the adaptable Ethernet bridge; receiving a given frame in the bridge; determining whether any of the rules are applicable to the given frame; and upon determining that one of the rules is applicable to the given frame, processing and switching the frame according to the applicable rule. The switching rules may be dynamically added, subtracted, or modified, for example, in response to changes in state or context information.

An important aspect of the invention is how the extended state vector can be utilized in new ways in the switching process. The switching process can be controlled and adapted in new ways allowing completely new ways to control networks.

BRIEF DESCRIPTION OF THE DRAWINGS

In the following, the essential features of the invention will be described in detail by showing preferred embodiments, with reference to the figures of the attached drawings.

FIG. 1 is a simplified block diagram of a Rule Engine utilized in the adaptable network bridge of the present invention;

FIG. 2 is a sequence diagram illustrating the steps of a process performed by the Rule Engine; and

FIG. 3 is an illustration of a network of adaptable Ethernet bridges controlled by a Network Management System (NMS).

DETAILED DESCRIPTION OF THE INVENTION

In the following description, for purposes of explanation and not limitation, specific details are set forth, such as particular embodiments, circuits, signal formats etc. in order to provide a thorough understanding of the present invention. It will be apparent to one skilled in the art that the present invention may be

practiced in other embodiments that depart from these specific details.

In existing Ethernet bridges, there are two identifiable phases in the switching process. In the first phase, parts of the frame header are checked for specific information. For example, the destination address of each frame may be checked against a forwarding table, priority bits, VLAN tags, and the like. In the second phase, different actions are taken within the bridge depending on the result of this check. For example, the frame may be sent to the port pointed to in the forwarding table, or the frame may be sent to all ports except the originating port, and so on. These steps characterize a process in which a frame is put through a filter that selects the action to take. Another important element is that the bridge stores and retrieves information in the form of a simple state vector in a forwarding table (i.e., memory), which the bridge is using in the process. However, the different parts of this process are fixed and may not be altered.

The present invention is based on a generalization of this switching process that adds several new key aspects to switching, enabling new ways of adapting and controlling the switching process as well as new ways to control Ethernet networks. The invention provides an adaptable Ethernet bridge having a Rule Engine that simplifies the addition or modification of functionality within the bridge. One of the purposes of the adaptable bridge is to provide a simple solution whereby anyone with reasonable skill can implement and add new functionality in a standardized way on standard equipment. Simplicity is achieved by specifying how Ethernet frames should be processed and switched in terms of switching rules. Flexibility is achieved by enabling dynamic loading and updating, inspection, and prioritization of rules in bridges via a protocol.

The generalized switching process consists of the Rule Engine processing Ethernet frames by applying the switching rules. A switching rule encapsulates a specific switching behavior that describes the filter and action processes together with the state vector. A switching rule thus consists of three parts:

1) A filter function that can filter any part or parts of a frame. The purpose of the filter is to decide if the rule is applicable to a given frame. In other words, the filter works as a predicate function on frames. In this decision process information in the state vector may be used.

2) An action function that describes how a frame header (or the entire frame) should be transformed and how the frame should be switched, i.e. to which port or ports the frame should be sent. In this process the action function may alter the state vector.

3) A state vector that the filter function and the action function can use in the switching process.

The switching rule can also contain other information such as version, priority specification, vendor information, and the like.

The state vector is an abstraction of several parts:

1) A local memory for a rule or a set of rules where rule-specific data can be stored and altered. Specific "management rules" make it possible to access that data from, for example, a management system.

2) A global bridge memory where bridge-specific data applicable to and accessible by all rules can be stored and altered. One example is the port-forwarding table.

3) Access to a set of bridge resources that can be implemented as library functions. Some examples are a function to retrieve the source port of the frame processed, a function that provides timestamps, counters and timers, and functions that describe the traffic flows and load in the bridge. The invention is not limited to these examples; other functions may be identified by those skilled in the art.

A switching rule describes the types of frames that match the rule and how these frames should be processed (e.g., switched or transformed before switching). Each incoming frame is checked against the set of rules stored in the bridge. If a rule matches, the frame is processed and switched according to the specification in the rule. If no rule matches the frame, a default rule may be utilized, or the frame may be dropped.

With the concept of rules, the switching process is defined as a simple and understandable programming model. Standardizing rule language and making it possible to dynamically load new rules in a bridge also makes it possible to add new functionality, adapting standard equipment to the needs at hand. If some functionality becomes obsolete, that functionality can be removed by removing the

rule or set of rules implementing it.

A powerful extension is to allow different rule sets to be valid for different ports enabling, for example, different network providers to share the same hardware but have different network behavior. Different embodiments of this extension include: a rules control mechanism that selects and prioritizes available rule sets based on port, time of day, load, and the like. Alternatively, this kind of information may be described in the filter functions, and the Rule Engine tries all rules until it finds a match.

Yet another powerful extension is to allow rules that monitor the state vector instead of being matched against Ethernet frames. That is, the filter function monitors the state vector, and when some condition is fulfilled, the action function is invoked. The monitoring process can be implemented by regularly processing these rules in the Rule Engine or in a separate Rule Engine.

Note that allowing the filter and action functions to have access to a state vector that goes beyond the forwarding table provides new and powerful ways to process and switch frames.

FIG. 1 is a simplified block diagram of an exemplary embodiment of the dynamically adaptable network bridge 10 of the present invention. The adaptable bridge includes many of the same components as existing bridges such as switch fabric 11, frame buffers 12a-12n and 13a-13n, and ports 14a-14n and 15a-15n. In addition, however, the adaptable bridge includes a Rule Engine 16, which replaces the implementation of some or all protocols. The adaptable bridge also includes a state vector memory 17 accessible by the Rule Engine, a rule control function 18, and a rules memory 19 accessible by the Rule Engine through the rule control function. Collectively, these new bridge components enable the adaptable bridge to receive and store rules from external units, store a state vector which may be dependent on external sources, and store rules that are imposed on other bridges.

Hybrid designs having both a Rule Engine and some of the protocols of existing bridges may also be implemented. The Rule Engine has processing capabilities, that is, the Rule Engine can (1) apply the filter function to a frame and calculate the result; (2) apply the action function if the rule matches; (3) access

the state vector; and (4) access the other resources of the switch. A frame entering the adaptable bridge through an input port such as port 14a and frame buffer 12a is acted upon by switching rules in the Rule Engine 16, including the filter function, action function, and state vector. The result is a calculated result frame, which is outputted through switch fabric 11, and an appropriate frame buffer and output port such as frame buffer 13a and output port 15a.

The rule control function 18 adds and removes rules dynamically, and orders or prioritizes THE set of rules. The rule control function may be implemented as protocols and secured against unwanted manipulation by using well known techniques intended for that purpose, for example, using digital signatures and encryption. Additionally, the rule control function may include a mechanism to transform (or translate) switching rules in human-readable form to an efficient form understandable by the Rule Engine 16.

The mechanisms to access different bridge resources and state information (i.e., state vectors), mechanisms to access library functions, and a protocol for externally accessing the state information are important parts of the adaptable bridge. The bridge resources and state information may be accessed via, for example, an Application Programming Interface (API). The protocol for externally accessing the state information may be implemented as a set of rules. The state of the bridge can affect switching rules dynamically, that is, a rule can behave differently depending on the current state. For example, frames can be switched differently depending on time of day, traffic pattern, traffic load, and the like. This behavior can also be controlled externally by, for example, a network management system (NMS).

FIG. 2 is a sequence diagram illustrating the steps of a method of processing an Ethernet frame by the Rule Engine 16. The process starts at step 21 and moves to step 22 where, for each incoming Ethernet frame, the port-specific rules associated with the Source Port are merged with the rules shared among all ports of the bridge (this merge can be done in advance by the priority module as a ordered list). The rules in the list are then tested to determine whether the frame matches the rule. This process begins at step 23, where a determination is made whether there are any more rules left to test. If not, and all

the rules have been tested, the frame is dropped at step 24. However, in the case of the first rule, there are rules left to test, so the method moves to step 25 where the filter function of the first rule in the list is applied to the frame. The filter function uses the information found in the frame together with the current state vector to determine whether the frame matches the rule. The filter function is a predicate, so the result at step 26 will be either true (the frame matches the rule) or false (the frame does not match the rule). If the result is false, the rule is discarded at step 27, and the method returns to step 23 where the test is performed on the next rule in the list.

The remainder of the rules are then tested until either no rules are left and the frame is dropped at step 24, or there is a match at step 26, and the method moves to step 28 where the action function of the rule is applied to the frame. The action function may use the library functions, the information in the frame, and the state vector to transform (or modify) the frame. In this process, the action function can use functions to modify the state vector. For example, a counter may be incremented, which counts the number of frames containing a specific high-level protocol coming from the port. At step 29, the action function determines the destination port or ports and sends the (possibly transformed) frame to the determined port or ports. The method then ends at step 30.

Another important aspect of the invention not shown in the sequence diagram is that other resources or modules in the bridge may continuously monitor and change the state vector in parallel. An extended Rule Engine, to the one described, may then invoke rules that, instead of being matched against Ethernet frames, monitor the state vector. When one or more predefined criteria are fulfilled (determined by the filter function), the action function generates new frames containing new rules to be sent to the bridge itself or to nearby bridges altering their behavior.

FIG. 3 is an illustration of a network 35 of adaptable Ethernet bridges controlled by a Network Management System (NMS) 36. The NMS monitors the state in the bridges and issues rules that adapt the behavior of certain bridges accordingly. For example, if $Bridge_{11}$ becomes heavily loaded with traffic, the NMS may issue a rule that will move some traffic from associated bridges (i.e.,

$Bridge_{21}$ – $Bridge_{23}$) to $Bridge_{12}$. For example, the rule may move every third frame or frames carrying a particular protocol to $Bridge_{12}$. Context information may also be added to the state vector, thereby allowing context to influence switching. Context information may contain, but is not limited to, information about traffic load in nearby bridges, forecasts or predictions of expected traffic, and the like. This feature is of particular interest in a mobile environment.

The NMS 36 is shown connected to the network by a dotted line because the adaptable bridges may automatically generate specific context rules and distribute them to nearby bridges as users move around in the network. When bridges automatically generate rules to other bridges (and to themselves), a self-managed network can be implemented in which bridges exchange rules dynamically. This capability can be used, for example, to setup a particular flow-path, distribute traffic during heavy load, or automatically redirect traffic in the case of network failures. For example, if Bridge11 receives a flow of high priority frames when it is already heavily loaded with traffic, it can issue a rule or set of rules to Bridge22 that move the flow to Bridge12 instead.

An Ethernet bridge is working in "wire-speed" and therefore it is desirable that the bridge introduce as little delay as possible. Adding a Rule Engine and basing the switching on rules may increase the delay. However, a number of techniques and mechanisms can be used to make the Rule Engine more efficient:

1) Filter functions can be tested in parallel, and a priority module can select the matching rule with highest priority. These computations can be performed because the filter functions are pure predicates, i.e. they contain no side effects.

2) Filter functions will in most cases be a conjunction (or disjunction) of comparisons. These can be directly mapped onto a pipelined matrix of comparison modules in hardware.

3) A similar pipelined hardware matrix realizing the language constructs used in the action functions can be implemented.

4) Some parts of the state vector, for example library functions reflecting bridge context and resources, can be implemented in hardware such as registers, thereby providing "instant" access to them.

In the preferred embodiment, the rule language used to describe the filter

and action functions is not a general purpose programming language. A programming language should only have the necessary constructs and mechanisms needed to express the solutions to the problems in the specified domain. If the language is too powerful, as a general-purpose language, it is possible to express an infinite loop (accidentally or on purpose) thereby increasing the delay to infinity. The preferred switching rule language, therefore, is a language primarily restricted to the domain of Ethernet switching (layer 2), with language constructs expressing the mechanisms used in the switching process. However, language constructs that are useful for other purposes like network management, constructs related to exchange of new rules, and so on, that are not in conflict with the basic requirements may also be included.

The present invention provides several advantages. The Ethernet standards can be implemented as switching rule sets, and as the standards evolve, the rule sets are updated. These updated rule sets can then easily be loaded into an adaptable network bridge. The rule sets can be vendor-independent and, for example, published by IEEE. Today, each vendor decides which protocols (standards) to support in a bridge, and unless the vendor issues an update whenever the standard is updated, users have to buy a new bridge if they need to follow the new standard.

In the present invention, the "programming language" of rules may be standardized thereby enabling anyone to add new functionality, as a switching rule, to any bridge with a Rule Engine. These rules can then be updated or removed if necessary. A proprietary solution can also be implemented as a rule or set of rules and loaded in any bridge with a Rule Engine. If a user later wants to use another solution, the rule set is simply replaced with a new rule set. It will be much less expensive to implement new functionality because, in most cases, the user only has to write some software in the form of a set of rules.

The present invention also greatly increases flexibility in networks because each bridge can be loaded with its own rule set down to the individual port level. The invention also provides the ability to adapt the functionality on each port based on time, user demands, service agreements, and the like. For example, different operators can share the same physical bridge (network) and provide

different services to their respective users.

Some examples of functionality that can be implemented with switching rules include the following:

1) Broadcast tunnels. It is simple to implement broadcast tunnels necessary for building Public Ethernet Networks.

2) Distributed router functionality added to a set of bridges.

3) Simple firewall functionality.

4) User adapted services on the port level.

5) Authentication and security mechanisms on the port level that are transparent to users.

6) Self-managed networks or semi-self managed networks.

This list is not intended to be exhaustive, and other functions may also be implemented with switching rules and loaded into any adaptable bridge with a Rule Engine.

As shown by the above examples, the state of the dynamically adaptable bridge can be utilized in the switching process to enable entirely new functionality and to provide new ways to use bridges in networks.

Although preferred embodiments of the present invention have been illustrated in the accompanying drawings and described in the foregoing Detailed Description, it is understood that the invention is not limited to the embodiments disclosed, but is capable of numerous rearrangements, modifications, and substitutions without departing from the scope of the invention. The specification contemplates any all modifications that fall within the scope of the invention defined by the following claims.

WHAT IS CLAIMED IS:

1. A dynamically adaptable network bridge comprising:

a Rule Engine that processes data frames according to a set of switching rules, wherein each switching rule identifies a type of frame to which the rule applies, and how the type of frame is to be processed; and

means for dynamically modifying the switching rules.

2. The dynamically adaptable network bridge of claim 1 wherein the means for dynamically modifying the switching rules includes a protocol for dynamically loading and updating the switching rules.

3. The dynamically adaptable network bridge of claim 1, wherein each switching rule includes a filter function for determining whether a given rule is applicable to a given frame.

4. The dynamically adaptable network bridge of claim 3, wherein each switching rule also includes an action function for specifying changes to the given frame's header and how the given frame should be switched.

5. The dynamically adaptable network bridge of claim 1, wherein each switching rule includes a state vector utilized for processing the type of frames to which each rule applies.

6. The dynamically adaptable network bridge of claim 5, wherein the state vector includes an abstraction of a local memory for storing rule-specific data for each switching rule.

7. The dynamically adaptable network bridge of claim 6, wherein specific management rules enable the Rule Engine to access the rule-specific data for each switching rule.

8. The dynamically adaptable network bridge of claim 6, wherein the state vector also includes an abstraction of a global bridge memory for storing bridge-specific data applicable to and accessible by all the switching rules.

9. The dynamically adaptable network bridge of claim 8, wherein the bridge-specific data includes a port-forwarding table.

10. The dynamically adaptable network bridge of claim 8, wherein the state vector also includes a set of library functions that implement bridge resources.

11. The dynamically adaptable network bridge of claim 10, wherein the set of library functions are selected from a group consisting of:

a function that retrieves a source port of a processed frame;

a function that provides timestamps;

a function that describes traffic flows in the bridge; and

a function that indicates the traffic load in the bridge.

12. The dynamically adaptable network bridge of claim 3, wherein the given frame is processed in accordance with the given rule if the filter function determines that the given rule is applicable to the given frame, and the given frame is dropped if the filter function determines that there is no rule applicable to the given frame.

13. The dynamically adaptable network bridge of claim 1, wherein the bridge includes a plurality of input/output ports, and the Rule Engine includes a plurality of sets of switching rules, wherein each set of switching rules is applicable to a different port in the bridge.

14. The dynamically adaptable network bridge of claim 1, wherein the bridge is an Ethernet bridge.

15. The dynamically adaptable network bridge of claim 1, wherein the means for dynamically modifying the switching rules includes a protocol that dynamically adds, subtracts, and modifies the switching rules.

16. The dynamically adaptable network bridge of claim 1, further comprising means for prioritizing rules within the set of switching rules.

17. The dynamically adaptable network bridge of claim 1, further comprising means for accessing state information for the bridge, wherein a given rule causes a given frame to be processed differently when the state information changes.

18. The dynamically adaptable network bridge of claim 17, wherein the given rule processes the given frame differently, depending on traffic load in the bridge.

19. The dynamically adaptable network bridge of claim 17, further comprising a port for receiving rules from a Network Management System (NMS), wherein the NMS monitors the state information for a plurality of adaptable bridges in a network, and issues rules that adapt the behavior of the bridge to optimize network performance.

20. The dynamically adaptable network bridge of claim 17, wherein the state information includes context information, wherein a given rule causes a given frame to be processed differently when the context information changes.

21. The dynamically adaptable network bridge of claim 1, wherein the bridge is one of a plurality of adaptable bridges in a network, and the bridge further comprises means for generating specific context rules and distributing the context rules to other adaptable bridges in the network.

22. A method of processing and switching data frames in an adaptable network bridge, said method comprising the steps of:

defining a set of switching rules for the adaptable bridge;

receiving a given frame in the bridge;

determining whether any of the rules are applicable to the given frame; and

upon determining that one of the rules is applicable to the given frame, processing and switching the frame according to the applicable rule.

23. The method of claim 22, further comprising dynamically adding, subtracting, and modifying the switching rules.

24. The method of claim 23, further comprising accessing state information for the bridge, wherein the step of dynamically adding, subtracting, and modifying the switching rules is performed in response to a change in the state information.

25. The method of claim 24, wherein the bridge is one of a plurality of adaptable bridges in a network, and the method further comprises the steps of:

receiving a new or modified switching rule from a Network Management System (NMS), wherein the NMS monitors the state information for a plurality of adaptable bridges in a network, and issues rules that adapt the behavior of the bridge to optimize network performance; and

processing and switching frames according to the new or modified switching rule received from the NMS.

26. The method of claim 24, further comprising adding context information to the state information, wherein the step of dynamically adding, subtracting, and modifying the switching rules is performed in response to a change in the context information.

27. The method of claim 22, wherein the bridge is one of a plurality of adaptable bridges in a network, and the method further comprises the steps of:

generating specific context rules; and

distributing the context rules to other adaptable bridges in the network.

28. The method of claim 22, wherein the bridge is an Ethernet bridge.

5

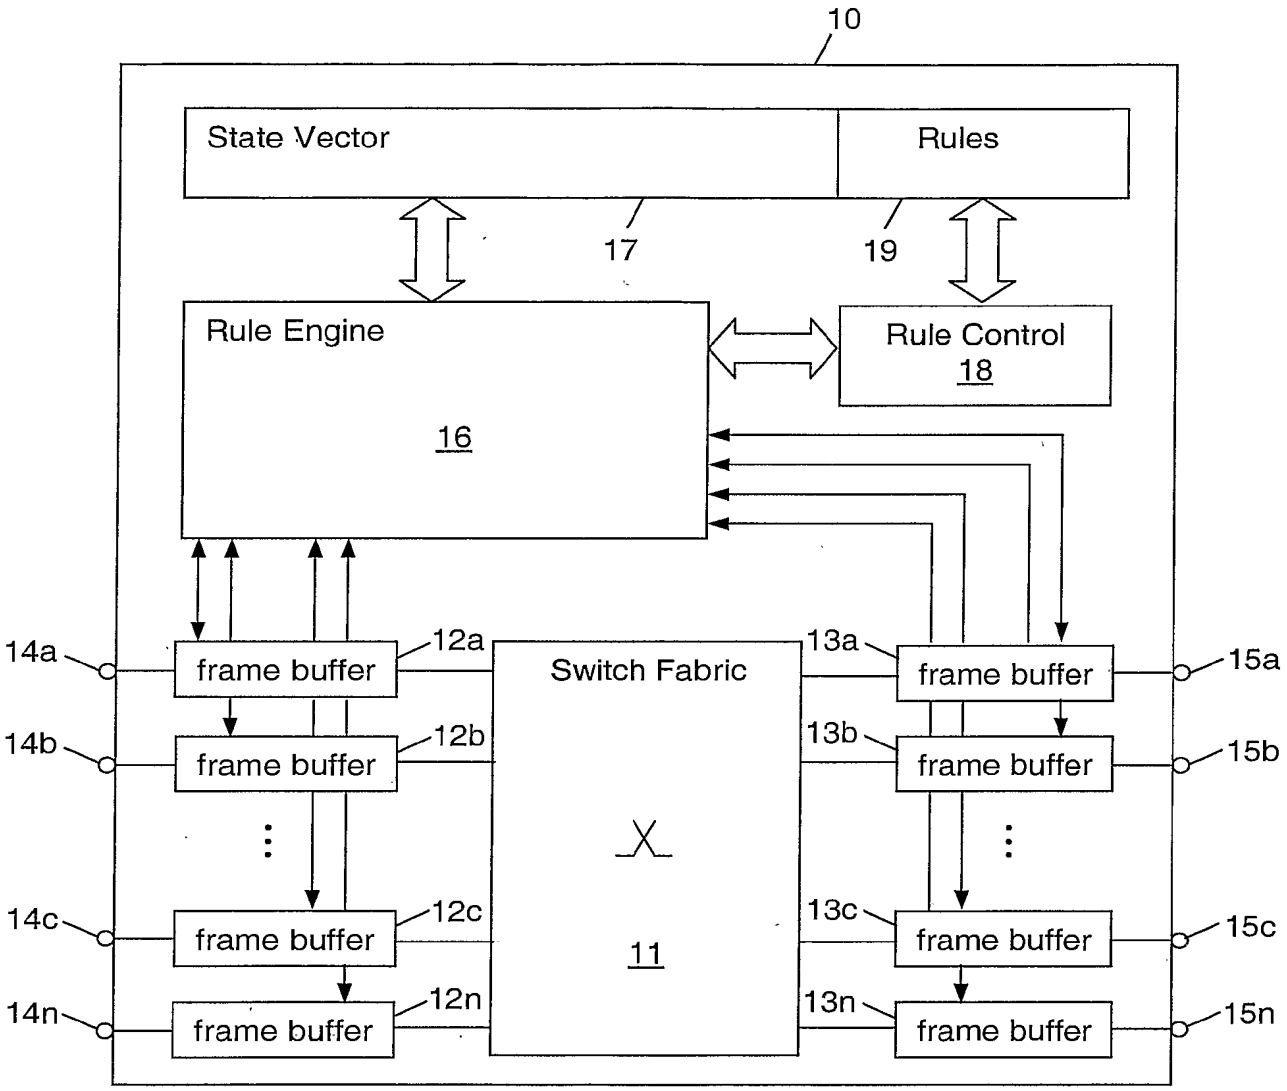

FIG. 1

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 173

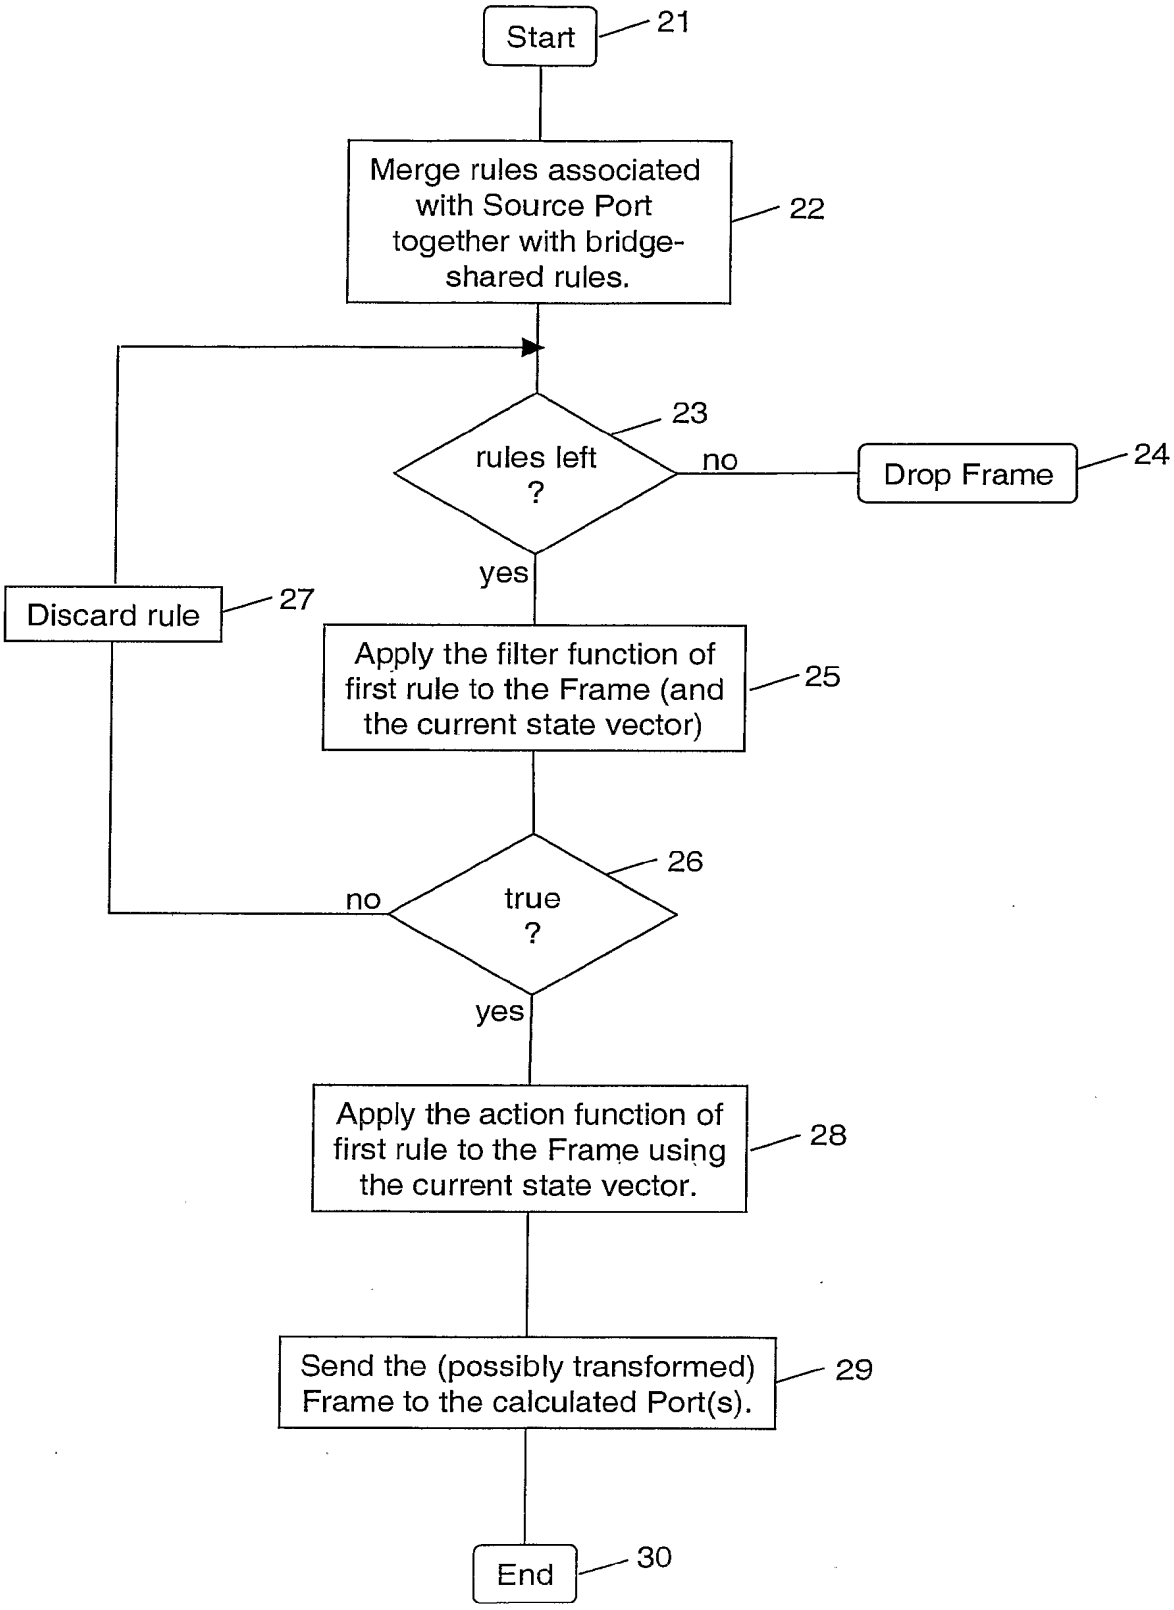

FIG. 2

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 174

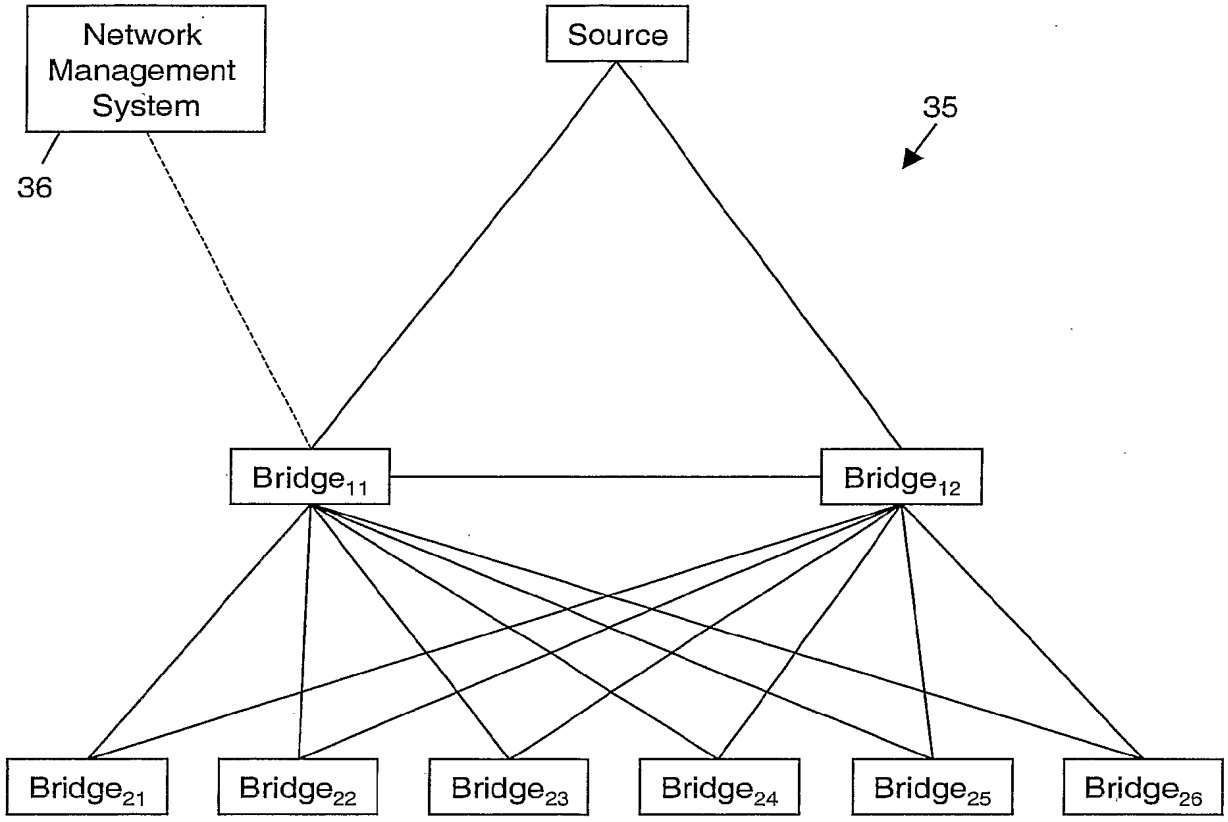

FIG. 3

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 175

INTERNATIONAL SEARCH REPORT

International Application No
PCT/SE2004/001611

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/56 H04L12/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 790 554 A (PITCHER ET AL) 4 August 1998 (1998-08-04) abstract column 1, line 15 - column 1, line 67 column 2, line 39 - column 3, line 25 column 3, line 48 - column 4, line 42 column 4, line 60 - column 8, line 4 column 8, line 24 - column 8, line 26; claims 1,2,14,22,28; figures 2,3 ----- -/--	1-5, 12-20, 22-25, 27,28

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
21 March 2005	31/03/2005

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Todorut, C

Form PCT/ISA/210 (second sheet) (January 2004)

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 176

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 396 493 A (SUGIYAMA ET AL) 7 March 1995 (1995-03-07) abstract column 1, line 63 - column 3, line 11 column 4, line 11 - column 4, line 51 column 5, line 67 - column 6, line 39; claims 1,17; figures 1-4	1-28
X	WO 00/56024 A (BROADCOM CORPORATION; KADAMBI, SHIRI; AMBE, SHEKHAR; KALKUNTE, MOHAN;) 21 September 2000 (2000-09-21) abstract	1-4,22, 23
A	page 53, line 12 - page 65, line 21; claims 1,2	5-21, 24-28
A	US 2003/123387 A1 (JACKSON ANDREW LLOYD 'GB!) 3 July 2003 (2003-07-03) abstract column 1, paragraph 2 - column 1, paragraph 9 column 2, paragraph 16 - column 3, paragraph 25; claims 1,6,9; figure 4	1-28
A	US 5 959 976 A (KUO ET AL) 28 September 1999 (1999-09-28) abstract column 1, line 5 - column 2, line 40 column 3, line 8 - column 3, line 18 column 5, line 56 - column 6, line 10; claim 1; figures 3-5	1-28
A	EP 0 508 886 A (DIGITAL EQUIPMENT CORPORATION) 14 October 1992 (1992-10-14) abstract column 1, line 40 - column 3, line 54 column 4, line 45 - column 5, line 16 column 5, line 55 - column 6, line 53; claims 1,23; figures 5-8	1-28
A	US 5 515 376 A (MURTHY ET AL) 7 May 1996 (1996-05-07) abstract column 6, line 1 - column 9, line 21	1-28

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
/SE2004/001611

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5790554 A	04-08-1998	NONE	
US 5396493 A	07-03-1995	JP 6224918 A	12-08-1994
		KR 9614983 B1	23-10-1996
WO 0056024 A	21-09-2000	US 2001043611 A1	22-11-2001
		AU 3529500 A	04-10-2000
		AU 3752000 A	04-10-2000
		AU 4325800 A	04-10-2000
		EP 1161817 A2	12-12-2001
		EP 1169809 A2	09-01-2002
		EP 1159805 A2	05-12-2001
		WO 0056024 A2	21-09-2000
		WO 0056011 A2	21-09-2000
		WO 0056013 A2	21-09-2000
		US 6707817 B1	16-03-2004
		US 6810037 B1	26-10-2004
		US 6850521 B1	01-02-2005
		US 6707818 B1	16-03-2004
		US 2004174898 A1	09-09-2004
		US 2005047411 A1	03-03-2005
		US 2002039365 A1	04-04-2002
		AU 5442700 A	12-12-2000
		AU 5586400 A	12-12-2000
		AU 5737300 A	31-01-2001
		AU 6334400 A	22-01-2001
		EP 1181791 A1	27-02-2002
		EP 1181792 A1	27-02-2002
		EP 1212867 A2	12-06-2002
		EP 1192767 A2	03-04-2002
		WO 0072533 A1	30-11-2000
		WO 0102965 A2	11-01-2001
		WO 0072531 A1	30-11-2000
		WO 0101724 A2	04-01-2001
		US 2002093974 A1	18-07-2002
		US 2002186705 A1	12-12-2002
		US 6813268 B1	02-11-2004
		US 6735679 B1	11-05-2004
		US 6792500 B1	14-09-2004
		US 6842457 B1	11-01-2005
		US 6430188 B1	06-08-2002
		US 6859454 B1	22-02-2005
		US 2004170176 A1	02-09-2004
		US 2005036488 A1	17-02-2005
		US 2002031090 A1	14-03-2002
US 2003123387 A1	03-07-2003	NONE	
US 5959976 A	28-09-1999	NONE	
EP 0508886 A	14-10-1992	DE 69217103 D1	13-03-1997
		DE 69217103 T2	21-08-1997
		EP 0508886 A1	14-10-1992
		JP 5114905 A	07-05-1993
		US 5515513 A	07-05-1996
US 5515376 A	07-05-1996	EP 0710415 A1	08-05-1996
		JP 3335358 B2	15-10-2002

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 178

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
/SE2004/001611

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5515376 A		JP 9500774 T	21-01-1997
		WO 9503659 A1	02-02-1995
		US 5610905 A	11-03-1997
		US 6545982 B1	08-04-2003

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 179

Electronic Acknowledgement Receipt

EFS ID:	19437021
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten/Rocko Bund
Filer Authorized By:	William Eugene Wooten
Attorney Docket Number:	007742.00029
Receipt Date:	27-JUN-2014
Filing Date:	16-APR-2014
Time Stamp:	14:50:02
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Non Patent Literature	NPL-ISR_16812869.PDF	1078548 bb0e661dcbcde2c2213fee0b75d5c81ce10349c9	no	9

Warnings:	
Information:	

2	Information Disclosure Statement (IDS) Form (SB08)	Information_Disclosure_Statement_Fillable_PDF_16811337.PDF	612284 917ab3589d278f868912b7bbd55dbff439d6be4c	no	4
Warnings:					
Information:					
3	Foreign Reference	EP1313290A1_16689905.PDF	1181323 0d673c0a51b64a7a9da6dec238ebe98d483c6031	no	18
Warnings:					
Information:					
4	Foreign Reference	WO2005046145_16811218.PDF	1189526 12556d88fc51d8167206348609ca3e7c4e53943b	no	27
Warnings:					
Information:					
			Total Files Size (in bytes):	4061681	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

PCT

To:
Wright, Bradley C.
BANNER & WITCOFF, LTD.
1100 13th Street, N.W. Suite 1200
Washington, DC 20005-4051
ETATS-UNIS D'AMERIQUE

NOTIFICATION OF TRANSMITTAL OF THE INTERNATIONAL SEARCH REPORT AND THE WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

Applicant's or agent's file reference 007742.00025	Date of mailing *(day/month/year)* 24 March 2014 (24-03-2014) **FOR FURTHER ACTION** See paragraphs 1 and 4 below
International application No. PCT/US2013/072566	International filing date *(day/month/year)* 2 December 2013 (02-12-2013)
Applicant CENTRIPETAL NETWORKS, INC.	

1. [X] The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

   Filing of amendments and statement under Article 19:
   The applicant is entitled, if he so wishes, to amend the claims of the International Application (see Rule 46):

   When? The time limit for filing such amendments is normally two months from the date of transmittal of the International Search Report.

   Where? Directly to the International Bureau of WIPO, 34 chemin des Colombettes 1211 Geneva 20, Switzerland, Fascimile No.: (41-22) 338.82.70

   For more detailed instructions, see *PCT Applicant's Guide*, International Phase, paragraphs 9.004 - 9.011.

2. [ ] The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.

3. [ ] **With regard to any protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

   [ ] the protest together with the decision thereon has been transmitted to the International Bureau together with any request to forward the texts of both the protest and the decision thereon to the designated Offices.

   [ ] no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. **Reminders**

   The applicant may submit comments on an informal basis on the written opinion of the International Searching Authority to the International Bureau. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established. Following the expiration of 30 months from the priority date, these comments will also be made available to the public.

   Shortly after the expiration of **18 months** from the priority date, the international application will be published by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau before completion of the technical preparations for international publication (Rules 90*bis*.1 and 90*bis*.3).

   Within **19 months** from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase **until 30 months** from the priority date (in some Offices even later); otherwise, the applicant must, **within 20 months** from the priority date, perform the prescribed acts for entry into the national phase before those designated Offices.

   In respect of other designated Offices, the time limit of **30 months** (or later) will apply even if no demand is filed within 19 months.

   For details about the applicable time limits, Office by Office, see www.wipo.int/pct/en/texts/time_limits.html and the *PCT Applicant's Guide*, National Chapters.

Name and mailing address of the International Searching Authority European Patent Office, P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk Tel. (+31-70) 340-2040 Fax: (+31-70) 340-3016	Authorized officer ROTHENBÜCHER, Anita Tel: +31 (0)70 340-9872

RECEIVED
MAR 28 2014
BANNER & WITCOFF, LTD.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 182

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference	FOR FURTHER ACTION	see Form PCT/ISA/220 as well as, where applicable, item 5 below.
007742.00025		
International application No.	International filing date *(day/month/year)*	(Earliest) Priority Date *(day/month/year)*
PCT/US2013/072566	2 December 2013 (02-12-2013)	11 January 2013 (11-01-2013)
Applicant		
CENTRIPETAL NETWORKS, INC.		

This international search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 4 sheets.

[X] It is also accompanied by a copy of each prior art document cited in this report.

1. **Basis of the report**
   a. With regard to the **language,** the international search was carried out on the basis of:
      [X] the international application in the language in which it was filed
      [ ] a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))
   b. [ ] This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6*bis*(a)).
   c. [ ] With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.
2. [ ] **Certain claims were found unsearchable** (See Box No. II)
3. [ ] **Unity of invention is lacking** (see Box No III)
4. With regard to the **title,**
   [X] the text is approved as submitted by the applicant
   [ ] the text has been established by this Authority to read as follows:
5. With regard to the **abstract,**
   [X] the text is approved as submitted by the applicant
   [ ] the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority
6. With regard to the **drawings,**
   a. the figure of the **drawings** to be published with the abstract is Figure No. 2
      [ ] as suggested by the applicant
      [X] as selected by this Authority, because the applicant failed to suggest a figure
      [ ] as selected by this Authority, because this figure better characterizes the invention
   b. [ ] none of the figures is to be published with the abstract

RECEIVED
MAR 2 8 2014

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 183

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2013/072566

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L29/06
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 313 290 A1 (STONESOFT CORP [FI]) 21 May 2003 (2003-05-21) abstract paragraph [0010] paragraph [0022] paragraph [0024] paragraph [0029] paragraph [0040] paragraph [0025] ----- -/--	1-24

[X] Further documents are listed in the continuation of Box C. [X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
14 March 2014	24/03/2014

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Oliveira, Joel

1

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2013/072566

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2005/046145 A1 (ERICSSON TELEFON AB L M [SE]) 19 May 2005 (2005-05-19) abstract page 3, line 14 - line 28 page 4, line 10 - line 11 page 5, line 10 page 5, line 23 - line 32 page 7, line 2 - line 8 page 7, line 30 - page 8, line 8 page 11, line 20 - line 25	1-24
A	US 2011/055916 A1 (AHN DAVID K [US]) 3 March 2011 (2011-03-03) paragraph [0017] paragraph [0021] - paragraph [0022]	1-24

1

Form PCT/ISA/210 (continuation of second sheet) (April 2005)

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 185

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2013/072566

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 1313290 A1	21-05-2003	AT 346446 T	15-12-2006
		DE 60216218 T2	08-03-2007
		EP 1313290 A1	21-05-2003
		US 2003097590 A1	22-05-2003
WO 2005046145 A1	19-05-2005	CN 1879361 A	13-12-2006
		EP 1683313 A1	26-07-2006
		JP 4323523 B2	02-09-2009
		JP 2007534219 A	22-11-2007
		US 2005111434 A1	26-05-2005
		WO 2005046145 A1	19-05-2005
US 2011055916 A1	03-03-2011	AU 2010297968 A1	22-03-2012
		EP 2471218 A2	04-07-2012
		US 2011055916 A1	03-03-2011
		US 2014007216 A1	02-01-2014
		WO 2011038420 A2	31-03-2011

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2013/072566	02.12.2013	11.01.2013

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

RECEIVED
MAR 26 2014
BANNER & WITCOFF, LTD.

Name and mailing address of the ISA:	Date of completion of this opinion	Authorized Officer
European Patent Office P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk - Pays Bas Tel. +31 70 340 - 2040 Fax: +31 70 340 - 3016	see form PCT/ISA/210	Oliveira, Joel Telephone No. +31 70 340-3334

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed

   ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43bis.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing filed or furnished:

   a. (means)

      ☐ on paper

      ☐ in electronic form

   b. (time)

      ☐ in the international application as filed

      ☐ together with the international application in electronic form

      ☐ subsequently to this Authority for the purposes of search

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	2, 3, 6-9, 11, 13, 14, 17-22, 24
	No:	Claims	1, 4, 5, 10, 12, 15, 16, 23
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-24
Industrial applicability (IA)	Yes:	Claims	1-24
	No:	Claims	

2. Citations and explanations

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

Reference is made to the following documents:

D1 EP 1 313 290 A1 (STONESOFT CORP [FI]) 21 May 2003 (2003-05-21)

D2 WO 2005/046145 A1 (ERICSSON TELEFON AB L M [SE]) 19 May 2005 (2005-05-19)

D3 US 2011/055916 A1 (AHN DAVID K [US]) 3 March 2011 (2011-03-03)

1 The present application does not meet the criteria of Article 33(2) PCT, because the subject-matter of claims **1, 12 and 23** is not new.

1.1 **D1** discloses a method comprising, at a network protection device **(abstract)**:

receiving a first rule set **(paragraph [0010], column3 lines 4-7)**;

receiving a second rule set **(paragraph [0010], column3 lines 4-7)**;

preprocessing the first rule set and the second rule set **(considered to be implicit from paragraphs [0010], [0013], [0040]: when a new set of rules is made available it must be made certain that they do not contain errors which could stop the filtering of the firewall)**;

configuring the network protection device to process packets in accordance with the first rule set **(paragraph [0010], column3 lines 13-18)**;

receiving packets **(column 6 lines 23-26)**;

processing a first portion of the packets in accordance with the first rule set **(paragraphs [0010], [0020])**;

reconfiguring the network protection device to process packets in accordance with the second rule set **(paragraphs [0010], [0024])**; and

processing a second portion of the packets in accordance with the second rule set **(paragraphs [0010], [0024])**.

1.2 The same reasoning applies, mutatis mutandis, to the subject-matter of the corresponding independent apparatus and computer-readable media claims **12 and 23**, which therefore are also considered not new.

1.3 Dependent claims **2-11, 13-22 and 24** do not contain any features which, in combination with the features of any claim to which they refer, meet the requirements of the PCT in respect of novelty (for claims **4, 5, 10, 15, 16**) and/or inventive step (for claims **2, 3, 6, 7, 8, 9, 11, 13, 14, 17, 18, 19, 20, 21, 22, 24**) see passages cited in the Search Report.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431
Examiner Name	Not Yet Assigned
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20060048142		2006-03-02	Roese et al.	
	2	20100011434		2010-01-14	Kay	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	2385676	EP	A1	2011-11-09	Alcatel Lucent		☐
	2	2498442	EP	A1	2012-09-12	Openet Telecom Ltd		☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2431
	Examiner Name	Not Yet Assigned
	Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	ISR off International Application No. PCT/US2014/027723, dated June 26, 2014.	☐
	2	"SBIR Case Study: Centripetal Networks Subtitle: How CNI Leveraged DHS S&T SBIR Funding to Launch a Successful Cyber Security Company, Cyber Security Division, 2012 Principal Investigators' Meeting"; Sean Moore, October 10, 2012.	☐
	3	"Control Plane Policing Implementation Best Practices"; Cisco Systems; March 13, 2013; <https://web.archive.org/web/20130313135143/http:www.cisco.com/web/about/security/intelligence/coppwp_gs.html>	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2431
	Examiner Name	Not Yet Assigned
	Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2014-07-03
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Espacenet

Bibliographic data: EP2385676 (A1) — 2011-11-09

Method for adapting security policies of an information system infrastructure

Inventor(s): CUPPENS FREDERIC [FR]; CUPPENS NORA [FR]; DUBUS SAMUEL [FR]; KANOUN WAEL [FR] ± (CUPPENS, FREDERIC, ; CUPPENS, NORA, ; DUBUS, SAMUEL, ; KANOUN, WAEL)

Applicant(s): ALCATEL LUCENT [FR]; INST TELECOM TELECOM BRETAGNE [FR] ± (ALCATEL LUCENT, ; INSTITUT TELECOM - TELECOM BRETAGNE)

Classification: **- international:** ***G06F21/57; H04L29/06***
- cooperative: **G06F21/577**; **H04L63/1408**; **H04L63/1433**; **H04L63/20**; H04L63/1441

Application number: EP20100290250 20100507

Priority number (s): EP20100290250 20100507

Also published as: US2013111548 (A1) KR20130005301 (A) JP2013525927 (A) WO2011138417 (A1) CN102934122 (A)

Abstract of EP2385676 (A1)

The present invention refers to a method for adapting security policies of an information system infrastructure in function of attacks wherein it comprises the steps of: - storing potential attacks and their associated risks in a data repository (125); - storing curative security policies (128) in response of the potential attacks in a data repository; - monitoring (101) entering contents representing data streams of the information system; - detecting (129) at least one attack in the information system; - assessing a success probability parameter (132) of the at least one detected attack and its associated cost impact parameter (133); -

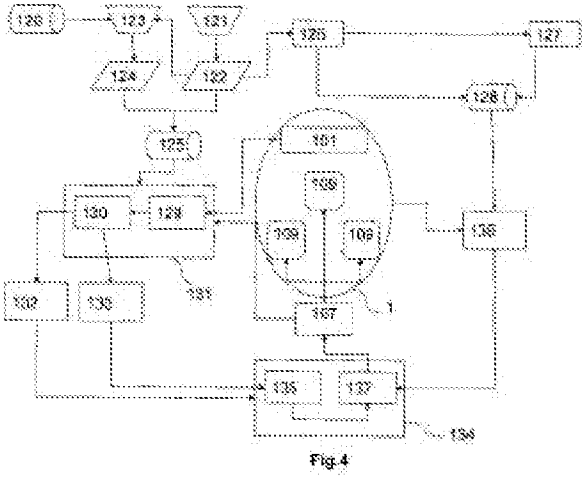

Fig 4

assessing an activation impact parameter (136) of at least one curative security policy in response to the at least one detected attack and its associated cost impact parameter; - deciding of the activation or deactivation (134) of a curative security policy in function of the success probability parameter of the, at least one, detected attack, of the activation impact parameter of at least one curative security policy and of the cost impact parameters of both the detected at least one attack and the at least one curative security policy.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 2 385 676 A1

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
09.11.2011 Bulletin 2011/45

(51) Int Cl.:
H04L 29/06 *(2006.01)* ***G06F 21/00*** *(2006.01)*

(21) Application number: **10290250.9**

(22) Date of filing: **07.05.2010**

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO SE SI SK SM TR
Designated Extension States:
BA ME RS

(71) Applicants:
- **Alcatel Lucent**
  75008 Paris (FR)
- **Institut Télécom - Télécom Bretagne**
  29238 BREST cedex 3 (FR)

(72) Inventors:
- **Cuppens, Frédéric**
  29238 Brest Cedex 3 (FR)
- **Cuppens, Nora**
  29238 Brest Cedex 3 (FR)
- **Dubus, Samuel**
  91620 Nozay (FR)
- **Kanoun, Wael**
  91620 Nozay (FR)

(74) Representative: **Croonenbroek, Thomas Jakob et al**
Cabinet Innovincia
11, avenue des Tilleuls
74200 Thonon-les-Bains (FR)

(54) Method for adapting security policies of an information system infrastructure

(57) The present invention refers to a method for adapting security policies of an information system infrastructure in function of attacks wherein it comprises the steps of:
- storing potential attacks and their associated risks in a data repository (125);
- storing curative security policies (128) in response of the potential attacks in a data repository;
- monitoring (101) entering contents representing data streams of the information system;
- detecting (129) at least one attack in the information system;
- assessing a success probability parameter (132) of the at least one detected attack and its associated cost impact parameter (133);
- assessing an activation impact parameter (136) of at least one curative security policy in response to the at least one detected attack and its associated cost impact parameter;
- deciding of the activation or deactivation (134) of a curative security policy in function of the success probability parameter of the, at least one, detected attack, of the activation impact parameter of at least one curative security policy and of the cost impact parameters of both the detected at least one attack and the at least one curative security policy.

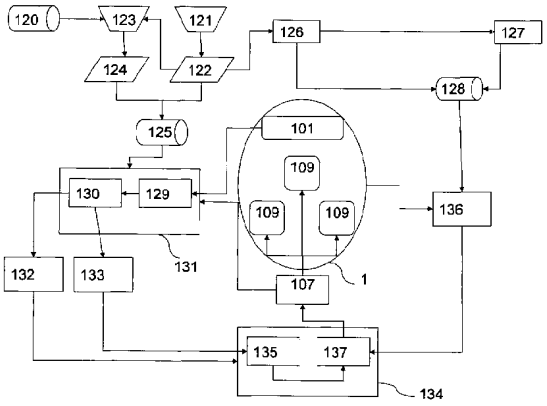

Fig.4

EP 2 385 676 A1

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 197

Description

[0001] The present invention relates to the field of information system protection and more precisely to the management of security policy in function of attacks undergone or being undergone by the information system.

[0002] Infrastructures of information systems need to be protected from harmful attacks leading to malicious events such as intrusions, data thefts, viruses or worms...

[0003] Due to the potential number of alerts generated by the existing attack detection systems with large information system infrastructures, it becomes impossible for operators to assess in real-time the risk of an attack and to decide of the suitable response to apply in response to the attack. Thus, automatic deployment of formally defined operational security policies starts to be considered in the protection of telecommunication and information infrastructures.

[0004] Fig.1 represents an example of such automatic protection of the state of the art.

[0005] The first step 101 corresponds to the detection of attacks toward the monitored information system which leads to the creation of elementary alerts (102). An alert correlation is then processed (103) to define correlated alerts (104) that are sent to a policy instantiation engine (105) to activate the appropriate security rules (106). These rules are sent to a policy decision point (107) which generates the configuration scripts (108) which are then used to configure policy enforcement points (109). Said policy enforcement points (109) are located in the information system 1 and apply the security rules in response to the detected attacks.

[0006] Such configuration of automatic policy activation suffers from drawbacks. Indeed, it is based only on correlated alerts and the number of correlated alerts may reach a very large number (up to thousands within a single day with large system) which would lead to thousands of security policy activations. Moreover, the deactivation of said security policies is not taken into account in the configurations of the state of the art such that a security policy may remain activated even if its impact on the users of the information system 1 is worth than the impact of the attack.

SUMMARY OF THE INVENTION

[0007] One object of the present invention is therefore to overcome the precited drawbacks of the state of the art and offer a method that allow to dynamically trig security policy activations only when it is necessary by taking into account a plurality of parameters influencing the activation decision and by defining a method that allow the deactivation of a security policy when it is necessary.

[0008] This is achieved by a method for adapting security policies of an information system infrastructure in function of attacks wherein it comprises the steps of:

- storing potential attacks and their associated risks in a data repository;
- storing curative security policies in response of the potential attacks in a data repository;
- monitoring entering contents representing data streams of the information system;
- detecting at least one attack in the information system;
- assessing a success probability parameter of the at least one detected attack and its associated cost impact parameter;
- assessing an activation impact parameter of at least one curative security policy in response to the at least one detected attack and its associated cost impact parameter;
- deciding of the activation or deactivation of a curative security policy in function of the success probability parameter of the, at least one, detected attack, of the activation impact parameter of at least one curative security policy and of the cost impact parameters of both the detected at least one attack and the at least one curative security policy.

[0009] According to another aspect of the invention, the step of storing potential attacks and their associated risks in a data repository comprises the steps of:

- defining the information system topology and attack detection signatures;
- defining a risk analysis of the information system that identifies potential attack objectives;
- specifying attack models to reach the identified attack objectives;
- storing said attack models in a data repository.

[0010] According to a further aspect of the invention, the step of storing curative security policies in response of the potential attacks in a data repository comprises:

- specifying at least one attack context;
- specifying curative security policies corresponding to the specified at least one attack context;
- storing said curative security policies in a data repository.

[0011] According to an additional aspect of the present invention, the step of assessing a success probability parameter of the at least one detected attack and its associated cost impact parameter comprises:

- generating attack strategy graphs based on the stored attack models and the detected at least one attack;
- assessing the probability for the attack to reach its objective;
- assessing the impact of the attack objective on the system security level and on the system quality of service (QoS) level;
- assessing the associated cost impact parameter of

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 198

the attack objective;

[0012] According to another aspect of the present invention, the step of assessing an activation impact parameter of at least one curative security policy in response to the at least one detected attack and its associated cost impact parameter is based on the stored curative security policies and the state of the monitored information system.

[0013] According to a further aspect of the invention, the step of deciding of the activation or deactivation of a curative security policy in function of the success probability parameter of the, at least one, detected attack, of the activation impact parameter of at least one curative security policy and of the cost impact parameters of both the detected at least one attack and the at least one curative security policy is applied dynamically based on the evolution of the state of the monitored system.

[0014] According to an additional aspect of the invention, the cost impact parameters comprise:

- a quality of service (QoS) impact and,
- a security level degradation impact.

[0015] The present invention also refers to a monitoring and protecting equipment comprising:

- at least one data repository for:
  - storing potential attacks and their associated risks;
  - storing curative security policies in response of the potential attacks;
- processing means for:
  - monitoring entering contents representing data streams of the information system;
  - detecting at least one attack in the information system;
  - assessing a success probability parameter of the detected at least one attack and its associated cost impact parameter;
  - assessing at least one activation impact parameter of at least one curative security policy in response to the at least one detected attack and its associated cost impact parameter;
  - deciding of the activation of a curative security policy in function of the success probability parameter of the at least one detected attack, of the activation impact parameter of at least one curative security policy and of the cost impact parameters of both the at least oneattack and the at least one curative policy,
  - activating at least one curative security policy.

[0016] The present invention also refers to a monitoring and protecting equipment comprising:

- at least one data repository for:
  - storing potential attacks and their associated risks;
  - storing curative security policies in response of the potential attacks;
- processing means for:
  - monitoring entering contents representing data streams of the information system;
  - detecting at least one attack in the information system;
  - assessing a success probability parameter of the detected at least one attack and its associated cost impact parameter;
  - assessing at least one activation impact parameter of at least one curative security policy in response to the at least one detected attack and its associated cost impact parameter;
  - deciding of the deactivation of a curative security policy in function of the success probability parameter of the at least one detected attack, of the activation impact parameter of at least one curative security policy and of the cost impact parameters of both the at least one attack and the at least one curative policy,
  - deactivating at least one curative security policy.

BRIEF DESCRIPTION OF THE DRAWINGS

[0017]

FIG.**1** is a synoptic diagram of the different steps of a protective method for launching security policies against attacks of an information system according to the state of the art;

FIG.**2** is a synoptic diagram of the different steps of a protective method for adapting security policies of an information system infrastructure against attacks according to the present invention;

FIG.**3** is a diagram of an example of attack graph comprising the different attack steps and the objectives associated with these attacks;

FIG.4 is a synoptic diagram of the different steps of a detailed protective method for adapting security policies of an information system infrastructure against attacks according to the present invention;

DETAILED DESCRIPTION OF THE INVENTION

[0018] As used herein, the term "attack" refers to an event in a system that transgress the normal authorized usage of the system or exploits deliberately or accidentally a vulnerability in the system as for example a network scanning, a password cracking, a sending of malicious email (also called spam), a sending of a malformed internet protocol (IP) packet... Moreover, the expression

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 199

attack with respect to an information system refers to attacks initiated from outside or inside of the information system (usually from an attacker machine) and directed toward said information system in order to produce dysfunctionings in said system.

[0019] As used herein, the term "SIP" refers to the acronym session Internet protocol.

[0020] As used herein, the term "IP" refers to the acronym Internet protocol.

[0021] As used herein, the term "QoS" refers to the acronym quality of service.

[0022] The embodiments of the present invention refer to a method for activating and deactivating security policies in order to protect an information system against malicious attacks wherein not only impacts of the attacks but also the success likelihood for the attack to reach its objective and impacts of security policies are taken into account in the activation/deactivation decision in order to minimize the impact, and in particular the cost, of attacks on the users of the information system.

[0023] Fig.2 represents the general steps of the method for adapting the protection of the information system 1.

[0024] Step 110 refers to the specification and the storing in a data repository of potential attacks and their associated risks. This first step allows to provide attack models that define the strategy and the objectives of the possible attacks.

[0025] Step 111 refers to specification and storing in a data repository of the foreseen security policies to apply in response to the possible attacks defined in the previous step.

[0026] Step 112 refers to the monitoring of entering content of the information system 1 corresponding to data streams and the attacks detection.

[0027] Step 113 refers to the assessment of a success probability of detected attacks and the cost impact associated with these attacks. Such assessment is achieved based on the attack models stored in step 110 and the attacks detected in step 112.

[0028] Step 114 refers to the assessment of the impact of an activation of a curative security policy in response to a detected attack and the cost impact associated with this activation. Such assessment is achieved based on the security policies stored in step 111 and the state of the monitored information system 1.

[0029] Step 115 refers to the decision of activating or deactivating a curative security policy in function of the assessments achieved in steps 113 and 114. Such decision corresponds to the comparison of the cumulative impacts of both the attack and the curative security policy on the information system 1.

[0030] The activation or deactivation decided in step 115 is then applied in step 116.

[0031] Moreover, it has to be noted that the steps 110 and 111 are preliminary steps which can be achieved offline whereas the steps 112, 113, 114, 115, 116 are dynamic steps which are achieved online such that any modification of the information system is taken into account in real time to decide of the activation or the deactivation of security policies.

[0032] Thus, the present invention, thanks to the assessments of both the attack consequences and the consequences of the activation of a security policy allows to determine if it is worthy to activate a curative security policy and when said policy is activated to determine if it is worthy to keep this policy activated in function of the evolution of the information system and the impact of this policy on said information system.

[0033] As described above, an attack is an event occurring in a system. In practice, an attack comprises different levels (or steps) that lead to the objective of the attack as described in fig.3 where an example of attack graph aimed at hacking a voice over internet protocol (VoIP) system is described.

[0034] The first level 201 corresponds to the sending of an email (malicious email or spam) toward a victim machine. Then the below (or next) levels represent the different steps to reach the objectives (213, 214, 215 and 216). The second level 202 refers to the gain of a remote shell in the victim machine (corresponding to the opening by the user of the malicious link or attachment of the email). The third level 203 corresponds to a bot (robot) infection which is the installation of a malicious software bot capable of simulating human activity that uses the remote shell. The bot can wait for future orders from the attacker to execute. As represented in fig.3, several remote shells and several corresponding bot infections may be develop in parallel in the information system.

[0035] The next level 204 corresponds to the discovery of the session initial protocol (SIP) discovery which is a scanning of the system (or network) to discover machines (computers in general) or servers using SIP protocol. Such attack may be done by a bot.

[0036] From this step, the attack may use two different ways depending on its objective.

[0037] The first way leads to step 205 which refers to the SIP fingerprinting which consists of identifying the type and the version of the software (operating systems, softphones, servers, etc...) installed on the SIP entities discovered at the previous level (204). This level may also be performed by a bot.

[0038] From step 205, the attack may have two different ways again, one leading to the objective of spam over IP (SPIT) 213 and comprising two levels, a discovery of the active users which is a scanning to determine the users of the VoIP system and which may be performed by a bot and a direct call 209 referring to a call to the victim user (by a bot) to perform spam over IP (SPIT).

[0039] The other possibility from level 205 is the discovery of the media access control (MAC) address of potential victim machines 207 (which may be performed by a bot).

[0040] The next step 208 is then the address resolution protocol (ARP) poisoning which forces the traffic (e.g. the established calls) between two victim users to pass through one bot. The attacker may then have access to

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 200

all the traffic between both victim users. This leads either to a step of sniff audio 211 to sniff the audio packets transmitted between both users or a step of inject audio 212 to inject audio packets in the transmission between both users for altering the established call. The corresponding respective objectives being the conversation tapping 214 and the conversation injection 215.

[0041] From level 204 the second attack solution consists in a server flooding 210 wherein the bots flood the main server of the VoIP system to cause a denial of service (DoS) which is the objective 216.

[0042] It has to be noted that the closer to the objective is the attack, the higher the probability to reach the objective. Thus, when a spam (level 201) is detected, it is generally not worth activating a security policy as the probability for the attack to reach one of the objectives (213, 214, 215 or 216) is still low whereas if the server is flooded (level 210) for example, then the probability of reaching the objective 216 is very high.

[0043] The graphs such as presented in fig.3 are achieved automatically based on attack models stored in data repository and are used in the assessment of the impact of an attack and its probability to reach its objective as described previously.

[0044] In order to better understand the different steps of the invention, a detailed configuration of a possible embodiment of the invention will now be described based on fig.4, it refers to a more detailed presentation of the organization presented in fig.2.

[0045] Step 120 corresponds to the storing in a data repository of the information system topology and the attack or intrusion detection signatures (IDS). Step 121 is a risk analysis for the monitored system which is conducted by an expert to identify potential attack objectives 122. Based on these attack objectives and the system topology (stored in step 120), attack models are specified (by an expert) 123 as well as an attack context specification 126. Attack models such as presented in fig.4 are then set up 124 and stored as elementary models in a data repository 125.

[0046] On the other hand, the specification of attack contexts 126 allows to specify response contexts 127. For each identified attack context an appropriate response context is specified with the associated security rules. Said security rules and attack contexts are stored as security policies in a data repository 128. Said security policies refer, for example, to firewall activation, authentication request activation or user account blocking or any actions allowing to cure the information system or reduce the impact of the attack. Based on the stored attack models and security policies to apply in response to each attack, the automatic and dynamic part of the invention may be implemented.

[0047] The information system 1 is monitored by intrusion detection systems (IDS) 101 which correspond to any types of sensors capable of detecting an attack and of generating and sending alerts. Said sensors may be made of specific electronic devices (micro-controllers, application specific integrated systems (ASICs)...) or in an information technology (IT) environment, it may simply be a software run by a computer, a server or a router.

[0048] The alerts sent by the IDS are then aggregated 129 in an online correlation engine 131 which uses these aggregated alerts and the attack models stored in a data repository to generate attack graphs such as presented in fig.3. The generated graphs are used, on one hand, for the assessment of the probability for the attack to reach its objective 132 and, on the other hand, for the assessment of the impact of the attack 133.

[0049] The determination of the probability to reach the objective corresponds to the success likelihood for an attack to reach an objective. Thus, the success likelihood level of each attainable objective (predefined in the attack graph) of the attack is calculated. This level shows how close the attack is with respect to its objective.

[0050] The attack impact assessment corresponds to the determination of the impact of the attack on the security and the quality of service (QoS) of the information system 1. This gathers the impact on the confidentiality, the integrity and the availability and any parameters influencing the using of the information system 1. Such assessment is also based on the state of the information system 1.

[0051] The results of both assessments are then sent to a policy instantiation engine (PIE) 134 where the cumulative impact of the attacks is determined.

[0052] Besides, in the same way as the assessment of the impact of the attacks, an assessment of the impact of the curative security policies which can be used in response of the attacks 136 is done based on the stored security policies. Again, impact on both the security and the QoS of the information system as well as as the state of said system are taken into account in the assessment. The results of the assessment are also transmitted to the PIE 134 where the cumulative impact of the curative security policies is determined.

[0053] The policy instantiation engine 134 receives the assessments corresponding on one side of the attack impact and on the other side of the curative security policy impact which are processed respectively by an attack handler module 135 and a response handler module 137. For each type of impact, at least one metric is defined, and in the general case, a plurality of metrics corresponding to different aspects of the impact (on security, on QoS...) are defined and the assessed value of all these impact contributions are added (practically by integrated their value over time).

[0054] Based on both cumulative impacts computed by the attack handler module and the response handler module as well as the probability for the attack to reach its objective (also called success likelihood) and the state of the monitored information system 1, the PIE 134 determines if curative security policies have to be activated or deactivated.

[0055] Different activation/deactivation rules may be used, for example, a security policy may be activated if

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 201

the success likelihood reached a given threshold, or if the attack impact reach a predetermined threshold. It can also be a combination of both (activation if the attack impact and the success likelihood reach given thresholds).

[0056] In general, a response policy must be activated when:

- the detected threat violated the security and operational policies and,
- the success likelihood (calculated dynamically, considering the attack progress and the state of the monitored system) of the threat exceeds a predefined threshold and,
- the impact of the threat exceeds a predefined threshold, and is greater than the cost of the associated response policy.

[0057] In the same way, if the success likelihood goes below a given threshold or the response impact goes over a predefined threshold or a combination of both, an activated security policy have to be deactivated. Such case may occur if the impact of a response policy induced more drawbacks than advantages with respect to the initial attack or if the risks associated with the attack have been dismissed or eradicated.

[0058] The order or command of activating or deactivating a security policy is then sent to response policy deployment point (PDP) 107 (also called policy decision point) wherein each security rule of a response policy is converted into scripts to configure the policy enforcement points (PEP) 109 which are located within the information system 1 and which are used to enforce the security policies (firewall activation or configuration, intrusion prevention systems activations, account permissions or access modifications...). The online steps (comprising the online correlation engine 131, the attack objective probability assessment 132, the attack impact assessment, the attack response assessment 136, the policy instantiation engine 134, the policy deployment point and the policy enforcement point) may be achieved by programs or software run by a computer or a server. Said online steps use the stored elements of the data repositories (attack model data repository 125 and response policy data repository 128) to determine dynamically the necessity of activation/deactivation of response security policy in function of the detected attacks and of the state of the monitored information system 1.

[0059] Thus, the present invention allows to provide a dynamic assessment of the impact of an attack as well as the impact of the security policy to apply in response to the attack providing therefore an optimized efficiency of the use of a security policy. Moreover, the idea of determining conditions for deactivating a security policy allows to improve the reactivity of the information system protection and to avoid an unnecessary degraded use of the information system. Indeed, security policy usually correspond to a degraded mode that limits the effects of the attack but also reduces the capacity of said information system and may disturb and affect the users of the information system leading to a loss of productivity or a reduction of quality of service (QoS) for the company using the information system.

Claims

1. Method for adapting security policies of an information system infrastructure in function of attacks wherein it comprises the steps of:

   - storing potential attacks and their associated risks in a data repository;
   - storing curative security policies in response of the potential attacks in a data repository;
   - monitoring entering contents representing data streams of the information system;
   - detecting at least one attack in the information system;
   - assessing a success probability parameter of the at least one detected attack and its associated cost impact parameter;
   - assessing an activation impact parameter of at least one curative security policy in response to the at least one detected attack and its associated cost impact parameter;
   - deciding of the activation or deactivation of a curative security policy in function of the success probability parameter of the, at least one, detected attack, of the activation impact parameter of at least one curative security policy and of the cost impact parameters of both the detected at least one attack and the at least one curative security policy.

2. Method for adapting security policies of an information system infrastructure in accordance with claim 1 wherein the step of storing potential attacks and their associated risks in a data repository comprises the steps of:

   - defining the information system topology and attack detection signatures;
   - defining a risk analysis of the information system that identifies potential attack objectives;
   - specifying attack models to reach the identified attack objectives;
   - storing said attack models in a data repository.

3. Method for adapting security policies of an information system infrastructure in accordance with claim 1 or 2 wherein the step of storing curative security policies in response of the potential attacks in a data repository comprises:

   - specifying at least one attack context;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 202

- specifying curative security policies corresponding to the specified at least one attack context;
- storing said curative security policies in a data repository.

4. Method for adapting security policies of an information system infrastructure in accordance with one of the previous claims wherein the step of assessing a success probability parameter of the at least one detected attack and its associated cost impact parameter comprises:

- generating attack strategy graphs based on the stored attack models and the detected at least one attack;
- assessing the probability for the attack to reach its objective;
- assessing the impact of the attack objective on the system security level and on the system quality of service (QoS) level;
- assessing the associated cost impact parameter of the attack objective;

5. Method for adapting security policies of an information system infrastructure in accordance with one of the previous claims wherein the step of assessing an activation impact parameter of at least one curative security policy in response to the at least one detected attack and its associated cost impact parameter is based on the stored curative security policies and the state of the monitored information system.

6. Method for adapting security policies of an information system infrastructure in accordance with one of the previous claims wherein the step of deciding of the activation or deactivation of a curative security policy in function of the success probability parameter of the, at least one, detected attack, of the activation impact parameter of at least one curative security policy and of the cost impact parameters of both the detected at least one attack and the at least one curative security policy is applied dynamically based on the evolution of the state of the monitored system.

7. Method for adapting security policies of an information system infrastructure in accordance with one of the previous claims wherein the cost impact parameters comprise:

- a quality of service (QoS) impact and,
- a security level degradation impact.

8. Monitoring and protecting equipment comprising:

- at least one data repository for:

- storing potential attacks and their associated risks;
- storing curative security policies in response of the potential attacks;

- processing means for:

- monitoring entering contents representing data streams of the information system;
- detecting at least one attack in the information system;
- assessing a success probability parameter of the detected at least one attack and its associated cost impact parameter;
- assessing at least one activation impact parameter of at least one curative security policy in response to the at least one detected attack and its associated cost impact parameter;
- deciding of the activation of a curative security policy in function of the success probability parameter of the at least one detected attack, of the activation impact parameter of at least one curative security policy and of the cost impact parameters of both the at least one attack and the at least one curative policy,
- activating at least one curative security policy.

9. Monitoring and protecting equipment comprising:

- at least one data repository for:

- storing potential attacks and their associated risks;
- storing curative security policies in response of the potential attacks;

- processing means for:

- monitoring entering contents representing data streams of the information system;
- detecting at least one attack in the information system;
- assessing a success probability parameter of the detected at least one attack and its associated cost impact parameter;
- assessing at least one activation impact parameter of at least one curative security policy in response to the at least one detected attack and its associated cost impact parameter;
- deciding of the deactivation of a curative security policy in function of the success probability parameter of the at least one detected attack, of the activation impact parameter of at least one curative security pol-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 203

icy and of the cost impact parameters of both the at least one attack and the at least one curative policy
- deactivating at least one curative security policy.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 204

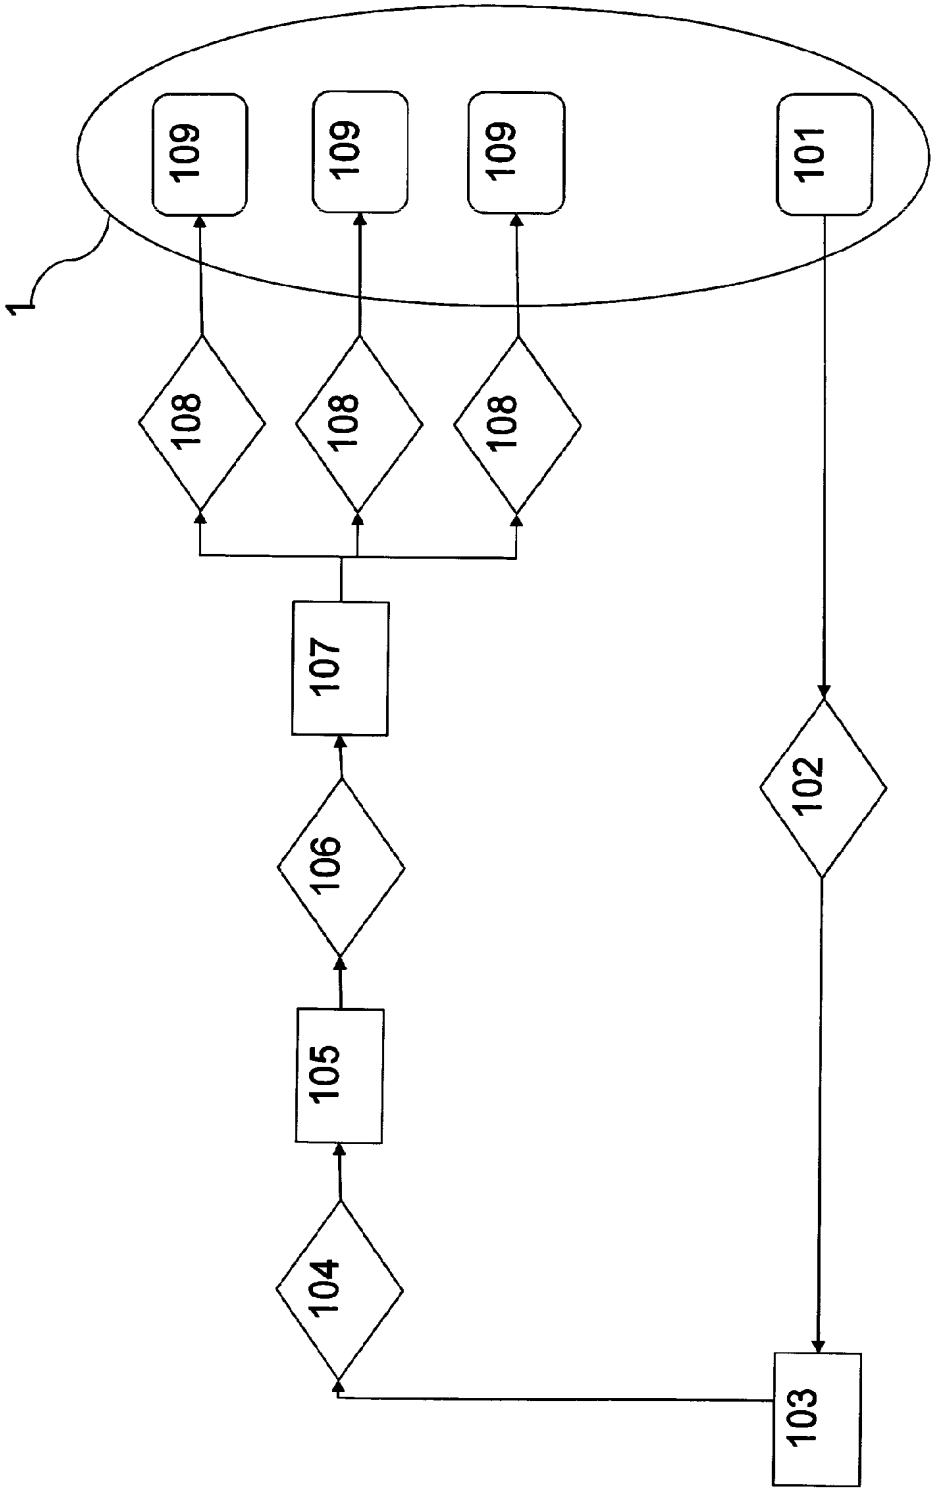

Fig.1

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 205

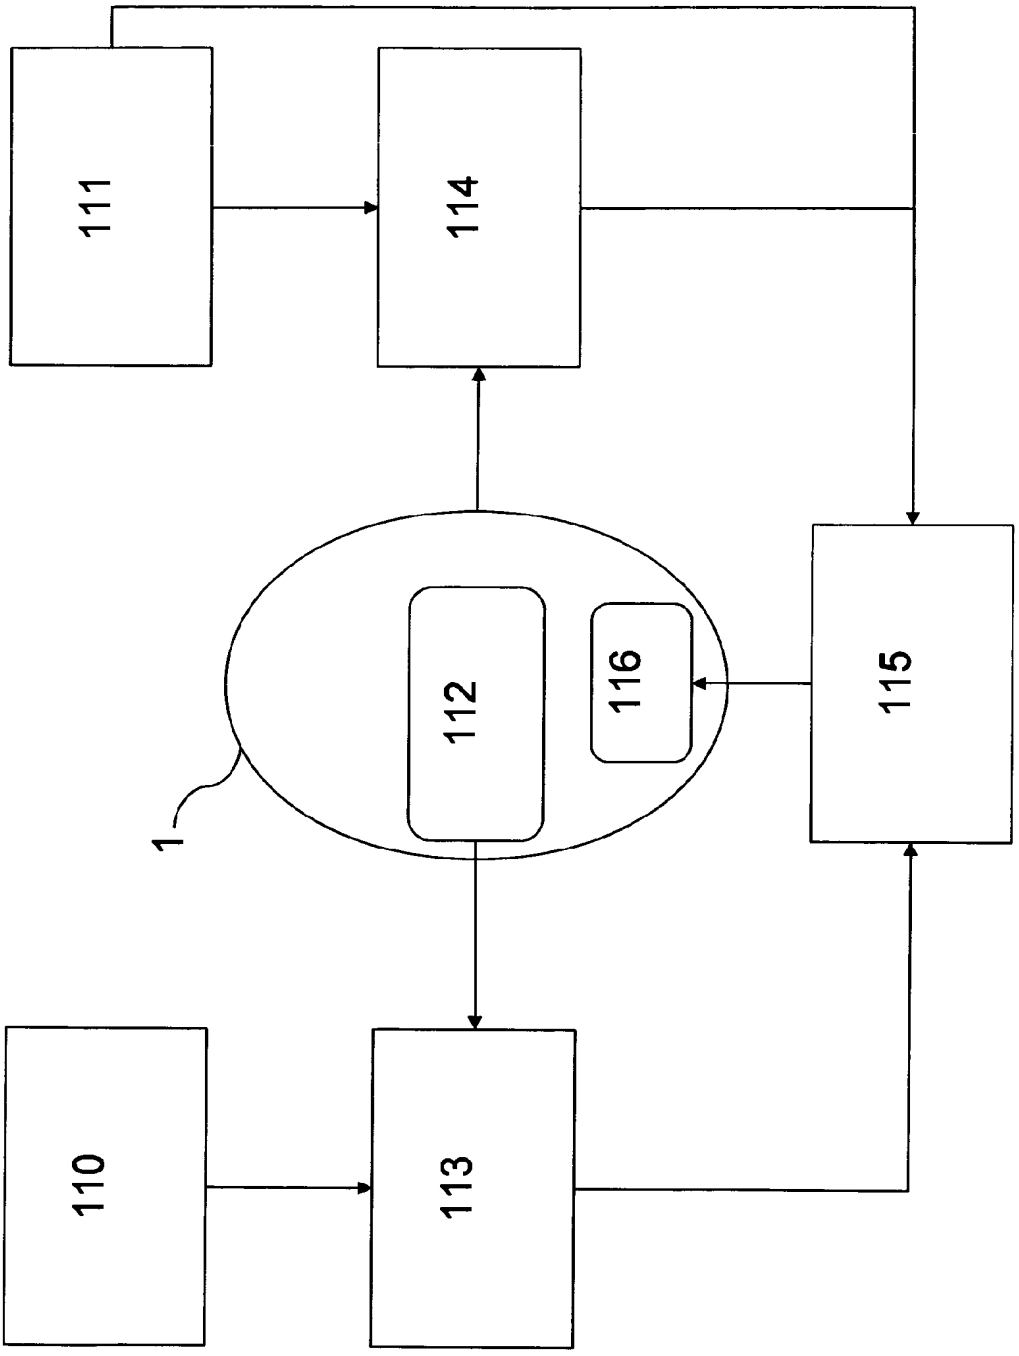

Fig.2

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 206

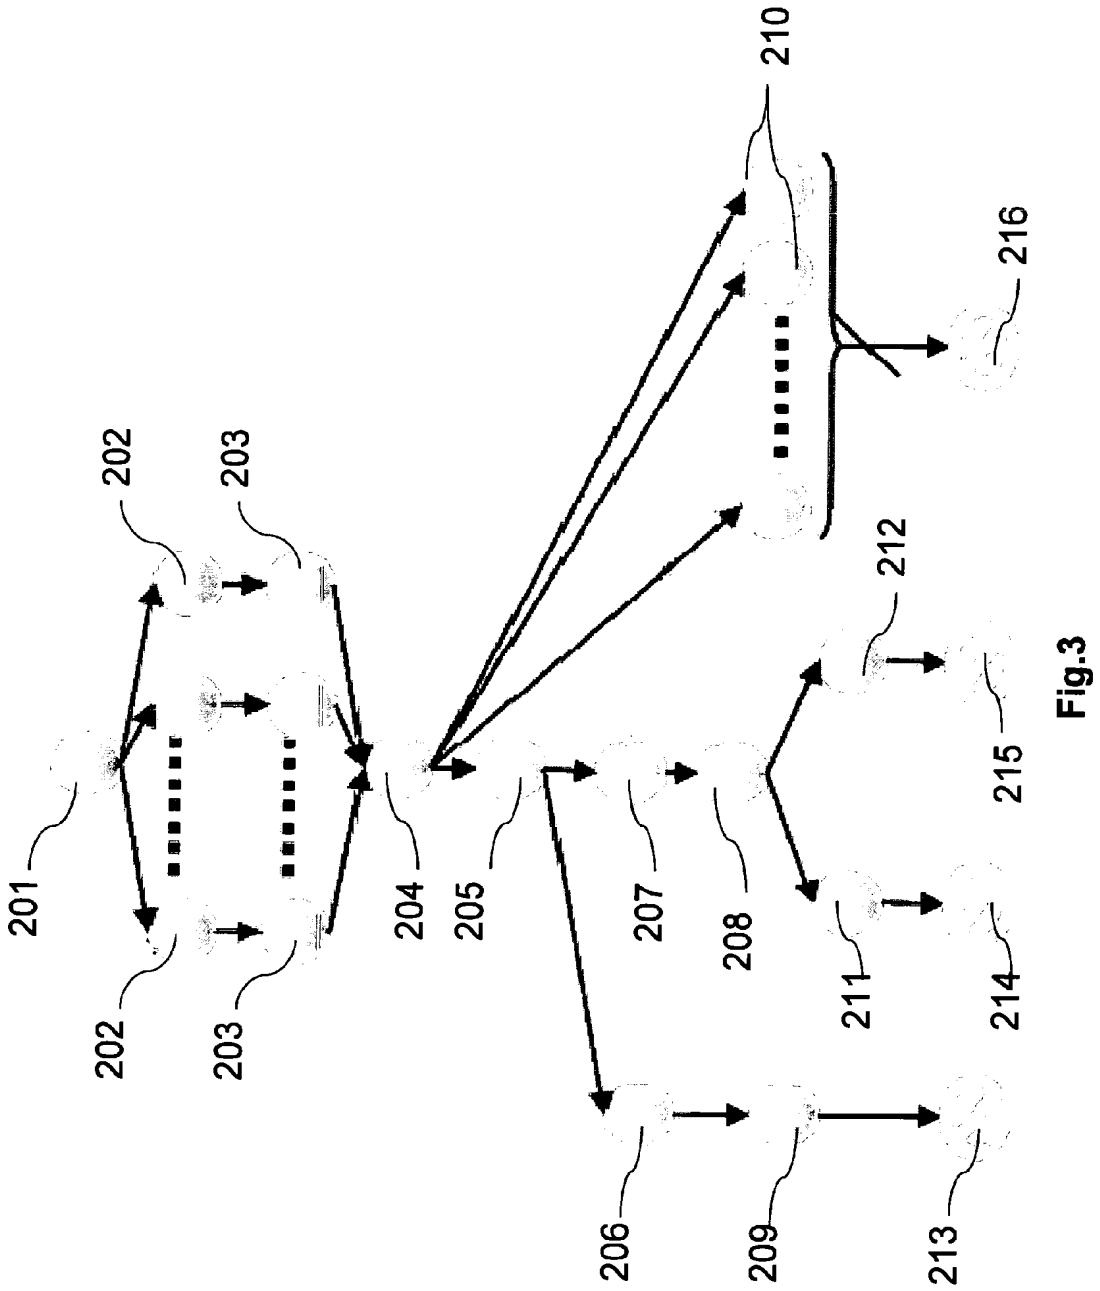

Fig.3

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 207

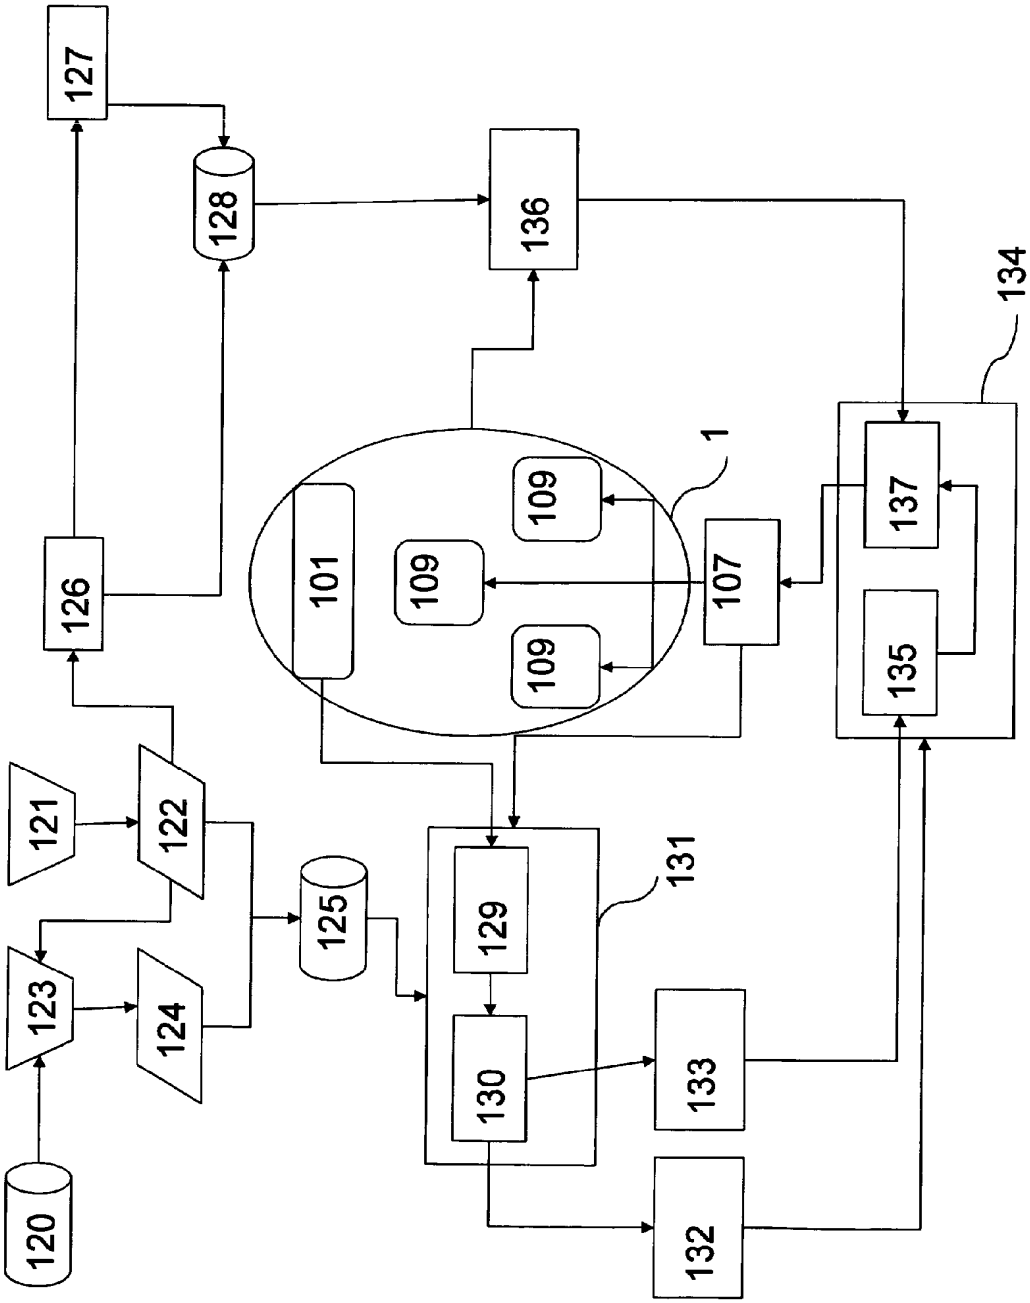

Fig.4

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 208

Europäisches Patentamt
European Patent Office
Office européen des brevets

EUROPEAN SEARCH REPORT

Application Number
EP 10 29 0250

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	WO 2004/015908 A2 (ALPHATECH INC [US]; KREIDL O PATRICK [US]; FRAZIER TIFFANY [US]) 19 February 2004 (2004-02-19) * abstract; figures 1-6; tables I, II * * paragraphs [0006], [0008], [0013] * * paragraphs [0021] - [0023] * * paragraphs [0037] - [0039] * * paragraphs [0052], [0053], [0062] * * paragraphs [0067], [0069], [0072] * * paragraphs [0073], [0076] * * paragraphs [0088] - [0110] * -----	1-9	INV. H04L29/06 G06F21/00
X	KANOUN W ET AL: "Success Likelihood of Ongoing Attacks for Intrusion Detection and Response Systems" COMPUTATIONAL SCIENCE AND ENGINEERING, 2009. CSE '09. INTERNATIONAL CONFERENCE ON, IEEE, PISCATAWAY, NJ, USA, 29 August 2009 (2009-08-29), pages 83-91, XP031543989 ISBN: 978-1-4244-5334-4 * abstract; figures 3,6,7 * * page 83, left-hand column, line 1 - right-hand column, line 37 * * page 84, left-hand column, line 42 - right-hand column, line 2 * * page 84, right-hand column, line 44 - page 85, left-hand column, line 15 * * page 87, left-hand column, line 35 - right-hand column, line 19 * -----	1-9	TECHNICAL FIELDS SEARCHED (IPC) H04L G06F G06N
A	WO 2007/027131 A2 (TELIASONERA AB [SE]; ENDERSZ VIKTOR [SE]) 8 March 2007 (2007-03-08) * figure 1 * * page 4, line 29 - page 8, line 34 * -----	1-9	
	-/--		

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
Munich	21 October 2010	Schossmaier, Klaus

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

4

EPO FORM 1503 03.82 (P04C01)

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 209

Europäisches Patentamt
European Patent Office
Office européen des brevets

EUROPEAN SEARCH REPORT

Application Number
EP 10 29 0250

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
A	MIGUEL JOHN: "COMPOSITE COST/BENEFIT/RISK ANALYSIS METHODOLOGY" COMPUTER SECURITY: A GLOBAL CHALLENGE; PROCEEDINGS / PROCEEDINGS OF THE IFIP INTERNATIONAL CONFERENCE ON COMPUTER SECURITY,, 1 January 1984 (1984-01-01), pages 307-311, XP009139953 ISBN: 978-0-444-87618-8 * the whole document * -----	1-9	

TECHNICAL FIELDS SEARCHED (IPC)

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
Munich	21 October 2010	Schossmaier, Klaus

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P04C01)

4

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 210

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO. EP 10 29 0250

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

21-10-2010

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2004015908 A2	19-02-2004	AU 2003278706 A1	25-02-2004
		CA 2495147 A1	19-02-2004
		EP 1535422 A2	01-06-2005
WO 2007027131 A2	08-03-2007	NONE	

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 211

Espacenet

Bibliographic data: EP2498442 (A1) — 2012-09-12

Methods, systems and devices for the detection and prevention of malware within a network

Inventor(s): RIESCHICK GARY [US]; DUNNE CAMERON ROSS [IE]; MCNAMEE ALAN [IE]; HOGAN JOE [IE] ± (RIESCHICK, GARY, ; DUNNE, CAMERON ROSS, ; MCNAMEE, ALAN, ; HOGAN, JOE)

Applicant(s): OPENET TELECOM LTD [IE] ± (OPENET TELECOM LTD)

Classification: **- international:** ***H04L12/24; H04L29/06; H04W12/08; H04W12/12;*** *H04W24/00; H04W88/16*
- cooperative: **H04L63/1441; H04W12/08; H04W12/12;** H04L63/1408; H04W88/16

Application number: EP20120158921 20120309

Priority number (s): US201161451775P 20110311 ; US201113159710 20110614

Also published as: US2012233656 (A1) US8726376 (B2)

Abstract of EP2498442 (A1)

Methods, systems and devices examine data flows in a communication system control network for known malware threats and suspicious properties typically associated with malware threats. A policy management system inside the control network accesses a user repository and a charging network, and performs pattern matching and/or observed behavior detection methods to determine if the data flows carry content (e.g., malware) that poses a security risk to network or wireless devices. The policy management system generates policy rules based on user preferences and risk-level. The policy management system sends the generated policy rules to a gateway/PCEF, which blocks the data flows, allows the data flows, or restricts the data flow based on the policy rules.

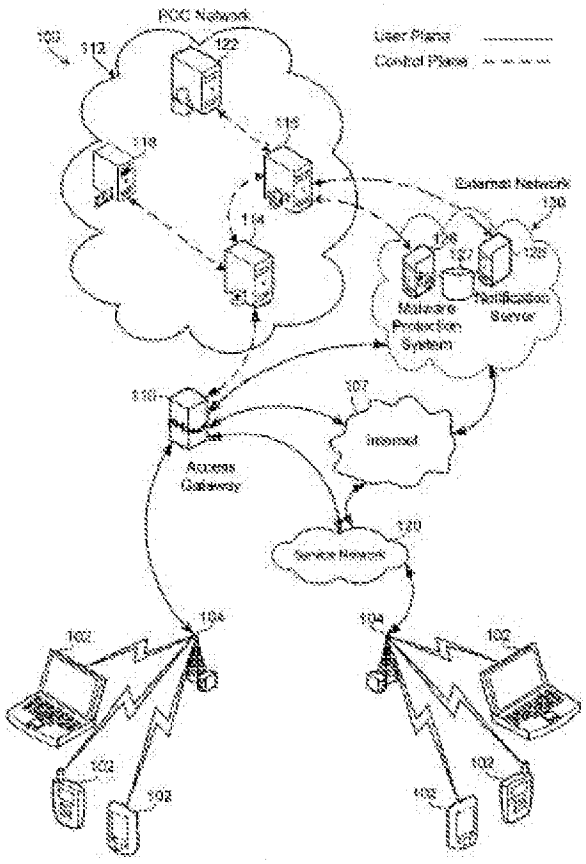
RNC Network
User Plane
Control Plane
External Network
Malware Protection System
Notification Server
Internet
Access Gateway
Service Network

FIG. 2

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 2 498 442 A1

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
12.09.2012 Bulletin 2012/37

(21) Application number: **12158921.2**

(22) Date of filing: **09.03.2012**

(51) Int Cl.:
H04L 12/24 *(2006.01)* ***H04L 29/06*** *(2006.01)*
H04W 12/08 *(2009.01)* ***H04W 12/12*** *(2009.01)*
H04W 24/00 *(2009.01)* *H04W 88/16* *(2009.01)*

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR
Designated Extension States:
BA ME

(30) Priority: **11.03.2011 US 201161451775 P**
14.06.2011 US 201113159710

(71) Applicant: **Openet Telecom Ltd.**
12 Dublin (IE)

(72) Inventors:
- **Rieschick, Gary**
  Louisburg, KS Kansas KS 66053 (US)
- **Dunne, Cameron Ross**
  County Kildare (IE)
- **McNamee, Alan**
  Dublin, 9 (IE)
- **Hogan, Joe**
  County Dublin (IE)

(74) Representative: **Kelly, Madeleine**
FRKelly
27 Clyde Road
Ballsbridge
Dublin 4 (IE)

(54) Methods, systems and devices for the detection and prevention of malware within a network

(57) Methods, systems and devices examine data flows in a communication system control network for known malware threats and suspicious properties typically associated with malware threats. A policy management system inside the control network accesses a user repository and a charging network, and performs pattern matching and/or observed behavior detection methods to determine if the data flows carry content (e.g., malware) that poses a security risk to network or wireless devices. The policy management system generates policy rules based on user preferences and risk-level. The policy management system sends the generated policy rules to a gateway/PCEF, which blocks the data flows, allows the data flows, or restricts the data flow based on the policy rules.

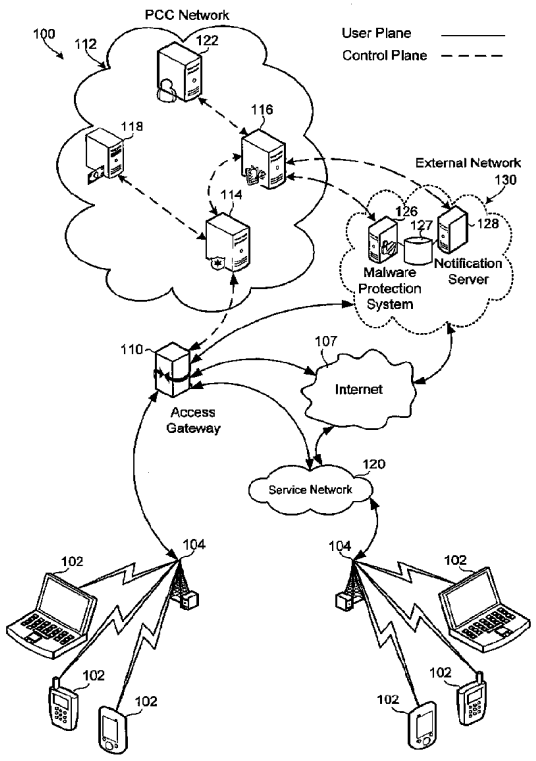

FIG. 2

EP 2 498 442 A1

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 214

Description

BACKGROUND

[0001] Recent advances in wireless and mobile technologies enable mobile-device users to access the Internet at any time and in any place. Internet usage habits are changing and many new types of Internet-enabled devices are emerging to meet the evolving user demands and changing Internet-usage habits. Many of these new and emerging Internet-enabled wireless devices use custom and/or non-traditional operating systems and access the Internet through mobile data networks. Further, many of these devices include powerful processors capable of executing third-party application software downloaded from digital application distribution platforms (e.g., Apple Inc.'s App Store, etc.). The emergence these new technologies and wireless devices presents additional security challenges for the telecommunications networks, mobile data networks and to the new wireless devices themselves.

[0002] Today, malicious software (herein "malware") may be used by thieves and hackers to steal confidential information, launch attacks on web servers and to flood networks with spam. Recently, there have been many high-profile cases of confidential customer information being stolen from businesses, and private wireless device users being subjected to online financial and identity thefts. Perhaps even more significantly, there has recently been an increase in wireless devices being infected by malware that forces the wireless device to participate in malicious activities without the user's knowledge or consent, causing users to incur higher usage-charges and leaving them vulnerable to potential legal actions. Recent press reports indicate this problem is increasing at an alarming rate. The emergence of these new types of Internet-enabled devices is likely to exacerbate these problems, as current malware detection and prevention techniques are ineffective at preventing or removing malware from wireless devices, especially the new breed of Internet-enabled wireless devices that have custom and/or non-traditional operating systems.

[0003] Further, the growth in popularity of application distribution platforms (herein "app ecosystems") allows wireless device users to more readily download third party software and execute the downloaded software with full administrative privileges on the mobile device. Such downloaded applications may be sourced from many different sources, such as email attachments, instant messaging systems, and device-to-device transfers. The emergence of app ecosystems presents additional security challenges that are distinct from the traditional security challenges presented by a malicious page being displayed within web browsers which could be managed by implementing a sandbox model.

[0004] Current malware detection and prevention techniques typically involve installing a protective software client application onto a traditional Internet-enabled device, such as a desktop computer or a laptop computer. Thus, current techniques aim to prevent malware by protecting user-level devices (e.g., laptops) rather than providing a comprehensive solution that can eliminate malware from within the network. Further, since many of the new and emerging Internet-enabled wireless devices do not have traditional operating systems, they cannot effectively employ existing malware detection and prevention techniques. As these new Internet-enabled devices grow in popularity, so will the importance of detecting and preventing malware from within the network.

SUMMARY

[0005] The various embodiments provide methods, systems and devices for receiving from one or more communication system control network components data flow parameters identifying characteristics of a data flow, performing a detection method to determine if the data flow carries malicious content, generating restrictive policy rules in response to determining the data flow carries malicious content, generating non-restrictive policy rules in response to determining the data flow does not carry malicious content, and pushing the generated policy rules to a communication system gateway configured to control the flow of packet-based data between a wireless device and an external network. In various embodiments, generating restrictive policy rules in response to determining the data flow carries malicious content includes generating restrictive policy rules that instruct the gateway to block the data flow and/or to block future data flows originating from the wireless device. The data flow may originate from a wireless device or an Internet server.

[0006] Various embodiment methods, systems and devices may perform charging operations that charge subscriber accounts for data flows allowed by the gateway and not charge subscriber accounts for the data flows blocked by the gateway.

[0007] Pattern matching detection methods may be performed in embodiment methods, systems and devices to determine if the data flow carries content that is a security risk. In an embodiment, pattern matching detection methods may be performed by comparing identified binary-file characteristic information with a binary-file characteristic of a known malware type to determine if the data flow carries content having binary-file characteristics of known malware, which may include comparing the identified binary-file characteristic information with a binary-file characteristic of a spam-sending Internet bot.

[0008] The embodiments include servers configured to implement the various methods through software programming, systems implementing such servers, and non-transitory computer-readable storage media on which are stored server-executable instructions implementing the various methods.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 215

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The accompanying drawings, which are incorporated herein and constitute part of this specification, illustrate exemplary aspects of the invention, and, together with the general description given above and the detailed description given below, serve to explain features of the invention.

[0010] FIG. 1 is a communication system block diagram illustrating network components of a communication system having a third generation partnership project (3GPP), evolved packet core (EPC) architecture suitable for use in the various embodiments.

[0011] FIG. 2 is a communication system block diagram illustrating network components and of an example communication-system having a policy and charging control (PCC) network in accordance with the various embodiments.

[0012] FIG. 3A is a communication system block diagram illustrating example communications between logical components in an example control network and an external network.

[0013] FIG. 3B is an communication system block diagram illustrating example communications between a malware protection system and external networks in accordance with various embodiments.

[0014] FIG. 4A is a process flow diagram illustrating an example method for receiving signaling flows and enforcing policy rules on the signaling flows in accordance with various embodiments.

[0015] FIG. 4B is a process flow diagram illustrating an example policy management system method for receiving parameters from a gateway and generating rules in accordance with the various embodiments.

[0016] FIG. 5 is a process flow diagram illustrating an example pattern matching detection method in accordance with an embodiment.

[0017] FIGs. 6A and 6B are process flow diagrams illustrating example observed behavior detection methods in accordance with various embodiments.

[0018] FIG. 7 is a component block diagram of a server suitable for use with an embodiment.

DETAILED DESCRIPTION

[0019] The various embodiments will be described in detail with reference to the accompanying drawings. Wherever possible, the same reference numbers will be used throughout the drawings to refer to the same or like parts. References made to particular examples and implementations are for illustrative purposes, and are not intended to limit the scope of the invention or the claims.

[0020] The word "exemplary" is used herein to mean "serving as an example, instance, or illustration." Any implementation described herein as "exemplary" is not necessarily to be construed as preferred or advantageous over other implementations.

[0021] As used herein, the terms "mobile device," "wireless device" and "user equipment (UE)" may be used interchangeably and refer to any one of various cellular telephones, smart-phones (e.g., iPhone®), personal data assistants (PDA's), palm-top computers, laptop computers, wireless electronic mail receivers (e.g., the Blackberry® and Treo® devices), VoIP phones, multimedia/Internet enabled cellular telephones, and similar electronic devices capable of sending and receiving wireless communication signals. A wireless device may include a programmable processor and memory.

[0022] As used herein, the term "bot" is used to refer to software applications that emulate human online-activity by automatically performing tasks typically performed by human computer users.

[0023] As used herein, the phrase "malware protection system" is used to refer to any system configured to detect, prevent and/or remove malware such as computer viruses, computer worms, Trojan horses, bots, spyware, adware, etc. Thus, a malware protection system may include a malware prevention system, a malware detection system, a malware removal system, or any other system that assists the detection, prevention and/or removal of malware, as well as combinations of such systems.

[0024] A number of different cellular and mobile communication services and standards are available or contemplated in the future, all of which may implement and benefit from the various embodiments. Such services and standards include, e.g., 3rd generation partnership project (3GPP), long term evolution (LTE) systems, 3rd generation wireless mobile communication technology (3G), 4th generation wireless mobile communication technology (4G), global system for mobile communications (GSM), universal mobile telecommunications system (UMTS), 3GSM, general packet radio service (GPRS), code division multiple access (CDMA) systems (e.g., cdmaone, CDMA2000TM), worldwide interoperability for microwave access (WiMAX), enhanced data rates for GSM evolution (EDGE), advanced mobile phone system (AMPS), digital AMPS (IS-136/TDMA), evolution-data optimized (EV-DO), digital enhanced cordless telecommunications (DECT), and integrated digital enhanced network (iden). Each of these technologies involves the transmission and reception of signaling and content messages. It should be understood that any references to terminology and/or technical details related to an individual standard or technology are for illustrative purposes only and are not intended to limit the scope of the claims to a particular broadcast communication system or technology unless specifically recited in the claim language.

[0025] As discussed above, Internet usage habits are changing and many new types of Internet-enabled devices are emerging to meet evolving user demands. The emergence of these devices presents new security challenges that cannot be addressed by existing antivirus/malware detection and removal systems. For example, current malware detection and prevention techniques attempt to prevent malware by protecting user-level devic-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 216

es (e.g., laptops) rather than from within the network. Since many new and emerging Internet-enabled devices do not having traditional operating systems, current malware detection and prevention techniques cannot efficiently detect or prevent malware.

[0026] The various embodiments provide a comprehensive security solution for detecting and preventing malware from within a mobile data network. The various embodiments protect mobile devices from becoming infected with malicious software (malware) and protect Internet servers and network components from attacks (e.g., denial of service attacks) and inconveniences (e.g., spam) caused by malware. The various embodiments detect both known and yet-to-be identified security threats by performing pattern matching and observed behavior detection methods. Various embodiments monitor network traffic for suspicious content or activities to detect malware and block suspicious data flows from within the network.

[0027] Various embodiments provide methods for detecting and preventing malware from within a data network by combining a policy management component, which has a global or near global view of network traffic and transactions, with a malware system detection component configured to recognize malware based upon patterns within the network traffic and transactions. The various embodiments may operate exclusively within the network and support any type of Internet device.

[0028] The various embodiments may detect transmissions of known and unknown malware in real time as the network patterns emerge. Various embodiments may enable defensive and ameliorative actions, such as modifying the transmissions of the malware or transmission generated by malware, which may include throttling malware transmissions. Various embodiments may prevent the transmission of malware once it has been detected. Various embodiments may alert wireless device users of detected malware. Various embodiments may alert telecommunications network operators of detected malware so they can take corrective actions and/or warn targets of an attack.

The various embodiments may include machine learning mechanisms to enable the malware protection system to learn about malware-generated transmissions and recognize malware patterns as they develop. Such machine learning capabilities may include tracking or correlating network activity levels of recognized threats as well as details of suspicious activities not yet recognized as threatening. Various embodiments may store user-specific malware detection (e.g., patterns in network traffic that occurred while a malware event was underway) and prevention information (e.g., countermeasures that proved effective in responding to a malware event), which may be used by the malware protection system for the detection and prevention of future malware events. Various embodiments may map specific malware threats identified to users of the network. Malware threats may be identified using any suitable mechanism, including the use of IP addresses.

[0029] The various embodiments may push client software capable of exchanging malware and security information with a control network to wireless devices. The malware and security information may include report details generated by the client, alerts generated by the system, instructions from the system to the client that relate to a specific threat to the client, and/or instructions that instruct the wireless device to change its behavior. In an embodiment, the malware and security information may include instructions that instruct a first application executing on the wireless device to cause a second application executing on the wireless device to change its behavior.

[0030] The various embodiments may detect spam by recognizing network traffic with certain characteristics, which may include destination information, protocols, and/or payload content. The various embodiments may protect networks by preventing wireless devices from being used as a source of malware attacks. Various embodiments may detect denial of service attacks by recognizing rapid increases in traffic addressed to a single or related destination from multiple sources. Various embodiments may prevent spam and denial of service attacks by throttling or blocking suspected malicious traffic.

[0031] The various embodiments may reduce the need for network operators to augment capacity and provide network congestion relief. The various embodiments may send detection and prevention information to telecommunications network operators, wireless devices and/or malware systems in real time.

[0032] The various embodiments may be implemented within a variety of mobile communication systems, an example of which is illustrated in FIG. 1. Wireless devices 102 may be configured to send and receive voice, data and control signals to and from a service network 36 (and ultimately the Internet) using a variety of communication systems/technologies (e.g., GPRS, UMTS, LTE, cdmaOne, CDMA2000™, etc). For example, general packet radio service (GPRS) data transmitted from a wireless device 102 is received by a base transceiver station (BTS) 12 and sent to a base station controller (BSC) and/or packet control unit (PCU) component (BSC/PCU) 18. Code division multiple access (CDMA) data transmitted from a wireless device 102 is received by a base transceiver station 12 and sent to a base station controller (BSC) and/or point coordination function (PCF) component (BSC/PCF) 22. Universal mobile telecommunications system (UMTS) data transmitted from a wireless device 102 is received by a NodeB 14 and sent to a radio network controller (RNC) 20. Long term evolution (LTE) data transmitted from a wireless device 102 is received by an eNodeB 16 and sent directly to a serving gateway (SGW) 28 located within the EPC 40.

[0033] The BSC/PCU 18, RNC 20 and BSC/PCF 22 components process the GPRS, UMTS and CDMS data, respectively, and send the processed data to a node within the EPC 40. More specifically, the BSC/PCU 18 and

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 217

RNC 20 units send the processed data to a serving GPRS support node (SGSN) 24 and the BSC/PCF 22 sends the processed data to a packet data serving node (PDSN) and/or high rate packet data serving gateway (HSGW) component (PDSN/HSGW) 30. The PDSN/HSGW 30 may act as a connection point between the radio access network and the IP based PCEF/PGW 32. The SGSN 24 is responsible for routing the data within a particular geographical service area and may send signaling (control plane) information (e.g., information pertaining to call set-up, security, authentication, etc.) to a mobility management entity (MME) 26. The MME may request user and subscription information from a home subscriber server (HSS) 42, perform various administrative tasks (e.g., user authentication, enforcement of roaming restrictions, etc), select a SGW 28 and send administrative and/or authorization information to the SGSN 24.

[0034] Upon receiving the authorization information from the MME 26 (e.g., authentication complete, identifier of a selected SGW, etc.), the SGSN 24 sends the GPRS/ UMTS data to a selected SGW 28. The SGW 28 stores information about the data (e.g., parameters of the IP bearer service, network internal routing information, etc.) and forwards user data packets to a policy control enforcement function (PCEF) and/or packet data network gateway (PGW) 32. The PCEF/ PGW 32 sends signaling information (control plane) to a policy control rules function 34. The PCRF 34 accesses subscriber databases, creates a set of intelligent policy rules and performs other specialized functions (e.g., interacts with online/offline charging systems, application functions, etc.) and sends the intelligent policy rules to the PCEF/ PGW 32 for enforcement. The PCEF/ PGW 32 implements the policy rules to control the bandwidth, the quality of service and the characteristics of the data, services and flows being communicated between the service network 36 and the end users.

[0035] FIG. 2 illustrates various logical components in an example communication system 100 suitable for implementing the various embodiments. Wireless devices 102 may be configured to communicate via cellular telephone network, a radio access network (e.g., UTRAN, RAN, etc.), WiFi network, WiMAX network, and/or other well known technologies (e.g., GPRS, UMTS, LTE, cdmaOne, CDMA2000TM). Wireless devices 102 may be configured to transmit and receive voice, data and control signals to and from a base station (e.g., base transceiver station, NodeB, eNodeB, etc.) 104, which may be coupled to a controller (e.g., cellular base station, radio network controller, service gateway, etc.) operable to communicate the voice, data, and control signals between mobile devices and to other network destinations. The base station 104 may communicate with an access gateway 110, which serves as the primary point of entry and exit of wireless device traffic and connects the wireless devices 102 to their immediate service provider and/or packet data networks (PDNs). The access gateway 110 may be a part of the policy control and charging network 112 or external to the policy control and charging network 112, and may include one or more of a controller, a gateway, a serving gateway (SGW), a packet data network gateway (PGW), an evolved packet data gateway (ePDG), a packet data serving node (PDSN), a serving general packet radio service support node (SGSN), a policy and charging enforcement function (PCEF), or any combination of the features/functions provided thereof. The access gateway 110 may be implemented in a single computing device or in many computing devices, either within a single network or across a wide area network, such as the Internet.

[0036] The access gateway 110 may forward the voice, data, and control signals to other network components as user data packets, provide connectivity to external packet data networks, manage and store contexts (e.g. network internal routing information, etc.) and act as an anchor between different technologies (e.g., 3GPP and non-3GPP systems). The access gateway 110 may also coordinate the transmission and reception of data to and from the Internet 107, and the transmission and reception of voice, data and control information to and from an external service network 120 connected to the Internet 107 and other base stations 104 and wireless devices 102. The access gateway 110 may also route control/signaling information to a policy and charging control (PCC) network 112, which may be a part of an implementation of an evolved packet core (EPC)/long term evolution (LTE) architecture.

[0037] The PCC network 112 may include a gateway and/or policy and charging enforcement function (PCEF) 114 component for controlling the flow of Internet protocol (IP) data between mobile devices 102 and external networks (e.g., Internet 107, service network 120, etc.). The gateway/PCEF 114 may contain a deep packet inspection (DPI) component (not illustrated) configured to obtain additional application-level information about a particular IP data flow and provide the information to other components in the PCC network 112. In various embodiments, the gateway/PCEF 114 may be integrated with, or a part of, the access gateway 110, either within the PCC network 112 or outside of the PCC network 112.

[0038] The policy and charging control network 112 may also include a policy management system 116, a charging system 118 and a user repository 122. The policy management system 116 may communicate with a malware protection system 126 and a notification server 128 in a network (External Network) 130 outside the telecommunications provider network. The notification server 128 may be configured to receive requests from the policy management system 116 (e.g., PCRF) and send messages to mobile devices 102 using various technologies/mechanisms, such as email, short message service (SMS) texts, instant messaging (IM) systems, interactive voice response (IVR) systems. The notification server 128 may also send messages in a mobile device-specific format, which may be unique to each mobile device type. The notification server 128 may choose

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 218

a communications technology/mechanism on a case-by-case basis. For example, the notification server 128 may choose a communications technology/mechanism depending on the user's preferences and/or mobile device type. Messages may be sent through the Internet 107 or the access gateway 110. In an embodiment, the notification server 128 may send messages to a client application running on the mobile devices 102 (e.g., via the Internet 107, access gateway 110, etc.), which may send a response to the notification server 128 for forwarding to the policy management system 116.

[0039] In an embodiment, the notification server 128 may be configured to send messages to mobile devices 102 "in-band" or "out-of-band." The in-band notifications may use the same communication mechanism used for receiving message and the out-of-band notifications may use a different communication mechanism than that used for receiving message. For example, if a user is browsing web pages and receives an in-band notification, then the notification may appear as a new web page.

[0040] The malware protection system 126 may be a dedicated malware system or server, which is sometime referred to herein as a "malware expert system." The malware protection system 126 may be implemented within one or a number of dedicated servers located within the communication network 100, or within one or a number of dedicated servers located external from but with communication links to the communication network 100 (as illustrated). In a further embodiment, some or all elements of the malware protection system 126 functionality may be implemented within a server performing another network function, such as the policy management system 116, a charging system 118 or a user repository 122.

[0041] The malware protection system 126 may maintain a repository of known security risks, gather real-time information regarding emerging security/malware threats, identify security/malware threats by monitoring traffic within the communication network 100, and generate security/malware threat counter-measures. Such security/malware threat counter-measures may include malware removal software and/or malware repair software. In an embodiment, the malware protection system 126 may be configured to generate security/malware threat counter-measures, malware removal software and/or malware repair software in a format specific to each mobile device 102 type. The mobile device formats may be defined by configuration or determined in real time by the policy control system 116, or by the malware protection system 126 based on information received in real time from the policy control system 116.

[0042] The malware protection system 126 may also include a local malware prevention memory or database 127 for each policy management system 116. Likewise, the policy management system 116 may include a local policy management memory (not illustrated), which may be populated by the policy management system 116, the malware protection system 126, or both. In an embodiment, the malware protection system 126 may push data to the policy management memory of the policy management system 116. In an embodiment, the policy management system 116 may periodically request information from the malware protection system 126 and populate the policy management memory with the received information.

[0043] FIG. 3 A is a system architecture diagram illustrating communications between logical components in an example control network 150 (e.g., PCC network) and an external network 155. FIG. 3 illustrates that the control network 150 may include one or more components for providing a policy and charging enforcement function (PCEF) 114, a policy management system 116, a charging system 118, and a user repository 122. The gateway/PCEF 114 may be a part of, or perform operations typically associated with, a gateway GPRS support node (GGSN), a packet data network gateway (PGW), an access gateway 110, or other similar components. In an embodiment, the policy management system 116 may include a policy and charging rules function (PCRF). In an embodiment, the charging system 118 may include an off-line charging system (OFCS) and an on-line charging system (OCS). In an embodiment, the user repository 122 may include a subscriber profile repository (SPR).

[0044] The gateway/PCEF 114 may serve as the primary enforcement point, as a gateway and as a router providing a routing mechanism between the Internet and the radio infrastructure/radio access network. The gateway/PCEF 114 may receive signaling flows and use information contained therein to select an optimal route for a particular type of communication, session, or quality of service (QoS). The gateway/PCEF 114 may also enforce various policies on the received flows, which may include querying, coordinating and/or adjusting various network resources based on a set of policy rules provided by the policy management system 116.

[0045] The policy management system 116 may identify appropriate policy rules for a given situation and send the identified policy rules to the gateway/PCEF 114 for enforcement. Each policy rule may govern the services, quality of service (QoS) and/or bandwidth that are made available to a particular subscriber in a particular data call and/or session. The policy rules may also govern when certain services are made accessible to the subscriber (e.g. weekdays from 9AM to 5PM, etc.) and how long the subscriber may access to those services (e.g., 15 minutes at a time, a total of two hours, etc.). The policy rules may further govern whether specific IP data flows should be "blocked," "allowed" or "allowed with constraints," based on current and potential security risks/malware threats, as discussed in further detail below.

[0046] In an embodiment, the policy management system 116 may generate, compile and select a set of business, technology and security rules that define the policies that are to be enforced. In doing so, the policy management system 116 may communicate with the user repository 122 and/or the charging system 118. The user

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 219

repository 112 may store subscriber information (e.g., customer IDs, preferences, subscription levels, etc.) and user-specific information, such as details of a user's subscription to malware prevention services and historical details of a user's exposure to specific malware threats. The charging system 118 may store information pertaining to how much wireless data is being used by a wireless subscriber and information pertaining to whether the subscriber has sufficient funds/credits/access units to receive a requested service/flow. The policy management system 116 may use the information stored by the user repository 122 and charging system 118 to identify potential security risks/malware threats and user preferences for responding to identified security risks/malware threats and make rule decisions based on a combination of the risk/threat and user preference. The policy management system 116 may make the rule decisions on a per-subscriber, per-session and/or per-service basis. For example, the policy management system 116 may use subscriber information (e.g., subscriber's city of residence), the current usage parameters (e.g., day of week, time of day, peak hours, etc.) and the subscriber's service level (e.g., Gold, Bronze, etc.) in conjunction with identified security risks (e.g., known virus, irregular usage patterns, etc.) and user preferences to generate and/or select one or more rules or rule groups, which are pushed to the gateway/PCEF 114 for enforcement.

[0047] The policy management system 116 may also communicate with external systems 155, such as a malware protection system 126 and a notification server 128. As mentioned above, the malware protection system 126 may maintain a repository of known malware threats, gather information regarding emerging malware threats, identify current and potential malware threats and generate malware threat counter-measures, malware removal software and/or malware repair software to neutralize the security threats. The policy management system 116 may send information to the malware protection system 126 to aid in the malware prevention operations. Likewise, the malware protection system 126 may send security related information to the policy management system 116 to aid in the rule making decisions/operations.

[0048] FIG. 3B is a system architecture diagram illustrating communications between policy management systems and malware protection systems. Specifically, FIG. 3B illustrates that the malware protection system 126 may receive information from the policy management system 116, and share information with other policy management systems (e.g., policy management systems 166) or external antivirus/malware protection systems (e.g., malware detection system 162).

[0049] In an embodiment, the policy management system 116 may provide malware detection feedback to the malware protection system 126. In an embodiment, the malware detection feedback may include detection rate information extracted from the IP data flows. The detection rate information may identify the detection rates of known malware threats and include additional parameters (e.g., wireless device type, location of a wireless device, etc.) readily accessible to components within the control network (e.g., policy management system 116, gateway/PCEF 114, etc.). These communication system control network components are generally unavailable to external services and servers (e.g., antivirus servers, malware detection systems, etc.). Such additional parameters may be carried by the IP data flows, extracted by the gateway/PCEF 114 and/or the policy management system 116 and sent to the malware protection system 126. The malware protection system 126 may then use this information to update its malware protection services (e.g., to update lists, generate new prevention software, etc.).

[0050] In an embodiment, the policy management system 116 may send the malware protection system 126 information relating to IP data flows with suspicious characteristics. This information may be used by the malware protection system 126 to identify emerging and unidentified malware threats that are not yet prevalent enough to be known to the malware protection system 126. The malware protection system 126 may also use this information to develop preemptive or preventative countermeasures for the identified malware threats.

[0051] In an embodiment, the policy management system 116 may send malware protection system sample data extracted from IP traffic flows within the detection feedback, which the malware protection system 126 may use for analyzing the malware threats.

[0052] In an embodiment, the malware protection system 126 can subsequently share information received from the policy management system 116 with other policy management systems (e.g., policy management systems 166) in other mobile data networks (e.g., roaming/home networks, external networks, etc.) and/or with independent malware protection systems (e.g., malware detection system 162) operating in private data networks, such as corporate networks, etc.

[0053] FIG. 4A illustrates an example gateway/PCEF method 400 for receiving signaling flows and enforcing policy rules on the signaling flows in accordance with the various embodiments. This embodiment method 400 may be implemented in a server configured with server-executable instructions configured to cause the server to perform the operations of the method steps. In step 402, the gateway/PCEF may receive a data flow originating from, or destined for, a wireless device. In step 406, the gateway/PCEF may examine the data flows and/or the network traffic, identify key characteristics of the data flows and/or network traffic, generate parameters based on the key characteristics and send the generated parameters to a policy managements system for analysis. In step 408, the gateway/PCEF may receive one or more policy rules from the policy management system. In step 410, the gateway/PCEF may use the received policy rules to enforce the various policies discussed above, such as blocking, restricting, or allowing

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 220

the data flow to be sent to, or received from, a wireless device.

[0054] FIG. 4B illustrates an example policy management system method 450 for receiving parameters from a gateway/PCEF and generating/selecting rules in accordance with the various embodiments. In step 452, the policy management system may receive parameters generated by the gateway/PCEF. In step 454, the policy management system may use the received parameters to determine if the data flows carry content (e.g., malware) that poses a security risk to network and/or wireless devices. For example, the policy management system may use the parameters to perform pattern matching detection methods or observed behavior detection methods to determine if the data flow carries malicious content or is otherwise a security risk. The policy management system may retrieve information from a local memory and/or request/receive information from external memories/systems/servers (e.g., user repository, charging system, a malware protection system, etc.) that may be used to determine if the data flows carry content (e.g., malware) associated with a potential or known security risk. In step 456, the policy management system may generate/select policy rules based on determining the data flows carry content that poses a "known risk," a "potential risk," or "no risk." In step 458, the policy management system may send the generated/selected policy rules to the gateway/PCEF for enforcement.

[0055] As mentioned above, the policy management system may retrieve information from a local memory as well as from external memories/systems/servers. In an embodiment, the policy management system may first access the local memory before requesting information from the external memories/systems/servers to reduce latency and improve efficiency. In an embodiment, if the local memory identifies the data flow as carrying a potential security risk, the policy management system may forgo accessing the external memories/systems/servers and generate rules based on the information in the local memory. This reduces latency and improves efficiency because fewer operations and external communications are required to make intelligent policy rule decisions.

[0056] FIG. 5 illustrates an example pattern matching detection method 500, which may be performed by the policy management system to determine if an IP data flow carries content (e.g., malware) that poses a security risk. In step 502, the policy management system may receive parameters identifying certain key characteristics of the IP data flows and/or network traffic. In step 504, the policy management system may examine the parameters to identify/extract information that may be used to determine if the data flow carries content that poses an unacceptable security risk to a wireless device, the network, or both. For example, in step 504, the policy management system may extract binary-file characteristic information, uniform resource locator (URL) information (e.g., from the HTTP-headers of IP packets), IP address information and/or simple mail transfer protocol (SMTP) information (e.g., from SMTP headers/bodies of outgoing emails) from the received parameters.

[0057] In determination step 506, the policy management system may compare the extracted binary-file characteristic information with binary-file characteristics of known malware types (e.g., bots, viruses, worms, Trojan Horses, etc.) to determine if the data flow carries content having binary-file characteristics of known malware (e.g., bots, viruses, worms, Trojan Horses, etc). If the policy management system determines that the data flow carries content that has binary-file characteristics in common with known malware (i.e., determination step 506 = "Yes"), in step 550, the policy management system may generate restrictive policies rules (e.g., "block content," "allow with constraints," etc.) based on the severity of the risk and push the generated rules to the gateway/PCEF. In an embodiment, the policy management system may generate restrictive policy rules that instruct the gateway/PCEF to block the IP data flow upon determining that an IP data flow carries content having binary-file characteristics of known malware threats (e.g., viruses, Trojan Horses, etc.).

[0058] If the policy management system determines that the data flow does not carry content having binary-file characteristics of known malware types (i.e., determination step 506 = "No"), in determination step 510, the policy management system may compare extracted URL information against a list of URLs known to be under attack (e.g., denial-of-service attack, etc.) to determine if the destination of data flow content is a server under attack. If the policy management system determines that the data flow carries or is likely to carry content destined for a server under attack (i.e., determination step 510 = "Yes"), in step 550, the policy management system may generate policy rules (e.g., "block content," "allow with constraints," etc.) based on various factors (e.g., severity of the risk, subscriber's usage history, subscription level, likelihood request is genuine, etc.) and push the generated rules to the gateway/PCEF. In an embodiment, the policy management system may generate an "allow with constraints" rule upon determining that the IP data flow represents a malware threat (e.g., carries content destined for a server under attack), but does not contain a binary that is a malware threat (e.g., carries content having binary-file characteristics of known malware). In an embodiment, the policy management system may generate rules that include constraints that set an upper threshold on the rate of IP data flows and/or introduce an artificial latency to the IP data flow. In this manner, the policy management system may generate intelligent policy rules to prevent attacks on servers, in real time and from within the network, without shutting off all communications to a particular Internet site.

[0059] If the policy management system determines that the data flow does not carry content destined for a server under attack (i.e., determination step 510 = "No"), in determination step 512, the policy management system may compare SMTP information of emails carried

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 221

by a data-flow with known spam heuristics (e.g., incorrect SMTP headers, origin IP address blacklisting, email not relayed, message body contains blacklisted keywords, message attachment scanning using Optical Character Recognition (OCR) techniques, Artificial Intelligence (AI) based methods such as Bayesian filtering, or any combination of these) to determine if the emails originate from a spam-sending bot. If the policy management system determines that the data-flow carries an email that originates from a spam-sending bot (i.e., determination step 512 = "Yes"), in step 550, the policy management system may generate restrictive policy rules that block or restrict the data flow and push the generated rules to the gateway/PCEF for enforcement. In an embodiment, the policy management system may generate an "allow with constraints" rule upon determining that the IP data flow carries content from a known spam bot. In this manner, the policy management system may generate intelligent policy rules to reduce spam, in real time and from within the network, without filtering desired emails/communications.

[0060] If the policy management system determines that the IP data flow does not carry an email that originates from a spam-sending bot (i.e., determination step 512 = "No"), in determination step 514, the policy management system may compare the IP address information with a list of known offender IP addresses to determine if the IP data flow originates from, or is destined for, a mobile device that is compromised (e.g., infected with malware, a virus, etc.) or blacklisted (e.g., known to intentionally engage in malicious activities). If the policy management system determines that the IP address information is associated with an IP address in the list of known offender IP addresses (i.e., determination step 514 = "Yes"), in step 550, the policy management system may generate policy rules that block or restrict the data flow and push the generated rules to the gateway/PCEF for enforcement. In an embodiment, the policy management system may generate an "allow with constraints" rule upon determining that the IP data flow carries content originating from, or destined for, a mobile device that is compromised or blacklisted. In this manner, the policy management system may generate intelligent policy rules to prevent the wireless device, a network server, or both from engaging in malicious activities, without blocking all the services/communications between a wireless device and a network component.

[0061] If the policy management system determines that the IP address information is not associated with an IP address in the list of known offender IP addresses (i.e., determination step 514 = "No"), in determination step 516, the policy management system may compare the URL information to a list of known offender network addresses (e.g., URLs) to determine if the data flow is destined for a website/server known to engage in malicious activities. If the policy management system determines that the URL information is associated with a URL in the list of known offender URLs (i.e., determination step 516 = "Yes"), in step 550, the policy management system may generate restrictive policy rules and push the generated rules to the gateway/PCEF for enforcement. In an embodiment, the policy management system may generate an "allow with constraints" rule upon determining that the IP data flow carries content destined for a website/server known to engage in malicious activities, without blocking all access to the website/server.

[0062] If the policy management system determines that the URL information is not associated with a URL in the list of known offender URLs (i.e., determination step 516 = "No"), in step 560, the policy management system may generate rules to allow the flow, or perform additional operations (e.g., observed behavior detection methods) to determine if the data flow carries content (e.g., malware) that poses a security risk to a wireless device, network server, or both.

[0063] In an embodiment, the policy management system may perform observed behavior detection methods to determine if the data flow carries content (e.g., malware) that may be harmful or a security risk to the network/wireless devices. The policy management system may perform the observed behavior detection methods instead of, or in addition to, the pattern matching detection methods discussed above. Observed behavior detection methods allow the policy management system to detect emerging and yet-unidentified security risks (e.g., malware threats) that are not yet prevalent enough to be known to the malware protection system or third party antivirus/security systems. The observed behavior detection methods may exploit the policy management system's access to the control network information to monitor IP data flows for suspicious characteristics and to identify security risks. For example, the policy management system may access information stored in a user repository and/or charging system in a control network to identify and neutralize security risks that cannot be identified or neutralized by systems that do not have access to user/charging information (e.g., components outside the control network).

[0064] FIGs. 6A-6B illustrate example observed behavior detection methods that may be performed by the policy management system to determine if an IP data flow is a security risk/malware threat. Performing the observed behavior oriented detection methods may include monitoring a number of data flows matching a specific behavior and/or detecting an unexpected rise in a number of requests matching a defined behavior from a specific wireless device. The specific and defined behaviors may include any of the behaviors described below, including unexpected rise in requests, requests to a specific port 80 of a specific IP address, increases in HTTP requests to a specific URL, unexpected rise in requests to a single destination, unexpected rise in outbound emails and/or DNS queries, etc.

[0065] FIG. 6A illustrates an example observed behavior detection method 600 for determining if an IP data flow carries content that may contribute to a denial of

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 222

service attack. In step 602, the policy management system may receive parameters identifying certain key characteristics of the IP data flows and/or network traffic. In step 604, the policy management system may analyze the parameters and update memories, counters and/or threshold values used to monitor sudden and unexpected rises in requests to a single destination. These counters may be used to store running totals, averages, and/or maximum values for one or more different time periods (e.g., the total number of requests to a destination IP address in the last 24 hours, or the average number of requests to a destination IP address per month). Therefore, the counters can be used to store or derive request rates. The counters may be specific to the user, or they may be generalized across groups of users. Therefore, a sudden and unexpected rise in requests to a single destination can be defined in terms of a counter value surpassing a pre-defined rate. This predefined rate can be expressed in terms of the counters (e.g., a sudden and unexpected rise can be defined as the current rate being greater than a tripling of the maximum hourly rate ever experienced). The associated thresholds may be defined as constants, or in terms of the counters (e.g., a threshold can be defined as a quadrupling of the maximum hourly rate ever experienced). Thus, there can be a sudden and unexpected rise in the number of requests to a single destination, but it may not exceed a threshold. In step 606, the policy management system may monitor the number of requests by performing IP level and application level monitoring operations. IP level monitoring operations may include monitoring requests to a specific port 80 of a specific IP address. Application level monitoring operations may include monitoring increases in HTTP requests to a specific URL. In determination step 608, the policy management system may determine if there has been a sudden and unexpected rise in requests to a single destination. This determination may be based upon a counter value surpassing a pre-defined rate, as described previously. If the policy management system determines that there has not been a sudden and unexpected rise in requests to a particular address (i.e., determination step 608 = "No"), in step 612, the policy management system may generate policy rules and push the generated rules to the gateway/PCEF. If the policy management system determines that there has been a sudden and unexpected rise in requests to a single destination (i.e., determination step 608 = "Yes"), in determination step 610, the policy management system may determine if the rise in requests exceeds a predetermined threshold. If the rise in requests does not exceed the predetermined threshold (i.e., determination step 610 = "No"), in step 612, the policy management system may generate policy rules and push the generated rules to the gateway/PCEF. If the rise in requests exceeds the predetermined threshold (i.e., determination step 610 = "Yes"), the policy management system may conclude that a denial of service attack is emerging, and in step 614, register the destination as being under attack, generate restrictive policy rules to block or limit current and future flows from accessing the registered destination, and push the generated rules to the gateway/PCEF for implementation.

[0066] FIG. 6B illustrates an example observed behavior detection method 650 for determining if an IP data flow carries content from spam engines. In step 652, the policy management system may receive parameters identifying certain key characteristics of the IP data flows and/or network traffic from a gateway/PCEF. In step 654, the policy management system may analyze the parameters and update local databases and/or counters with information that may be used to monitor sudden rises in outbound emails and/or sudden rises in domain name system (DNS) queries for mail exchanger (MX) records by a single mobile device. In step 656, the policy management system may monitor the number of outbound emails and/or DNS queries by performing IP-level and application-level monitoring operations. IP-level monitoring operations performed by the policy management system may include monitoring/counting the number of requests by a specific mobile device to port 25 of different IP addresses. Application-level monitoring operations may include monitoring the number of SMTP requests being sent directly to an email recipient's SMTP server from a single mobile device. The IP-level and application-level monitoring operations may be used to identify data flows carrying content originating from a spam engine because "well-behaved" email clients generally send all their outbound email via a single SMTP server, whereas spam engines typically send outbound emails directly to the email recipient's SMTP server.

[0067] In determination step 658, the policy management system may determine if there has been a sudden and unexpected rise in outbound emails and/or in domain name system (DNS) queries for mail exchanger (MX) records by a single mobile device. This determination may be based upon a counter value surpassing a pre-defined rate, as described previously. If the policy management system determines that there has not been a sudden and unexpected rise in outbound emails/queries (i.e., determination step 658 = "No"), in step 662, the policy management system may generate policy rules and push the rules to the gateway/PCEF. If the policy management system determines that there has been a sudden and unexpected rise in outbound emails and/or in domain name system (DNS) queries (i.e., determination step 658 = "Yes"), in determination step 660, the policy management system may determine if the rise in requests exceeds a predetermined threshold. If the rise exceeds the predetermined threshold (i.e., determination step 660 = "Yes"), in step 664, the policy management system may identify the data flow as carrying content from a spam engine, register the spam engine (e.g., update a spam engine list in memory, etc.), generate restrictive policy rules to block or limit data flows originating from the spam engine, and push the rules to the gateway/PCEF for enforcement. If the rise does not exceed the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 223

predetermined threshold (i.e., determination step 660 = "No"), in step 664, the policy management system may generate policy rules (e.g., allow, allow with constraints, etc.) and push the rules to the gateway/PCEF.

[0068] As mentioned above, the policy management system may perform detection methods (e.g., pattern matching, observed behavior, etc.) and generate restrictive policy rules based on the type of malware detected. For example, the policy management system may generate a "block" rule if an IP data flow carries content having binary-file characteristics of known malware threats (e.g., viruses, Trojan Horses, etc.) and an "allow with constraints" rule if the IP data flow represents a malware threat, but does not contain a binary that is a malware threat (e.g., the IP data flow is part of a denial of service attack, the IP data flow originates from a spam bot, etc.). The "allow with constraints" rule may define requirements (i.e., constraints) for allowing certain IP data flows to pass through the gateway/PCEF. For example, the constraints may require the gateway/PCEF to introduce an artificial latency to certain IP data flows or to set an upper threshold on certain IP data flow rates.

[0069] The "allow with constraints" type rules, and the constraints they impose, allow the policy management system to customize each rule to the preferences of a particular subscriber, wireless device type or type value, network component or network client. For example, the policy management system may access user preference parameters that define which constraints are to be used in various situations, as well as the actions, conditions and variables (e.g., risk levels, thresholds, data usage levels, costs, etc.) associated with each constraint. In an embodiment, a unique/separate set of user preference parameters may be maintained for each wireless subscriber, device, device type, network component, or client. In this manner telecommunications network operators and their customers may customize responses to malware events, which may be implemented automatically and in real time, consistent with the type of threat and how the network operator and/or customer prefers to respond. For example, some organizations cannot to afford to shut off all requests for service during a denial of service attack, and thus would prefer constraints on the volume or data rates imposed on incoming service requests.

[0070] In an embodiment, the constraints imposed by the policy management system may be proportional to a malware threat or threat-level as determined by the malware protection system or specified by an operator. Methods for determining a malware threat-level within malware protection system are disclosed herein. In an embodiment, the constraints imposed by the policy management system may become progressively more restrictive or constrained, in real-time, as a malware threat or threat-level rises as determined by malware protection system, an operator or other systems configured to identify the threat level.

[0071] User preference parameters may define the actions and conditions associated with each constraint in "allow with constraints" type rules. Thus, the user preference parameters may be used by the policy management system to generate policy rules that instruct the gateway/PCEF to perform several different types of operations (i.e., actions) as well as set conditions for performing those actions. For example, the policy management system may access the user preference parameters to determine that the correct course of action (i.e., the correct rule) for the gateway/PCEF upon the detection of low-risk malware is to send an alert message to the wireless device without blocking the IP data flow. Likewise, the user preference parameters may be used by the policy management system to generate rules that instruct the gateway/PCEF to block the current IP data flow and all future IP data flows until the malware has been manually removed from the mobile device.

[0072] In an embodiment, the policy management system may perform quarantining operations by storing details of IP data flows, aggregated details of the IP data flows, or sequestering the actual content of the IP data flows, that were allowed, constrained, or blocked. The policy management system may store source information, destination information and time information for each IP data flow that is constrained. For example, the policy management system may track the number of IP data flows containing infected binaries within a predetermined time period (e.g., within the last month) and store the number value in a memory. The policy management system may also store emails deemed to be spam. This information may be stored such that a wireless device user whose mobile device is associated with these IP data flows may subsequently view the quarantining details. In an embodiment, the policy management system may be configured to receive messages from the wireless device user that trigger further quarantining actions (e.g., removing a stored IP data flow, un-quarantining stored IP data flow, etc.). For example, the user may send the policy management system a message to reclassify an email as not spam, in which case the email may be pushed to the mobile device by the policy management system. In an embodiment, the user may classify the IP data flow in a manner that causes the policy management system to allow delivery of future similar IP data flows.

[0073] In the various embodiments, user preference parameters may be generated from information from a subscriber account in the user repository, by interacting with the user (e.g., via SMS messages sent through the notification server 126, etc.) to determine his/her preferences in real-time, by using contextual information (e.g., roaming in third party network, home network, etc.), or any combination thereof. For example, a wireless device user may upload user preference parameters to the user repository that instructs the policy management system to send an alert to the user's wireless device warning of potential risks if the wireless device is not roaming in a third-party mobile data network and to block the current IP data flow and all future IP data flows if the wireless

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 224

device is roaming in a third-party mobile data network. In this manner, cost conscientious users may prevent a potential malware program from accessing online services while roaming in an expensive third-party mobile data network.

[0074] In an embodiment, the user preference parameters may be based on a malware threat model. The malware threat model may be defined in terms of multiple independent factors, such as the probability of infection, the type of infection, the severity of the threat posed by the infection, and the difficulty associated with removing the infection. These factors may be weighted and combined to produce an overall threat level. The user may then specify his/her preference in terms of this overall threat level. Alternatively, the user may specify his/her preference in terms of one of the factors within the threat model (e.g., a user may be willing to be exposed to a type of infection that installs Adware, whereas he/she may not be willing to be exposed to a type of infection that corrupts or destroys data stored on the wireless device).

[0075] In an embodiment, the policy management system and/or gateway/PCEF may be configured to charge users/subscribers differently for the IP data flows depending on the actions. For example, a user/subscriber may not be charged for any IP data flows that were blocked by the gateway.

[0076] In an embodiment, a malware prevention and removal client (MPRC) software application may be pushed to the wireless devices. The policy management system may, upon determining that an IP data flow is associated with a malware threat, notify the malware protection system of the wireless device's identity and the detected malware. The malware protection system may then update, or instruct, the MPRC to take preventative, recovery, or removal actions specific to the wireless device type or type value and the detected malware threat. For example, the malware protection system may instruct the MPRC to disable the wireless device's data capabilities and/or isolate the mobile device from the mobile data network until the malware is removed or is otherwise no longer a threat. In an embodiment, the policy management system may send a notification based on the status of the user's subscription to the malware prevention service.

[0077] As discussed above with reference to FIG. 3, the policy management system may exchange information with both a user repository and a charging system. In an embodiment, the policy management system may directly charge the subscription accounts of beneficiaries (e.g., wireless device users, web servers, telecommunication clients, etc.) of the above mentioned malware detection/prevention services by sending a charging message to deduct funds from each subscriber's service account. For example, the policy management system may send a charge message to charge a wireless device user's wireless subscription account each time malware is detected and/or removed from the user's wireless device.

[0078] The policy management system may support multiple different charging models to support any required business models. The charging models may include providing the above mentioned services free-of-charge, charging on an annual basis, on a pay-as-you-consume basis, or any combination thereof. If the services are charged on an annual basis, the policy management system may determine the status of the users' annual subscriptions (e.g., last time payment was made, account status, subscription level, etc.) by accessing the user repository and sending charging messages to the charging system to charge the customers as required. If the services are charged on a pay-as-you-consume basis, the policy management system may send a charging request to the charging system each time the service is provided (e.g., each time a message is send waning of potential malware). In an embodiment, each service may be charged separately. For example, the policy management system may send charging messages to the charging system annually for the detection services and send charging messages to the charging system each time malware is removed.

[0079] In an embodiment, the policy management system may be configured to automatically adapt to a network wide stance if the detection of malware indicates that the network is under attack. The adaption of the network wide stance may be in stages, such that the network-wide protections are progressive. In an embodiment, the policy management system may change and/or bulk push policies/rules into the network if the network the detection of malware indicates that the network is under attack. Policies/rules may be pushed to network elements upon detection of a new malware threat or upon a malware threat level achieving a predefined threshold such that the network adopts an "under attack" pre-defined configuration. The initiation of such "under attack" responses or configuration may be based on the detection of malware and/or policy counters exceeding a predefined threshold number or instances of attack.

[0080] The various embodiments and logical components may be implemented on any of a variety of commercially available server devices, such as the server 700 illustrated in FIG. 7. Such a server 700 typically includes a processor 701 coupled to volatile memory 702 and a large capacity nonvolatile memory, such as a disk drive 703. The server 700 may also include a floppy disc drive, compact disc (CD) or DVD disc drive 706 coupled to the processor 701. The server 700 may also include network access ports 704 coupled to the processor 701 for establishing data connections with a network 705, such as a local area network coupled to other broadcast system computers and servers.

[0081] The processor 701 may be any programmable microprocessor, microcomputer or multiple processor chip or chips that can be configured by software instructions (applications) to perform a variety of functions, including the functions of the various embodiments described below. Multiple processors 701 may be provided,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 225

such as one processor dedicated to wireless communication functions and one processor dedicated to running other applications. Typically, software applications may be stored in the internal memory 702, 703 before they are accessed and loaded into the processor 701. The processor 701 may include internal memory sufficient to store the application software instructions.

[0082] The foregoing method descriptions and the process flow diagrams are provided merely as illustrative examples and are not intended to require or imply that the steps of the various embodiments must be performed in the order presented. As will be appreciated by one of skill in the art the order of steps in the foregoing embodiments may be performed in any order. Words such as "thereafter," "then," "next," etc. are not intended to limit the order of the steps; these words are simply used to guide the reader through the description of the methods. Further, any reference to claim elements in the singular, for example, using the articles "a," "an" or "the" is not to be construed as limiting the element to the singular.

[0083] The various illustrative logical blocks, modules, circuits, and algorithm steps described in connection with the embodiments disclosed herein may be implemented as electronic hardware, computer software, or combinations of both. To clearly illustrate this interchangeability of hardware and software, various illustrative components, blocks, modules, circuits, and steps have been described above generally in terms of their functionality. Whether such functionality is implemented as hardware or software depends upon the particular application and design constraints imposed on the overall system. Skilled artisans may implement the described functionality in varying ways for each particular application, but such implementation decisions should not be interpreted as causing a departure from the scope of the present invention.

[0084] The hardware used to implement the various illustrative logics, logical blocks, modules, and circuits described in connection with the aspects disclosed herein may be implemented or performed with a general purpose processor, a digital signal processor (DPC), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, discrete hardware components, or any combination thereof designed to perform the functions described herein. A general-purpose processor may be a microprocessor, but, in the alternative, the processor may be any conventional processor, controller, microcontroller, or state machine. A processor may also be implemented as a combination of computing devices, e.g., a combination of a DPC and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DPC core, or any other such configuration. Alternatively, some steps or methods may be performed by circuitry that is specific to a given function.

[0085] In one or more exemplary aspects, the functions described may be implemented in hardware, software, firmware, or any combination thereof. If implemented in software, the functions may be stored on or transmitted over as one or more instructions or code on a non-transitory computer-readable medium. The steps of a method or algorithm disclosed herein may be embodied in a processor-executable software module which may reside on a non-transitory computer-readable storage medium. Non-transitory computer-readable storage media may be any available media that may be accessed by a computer (e.g., a server) or processor of a computer or server. By way of example, and not limitation, non-transitory computer-readable media may include RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other storage medium that may be used to store desired program code in the form of processor-executable instructions or data structures and that may be accessed by a computer. Disk and disc, as used herein, includes compact disc (CD), laser disc, optical disc, digital versatile disc (DVD), floppy disk, and blu-ray disc where disks usually reproduce data magnetically, while discs reproduce data optically with lasers. Combinations of the above should also be included within the scope of non-transitory computer-readable media. Additionally, the operations of a method or algorithm may reside as one or any combination or set of codes and/or instructions on a non-transitory machine readable medium and/or computer-readable medium, which may be incorporated into a computer program product.

[0086] The preceding description of the disclosed embodiments is provided to enable any person skilled in the art to make or use the present invention. Various modifications to these embodiments will be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other embodiments without departing from the scope of the invention. Thus, the present invention is not intended to be limited to the embodiments shown herein but is to be accorded the widest scope consistent with the following claims and the principles and novel features disclosed herein.

Claims

1. A communication method, comprising:

receiving in a server from one or more communication system control network components data flow parameters identifying characteristics of a data flow;

performing in the server a detection method to determine if the data flow carries malicious content;

generating restrictive policy rules in response to determining the data flow carries malicious content;

generating non-restrictive policy rules in response to determining the data flow does not

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 226

carry malicious content; and
pushing the generated policy rules to a communication system gateway configured to control a flow of packet-based data between a wireless device and an external network.

2. The method of claim 1, wherein generating restrictive policy rules in response to determining the data flow carries malicious content comprises generating restrictive policy rules based on a subscriber's usage history and a subscription level.

3. The method of claim 1, wherein generating restrictive policy rules in response to determining the data flow carries malicious content comprises generating constraints that are proportional to a threat-level associated with the malicious content and generating restrictive policy rules that impose the generated constraints on the communication system gateway.

4. The method of claim 1, further comprising transmitting malware detection feedback to a network component in a second external network.

5. The method of claim 1, further comprising interacting with a wireless device user to determine user preferences in real-time via alert messages sent through a notification server to receive user preference parameters.

6. The method of claim 1, wherein the data flow originates from the wireless device, and generating restrictive policy rules in response to determining the data flow carries malicious content comprises generating restrictive policy rules that instruct the communication system gateway to block future data flows originating from the wireless device.

7. The method of claim 1, wherein performing a detection method to determine if the data flow carries content that is a security risk comprises performing a pattern matching detection method.

8. The method of claim 7, wherein:

receiving in a server from one or more communication system control network components data flow parameters identifying characteristics of a data flow comprises receiving parameters identifying binary-file characteristic information; and
performing a pattern matching detection method comprises comparing the identified binary-file characteristic information with a binary-file characteristic of a known malware type to determine if the data flow carries content having binary-file characteristics of known malware.

9. The method of claim 7, wherein:

receiving in a server from one or more communication system control network components data flow parameters identifying characteristics of a data flow comprises receiving parameters identifying source and destination information; and
performing in the server a pattern matching detection method comprises comparing the identified source and destination information against a list of known offender sources and destinations to determine if the data flow originates from a wireless device known to engage in malicious activity and if the data flow is destined for a website/server known to engage in malicious activities.

10. The method of claim 7, wherein:

receiving in a server from one or more communication system control network components data flow parameters identifying characteristics of a data flow comprises receiving parameters identifying destination information; and
performing in the server a pattern matching detection method comprises comparing the identified destination information against a list of destinations known to be under attack to determine if a destination of the data flow is a network server under attack.

11. The method of claim 7, wherein:

receiving in a server from one or more communication system control network components data flow parameters identifying characteristics of a data flow comprises receiving parameters identifying messaging information; and
performing in the server a pattern matching detection method comprises comparing the identified messaging information with spam heuristics.

12. The method of claim 1, wherein performing in the server a detection method to determine if the data flow carries content that is a security risk comprises performing a behavior oriented detection method.

13. The method of claim 12, wherein performing a behavior oriented detection method comprises accessing information stored in one of a user repository of a control network and a charging system of a control network.

14. The method of claim 12, wherein performing a behavior oriented detection method comprises updating a counter value to monitor a number of data flows

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 227

matching a specific behavior.

15. The method of claim 12, wherein performing a behavior oriented detection method comprises performing monitoring operations in order to detect an unexpected rise in a number of requests matching a defined behavior from a specific wireless device.

16. A server, comprising:

a memory; and
a processor coupled to the memory, wherein the processor is configured with processor-executable instructions to perform operations comprising:

receiving data flow parameters identifying characteristics of a data flow;
performing a detection method to determine if the data flow carries malicious content;
generating restrictive policy rules in response to determining the data flow carries malicious content;
generating non-restrictive policy rules in response to determining the data flow does not carry malicious content; and
pushing the generated policy rules to a communication system gateway configured to control a flow of packet-based data between a wireless device and an external network.

17. A non-transitory computer readable medium having stored thereon processor-executable software instructions configured to cause a processor to perform operations comprising:

receiving from one or more communication system control network components data flow parameters identifying characteristics of a data flow;
performing a detection method to determine if the data flow carries malicious content;
generating restrictive policy rules in response to determining the data flow carries malicious content;
generating non-restrictive policy rules in response to determining the data flow does not carry malicious content; and
pushing the generated policy rules to a communication system gateway configured to control a flow of packet-based data between a wireless device and an external network.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 228

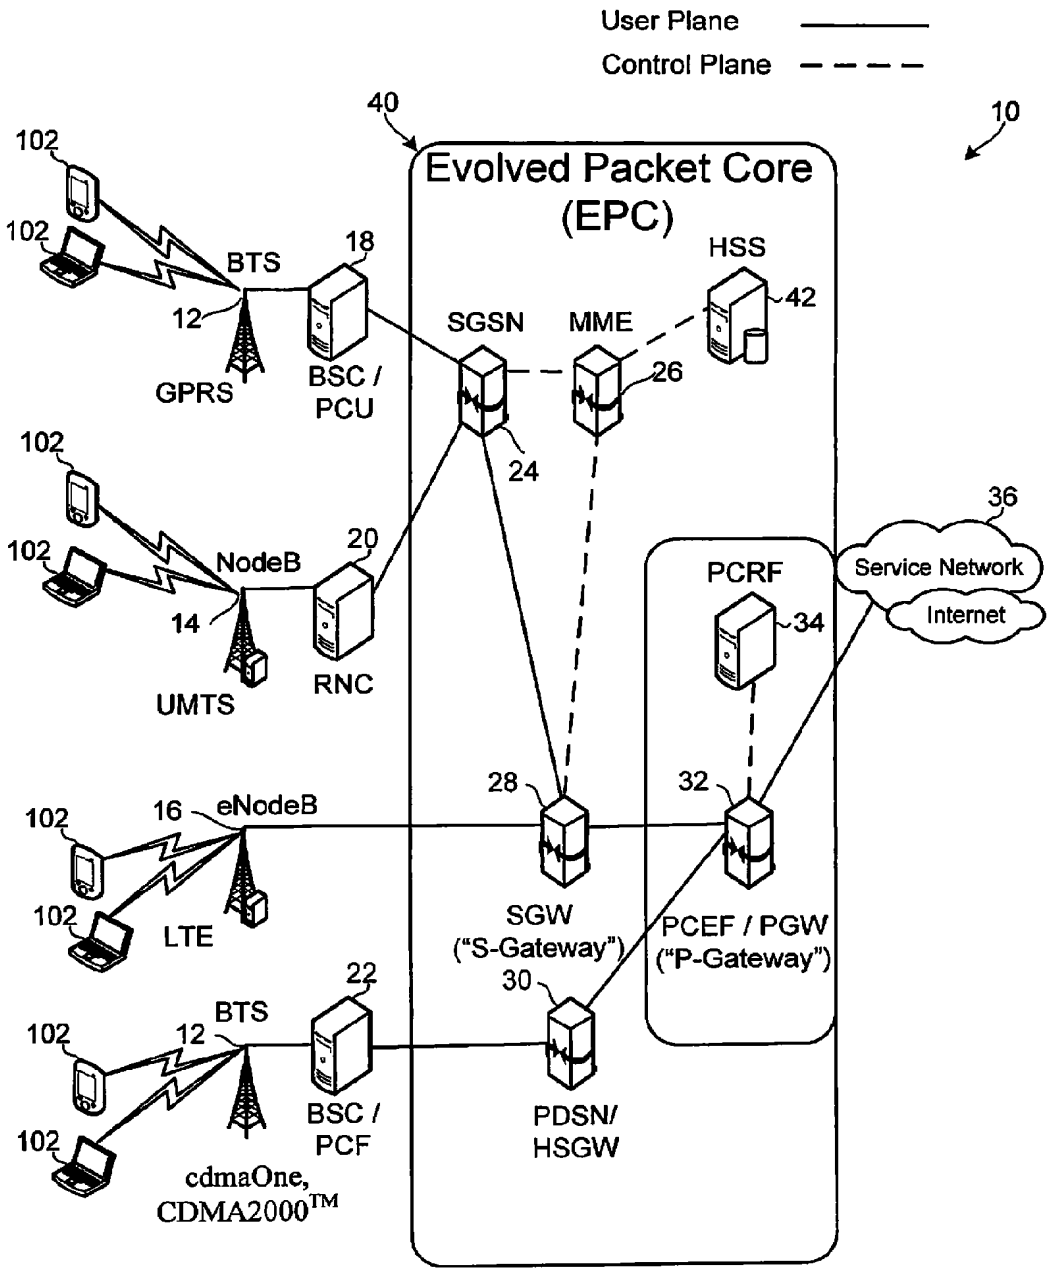

FIG. 1

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 229

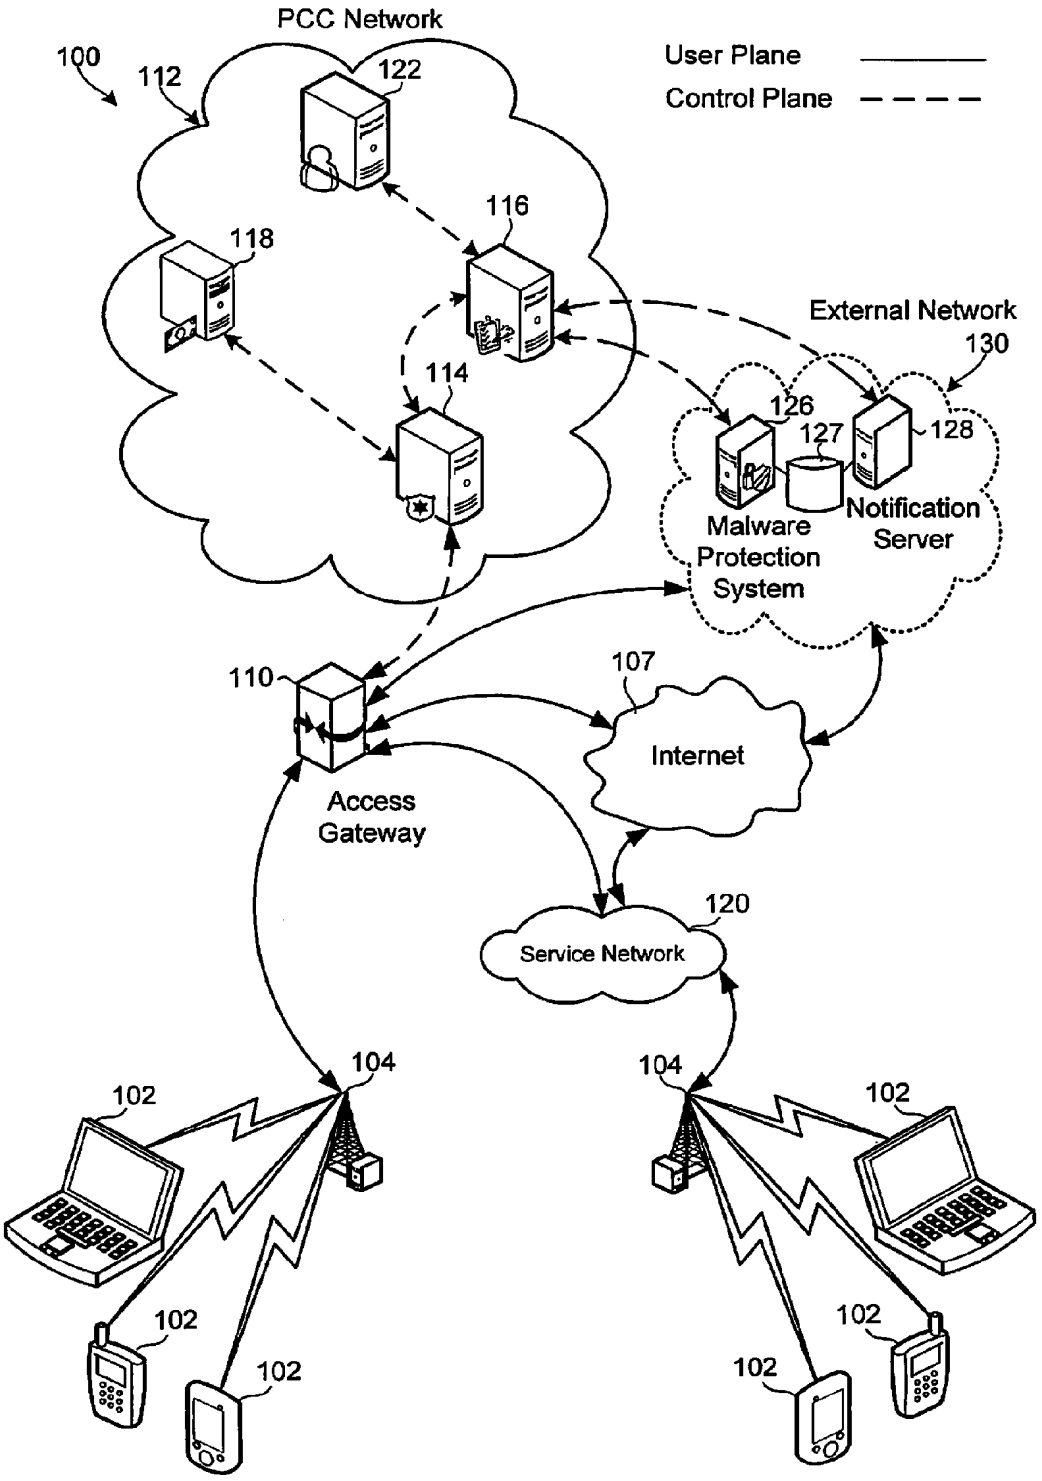

FIG. 2

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 230

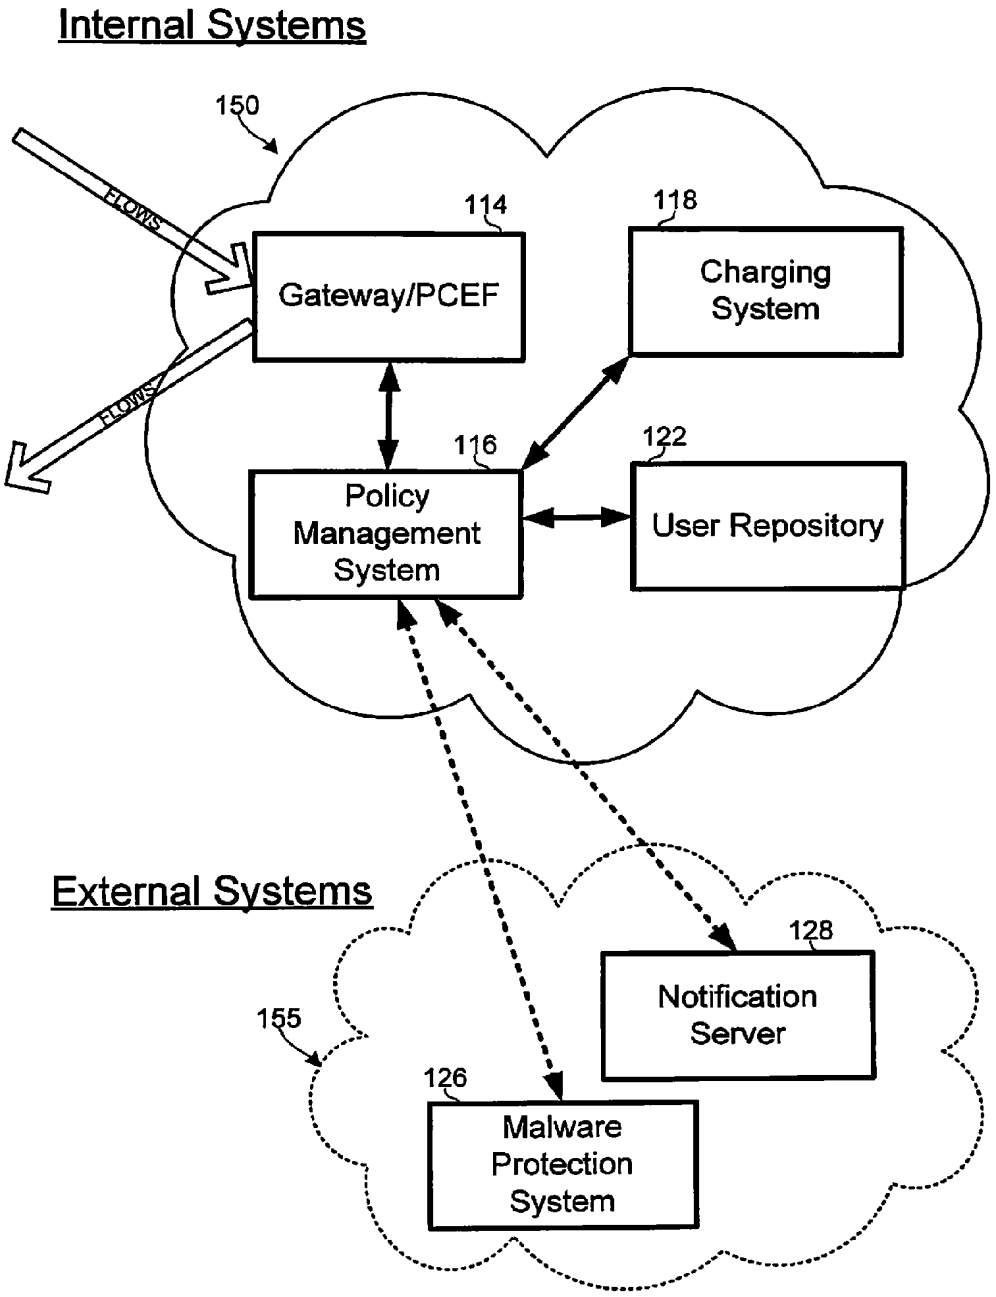

FIG. 3A

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 231

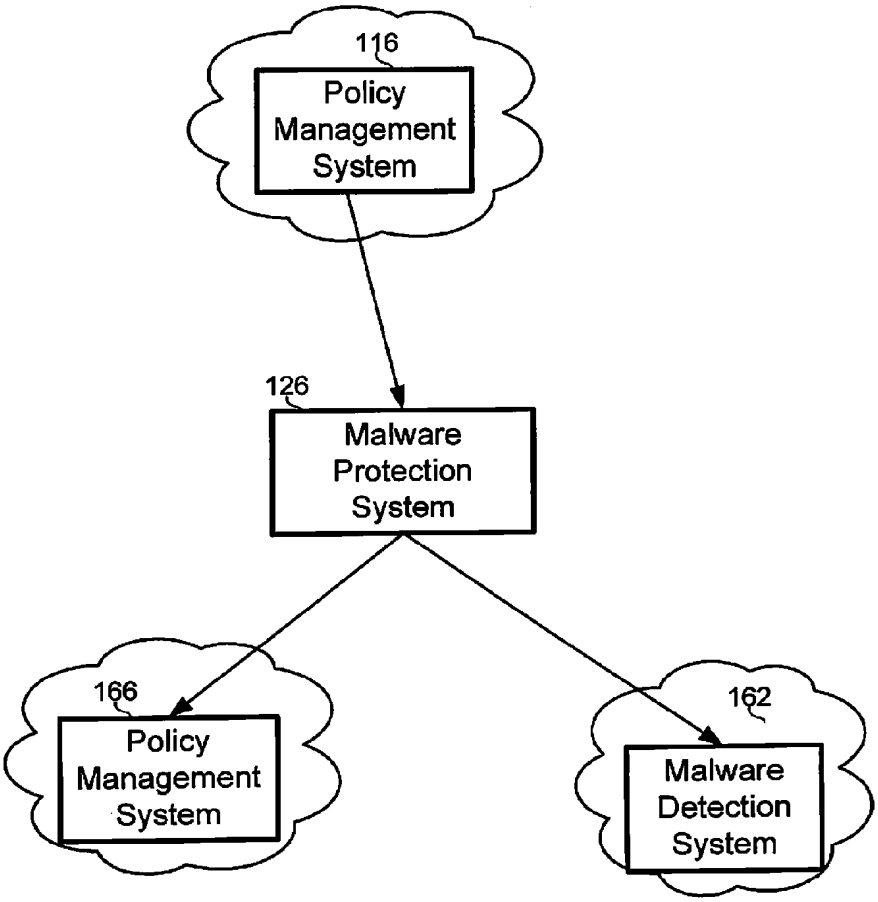

FIG. 3B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 232

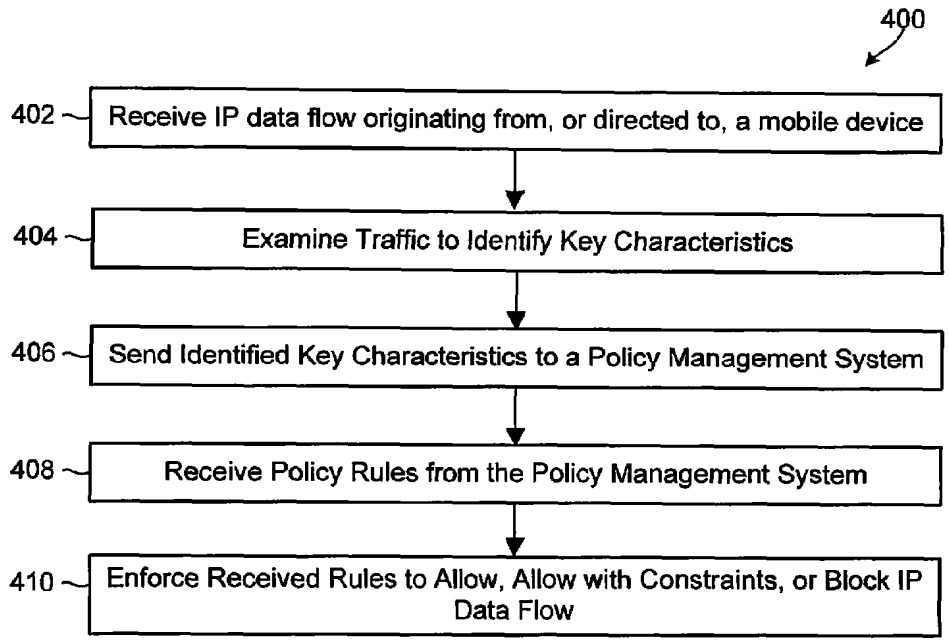

FIG. 4A

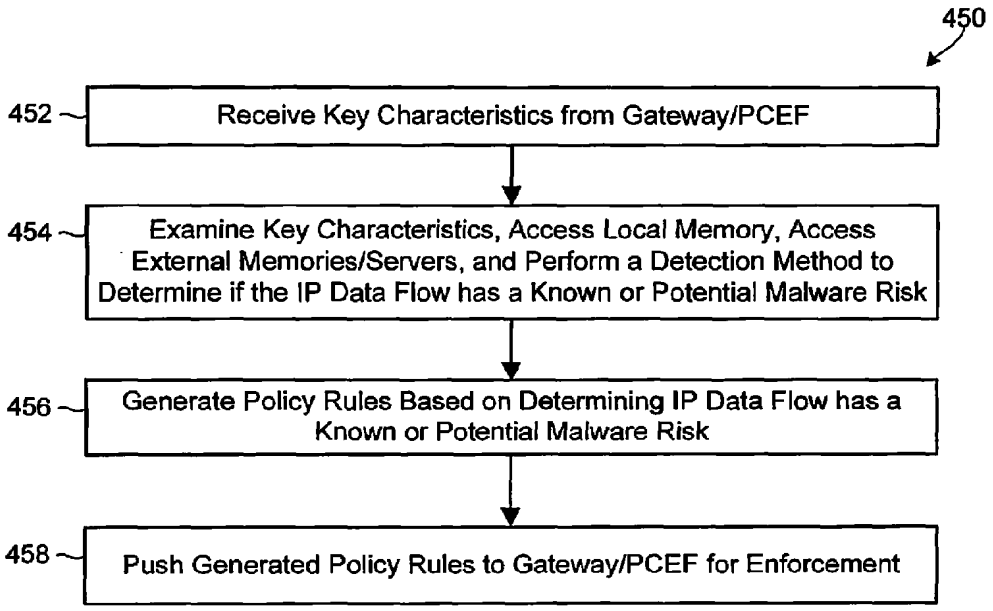

FIG. 4B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 233

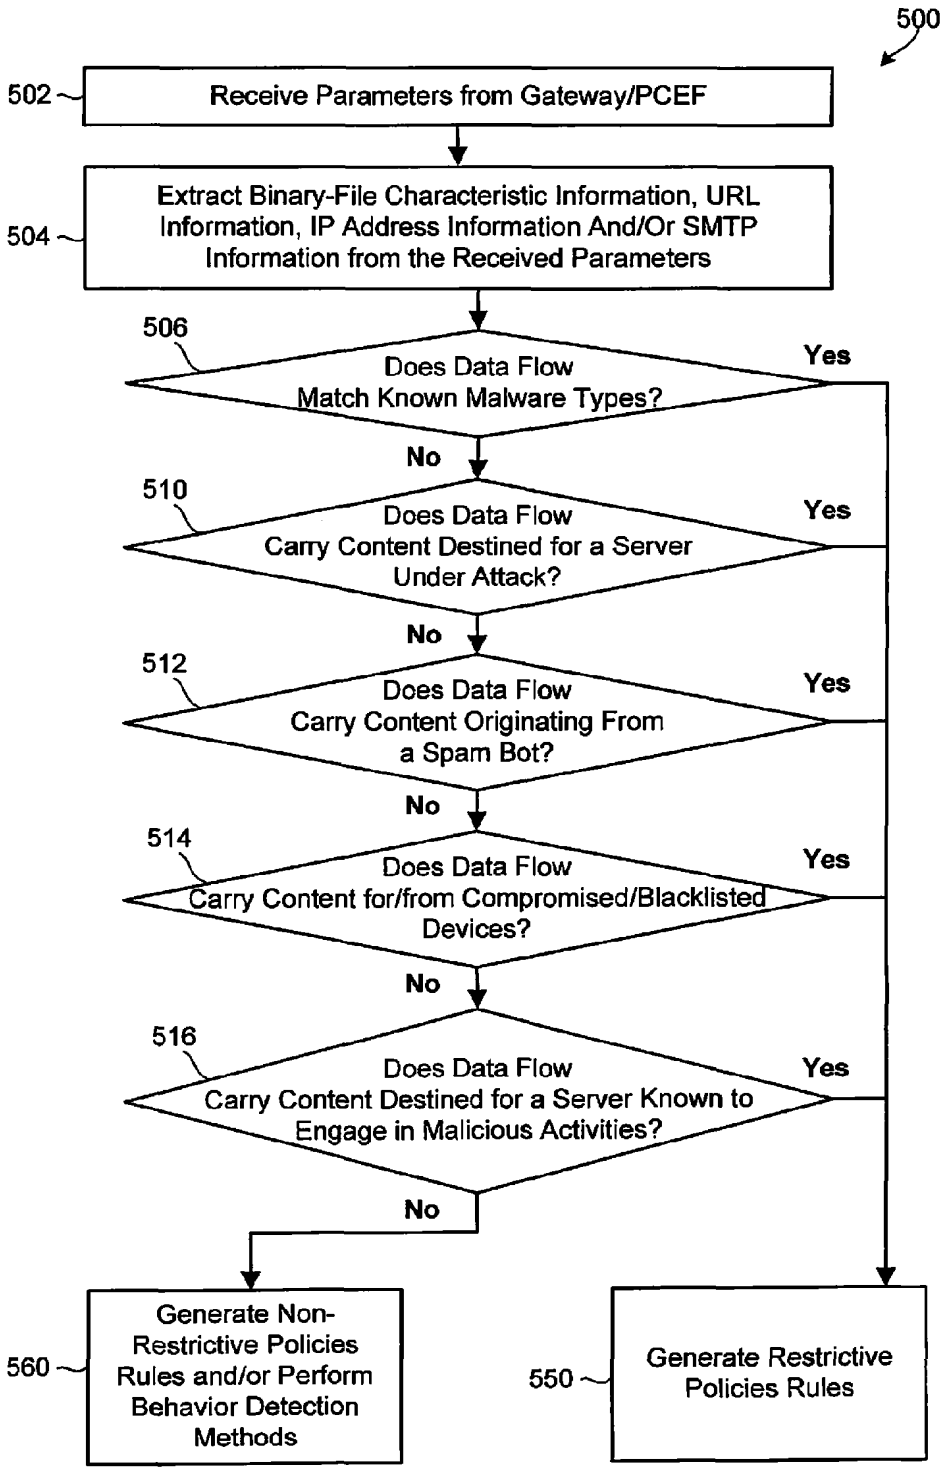

FIG. 5

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 234

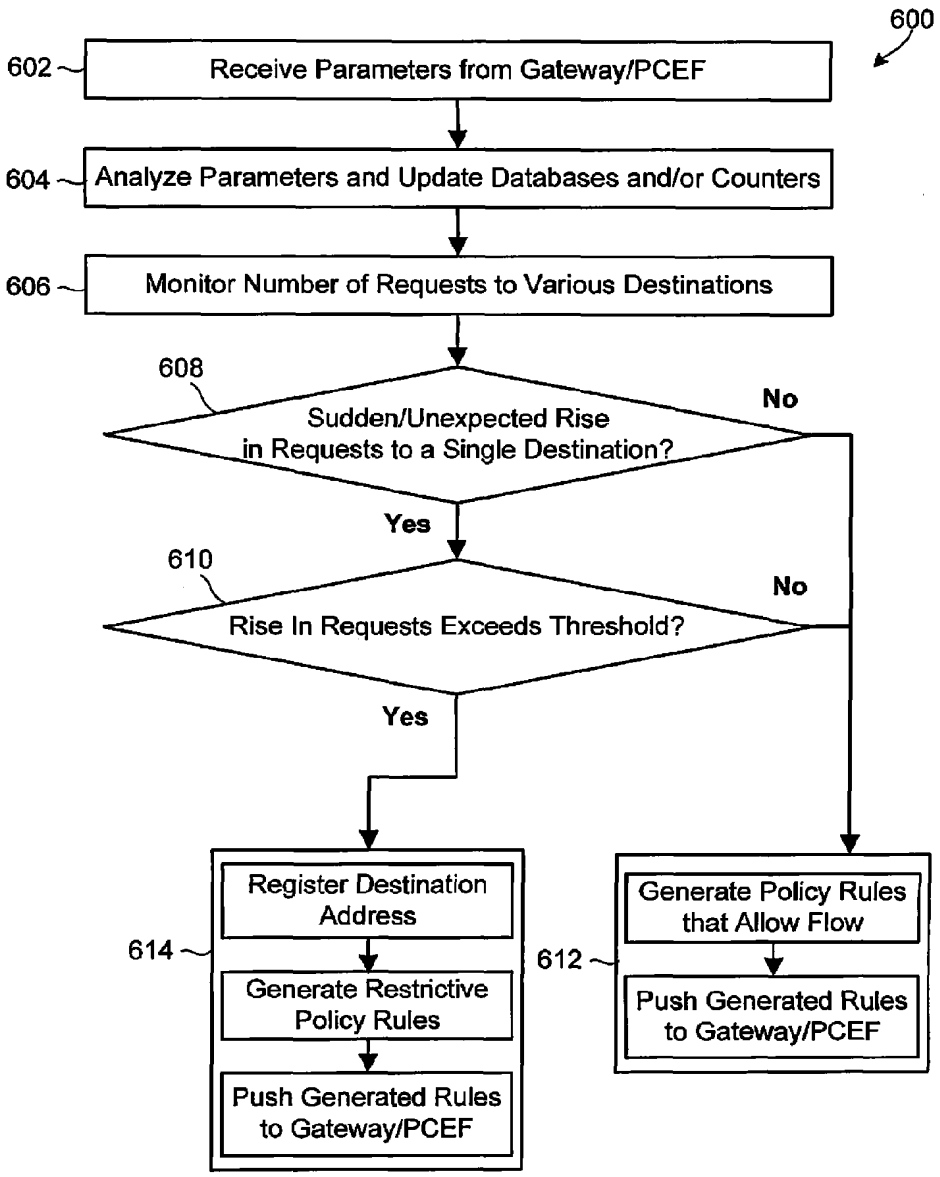

FIG. 6A

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 235

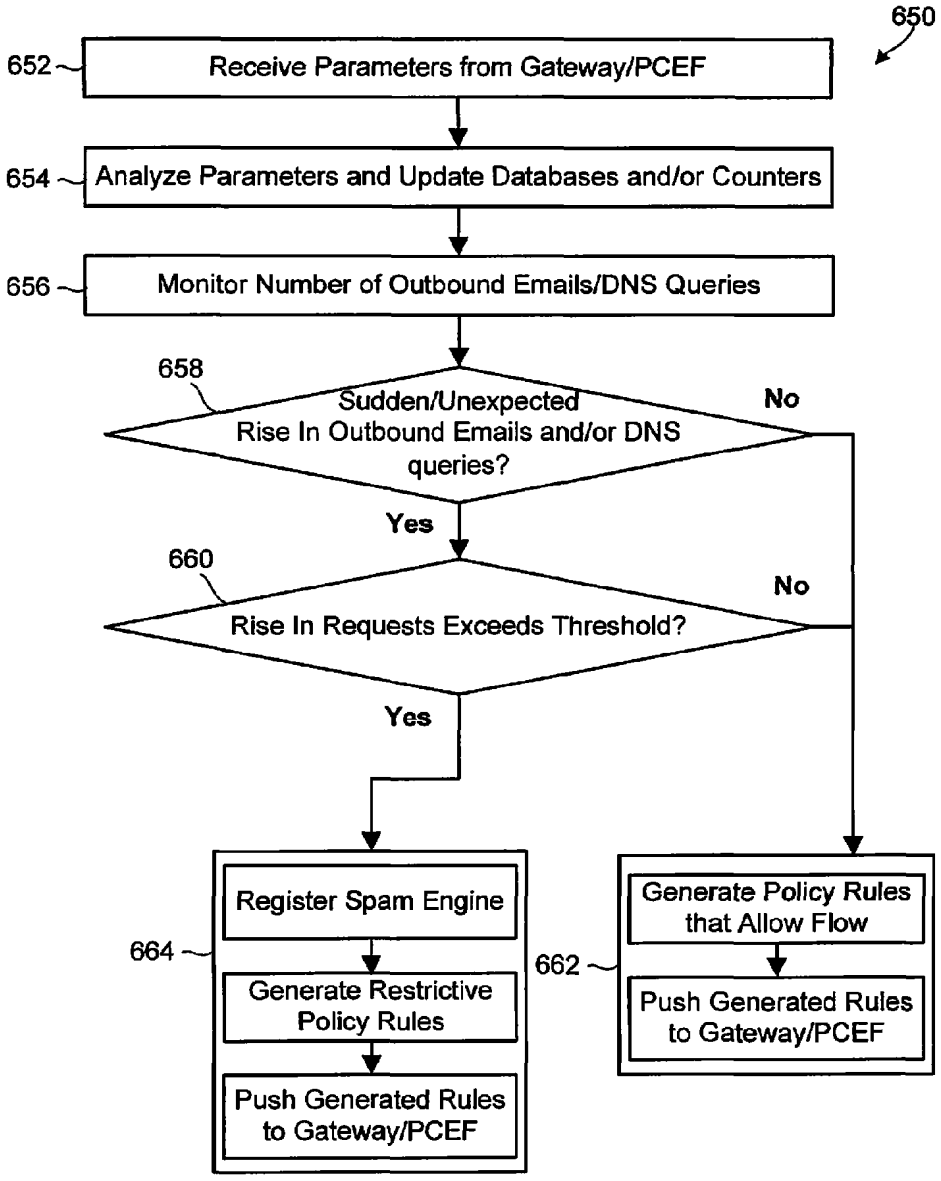

FIG. 6B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 236

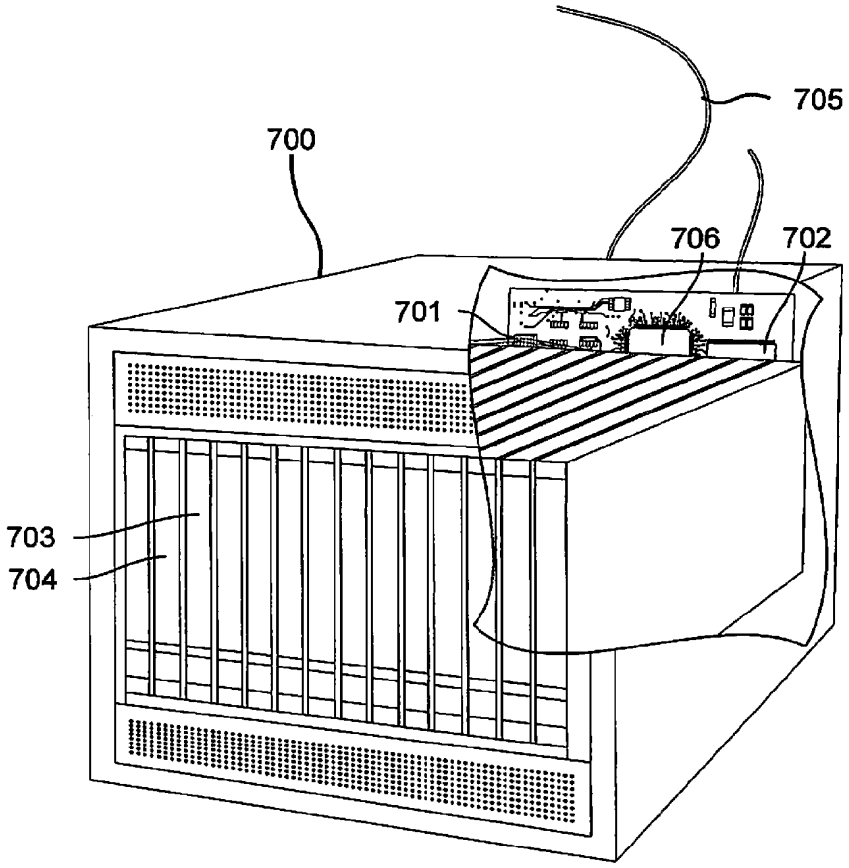

FIG. 7

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 237

Europäisches Patentamt
European Patent Office
Office européen des brevets

EUROPEAN SEARCH REPORT

Application Number
EP 12 15 8921

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	WO 2009/132700 A1 (ERICSSON TELEFON AB L M [SE]; STENFELT JOHN [SE]) 5 November 2009 (2009-11-05) * page 6, lines 4-10 * * page 8, lines 6-10 * * page 15, line 25 - page 17, line 16 * * page 17, lines 29-31 * * claim 12; figures 1,4 *	1-17	INV. H04L12/24 H04L29/06 H04W12/08 H04W12/12 ADD. H04W24/00 H04W88/16
X	US 2004/193912 A1 (LI HONG C [US] ET AL) 30 September 2004 (2004-09-30) * paragraphs [0024], [0039] - [0041] * * figure 3 *	1-17	
A	US 2007/006309 A1 (HERBERT HOWARD C [US] ET AL SMITH NED M [US] ET AL) 4 January 2007 (2007-01-04) * paragraph [0026]; figure 2 *	1-17	

TECHNICAL FIELDS SEARCHED (IPC)

H04L
H04W

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
Munich	20 June 2012	Pajatakis, Emmanouil

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

EPO FORM 1503 03.82 (P04C01)

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 238

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO. EP 12 15 8921

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

20-06-2012

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2009132700 A1	05-11-2009	EP 2304915 A1	06-04-2011
		US 2011041182 A1	17-02-2011
		WO 2009132700 A1	05-11-2009
US 2004193912 A1	30-09-2004	CN 1768518 A	03-05-2006
		DE 112004000428 B4	26-11-2009
		DE 112004000428 T5	09-02-2006
		GB 2412540 A	28-09-2005
		JP 2006521598 A	21-09-2006
		JP 2009187587 A	20-08-2009
		KR 20050118223 A	15-12-2005
		TW I253571 B	21-04-2006
		US 2004193912 A1	30-09-2004
		WO 2004095801 A1	04-11-2004
US 2007006309 A1	04-01-2007	TW I323102 B	01-04-2010
		US 2007006309 A1	04-01-2007
		WO 2007002875 A1	04-01-2007

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 239

Electronic Acknowledgement Receipt

EFS ID:	19487021
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten/Rocko Bund
Filer Authorized By:	William Eugene Wooten
Attorney Docket Number:	007742.00029
Receipt Date:	03-JUL-2014
Filing Date:	16-APR-2014
Time Stamp:	10:16:31
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Non Patent Literature	NPL-ISR_dated_260614_16818779.PDF	744984 e1ca96f7cd5b5a2fa2cccc54367d9a3ea7ad69bc	no	11

Warnings:	
Information:	

2	Non Patent Literature	NPL-Moore_16818826.PDF	1932721 7f7362dd6ca63b62d069e3a66279ca0175c2da98	no	9
Warnings:					
Information:					
3	Non Patent Literature	NPL-Cisco_16818946.PDF	2472590 e14c692da0b6acaa6bb75cbae061bdd60efaea86	no	26
Warnings:					
Information:					
4	Information Disclosure Statement (IDS) Form (SB08)	Information_Disclosure_Statement_Fillable_PDF_16819331.PDF	612619 d7e4477424fe026117858db1871e4d280e9b3e56	no	4
Warnings:					
Information:					
5	Foreign Reference	EP2385676A1_16818890.PDF	780191 b81189390d037dff528b0afcc8f407a761e32de5	no	17
Warnings:					
Information:					
6	Foreign Reference	EP2498442A1_16818909.PDF	1749069 e86e61d1c020da22685c864e51bdf4c636777f30	no	28
Warnings:					
Information:					
		Total Files Size (in bytes):	8292174		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

To:
Wright, Bradley C.
BANNER & WITCOFF, LTD.
1100 13th Street, N.W. Suite 1200
Washington, DC 20005-4051
ETATS-UNIS D'AMERIQUE

PCT

NOTIFICATION OF TRANSMITTAL OF THE INTERNATIONAL SEARCH REPORT AND THE WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

	Date of mailing *(day/month/year)* 26 June 2014 (26-06-2014)
Applicant's or agent's file reference 007742.00031	**FOR FURTHER ACTION** See paragraphs 1 and 4 below
International application No. PCT/US2014/027723	International filing date *(day/month/year)* 14 March 2014 (14-03-2014)
Applicant CENTRIPETAL NETWORKS, INC.	

1. [X] The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

   Filing of amendments and statement under Article 19:
   The applicant is entitled, if he so wishes, to amend the claims of the International Application (see Rule 46):

   When? The time limit for filing such amendments is normally two months from the date of transmittal of the International Search Report.

   Where? Directly to the International Bureau of WIPO, 34 chemin des Colombettes 1211 Geneva 20, Switzerland, Fascimile No.: (41-22) 338.82.70

   For more detailed instructions, see *PCT Applicant's Guide,* International Phase, paragraphs 9.004 - 9.011.

2. [ ] The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.

3. [ ] **With regard to any protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

   [ ] the protest together with the decision thereon has been transmitted to the International Bureau together with any request to forward the texts of both the protest and the decision thereon to the designated Offices.

   [ ] no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. **Reminders**

   The applicant may submit comments on an informal basis on the written opinion of the International Searching Authority to the International Bureau. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established. Following the expiration of 30 months from the priority date, these comments will also be made available to the public.

   Shortly after the expiration of **18 months** from the priority date, the international application will be published by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau before completion of the technical preparations for international publication (Rules 90*bis*.1 and 90*bis*.3).

   Within **19 months** from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase **until 30 months** from the priority date (in some Offices even later); otherwise, the applicant must, **within 20 months** from the priority date, perform the prescribed acts for entry into the national phase before those designated Offices.

   In respect of other designated Offices, the time limit of **30 months** (or later) will apply even if no demand is filed within 19 months.

   For details about the applicable time limits, Office by Office, see www.wipo.int/pct/en/texts/time_limits.html and the *PCT Applicant's Guide*, National Chapters.

Name and mailing address of the International Searching Authority European Patent Office, P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk Tel. (+31-70) 340-2040 Fax: (+31-70) 340-3016	Authorized officer GOTTAR, Chantal Tel: +49 (0)89 2399-3203

RECEIVED
JUL 01 2014
BANNER & WITCOFF LTD

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference	**FOR FURTHER ACTION**	see Form PCT/ISA/220 as well as, where applicable, item 5 below.
007742.00031		
International application No.	International filing date *(day/month/year)*	(Earliest) Priority Date *(day/month/year)*
PCT/US2014/027723	14 March 2014 (14-03-2014)	15 March 2013 (15-03-2013)
Applicant		
CENTRIPETAL NETWORKS, INC.		

This international search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 4 sheets.

[X] It is also accompanied by a copy of each prior art document cited in this report.

1. **Basis of the report**
   a. With regard to the **language,** the international search was carried out on the basis of:
      [X] the international application in the language in which it was filed
      [ ] a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))
   b. [ ] This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6*bis*(a)).
   c. [ ] With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.

2. [ ] **Certain claims were found unsearchable** (See Box No. II)

3. [ ] **Unity of invention is lacking** (see Box No III)

4. With regard to the **title,**
   [X] the text is approved as submitted by the applicant
   [ ] the text has been established by this Authority to read as follows:

5. With regard to the **abstract,**
   [X] the text is approved as submitted by the applicant
   [ ] the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority

6. With regard to the **drawings,**
   a. the figure of the **drawings** to be published with the abstract is Figure No. 1
      [X] as suggested by the applicant
      [ ] as selected by this Authority, because the applicant failed to suggest a figure
      [ ] as selected by this Authority, because this figure better characterizes the invention
   b. [ ] none of the figures is to be published with the abstract

RECEIVED
JUL 01 2014
BANNER & WITCOFF LTD

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/027723

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L29/06
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2006/048142 A1 (ROESE JOHN J [US] ET AL) 2 March 2006 (2006-03-02) figures 1-4 paragraph [0002] - paragraph [0027] paragraph [0037] - paragraph [0042] paragraph [0046] - paragraph [0058] ----- -/--	1-20

[X] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
13 June 2014	26/06/2014

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Olaechea, Javier

1

Form PCT/ISA/210 (second sheet) (April 2005)

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Sean Moore: "SBIR Case Study: Centripetal Networks Subtitle: How CNI Leveraged DHS S&T SBIR Funding to Launch a Successful Cyber Security Company Cyber Security Division 2012 Principal Investigators' Meeting 10-October-2012", , 10 October 2012 (2012-10-10), XP055122227, Retrieved from the Internet: URL:http://www.dhs.gov/sites/default/files/publications/csd-day-2-04-sbir-centripetal-price.pdf [retrieved on 2014-06-06] the whole document	1-20
A	EP 2 385 676 A1 (ALCATEL LUCENT [FR]; INST TELECOM TELECOM BRETAGNE [FR]) 9 November 2011 (2011-11-09) paragraph [0044] - paragraph [0059]	1-20
A	EP 2 498 442 A1 (OPENET TELECOM LTD [IE]) 12 September 2012 (2012-09-12) paragraph [0053] - paragraph [0079]	1-20
A	Cisco: "Control Plane Policing Implementation Best Practices", , 13 March 2013 (2013-03-13), XP002725728, Retrieved from the Internet: URL:https://web.archive.org/web/20130313135143/http://www.cisco.com/web/about/security/intelligence/coppwp_gs.html [retrieved on 2014-06-12] the whole document	1-20
A	US 2010/011434 A1 (KAY RONY [US]) 14 January 2010 (2010-01-14) paragraph [0034] - paragraph [0045]	1-20

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2014/027723

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2006048142 A1	02-03-2006	NONE	
EP 2385676 A1	09-11-2011	CN 102934122 A	13-02-2013
		EP 2385676 A1	09-11-2011
		JP 2013525927 A	20-06-2013
		KR 20130005301 A	15-01-2013
		US 2013111548 A1	02-05-2013
		WO 2011138417 A1	10-11-2011
EP 2498442 A1	12-09-2012	EP 2498442 A1	12-09-2012
		US 2012233656 A1	13-09-2012
US 2010011434 A1	14-01-2010	NONE	

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY
(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2014/027723	14.03.2014	15.03.2013

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☒ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

RECEIVED
JUL 01 2014
BANNER & WITCOFF LTD

Name and mailing address of the ISA:	Date of completion of this opinion	Authorized Officer
European Patent Office D-80298 Munich Tel. +49 89 2399 - 0 Fax: +49 89 2399 - 4465	see form PCT/ISA/210	Olaechea, Javier Telephone No. +49 89 2399-8233

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed

   ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43bis.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing filed or furnished:

   a. (means)

      ☐ on paper

      ☐ in electronic form

   b. (time)

      ☐ in the international application as filed

      ☐ together with the international application in electronic form

      ☐ subsequently to this Authority for the purposes of search

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	
	No:	Claims	1-20
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-20
Industrial applicability (IA)	Yes:	Claims	1-20
	No:	Claims	

2. Citations and explanations

see separate sheet

Box No. VII Certain defects in the international application

The following defects in the form or contents of the international application have been noted:

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1 Reference is made to the following document:

D1 US 2006/048142 A1 (ROESE JOHN J [US] ET AL) 2 March 2006 (2006-03-02)

D2 Sean Moore: "SBIR Case Study: Centripetal Networks Subtitle: How CNI Leveraged DHS S&T SBIR Funding to Launch a Successful Cyber Security Company Cyber Security Division 2012 Principal Investigators' Meeting 10-October-2012", , 10 October 2012 (2012-10-10), XP055122227, Retrieved from the Internet: URL:http://www.dhs.gov/sites/default/files/publications/csd-day-2-04-sbir-centripetal-price.pdf [retrieved on 2014-06-06]

D3 Cisco: "Control Plane Policing Implementation Best Practices", , 13 March 2013 (2013-03-13), XP002725728, Retrieved from the Internet: URL:https://web.archive.org/web/20130313135143/http://www.cisco.com/web/about/security/intelligence/coppwp_gs.html [retrieved on 2014-06-12]

2 The present application does not meet the criteria of Article 33(1) PCT, because the subject-matter of claims 1, 11 and 16 is not new in the sense of Article 33(2) PCT.

2.1 The document D1 discloses (the references in parentheses applying to this document) in terms of claim 1:

A method, comprising:

at a computing platform comprising at least one processor, memory, and communication interface (*network device in Figures 1 and 2*):

receiving, via the communication interface, a plurality of packets (paragraph 55; *traffic is received by the network device*);
responsive to a determination that an overload condition has occurred in one or more networks associated with the computing platform, applying, by the processor and to at least some of the plurality of packets, a first group of packet-filtering rules that comprises at least one five-tuple indicating a first set of packets that should be allowed to continue toward their respective destinations, wherein applying the first group of packet-filtering rules includes allowing at least a first portion of the plurality of packets that fall within the first set of packets to continue toward their respective destinations (paragraphs 24, 56 and table in pages 9-10; *in response to an attack trigger, a rule set is activated by the network device, said rule set including disabling all traffic excluding the management traffic, RR1*); and

responsive to a determination that the overload condition has been mitigated to a first degree, applying, to at least some of the plurality of packets, a second group of packet-filtering rules that comprises at least one five-tuple indicating a second set of packets that should be allowed to continue toward their respective destinations, wherein applying the second group of packet-filtering rules includes allowing at least a second portion of the plurality of packets that fall within the second set of packets to continue toward their respective destinations, and wherein the second group of packet-filtering rules is less restrictive than the first group of packet-filtering rules (paragraph 24; *upon detection of a new triggering condition, a less restrictive set of filtering rules is implemented by the networking device*).

Since all the features of claim 1 are known in combination from document D1, the subject-matter of claim 1 is not new (Article 33(2) PCT). Therefore, claim 1 is not allowable (Article 33(1) PCT).

2.2 The same reasoning applies, mutatis mutandis, to the subject-matter of the corresponding independent claims 11 and 16, which therefore are also considered not new (Article 33(2) PCT).

2.3 Furthermore, it is noted that the disclosure of document D2 (see page 4) would also render the subject-matter of claims 1, 11 and 16 not new or at least not inventive.

3 None of the dependent claims 2-10, 12-15 and 17-20 contain additional features which lead the subject-matter to be both new and inventive (see documents D1-D3 and the corresponding passages cited in the search report).

Re Item VII

Certain defects in the international application

The independent claims are not in the two-part form in accordance with Rule 6.3(b) PCT, which in the present case would be appropriate, with those features known in combination from the prior art being placed in the preamble (Rule 6.3(b)(i) PCT) and with the remaining features being included in the characterising part (Rule 6.3(b)(ii) PCT).

The features of the claims are not provided with reference signs placed in parentheses (Rule 6.2(b) PCT).

The description is not in conformity with the claims as required by Rule 5.1(a)(iii) PCT.

Contrary to the requirements of Rule 5.1(a)(ii) PCT, the relevant background art disclosed in the documents D1-D3 is not mentioned in the description, nor are these documents identified therein.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431
Examiner Name	TBA
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7721084		2010-05-18	Salminen et al.	
	2	7684400		2010-03-23	Govindarajan et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20120113987	A1	2012-05-10	Riddoch et al.	
	2	20110270956	A1	2011-11-03	McDysan et al.	
	3	20100211678	A1	2010-08-19	McDysan et al.	
	4	20050117576		2005-06-02	McDysan et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2431
	Examiner Name	TBA
	Attorney Docket Number	007742.00029

	5	20040151155	A1	2004-08-05	Jouppi	
	6	20030212900	A1	2003-11-13	Liu et al.	
	7	20030188192	A1	2003-10-02	Tang et al.	
	8	20120264443	A1	2012-10-18	Ng et al.	
	9	20130059527	A1	2013-03-07	HASESAKA et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2431
	Examiner Name	TBA
	Attorney Docket Number	007742.00029

EXAMINER SIGNATURE			
Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431
Examiner Name	TBA
Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2014-09-09
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	20086757
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	09-SEP-2014
Filing Date:	16-APR-2014
Time Stamp:	14:38:27
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	612521 d47a0b5a4380b0cfa92082b760cc7dc82acb5c7f	no	5

Warnings:	
Information:	

Total Files Size (in bytes):	612521

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431
Examiner Name	TBA
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20010039624	A1	2001-11-08	Kellum	
	2	20030035370	A1	2003-02-20	Brustoloni	
	3	20030142681	A1	2003-07-31	Chen et al.	
	4	20030145225	A1	2003-07-31	Bruton et al.	
	5	20050138204	A1	2005-06-23	Iyer et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2431
	Examiner Name	TBA
	Attorney Docket Number	007742.00029

	6	20050141537	A1	2005-06-30	Kumar et al.	
	7	20060212572	A1	2006-09-21	Afek et al.	
	8	20080072307	A1	2008-03-20	Maes	
	9	20100082811	A1	2010-04-01	Van Der Merwe et al.	
	10	20100296441	A1	2010-11-25	Barkan	
	11	20110088092	A1	2011-04-14	Nguyen et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2431
	Examiner Name	TBA
	Attorney Docket Number	007742.00029

	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431
Examiner Name	TBA
Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2014-10-27
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	20518285
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	27-OCT-2014
Filing Date:	16-APR-2014
Time Stamp:	09:54:49
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	612611 1ddcf634556c19d9f87fb4ddb49cc30aa21706f3	no	5

Warnings:	
Information:	

Total Files Size (in bytes):	612611

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2431
Examiner Name	TBA
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7539186	B2	2009-05-26	Aerrabotu et al.	
	2	7710885	B2	2010-05-04	Ilnicki et al.	
	3	8856926	B2	2014-10-07	Narayanaswamy et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (**Not for submission under 37 CFR 1.99)**		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven ROGERS
	Art Unit	2431
	Examiner Name	TBA
	Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	REUMANN, John et al., "Adaptive Packet Filters," Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, MI, IEEE, 2001.	☐
	2	GREENWALD, Michael et al., "Designing an Academic Firewall: Policy, Practice, and Experience with SURF," Department of Computer Science, Stanford University, Stanford, CA, IEEE, Proceedings of SNDSS, 1996.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2431
Examiner Name	TBA
Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2014-11-25
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	20802739
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	25-NOV-2014
Filing Date:	16-APR-2014
Time Stamp:	18:27:19
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	612531 c81143b78f05023648e756dd5bfec406ea999853	no	4

Warnings:	
Information:	

2	Non Patent Literature	NPL1.pdf	507671 6740bd120958939e87fc90a8262baad96e26f861	no	5
Warnings:					
Information:					
3	Non Patent Literature	NPL2.pdf	1464667 f99fa81447cc68e17deee719d658cdb0a3edfb87	no	14
Warnings:					
Information:					
			Total Files Size (in bytes):	2584869	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2432
Examiner Name	Barron Jr., Gilberto
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20090172800	A1	2009-07-02	Wool	
	2	20080005795	A1	2008-01-03	Acharya et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven ROGERS
	Art Unit	2432
	Examiner Name	Barron Jr., Gilberto
	Attorney Docket Number	007742.00029

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2432
Examiner Name	Barron Jr., Gilberto
Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2015-02-18
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	21534124
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	18-FEB-2015
Filing Date:	16-APR-2014
Time Stamp:	19:17:01
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	612226 72f3804460847098 4a2233a674910d6df14a8e70	no	4

Warnings:	
Information:	

Total Files Size (in bytes):	612226

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2432
Examiner Name	Barron Jr., Gilberto
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	8806638	B1	2014-08-12	Mani	
	2	6611875	B1	2003-08-26	Chopra et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20090328219	A1	2009-12-31	Narayanaswamy	
	2	20120240135	A1	2012-09-20	Risbood et al.	
	3	20120314617	A1	2012-12-13	Erichsen et al.	
	4	20030097590	A1	2003-05-22	Syvanne	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven ROGERS
	Art Unit	2432
	Examiner Name	Barron Jr., Gilberto
	Attorney Docket Number	007742.00029

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2432
Examiner Name	Barron Jr., Gilberto
Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2015-03-13
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	21758041
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	13-MAR-2015
Filing Date:	16-APR-2014
Time Stamp:	02:47:20
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	612415 0052b9e9ed99092ff87d1874ad579629dfca4e7a	no	4

Warnings:	
Information:	

Total Files Size (in bytes):	612415

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20040088542	A1	2004-05-06	Daude et al.	
	2	20080235755	A1	2008-09-25	Blaisdell et al.	
	3	20100132027	A1	2010-05-27	Ou	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	Mizuno et al., A New Remote Configurable Firewall System for Home-use Gateways, NTT Information Sharing Platform Laboratories, IEEE, 2004.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven ROGERS
	Art Unit	2497
	Examiner Name	Huang, Cheng-Feng
	Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2015-04-20
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	22116782
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	20-APR-2015
Filing Date:	16-APR-2014
Time Stamp:	23:32:42
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	612404 1a7070ab39ed31b499e4d509c793c1f35475bcf4	no	4

Warnings:	
Information:	

2	Non Patent Literature	NPL.pdf	367578 11b72904243b6c17e13b0668dad350cd5503abbd	no	3

Warnings:	
Information:	
Total Files Size (in bytes):	979982

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7478429	B2	2009-01-13	Lyon	
	2	7818794	B2	2010-10-19	Wittman	
	3	8935785	B2	2015-01-13	Pandrangi	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven ROGERS
	Art Unit	2497
	Examiner Name	Huang, Cheng-Feng
	Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button			Add
NON-PATENT LITERATURE DOCUMENTS			Remove
Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐
If you wish to add additional non-patent literature document citation information please click the Add button			Add

EXAMINER SIGNATURE			
Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2015-06-19
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	22691411
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	19-JUN-2015
Filing Date:	16-APR-2014
Time Stamp:	21:24:46
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	612041 75c8fec12b11ba70ddaa3bfae6c837434d9fef2f	no	4

Warnings:	
Information:	

Total Files Size (in bytes):	612041

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	8004994	B1	2011-08-23	Darisi et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20030154297	A1	2003-08-14	Suzuki et al.	
	2	20060114899	A1	2006-06-01	Toumura et al.	
	3	20110185055	A1	2011-07-28	Nappier et al.	
	4	20120331543	A1	2012-12-27	Bostrom et al.	
	5	20130254766	A1	2013-09-26	Zuo et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven ROGERS
	Art Unit	2497
	Examiner Name	Huang, Cheng-Feng
	Attorney Docket Number	007742.00029

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2015-07-21
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	22971206
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	21-JUL-2015
Filing Date:	16-APR-2014
Time Stamp:	05:06:56
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	1035174 f3b92abcb4517fcbda8186c533d820ef7139ec50	no	4

Warnings:	
Information:	

Total Files Size (in bytes):	1035174

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	6317837	B1	2001-11-13	Kenworthy	
	2	8306994	B2	2012-11-06	Kenworthy	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20020049899	A1	2002-04-25	Kenworthy	
	2	20060136987	A1	2006-06-22	Okuda	
	3	20060146879	A1	2006-07-06	Anthias et al.	
	4	20070240208	A1	2007-10-11	Yu et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven ROGERS
	Art Unit	2497
	Examiner Name	Huang, Cheng-Feng
	Attorney Docket Number	007742.00029

	5	20100242098	A1	2010-09-23	Kenworthy	
	6	20130061294	A1	2013-03-07	Kenworthy	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] ¡	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	Communication Relating to the Results of the Partial International Search for International App. No. PCT/ US2015/024691, dated July 10, 2015.	☐
	2	International Preliminary Report on Patentability for International App. No. PCT/US2013/072566, dated July 23, 2015.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven ROGERS
	Art Unit	2497
	Examiner Name	Huang, Cheng-Feng
	Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2015-07-24
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

PCT

To:
Wright, Bradley C.
BANNER & WITCOFF, LTD.
1100 13th Street, N.W. Suite 1200
Washington, DC 20005-4051
ETATS-UNIS D'AMERIQUE

INVITATION TO PAY ADDITIONAL FEES AND, WHERE APPLICABLE, PROTEST FEE

(PCT Article 17(3)(a) and Rule 40.1 and 40.2(e))

Applicant's or agent's file reference 007742.00043	Date of mailing *(day/month/year)* 10 July 2015 (10-07-2015)
International application No. PCT/US2015/024691	**PAYMENT DUE** within **ONE MONTH** from the above date of mailing
Applicant CENTRIPETAL NETWORKS, INC.	International filing date *(day/month/year)* 7 April 2015 (07-04-2015)

1. This International Searching Authority

(i) considers that there are 3 *(number of)* inventions claimed in the international application covered by the claims indicated on an extra sheet:

(ii) therefore considers that **the international application does not comply with the requirements of unity of invention** (Rules 13.1, 13.2 and 13.3) for the reasons indicated on an extra sheet:

(iii) [X] has carried out a partial international search (see Annex) [ ] will establish the international search report on those parts of the international application which relate to the invention first mentioned in claims Nos.:

see extra sheet

(iv) will establish the international search report on the other parts of the international application only if, and to the extent to which, additional fees are paid.

2. Consequently, the applicant is hereby **invited to pay,** within the time limit indicated above, the amount indicated below:

EUR 1.875,00	x	2	=	EUR 3.750,00
Fee per additional invention		number of additional inventions		currency/total amount of additional fees

3. The applicant is informed that, according to Rule 40.2(c), **the payment of any additional fee may be made under protest,** i.e., a reasoned statement to the effect that the international application complies with the requirement of unity of invention or that the amount of the required additional fee is excessive, where applicable, subject to the payment of a protest fee.

Where the applicant pays additional fees under protest, the applicant is hereby invited, within the time limit indicated above, to pay a protest fee (Rule 40.2(e)) in the amount of EUR 865,00 *(currency/amount)*

Where the applicant has not, within the time limit indicated above, paid the required protest fee, the protest will be considered not to have been made and the International Searching Authority will so declare.

4. [ ] Claim(s) Nos. _____ have been found to be unsearchable under Article 17(2)(b) because of defects under Article 17(2)(a) and therefore have not been included with any invention.

Name and mailing address of the International Searching Authority	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk Tel. (+31-70) 340-2040 Fax: (+31-70) 340-3016	BARRIO BARANANO, Ainhoa Tel: +49 (0)89 2399-8621

RECEIVED
JUL 21 2015
BANNER & WITCOFF LTD

Form PCT/ISA/206 (April 2005)
Registered Mail

INVITATION TO PAY ADDITIONAL FEES

International application No.
PCT/US2015/024691

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-9

A method for packet filtering based on application-layer packet header information

---

2. claims: 10-18

A method for selective packet digest logging

---

3. claims: 19, 20

A method for security policy update and management

The lack of unity becomes apparent a posteriori, i.e. for the following reasoning document D1 (US20060136987) is taken into account. From this prior art document what is known is a method comprising a security policy management server ("the controller 13 provides the identifying rule and the processing rule" in paragraph [0093]; "the controller may delete from the identifying rule/processing rule setting table the identifying rule and the processing rule which have not been accessed for a predetermined time." in paragraph [0036]) as claimed in claim 1. These features are the only features common to all the groups of inventions.
With the reference to the prior art document D1:
The third group yields the potential special technical feature of receiving, by a security policy management server and from a computing device, a security update comprising a set of network addresses; updating, one or more rules; receiving, from a different computing device, a security update comprising a different set of network addresses; determining that the different set of network addresses includes at least a portion of network addresses included in the set of network addresses; and responsive to that identifying the at least a portion of network addresses included in the set of network addresses, at least one of the one or more rules stored in the memory that specifies a range of network addresses comprising the at least a portion of network addresses; and updating the at least one of the one or more rules to include one or more other network addresses included in the different set of network addresses, hence solving the objective problem of how to manage security policies based on information received from different sources comprising overlapping network addresses.
Moreover, the first and second invention have further common features, which are also disclosed in document D1 as shown in the following. From document D1 what is further known is:
receiving ("receive the identifying rule and the processing rule from another communication apparatus" in paragraph [0040]), by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server ("the controller 13 provides the identifying rule and the processing rule" in paragraph [0093]; "while the notifying destination of the notifying

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 309

message (identifying rule and the processing rule) 740 is made the router 100C in FIG. 8, the gateway apparatus, a switch, and a management system (not shown) may be made a notifying destination." in paragraph [0095]), a dynamic security policy ("receive the identifying rule and the processing rule from another communication apparatus" in paragraph [0040]; "the controller may delete from the identifying rule/processing rule setting table the identifying rule and the processing rule which have not been accessed for a predetermined time." in paragraph [0036]) that includes at least one rule specifying a packet transformation function to be performed on packets

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway ("received packet" in the abstract);

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets ("identify a protocol type of a communication packet transmitted/received by a device" in paragraph [0028]); and

performing, by the packet security gateway and on a packet-by-packet basis, the packet transformation function on each of the one or more packets ("to automatically set the processing rule of the packet" in paragraph [0028]; "the processing policy may include a policy concerning at least one of a service quality class, filtering, and routing of the packet." in paragraph [0029]) as claimed in claim 1. These features are the only features common to the first two groups of inventions.

With the reference to the prior art document D1:

The first group yields the potential special technical feature of using application-layer packet-header information as the packet identification at the gateway, hence solving the objective problem of how to improve the scalability issues in high resolution filtering.

The second group yields the potential special technical feature of using a packet digest logging function to the identified packets, hence solving the objective problem of how to enable security awareness information in a network.

Consequently, neither the objective problems underlying the subjects of the three claimed inventions, nor the solutions as defined by the special technical features described allow for the link of a common inventive concept to be established between said inventions in the meaning of Rule 13.2 PCT. In conclusion, the three groups of claims are not linked by a single general inventive concept. The application hence does not meet the requirements of unity of invention as defined in Rule 13.1 PCT.

Annex to Form PCT/ISA/206
COMMUNICATION RELATING TO THE RESULTS OF THE PARTIAL INTERNATIONAL SEARCH

International Application No
PCT/US2015/024691

1. The present communication is an Annex to the invitation to pay additional fees (Form PCT/ISA/206). It shows the results of the international search established on the parts of the international application which relate to the invention first mentioned in claims Nos.:
see 'Invitation to pay additional fees'
2. This communication is not the international search report which will be established according to Article 18 and Rule 43.

3. If the applicant does not pay any additional search fees, the information appearing in this communication will be considered as the result of the international search and will be included as such in the international search report.

4. If the applicant pays additional fees, the international search report will contain both the information appearing in this communication and the results of the international search on other parts of the international application for which such fees will have been paid.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2006/136987 A1 (OKUDA MASATO [JP]) 22 June 2006 (2006-06-22) paragraph [0028] - paragraph [0030]; figures 1,21 paragraphs [0033], [0036], [0040], [0093], [0095], [0135]	1-9
X	US 2007/240208 A1 (YU MING-CHE [TW] ET AL) 11 October 2007 (2007-10-11) abstract paragraphs [0021], [0024] - paragraph [0026] paragraph [0033] - paragraph [0036] paragraph [0039] - paragraph [0040]	1-9
A	US 2006/146879 A1 (ANTHIAS TEFCROS [US] ET AL) 6 July 2006 (2006-07-06) paragraph [0074] - paragraph [0076] paragraph [0112] - paragraph [0119] paragraphs [0207], [0208], [0240]	1-9
A	US 8 306 994 B2 (KENWORTHY STACY [US]) 6 November 2012 (2012-11-06) abstract; figures 4,6 column 2 - column 4 column 6 - column 9	1-9

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

1

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 311

Patent Family Annex

Information on patent family members

International Application No
PCT/US2015/024691

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2006136987 A1	22-06-2006	JP 2006174350 A	29-06-2006
		US 2006136987 A1	22-06-2006
US 2007240208 A1	11-10-2007	NONE	
US 2006146879 A1	06-07-2006	CN 101099345 A	02-01-2008
		EP 1834453 A2	19-09-2007
		US 2006146879 A1	06-07-2006
		WO 2006073740 A2	13-07-2006
US 8306994 B2	06-11-2012	US 6317837 B1	13-11-2001
		US 2002049899 A1	25-04-2002
		US 2010242098 A1	23-09-2010
		US 2013061294 A1	07-03-2013

European Patent Organisation

Account details

N° 3 338 800 00 (BLZ 700 800 00)
Commerzbank
Promenadeplatz 7
D-80273 München
SWIFT Code: DRESDEFF700
IBAN: DE20 7008 0000 0333 880000
BIC: DRESDEFF

BNSDOCID: <XS_____20070401OCR_1_>

PATENT COOPERATION TREATY

From the INTERNATIONAL BUREAU

PCT

NOTIFICATION CONCERNING TRANSMITTAL OF COPY OF INTERNATIONAL PRELIMINARY REPORT ON PATENTABILITY (CHAPTER I OF THE PATENT COOPERATION TREATY)
(PCT Rule 44bis.1(c))

To:

WRIGHT, Bradley C.
1100 13th Street, N.W.
Suite 1200
Washington, District of Columbia 20005-4051
ETATS-UNIS D'AMERIQUE

Date of mailing *(day/month/year)*	
23 July 2015 (23.07.2015)	
Applicant's or agent's file reference: 007742.00025	IMPORTANT NOTICE
International application No.: PCT/US2013/072566	International filing date *(day/month/year)*: 02 December 2013 (02.12.2013)
Priority date *(day/month/year)*: 11 January 2013 (11.01.2013)	
Applicant: CENTRIPETAL NETWORKS, INC.	

The International Bureau transmits herewith a copy of the international preliminary report on patentability (Chapter I of the Patent Cooperation Treaty)

The International Bureau of WIPO
34, chemin des Colombettes
1211 Geneva 20, Switzerland

Facsimile No. +41 22 338 82 70

Authorized officer: Simin Baharlou

e-mail: pt09.pct@wipo.int

PATENT COOPERATION TREATY

PCT

INTERNATIONAL PRELIMINARY REPORT ON PATENTABILITY

(Chapter I of the Patent Cooperation Treaty)

(PCT Rule 44*bis*)

Applicant's or agent's file reference 007742.00025	**FOR FURTHER ACTION**	See item 4 below
International application No. PCT/US2013/072566	International filing date *(day/month/year)* 02 December 2013 (02.12.2013)	Priority date *(day/month/year)* 11 January 2013 (11.01.2013)
International Patent Classification (8th edition unless older edition indicated) See relevant information in Form PCT/ISA/237		
Applicant CENTRIPETAL NETWORKS, INC.		

1. This international preliminary report on patentability (Chapter I) is issued by the International Bureau on behalf of the International Searching Authority under Rule 44 *bis*.1(a).

2. This REPORT consists of a total of 5 sheets, including this cover sheet.

   In the attached sheets, any reference to the written opinion of the International Searching Authority should be read as a reference to the international preliminary report on patentability (Chapter I) instead.

3. This report contains indications relating to the following items:

☒	Box No. I	Basis of the report
☐	Box No. II	Priority
☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
☐	Box No. IV	Lack of unity of invention
☒	Box No. V	Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
☐	Box No. VI	Certain documents cited
☐	Box No. VII	Certain defects in the international application
☐	Box No. VIII	Certain observations on the international application

4. The International Bureau will communicate this report to designated Offices in accordance with Rules 44*bis*.3(c) and 93*bis*.1 but not, except where the applicant makes an express request under Article 23(2), before the expiration of 30 months from the priority date (Rule 44*bis* .2).

	Date of issuance of this report 14 July 2015 (14.07.2015)
The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland Facsimile No. +41 22 338 82 70	Authorized officer Simin Baharlou e-mail: pt09.pct@wipo.int

Form PCT/IB/373 (January 2004)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2013/072566	02.12.2013	11.01.2013

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA:	Date of completion of this opinion	Authorized Officer
European Patent Office P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk - Pays Bas Tel. +31 70 340 - 2040 Fax. +31 70 340 - 3016	see form PCT/ISA/210	Oliveira, Joel Telephone No. +31 70 340-3334

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 316

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☑ the international application in the language in which it was filed

   ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43bis.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing filed or furnished:

   a. (means)

      ☐ on paper

      ☐ in electronic form

   b. (time)

      ☐ in the international application as filed

      ☐ together with the international application in electronic form

      ☐ subsequently to this Authority for the purposes of search

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	2, 3, 6-9, 11, 13, 14, 17-22, 24
	No:	Claims	1, 4, 5, 10, 12, 15, 16, 23
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-24
Industrial applicability (IA)	Yes:	Claims	1-24
	No:	Claims	

2. Citations and explanations

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

Reference is made to the following documents:

D1 EP 1 313 290 A1 (STONESOFT CORP [FI]) 21 May 2003 (2003-05-21)

D2 WO 2005/046145 A1 (ERICSSON TELEFON AB L M [SE]) 19 May 2005 (2005-05-19)

D3 US 2011/055916 A1 (AHN DAVID K [US]) 3 March 2011 (2011-03-03)

1 The present application does not meet the criteria of Article 33(2) PCT, because the subject-matter of claims **1, 12 and 23** is not new.

1.1 **D1** discloses a method comprising, at a network protection device **(abstract)**:

receiving a first rule set **(paragraph [0010], column3 lines 4-7)**;

receiving a second rule set **(paragraph [0010], column3 lines 4-7)**;

preprocessing the first rule set and the second rule set **(considered to be implicit from paragraphs [0010], [0013], [0040]: when a new set of rules is made available it must be made certain that they do not contain errors which could stop the filtering of the firewall)**;

configuring the network protection device to process packets in accordance with the first rule set **(paragraph [0010], column3 lines 13-18)**;

receiving packets **(column 6 lines 23-26)**;

processing a first portion of the packets in accordance with the first rule set **(paragraphs [0010], [0020])**;

reconfiguring the network protection device to process packets in accordance with the second rule set **(paragraphs [0010], [0024])**; and

processing a second portion of the packets in accordance with the second rule set **(paragraphs [0010], [0024])**.

1.2 The same reasoning applies, mutatis mutandis, to the subject-matter of the corresponding independent apparatus and computer-readable media claims **12 and 23**, which therefore are also considered not new.

1.3 Dependent claims **2-11, 13-22 and 24** do not contain any features which, in combination with the features of any claim to which they refer, meet the requirements of the PCT in respect of novelty (for claims **4, 5, 10, 15, 16**) and/or inventive step (for claims **2, 3, 6, 7, 8, 9, 11, 13, 14, 17, 18, 19, 20, 21, 22, 24)** see passages cited in the Search Report.

Electronic Acknowledgement Receipt

EFS ID:	23019467
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	24-JUL-2015
Filing Date:	16-APR-2014
Time Stamp:	16:07:33
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	1035471 5d9c1692620023e9635121612d9c27d25dceee15	no	5

Warnings:	
Information:	

2	Non Patent Literature	NPL1.pdf	815786 eb48aa2043fa84c9e8370fe50f66417815b641ca	no	6
Warnings:					
Information:					
3	Non Patent Literature	NPL2.pdf	4664429 f85b52c4383706df599c6d06af9ca94e83f98034	no	6
Warnings:					
Information:					
			Total Files Size (in bytes):	6515686	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	8726379	B1	2014-05-13	Stiansen et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20140366132	A1	2014-12-11	Stiansen et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven ROGERS
	Art Unit	2497
	Examiner Name	Huang, Cheng-Feng
	Attorney Docket Number	007742.00029

	1	International Search Report and Written Opinion for International App. No. PCT/US2015/024691, dated September 16, 2015.	☐
	2	International Preliminary Report on Patentability for International App. No. PCT/US2013/057502, dated May 7, 2015.	☐
	3	International Preliminary Report on Patentability for International App. No. PCT/US2014/023286, dated September 24, 2015.	☐
	4	International Preliminary Report on Patentability for International App. No. PCT/US2014/027723, dated September 24, 2015.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2015-09-30
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

To:
Wright, Bradley C.
BANNER & WITCOFF, LTD.
1100 13th Street, N.W. Suite 1200
Washington, DC 20005-4051
ETATS-UNIS D'AMERIQUE

PCT

NOTIFICATION OF TRANSMITTAL OF THE INTERNATIONAL SEARCH REPORT AND THE WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

Applicant's or agent's file reference 007742.00043	Date of mailing *(day/month/year)* 16 September 2015 (16-09-2015)
International application No. PCT/US2015/024691	**FOR FURTHER ACTION** See paragraphs 1 and 4 below
Applicant CENTRIPETAL NETWORKS, INC.	International filing date *(day/month/year)* 7 April 2015 (07-04-2015)

1. [X] The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

   Filing of amendments and statement under Article 19:
   The applicant is entitled, if he so wishes, to amend the claims of the International Application (see Rule 46):

   When? The time limit for filing such amendments is normally two months from the date of transmittal of the International Search Report.

   How? Directly to the International Bureau of WIPO, 34 chemin des Colombettes 1211 Geneva 20, Switzerland, Fascimile No.: (41-22) 338.82.70

   For more detailed instructions, see *PCT Applicant's Guide*, International Phase, paragraphs 9.004 - 9.011.

2. [ ] The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.

3. [ ] **With regard to any protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

   [ ] the protest together with the decision thereon has been transmitted to the International Bureau together with any applicant's request to forward the texts of both the protest and the decision thereon to the designated Offices.

   [ ] no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. **Reminders**

   The applicant may **submit comments on an informal basis on the written opinion of the International Searching Authority** to the International Bureau. These comments will be made available to the public after international publication. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established.

   Shortly after the expiration of **18 months from the priority date, the international application will be published** by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau before the completion of the technical preparations for international publication (Rules 90*bis*.1 and 90*bis*.3).

   Within **19 months** from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase **until 30 months** from the priority date (in some Offices even later); otherwise, the applicant must, **within 20 months** from the priority date, perform the prescribed acts for **entry into the national phase** before those designated Offices. In respect of other designated Offices, the time limit of **30 months** (or later) will apply even if no demand is filed within 19 months. For details about the applicable time limits, Office by Office, see www.wipo.int/pct/en/texts/time_limits.html and the *PCT Applicant's Guide*, National Chapters.

   Within **19 months from the priority date, the applicant may request that a supplementary international search be carried out** by a different International Searching Authority that offers this service (Rule 45*bis*.1). The procedure for requesting supplementary international search is described in the *PCT Applicant's Guide*, International Phase, paragraphs 8.006-8.032.

Name and mailing address of the International Searching Authority European Patent Office, P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk Tel. (+31-70) 340-2040 Fax: (+31-70) 340-3016	Authorized officer FILIP, Hannelore Tel: +49 (0)89 2399-2861

SEP 24 2015

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 326

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference	FOR FURTHER ACTION	see Form PCT/ISA/220 as well as, where applicable, item 5 below.
007742.00043		
International application No.	International filing date *(day/month/year)*	(Earliest) Priority Date *(day/month/year)*
PCT/US2015/024691	7 April 2015 (07-04-2015)	16 April 2014 (16-04-2014)
Applicant		
CENTRIPETAL NETWORKS, INC.		

This international search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 6 sheets.

☐ It is also accompanied by a copy of each prior art document cited in this report.

1. **Basis of the report**
   a. With regard to the **language,** the international search was carried out on the basis of:
      ☒ the international application in the language in which it was filed
      ☐ a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))
   b. ☐ This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6*bis*(a)).
   c. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.

2. ☐ **Certain claims were found unsearchable** (See Box No. II)

3. ☒ **Unity of invention is lacking** (see Box No III)

4. With regard to the **title,**
   ☒ the text is approved as submitted by the applicant
   ☐ the text has been established by this Authority to read as follows:

5. With regard to the **abstract,**
   ☒ the text is approved as submitted by the applicant
   ☐ the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority

6. With regard to the **drawings,**
   a. the figure of the **drawings** to be published with the abstract is Figure No. 11
      ☐ as suggested by the applicant
      ☐ as selected by this Authority, because the applicant failed to suggest a figure
      ☒ as selected by this Authority, because this figure better characterizes the invention
   b. ☐ none of the figures is to be published with the abstract

SEP 24 2015

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2015/024691

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

see additional sheet

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.

2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.

3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

1-9

Remark on Protest

☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.

☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.

☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/024691

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L29/06
ADD. H04L29/08

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2006/136987 A1 (OKUDA MASATO [JP]) 22 June 2006 (2006-06-22) paragraph [0028] - paragraph [0030]; figures 1,21 paragraphs [0033], [0036], [0040], [0093], [0095], [0135]	1-9
X	US 2007/240208 A1 (YU MING-CHE [TW] ET AL) 11 October 2007 (2007-10-11) abstract paragraphs [0021], [0024] - paragraph [0026] paragraph [0033] - paragraph [0036] paragraph [0039] - paragraph [0040]	1-9

-/--

[X] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
3 July 2015	16/09/2015

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Barla Harter, I

1

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/024691

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2006/146879 A1 (ANTHIAS TEFCROS [US] ET AL) 6 July 2006 (2006-07-06) paragraph [0074] - paragraph [0076] paragraph [0112] - paragraph [0119] paragraphs [0207], [0208], [0240] -----	1-9
A	US 8 306 994 B2 (KENWORTHY STACY [US]) 6 November 2012 (2012-11-06) abstract; figures 4,6 column 2 - column 4 column 6 - column 9 -----	1-9

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2015/024691

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2006136987 A1	22-06-2006	JP 2006174350 A	29-06-2006
		US 2006136987 A1	22-06-2006
US 2007240208 A1	11-10-2007	NONE	
US 2006146879 A1	06-07-2006	CN 101099345 A	02-01-2008
		EP 1834453 A2	19-09-2007
		US 2006146879 A1	06-07-2006
		WO 2006073740 A2	13-07-2006
US 8306994 B2	06-11-2012	US 6317837 B1	13-11-2001
		US 2002049899 A1	25-04-2002
		US 2010242098 A1	23-09-2010
		US 2013061294 A1	07-03-2013

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-9

   A method for packet filtering based on application-layer packet header information

   ---

2. claims: 10-18

   A method for selective packet digest logging

   ---

3. claims: 19, 20

   A method for security policy update and management

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2015/024691	07.04.2015	16.04.2014

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06 ADD. H04L29/08

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☒ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☒ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

SEP 2 - 2015

Name and mailing address of the ISA:	Date of completion of this opinion	Authorized Officer
European Patent Office D-80298 Munich Tel. +49 89 2399 - 0 Fax: +49 89 2399 - 4465	see form PCT/ISA/210	Barla Harter, I Telephone No. +49 89 2399-0

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed.

   ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43*bis*.1(a))

3. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing:

   a. ☐ forming part of the international application as filed:

      ☐ in the form of an Annex C/ST.25 text file.

      ☐ on paper or in the form of an image file.

   b. ☐ furnished together with the international application under PCT Rule 13*ter*.1(a) for the purposes of international search only in the form of an Annex C/ST.25 text file.

   c. ☐ furnished subsequent to the international filing date for the purposes of international search only:

      ☐ in the form of an Annex C/ST.25 text file (Rule 13*ter*.1(a)).

      ☐ on paper or in the form of an image file (Rule 13*ter*.1(b) and Administrative Instructions, Section 713).

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that forming part of the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability

The questions whether the claimed invention appears to be novel, to involve an inventive step (to be non obvious), or to be industrially applicable have not been examined in respect of

☐ the entire international application

☒ claims Nos. 10-20

because:

☐ the said international application, or the said claims Nos. relate to the following subject matter which does not require an international search *(specify)*:

☐ the description, claims or drawings *(indicate particular elements below)* or said claims Nos. are so unclear that no meaningful opinion could be formed *(specify)*:

☐ the claims, or said claims Nos. are so inadequately supported by the description that no meaningful opinion could be formed *(specify)*:

☒ no international search report has been established for the whole application or for said claims Nos. 10-20

☐ a meaningful opinion could not be formed without the sequence listing; the applicant did not, within the prescribed time limit:

- ☐ furnish a sequence listing in the form of an Annex C/ST.25 text file, and such listing was not available to the International Searching Authority in the form and manner acceptable to it; or the sequence listing furnished did not comply with the standard provided for in Annex C of the Administrative Instructions.
- ☐ furnish a sequence listing on paper or in the form of an image file complying with the standard provided for in Annex C of the Administrative Instructions, and such listing was not available to the International Searching Authority in the form and manner acceptable to it; or the sequence listing furnished did not comply with the standard provided for in Annex C of the Administrative Instructions.
- ☐ pay the required late furnishing fee for the furnishing of a sequence listing in response to an invitation under Rule 13*ter*.1(a) or (b).

☐ See Supplemental Box for further details

Box No. IV Lack of unity of invention

1. ☒ In response to the invitation (Form PCT/ISA/206) to pay additional fees, the applicant has, within the applicable time limit:

   ☐ paid additional fees

   ☐ paid additional fees under protest and, where applicable, the protest fee

   ☐ paid additional fees under protest but the applicable protest fee was not paid

   ☒ not paid additional fees

2. ☐ This Authority found that the requirement of unity of invention is not complied with and chose not to invite the applicant to pay additional fees.

3. This Authority considers that the requirement of unity of invention in accordance with Rule 13.1, 13.2 and 13.3 is

   ☐ complied with

   ☒ not complied with for the following reasons:

   see separate sheet

4. Consequently, this report has been established in respect of the following parts of the international application:

   ☐ all parts.

   ☒ the parts relating to claims Nos. 1-9

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	
	No:	Claims	1-9
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-9
Industrial applicability (IA)	Yes:	Claims	1-9
	No:	Claims	

2. Citations and explanations

see separate sheet

Re Item IV

Lack of unity of invention

1 This Authority considers that the application does not meet the requirements of unity of invention and that there are three inventions covered by the claims indicated as follows:

claims: 1-9

A method for packet filtering based on application-layer packet header information

claims: 10-18

A method for selective packet digest logging

claims: 19, 20

A method for security policy update and management

2 The reasons for which the inventions are not so linked as to form a single general inventive concept, as required by Rule 13.1 PCT, are as follows:

The lack of unity becomes apparent a posteriori, i.e. for the following reasoning document D1 (US20060136987) is taken into account. From this prior art document what is known is a method comprising a security policy management server ("the controller 13 provides the identifying rule and the processing rule" in paragraph [0093]; "the controller may delete from the identifying rule/processing rule setting table the identifying rule and the processing rule which have not been accessed for a predetermined time." in paragraph [0036]) as claimed in claim 1. These features are the only features common to all the groups of inventions.

With the reference to the prior art document D1:

The third group yields the potential special technical feature of receiving, by a security policy management server and from a computing device, a security update comprising a set of network addresses; updating, one or more rules; receiving, from a different computing device, a security update comprising a different set of network addresses; determining that the different set of network addresses includes at least a portion of network addresses included in the set of network addresses; and responsive to that identifying the at least a portion of network addresses included in the set of network addresses, at least one of the one or more rules stored in the memory that specifies a range of network addresses comprising the at least a portion of network addresses;

and updating the at least one of the one or more rules to include one or more other network addresses included in the different set of network addresses, hence solving the objective problem of how to manage security policies based on information received from different sources comprising overlapping network addresses.

Moreover, the first and second invention have further common features, which are also disclosed in document D1 as shown in the following. From document D1 what is further known is:

receiving ("receive the identifying rule and the processing rule from another communication apparatus" in paragraph [0040]), by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server ("the controller 13 provides the identifying rule and the processing rule" in paragraph [0093]; "while the notifying destination of the notifying message (identifying rule and the processing rule) 740 is made the router 100C in FIG. 8, the gateway apparatus, a switch, and a management system (not shown) may be made a notifying destination." in paragraph [0095]), a dynamic security policy ("receive the identifying rule and the processing rule from another communication apparatus" in paragraph [0040]; "the controller may delete from the identifying rule/processing rule setting table the identifying rule and the processing rule which have not been accessed for a predetermined time." in paragraph [0036]) that includes at least one rule specifying a packet transformation function to be performed on packets

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway ("received packet" in the abstract);

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets ("identify a protocol type of a communication packet transmitted/received by a device" in paragraph [0028]); and

performing, by the packet security gateway and on a packet-by-packet basis, the packet transformation function on each of the one or more packets ("to automatically set the processing rule of the packet" in paragraph [0028]; "the processing policy may include a policy concerning at least one of a service

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 338

quality class, filtering, and routing of the packet." in paragraph [0029]) as claimed in claim 1. These features are the only features common to the first two groups of inventions.

With the reference to the prior art document D1:

The first group yields the potential special technical feature of using application-layer packet-header information as the packet identification at the gateway, hence solving the objective problem of how to improve the scalability issues in high resolution filtering.

The second group yields the potential special technical feature of using a packet digest logging function to the identified packets, hence solving the objective problem of how to enable security awareness information in a network.

Consequently, neither the objective problems underlying the subjects of the three claimed inventions, nor the solutions as defined by the special technical features described allow for the link of a common inventive concept to be established between said inventions in the meaning of Rule 13.2 PCT. In conclusion, the three groups of claims are not linked by a single general inventive concept. The application hence does not meet the requirements of unity of invention as defined in Rule 13.1 PCT.

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1 Reference is made to the following documents:

D1	US 2006/136987 A1 (OKUDA MASATO [JP]) 22 June 2006 (2006-06-22)
D2	US 2007/240208 A1 (YU MING-CHE [TW] ET AL) 11 October 2007 (2007-10-11)

2 The present application does not meet the criteria of Article 33(1) PCT, because the subject-matter of claim 1 is not new in the sense of Article 33(2) PCT.

2.1 **Claim 1:**

2.1.1 Document **D1** discloses all the features of claim 1 (references in parentheses applying to document D1 and using the wording of this claim):

A method, comprising:
receiving ("receive the identifying rule and the processing rule from another communication apparatus" in paragraph [0040]),

by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server ("the controller 13 provides the identifying rule and the processing rule" in paragraph [0093]; "while the notifying destination of the notifying message (identifying rule and the processing rule) 740 is made the router 100C in FIG. 8, the gateway apparatus, a switch, and a management system (not shown) may be made a notifying destination." in paragraph [0095]),

a dynamic security policy ("receive the identifying rule and the processing rule from another communication apparatus" in paragraph [0040]; "the controller may delete from the identifying rule/processing rule setting table the identifying rule and the processing rule which have not been accessed for a predetermined time." in paragraph [0036])

that includes at least one rule specifying application-layer packet-header information and a packet transformation function to be performed on packets comprising the application-layer packet-header information ("filtering can be realized with a URL (Uniform Resource Locator) within an HTTP message" in paragraph [0135]);
receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway ("received packet" in the abstract);
identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets comprising the application-layer packet-header information ("identify a protocol type of a communication packet transmitted/received by a device" in paragraph [0028]; "the identifying rule/processing rule setting table may have at least one of a transmitting source address, a destination address, a protocol type, a transmitting source port number, and a destination port number value identified by the identifying policy as the identifying rule." in paragraph [0030], "filtering can be realized with a URL (Uniform Resource Locator) within an HTTP message" in

paragraph [0135]); and
performing, by the packet security gateway and on a packet-by-packet basis, the packet transformation function on each of the one or more packets comprising the application-layer packet-header information ("to automatically set the processing rule of the packet" in paragraph [0028]; "the processing policy may include a policy concerning at least one of a service quality class, filtering, and routing of the packet." in paragraph [0029]).

The subject-matter of claim 1 is therefore not new (Article 33(2) PCT).

2.1.2 Moreover, document **D2** discloses in paragraphs [0021], [0024], [0026], [0033]-[0036] and [0039] a proxy server filtering HTTP packets based on various information in the packet header like the IP address, port numbers and URL information. In D2 a single proxy server is disclosed attached to a LAN, which is configured by an administrator. It would be part of normal operation to use multiple such proxies for multiple LANs and to send the configuration information from a central administration device. Therefore, the subject-matter of claim 1 does not involve an inventive step in the sense of Article 33(3) PCT.

2.2 **Dependent claims 2-9:**

The additional features of dependent claims 2-9 do not add anything new or inventive to the independent claims, because these features are either known from the above mentioned prior art (D1 further discloses in paragraphs [0030], [0033] and [0135], and in figure 21 the filtering of the packets based on the 5-tuple of IP addresses, port numbers, protocol types and URL information in the HTTP packets, based on which the packets are either discarded or passed. Furthermore, D2 also discloses the filtering of the HTTP packets based on URL and IP address and/or port information and discarding or sending them to their destination based on the policy result.) or common measures.

PATENT COOPERATION TREATY

From the INTERNATIONAL BUREAU

PCT

NOTIFICATION CONCERNING TRANSMITTAL OF COPY OF INTERNATIONAL PRELIMINARY REPORT ON PATENTABILITY (CHAPTER I OF THE PATENT COOPERATION TREATY)
(PCT Rule 44bis.1(c))

To:

WRIGHT, Bradley
Banner & Witcoff, Ltd.
1100 13th Street, N.W.
Suite 1200
Washington, District of Columbia 20005-4051
ETATS-UNIS D'AMERIQUE

Date of mailing *(day/month/year)* 07 May 2015 (07.05.2015)	
Applicant's or agent's file reference 007742.00023	**IMPORTANT NOTICE**

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2013/057502	30 August 2013 (30.08.2013)	22 October 2012 (22.10.2012)

Applicant
CENTRIPETAL NETWORKS, INC.

The International Bureau transmits herewith a copy of the international preliminary report on patentability (Chapter I of the Patent Cooperation Treaty)

The International Bureau of WIPO
34, chemin des Colombettes
1211 Geneva 20, Switzerland

Facsimile No. +41 22 338 82 70

Authorized officer
Kihwan Moon

e-mail: pt01.pct@wipo.int

PATENT COOPERATION TREATY

PCT

INTERNATIONAL PRELIMINARY REPORT ON PATENTABILITY

(Chapter I of the Patent Cooperation Treaty)

(PCT Rule 44*bis*)

Applicant's or agent's file reference 007742.00023	**FOR FURTHER ACTION**	See item 4 below
International application No. PCT/US2013/057502	International filing date *(day/month/year)* 30 August 2013 (30.08.2013)	Priority date *(day/month/year)* 22 October 2012 (22.10.2012)
International Patent Classification (8th edition unless older edition indicated) See relevant information in Form PCT/ISA/237		
Applicant CENTRIPETAL NETWORKS, INC.		

1. This international preliminary report on patentability (Chapter I) is issued by the International Bureau on behalf of the International Searching Authority under Rule 44 *bis*.1(a).

2. This REPORT consists of a total of 5 sheets, including this cover sheet.

   In the attached sheets, any reference to the written opinion of the International Searching Authority should be read as a reference to the international preliminary report on patentability (Chapter I) instead.

3. This report contains indications relating to the following items:

☒	Box No. I	Basis of the report
☐	Box No. II	Priority
☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
☐	Box No. IV	Lack of unity of invention
☒	Box No. V	Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
☐	Box No. VI	Certain documents cited
☐	Box No. VII	Certain defects in the international application
☐	Box No. VIII	Certain observations on the international application

4. The International Bureau will communicate this report to designated Offices in accordance with Rules 44*bis*.3(c) and 93*bis*.1 but not, except where the applicant makes an express request under Article 23(2), before the expiration of 30 months from the priority date (Rule 44*bis* .2).

	Date of issuance of this report 28 April 2015 (28.04.2015)
The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland Facsimile No. +41 22 338 82 70	Authorized officer Kihwan Moon e-mail: pt01.pct@wipo.int

Form PCT/IB/373 (January 2004)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2013/057502	30.08.2013	22.10.2012

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA:	Date of completion of this opinion	Authorized Officer
European Patent Office P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk - Pays Bas Tel. +31 70 340 - 2040 Fax. +31 70 340 - 3016	see form PCT/ISA/210	Veen, Gerardus Telephone No. +31 70 340-3811

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 344

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☑ the international application in the language in which it was filed

   ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43bis.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing filed or furnished:

   a. (means)

      ☐ on paper

      ☐ in electronic form

   b. (time)

      ☐ in the international application as filed

      ☐ together with the international application in electronic form

      ☐ subsequently to this Authority for the purposes of search

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	
	No:	Claims	1-20
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-20
Industrial applicability (IA)	Yes:	Claims	1-20
	No:	Claims	

2. Citations and explanations

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1 The following documents are referred to:

D1 EP 1 677 484 A2 (MICROSOFT CORP [US]) 5 July 2006 (2006-07-05)

D2 US 6 226 372 B1 (BEEBE TODD [US] ET AL) 1 May 2001 (2001-05-01)

2 The present application does not meet the criteria of Article 33(2) PCT, because the subject-matter of claims 1-20 is not new.

2.1 In the words of independent claim 1, D1 discloses:

'A method (column 4, line 13), comprising:

at each of one or more packet security gateways (c.4, l.16, 'enforcement engine') associated with a security policy management server (c.4, l.15, 'policy server'):

receiving a dynamic security policy from the security policy management server (c.4, l.16-19);

receiving packets associated with a network protected by each respective packet security gateway; and

performing, on a packet by packet basis, at least one of multiple packet transformation functions specified by the dynamic security policy on the packets (c.4, l.19-20, 'enforcing the security policies'), wherein performing the at least one of multiple packet transformation functions specified by the dynamic security policy on the packets comprises performing at least one packet transformation function other than forwarding or dropping the packets (c.5, l.30-36)'.

Claim 1 is thus not new (Art. 33(2) PCT).

2.2 The same reasoning applies, mutatis mutandis, to independent claims 19 and 20 which are, therefore, also not new.

2.3 Dependent claims 2-18 do not contain any features which, in combination with the features of the claims to which they refer, meet the requirements of the PCT in respect of novelty and inventive step, see D1, D2 and the corresponding citations in the search report.

PATENT COOPERATION TREATY

From the INTERNATIONAL BUREAU

PCT

NOTIFICATION CONCERNING
TRANSMITTAL OF COPY OF INTERNATIONAL
PRELIMINARY REPORT ON PATENTABILITY
(CHAPTER I OF THE PATENT COOPERATION
TREATY)
(PCT Rule 44bis.1(c))

To:

WRIGHT, Bradley C.
Banner & Witcoff, Ltd.
1100 13th Street, N.W.
Suite 1200
Washington, District of Columbia 20005-4051
ETATS-UNIS D'AMERIQUE

Date of mailing *(day/month/year)* 24 September 2015 (24.09.2015)	
Applicant's or agent's file reference 007742.00030	**IMPORTANT NOTICE**

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2014/023286	11 March 2014 (11.03.2014)	12 March 2013 (12.03.2013)

Applicant
CENTRIPETAL NETWORKS, INC.

The International Bureau transmits herewith a copy of the international preliminary report on patentability (Chapter I of the Patent Cooperation Treaty).

The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland	Authorized officer Philippe Bécamel
Facsimile No. +41 22 338 82 70	e-mail: pt03.pct@wipo.int

Form PCT/IB/326 (January 2004)

PATENT COOPERATION TREATY

PCT

INTERNATIONAL PRELIMINARY REPORT ON PATENTABILITY
(Chapter I of the Patent Cooperation Treaty)

(PCT Rule 44*bis*)

Applicant's or agent's file reference 007742.00030	**FOR FURTHER ACTION**	See item 4 below
International application No. PCT/US2014/023286	International filing date *(day/month/year)* 11 March 2014 (11.03.2014)	Priority date *(day/month/year)* 12 March 2013 (12.03.2013)
International Patent Classification (8th edition unless older edition indicated) See relevant information in Form PCT/ISA/237		
Applicant CENTRIPETAL NETWORKS, INC.		

1. This international preliminary report on patentability (Chapter I) is issued by the International Bureau on behalf of the International Searching Authority under Rule 44 *bis*.1(a).

2. This REPORT consists of a total of 6 sheets, including this cover sheet.

   In the attached sheets, any reference to the written opinion of the International Searching Authority should be read as a reference to the international preliminary report on patentability (Chapter I) instead.

3. This report contains indications relating to the following items:

☒	Box No. I	Basis of the report
☐	Box No. II	Priority
☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
☐	Box No. IV	Lack of unity of invention
☒	Box No. V	Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
☐	Box No. VI	Certain documents cited
☐	Box No. VII	Certain defects in the international application
☒	Box No. VIII	Certain observations on the international application

4. The International Bureau will communicate this report to designated Offices in accordance with Rules 44*bis*.3(c) and 93*bis*.1 but not, except where the applicant makes an express request under Article 23(2), before the expiration of 30 months from the priority date (Rule 44*bis* .2).

	Date of issuance of this report 15 September 2015 (15.09.2015)
The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland Facsimile No. +41 22 338 82 70	Authorized officer Philippe Bécamel e-mail: pt03.pct@wipo.int

Form PCT/IB/373 (January 2004)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2014/023286	11.03.2014	12.03.2013

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☐ Box No. VII Certain defects in the international application
- ☒ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA:	Date of completion of this opinion	Authorized Officer
European Patent Office P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk - Pays Bas Tel. +31 70 340 - 2040 Fax: +31 70 340 - 3016	see form PCT/ISA/210	Veen, Gerardus Telephone No. +31 70 340-3811

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed

   ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43bis.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing filed or furnished:

   a. (means)

      ☐ on paper

      ☐ in electronic form

   b. (time)

      ☐ in the international application as filed

      ☐ together with the international application in electronic form

      ☐ subsequently to this Authority for the purposes of search

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	
	No:	Claims	1-20
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-20
Industrial applicability (IA)	Yes:	Claims	1-20
	No:	Claims	

2. Citations and explanations

see separate sheet

Box No. VIII Certain observations on the international application

The following observations on the clarity of the claims, description, and drawings or on the question whether the claims are fully supported by the description, are made:

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

The following documents are referred to:

D1 US 2003/123456 A1 (DENZ PETER R [US] ET AL) 3 July 2003 (2003-07-03)

D2 US 2007/083924 A1 (LU HONGQIAN K [US]) 12 April 2007 (2007-04-12)

1 The present application does not meet the criteria of Article 33(2) PCT, because the subject-matter of claims 1-20 is not new.

1.1 Independent claim 1 (as far as it is clear, see **Item VIII**) defines a method for filtering a packet, the method having two stages:

1) in the first stage, a first operation ('operator') is performed upon the packet, based on a first header field value;

2) in the second stage, a second operation ('transformation') is performed upon the packet, based on a second header field value.

Moreover, these two steps are interrelated in the sense that the criteria, with which the second header field value is compared during the second stage, are determined during the first stage.

Such a method is also disclosed in D1 (see the passages cited in the search report).

Present claim 1 is thus not new according to Art. 33(2) PCT.

1.2 The same reasoning applies, mutatis mutandis, to independent claims 11 and 20 which are, therefore, also not new.

1.3 Dependent claims 2-10 and 12-19 do not appear to contain any features which, in combination with the features of the claims to which they refer, meet the requirements of the PCT in respect of novelty and inventive step, see D1, D2 and the corresponding references cited in the search report.

1.4 It is further noted that a reasoning based on D2 would lead to very similar conclusions.

Re Item VIII

Certain observations on the international application

2 Claims 1-20 do not meet the requirements of Article 6 PCT for, at least, the following reasons:

- in claim 1, the expressions 'header field' and 'application header field', and especially how they differ from each other, are not clear;

- a similar remark applies to the expressions 'operator' and 'packet transformation function', respectively.

Similar observations apply to independent claims 11 and 20.

PATENT COOPERATION TREATY

From the INTERNATIONAL BUREAU

PCT

NOTIFICATION CONCERNING
TRANSMITTAL OF COPY OF INTERNATIONAL
PRELIMINARY REPORT ON PATENTABILITY
(CHAPTER I OF THE PATENT COOPERATION
TREATY)
(PCT Rule 44bis.1(c))

To:

WRIGHT, Bradley C.
Banner & Witcoff, Ltd.
1100 13th Street, N.W.
Suite 1200
Washington, District of Columbia 20005-4051
ETATS-UNIS D'AMERIQUE

Date of mailing *(day/month/year)* 24 September 2015 (24.09.2015)	
Applicant's or agent's file reference 007742.00031	**IMPORTANT NOTICE**

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2014/027723	14 March 2014 (14.03.2014)	15 March 2013 (15.03.2013)

Applicant
CENTRIPETAL NETWORKS, INC.

The International Bureau transmits herewith a copy of the international preliminary report on patentability (Chapter I of the Patent Cooperation Treaty).

The International Bureau of WIPO
34, chemin des Colombettes
1211 Geneva 20, Switzerland

Facsimile No. +41 22 338 82 70

Authorized officer
Nora Lindner

e-mail: pt05.pct@wipo.int

PATENT COOPERATION TREATY

PCT

INTERNATIONAL PRELIMINARY REPORT ON PATENTABILITY
(Chapter I of the Patent Cooperation Treaty)

(PCT Rule 44*bis*)

Applicant's or agent's file reference 007742.00031	**FOR FURTHER ACTION**	See item 4 below
International application No. PCT/US2014/027723	International filing date *(day/month/year)* 14 March 2014 (14.03.2014)	Priority date *(day/month/year)* 15 March 2013 (15.03.2013)
International Patent Classification (8th edition unless older edition indicated) See relevant information in Form PCT/ISA/237		
Applicant CENTRIPETAL NETWORKS, INC.		

1. This international preliminary report on patentability (Chapter I) is issued by the International Bureau on behalf of the International Searching Authority under Rule 44 *bis*.1(a).

2. This REPORT consists of a total of 7 sheets, including this cover sheet.

   In the attached sheets, any reference to the written opinion of the International Searching Authority should be read as a reference to the international preliminary report on patentability (Chapter I) instead.

3. This report contains indications relating to the following items:

☒	Box No. I	Basis of the report
☐	Box No. II	Priority
☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
☐	Box No. IV	Lack of unity of invention
☒	Box No. V	Reasoned statement under Article 35(2) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
☐	Box No. VI	Certain documents cited
☒	Box No. VII	Certain defects in the international application
☐	Box No. VIII	Certain observations on the international application

4. The International Bureau will communicate this report to designated Offices in accordance with Rules 44*bis*.3(c) and 93*bis*.1 but not, except where the applicant makes an express request under Article 23(2), before the expiration of 30 months from the priority date (Rule 44*bis* .2).

	Date of issuance of this report 15 September 2015 (15.09.2015)
The International Bureau of WIPO 34, chemin des Colombettes 1211 Geneva 20, Switzerland Facsimile No. +41 22 338 82 70	Authorized officer Nora Lindner e-mail: pt05.pct@wipo.int

Form PCT/IB/373 (January 2004)

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2014/027723	14.03.2014	15.03.2013

International Patent Classification (IPC) or both national classification and IPC
INV. H04L29/06

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☒ Box No. VII Certain defects in the international application
- ☐ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA:	Date of completion of this opinion	Authorized Officer
European Patent Office D-80298 Munich Tel. +49 89 2399 - 0 Fax: +49 89 2399 - 4465	see form PCT/ISA/210	Olaechea, Javier Telephone No. +49 89 2399-8233

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 357

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed

   ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43bis.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing filed or furnished:

   a. (means)

      ☐ on paper

      ☐ in electronic form

   b. (time)

      ☐ in the international application as filed

      ☐ together with the international application in electronic form

      ☐ subsequently to this Authority for the purposes of search

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	
	No:	Claims	1-20
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-20
Industrial applicability (IA)	Yes:	Claims	1-20
	No:	Claims	

2. Citations and explanations

see separate sheet

Box No. VII Certain defects in the international application

The following defects in the form or contents of the international application have been noted:

see separate sheet

Re Item V

Reasoned statement with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1 Reference is made to the following document:

D1 US 2006/048142 A1 (ROESE JOHN J [US] ET AL) 2 March 2006 (2006-03-02)

D2 Sean Moore: "SBIR Case Study: Centripetal Networks Subtitle: How CNI Leveraged DHS S&T SBIR Funding to Launch a Successful Cyber Security Company Cyber Security Division 2012 Principal Investigators' Meeting 10-October-2012", , 10 October 2012 (2012-10-10), XP055122227, Retrieved from the Internet: URL:http://www.dhs.gov/sites/default/files/publications/csd-day-2-04-sbir-centripetal-price.pdf [retrieved on 2014-06-06]

D3 Cisco: "Control Plane Policing Implementation Best Practices", , 13 March 2013 (2013-03-13), XP002725728, Retrieved from the Internet: URL:https://web.archive.org/web/20130313135143/http://www.cisco.com/web/about/security/intelligence/coppwp_gs.html [retrieved on 2014-06-12]

2 The present application does not meet the criteria of Article 33(1) PCT, because the subject-matter of claims 1, 11 and 16 is not new in the sense of Article 33(2) PCT.

2.1 The document D1 discloses (the references in parentheses applying to this document) in terms of claim 1:

A method, comprising:

at a computing platform comprising at least one processor, memory, and communication interface (*network device in Figures 1 and 2*):

receiving, via the communication interface, a plurality of packets (paragraph 55; *traffic is received by the network device*);
responsive to a determination that an overload condition has occurred in one or more networks associated with the computing platform, applying, by the processor and to at least some of the plurality of packets, a first group of packet-filtering rules that comprises at least one five-tuple indicating a first set of packets that should be allowed to continue toward their respective destinations, wherein applying the first group of packet-filtering rules includes allowing at least a first portion of the plurality of packets that fall within the first set of packets to continue toward their respective destinations (paragraphs 24, 56 and table in pages 9-10; *in response to an attack trigger, a rule set is activated by the network device, said rule set including disabling all traffic excluding the management traffic, RR1*); and

responsive to a determination that the overload condition has been mitigated to a first degree, applying, to at least some of the plurality of packets, a second group of packet-filtering rules that comprises at least one five-tuple indicating a second set of packets that should be allowed to continue toward their respective destinations, wherein applying the second group of packet-filtering rules includes allowing at least a second portion of the plurality of packets that fall within the second set of packets to continue toward their respective destinations, and wherein the second group of packet-filtering rules is less restrictive than the first group of packet-filtering rules (paragraph 24; *upon detection of a new triggering condition, a less restrictive set of filtering rules is implemented by the networking device*).

Since all the features of claim 1 are known in combination from document D1, the subject-matter of claim 1 is not new (Article 33(2) PCT). Therefore, claim 1 is not allowable (Article 33(1) PCT).

2.2 The same reasoning applies, mutatis mutandis, to the subject-matter of the corresponding independent claims 11 and 16, which therefore are also considered not new (Article 33(2) PCT).

2.3 Furthermore, it is noted that the disclosure of document D2 (see page 4) would also render the subject-matter of claims 1, 11 and 16 not new or at least not inventive.

3 None of the dependent claims 2-10, 12-15 and 17-20 contain additional features which lead the subject-matter to be both new and inventive (see documents D1-D3 and the corresponding passages cited in the search report).

Re Item VII

Certain defects in the international application

The independent claims are not in the two-part form in accordance with Rule 6.3(b) PCT, which in the present case would be appropriate, with those features known in combination from the prior art being placed in the preamble (Rule 6.3(b)(i) PCT) and with the remaining features being included in the characterising part (Rule 6.3(b)(ii) PCT).

The features of the claims are not provided with reference signs placed in parentheses (Rule 6.2(b) PCT).

The description is not in conformity with the claims as required by Rule 5.1(a)(iii) PCT.

Contrary to the requirements of Rule 5.1(a)(ii) PCT, the relevant background art disclosed in the documents D1-D3 is not mentioned in the description, nor are these documents identified therein.

Electronic Acknowledgement Receipt

EFS ID:	23643542
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	30-SEP-2015
Filing Date:	16-APR-2014
Time Stamp:	08:23:20
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	1035224 044b9e2edd5b4323b66ae4a44a3cbe01587762f2	no	4

Warnings:	
Information:	

2	Non Patent Literature	NPL1.pdf	1165116 696ebe36087e31a3b026e354c2144c4ea999d2f9	no	16
Warnings:					
Information:					
3	Non Patent Literature	NPL2.pdf	4652984 1e4a73420967afeb7605a1ae68dd2351e973c318	no	6
Warnings:					
Information:					
4	Non Patent Literature	NPL3.pdf	205308 d516285c1596bc34a6ed7ffda89d30875579b9c1	no	7
Warnings:					
Information:					
5	Non Patent Literature	NPL4.pdf	265145 57a66a25537b108cae0d2c77f61ced460edd2263	no	8
Warnings:					
Information:					
		Total Files Size (in bytes):	7323777		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PLUS Search Results for S/N 14253992, Searched Thu Oct 08 15:29:07 EDT 2015

The Patent Linguistics Utility System (PLUS) is a USPTO automated search system for U.S. Patents from 1971 to the present PLUS is a query-by-example search system which produces a list of patents that are most closely related linguistically to the application searched. This search was prepared by the staff of the Scientific and Technical Information Center, SIRA.

5940590 99
6078586 99
4885789 97
5414833 97
5455855 97
5694335 97
5710760 97
5742756 97
5805801 97
5905859 97
5978459 97
5982507 97
5987520 97
5991401 97
6161136 97
4348661 95
4423287 95
4578530 95
4768848 95
4774658 95
4815128 95
4815741 95
4825195 95
4837799 95
4873517 95
4881263 95
4907260 95
4920567 95
4924513 95
5210795 95
5276735 95
5280529 95
5309564 95
5329623 95
5349643 95
5394471 95
5404395 95
5408618 95
5416781 95
5442708 95
5444782 95
5448765 95
5481265 95
5491752 95
5504810 95
5539824 95
5550981 95
5568126 95
5598456 95
5602973 95

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
14/253,992	04/16/2014	Steven Rogers	007742.00029	5163

22907 7590 10/15/2015

BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051

EXAMINER	
HUANG, CHENG-FENG	

ART UNIT	PAPER NUMBER
2497	

MAIL DATE	DELIVERY MODE
10/15/2015	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No. 14/253,992	Applicant(s) ROGERS ET AL.	
	Examiner CHENG-FENG HUANG	Art Unit 2497	AIA (First Inventor to File) Status Yes

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 2 MONTHS FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *04/16/2014*.

☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2a)☐ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

4)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims*

5)☒ Claim(s) *1-20* is/are pending in the application.

5a) Of the above claim(s) _____ is/are withdrawn from consideration.

6)☐ Claim(s) _____ is/are allowed.

7)☐ Claim(s) _____ is/are rejected.

8)☐ Claim(s) _____ is/are objected to.

9)☒ Claim(s) *1-20* are subject to restriction and/or election requirement.

* If any claims have been determined allowable, you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to PPHfeedback@uspto.gov.

Application Papers

10)☐ The specification is objected to by the Examiner.

11)☒ The drawing(s) filed on *04/16/2014* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

a)☐ All b)☐ Some** c)☐ None of the:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

** See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)

2) ☒ Information Disclosure Statement(s) (PTO/SB/08a and/or PTO/SB/08b) Paper No(s)/Mail Date *See Continuation Sheet*.

3) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .

4) ☐ Other: _____.

Continuation of Attachment(s) 2). Information Disclosure Statement(s) (PTO/SB/08), Paper No(s)/Mail Date :04/16/2014, 06/25/2014, 06/27/2014, 07/03/2014, 09/09/2014, 10/27/2014, 11/25/2014, 02/18/2015, 03/13/2015, 04/20/2015, 06/19/2015, 07/21/2015, 07/24/2015, 09/30/2015. .

DETAILED ACTION

1. In communications filed on 04/16/2014, claims 1-20 are pending in this examination.

Election/Restrictions

2. Restriction to one of the following inventions is required under 35 U.S.C. 121:

I. **Group 1**: Claims 1-18 are drawn to managing network security proactively by applying dynamic network security policies based on packet-identification criteria and packet transformation function; classified in CPC H04L 63/20.

II. **Group 2**: Claims 19-20 are drawn to updating network security rules by identifying and determining different set of network addresses; classified in CPC H04L 63/0263.

3. The inventions are distinct, each from the other because of the following reasons:

Inventions group 1 and group 2 are related as subcombinations disclosed as usable together in a single combination. The subcombinations are distinct if they do not overlap in scope and are not obvious variants, and if it is shown that at least one subcombination are separately usable. In the instant case, subcombination group 1 has separate utility of providing a proactive solution for protecting secured network by applying dynamic network security policies whereas subcombination group 2 has separate utility of managing network security rules and details related to receiving/updating/combining network security rules based on identifying and

determining different set of network addresses. Hence they are not obvious variants of each other and they don't overlap in scope. See MPEP § 806.05(d).

The examiner has required restriction between subcombinations usable together. Where applicant elects a subcombination and claims thereto are subsequently found allowable, any claim(s) depending from or otherwise requiring all the limitations of the allowable subcombination will be examined for patentability in accordance with 37 CFR 1.104. See MPEP § 821.04(a). Applicant is advised that if any claim presented in a continuation or divisional application is anticipated by, or includes all the limitations of, a claim that is allowable in the present application, such claim may be subject to provisional statutory and/or nonstatutory double patenting rejections over the claims of the instant application.

4. Restriction for examination purposes as indicated is proper because all these inventions listed in this action are independent or distinct for the reasons given above and there would be a serious search and examination burden if restriction were not required because one or more of the following reasons apply:

(a) the inventions have acquired a separate status in the art in view of their different classification;

(b) the inventions have acquired a separate status in the art due to their recognized divergent subject matter;

(c) the inventions require a different field of search (for example, searching different classes/subclasses or electronic resources, or employing different search queries);

(d) the prior art applicable to one invention would not likely be applicable to another invention;

(e) the inventions are likely to raise different non-prior art issues under 35 U.S.C. 101 and/or 35 U.S.C. 112, first paragraph.

5. Applicant is advised that the reply to this requirement to be complete must include (i) an election of a species or invention to be examined even though the requirement may be traversed (37 CFR 1.143) and (ii) identification of the claims encompassing the elected invention.

The election of an invention may be made with or without traverse. To preserve a right to petition, the election must be made with traverse. If the reply does not distinctly and specifically point out supposed errors in the restriction requirement, the election shall be treated as an election without traverse. Traversal must be presented at the time of election in order to be considered timely. Failure to timely traverse the requirement will result in the loss of right to petition under 37 CFR 1.144. If claims are added after the election, applicant must indicate which of these claims are readable on the elected invention.

If claims are added after the election, applicant must indicate which of these claims are readable upon the elected invention.

Should applicant traverse on the ground that the inventions are not patentably distinct, applicant should submit evidence or identify such evidence now of record showing the inventions to be obvious variants or clearly admit on the record that this is the case. In either instance, if the examiner finds one of the inventions unpatentable

over the prior art, the evidence or admission may be used in a rejection under 35 U.S.C. 103(a) of the other invention.

6. Applicant is reminded that upon the cancellation of claims to a non-elected invention, the inventorship must be amended in compliance with 37 CFR 1.48(b) if one or more of the currently named inventors in no longer an inventor of at least one claim remaining in the application. Any amendment of inventorship must be accompanied by a petition under 37 CFR 1.48(b) and by the fee required under 37 CFR 1.17(1).

Conclusion

7. A shortened statutory period for reply to this action is set to expire TWO MONTHS from the mailing date of this action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to CHENG-FENG HUANG whose telephone number is (571)272-6186. The examiner can normally be reached on 7:30 am ~ 5:00 pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Hadi Armouche can be reached on 571-270-3618. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/CHENG-FENG HUANG/
Examiner, Art Unit 2497

/HADI ARMOUCHE/
Supervisory Patent Examiner, Art Unit 2497

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7721084		2010-05-18	Salminen et al.	
	2	7684400		2010-03-23	Govindarajan et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20120113987	A1	2012-05-10	Riddoch et al.	
	2	20110270956	A1	2011-11-03	McDysan et al.	
	3	20100211678	A1	2010-08-19	McDysan et al.	
	4	20050117576		2005-06-02	McDysan et al.	

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

	5	20040151155	A1	2004-08-05	Jouppi	
	6	20030212900	A1	2003-11-13	Liu et al.	
	7	20030188192	A1	2003-10-02	Tang et al.	
	8	20120264443	A1	2012-10-18	Ng et al.	
	9	20130059527	A1	2013-03-07	HASESAKA et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

EXAMINER SIGNATURE			
Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2431 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7539186	B2	2009-05-26	Aerrabotu et al.	
	2	7710885	B2	2010-05-04	Ilnicki et al.	
	3	8856926	B2	2014-10-07	Narayanaswamy et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2431 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	REUMANN, John et al., "Adaptive Packet Filters," Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, MI, IEEE, 2001.	☐
	2	GREENWALD, Michael et al., "Designing an Academic Firewall: Policy, Practice, and Experience with SURF," Department of Computer Science, Stanford University, Stanford, CA, IEEE, Proceedings of SNDSS, 1996.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	6317837	B1	2001-11-13	Kenworthy	
	2	8306994	B2	2012-11-06	Kenworthy	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20020049899	A1	2002-04-25	Kenworthy	
	2	20060136987	A1	2006-06-22	Okuda	
	3	20060146879	A1	2006-07-06	Anthias et al.	
	4	20070240208	A1	2007-10-11	Yu et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

	5	20100242098	A1	2010-09-23	Kenworthy	
	6	20130061294	A1	2013-03-07	Kenworthy	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	Communication Relating to the Results of the Partial International Search for International App. No. PCT/ US2015/024691, dated July 10, 2015.	☐
	2	International Preliminary Report on Patentability for International App. No. PCT/US2013/072566, dated July 23, 2015.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	8726379	B1	2014-05-13	Stiansen et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20140366132	A1	2014-12-11	Stiansen et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

	1	International Search Report and Written Opinion for International App. No. PCT/US2015/024691, dated September 16, 2015.	☐
	2	International Preliminary Report on Patentability for International App. No. PCT/US2013/057502, dated May 7, 2015.	☐
	3	International Preliminary Report on Patentability for International App. No. PCT/US2014/023286, dated September 24, 2015.	☐
	4	International Preliminary Report on Patentability for International App. No. PCT/US2014/027723, dated September 24, 2015.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7478429	B2	2009-01-13	Lyon	
	2	7818794	B2	2010-10-19	Wittman	
	3	8935785	B2	2015-01-13	Pandrangi	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2432 2497
Examiner Name	Barron Jr., Gilberto Cheng-Feng Huang
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	8806638	B1	2014-08-12	Mani	
	2	6611875	B1	2003-08-26	Chopra et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20090328219	A1	2009-12-31	Narayanaswamy	
	2	20120240135	A1	2012-09-20	Risbood et al.	
	3	20120314617	A1	2012-12-13	Erichsen et al.	
	4	20030097590	A1	2003-05-22	Syvanne	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2432 2497
Examiner Name	Barron Jr., Gilberto Cheng-Feng Huang
Attorney Docket Number	007742.00029

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2]	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20040088542	A1	2004-05-06	Daude et al.	
	2	20080235755	A1	2008-09-25	Blaisdell et al.	
	3	20100132027	A1	2010-05-27	Ou	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	Mizuno et al., A New Remote Configurable Firewall System for Home-use Gateways, NTT Information Sharing Platform Laboratories, IEEE, 2004.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	8004994	B1	2011-08-23	Darisi et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20030154297	A1	2003-08-14	Suzuki et al.	
	2	20060114899	A1	2006-06-01	Toumura et al.	
	3	20110185055	A1	2011-07-28	Nappier et al.	
	4	20120331543	A1	2012-12-27	Bostrom et al.	
	5	20130254766	A1	2013-09-26	Zuo et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

BIB DATA SHEET

CONFIRMATION NO. 5163

SERIAL NUMBER	FILING or 371(c) DATE RULE	CLASS	GROUP ART UNIT	ATTORNEY DOCKET NO.
14/253,992	04/16/2014	726	2497	007742.00029

APPLICANTS
Centripetal Networks, Inc., Reston, VA;

INVENTORS
Steven Rogers, Leesburg, VA;
Sean Moore, Hollis, NH;
David K. Ahn, Winston-Salem, NC;
Peter P. Geremia, Portsmouth, NH;

** **CONTINUING DATA** **********

** **FOREIGN APPLICATIONS** **********

** **IF REQUIRED, FOREIGN FILING LICENSE GRANTED** ** SMALL ENTITY **
05/02/2014

Foreign Priority claimed ☐ Yes ☑ No 35 USC 119(a-d) conditions met ☐ Yes ☑ No Verified and Acknowledged /Cheng-Feng Huang/ Examiner's Signature	☐ Met after Allowance Initials	STATE OR COUNTRY	SHEETS DRAWINGS	TOTAL CLAIMS	INDEPENDENT CLAIMS
		VA	13	20	3

ADDRESS

BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051
UNITED STATES

TITLE

METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

FILING FEE RECEIVED		
730	FEES: Authority has been given in Paper No._____ to charge/credit DEPOSIT ACCOUNT No._____ for following:	☐ All Fees ☐ 1.16 Fees (Filing) ☐ 1.17 Fees (Processing Ext. of time) ☐ 1.18 Fees (Issue) ☐ Other _____ ☐ Credit

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	14253992	ROGERS ET AL.
	Examiner	Art Unit
	CHENG-FENG HUANG	2497

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	10/09/2015								
	1	÷								
	2	÷								
	3	÷								
	4	÷								
	5	÷								
	6	÷								
	7	÷								
	8	÷								
	9	÷								
	10	÷								
	11	÷								
	12	÷								
	13	÷								
	14	÷								
	15	÷								
	16	÷								
	17	÷								
	18	÷								
	19	÷								
	20	÷								

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 393

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	Not Yet Assigned Cheng-Feng Huang
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20030123456	A1	2003-07-03	Denz et al.	
	2	20070083924	A1	2007-04-12	Lu	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	Not Yet Assigned Cheng-Feng Huang
Attorney Docket Number	007742.00029

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	International Search Report off International Application No. PCT/US2014/023286, dated June 24, 2014.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	Not Yet Assigned Cheng-Feng Huang
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20060195896		2006-08-31	Fulp et al.	
	2	20060248580		2006-11-02	Fulp et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	1313290	EP	A1	2003-05-21	Stonesoft Corp		☐
	2	2005046145	WO	A1	2005-05-19	Ericsson Telefon Ab L M		☐

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	Not Yet Assigned Cheng-Feng Huang
Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	ISR off International Application No. PCT/US2013/072566, dated March 24, 2014.	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

Doc code: IDS
Doc description: Information Disclosure Statement (IDS) Filed

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2432 2497
Examiner Name	Barron Jr., Gilberto Cheng-Feng Huang
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20090172800	A1	2009-07-02	Wool	
	2	20080005795	A1	2008-01-03	Acharya et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2432 2497
Examiner Name	Barron Jr., Gilberto Cheng-Feng Huang
Attorney Docket Number	007742.00029

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	Not Yet Assigned Cheng-Feng Huang
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20060048142		2006-03-02	Roese et al.	
	2	20100011434		2010-01-14	Kay	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	2385676	EP	A1	2011-11-09	Alcatel Lucent		☐
	2	2498442	EP	A1	2012-09-12	Openet Telecom Ltd		☐

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	Not Yet Assigned Cheng-Feng Huang
Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	ISR off International Application No. PCT/US2014/027723, dated June 26, 2014.	☐
	2	"SBIR Case Study: Centripetal Networks Subtitle: How CNI Leveraged DHS S&T SBIR Funding to Launch a Successful Cyber Security Company, Cyber Security Division, 2012 Principal Investigators' Meeting"; Sean Moore, October 10, 2012.	☐
	3	"Control Plane Policing Implementation Best Practices"; Cisco Systems; March 13, 2013; <https://web.archive.org/web/20130313135143/http:www.cisco.com/web/about/security/intelligence/coppwp_gs.html>	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	TBA 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	6226372		2001-05-01	Beebe et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	US200601958		2006-08-31	Fulp et al.	
	2	US200602485		2006-11-02	Fulp et al.	
	3	20110055916		2011-03-03	Ahn	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	1677484	EP	A2	2006-07-05	Microsoft Corp [us]		☐

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	TBA 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	John Kindervag; "Build Security Into Your Network's DNA: The Zero Trust Network Architecture", Forrester Research Inc.; Novermber 5, 2010, pps: 1-26.	☐
	2	Palo Alto Networks; "Designing A Zero Trust Network With Next-Generation Firewalls"; pps. 1-10; last viewed on October 21, 2012.	☐
	3	Copy of International Search Report for PCT/US2013/057502, dated November 7, 2013	☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20010039624	A1	2001-11-08	Kellum	
	2	20030035370	A1	2003-02-20	Brustoloni	
	3	20030142681	A1	2003-07-31	Chen et al.	
	4	20030145225	A1	2003-07-31	Bruton et al.	
	5	20050138204	A1	2005-06-23	Iyer et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

	6	20050141537	A1	2005-06-30	Kumar et al.	
	7	20060212572	A1	2006-09-21	Afek et al.	
	8	20080072307	A1	2008-03-20	Maes	
	9	20100082811	A1	2010-04-01	Van Der Merwe et al.	
	10	20100296441	A1	2010-11-25	Barkan	
	11	20110088092	A1	2011-04-14	Nguyen et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2431 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	10/08/2015

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
14/253,992	04/16/2014	Steven Rogers	007742.00029

CONFIRMATION NO. 5163

22907
BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051

PUBLICATION NOTICE

OC000000078212365

Title:METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

Publication No.US-2015-0304354-A1
Publication Date:10/22/2015

NOTICE OF PUBLICATION OF APPLICATION

The above-identified application will be electronically published as a patent application publication pursuant to 37 CFR 1.211, et seq. The patent application publication number and publication date are set forth above.

The publication may be accessed through the USPTO's publically available Searchable Databases via the Internet at www.uspto.gov. The direct link to access the publication is currently http://www.uspto.gov/patft/.

The publication process established by the Office does not provide for mailing a copy of the publication to applicant. A copy of the publication may be obtained from the Office upon payment of the appropriate fee set forth in 37 CFR 1.19(a)(1). Orders for copies of patent application publications are handled by the USPTO's Office of Public Records. The Office of Public Records can be reached by telephone at (703) 308-9726 or (800) 972-6382, by facsimile at (703) 305-8759, by mail addressed to the United States Patent and Trademark Office, Office of Public Records, Alexandria, VA 22313-1450 or via the Internet.

In addition, information on the status of the application, including the mailing date of Office actions and the dates of receipt of correspondence filed in the Office, may also be accessed via the Internet through the Patent Electronic Business Center at www.uspto.gov using the public side of the Patent Application Information and Retrieval (PAIR) system. The direct link to access this status information is currently http://pair.uspto.gov/. Prior to publication, such status information is confidential and may only be obtained by applicant using the private side of PAIR.

Further assistance in electronically accessing the publication, or about PAIR, is available by calling the Patent Electronic Business Center at 1-866-217-9197.

Office of Data Managment, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 407

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	9137205	B2	2015-09-15	Rogers et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven ROGERS
	Art Unit	2497
	Examiner Name	Huang, Cheng-Feng
	Attorney Docket Number	007742.00029

	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

☐ That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

☐ See attached certification statement.

☐ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/William E. Wooten/	Date (YYYY-MM-DD)	2015-10-23
Name/Print	William E. Wooten	Registration Number	60049

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	23876287
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	23-OCT-2015
Filing Date:	16-APR-2014
Time Stamp:	14:51:39
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Information Disclosure Statement (IDS) Form (SB08)	IDS.pdf	1034940 458e53a2eea870ab6bfb742654341169b567c7f7	no	4

Warnings:	
Information:	

Total Files Size (in bytes):	1034940

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re the Application of:		Atty. Docket No.:	007742.00029
Steven ROGERS *et al.*			
Serial No.:	14/253,992	Group Art Unit:	2497
Filed:	April 16, 2014	Examiner:	Huang, Cheng-Feng
For:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK	Confirmation No.:	5163

AMENDMENT AND RESPONSE TO RESTRICTION REQUIREMENT

Mail Stop Amendment
Commissioner for Patents
Randolph Building
401 Dulany Street
Alexandria, VA 22314

Dear Sir:

In response to the Office Action dated October 15, 2015, please enter the following:

Amendments to the Claims reflected in the Listing of Claims, which begins at page 2 of this paper; and

Remarks, which begin at page 8 of this paper.

Applicant hereby petitions for any applicable extension of time that may be required. If any fees are required or if an overpayment is made, the Commissioner is authorized to debit or credit Deposit Account No. 19-0733, accordingly.

This Listing of Claims will replace all prior versions and listings of claims in this application.

Listing of Claims:

1. (Original) A method, comprising:

receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy that includes at least one rule specifying application-layer packet-header information and a packet transformation function to be performed on packets comprising the application-layer packet-header information;

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway;

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets comprising the application-layer packet-header information; and

performing, by the packet security gateway and on a packet-by-packet basis, the packet transformation function on each of the one or more packets comprising the application-layer packet-header information.

2. (Original) The method of claim 1, wherein the at least one rule indicates that packets having the application-layer packet-header information should have an accept packet transformation function performed on them, and wherein performing the packet transformation function on each of the one or more packets comprising the application-layer packet-header information comprises the packet security gateway forwarding each of the one or more packets comprising the application-layer packet-header information toward its respective destination.

3. (Original) The method of claim 1, wherein the at least one rule indicates that packets having the application-layer packet-header information should have a deny packet

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 415

transformation function performed on them, and wherein performing the packet transformation function on each of the one or more packets comprising the application-layer packet-header information comprises the packet security gateway dropping each of the one or more packets comprising the application-layer packet-header information.

4. (Original) The method of claim 1, wherein the application-layer packet-header information identifies one or more hypertext transfer protocol (HTTP) packets, and wherein identifying the one or more packets comprising the application-layer packet-header information comprises determining by the packet security gateway that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets.

5. (Original) The method of claim 4, wherein the one or more HTTP packets comprise an HTTP GET method call, and wherein determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets comprises determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP GET method call.

6. (Original) The method of claim 5, wherein the HTTP GET method call specifies a uniform resource identifier (URI), and wherein determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP GET method call comprises determining that the one or more packets comprising the application-layer packet-header information originated from or are destined for a network address corresponding to the URI.

7. (Original) The method of claim 4, wherein the one or more HTTP packets comprise an HTTP PUT method call, and wherein determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 416

packets comprises determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP PUT method call.

8. (Original) The method of claim 7, wherein the HTTP PUT method call specifies a uniform resource identifier (URI), and wherein determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP PUT method call comprises determining that the one or more packets comprising the application-layer packet-header information originated from or are destined for a network address corresponding to the URI.

9. (Original) The method of claim 1, wherein the at least one rule comprises a five-tuple specifying one or more transport-layer protocols, a range of source addresses, a range of source ports, a range of destination addresses, and a range of destination ports, and wherein identifying the one or more packets comprising the application-layer packet-header information comprises determining by the packet security gateway that each of the one or more packets comprising the application-layer packet-header information corresponds to at least one of the one or more transport-layer protocols, has a source address within the range of source addresses, a source port within the range of source ports, a destination address within the range of destination addresses, and a destination port within the range of destination ports.

10. (Original) A method, comprising:

receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy that includes at least one rule specifying packet-identification criteria and a packet transformation function comprising a packet digest logging function to be performed on packets corresponding to the packet-identification criteria;

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 417

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets corresponding to the packet-identification criteria; and

performing, by the packet security gateway and on a packet-by-packet basis, the packet digest logging function on each of the one or more packets corresponding to the packet-identification criteria.

11. (Original) The method of claim 10, wherein performing the packet digest logging function on each of the one or more packets corresponding to the packet-identification criteria comprises:

identifying a subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria; and

generating, for each of the one or more packets corresponding to the packet-identification criteria, a record comprising the subset of information specified by the packet digest logging function.

12. (Original) The method of claim 11, wherein the subset of information comprises at least one of a portion of data from the packet, a source address of the packet, a source port of the packet, a destination address of the packet, a destination port of the packet, a transport-protocol type of the packet, a uniform resource identifier (URI) from the packet, an arrival time of the packet, a size of the packet, a flow direction of the packet, an identifier of an interface of the packet security gateway that received the packet, or one or more media access control (MAC) addresses associated with the packet.

13. (Original) The method of claim 11, comprising:

temporarily storing data in a memory of the packet security gateway, the data comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria; and

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 418

utilizing, by the packet security gateway, the data to generate a message comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria.

14. (Original) The method of claim 13 comprising communicating, by the packet security gateway and to the security policy management server, the message comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria.

15. (Original) The method of claim 11, comprising reformatting, for each of the one or more packets corresponding to the packet-identification criteria, the subset of information specified by the packet digest logging function in accordance with a logging system standard.

16. (Original) The method of claim 15, wherein reformatting the subset of information specified by the packet digest logging function in accordance with the logging system standard comprises reformatting the subset of information specified by the packet digest logging function in accordance with a syslog logging system standard.

17. (Original) The method of claim 10, wherein the packet-identification criteria comprises a five-tuple specifying one or more transport-layer protocols, a range of source addresses, a range of source ports, a range of destination addresses, and a range of destination ports, and wherein identifying the one or more packets corresponding to the packet-identification criteria comprises determining by the packet security gateway that each of the one or more packets corresponding to the packet-identification criteria corresponds to at least one of the one or more transport-layer protocols, has a source address within the range of source addresses, a source port within the range of source ports, a destination address within the range of destination addresses, and a destination port within the range of destination ports.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 419

18. (Original) The method of claim 10, wherein the packet-identification criteria comprises application-layer packet-header information, and wherein identifying the one or more packets corresponding to the packet-identification criteria comprises determining by the packet security gateway that each of the one or more packets corresponding to the packet-identification criteria comprises at least a portion of the application-layer packet-header information.

19-20. (Canceled)

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 420

REMARKS

The Office Action dated October 15, 2015 ("Office Action"), has been reviewed, and these remarks are responsive thereto. By this amendment, claims 19 and 20 are canceled. Upon entry of this amendment, claims 1-18 will be pending in this application. Reconsideration and allowance of the instant application are respectfully requested.

Election/Restrictions

The Office Action indicates the instant application contains claims directed to two distinct inventions, grouped as:

I. **Claims 1-18**, "drawn to managing network security proactively by applying dynamic network security policies based on packet-identification criteria and packet transformation function," classified in CPC H04L 63/20; and

II. **Claims 19 and 20**, "drawn to updating network security rules by identifying and determining different set[s] of network addresses," classified in CPC H04L 63/0263.

Office Action, p. 2.

Applicant elects, ***without traverse***, to proceed with examination of ***Group I***.

[CONCLUSION AND SIGNATURE FOLLOW ON SEPARATE PAGE.]

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 421

CONCLUSION

If any fees are required or an overpayment is made, the Commissioner is authorized to debit or credit Deposit Account No. 19-0733, accordingly.

Applicant respectfully submits the instant application is in condition for allowance and respectfully solicits prompt notification of the same.

Respectfully submitted,

BANNER & WITCOFF, LTD.

Dated: December 15, 2015

By: /William E. Wooten/
William E. Wooten
Registration No. 60,049

Banner & Witcoff, Ltd.
1100 13th Street NW, Suite 1200
Washington, D.C. 20005-4051
Telephone: (202) 824-3000
Facsimile: (202) 824-3001

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 422

Electronic Acknowledgement Receipt

EFS ID:	24374120
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	William Eugene Wooten
Filer Authorized By:	
Attorney Docket Number:	007742.00029
Receipt Date:	15-DEC-2015
Filing Date:	16-APR-2014
Time Stamp:	23:36:14
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Response to Election / Restriction Filed	Amendment.pdf	186479 8b594840593eab4bdf21dc9290e186633866bb2d	no	9

Warnings:	
Information:	

Total Files Size (in bytes):	186479

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 14/253,992	Filing Date 04/16/2014	☐ To be Mailed

ENTITY: ☐ LARGE ☒ SMALL ☐ MICRO

APPLICATION AS FILED – PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	

APPLICATION AS AMENDED – PART II

AMENDMENT **12/15/2015**

	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
Total (37 CFR 1.16(i))	* 18	Minus	** 20	= 0	X $40 =	0
Independent (37 CFR 1.16(h))	* 3	Minus	***3	= 0	X $210 =	0
☐ Application Size Fee (37 CFR 1.16(s))						
☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
					TOTAL ADD'L FEE	0

AMENDMENT

	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =	
Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =	
☐ Application Size Fee (37 CFR 1.16(s))						
☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
					TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

LIE
/JAMES TUNSTALL/

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**
If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PLUS Search Results for S/N 14253992, Searched Fri Feb 12 14:01:49 EST 2016

The Patent Linguistics Utility System (PLUS) is a USPTO automated search system for U.S. Patents from 1971 to the present PLUS is a query-by-example search system which produces a list of patents that are most closely related linguistically to the application searched. This search was prepared by the staff of the Scientific and Technical Information Center, SIRA.

6137792 99
6084858 97
6055236 98
5706348 98
5710760 98
5805705 98
5987520 98
5991401 98
6011910 98
6026085 98
6147987 98
6233232 98
4774658 97
4799153 97
4901348 97
4965804 97
5247513 97
5251203 97
5280529 97
5442708 97
5444782 97
5485455 97
5490252 97
5491694 97
5521910 97
5615340 97
5640456 97
5642421 97
5657452 97
5684800 97
5729594 97
5737318 97
5751952 97
5768525 97
5781550 97
5790546 97
5796942 97
5799285 97
5828846 97
5864683 97
5896499 97
5913037 97
5935249 97
5946308 97
5950172 97
5953337 97
6009528 97
6035405 97
6070243 97
6081524 97

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
14/253,992	04/16/2014	Steven Rogers	007742.00029	5163

22907 7590 02/26/2016

BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051

EXAMINER	
HUANG, CHENG-FENG	

ART UNIT	PAPER NUMBER
2497	

NOTIFICATION DATE	DELIVERY MODE
02/26/2016	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

eofficeaction@bannerwitcoff.com

Office Action Summary	Application No. 14/253,992	Applicant(s) ROGERS ET AL.	
	Examiner CHENG-FENG HUANG	Art Unit 2497	AIA (First Inventor to File) Status Yes

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTHS FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *12/15/2015*.

☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

4)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims*

5)☒ Claim(s) *1-18* is/are pending in the application.

5a) Of the above claim(s) _____ is/are withdrawn from consideration.

6)☐ Claim(s) _____ is/are allowed.

7)☒ Claim(s) *1-18* is/are rejected.

8)☐ Claim(s) _____ is/are objected to.

9)☐ Claim(s) _____ are subject to restriction and/or election requirement.

* If any claims have been determined allowable, you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to PPHfeedback@uspto.gov.

Application Papers

10)☐ The specification is objected to by the Examiner.

11)☒ The drawing(s) filed on *04/16/2014* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

a)☐ All b)☐ Some** c)☐ None of the:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

** See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)

2) ☒ Information Disclosure Statement(s) (PTO/SB/08a and/or PTO/SB/08b) Paper No(s)/Mail Date *10/23/2015.*.

3) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .

4) ☐ Other: _____.

DETAILED ACTION

1. The present application, filed on or after March 16, 2013, is being examined under the first inventor to file provisions of the AIA.

2. This is a reply to the application filed on 12/15/2015, in which, claim(s) 1-18 are pending. Claim(s) 1 and 10 are independent.

When making claim amendments, the applicant is encouraged to consider the references in their entireties, including those portions that have not been cited by the examiner and their equivalents as they may most broadly and appropriately apply to any particular anticipated claim amendments.

Response to Election

3. Applicant's election without traverse of Claims 1-18 in the reply filed on 12/15/2015 is acknowledged. Claims 1-18 will be examined on the merits in this Non-Final Office Action.

Information Disclosure Statement

4. The information disclosure statements (IDS) submitted on 04/16/2014, 06/25/2014, 06/27/2014, 07/03/2014, 09/09/2014, 10/27/2014, 11/25/2014, 02/18/2015, 03/13/2015, 04/20/2015, 06/19/2015, 07/21/2015, 07/24/2015, 09/30/2015, 10/23/2015, have been reviewed. The submission is in compliance with the provisions of 37 CFR 1.97. Accordingly, the examiner is considering the information disclosure statement.

Drawings

5. The drawings filed on 04/16/2014 is/are accepted by The Examiner.

Double Patenting

6. The nonstatutory double patenting rejection is based on a judicially created doctrine grounded in public policy (a policy reflected in the statute) so as to prevent the unjustified or improper time-wise extension of the "right to exclude" granted by a patent and to prevent possible harassment by multiple assignees. A nonstatutory double patenting rejection is appropriate where the claims at issue are not identical, but at least one examined application claim is not patentably distinct from the reference claim(s) because the examined application claim is either anticipated by, or would have been obvious over, the reference claim(s). See, e.g., *In re Berg*, 140 F.3d 1428, 46 USPQ2d 1226 (Fed. Cir. 1998); *In re Goodman*, 11 F.3d 1046, 29 USPQ2d 2010 (Fed. Cir. 1993); *In re Longi*, 759 F.2d 887, 225 USPQ 645 (Fed. Cir. 1985); *In re Van Ornum*, 686 F.2d 937, 214 USPQ 761 (CCPA 1982); *In re Vogel*, 422 F.2d 438, 164 USPQ 619 (CCPA 1970); and *In re Thorington*, 418 F.2d 528, 163 USPQ 644 (CCPA 1969).

A timely filed terminal disclaimer in compliance with 37 CFR 1.321(c) or 1.321(d) may be used to overcome an actual or provisional rejection based on a nonstatutory double patenting ground provided the reference application or patent either is shown to be commonly owned with this application, or claims an invention made as a result of

activities undertaken within the scope of a joint research agreement. A terminal disclaimer must be signed in compliance with 37 CFR 1.321(b).

The USPTO internet Web site contains terminal disclaimer forms which may be used. Please visit http://www.uspto.gov/forms/. The filing date of the application will determine what form should be used. A web-based eTerminal Disclaimer may be filled out completely online using web-screens. An eTerminal Disclaimer that meets all requirements is auto-processed and approved immediately upon submission. For more information about eTerminal Disclaimers, refer to http://www.uspto.gov/patents/process/file/efs/guidance/eTD-info-I.jsp.

7. Claims 1-18 are non-provisionally rejected on the ground of nonstatutory obviousness-type double patenting as being unpatentable over:

Claims 1 and 10 of Patent 9,137,205.

Patent application No. US 9,137,205 (13/657,010)	Current Application No.(14/253992)
Claim 1. A method, comprising: at each packet security gateway of one or more packet security gateways associated with a security policy management server: receiving a plurality of dynamic security policies from the security policy management server, wherein receiving the plurality of dynamic security policies comprises: receiving at least one rule specifying a set of network addresses for which associated packets should be forwarded and at least one rule specifying that all packets associated	Claim 1. A method, comprising: receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy that includes at least one rule specifying application-layer packet-header information and a packet transformation function to be performed on packets comprising the application-layer

with network addresses outside the set of network addresses for which packets should be forwarded should be dropped; receiving, at a first time, a dynamic security policy specifying a first set of network addresses for which packets should be forwarded; receiving, at a second time, a dynamic security policy specifying a second set of network addresses for which packets should be forwarded; and receiving, at a third time, a dynamic security policy specifying a third set of network addresses for which packets should be forwarded, the second time being after the first time, the third time being after the second time, the second set of network addresses including more network addresses than the first set of network addresses, and the third set of network addresses including more network addresses than the second set of network addresses;	packet-header information;
receiving packets associated with a network protected by the packet security gateway; and	receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway;
performing, on a packet by packet basis, at least one of multiple packet transformation functions specified by the plurality of dynamic security policies on the packets associated with the network protected by the packet security gateway, wherein performing the at least one of the multiple packet transformation functions specified by the plurality of dynamic security policies on the packets comprises performing at least one packet transformation function other than forwarding or dropping the packets.	identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets comprising the application-layer packet-header information; and performing, by the packet security gateway and on a packet-by-packet basis, the packet transformation function on each of the one or more packets comprising the application-layer packet-header information.

Claim 10. The method of claim 1, wherein receiving the plurality of dynamic security policies comprises receiving at least one rule comprising a five-tuple specifying one or more protocols, one or more Internet Protocol (IP) source addresses, one or more source port values, one or more IP destination addresses, and one or more destination port values.	Claim 9. The method of claim 1, wherein the at least one rule comprises a five-tuple specifying one or more transport-layer protocols, a range of source addresses, a range of source ports, a range of destination addresses, and a range of destination ports, and wherein identifying the one or more packets comprising the application-layer packet-header information comprises determining by the packet security gateway that each of the one or more packets comprising the application-layer packet-header information corresponds to at least one of the one or more transport-layer protocols, has a source address within the range of source addresses, a source port within the range of source ports, a destination address within the range of destination addresses, and a destination port within the range of destination ports.

Claim Rejections - 35 USC § 103

8. The following is a quotation of 35 U.S.C. 103 which forms the basis for all obviousness rejections set forth in this Office action:

> A patent for a claimed invention may not be obtained, notwithstanding that the claimed invention is not identically disclosed as set forth in section 102 of this title, if the differences between the claimed invention and the prior art are such that the claimed invention as a whole would have been obvious before the effective filing date of the claimed invention to a person having ordinary skill in the art to which the claimed invention pertains. Patentability shall not be negated by the manner in which the invention was made.

In the event the determination of the status of the application as subject to AIA 35 U.S.C. 102 and 103 (or as subject to pre-AIA 35 U.S.C. 102 and 103) is incorrect, any correction of the statutory basis for the rejection will not be considered a new ground of rejection if the prior art relied upon, and the rationale supporting the rejection, would be the same under either status.

9. Claims 1-6, 9-10, and 17-18 are rejected under 35 U.S.C. 103 as being unpatentable over Krishna Narayanaswamy (US 2009/0328219 A1, cited by the applicant in the 03/13/2015 IDS) in view of Kapoor et al. (US 2008/0229415 A1).

10. Regarding Claim 1, Narayanaswamy discloses ***A method, comprising***:

receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy ([0105], "a plurality of IDP devices 110A-110N are configured to perform the dynamic (security) policy provisioning techniques", [0108], "Policy distribution module 126 may distribute policies 124 to each of IDP devices 110 for application to incoming network traffic") ***that includes at least one rule*** ([0098], "a plurality of policies 32 (or rules)") ***and a packet transformation function to be performed on packets*** ([0102], "process the packet,…in accordance with the one of policies 32");

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway ([0105], "a plurality of IDP devices 110A-110N are configured to perform the dynamic (security) policy provisioning techniques", [0108], "Policy distribution module 126 may distribute policies 124 to each of IDP devices 110 for application to incoming network traffic");

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets ([0102], "Servicing engine 28 may

apply the one of policies 32 by applying a set of patterns 39 identified by the associated one of policies 32 to determine whether one of the set of patterns 39 match the pattern presented by the packet and subsequent transaction (on a packet-by-packet basis) ”); and

performing, by the packet security gateway and on a packet-by-packet basis, the packet transformation function on each of the one or more packets ([0102], “process the packet,…in accordance with the one of policies 32”, “Servicing engine 28 may apply the one of policies 32 by applying a set of patterns 39 identified by the associated one of policies 32 to determine whether one of the set of patterns 39 match the pattern presented by the packet and subsequent transaction”).

Although Narayanaswamy teaches packets and policy includes at least one rule, Narayanaswamy does not explicitly teach the rule specifying ***application-layer packet-header information,*** and the packets comprising ***the application-layer packet-header information.***

Kapoor teaches ***application-layer packet-header information*** ([0525], “an application layer packet header”).

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to receive a dynamic security policy that includes at least one rule (as disclosed by Narayanaswamy) specifying application-layer packet-header information (as taught by Kapoor) and a packet transformation function to be performed on packets (as disclosed by

Narayanaswamy) comprising the application-layer packet-header information (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

11. Regarding Claim 2, Narayanaswamy discloses ***wherein the at least one rule indicates that packets should have an accept packet transformation function performed on them, and wherein performing the packet transformation function on each of the one or more packets comprises the packet security gateway forwarding each of the one or more packets toward its respective destination*** ([0102], "Servicing engine 28 may apply the one of policies 32 by applying a set of patterns 39 identified by the associated one of policies 32 to determine whether one of the set of patterns 39 match the pattern presented by the packet and subsequent transaction", [0103],"If not (considered as accept packet transformation function), servicing engine 28 may forward the traffic as secure outgoing traffic 44A (to its respective destination)").

Narayanaswamy does not explicitly teach ***packets having (comprising) the application-layer packet-header information***,

Kapoor teaches ***application-layer packet-header information*** ([0525], "an application layer packet header").

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to perform accept packet

transformation function and forward the packets (as disclosed by Narayanaswamy) which comprising application-layer packet-header information (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

12. Regarding Claim 3, Narayanaswamy discloses ***wherein the at least one rule indicates that packets should have a deny packet transformation function performed on them, and wherein performing the packet transformation function on each of the one or more packets comprises the packet security gateway dropping each of the one or more packets*** ([0102], "Servicing engine 28 may apply the one of policies 32 by applying a set of patterns 39 identified by the associated one of policies 32 to determine whether one of the set of patterns 39 match the pattern presented by the packet and subsequent transaction", [0103], "If such a match occurs (considered as deny packet transformation function), servicing engine 28 may quarantine and/or drop the packet (104)")

Narayanaswamy does not explicitly teach ***packets having (comprising) the application-layer packet-header information***,

Kapoor teaches ***application-layer packet-header information*** ([0525], "an application layer packet header").

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to perform deny packet

transformation function and drop the packets (as disclosed by Narayanaswamy) which comprising application-layer packet-header information (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

13. Regarding Claim 4, Narayanaswamy discloses ***the packet security gateway*** ([0105], "IDP devices"),

Narayanaswamy does not explicitly teach ***wherein the application-layer packet-header information identifies one or more hypertext transfer protocol (HTTP) packets, and wherein identifying the one or more packets comprising the application-layer packet-header information comprises determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets***.

Kapoor teaches ***wherein the application-layer packet-header information*** ([0461], "application layer packet header") ***identifies one or more hypertext transfer protocol (HTTP) packets*** ([0458], "HTTP GET"), ***and wherein identifying the one or more packets comprising the application-layer packet-header information*** ([0525], "inspecting (identifying) an application layer packet header") ***comprises determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets*** ([0458], "HTTP GET").

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in

the art before the effective filing date of the claimed invention to identify the packets comprising the application-layer packet-header information and determine that the packets are HTTP packets (as taught by Kapoor) by the packet security gateway (as taught by Narayanaswamy). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

14. Regarding Claim 5, Narayanaswamy does not explicitly teach ***wherein the one or more HTTP packets comprise an HTTP GET method call, and wherein determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets comprises determining that the one or more packets comprising the application-layer packet- header information are associated with the HTTP GET method call***.

Kapoor teaches ***wherein the one or more HTTP packets comprise an HTTP GET method call*** ([0458], "HTTP GET"), ***and wherein determining that the one or more packets comprising the application-layer packet-header information*** ([0461], "application layer packet header") ***are amongst the one or more HTTP packets comprises determining that the one or more packets comprising the application-layer packet- header information*** ([0525], "inspecting (determining) an application layer packet header") ***are associated with the HTTP GET method call*** ([0458], "HTTP GET").

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in

the art before the effective filing date of the claimed invention to determine by the packet security gateway (as taught by Narayanaswamy) and identify the packets comprising the application-layer packet-header information are associated with the HTTP GET method call (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

15. Regarding Claim 6, Narayanaswamy does not explicitly teach ***wherein the HTTP GET method call specifies a uniform resource identifier (URI), and wherein determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP GET method call comprises determining that the one or more packets comprising the application-layer packet- header information originated from or are destined for a network address corresponding to the URI.***

Kapoor teaches ***wherein the HTTP GET method call*** ([0458], "HTTP GET") ***specifies a uniform resource identifier (URI)*** ([0458], "recognize patterns that suggest the presence of a URL (URI) or a request to access to a URL (URI)", "the request is embodied as an HTTP GET"), ***and wherein determining that the one or more packets comprising the application-layer packet-header information*** ([0461], "application layer packet header") ***are associated with the HTTP GET method call*** ([0458], "HTTP GET") ***comprises determining that the one or more packets comprising the application-layer packet-header information*** ([0525], "inspecting (determining) an application layer packet header") ***originated from or are destined for***

a network address corresponding to the URI ([0458], "recognize patterns that suggest the presence of a URL (URI) or a request to access to a URL (URI)", "the request is embodied as an HTTP GET").

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to determine by the packet security gateway (as taught by Narayanaswamy) and identify the packets comprising the application-layer packet-header information are associated with the HTTP GET method call which specifies a uniform resource identifier (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

16. Regarding Claim 9, Narayanaswamy discloses, ***wherein the at least one rule*** ([0098], "a plurality of policies 32 (or rules)") ***comprises a five-tuple specifying one or more transport-layer protocols, a range of source addresses, a range of source ports, a range of destination addresses, and a range of destination ports*** ([0101], "a five-tuple extracted from each fast path packet", [0048], "A five-tuple comprises a source Internet Protocol (IP) address, a destination IP address, a source port, a destination port, and a protocol") ***and wherein identifying the one or more packets comprises determining by the packet security gateway*** ([0105], "IDP devices") ***that each of the one or more packets corresponds to at least one of the one or more transport-layer protocols*** ([0050], "Real-time Transportation Protocol")***, has a source***

address within the range of source addresses, a source port within the range of source ports, a destination address within the range of destination addresses, and a destination port within the range of destination ports ([0101], "a five-tuple extracted from each fast path packet", [0048], "a source Internet Protocol (IP) address, a destination IP address, a source port, a destination port").

Narayanaswamy does not explicitly teach ***packets comprising the application-layer packet-header information***.

Kapoor teaches ***application-layer packet-header information*** ([0525], "an application layer packet header").

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to have at least one rule comprising a five-tuple for identifying the packets (as taught by Narayanaswamy) which comprises the application-layer packet-header information (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

17. Regarding Claim 10, Narayanaswamy discloses ***A method, comprising***:

receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy ([0105], "a plurality of IDP devices 110A-110N are configured to perform the dynamic (security) policy provisioning

techniques", [0108], "Policy distribution module 126 may distribute policies 124 to each of IDP devices 110 for application to incoming network traffic") ***that includes at least one rule*** ([0098], "a plurality of policies 32 (or rules)") ***and a packet transformation function comprising a packet digest logging function*** ([0091], "five-tuple extracted from each fast path packet", [0048], "extract the five-tuple from the header of each of the packets", as the packet digest logging function) ***to be performed on packets***;

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway ([0105], "a plurality of IDP devices 110A-110N are configured to perform the dynamic (security) policy provisioning techniques", [0108], "Policy distribution module 126 may distribute policies 124 to each of IDP devices 110 for application to incoming network traffic");

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets ([0102], "Servicing engine 28 may apply the one of policies 32 by applying a set of patterns 39 identified by the associated one of policies 32 to determine whether one of the set of patterns 39 match the pattern presented by the packet and subsequent transaction (on a packet-by-packet basis) ");

and

performing, by the packet security gateway and on a packet-by-packet basis, the packet digest logging function ([0091], "five-tuple extracted from each fast path packet", [0048], "extract the five-tuple from the header of each of the packets", as

the packet digest logging function) ***on each of the one or more packets*** ([0102], "process the packet,…in accordance with the one of policies 32", "Servicing engine 28 may apply the one of policies 32 by applying a set of patterns 39 identified by the associated one of policies 32 to determine whether one of the set of patterns 39 match the pattern presented by the packet and subsequent transaction").

Although Narayanaswamy teaches packets and policy includes at least one rule, Narayanaswamy does not explicitly teach the rule specifying ***packet-identification criteria,*** and the packets corresponding to ***the packet-identification criteria***.

Kapoor teaches ***packet-identification criteria*** ([0525], "an application layer packet header", considered as packet-identification criteria).

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to receive a dynamic security policy that includes at least one rule (as disclosed by Narayanaswamy) specifying packet-identification criteria (as taught by Kapoor) and a packet transformation function comprising a packet digest logging function to be performed on packets (as disclosed by Narayanaswamy) comprising the packet-identification criteria (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

18. Regarding Claim 17, Narayanaswamy discloses ***wherein a five-tuple specifying one or more transport-layer protocols, a range of source addresses, a***

range of source ports, a range of destination addresses, and a range of destination ports ([0101], "a five-tuple extracted from each fast path packet", [0048], "A five-tuple comprises a source Internet Protocol (IP) address, a destination IP address, a source port, a destination port, and a protocol"), ***and wherein identifying the one or more packets comprises determining by the packet security gateway that each of the one or more packets corresponds to at least one of the one or more transport-layer protocols*** ([0050], "Real-time Transportation Protocol")***, has a source address within the range of source addresses, a source port within the range of source ports, a destination address within the range of destination addresses, and a destination port within the range of destination ports*** ([0101], "a five-tuple extracted from each fast path packet", [0048], "a source Internet Protocol (IP) address, a destination IP address, a source port, a destination port").

Narayanaswamy does not explicitly teach ***the packet-identification criteria, packets corresponding to the packet-identification criteria***.

Kapoor teaches ***the packet-identification criteria*** ([0525], "an application layer packet header", considered as packet-identification criteria).

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to have the packet-identification criteria (as taught by Kapoor) comprising a five-tuple for identifying the packets (as taught by Narayanaswamy) corresponding to the packet-identification criteria (as taught by Kapoor). The motivation/suggestion would have been to address

patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

19. Regarding Claim 18, Narayanaswamy discloses ***the packet security gateway*** ([0105], "IDP devices"),

Narayanaswamy does not explicitly teach ***wherein the packet-identification criteria comprises application-layer packet-header information, and wherein identifying the one or more packets corresponding to the packet-identification criteria comprises determining that each of the one or more packets corresponding to the packet-identification criteria comprises at least a portion of the application-layer packet-header information.***

Kapoor teaches ***wherein the packet-identification criteria comprises application-layer packet-header information*** ([0461], "application layer packet header"), ***and wherein identifying the one or more packets corresponding to the packet-identification criteria comprises determining that each of the one or more packets corresponding to the packet-identification criteria*** ([0525], "inspecting (determining) an application layer packet header") ***comprises at least a portion of the application-layer packet-header information*** ([0525], "an application layer packet header").

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to identify the packets corresponding to the packet-identification criteria and determine that the criteria

comprises the application-layer packet-header information (as taught by Kapoor) by the packet security gateway (as taught by Narayanaswamy). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

20. Claims 7-8 and 11-14 are rejected under 35 U.S.C. 103 as being unpatentable over Krishna Narayanaswamy (US 2009/0328219 A1, cited by the applicant in the 03/13/2015 IDS) in view of Kapoor et al. (US 2008/0229415 A1) further in view of Erb et al. (US 2014/0215574 A1).

21. Regarding Claim 7, Narayanaswamy does not explicitly teach ***wherein the one or more HTTP packets comprise an HTTP PUT method call, and wherein determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets comprises determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP PUT method call***.

Kapoor teaches ***wherein the one or more HTTP packets*** ([0458], "HTTP" packets), ***and wherein determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets*** ([0458], "HTTP" packets) ***comprises determining that the one or more packets comprising the application-layer packet-header information*** ([0461], "application layer packet header").

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to determine by the packet security gateway (as taught by Narayanaswamy) and identify the packets comprising the application-layer packet-header information are associated with the HTTP packets (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

The combined teaching of Narayanaswamy and Kapoor does not explicitly teach ***packets comprise an HTTP PUT method call, packets associated with the HTTP PUT method call***.

Erb teaches ***an HTTP PUT method call*** ([0053], “The request can include a HTTP PUT”).

Erb, Narayanaswamy and Kapoor are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to determine by the packet security gateway and identify the packets comprising the application-layer packet-header information are associated with the HTTP packets (as taught by the combined teaching of Narayanaswamy and Kapoor) which comprise and are associated with an HTTP PUT method call (as taught by Erb). The motivation/suggestion would have been to track and protect against denial of service (DOS) attacks (Erb, [0027]).

22. Regarding Claim 8, Narayanaswamy does not explicitly teach ***wherein the HTTP PUT method call specifies a uniform resource identifier (URI), and wherein determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP PUT method call comprises determining that the one or more packets comprising the application-layer packet- header information originated from or are destined for a network address corresponding to the URI.***

Kapoor teaches ***wherein determining that the one or more packets comprising the application-layer packet-header information*** ([0461], "application layer packet header") ***comprises determining that the one or more packets comprising the application-layer packet-header information originated from or are destined for a network address corresponding to the URI*** ([0458], "recognize patterns that suggest the presence of a URL (URI) or a request to access to a URL (URI)").

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to determine by the packet security gateway (as taught by Narayanaswamy) and identify the packets comprising the application-layer packet-header information originated from or are destined for a network address corresponding to the URI (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of

types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

The combined teaching of Narayanaswamy and Kapoor does not explicitly teach ***the HTTP PUT method call, packets associated with the HTTP PUT method call***.

Erb teaches ***an HTTP PUT method call*** ([0053], "The request can include a HTTP PUT").

Erb, Narayanaswamy and Kapoor are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to determine by the packet security gateway and identify the packets comprising the application-layer packet-header information are associated with the HTTP packets (as taught by the combined teaching of Narayanaswamy and Kapoor) which comprise and are associated with an HTTP PUT method call (as taught by Erb) which specifies a uniform resource identifier (as taught by the combined teaching of Narayanaswamy and Kapoor). The motivation/suggestion would have been to track and protect against denial of service (DOS) attacks (Erb, [0027]).

23. Regarding Claim 11, Narayanaswamy discloses

identifying a subset of information ([0048], "a source Internet Protocol (IP) address" of the packet as a subset of information specified in instant specification paragraph 86) ***specified by the packet digest logging function for each of the one or more packets*** ([0091], "five-tuple extracted from each fast path packet", [0048],

"extract the five-tuple from the header of each of the packets", as the packet digest logging function);

Narayanaswamy does not explicitly teach ***packets corresponding to the packet-identification criteria, generating, for each of the one or more packets, a record***.

Kapoor teaches ***packet-identification criteria*** ([0525], "an application layer packet header", considered as packet-identification criteria).

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to receive a dynamic security policy that includes at least one rule (as disclosed by Narayanaswamy) specifying packet-identification criteria (as taught by Kapoor) and identify a subset of information specified by the packet digest logging function (as disclosed by Narayanaswamy). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

The combined teaching of Narayanaswamy and Kapoor does not explicitly teach ***generating, for each of the one or more packets, a record***.

Erb teaches ***generating, for each of the one or more packets, a record*** ([0105], "be created for the records").

Erb, Narayanaswamy and Kapoor are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill

in the art before the effective filing date of the claimed invention to receive a dynamic security policy that includes at least one rule specifying packet-identification criteria and identify a subset of information specified by the packet digest logging function (as taught by the combined teaching of Narayanaswamy and Kapoor) and generate a record (as taught by Erb) comprising the subset of information specified by the packet digest logging function (as taught by the combined teaching of Narayanaswamy and Kapoor). The motivation/suggestion would have been to track and protect against denial of service (DOS) attacks (Erb, [0027]).

24. Regarding Claim 12, the combined teaching of Narayanaswamy, Kapoor and Erb teaches ***wherein the subset of information comprises at least one of a source address of the packet*** (Narayanaswamy, ([0048], "a source Internet Protocol (IP) address" of the packet as a subset of information specified in instant specification paragraph 86).

25. Regarding Claim 13, Narayanaswamy discloses

temporarily storing data in a memory of the packet security gateway ([0093], "(temporarily) placing the packet in a portion of memory"), ***the data comprising the subset of information*** ([0048], "a source Internet Protocol (IP) address" of the packet as a subset of information specified in instant specification paragraph 86) ***specified by the packet digest logging function for each of the one or more packets*** ([0091], "five-tuple extracted from each fast path packet", [0048], "extract the five-tuple from the header of each of the packets", as the packet digest logging function);

Narayanaswamy does not explicitly teach ***packets corresponding to the packet-identification criteria***, and ***utilizing the data to generate a message***.

Kapoor teaches ***packet-identification criteria*** ([0525], "an application layer packet header", considered as packet-identification criteria).

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to temporarily store data in a memory where the data comprising the subset of information specified by the packet digest logging function for the packets (as disclosed by Narayanaswamy) corresponding to packet-identification criteria (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

The combined teaching of Narayanaswamy and Kapoor does not explicitly teach ***utilizing the data to generate a message***.

Erb teaches ***utilizing the data to generate a message*** ([0105], "generate upload messages containing the (data) records").

Erb, Narayanaswamy and Kapoor are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to temporarily store data in a memory where the data comprising the subset of information specified by the packet digest logging function for the packets (as taught by the combined teaching of

Narayanaswamy and Kapoor) and utilize the data to generate a message (as taught by Erb) comprising the subset of information specified by the packet digest logging function (as taught by the combined teaching of Narayanaswamy and Kapoor). The motivation/suggestion would have been to track and protect against denial of service (DOS) attacks (Erb, [0027]).

26. Regarding Claim 14, Narayanaswamy discloses

the packet security gateway ([0105], "IDP devices"), ***the security policy management server*** ([0108], "Policy distribution module"), ***and***

the subset of information ([0048], "a source Internet Protocol (IP) address" of the packet as a subset of information specified in instant specification paragraph 86) ***specified by the packet digest logging function for each of the one or more packets*** ([0091], "five-tuple extracted from each fast path packet", [0048], "extract the five-tuple from the header of each of the packets", as the packet digest logging function).

Narayanaswamy does not explicitly teach ***packets corresponding to the packet-identification criteria*** and ***communicating the message***.

Kapoor teaches ***packet-identification criteria*** ([0525], "an application layer packet header", considered as packet-identification criteria).

Kapoor and Narayanaswamy are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to temporarily store data in a memory where the data comprising the subset of information specified by the

packet digest logging function for the packets (as disclosed by Narayanaswamy) corresponding to packet-identification criteria (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

The combined teaching of Narayanaswamy and Kapoor does not explicitly teach ***communicating the message***.

Erb teaches ***communicating the message*** ([0105], "send the upload messages").

Erb, Narayanaswamy and Kapoor are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to communicate the message (as taught by Erb) comprising the subset of information specified by the packet digest logging function from the packet security gateway to the security policy management server (as taught by the combined teaching of Narayanaswamy and Kapoor). The motivation/suggestion would have been to track and protect against denial of service (DOS) attacks (Erb, [0027]).

27. Claims 15 and 16 are rejected under 35 U.S.C. 103 as being unpatentable over Krishna Narayanaswamy (US 2009/0328219 A1, cited by the applicant in the 03/13/2015 IDS) in view of Kapoor et al. (US 2008/0229415 A1) further in view of Erb et

al. (US 2014/0215574 A1) and further in view of Nappier et al. (US 2011/0185055 A1, cited by the applicant in the 07/21/2015 IDS).

28. Regarding Claim 15, the combined teaching of Narayanaswamy and Erb teaches ***the subset of information*** ([0048], "a source Internet Protocol (IP) address" of the packet as a subset of information specified in instant specification paragraph 86) ***specified by the packet digest logging function*** ([0091], "five-tuple extracted from each fast path packet", [0048], "extract the five-tuple from the header of each of the packets", as the packet digest logging function)

The combined teaching of Narayanaswamy and Erb does not explicitly teach ***packets corresponding to the packet-identification criteria*** and ***reformatting with a logging system standard.***

Kapoor teaches ***packet-identification criteria*** ([0525], "an application layer packet header", considered as packet-identification criteria).

Kapoor, Narayanaswamy and Erb are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to have the subset of information specified by the packet digest logging function for the packets (as taught by the combined teaching of Narayanaswamy and Erb) corresponding to packet-identification criteria (as taught by Kapoor). The motivation/suggestion would have been to address patterns relevant to a variety of types of network and data integrity threats that relate to computer systems (Kapoor, Abstract).

The combined teaching of Narayanaswamy, Erb and Kapoor does not explicitly teach ***reformatting with a logging system standard***.

Nappier teaches ***reformatting with a logging system standard*** ([0023], "normalize the information contained in the logs into a common format" as with a logging system standard)

Nappier, Narayanaswamy, Erb and Kapoor are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to have the subset of information specified by the packet digest logging function for the packets corresponding to the packet-identification criteria (as taught by the combined teaching of Narayanaswamy, Erb and Kapoor) reformatting with a logging system standard (as taught by Nappier). The motivation/suggestion would have been to learn new or changed relationships between network identities and network addresses, and furthermore, to scanning hardware devices in the network to learn network device addresses that can be used to track and control changes and activity in the network (Nappier, [0002]).

29. Regarding Claim 16, the combined teaching of Narayanaswamy, Erb and Kapoor teaches ***the subset of information*** (Narayanaswamy, [0048], "a source Internet Protocol (IP) address" of the packet as a subset of information specified in instant specification paragraph 86) ***specified by the packet digest logging function*** (Narayanaswamy, [0091], "five-tuple extracted from each fast path packet", [0048],

"extract the five-tuple from the header of each of the packets", as the packet digest logging function)

The combined teaching of Narayanaswamy, Erb and Kapoor does not explicitly teach ***reformatting in accordance with the logging system standard comprises reformatting in accordance with a syslog logging system standard***.

Nappier teaches ***reformatting in accordance with the logging system standard comprises reformatting in accordance with a syslog logging system standard*** ([0023], "normalize the information contained in the logs into a common format", ""normalize the logs into "syslog" messages"" as in accordance with a syslog logging system standard)

Nappier, Narayanaswamy, Erb and Kapoor are analogous art as they are in the same field of endeavor of information security. It would have been obvious to one of ordinary skill in the art before the effective filing date of the claimed invention to have the subset of information specified by the packet digest logging function for the packets corresponding to the packet-identification criteria (as taught by the combined teaching of Narayanaswamy, Erb and Kapoor) reformatting with a logging system standard in accordance with a syslog logging system standard (as taught by Nappier). The motivation/suggestion would have been to learn new or changed relationships between network identities and network addresses, and furthermore, to scanning hardware devices in the network to learn network device addresses that can be used to track and control changes and activity in the network (Nappier, [0002]).

Conclusion

30. Any inquiry concerning this communication or earlier communications from the examiner should be directed to CHENG-FENG HUANG whose telephone number is (571)272-6186. The examiner can normally be reached on 7:30 am ~ 5:00 pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Hadi Armouche can be reached on 571-270-3618. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/CHENG-FENG HUANG/
Examiner, Art Unit 2497

/HADI ARMOUCHE/
Supervisory Patent Examiner, Art Unit 2497

Notice of References Cited	Application/Control No. 14/253,992	Applicant(s)/Patent Under Reexamination ROGERS ET AL.	
	Examiner CHENG-FENG HUANG	Art Unit 2497	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	CPC Classification	US Classification
*	A	US-2008/0229415 A1	09-2008	Kapoor; Harsh	G06F21/55	726/22
*	B	US-2014/0215574 A1	07-2014	Erb; David	H04L63/10	726/4
	C	US-				
	D	US-				
	E	US-				
	F	US-				
	G	US-				
	H	US-				
	I	US-				
	J	US-				
	K	US-				
	L	US-				
	M	US-				

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	CPC Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	9137205	B2	2015-09-15	Rogers et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							☐

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(**Not for submission under 37 CFR 1.99)**

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven ROGERS
Art Unit	2497
Examiner Name	Huang, Cheng-Feng
Attorney Docket Number	007742.00029

	1		☐

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	02/18/2016

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	14253992	ROGERS ET AL.
	Examiner	**Art Unit**
	CHENG-FENG HUANG	2497

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	10/09/2015	02/18/2016							
	1	÷	✓							
	2	÷	✓							
	3	÷	✓							
	4	÷	✓							
	5	÷	✓							
	6	÷	✓							
	7	÷	✓							
	8	÷	✓							
	9	÷	✓							
	10	÷	✓							
	11	÷	✓							
	12	÷	✓							
	13	÷	✓							
	14	÷	✓							
	15	÷	✓							
	16	÷	✓							
	17	÷	✓							
	18	÷	✓							
	19	÷	-							
	20	÷	-							

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S30	12	(("9137205") or ("2006136987") or ("2007240208")).PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2016/02/05 15:02
S31	4	(("20060136987") or ("20070240208")).PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2016/02/05 15:06
S38	2	("20150304354").PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2016/02/12 15:38
S39	1372	((Steven) near2 (Rogers)).INV.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 15:41
S40	256	((Sean) near2 (Moore)).INV.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 15:41
S41	153	((David) near2 (Ahn)).INV.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 15:42
S42	2	((Peter) near2 (Geremia)).INV.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 15:43
S43	15	((Centripetal) near (Networks)).AS.	US-PGPUB;	OR	ON	2016/02/12

			USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB			15:43
S44	1757	S39 S40 S41 S42 S43	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 15:44
S45	2	S44 and ((packet adj transformation adj function) with (((application-$layer) (application near layer)) near ((packet-$header) (packet near header))))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 16:15
S46	66794	(H04L63/20 or H04L67/02 or H04L63/0209).cpc.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 16:18
S47	2	S46 and ((packet adj transformation adj function) with (((application-$layer) (application near layer)) near ((packet-$header) (packet near header))))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 16:19
S48	183	("20120113987" | "20010039624" | "20030142681" | "20100082811" | "20080005795" | "20120314617" | "8806638" | "20110185055" | "20120331543" | "20070240208" | "6317837" | "20060248580" | "7684400" | "20050138204" | "20060212572" | "20100296441" | "7539186" | "20060146879" | "8306994" | "20110270956" | "20130059527" | "20030035370" | "20110088092" | "20090328219" | "8935785" | "20030154297" | "20020049899" | "9137205" | "20060195896" | "20100011434" | "20120264443" | "7818794" | "20140366132" | "6226372" | "US20060195896" | "20040151155" | "20050141537" | "20080072307" | "20100132027" | "20060114899" | "20130061294" | "8726379" | "20110055916" | "20030212900" | "20050117576" | "20030145225" | "20030097590" | "7478429" | "8004994" | "20100242098" | "20060048142" | "20030188192" | "7721084" | "7710885" | "20040088542" | "20080235755" |	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 16:20

		"US20060248580" | "20030123456" | "20070083924" | "20100211678" | "8856926" | "20090172800" | "20120240135" | "6611875" | "20130254766" | "20060136987").PN.				
S49	0	S48 and ((packet adj transformation adj function) with (((application-$layer) (application near layer)) near ((packet-$header) (packet near header))))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 16:20
S50	0	S48 and ((((application-$layer) (application near layer)) near ((packet-$header) (packet near header))))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 16:21
S51	2	((packet adj transformation adj function) with (((application-$layer) (application near layer)) near ((packet-$header) (packet near header))))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 16:22
S52	0	S48 and ((((application-$layer) (application near layer)) near ((packet-$header) (packet near header))))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 16:23
S53	3	S48 and ((((application-$layer) (application near layer)) with ((packet-$header) (packet near header))))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/12 16:24
S54	183	("20120113987" | "20010039624" | "20030142681" | "20100082811" | "20080005795" | "20120314617" | "8806638" | "20110185055" | "20120331543" | "20070240208" | "6317837" | "20060248580" | "7684400" | "20050138204" | "20060212572" | "20100296441" | "7539186" | "20060146879" | "8306994" | "20110270956" | "20130059527" | "20030035370" | "20110088092" | "20090328219" | "8935785" | "20030154297" | "20020049899" | "9137205" | "20060195896" | "20100011434" | "20120264443" | "7818794" | "20140366132" | "6226372" | "US20060195896" | "20040151155" | "20050141537" | "20080072307" | "20100132027" | "20060114899" | "20130061294" | "8726379" |	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 11:32

		"20110055916" | "20030212900" | "20050117576" | "20030145225" | "20030097590" | "7478429" | "8004994" | "20100242098" | "20060048142" | "20030188192" | "7721084" | "7710885" | "20040088542" | "20080235755" | "US20060248580" | "20030123456" | "20070083924" | "20100211678" | "8856926" | "20090172800" | "20120240135" | "6611875" | "20130254766" | "20060136987").PN.				
S55	0	S54 and ((application adj layer) near (packet adj header))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 11:35
S56	39	S54 and ((packet adj header))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 11:35
S57	39	S56 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 11:36
S58	4	S54 and ((packet adj header) with application)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 11:37
S59	4	S58 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 11:37
S60	2	("20090328219").PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2016/02/16 11:47
S61	2	S54 and S60	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 11:47

S62	50	(US-6137792-$ or US-6055236-$ or US-5706348-$ or US-5710760-$ or US-5805705-$ or US-5987520-$ or US-5991401-$ or US-6011910-$ or US-6026085-$ or US-6147987-$ or US-6233232-$ or US-4774658-$ or US-4799153-$ or US-4901348-$ or US-4965804-$ or US-5247513-$ or US-5251203-$ or US-5280529-$ or US-5442708-$ or US-5444782-$ or US-5485455-$ or US-5490252-$ or US-5491694-$ or US-5521910-$ or US-5615340-$ or US-5640456-$).did. or (US-5642421-$ or US-5657452-$ or US-5684800-$ or US-5729594-$ or US-5737318-$ or US-5751952-$ or US-5768525-$ or US-5781550-$ or US-5790546-$ or US-5796942-$ or US-5799285-$ or US-5828846-$ or US-5864683-$ or US-5896499-$ or US-5913037-$ or US-5935249-$ or US-5946308-$ or US-5950172-$ or US-5953337-$ or US-6009528-$ or US-6035405-$ or US-6070243-$ or US-6081524-$ or US-6084858-$).did.	USPAT	OR	ON	2016/02/16 12:01
S63	0	S62 and ((packet adj header) with application)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 12:01
S64	19	S62 and ((packet adj header))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 12:02
S65	19	S64 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 12:02
S66	0	S60 and (log logging)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 12:05
S67	0	S60 and (digest near logging)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 12:06
S68	0	S62 and (digest near logging)	US-PGPUB; USPAT;	OR	ON	2016/02/16 12:06

			USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB			
S69	0	S54 and (digest near logging)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 12:06
S70	2	S60 and memory	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 12:15
S71	57	((((application-$layer) (application near layer)) near ((packet-$header) (packet near header))))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 12:58
S72	1	S71 and syslog	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:00
S73	28	S71 and HTTP	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:00
S74	23	S71 and (denial near service)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:02
S75	21	S74 and (HTTP near GET)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:02
S76	20	S75 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:02

S77	0	S76 and (HTTP near PUT)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:06
S78	2	S54 and (HTTP near GET)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:10
S79	0	S54 and S71	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:12
S80	2	("20080229415").PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2016/02/16 13:12
S81	0	S54 and S80	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:12
S82	1950368	((denial near service) Dos DDos)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:13
S83	26	S71 and S82	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:14
S84	1	S83 and (HTTP near PUT)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:15
S85	21	S83 and (HTTP near GET)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO;	OR	ON	2016/02/16 13:15

			DERWENT; IBM_TDB			
S86	2	("9137205").PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2016/02/16 13:32
S87	0	S71 and S86	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 13:32
S88	2	("20080229415").PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2016/02/16 14:33
S89	0	S54 and S88	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/16 14:33
S90	57	((((application-$layer) (application near layer)) near ((packet-$header) (packet near header))))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 09:18
S91	1950368	((denial near service) Dos DDos)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 09:18
S92	26	S90 and S91	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 09:18
S93	21	S92 and (HTTP near GET)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 09:18
S94	1	S93 and ((generating creat$3 produc$3) with (record log message))	US-PGPUB; USPAT; USOCR;	OR	ON	2016/02/17 15:00

			FPRS; EPO; JPO; DERWENT; IBM_TDB			
S95	32783	S91 and ((generating creat$3 produc$3) with (record log message digest))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:01
S96	19787	S91 and ((generating creat$3 produc$3) near3 (record log message digest))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:04
S97	10090	S91 and ((generating creat$3 produc$3) near (record log message digest))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:04
S98	204	S97 and (HTTP near GET)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:05
S99	191	S98 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:05
S100	12	S99 and (HTTP near PUT)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:06
S101	66961	(H04L63/20 or H04L67/02 or H04L63/0209).cpc.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:14
S102	3382	S101 and ((generating creat$3 produc$3) near (record log message digest))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:15
S103	21	S102 and (HTTP near PUT)	US-PGPUB;	OR	ON	2016/02/17

			USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB			15:15
S104	19	S103 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:15
S105	465	S95 and S101 and ((generating creat$3 produc$3) near (record log message digest))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:17
S106	3	S105 and (HTTP near PUT)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:17
S107	2	S106 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:17
S108	20585	((denial near service) DDos)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:19
S109	6311	S108 and ((generating creat$3 produc$3) with (record log message digest))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:20
S110	2274	S108 and ((generating creat$3 produc$3) near (record log message digest))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/17 15:21
S111	9	S110 and (HTTP near PUT)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT;	OR	ON	2016/02/17 15:21

			IBM_TDB			
S112	1950516	((denial near service) Dos DDos)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 11:06
S113	10094	S112 and ((generating creat$3 produc$3) near (record log message digest))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 11:06
S114	204	S113 and (HTTP near GET)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 11:06
S115	191	S114 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 11:06
S116	12	S115 and (HTTP near PUT)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 11:06
S117	12	S116 and message	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 11:06
S118	1	S116 and syslog	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 13:04
S119	117	S113 and syslog	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 13:05
S120	183	("20120113987" | "20010039624" | "20030142681" | "20100082811" | "20080005795" | "20120314617" | "8806638" | "20110185055" |	US-PGPUB; USPAT; USOCR; FPRS;	OR	ON	2016/02/18 13:21

		"20120331543" | "20070240208" | "6317837" | "20060248580" | "7684400" | "20050138204" | "20060212572" | "20100296441" | "7539186" | "20060146879" | "8306994" | "20110270956" | "20130059527" | "20030035370" | "20110088092" | "20090328219" | "8935785" | "20030154297" | "20020049899" | "9137205" | "20060195896" | "20100011434" | "20120264443" | "7818794" | "20140366132" | "6226372" | "US20060195896" | "20040151155" | "20050141537" | "20080072307" | "20100132027" | "20060114899" | "20130061294" | "8726379" | "20110055916" | "20030212900" | "20050117576" | "20030145225" | "20030097590" | "7478429" | "8004994" | "20100242098" | "20060048142" | "20030188192" | "7721084" | "7710885" | "20040088542" | "20080235755" | "US20060248580" | "20030123456" | "20070083924" | "20100211678" | "8856926" | "20090172800" | "20120240135" | "6611875" | "20130254766" | "20060136987").PN.	EPO; JPO; DERWENT; IBM_TDB			
S123	4	S120 and (five adj tuple)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 13:23
S124	4	S120 and syslog	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 13:24
S125	0	S123 and syslog	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 13:24
S126	2	"13/657010".app.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/18 16:02
S127	0	S126 and ((application adj layer) near (packet adj header))	US-PGPUB; USPAT; USOCR; FPRS;	OR	ON	2016/02/18 16:02

			EPO; JPO; DERWENT; IBM_TDB			
S128	183	("20120113987" | "20010039624" | "20030142681" | "20100082811" | "20080005795" | "20120314617" | "8806638" | "20110185055" | "20120331543" | "20070240208" | "6317837" | "20060248580" | "7684400" | "20050138204" | "20060212572" | "20100296441" | "7539186" | "20060146879" | "8306994" | "20110270956" | "20130059527" | "20030035370" | "20110088092" | "20090328219" | "8935785" | "20030154297" | "20020049899" | "9137205" | "20060195896" | "20100011434" | "20120264443" | "7818794" | "20140366132" | "6226372" | "US20060195896" | "20040151155" | "20050141537" | "20080072307" | "20100132027" | "20060114899" | "20130061294" | "8726379" | "20110055916" | "20030212900" | "20050117576" | "20030145225" | "20030097590" | "7478429" | "8004994" | "20100242098" | "20060048142" | "20030188192" | "7721084" | "7710885" | "20040088542" | "20080235755" | "US20060248580" | "20030123456" | "20070083924" | "20100211678" | "8856926" | "20090172800" | "20120240135" | "6611875" | "20130254766" | "20060136987").PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/22 11:27
S129	22	(packet adj security adj gateway).clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/22 11:28
S130	1	S128 and S129	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/02/22 11:28

EAST Search History (Interference)

< This search history is empty>

2/ 22/ 2016 12:47:45 PM
C:\ Users\ chuang2\ Documents\ EAST\ Workspaces\ 14253992.wsp

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	14253992	ROGERS ET AL.
	Examiner	Art Unit
	CHENG-FENG HUANG	2497

CPC- SEARCHED		
Symbol	Date	Examiner
H04L63/20 or H04L67/02 or H04L63/0209	10/08/2015	C. H.

CPC COMBINATION SETS - SEARCHED		
Symbol	Date	Examiner

US CLASSIFICATION SEARCHED			
Class	Subclass	Date	Examiner

SEARCH NOTES		
Search Notes	Date	Examiner
EAST Searches (including inventor and assignee search)	10/08/2015	C. H.
EAST classification search (keywords combined with CPC)	10/08/2015	C. H.
NPL Searches	10/08/2015	C. H.
PLUS Search Performed	10/08/2015	C. H.
Update EAST Searches	02/18/2016	C. H.
Update NPL Searches	02/18/2016	C. H.

INTERFERENCE SEARCH			
US Class/ CPC Symbol	US Subclass / CPC Group	Date	Examiner

/CHENG-FENG HUANG/ Examiner.Art Unit 2497	

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	6098172		2000-08-01	Coss et al.	
	2	7331061	B1	2008-02-12	Ramsey et al.	
	3	7299353	B2	2007-11-20	Le Pennec et al.	
	4	7227842	B1	2007-06-05	Ji et al.	
	5	7107613	B1	2006-09-12	Chen et al.	
	6	7263099	B1	2007-08-28	Woo et al.	
	7	7237267	B2	2007-06-26	Rayes et al.	
	8	7089581	B1	2006-08-08	Nagai et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	9	6662235	B1	2003-12-09	Callis et al.	
	10	6484261	B1	2002-11-19	Wiegel	
	11	8042167	B2	2011-10-18	Fulp et al.	
	12	8037517	B2	2011-10-11	Fulp et al.	
	13	7954143	B2	2011-05-31	Aaron	
	14	8495725	B2	2013-07-23	Ahn	
	15	9124552	B2	2015-09-01	Moore	
	16	9160713	B2	2015-10-13	Moore	
	17	9094445	B2	2015-07-28	Moore et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	1	20060070122	A1	2006-03-30	Bellovin	
	2	20050010765	A1	2005-01-13	Swander et al.	
	3	20020198981	A1	2002-12-26	Corl et al.	
	4	20090138938	A1	2009-05-28	Harrison et al.	
	5	20060104202	A1	2006-05-18	Reiner	
	6	20050229246	A1	2005-10-13	Rajagopal et al.	
	7	20040177139	A1	2004-09-09	Schuba et al.	
	8	20040010712	A1	2004-01-15	Hui et al.	
	9	20030120622	A1	2003-06-26	Nurmela et al.	
	10	20080301765	A1	2008-12-04	Nicol et al.	
	11	20050251570	A1	2005-11-10	Heasman et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	12	20050183140	A1	2005-08-18	Goddard	
	13	20050125697	A1	2005-06-09	Tahara	
	14	20040193943	A1	2004-09-30	Angelino et al.	
	15	20020165949	A1	2002-11-07	Na et al.	
	16	20020038339	A1	2002-03-28	Xu	
	17	20050114704	A1	2005-05-26	Swander	
	18	20100268799	A1	2010-10-21	Maestas	
	19	20110055923	A1	2011-03-03	Thomas	
	20	20100303240	A1	2010-12-02	Beachem et al.	
	21	20100199346	A1	2010-08-05	Ling et al.	
	22	20100011433	A1	2010-01-14	HARRISON et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	23	20080077705	A1	2008-03-27	Li et al.	
	24	20050286522	A1	2005-12-29	Paddon et al.	
	25	20040205360	A1	2004-10-14	Norton et al.	
	26	20140281030	A1	2014-09-18	Cui et al.	
	27	20040073655	A1	2004-04-15	Kan et al.	
	28	20150237012	A1	2015-08-20	Moore	
	29	20150304354	A1	2015-10-22	Rogers et al.	
	30	20140115654	A1	2014-04-24	Rogers et al.	
	31	20140201123	A1	2014-07-17	Ahn et al.	
	32	20140283004	A1	2014-09-18	Moore	
	33	20140283030	A1	2014-09-18	Moore et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	20010079361	KR	A	2001-08-22	Kim, Sang Wook		
	2	1006701	EP	A2	2000-06-07	Lucent Technologies Inc		
	3	2006105093	WO	A2	2006-10-05	Univ Wake Forest		
	4	2006093557	WO	A2	2006-09-08	Univ Wake Forest		
	5	2005328336	AU	B2	2011-09-15	Univ Wake Forest		
	6	2011038420	WO	A2	2011-03-31	Great Wall Systems		
	7	1484884	EP	A2	2004-12-08	Microsoft Corp		
	8	2006230171	AU	B2	2012-06-21	Univ Wake Forest		
	9	1864226	EP	B1	2013-05-15	Univ Wake Forest		

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	10	2600236	CA	A1	2006-10-05	Univ Wake Forest		
	11	2012146265	WO	A1	2012-11-01	Voipfuture Gmbh		

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	2011 Apr 27 - (WO) International Search Report and Written Opinion - App PCT/US2010/054520	
	2	2011 Mar 04 - (US) Notice of Allowance - App 11/316,331	
	3	2011 Mar 03 - (EP) Communication Pursuant to Rules 70(2) and 70a(2) - App 06758213.0	
	4	2011 Feb 14 - (EP) Search Report - App 06758213.0	
	5	Fulp, Errin: "Errin Fulp," XP002618346, www.cs.wfu.edu/fulp/ewfPub.html, pgs. 1-5 (Copyright 2010)	
	6	2010 Sept 30 - (US) Office Action - App 11/390,976	
	7	2010 Sept 10 - (AU) Office Action - App 2006230171	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	8	2010 Aug 20 - (AU) Office Action - App 2005328336	
	9	2010 Jun 23 - (US) Final Rejection - App 11/316,331	
	10	2010 Apr 29 - (US) Interview Summary - App 11/390,976	
	11	2010 Mar 26 - (US) Final Rejection - App 11/390,976	
	12	2009 Sept 14 (US) Office Action - App 11/316,331	
	13	2009 Jun 24 - (US) Office Action - App 11/390,976	
	14	2008 Jul 03 - (WO) Written Opinion of the International Searching Authority - App PCT/US06/11291	
	15	2007 Aug 31 - (EP) Communication Pursuant to Rules 109 and 110 - App 05857614.1	
	16	Acharya et al, "OPTWALL: A Hierarchical Traffic-Aware Firewall," Department of Computer Science, Telecommunications Program, University of Pittsburgh, pgs. 1-11 (2007)	
	17	2006 Sept 11 - (WO) Written Opinion of the International Searching Authority - App PCT/US05/47008	
	18	Tarsa et al., "Balancing Trie-Based Policy representations for Network Firewalls," Department of Computer Science, Wake Forest University, pgs. 1-6 (2006)	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	19	Fulp, "Trie-Based Policy Representations for Network Firewalls," Proceedings of the IEEE International Symposium on Computer Communications (2005)	
	20	E. Fulp, "Optimization of Network Firewall Policies Using Ordered Sets and Directed Acyclical Graphs", Technical Report, Computer Scient Department, Wake Forest University, January 2004	
	21	E. Fulp et al., "Network Firewall Policy Tries", Technical Report, Computer Science Department, Wake Forest University, 2004	
	22	E. Al-Shaer et al., "Modeling and Management of Firewall Policies", IEEE Transactions on Network and Service Management, 1(1): 2004	
	23	E.W. Fulp, "Firewall Architectures for High Speed Networks", U.S. Department of Energy Grant Application, Funded September 2003	
	24	E. Al-Shaer et al., "Firewall Policy Advisor for Anomaly Discovery and Rule Editing", Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management, 2003	
	25	V.P. Ranganath, "A Set-Based Approach to Packet Classification", Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems, 889-894, 2003	
	26	M. Christiansen et al., "Using IDDsfor Packet Filtering", Technical Report, BRICS, October 2002	
	27	Lee et al., "Development Framework for Firewall Processors," IEEE, pgs. 352-355 (2002)	
	28	L. Qui et al., "Fast Firewall Implementations for Software and Hardware-Based Routers", Proceedings of ACM Sigmetrics, June 2001	
	29	D. Eppstein et al., "Internet Packet Filter Management and Rectangle Geometry", Proceedings of the Symposium on Discrete Algorithms, 827-835, 2001	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	30	E. Fulp, "Preventing Denial of Service Attacks on Quality of Service", Proceedings of the 2001 DARPA Information Survivability Conference and Exposition II, 2001	
	31	S. Goddard et al., "An Unavailability Analysis of Firewall Sandwich Configurations", Proceedings of the 6th IEEE Symposium on High Assurance Systems Engineering, 2001	
	32	G.V. Rooij, "Real Stateful TCP Packet Filtering in IP Filter", Proceedings of the 10th USENIX Security Symposium, 2001	
	33	P. Warkhede et al., "Fast Packet Classification for Two-Dimensional Conflict-Free Filters", Proceedings of IEEE INFOCOM, 1434-1443, 2001	
	34	D. Decasper et al., "Router Plugins: A Software Architecture for Next-Generation Routers", IEEE/ACM Transactions on Networking, 8(1): February 2000	
	35	A. Feldmann et al., "Tradeoffs for Packet Classification", Proceedings of the IEEE INFOCOM, 397-413, 2000	
	36	X. Gan et al., "LSMAC vs. LSNAT: Scalable Cluster-based Web servers", Journal of Networks, Software Tools, and Applications, 3(3): 175-185, 2000	
	37	A. Hari et al., "Detecting and Resolving Packet Filter Conflicts", Proceedings of IEEE INFOCOM, 1203-1212, 2000	
	38	O. Paul et al., "A full Bandwidth ATM Firewall", Proceedings of the 6th European Symposium on Research in Computer Security ESORICS'2000, 2000	
	39	J. Xu et al., "Design and Evaluation of a High-Performance ATM Firewall Switch and Its Applications", IEEE Journal on Selected Areas in Communications, 17(6): 1190-1200, June 1999	
	40	C. Benecke, "A Parallel Packet Screen for High Speed Networks", Proceedings of the 15th Annual Computer Security Applications Conference, 1999	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	41	R. Funke et al., "Performance Evaluation of Firewalls in Gigabit-Networks", Proceedings of the Symposium on Performance Evaluation of Computer and Telecommunication Systems, 1999
	42	S. Suri et al., "Packet Filtering in High Speed Networks", Proceedings of the Symposium on Discrete Algorithms, 969-970, 1999
	43	U. Ellermann et al., "Firewalls for ATM Networks", Proceedings of INFOSEC'COM, 1998
	44	V. Srinivasan et al., "Fast and Scalable Layer Four Switching", Proceedings of ACM SIGCOMM, 191-202, 1998
	45	M. Degermark et al., "Small Forwarding Tables for Fast Routing Lookups", Proceedings of ACM SIGCOMM, 4-13, 1997
	46	S,M. Bellovin et al., "Network Firewalls", IEEE Communications Magazine, 50-57, 1994
	47	W.E. Leland et al., "On the Self-Similar Nature of Ethernet Traffic", IEEE Transactions on Networking, 2(1); 15, 1994
	48	G. Brightwell et al., "Counting Linear Extensions is #P-Complete", Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing, 1991
	49	M. Al-Suwaiyel et al., "Algorithms for Trie Compaction", ACM Transactions on Database Systems, 9(2): 243-263, June 1984
	50	D. Comer, "Analysis of a Heuristic for Full Trie Minimization", ACM Transactions on Database Systems, 6(3): 513-537, Sept. 1981

If you wish to add additional non-patent literature document citation information please click the Add button Add

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

EXAMINER SIGNATURE			
Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

☒ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Bradley C. Wright/	Date (YYYY-MM-DD)	2016-03-11
Name/Print	Bradley C. Wright	Registration Number	38061

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	1	R.L. Graham et al., "Optimization and Approximation in Deterministic Sequencing and Scheduling: A Survey", Annals of Discrete Mathematics, 5: 287-326, 1979	
	2	E.L. Lawler, "Sequencing Jobs to Minimize Total Weighted Completion oTime Subject to Precedence Constraints", Annals of Discrete Mathematics, 2: 75-90, 1978	
	3	J.K. Lenstra et al., "Complexity of Scheduling Under Precedence Constraints", Operations Research, 26(1): 22-35,1978	
	4	R. Rivest, "On Self-Organizing Sequential Search Heuristics", Communications of the ACM, 19(2): 1976	
	5	W.E. Smith, "Various Optimizers for Single-Stage Productions", Naval Research Logistics Quarterly, 3: 59-66, 1956	
	6	Bellion, "HIgh Performance Packet Classification", http://www.hipac.org (Publication Date Unknown)	
	7	2011 Oct 18 - (EP) Communication Pursuant to Article 94(3) - App 06 758 213.0	
	8	2011 Jun 09 - (US) Notice of Allowance - App 11/390,976	
	9	2012 Jun 26 - (EP) Extended Search Report - App 05857614.1	
	10	2012 Jun 09 - (AU) Notice of Acceptance - App 2006230171	
	11	2011 Nov 11 - (AU) Second Office Action - App 2006230171	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	12	2013 Jan 17 - (CA) Office Action - App 2,600,236	
	13	2013 Jan 16 - (CA) Office Action - App 2,594,020	
	14	2012 Nov 20 - (EP) Communication under rule 71(3) - App 06 758 213.0	
	15	2013 Apr 18 - (EP) Decision to Grant a European Patent - App 06758212.0	
	16	2011 Aug 25 - (US) Non Final Rejection - App 12/871,806	
	17	2012 Feb 06 - (US) Final Rejection - App 12/871,806	
	18	2012 Aug 07 - (US) Non Final Rejection - App 12/871,806	
	19	2012 Nov 26 - (US) Final Rejection - App 12/871,806	
	20	2013 Apr 04 - (US) Notice of Allowance - App 12/871,806	
	21	2015 Jan 14 - (EP) Extended Search Report - App 10819667.6	
	22	2014 May 26 - (CA) Office Action - App 2010297968	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	23	2015 May 25 - (AU) Notice of Acceptance - App 2010297968	
	24	2015 May 14 - (US) Non Final Rejection - App 13/940,240	
	25	2015 Nov 27 - (US) Final Rejection - App 13/940,240	
	26	2016 Jan 28 - (WO) International Search Report and Written Opinion - App PCT/US2015/062691	
	27	Statement RE: Related Application, dated July 24, 2015	
	28	2015 Dec 22 - (US) Final Office Action - App 14/714,207	

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

☒ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Bradley C. Wright/	Date (YYYY-MM-DD)	2016-03-11
Name/Print	Bradley C. Wright	Registration Number	38061

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Espacenet

Bibliographic data: KR20010079361 (A) — 2001-08-22

FIRE WALL SYSTEM BASED ON NETWORK STATUS AND METHOD THEREOF

Inventor(s): JANG HUI JIN [KR]; KIM SANG WOOK [KR]; PARK BO SEOK [KR]; PARK IN SEONG [KR]; SONG BYEONG UK [KR] ± (JANG, HUI JIN, ; KIM, SANG WOOK, ; PARK, BO SEOK, ; PARK, IN SEONG, ; SONG, BYEONG UK)

Applicant(s): KIM SANG WOOK [KR] ± (KIM, SANG WOOK)

Classification: **- international:** ***H04L12/22; H04L12/26; H04L12/70;*** (IPC1-7): H04L12/56

- cooperative:

Application number: KR20010040894 20010709

Priority number(s): KR20010040894 20010709

Abstract of KR20010079361 (A)

PURPOSE: A fire wall system based on network status and a method thereof are provided to interrupt all illegal packet or traffic flowing into an internal network through an external information communication networks by analyzing the status of network traffic from various angles and differentially reflecting a traffic control policy together with recognition for status through a status graph. CONSTITUTION: A fire wall system based on network status consists of a network traffic collection unit(100), a static status information collection unit(200), a dynamic status information analysis unit(300), a traffic monitor unit(400), a rule generation unit(500), and a traffic control unit(600). The network traffic collection unit(100) classifies the network packets collected from an external information communication network by hosts, protocols and connections. The static status information collection unit(200), comprised of a statistics analyzer(230), a status unit table(220) and a packet manager(210), analyzes the information of the network packets transmitted from the network traffic collection unit(100), updates the information of the status unit table(220), and processes traffic information. The dynamic status information analysis unit(300), containing a traffic log table(310) and a traffic history analyzer(320), analyzes and records the long-term inclination and pattern of the network traffic of the traffic information processed at the static status information collection unit(200). The traffic monitor unit(400) monitors the trend of network traffic using each information provided from the static status information collection unit(200) and the dynamic status information analysis unit(300), recognizes an abnormal situation and outputs a state transition signal. The rule generation unit(500), composed of a policy manager(510), a status graph table(520) and an action manager(530), receives the state transition signal from the traffic monitor unit(400), creates a status graph, generates a packet filtering rule using the graph, and transfers the rule to the

traffic control unit(600) so that the traffic control unit(600) can cope with the abnormal situation.

(19) 대한민국특허청(KR)
(12) 공개특허공보(A)

(11) 공개번호 특2001-0079361
(43) 공개일자 2001년08월22일

(51) 국제분류코드
H04L 12/56 (2006.01)

(21) 출원번호 10-2001-0040894

(22) 출원일자 2001년07월09일

(71) 출원인
김상욱
대구광역시 수성구 수성1가 신세계타운 10동 501호, 대한민국

(72) 발명자
김상욱
대구광역시 수성구 수성1가 신세계타운 10동 501호, 대한민국
송병욱
대구광역시북구산격3동1313-37, 대한민국
박인성
경상북도경주시석장동858번지, 대한민국
장희진
대구광역시수성구지산동서한청산맨션102동1105호, 대한민국
박보석
대구광역시동구신암1동652-11, 대한민국

(74) 대리인
이영필
이해영

(77) 심사관
있음

전체 청구항 수 : 총 9 항

(54) 네트워크 상태 기반의 방화벽 장치 및 그 방법

(57)요 약

대 표 도 - 도 2

본 발명은 외부 정보 통신망으로부터 유입되는 모든 패킷의 정보를 분석하여 유효성과 불법성을 분석하고, 내부 네트워크의 비정상적인 상황을 유도하는 트래픽 패턴을 인식하며, 비정상적 상황을 파악하고 네트워크 관리 정책에 따라 대응하기 위한 네트워크 상태 기반의 방화벽 장치 및 그 방법에 관한 것으로,

본 발명의 장치는, 외부 정보 통신망으로부터 수집한 네트워크 패킷을 호스트, 프로토콜, 커넥션별로 구분하여 출력하는 네트워크트래픽수집수단; 상기 네트워크트래픽수집수단에서 전송되는 상기 네트워크 패킷의 정보를 분석하고, 상태 유닛 테이블 정보를 갱신하여 트래픽 정보를 가공하는 정적상태정보수집수단; 상기 정적상태정보수집수단에서 가공된 상기 트래픽 정보의 네트워크 트래픽의 장기적인 성향과 패턴을 분석하고 기록하는 동적상태정보분석수단; 상기 정적상태정보수집수단과 상기 동적상태정보분석수단에서 각각 제공하는 정보를 이용하여 상기 네트워크 트래픽의 동향을 감시하고 비정상적인 상황을 인식하여 상태 전이 신호를 출력하는 트래픽모니터수단; 및 상기 트래픽모니터수단에서 출력되는 상태 전이 신호를 받아서 생성된 상태 그래프 테이블을 이용하여 비정상적인 네트워크 상황을 정상적으로 유도하기 위한 패킷 필터링 룰을 생성하는 룰생성수단를 포함하는데 특징이 있다. 또한, 본 발명의 방법은, 게이트웨이를 통과하는 네트워크 패킷을 수집하여 상태 유닛별로 분류하는 제 1 과정; 상기 제 1 과정에서 수집된 상기 네트워크 패킷을 통계 및 패턴 분석 처리하여 현재의 트래픽 상황에 대한 정보를 생성하는 제 2 과정; 상기 제 2 과정에서 생성된 트래픽 상황에 대한 정보를 분석하고 감시하여 네트워크의 비정상적인 상황의 발생을 인지하여 상태 전이 신호를 생성하는 제 3 과정; 상기 제 3 과정에서 생성된 상태 전이 신호에 따라 상태 전이를 발생시켜, 그에 따른 정책을 채택하고 새로운 패킷 필터링 룰을 생성하는 제 4 과정; 및 상기 제 4 과정이 완료되면 현재 적용되고 있는 정책과 필터링 룰을 기록하고 상태 정보와 로그를 저장하는 제 5 과정을 포함하는데 특징이 있다.

키워드 : 트래픽, 상태 유닛, 통계 분석, 히스토리, 로그, 정책, 상태 그래프, 액션

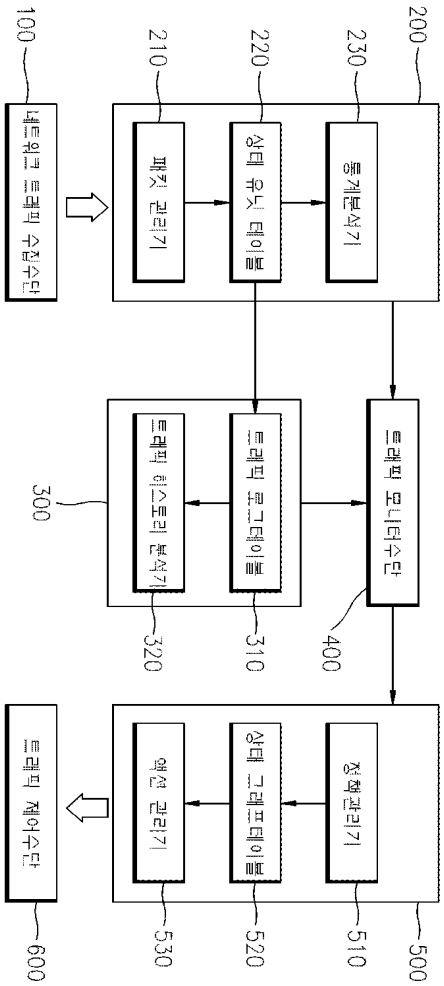

This Facsimile Page has been artificially created from KIPO data.

청구의 범위

청구항 1

외부 정보 통신망으로부터 수집한 네트워크 패킷을 호스트, 프로토콜, 커넥션별로 구분하여 출력하는 네트워크트래픽수집수단;

상기 네트워크트래픽수집수단에서 전송되는 상기 네트워크 패킷의 정보를 분석하고, 상태 유닛 테이블 정보를 갱신하여 트래픽 정보를 가공하는 정적상태정보수집수단;

상기 정적상태정보수집수단에서 가공된 상기 트래픽 정보의 네트워크 트래픽의 장기적인 성향과 패턴을 분석하고 기록하는 동적상태정보분석수단;

상기 정적상태정보수집수단과 상기 동적상태정보분석수단에서 각각 제공하는 정보를 이용하여 상기 네트워크 트래픽의 동향을 감시하고 비정상적인 상황을 인식하여 상태 전이 신호를 출력하는 트래픽모니터수단; 및

상기 트래픽모니터수단에서 출력되는 상태 전이 신호를 받아서 생성된 상태 그래프 테이블을 이용하여 비정상적인 네트워크 상황을 정상적으로 유도하기 위한 패킷 필터링 룰을 생성하는 룰생성수단를 포함하는 것을 특징으로 하는 네트워크 상태 기반의 방화벽 장치.

청구항 2

제 1 항에 있어서, 상기 정적상태정보수집수단은

상기 네트워크 패킷의 정보를 분석하는 패킷관리기; 및

상기 패킷관리기에서 분석된 상기 네트워크 패킷의 정보에 따라 상태 유닛 테이블의 정보를 갱신하는 통계분석기를 포함함을 특징으로 하는 네트워크 상태 기반의 방화벽 장치.

청구항 3

제 1 항에 있어서, 상기 동적상태정보분석수단은

발생하는 모든 이벤트에 대한 로그를 기록하여 관리자에게 네트워크와 방화벽의 관리에 대한 참고 자료를 제공하는 트래픽로그테이블; 및

상기 트래픽로그테이블에서 제공되는 참고 자료를 받아서 상기 네트워크 트래픽의 패턴과 그에 따른 대응 방식을 기록, 분석하는 트래픽히스토리분석기를 포함함을 특징으로 하는 네트워크 상태 기반의 방화벽 장치.

청구항 4

제 1 항에 있어서, 상기 룰생성수단은

상기 트래픽모니터수단에 의하여 생성된 상태 전이 신호를 받아 상태 그래프 테이블을 생성하는 정책관리기; 및

상기 정책관리기에서 생성된 상태 그래프 테이블에서 일어난 상태 전이에 따라 현재의 상황 지시자가 새로운 상태 노드로 이동하면 채택되는 새로운 네트워크 트래픽 관리 정책에 따라 현재 적용되고 있는 패킷 데이터 룰을 제거하거나 새로운 룰을 추가하는 액션관리기를 포함함을 특징으로 하는 네트워크 상태 기반의 방화벽 장치.

청구항 5

제 1 항에 있어서,

비정상적인 네트워크 상황을 정상적으로 유도하기 위한 패킷 필터링 룰을 전달받아 네트워크의 비정상적 상황에 대응하는 트래픽제어수단을 더 포함함을 특징으로 하는 네트워크 상태 기반의 방화벽 장치.

청구항 6

게이트웨이를 통과하는 네트워크 패킷을 수집하여 상태 유닛별로 분류하는 제 1 과정;

상기 제 1 과정에서 수집된 상기 네트워크 패킷을 통계 및 패턴 분석 처리하여 현재의 트래픽 상황에 대한 정보를 생성하는 제 2 과정;

상기 제 2 과정에서 생성된 트래픽 상황에 대한 정보를 분석하고 감시하여
네트워크의 비정상적인 상황의 발생을 인지하여 상태 전이 신호를 생성하는 제 3 과정;

상기 제 3 과정에서 생성된 상태 전이 신호에 따라 상태 전이를 발생시켜, 그에 따른 정책을 채택하고 새로운 패킷 필터링 룰을 생성하는 제 4 과정; 및

상기 제 4 과정이 완료되면 현재 적용되고 있는 정책과 필터링 룰을 기록하고 상태 정보와 로그를 저장하는 제 5 과정을 포함하는 것을 특징으로 하는 네트워크 상태 기반의 방화 방법.

청구항 7

제 6 항에 있어서, 상기 제 2 과정은

수집되는 상기 네트워크 패킷을 서브 네트워크에 포함되어 있는 호스트, 프로토콜, TCP 및 UDP 커넥션별로 분류하고, 각각의 특성에 따른 정보를 차별하여 분석하는 제 1 단계;

수집되는 상기 네트워크 패킷이 이미 생성되어 있는 상태 유닛에 포함된다면 상기 네트워크 패킷의 정보를 추출하여 상태 유닛의 정보를 갱신하고, 상태 유닛 테이블에 존재하지 않는 것이라면 상기 네트워크 패킷의 정보에 따라 상태 유닛을 생성하는 제 2 단계;

상기 네트워크 패킷의 시간적, 논리적 상관 관계를 트래픽 히스토리 정보 테이블을 이용하여 유효성을 검사하고 논리적 오류가 발견되면 상기 네트워크 패킷을 패기하는 제 3 단계;

트래픽의 전체적인 형태를 파악하기 위하여 유효성이 검사된 상기 네트워크

패킷의 정보를 대상으로 통계 정보를 생성하는 제 4 단계;

트래픽의 전체적 형태와 장기간적 동향을 파악하기 위하여 트래픽 패턴의 기록 및 그에 따른 대응 형태 등을 트래픽 히스토리 테이블에 기록하는 제 5 단계;

네트워크와 방화벽의 관리에 참조하기 위하여 상기 네트워크 패킷의 유효성 검사와 통계 정보 생성, 상태 유닛 테이블의 관리에 수반되는 로그를 로그 테이블에 기록하는 제 6 단계; 및

네트워크 상태의 비정상적 상황을 인식하기 위한 트래픽모니터수단에 트래픽 상태 정보를 제공하는 제 7 단계를 포함하는 것을 특징으로 하는 네트워크 상태 기반의 방화 방법.

청구항 8

제 6 항에 있어서, 상기 제 3 과정은

정적상태정보수집수단으로부터 제공되는 트래픽 상태 정보에서 정적으로 설정되어 있는 트래픽 정상 조건의 보정을 위한 평균 트래픽 정보를 추출하는 제 1 단계;

상기 제 1 단계에서 추출된 평균 트랙픽 정보를 이용하여 트래픽 정상 조건을 보정하는 제 2 단계;

트래픽 정상 조건과 트래픽 상태 정보를 비교하여 현재 네트워크의 상황의 비정상 여부를 판단하는 제 3 단계;

네트워크의 상태가 비정상적 상황임이 판명될 경우 새로운 트래픽 제어 정책의 채택을 위한 상태 전이 신호를 생성하는 제 4 단계; 및

상기 제 4 단계에서 생성된 상태 전이 신호를 룰생성수단의 상태 그래프에 전달하는 제 5 단계를 포함함을 특징으로 하는 네트워크 상태 기반의 방화 방법.

청구항 9

제 6 항에 있어서, 상기 제 4 과정은 트래픽모니터수단에서 전달된 상태 전이 신호에서 현재 네트워크 상태에 대한 정보를 추출하고, 비정상적인 요인을 추출하여 상태 그래프 테이블에 적용하는 제 1 단계;

상태 그래프 테이블에서 현재 상태 노드에서 전이 가능한 노드를 선별하여, 그 중에 상태 전이 신호의 정보를 수용할 수 있는 노드를 추출한 후 상태 전이를 수행하는 제 2 단계;

상태 그래프 테이블을 참조하여 새로운 상태 노드의 트래픽 제어 정책을 채택하고, 그에 따른 패킷 필터링 룰을 생성하는 제 3 단계; 및

상기 제 3 단계에서 새롭게 생성된 패킷 필터링 룰을 트래픽제어수단에 전달하여 네트워크의 비정상적 상황에 대응하는 제 4 단계를 포함함을 특징으로 하는 네트워크 상태 기반의 방화 방법.

명세서

발명의 명칭

네트워크 상태 기반의 방화벽 장치 및 그 방법{Apparatus for firewall of network status based Method thereof}

도면의 간단한 설명

[0001] 도 1은 일반적인 홈 네트워크 구조에서의 방화벽 배치도,

[0002] 도 2는 본 발명에 따른 네트워크 상태 기반의 방화벽 장치의 구성도,

[0003] 도 3은 본 발명에 따른 네트워크 상태 기반의 방화벽 장치를 구현하는 방법의 전체 흐름도,

[0004] 도 4는 도 3에 도시된 네트워크 상태 기반의 방화벽 장치를 구현하는 방법의 상태 정보 수집을 위한 흐름도,

[0005] 도 5는 도 3에 도시된 네트워크 상태 기반의 방화벽 장치를 구현하는 방법의 상태 전이 신호 생성을 위한 흐름도,

[0006] 도 6은 도 3에 도시된 네트워크 상태 기반의 방화벽 장치를 구현하는 방법의 패킷 필터링 룰 생성을 위한 흐름도.

[0007] <도면의 주요 부분에 대한 부호의 설명>

[0008] 100 : 네트워크트래픽수집수단 200 : 정적상태정보수집수단

[0009] 300 : 동적상태정보분석수단 400 : 트래픽모니터수단

[0010] 500 : 룰생성수단 600 : 트래픽제어수단

발명의 상세한 설명

[0011] 본 발명은 네트워크 상태 기반의 방화벽 장치 및 그 방법에 관한 것으로, 보다 상세하게는 외부 정보 통신망으로부터 유입되는 모든 패킷의 정보를 분석하여 유효성과 불법성을 분석하고, 내부 네트워크의 비정상적인 상황을 유도하는 트래픽 패턴을 인식하며, 비정상적 상황을 파악하고 네트워크 관리 정책에 따라 대응하기 위한 네트워크 상태 기반의 방화벽 장치 및 그 방법에 관한 것이다.

[0012] 최근에는 네트워크가 가정으로까지 확산되어 블루투스, 하비, 지니와 같은 홈네트워킹에 필요한 제반 기술 및 응용 분야가 개발되고 있다. 이와 같은 홈네트워크가 네트워크 분야에서 한 주류를 형성하면서 이질적인 기기들의 집합체인 홈네트워크를 제어, 관리하기 위한 에이전트가 필요한데 반하여, 에이전트로의 불법적 접근은 가정의 안전을 해칠 수도 있다. 이러한 기업, 가정 및 그 밖의 네트워크에서 외부로부터의 안전을 지키기 위해서는 네트워크의 특성에 따라 침입탐지, 침입차단, 역추적, 바이러스 방지 등의 많은 기술을 요구한다. 그 중 가장 기본적이고 모든 네트워크에 적용할 수 있는 것이 침입 차단 기술을 이용한 방화벽 시스템인 바, 방화벽은 외부로부터 내부 네트워크 자원을 보호하기 위해 필수적인 요소이며, 네트워크 내부의 구조와는 관계없이 외부와 내부의 접점에서 내부로의 패킷 흐름을 통제하고, 내부 사용자의 통신에 투명성을 제공하여야 한다. 도 1은 일반적인 홈 네트워크 구조에서의 방화벽 배치도

로서, 홈 네트워크 구조는 액세스 네트워크(10), 홈 게이트웨이 및 방화벽(11), 홈서버(12), PLC(13), DTV(14),
IEEE1394카메라(15), DVCR(16), 디지털오디오(17), 게임기(18), 수도(19), 보안기(20), 불루투스카메라(21), 팬(Fan)(22), 전등(23), 세탁기(24), 냉장고(25), 윈도우/커튼(26), 웹패드(27), PDA(28), 휴대폰(29), 화상전화기(30), 프린터(31), 스캐너(32), 팩스(33), PC(34), 콘센트(35) 등으로 구성되는데, 방화벽(11)은 그 발전 과정에 따라 세대별로 구분할 수 있는데, 제 1 세대 방화벽인 패킷 필터링은 각 패킷의 헤더 정보들에 기반하여 네트워크 사이의 전달을 허가하거나 거부하는 프로세스이다. 이와 같은 패킷 필터링은 한 서비스가 외부 네트워크에 항상 열려 있기 때문에 정적 상태를 갖으며, 장점으로는 오버헤드가 적고, 높은 성능을 보이며 구현하는데 비용이 적거나 들지 않고 트래픽 관리가 쉬운데 반하여, 단점으로는 서비스를 위해 열려있는 포트가 외부 클라이언트로부터 직접적인 연결을 허가하며, 서비스를 제공하는 포트가 외부 클라이언트에게 항상 열려있고, 복잡한 환경에서 신속한 관리가 어려우며, 특정 IP를 막더라도 IP 눈속임으로 공격당하기 쉽고, 사용자 인증과 투명성을 제공하지 않는다. 또한, 제 1 세대 방화벽과 비교되는 제 2 세대 방화벽의 가장 큰 차이점은 제 1 세대 방화벽에서는 방화벽(11)쪽의 서버 프로그램이 클라이언트쪽의 클라이언트 소프트웨어와 통신하면서 클라이언트의 존재를 감지해야만 하는 바, 서버(12)와 클라이언트들 각각에 대해 소프트웨어를 설치해야하는 번거러움이 있으며, 방화벽은 인지된 클라이언트들에 대해서만 방화벽 기능을 제공하는데 반하여, 제 2 세대 방화벽은 이러한 단점을 해결하여 클라이언트에 존재를 감지할 필요 없이 내부 호스트들에 대한 방화벽 기능을 수행하는 투명성을 제공하며, 또 하나의 제 1 세대 방화벽의 단점인 지속적 포트 개방
으로 인한 공격자의 공격 기회를 줄이기 위해 동적으로 제어하도록 설계되는데, 이와 같은 제 2 세대 방화벽 역시 제 1 세대 방화벽과 마찬가지로 사용자의 인증을 제공할 수 없고 IP 눈속임 공격에 대응하지 못하는 문제점이 있다.

[0013] 따라서, 본 발명은 이와 같은 문제점을 해결하기 위한 것으로서, 본 발명의 목적은 네트워크 트래픽의 상태를 다양한 관점에서 분석하고 상태 그래프를 통하여 상황에 대한 인식과 함께 트래픽 제어 정책을 차별적으로 반영함으로써, 외부 정보 통신망을 통하여 내부 네트워크로 유입되는 모든 불법적인 패킷 또는 패킷의 연속으로 이루어지는 트래픽을 차단하고, 내부 네트워크의 비정상적인 상황을 초래하는 패킷 흐름을 제어하여 원활하고 안전한 네트워크 상태를 유지할 수 있는 네트워크 상태 기반의 방화벽 장치 및 그 방법을 제공하는데 있다.

[0014] 본 발명의 다른 목적은 네트워크의 전반적인 상황을 파악할 수 있는 다양한 정보 및 로그를 제공하는 네트워크 상태 기반의 방화벽 장치 및 그 방법을 제공하는데 있다.

[0015] 이와 같은 목적을 달성하기 위한 본 발명의 특징은, 외부 정보 통신망으로부터 수집한 네트워크 패킷을 호스트, 프로토콜, 커넥션별로 구분하여 출력하는 네트워크트래픽수집수단; 상기 네트워크트래픽수집수단에서 전송되는 상기 네트워크 패킷의 정보를 분석하고, 상태 유닛 테이블 정보를 갱신하여 트래픽 정보를 가공하는 정적상태정보수집수단; 상기 정적상태정보수집수단에서 가공된 상기 트래픽 정보의
네트워크 트래픽의 장기적인 성향과 패턴을 분석하고 기록하는 동적상태정보분석수단; 상기 정적상태정보수집수단과 상기 동적상태정보분석수단에서 각각 제공하는 정보를 이용하여 상기 네트워크 트래픽의 동향을 감시하고 비정상적인 상황을 인식하여 상태 전이 신호를 출력하는 트래픽모니터수단; 및 상기 트래픽모니터수단에서 출력되는 상태 전이 신호를 받아서 생성된 상태 그래프 테이블을 이용하여 비정상적인 네트워크 상황을 정상적으로 유도하기 위한 패킷 필터링 룰을 생성하는 룰생성수단를 포함하는데 있다.

[0016] 바람직하게, 상기 정적상태정보수집수단은 상기 네트워크 패킷의 정보를 분석하는 패킷관리기; 및 상기 패킷관리기에서 분석된 상기 네트워크 패킷의 정보에 따라 상태 유닛 테이블의 정보를 갱신하는 통계분석기를 포함함을 특징으로 한다.

[0017] 바람직하게, 상기 동적상태정보분석수단은 발생하는 모든 이벤트에 대한 로그를 기록하여 관리자에게 네트워크와 방화벽의 관리에 대한 참고 자료를 제공하는 트래픽로그테이블; 및 상기 트래픽로그테이블에서 제공되는 참고 자료를 받아서 상기 네트워크 트래픽의 패턴과 그에 따른 대응 방식을 기록, 분석하는 트래픽히스토리분석기를 포함함을 특징으로 한다.

[0018] 바람직하게, 상기 룰생성수단은 상기 트래픽모니터수단에 의하여 생성된 상태 전이 신호를 받아 상태 그래프 테이블을 생성하는 정책관리기; 및 상기 정책관리기에서 생성된 상태 그래프 테이블에서 일어난 상태 전이에 따라 현재의 상황 지시자가 새로운 상태 노드로 이동하면 채택되는 새로운 네트워크 트래픽 관리 정책에 따라 현재 적용되고 있는 패킷 데이터 룰을 제거하거나 새로운 룰을 추가하는
액션관리기를 포함함을 특징으로 한다.

[0019] 바람직하게, 비정상적인 네트워크 상황을 정상적으로 유도하기 위한 패킷 필터링 룰을 전달받아 네트워크의 비정상적 상황에 대응하는 트래픽제어수단을 더 포함함을 특징으로 한다.

[0020] 본 발명의 다른 특징은, 게이트웨이를 통과하는 네트워크 패킷을 수집하여 상태 유닛별로 분류하는 제 1 과정; 상기 제 1 과정에서 수집된 상기 네트워크 패킷을 통계 및 패턴 분석 처리하여 현재의 트래픽 상황에 대한 정보를 생성하는 제 2 과정; 상기 제 2 과정에서 생성된 트래픽 상황에 대한 정보를 분석하고 감시하여 네트워크의 비정상적인 상황의 발생을 인지하여 상태 전이 신호를 생성하는 제 3 과정; 상기 제 3 과정에서 생성된 상태 전이 신호에 따라 상태 전이를 발생시켜, 그에 따른 정책을 채택하고 새로운 패킷 필터링 룰을 생성하는 제 4 과정; 및 상기 제 4 과정이 완료되면 현재 적용되고 있는 정책과 필터링 룰을 기록하고 상태 정보와 로그를 저장하는 제 5 과정을 포함하는데 있다.

[0021] 바람직하게, 상기 제 2 과정은 수집되는 상기 네트워크 패킷을 서브 네트워크에 포함되어 있는 호스트, 프로토콜, TCP 및 UDP 커넥션별로 분류하고, 각각의 특성에 따른 정보를 차별하여 분석하는 제 1 단계; 수집되는 상기 네트워크 패킷이 이미 생성되어 있는 상태 유닛에 포함된다면 상기 네트워크 패킷의 정보를 추출하여 상태 유닛의 정보를 갱신하고, 상태 유닛 테이블에 존재하지 않는 것이라면 상기 네트워크 패킷의 정보에 따라 상태 유닛을 생성하는 제 2 단계; 상기 네트워크 패킷의 시간적, 논리적 상관 관계를 트래픽 히스토리 정보 테이블을 이용하여 유효성을 검사하고 논리적 오류가 발견되면 상기 네트워크 패킷을 패기하는 제 3 단계; 트래픽의 전체적인 형태를 파악하기 위하여 유효성이 검사된 상기 네트워크 패킷의 정보를 대상으로 통계 정보를 생성하는 제 4 단계; 트래픽의 전체적 형태와 장기간적 동향을 파악하기 위하여 트래픽 패턴의 기록 및 그에 따른 대응 형태 등을 트래픽 히스토리 테이블에 기록하는 제 5 단계; 네트워크와 방화벽의 관리에 참조하기 위하여 상기 네트워크 패킷의 유효성 검사와 통계 정보 생성, 상태 유닛 테이블의 관리에 수반되는 로그를 로그 테이블에 기록하는 제 6 단계; 및 네트워크 상태의 비정상적 상황을 인식하기 위한 트래픽모니터수단에 트래픽 상태 정보를 제공하는 제 7 단계를 포함하는데 있다.

[0022] 바람직하게, 상기 제 3 과정은 정적상태정보수집수단으로부터 제공되는 트래픽 상태 정보에서 정적으로 설정되어 있는 트래픽 정상 조건의 보정을 위한 평균 트래픽 정보를 추출하는 제 1 단계; 상기 제 1 단계에서 추출된 평균 트랙픽 정보를 이용하여 트래픽 정상 조건을 보정하는 제 2 단계; 트래픽 정상 조건과 트래픽 상태 정보를 비교하여 현재 네트워크의 상황의 비정상 여부를 판단하는 제 3 단계; 네트워크의 상태가 비정상적 상황임이 판명될 경우 새로운 트래픽 제어 정책의 채택을 위한 상태 전이 신호를 생성하는 제 4 단계; 및 상기 제 4 단계에서 생성된 상태 전이 신호를 룰생성수단의 상태 그래프에 전달하는 제 5 단계를 포함함을 특징으로 한다.

[0023] 바람직하게, 상기 제 4 과정은 트래픽모니터수단에서 전달된 상태 전이 신호에서 현재 네트워크 상태에 대한 정보를 추출하고, 비정상적인 요인을 추출하여 상

태 그래프 테이블에 적용하는 제 1 단계; 상태 그래프 테이블에서 현재 상태 노드에서 전이 가능한 노드를 선별하여, 그 중에 상태 전이 신호의 정보를 수용할 수 있는 노드를 추출한 후 상태 전이를 수행하는 제 2 단계; 상태 그래프 테이블을 참조하여 새로운 상태 노드의 트래픽 제어 정책을 채택하고, 그에 따른 패킷 필터링 룰을 생성하는 제 3 단계; 및 상기 제 3 단계에서 새롭게 생성된 패킷 필터링 룰을 트래픽제어수단에 전달하여 네트워크의 비정상적 상황에 대응하는 제 4 단계를 포함함을 특징으로 한다.

[0024] 이하, 본 발명의 바람직한 실시예에 대하여 첨부 도면을 참조하여 상세히 설명한다.

[0025] 도 2는 본 발명에 따른 네트워크 상태 기반의 방화벽 장치의 구성도이다.

[0026] 도 2에 있어서, 본 발명은 네트워크트래픽수집수단(100), 정적상태정보수집수단(200), 동적상태정보분석수단(300), 트래픽모니터수단(400), 룰생성수단(500), 트래픽제어수단(600)으로 구성된다.

[0027] 네트워크트래픽수집수단(100)은 도 1의 액세스네트워크(10)인 외부 정보 통신망으로부터 수집한 네트워크 패킷을 호스트, 프로토콜, 커넥션별로 구분하여 출력하며, 정작상태정보수집수단(200)은 네트워크 패킷의 정보를 분석하는 패킷관리기(210) 및 패킷관리기(210)에서 분석된 네트워크 패킷의 정보에 따라 상태 유닛 테이블(220)의 정보를 갱신하는 통계분석기(230)를 포함하여, 네트워크트래픽수집수단(100)에서 전송되는 네트워크 패킷의 정보를 분석하고, 상태 유닛 테이블 정보를 갱신하여 트래픽 정보를 가공한다.

[0028] 동적상태정보분석수단(300)은 발생하는 모든 이벤트에 대한 로그를 기록하여 관리자에게 네트워크와 방화벽의 관리에 대한 참고 자료를 제공하는 트래픽로그테이블(310) 및 트래픽로그테이블(310)에서 제공되는 참고 자료를 받아서 네트워크 트래픽의 패턴과 그에 따른 대응 방식을 기록, 분석하는 트래픽히스토리분석기(320)를 포함하여, 정적상태정보수집수단(200)에서 가공된 트래픽 정보의 네트워크 트래픽의 장기적인 성향과 패턴을 분석하고 기록하며, 트래픽모니터수단(400)은 정적상태정보수집수단(200)과 동적상태정보분석수단(300)에서 각각 제공하는 정보를 이용하여 네트워크 트래픽의 동향을 감시하고 비정상적인 상황을 인식하여 상태 전이 신호를 출력한다.

[0029] 한편, 룰생성수단(500)은 트래픽모니터수단(400)에 의하여 생성된 상태 전이 신호를 받아 상태 그래프를 생성하는 정책관리기(510) 및 정책관리기(510)에서 생성된 상태 그래프 테이블(520)에서 일어난 상태 전이에 따라 현재의 상황 지시자가 새로운 상태 노드로 이동하면 채택되는 새로운 네트워크 트래픽 관리 정책에 따라 현재 적용되고 있는 패킷 데이터 룰을 제거하거나 새로운 룰을 추가하는 액션관리기(530)를 포함하여, 트래픽모니터수단(400)에서 출력되는 상태 전이 신호를 받아서 생성된 상태 그래프를 이용하여 비정상적인 네트워크 상황을 정상적으로 유도하기 위한 패킷 필터링 룰을 생성하여 트래픽제어수단(600)에 전달하여, 트래픽제어수단(600)으로 하여금 네트워크의 비정상적 상황에 대응하도록 한다.

[0030] 도 3은 본 발명에 따른 네트워크 상태 기반의 방화벽 장치를 구현하는 방법의 전체 흐름도이다.

[0031] 도 3에 있어서, 도 1에 도시된 게이트웨이(11)를 통과하는 네트워크 패킷을 수집하여 상태 유닛별로 분류하여(단계 S100), S100 단계에서 수집된 네트워크 패킷의 상태 정보를 수집하여, 통계 및 패턴 분석 처리하고 현재의 트래픽 상황에 대한 정보를 생성한다(단계 200). S200 단계에서 생성된 트래픽 상황에 대한 정보를 분석하고 감시하여 네트워크의 비정상적인 상황의 발생을 인지하여 상태 전이 신호를 생성하며(단계 S300), S300 단계에서 생성된 상태 전이 신호에 따라 상태 전이를 발생시켜, 그에 따른 정책을 채택하고 새로운 패킷 필터링 룰을 생

성한다(단계 S400). S400 단계에서 새로운 패킷 필터링 룰이 생성되면 현재 적용되고 있는 정책과 필터링 룰을 기록하고 상태 정보와 로그를 저장한다(단계 S500).

[0032] 도 4는 도 3에 도시된 네트워크 상태 기반의 방화벽 장치를 구현하는 방법의 상태 정보 수집을 위한 흐름도이다.

[0033] 도 4에 있어서, 수집되는 네트워크 패킷을 서브 네트워크에 포함되어 있는 호스트, 프로토콜, TCP 및 UDP 커넥션별로 분류하고, 각각의 특성에 따른 정보를 차별하여 분석한다(단계 S201). S201 단계에서 네트워크 정보를 분석한 후에, 수집되는 네트워크 패킷이 이미 생성되어 있는 상태 유닛에 포함된다면 네트워크 패킷의 정보를 추출하여 상태 유닛의 정보를 갱신하고, 상태 유닛 테이블(220)에 존재하지 않는 것이라면 네트워크 패킷의 정보에 따라 상태 유닛을 생성한다(단계 S202).

[0034] S202 단계에서 상태 유닛이 생성되면 네트워크 패킷의 시간적, 논리적 상관 관계를 트래픽 히스토리분석기(320)에 내장된 트래픽 히스토리 정보 테이블을 이용
하여 유효성을 검사하고(단계 S203), S203 단계에서 논리적 오류가 발견되는가에 따라 유효한가 판단한다(단계 S204). S204 단계에서 유효하지 않으면 네트워크 패킷을 패기하고(단계 S205), S204 단계에서 유효하면 트래픽의 전체적인 형태를 파악하기 위하여 유효성이 검사된 상기 네트워크 패킷의 정보를 대상으로 통계 정보를 생성한다(단계 206).

[0035] S206 단계에서 통계 정보를 처리하면 트래픽의 전체적 형태와 장기간적 동향을 파악하기 위하여 트래픽 패턴의 기록 및 그에 따른 대응 형태 등을 트래픽히스토리분석기(320)에 내장된 트래픽 히스토리 테이블에 기록하고(단계 S207), 네트워크와 방화벽의 관리에 참조하기 위하여 네트워크 패킷의 유효성 검사와 통계 정보 생성, 상태 유닛 테이블의 관리에 수반되는 로그를 트래픽 로그 테이블(310)에 기록한다(단계 S208). S208 단계에서 상태 정보 생성 로그 처리가 되면 네트워크 상태의 비정상적 상황을 인식하기 위한 트래픽모니터수단(400)에 트래픽 상태 정보를 제공한다(단계 S209).

[0036] 도 5는 도 3에 도시된 네트워크 상태 기반의 방화벽 장치를 구현하는 방법의 상태 전이 신호 생성을 위한 흐름도이다.

[0037] 도 5에 있어서, 정적상태정보수집수단(200)으로부터 제공되는 트래픽 상태 정보에서 정적으로 설정되어 있는 트래픽 정상 조건의 보정을 위한 평균 트래픽 정보를 추출하여(단계 S301), S301 단계에서 수집, 추출된 평균 트랙픽 정보를 이용하여 트래픽 정상 조건을 보정한다(단계 S302).

[0038] S302 단계에서 트래픽 정상 조건이 보정되면 트래픽 비정상 검사를 수행함으
로써(단계 S303), 트래픽 정상 조건과 트래픽 상태 정보를 비교하여 현재 네트워크의 상황의 비정상 여부를 판단한다(단계 S304). S304 단계에서 네트워크의 상태가 비정상적 상황이 아니면 S301 단계로 복귀하고, S304 단계에서 네트워크의 상태가 비정상적 상황임이 판명될 경우 새로운 트래픽 제어 정책의 채택을 위한 상태 전이 신호를 생성하여 룰생성수단(500)의 상태 그래프(520)에 전달한다(단계 S305).

[0039] 도 6은 도 3에 도시된 네트워크 상태 기반의 방화벽 장치를 구현하는 방법의 패킷 필터링 룰 생성을 위한 흐름도이다.

[0040] 도 6에 있어서, 트래픽모니터수단(400)에서 룰생성수단(500)에 전달된 상태 전이 신호를 분석하고(단계 S401), S401 단계에서 상태 전이 신호가 분석되면 현재 네트워크 상태에 대한 정보를 추출하여(단계 S402), 비정상적인 요인을 추출하여 상태 그래프 테이블에 적용한다. S402 단계에서 트래픽 상황 정보가 추출되면 상태 그래프 테이블(520)에서 현재 상태 노

드에서 전이 가능한 노드를 선별하여(단계 S403), 그 중에 상태 전이 신호의 정보를 수용할 수 있는 노드를 추출한 후 상태 전이가 가능한가 판단한다(단계 S404).

[0041] S404 단계에서 상태 전이가 가능하면 상태 전이를 수행하여(단계 S405), 상태 그래프 테이블(520)을 참조하여 새로운 상태 노드의 트래픽 제어 정책을 채택하고, 그에 따른 패킷 필터링 룰을 생성한다(단계 S406). S406 단계에서 새롭게 생성된 패킷 필터링 룰을 트래픽제어수단(600)에 전달하여 네트워크의 비정상적 상황에 대응하기 위하여 트래픽 제어를 수행한다(단계 S407).

[0042] 이와 같이, 본 발명의 상세한 설명에서는 구체적인 실시예에 관해 설명하였으나, 본 발명의 범주에서 벗어나지 않는 한도 내에서 여러가지 변형이 가능함은 물론이다. 그러므로, 본 발명의 범위는 설명된 실시예에 국한되어 정해져서는 안되며 후술하는 특허청구범위 뿐만 아니라 이 특허청구범위와 균등한 것들에 의해 정해져야 한다.

[0043] 결국, 본 발명에 의한 네트워크 상태 기반의 방화벽 장치 및 그 방법에 따르면 다음과 같은 이점이 발생한다.

[0044] 즉, 불법적인 패킷 조작을 통한 해킹 및 침입 시도를 차단하고 비정상적 패킷 흐름을 이용하여 내부 네트워크의 혼란을 일으키는 서비스 거부 공격과 같은 공격에 신속하고 효과적으로 대응함으로써 내부 호스트와 네트워크의 안전 및 신뢰성을 보장할 수 있다.

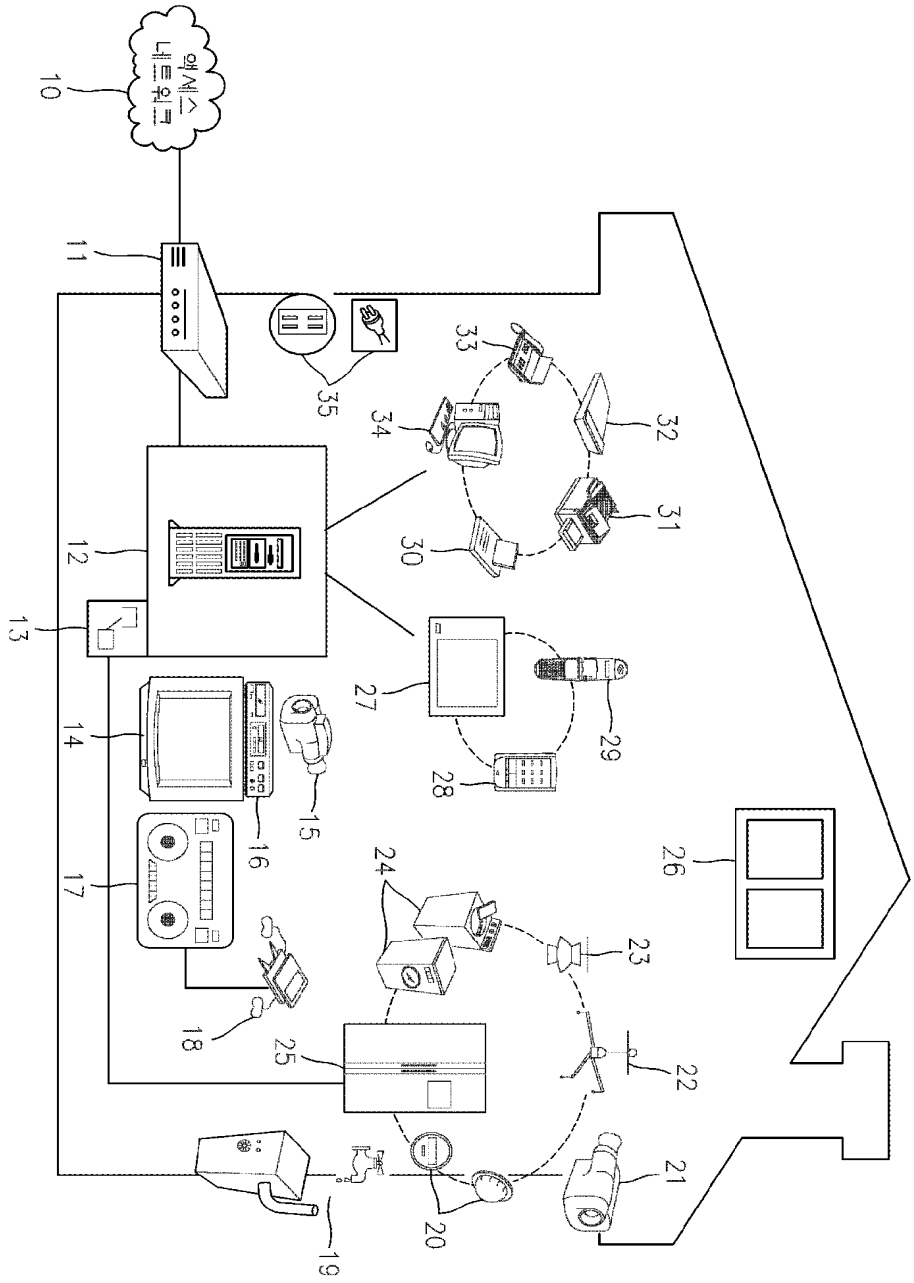
액세스
네트워크
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35

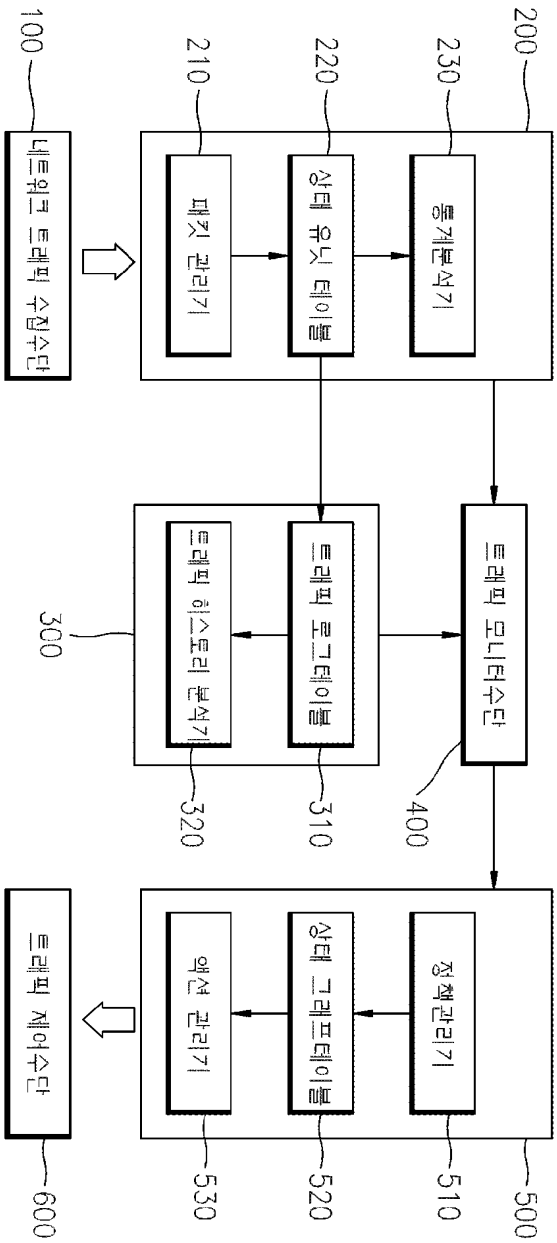
100
네트워크 트래픽 수집수단
200
210
패킷 관리기
220
상태 유닛 테이블
230
통계분석기
300
트래픽 히스토리 분석기
320
트래픽 로그테이블
310
트래픽 모니터수단
400
500
정책관리기
510
상태 그래프테이블
520
액션 관리기
530
트래픽 제어수단
600

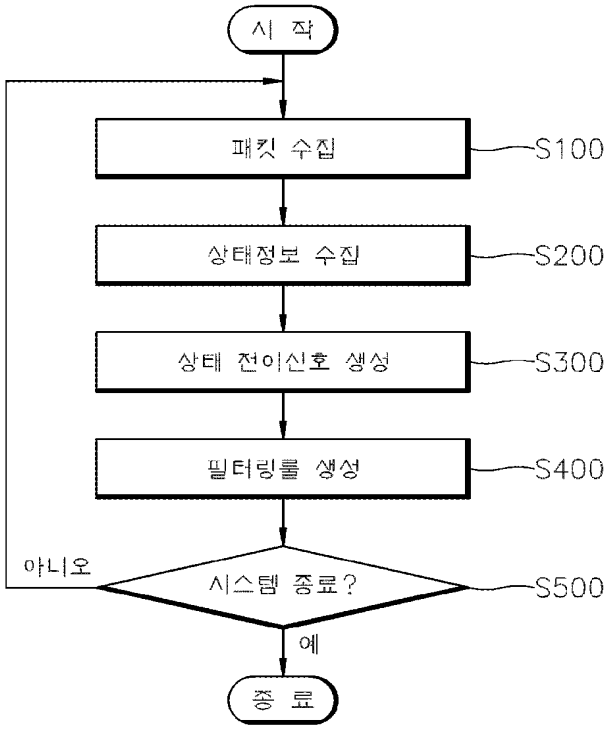
시 작
패킷 수집
S100
상태정보 수집
S200
상태 전이신호 생성
S300
필터링룰 생성
S400
아니오
시스템 종료?
S500
예
종 료

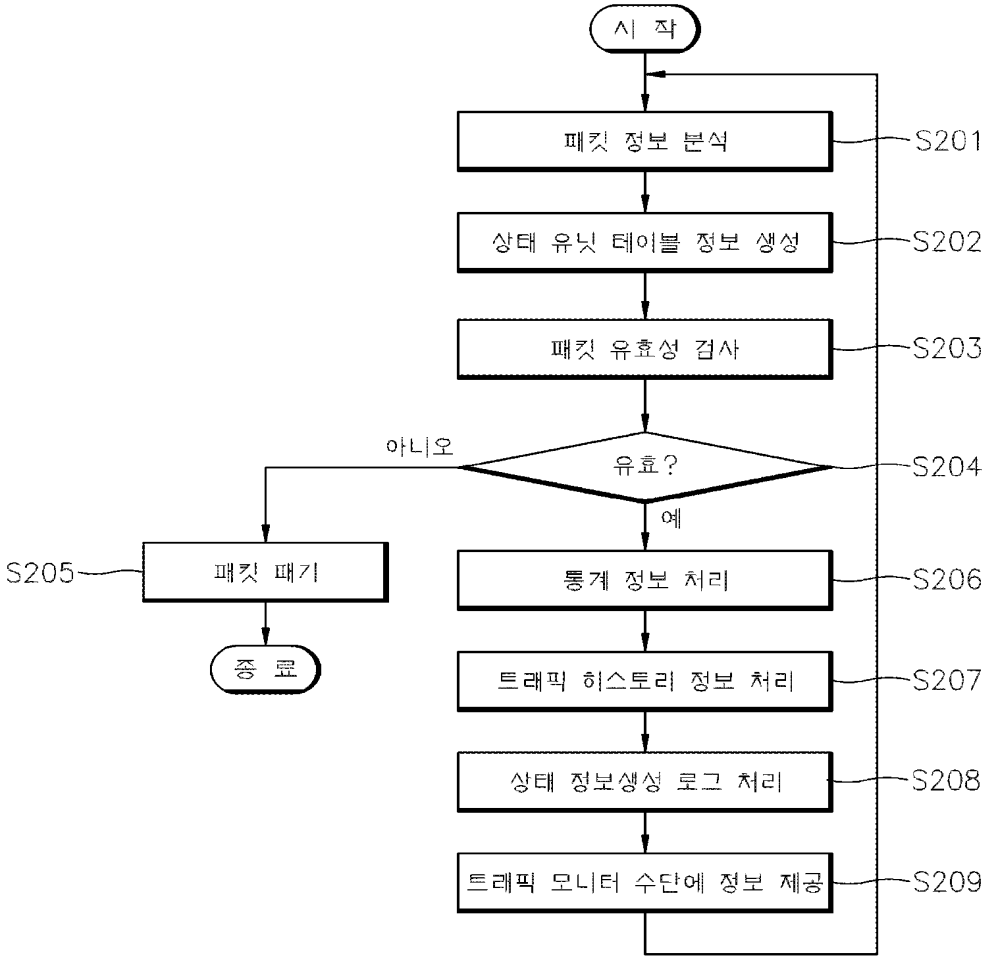
시 작
패킷 정보 분석
S201
상태 유닛 테이블 정보 생성
S202
패킷 유효성 검사
S203
유효?
S204
아니오
예
S205
패킷 폐기
종 료
통계 정보 처리
S206
트래픽 히스토리 정보 처리
S207
상태 정보생성 로그 처리
S208
트래픽 모니터 수단에 정보 제공
S209

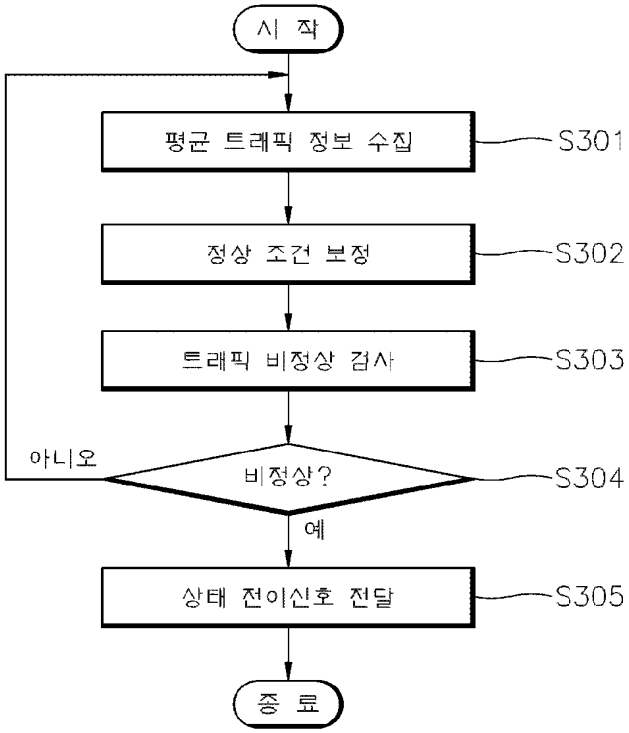
시 작
평균 트래픽 정보 수집
S301
정상 조건 보정
S302
트래픽 비정상 검사
S303
아니오
비정상?
S304
예
상태 전이신호 전달
S305
종 료

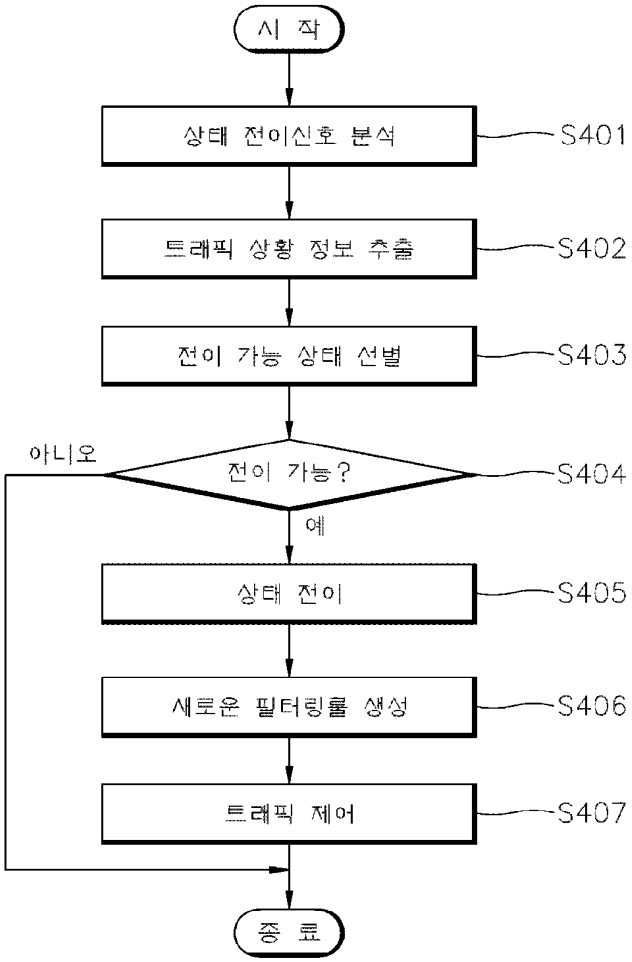
시 작
상태 전이신호 분석
S401
트래픽 상황 정보 추출
S402
전이 가능 상태 선별
S403
아니오
전이 가능?
S404
예
상태 전이
S405
새로운 필터링룰 생성
S406
트래픽 제어
S407
종 료

Electronic Patent Application Fee Transmittal

Application Number:	14253992
Filing Date:	16-Apr-2014
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Filer:	Bradley C. Wright/Holly Parker
Attorney Docket Number:	007742.00029

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Miscellaneous:				
Submission- Information Disclosure Stmt	2806	1	90	90
		Total in USD ($)		**90**

Electronic Acknowledgement Receipt

EFS ID:	25169367
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	Bradley C. Wright/Holly Parker
Filer Authorized By:	Bradley C. Wright
Attorney Docket Number:	007742.00029
Receipt Date:	11-MAR-2016
Filing Date:	16-APR-2014
Time Stamp:	11:20:26
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Deposit Account
Payment was successfully received in RAM	$90
RAM confirmation Number	8163
Deposit Account	190733
Authorized User	WRIGHT, BRADLEY

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

Charge any Additional Fees required under 37 CFR 1.17 (Patent application and reexamination processing fees)

Charge any Additional Fees required under 37 CFR 1.19 (Document supply fees)

Charge any Additional Fees required under 37 CFR 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Non Patent Literature	Non-Final_Rejection__892__1449_17666796.pdf	3313609 f94b21bf5072e9131e28903be6edbe2c064d83d3	no	82
Warnings:					
Information:					
2	Non Patent Literature	12871806_17760757.pdf	133950 d087d3fbbe5474766fb472b2fee3ab27b0b1d25c	no	3
Warnings:					
Information:					
3	Non Patent Literature	128718063_17748716.pdf	1885109 ff333809c196b7f7a899c06c24df7be6148b7b74	no	27
Warnings:					
Information:					
4	Non Patent Literature	128718063_17748705.pdf	72597 651be9b6f8516e40bd49cb8fea3b32608f57ad50	no	1
Warnings:					
Information:					
5	Non Patent Literature	128718063_17748704.pdf	280016 221a9d9fc8b69cd190f9a5e0c1fa7680eb10d7d4	no	8
Warnings:					
Information:					
6	Non Patent Literature	128718063_17748700.pdf	219866 2ff7ed11de580c33fbb1b4952953d81e9b627c07	no	6
Warnings:					
Information:					
7	Non Patent Literature	128718063_17748691.pdf	1056941 bfadf021c891e2d729069fdda5bb30ab7fbf6763	no	16
Warnings:					
Information:					

8	Non Patent Literature	128718063_17748689.pdf	93553 aeebc179c8fe4b8652ce76890598f0a71ce1aaa8	no	2
Warnings:					
Information:					
9	Non Patent Literature	128718063_17748686.pdf	66979 3884e538b6229e78e030c0dadc818f3747c9c505	no	1
Warnings:					
Information:					
10	Non Patent Literature	128718063_17748681.pdf	996456 64c55e7adf26bce0057805bbce8893d9484d48a2	no	30
Warnings:					
Information:					
11	Non Patent Literature	128718063_17748678.pdf	156743 31b7acfb7a47df115a561685241059d1d659011d	no	3
Warnings:					
Information:					
12	Non Patent Literature	128718063_17748674.pdf	954410 43e217a062cbd30d39f335589611d5c666afc0b0	no	22
Warnings:					
Information:					
13	Non Patent Literature	128718063_17748670.pdf	2407621 8036abd1045aaa8ace8ddc328d101c5e3756a533	no	40
Warnings:					
Information:					
14	Non Patent Literature	128718063_17748667.pdf	2047479 88a901174011cc43df79b97eefab85b253c4402e	no	29
Warnings:					
Information:					
15	Non Patent Literature	WO2006105093_17760770.pdf	130515 ffedf6884c4fc8189503e9ff67e119dd3c780109	no	3
Warnings:					
Information:					
16	Non Patent Literature	128718063_17748664.pdf	237193 7e270b000e2db4b9499e9c0955d46568991f37b2	no	4
Warnings:					
Information:					

17	Non Patent Literature	128718063_17748659.pdf	954305 3f4b7c0c39cc5a790a91fb62c5377170c2dd5d5b	no	11
Warnings:					
Information:					
18	Non Patent Literature	WO2006093557_17760776.pdf	330618 8b1a7002192e2958fc5a052a07ea100d87f170f4	no	6
Warnings:					
Information:					
19	Non Patent Literature	128718063_17748656.pdf	680644 f6452d529346441b5a7b66f39d8cb7b835fc821f	no	6
Warnings:					
Information:					
20	Non Patent Literature	128718063_17748652.pdf	809858 22ec24e861a9992648686b98ef88ea8c83a07bab	no	8
Warnings:					
Information:					
21	Non Patent Literature	128718063_17748644.pdf	843872 c9065af421ba9ec73f59a3c0abbae1bf99b75614	no	12
Warnings:					
Information:					
22	Non Patent Literature	128718063_17748638.pdf	1048051 62d2f5d39bb17ff27c66337e49aee4f6c9830c66	no	17
Warnings:					
Information:					
23	Non Patent Literature	128718063_17748635.pdf	910845 6e552f84cff6cf8d14434127ae71adb7df616dec	no	10
Warnings:					
Information:					
24	Non Patent Literature	128718063_17748632.pdf	1400185 d1485e17039f0972c9bca74d770ed585d940bd53	no	20
Warnings:					
Information:					
25	Non Patent Literature	128718063_17748622.pdf	1289250 fcdde68b1fc3570582a7a07ca4720c19537e39da	no	14
Warnings:					
Information:					

26	Non Patent Literature	128718063_17748618.pdf	585331 8f5ff661f447eb4f39db56058ccd1229bcaa0dfc	no	6
Warnings:					
Information:					
27	Non Patent Literature	128718063_17748614.pdf	1233391 f84378a5a084fed2aaf6cba80ececce4a98f2ba0	no	28
Warnings:					
Information:					
28	Non Patent Literature	128718063_17748610.pdf	266637 d6500988f54cbcebcbe36d73eca657359214084b	no	4
Warnings:					
Information:					
29	Non Patent Literature	128718063_17748605.pdf	888581 179b281ea068d86062d3084a0dbbc4632778ce46	no	10
Warnings:					
Information:					
30	Non Patent Literature	128718063_17748600.pdf	943167 06b16c35524d8f1d1a800b00e6f9293a6e11f14b	no	9
Warnings:					
Information:					
31	Non Patent Literature	128718063_17748594.pdf	1313706 cb35a896663a5ce778fd6ac66689f1bbdee57433	no	14
Warnings:					
Information:					
32	Non Patent Literature	128718063_17748590.pdf	1117022 140263fcd5145fb1e2e2a099b1e98f18b8d6d360	no	10
Warnings:					
Information:					
33	Non Patent Literature	128718063_17748585.pdf	897641 06520e07f43a9afbc0f7b7ad730dff8247a967af	no	15
Warnings:					
Information:					
34	Non Patent Literature	128718063_17748581.pdf	1165960 0c1b258923b72f80b1dc6cdf7b6363c8a95a44dd	no	10
Warnings:					
Information:					

35	Non Patent Literature	128718063_17748578.pdf	1551967 88a0f067f9a24decfbe884fa2df29cc25ab0d0b5	no	14
Warnings:					
Information:					
36	Non Patent Literature	128718063_17748577.pdf	1143829 25ac578781f5ed2e25f8b34aa201a16e82b83daa	no	10
Warnings:					
Information:					
37	Non Patent Literature	128718063_17748574.pdf	1347663 c2f845908e887514e60f22eea11e9ec7546ca129	no	11
Warnings:					
Information:					
38	Non Patent Literature	128718063_17748567.pdf	1222050 3962296f0c45ec3f80f4d5cbf25eedfe9ef404b1	no	10
Warnings:					
Information:					
39	Non Patent Literature	128718063_17748564.pdf	1118094 8e2f6a7b03a382382e26a106186d64dd475fcedb	no	16
Warnings:					
Information:					
40	Non Patent Literature	128718063_17748559.pdf	959293 0848229de6e141860751e8e14d6e021a7da1b25f	no	11
Warnings:					
Information:					
41	Non Patent Literature	128718063_17748527.pdf	655973 538c772956eacec011f7a184a5ac5a01ee34d6bd	no	8
Warnings:					
Information:					
42	Non Patent Literature	128718063_17748517.pdf	527673 ec48fa1268dddc8f6dcd05783730b341339a1371	no	5
Warnings:					
Information:					
43	Non Patent Literature	128718063_17748507.pdf	214277 aa3925bbab5f9517b3d068f7238a2985fbc4eec9	no	2
Warnings:					
Information:					

44	Non Patent Literature	128718063_17748503.pdf	1104955 9a812673586643b63a12e826cdd6cb53c332b895	no	18
Warnings:					
Information:					
45	Non Patent Literature	128718063_17748500.pdf	1428072 2b7e6fb1104bdca0e149137487d7ed7da7f11317	no	12
Warnings:					
Information:					
46	Non Patent Literature	128718063_17748492.pdf	1206677 9796972ed4e3a9278f61305fd084edf6b1d8e87c	no	12
Warnings:					
Information:					
47	Non Patent Literature	128718063_17748477.pdf	820708 3b06c3828c0a94ea80bec34368675ee9b0d9661f	no	8
Warnings:					
Information:					
48	Non Patent Literature	128718063_17748473.pdf	1310896 32621dc86ff0594ecbf95b9eae06e336a994092c	no	15
Warnings:					
Information:					
49	Non Patent Literature	128718063_17748470.pdf	732085 2a34455f7f90840a6f7a21d383c468f4c799cac2	no	7
Warnings:					
Information:					
50	Non Patent Literature	128718063_17748468.pdf	1693547 ba5c31882f9787e2ecd23ed1ce54a0f53f07b8df	no	21
Warnings:					
Information:					
51	Non Patent Literature	128718063_17748464.pdf	1210320 146d8a2fe7e3bb0cf21b700f047ad15909315d15	no	25
Warnings:					
Information:					
52	Non Patent Literature	128718063_17748458.pdf	3239881 e395eba28f8604973c9bb837217d4d755d203ca3	no	40
Warnings:					
Information:					

53	Non Patent Literature	128718063_17748456.pdf	920276 190c9f0448edc31547b0cc19fbb60ebdc4103a57	no	16
Warnings:					
Information:					
54	Non Patent Literature	128718063_17748450.pdf	1331868 48b1a01defbf29b5f288ee8887e8c4467304d3d8	no	14
Warnings:					
Information:					
55	Non Patent Literature	128718063_17748446.pdf	463011 fbdf3061cc05d35a338f7f55aea7b621fea7fa9f	no	5
Warnings:					
Information:					
56	Non Patent Literature	Corr_from_WIPO__-_N__of_Transmittal_of_the_International_Search_Report_17713511.PDF	8618992 3d1c384aa26e1d72a68822e3e4c8bd9f2de5249f	no	126
Warnings:					
Information:					
57	Non Patent Literature	Final_Rejection__1449_17645159.pdf	2839706 a8a9e7d43df581cbe3ff549c6bb5d8f0996c4060	no	68
Warnings:					
Information:					
58	Information Disclosure Statement (IDS) Form (SB08)	Information_Disclosure_Statement_Fillable_PDF_1_of_2_17760799.pdf	1039243 3d9a45dd9fe973a7e97a79be6702570076824e80	no	14
Warnings:					
Information:					
59	Information Disclosure Statement (IDS) Form (SB08)	Information_Disclosure_Statement_Fillable_PDF_2_of_2_17760800.pdf	1036129 06b377a9bb526d29ad1a46bf8cdb42776cbc06d2	no	6
Warnings:					
Information:					
A U.S. Patent Number Citation or a U.S. Publication Number Citation is required in the Information Disclosure Statement (IDS) form for autoloading of data into USPTO systems. You may remove the form to add the required data in order to correct the Informational Message if you are citing U.S. References. If you chose not to include U.S. References, the image of the form will be processed and be made available within the Image File Wrapper (IFW) system. However, no data will be extracted from this form. Any additional data such as Foreign Patent Documents or Non Patent Literature will be manually reviewed and keyed into USPTO systems.					
60	Foreign Reference	KR20010079361A_17690002.pdf	755268 996a657ccd2abe1d4285baa60bb510a5b1dccd29	no	19
Warnings:					
Information:					

61	Fee Worksheet (SB06)	fee-info.pdf	31033 ac5b8a831c685d97d3ff9c74bc16987483d50cf0	no	2

Warnings:	
Information:	
Total Files Size (in bytes):	69255587

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference	**FOR FURTHER ACTION**	see Form PCT/ISA/220 as well as, where applicable, item 5 below.
1743/2 PCT		
International application No.	International filing date *(day/month/year)*	(Earliest) Priority Date *(day/month/year)*
PCT/US2010/054520	**28 OCTOBER 2010 (28.10.2010)**	28 AUGUST 2009 (28.08.2009)

Applicant

GREAT WALL SYSTEMS et al

This International search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of **3** sheets.

☐ It is also accompanied by a copy of each prior art document cited in this report.

1. **Basis of the report**
   a. With regard to the **language**, the international search was carried out on the basis of :
      ☒ the international application in the language in which it was filed
      ☐ a translation of the international application into _____ , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))
   b. ☐ This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6*bis*(a)).
   c. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.
2. ☐ **Certain claims were found unsearchable** (See Box No. II)
3. ☐ **Unity of invention is lacking** (See Box No. III)
4. With regard to the **title,**
   ☒ the text is approved as submitted by the applicant.
   ☐ the text has been established by this Authority to read as follows:
5. With regard to the **abstract,**
   ☒ the text is approved as submitted by the applicant.
   ☐ the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority.
6. With regard to the drawings,
   a. the figure of the **drawings** to be published with the abstract is Figure No. **1**
      ☒ as suggested by the applicant.
      ☐ as selected by this Authority, because the applicant failed to suggest a figure.
      ☐ as selected by this Authority, because this figure better characterizes the invention.
   b. ☐ none of the figure is to be published with the abstract.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 528

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2010/054520

A. CLASSIFICATION OF SUBJECT MATTER

H04L 12/22(2006.01)i, H04L 12/26(2006.01)i, H04L 12/56(2006.01)i, H04L 29/06(2006.01)i

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H04L 12/22; G06F 15/16; H04L 9/00; G06F 9/00; H04L 12/56

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: firewall, filter, rule, policy

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 6098172 A1 (MICHAEL JOHN COSS et al.) 01 August 2000 See figures 5A, 5B and their corresponding explanation.	1-31
A	US 2006-0070122 A1 (STEVEN MICHAEL BELLOVIN) 30 March 2006 See paragraphs [0007], [0049]-[0055].	1-31
A	US 2005-0010765 A1 (BRIAN D. SWANDER et al.) 13 January 2005 See figures 9, 10.	1-31
A	KR 10-2001-0079361 A (S.W. KIM) 22 August 2001 See abstract and figure 2.	1-31

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents,such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
26 APRIL 2011 (26.04.2011)	**27 APRIL 2011 (27.04.2011)**

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
Government Complex-Daejeon, 189 Cheongsa-ro,
Seo-gu, Daejeon 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

CHUN, DAE NYUNG

Telephone No. 82-42-481-5991

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2010/054520

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6098172 A1	01.08.2000	EP 0909073 A2	14.04.1999
		JP 03-298832 B2	19.04.2002
		JP 11-167537 A	22.06.1999
		JP 2002-215478 A	02.08.2002
		JP 2009-295187 A	17.12.2009
		JP 3298832 B2	08.07.2002
US 2006-0070122 A1	30.03.2006	US 7051365 B1	23.05.2006
US 2005-0010765 A1	13.01.2005	CN 1574791 A	02.02.2005
		EP 1484886 A2	08.12.2004
		JP 2004-364315 A	24.12.2004
		KR 10-2004-0110981 A	31.12.2004
		US 2007-0204154 A1	30.08.2007
		US 7308711 B2	11.12.2007
		US 7761708 B2	20.07.2010
KR 10-2001-0079361 A	22.08.2001	None	

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 530

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:
GREGORY A. HUNT
JENKINS, WILSON, TAYLOR & HUNT, P.A.
SUITE 1200 UNIVERSITY TOWER
3100 TOWER BOULEVARD
DURHAM, NC 27707

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing *(day/month/year)*	03 JUL 2008
Applicant's or agent's file reference 1498/3/2 PCT	**FOR FURTHER ACTION** See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US06/11291	28 March 2006 (28.03.2006)	28 March 2005 (28.03.2005)

International Patent Classification (IPC) or both national classification and IPC

IPC: **H04L 9/00**(2006.01),**G06F 9/00,15/16/17/00,15/18**(2006.01)
USPC: 726/1,11,25;713/163,166

Applicant

WAKE FOREST UNIVERSITY

1. This opinion contains indications relating to the following items:

- ☒ Box No. I — Basis of the opinion
- ☐ Box No. II — Priority
- ☐ Box No. III — Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV — Lack of unity of invention
- ☒ Box No. V — Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI — Certain documents cited
- ☐ Box No. VII — Certain defects in the international application
- ☐ Box No. VIII — Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1*bis(b)* that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

3. For further details, see notes to Form PCT/ISA/220.

Name and mailing address of the ISA/ US	Date of completion of this opinion	Authorized officer
Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450 Alexandria, Virginia 22313-1450 Facsimile No. (571) 273-3201	24 June 2008 (24.06.2008)	Ayaz Sheikh DEBORAH A. THOMAS PARALEGAL SPECIALIST Telephone No. 571-272-3000

Form PCT/ISA/237 (cover sheet) (April 2007)

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US06/11291

Box No. I Basis of this opinion

1. With regard to the **language**, this opinion has been established on the basis of:

☒ the international application in the language in which it was filed

☐ a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43*bis*.1(a))

3. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of:

a. type of material

☐ a sequence listing

☐ table(s) related to the sequence listing

b. format of material

☐ on paper

☐ in electronic form

c. time of filing/furnishing

☐ contained in the international application as filed.

☐ filed together with the international application in electronic form.

☐ furnished subsequently to this Authority for the purposes of search.

4. ☐ In addition, in the case that more than one version or copy of a sequence listing and/or table(s) relating thereto has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 532

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US06/11291

Box No. V Reasoned statement under Rule 43 *bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Claims 1-41	YES
	Claims NONE	NO
Inventive step (IS)	Claims 1-41	YES
	Claims NONE	NO
Industrial applicability (IA)	Claims 1-41	YES
	Claims NONE	NO

2. Citations and explanations:

Claims 1-41 meet the criteria set out in PCT Article 33(2)-(3), because the prior art does not teach or fairly suggest optimization of firewall policy management for intrusion detection and intrusion prevention while maintaining the integrity of the original policy. Prior art does not teach a heuristic to order firewall rules the reduces average number of comparison by calculating a probability indication a likelihood of receiving a packet matching the rule sorted in order of non-increasing probability while maintaining the integrity of firewall policy.

Claims 1-41 meet the criteria set out in PCT Article 33 (4) because the subject matter claimed can be made or used in industry.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 533

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:
GREGORY A. HUNT
JENKINS, WILSON & TAYLOR, P.A.
SUITE 1400 UNIVERSITY TOWER
3100 TOWER BOULEVARD
DURHAM, NC 27707

PCT

WRITTEN OPINION OF THE
INTERNATIONAL SEARCHING AUTHORITY

(PCT Rule 43*bis*.1)

Date of mailing (*day/month/year*)	11 SEP 2006
Applicant's or agent's file reference 1498/3/3 PCT	FOR FURTHER ACTION See paragraph 2 below

International application No.	International filing date (*day/month/year*)	Priority date (*day/month/year*)
PCT/US05/47008	22 December 2005 (22.12.2005)	22 December 2004 (22.12.2004)

International Patent Classification (IPC) or both national classification and IPC

IPC: **G06F 17/00**(2006.01) G06F 17/00(2006.01)
USPC: 726/14

Applicant

WAKE FOREST UNIVERSITY

1. This opinion contains indications relating to the following items:

☒	Box No. I	Basis of the opinion
☐	Box No. II	Priority
☐	Box No. III	Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
☐	Box No. IV	Lack of unity of invention
☒	Box No. V	Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement
☐	Box No. VI	Certain documents cited
☐	Box No. VII	Certain defects in the international application
☐	Box No. VIII	Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notified the International Bureau under Rule 66.1*bis(b)* that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

3. For further details, see notes to Form PCT/ISA/220.

Name and mailing address of the ISA/ US	Date of completion of this opinion	Authorized officer
Mail Stop PCT, Attn: ISA/US Commissioner for Patents P.O. Box 1450 Alexandria, Virginia 22313-1450 Facsimile No. (571) 273-3201	16 August 2006 (16.08.2006)	Nasser G Moazzami Telephone No. 571-272-2100

Form PCT/ISA/237 (cover sheet) (April 2005)

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US05/47008

Box No. I Basis of this opinion

1. With regard to the **language**, this opinion has been established on the basis of:

☒ the international application in the language in which it was filed

☐ a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b)).

2. With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application and necessary to the claimed invention, this opinion has been established on the basis of:

a. type of material

☐ a sequence listing

☐ table(s) related to the sequence listing

b. format of material

☐ on paper

☐ in electronic form

c. time of filing/furnishing

☐ contained in the international application as filed.

☐ filed together with the international application in electronic form.

☐ furnished subsequently to this Authority for the purposes of search.

3. ☐ In addition, in the case that more than one version or copy of a sequence listing and/or table(s) relating thereto has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that in the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

4. Additional comments:

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY

International application No.
PCT/US05/47008

Box No. V Reasoned statement under Rule 43 *bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Claims NONE	YES
	Claims 1-52	NO
Inventive step (IS)	Claims NONE	YES
	Claims 1-52	NO
Industrial applicability (IA)	Claims 1-52	YES
	Claims NONE	NO

2. Citations and explanations:

Please See Continuation Sheet

Supplemental Box
In case the space in any of the preceding boxes is not sufficient.

V. 2. Citations and Explanations:
Claims 1-52 lack novelty under PCT Article 33(2) as being anticipated by Xu (U.S. Patent Pub. No. 2002/0038339 A1).

Regarding claims 1 and 31, Xu teaches a function-parallel firewall/system comprising:
a first firewall node for filtering received packets using a first portion of a rule set including a plurality of rules, the first portion including less than all of the rules in the rule set (figure 13, reference number 1361); and
at least one second firewall node for filtering packets using a second portion of the rule set, the second portion including at least one rule in the rule set that is not present in the first portion (figure 13, reference number 1362),
wherein the first and second portions together include all of the rules in the rule set (paragraph 0170).

Regarding claim 2, Xu teaches wherein the first and second portions respectively include rules present in first and second levels of a firewall hierarchy (figure 9E).

Regarding claim 3, Xu teaches wherein the first and second firewall nodes are connected to each other such that the rules in the first level are applied to received packets before the rules in the second level (paragraph 0092).

Regarding claim 4, Xu teaches wherein the first firewall node is adapted to forward packets that pass one of the rules in the first portion to an internal network in a manner that bypasses the second firewall node (paragraph 0122).

Regarding claim 5, Xu teaches wherein the first firewall node is adapted to discard packets that fail one of the rules in the first portion (paragraph 0171).

Regarding claim 6, Xu teaches wherein the rules in the second portion are more processor-intensive than the rules in the first portion (paragraph 0170).

Regarding claim 7, Xu teaches comprising:
a packet distributor for replicating each packet to the first and second firewall nodes (figure 13, reference number 1390); and
a gate machine for collecting results of applying the first and second portions of the rule set to each packet from the first and second firewall nodes and for making a forwarding decision based on the results (figure 13, reference number 1300).

Supplemental Box
In case the space in any of the preceding boxes is not sufficient.

Regarding claim 8, Xu teaches wherein each of the first and second firewall nodes is adapted to process the packets in a pipelined manner and wherein the gate machine is adapted to maintain state information for multiple packets simultaneously (paragraph 0168).

Regarding claim 9, Xu teaches wherein the gate machine is adapted to signal the second firewall node when processing of a packet is no longer needed due to matching of the packet with one of the rules used by the first firewall node (paragraph 0119).

Regarding claim 10, Xu teaches wherein the first and second firewall nodes implement rules of the same hierarchical level (figure 9E).

Regarding claim 11, Xu teaches wherein at least one of the rules in the first portion is present in the second portion to provide redundancy (paragraph 0177).

Regarding claims 12 and 22, Xu teaches wherein the first and second firewall nodes are configured to implement a gateless design wherein the rules are distributed among the first and second firewall nodes such that for any given packet, one of the first and second firewall nodes will accept the packet and the other of the first and second firewall nodes will deny the packet (paragraph 0201).

Regarding claim 13, Xu teaches wherein the first and second firewall nodes are implemented on different physical machines (figure 13, reference number 1361, 1362, and 1363).

Regarding claim 14, Xu teaches wherein the first and second firewall nodes are implemented on different processors in the same machine (paragraph 0244).

Regarding claims 15 and 28, Xu teaches wherein the first and second firewall nodes are adapted to simultaneously process copies of the same packets (figure 12, reference number 1261-1263).

Regarding claims 16, 24, 32, and 36, Xu teaches wherein the rule set includes firewall rules for filtering packets based on packet headers (paragraph 0091).

Regarding claims 17, 25, 33, and 37, Xu teaches wherein the rule set includes intrusion detection rules for filtering packets based on packet payloads (paragraph 0091).

Regarding claims 18, 26, 34, and 38, Xu teaches wherein the rule set includes intrusion protection rules for filtering packets based on a combination of packet headers and packet payloads (paragraph 0091).

Regarding claim 19, Xu teaches a stateful function-parallel firewall system comprising:
a first function-parallel firewall subsystem including a plurality of firewall nodes implementing a first rule set for filtering packets arriving in a network, wherein the first function-parallel firewall subsystem includes at least first and second firewall nodes respectively implementing first and second portions of the first rule set, the first portion including at least one rule that is not present in the second portion (figure 13, reference number 1361 and figure 8, INBOUND PACKETS); and
a second function-parallel firewall sub-system including a second set of firewall nodes implementing a second rule set for filtering packets departing from the network, wherein the second function-parallel firewall subsystem includes at least third and fourth firewall nodes, the third and fourth firewall nodes implementing first and second portions of the second rule set, the first portion of the second rule set including at least one rule that is not present in the second portion of the second rule set, wherein the first and second function-parallel firewall subsystems share state information regarding connections established through the first and second sets of firewall nodes (figure 13, reference number 1362 and figure 8, OUTBOUND PACKETS).

Regarding claim 20, Xu teaches comprising a database node coupled to the first and second function-parallel firewall subsystems, wherein the database node stores the state information, wherein the first set of firewall nodes obtains state information regarding connections established through the second set of firewall nodes from the database node, and wherein the second set of firewall nodes obtains state information regarding connections established through the first set of firewall nodes from the database node (figure 14, reference number 1471-1473).

Regarding claim 21, Xu teaches wherein the first and second sets of firewall nodes store copies of the state information maintained by the database node (paragraph 0014).

Regarding claim 23, Xu teaches wherein the firewall nodes of the first subsystem are adapted to simultaneously process copies of the same packets arriving in the network and wherein the firewall nodes in the second subsystem are adapted to simultaneously process copies of packets departing from the network (figure 12, reference number 1261-1263).

Regarding claim 27, Xu teaches a firewall grid comprising:
a plurality of firewall nodes being physically connected to each other via a network for filtering packets (figure 12, reference

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 538

Supplemental Box
In case the space in any of the preceding boxes is not sufficient.

number 1261, 1262, and 1263); and
a controller for controlling logical connections between the firewall nodes, wherein the controller is adapted to configure the firewall nodes in a logical firewall hierarchy without changing physical connections between the firewall nodes (paragraph 0165-0167).

Regarding claim 29, Xu teaches wherein the logical firewall hierarchy comprises a function-parallel hierarchy where individual firewall nodes implement different portions of a rule set (figure 13, reference number 1361-1363).

Regarding claim 30, Xu teaches wherein the controller is adapted to detect congestion at a first level in the firewall hierarchy and is adapted to reconfigure the logical connections between the firewall nodes to add a firewall node from another level in the firewall hierarchy to the level at which congestion was detected (figure 17).

Regarding claims 35 and 50, Xu teaches a method/computer program product for controlling access to a network based on a set of packet filtering rules distributed in a function-parallel manner, the method comprising:
distributing packet filtering rules of a rule set among a plurality of different nodes in a function-parallel manner so that at least some of the different nodes implement different portions of the rule set (figure 13, reference number 1361-1363);
replicating packets to each of the nodes (paragraph 0168); and
applying the rules to filter the packets (paragraph 0170).

Regarding claim 39, Xu teaches comprising, receiving a first packet at a first node, determining whether the packet matches a first rule implemented by the first node, and, in response to determining that a match occurs, notifying the other nodes of the match (paragraph 0119).

Regarding claim 40, Xu teaches comprising, at the other nodes, ceasing processing of the first packet (paragraph 0119).

Regarding claims 41 and 51, Xu teaches a method/computer program product for distributing rules in a function-parallel firewall, the method comprising:
defining a rule set for a function-parallel firewall (paragraph 0168);
assigning rules in the rule set to nodes and branches in a trie data structure, wherein each node in the trie data structure corresponds to a data field and each branch represents a value for each data field to be compared to values in corresponding fields in received packets (figure 13, reference number 1361-1363 and paragraph 0170); and
pruning the trie data structure in a manner that preserves ordering of the rules in the rule set (figure 20, deleting certain entries in 2091, 2092, and 2093).

Regarding claim 42, Xu teaches wherein pruning the trie data structure includes pushing a rule to its descendants and deleting the original rule (figure 20, deleting entries in 2091-2093).

Regarding claim 43, Xu teaches comprising reducing the depth of the trie using level reordering (paragraph 0246).

Regarding claim 44, Xu teaches wherein assigning rules in the rule set to nodes and branches in the trie data structure includes assigning rules requiring longer processing time to lower levels in the trie (paragraph 0170).

Regarding claims 45 and 52, Xu teaches a method/computer program product for assigning rules to a plurality of firewall nodes, the method comprising:
assigning packet filtering rules to nodes in a directed acyclical graph (DAG) (paragraph 0168);
representing relationships between rules by edges in the DAG (figure 9E); and
distributing the rules among a plurality of firewall nodes using the DAG such that different firewall nodes implement different rules and such that relationships between the rules specified in the DAG are preserved (figure 13, reference number 1361-1363 and paragraph 0170).

Regarding claim 46, Xu teaches wherein distributing the rules among a plurality of firewall nodes using the DAG includes distributing the rules such that intersecting accept rules are not present in more than one firewall node (figure 9E, rules that overlap are blacked out).

Regarding claim 47, Xu teaches wherein distributing the rules among a plurality of firewall nodes using the DAG includes distributing the rules such that the edges to not span multiple firewall nodes (figure 9E, rules that overlap are blacked out).

Regarding claim 48, Xu teaches wherein distributing the rules among a plurality of firewall nodes using the DAG includes replicating deny rules in each of the firewall nodes (paragraph 0177).

Regarding claim 49, Xu teaches comprising eliminating a deny rule implemented by a first firewall node in response to determining that the deny rule for the first firewall node includes no inbound edges and outbound edges from the first deny rule terminate only in a superset deny rule (paragraph 0177).

PATENT COOPERATION TREATY

From the INTERNATIONAL SEARCHING AUTHORITY

To:
Wright, Bradley
BANNER & WITCOFF, LTD.
1100 13th Street, N.W. Suite 1200
Washington, DC 20005-4051
ETATS-UNIS D'AMERIQUE

PCT

NOTIFICATION OF TRANSMITTAL OF THE INTERNATIONAL SEARCH REPORT AND THE WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY, OR THE DECLARATION

(PCT Rule 44.1)

	Date of mailing *(day/month/year)* 28 January 2016 (28-01-2016)
Applicant's or agent's file reference 007742.00064	**FOR FURTHER ACTION** See paragraphs 1 and 4 below
International application No. PCT/US2015/062691	International filing date *(day/month/year)* 25 November 2015 (25-11-2015)
Applicant CENTRIPETAL NETWORKS, INC.	

1. [X] The applicant is hereby notified that the international search report and the written opinion of the International Searching Authority have been established and are transmitted herewith.

   Filing of amendments and statement under Article 19:
   The applicant is entitled, if he so wishes, to amend the claims of the International Application (see Rule 46):

   When? The time limit for filing such amendments is normally two months from the date of transmittal of the International Search Report.

   How? Directly to the International Bureau of WIPO, 34 chemin des Colombettes 1211 Geneva 20, Switzerland, Fascimile No.: (41-22) 338.82.70

   For more detailed instructions, see *PCT Applicant's Guide,* International Phase, paragraphs 9.004 - 9.011.

2. [ ] The applicant is hereby notified that no international search report will be established and that the declaration under Article 17(2)(a) to that effect and the written opinion of the International Searching Authority are transmitted herewith.

3. [ ] **With regard to any protest** against payment of (an) additional fee(s) under Rule 40.2, the applicant is notified that:

   [ ] the protest together with the decision thereon has been transmitted to the International Bureau together with any applicant's request to forward the texts of both the protest and the decision thereon to the designated Offices.

   [ ] no decision has been made yet on the protest; the applicant will be notified as soon as a decision is made.

4. **Reminders**

   The applicant may **submit comments on an informal basis on the written opinion of the International Searching Authority** to the International Bureau. These comments will be made available to the public after international publication. The International Bureau will send a copy of such comments to all designated Offices unless an international preliminary examination report has been or is to be established.

   Shortly after the expiration of **18 months from the priority date, the international application will be published** by the International Bureau. If the applicant wishes to avoid or postpone publication, a notice of withdrawal of the international application, or of the priority claim, must reach the International Bureau before the completion of the technical preparations for international publication (Rules 90*bis*.1 and 90*bis*.3).

   Within **19 months** from the priority date, but only in respect of some designated Offices, a demand for international preliminary examination must be filed if the applicant wishes to postpone the entry into the national phase **until 30 months** from the priority date (in some Offices even later); otherwise, the applicant must, **within 20 months** from the priority date, perform the prescribed acts for **entry into the national phase** before those designated Offices. In respect of other designated Offices, the time limit of **30 months** (or later) will apply even if no demand is filed within 19 months. For details about the applicable time limits, Office by Office, see www.wipo.int/pct/en/texts/time_limits.html and the *PCT Applicant's Guide*, National Chapters.

   Within **19 months from the priority date, the applicant may request that a supplementary international search be carried out** by a different International Searching Authority that offers this service (Rule 45*bis*.1). The procedure for requesting supplementary international search is described in the *PCT Applicant's Guide*, International Phase, paragraphs 8.006-8.032.

Name and mailing address of the International Searching Authority European Patent Office, P.B. 5818 Patentlaan 2 NL-2280 HV Rijswijk Tel. (+31-70) 340-2040 Fax: (+31-70) 340-3016	Authorized officer FINNIE, Alistair Tel: +49 (0)89 2399-8251 FEB 8 - 2016

Form PCT/ISA/220 (July 2014)

PATENT COOPERATION TREATY

PCT

INTERNATIONAL SEARCH REPORT

(PCT Article 18 and Rules 43 and 44)

Applicant's or agent's file reference	**FOR FURTHER ACTION**	see Form PCT/ISA/220 as well as, where applicable, item 5 below.
007742.00064		
International application No.	International filing date *(day/month/year)*	(Earliest) Priority Date *(day/month/year)*
PCT/US2015/062691	25 November 2015 (25-11-2015)	10 February 2015 (10-02-2015)
Applicant		
CENTRIPETAL NETWORKS, INC.		

This international search report has been prepared by this International Searching Authority and is transmitted to the applicant according to Article 18. A copy is being transmitted to the International Bureau.

This international search report consists of a total of 3 sheets.

[X] It is also accompanied by a copy of each prior art document cited in this report.

1. **Basis of the report**
   a. With regard to the **language,** the international search was carried out on the basis of:
      [X] the international application in the language in which it was filed
      [ ] a translation of the international application into _____, which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1(b))
   b. [ ] This international search report has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43.6*bis*(a)).
   c. [ ] With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, see Box No. I.

2. [ ] **Certain claims were found unsearchable** (See Box No. II)

3. [ ] **Unity of invention is lacking** (see Box No III)

4. With regard to the **title,**
   [X] the text is approved as submitted by the applicant
   [ ] the text has been established by this Authority to read as follows:

5. With regard to the **abstract,**
   [X] the text is approved as submitted by the applicant
   [ ] the text has been established, according to Rule 38.2, by this Authority as it appears in Box No. IV. The applicant may, within one month from the date of mailing of this international search report, submit comments to this Authority

6. With regard to the **drawings,**
   a. the figure of the **drawings** to be published with the abstract is Figure No. 1
      [X] as suggested by the applicant
      [ ] as selected by this Authority, because the applicant failed to suggest a figure
      [ ] as selected by this Authority, because this figure better characterizes the invention
   b. [ ] none of the figures is to be published with the abstract

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 541

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/062691

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/701 H04L12/26
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014/281030 A1 (CUI LIANG [CN] ET AL) 18 September 2014 (2014-09-18) paragraph [0002] - paragraph [0003] paragraph [0004] paragraph [0034] paragraph [0046] - paragraph [0048] paragraph [0053] -----	1-25
X	US 2004/073655 A1 (KAN CHAO [US] ET AL) 15 April 2004 (2004-04-15)	1,24,25
A	paragraph [0006] paragraph [0037] -----	2-23
A	WO 2012/146265 A1 (VOIPFUTURE GMBH [DE]; SCHOLZ HENDRIK [DE]; WALLBAUM MICHAEL [DE]; KRUE) 1 November 2012 (2012-11-01) abstract -----	1-25

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
21 January 2016	28/01/2016

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Pais Gonçalves, A

1

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2015/062691

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014281030 A1	18-09-2014	NONE	
US 2004073655 A1	15-04-2004	EP 1418705 A2	12-05-2004
		US 2004073655 A1	15-04-2004
WO 2012146265 A1	01-11-2012	CA 2834168 A1	01-11-2012
		EP 2702740 A1	05-03-2014
		US 2014089504 A1	27-03-2014
		WO 2012146265 A1	01-11-2012

Information on Search Strategy - Pilot phase (see OJ 2015, A86)
The type of information contained in this sheet may change during the pilot for improving the usefulness of this new service.

Application Number
PCT/US2015/062691

TITLE: CORRELATING PACKETS IN COMMUNICATIONS NETWORKS

APPLICANT: CENTRIPETAL NETWORKS, INC.

IPC CLASSIFICATION: H04L12/701, H04L12/26

EXAMINER: Pais Gonçalves, A

CONSULTED DATABASES: WPI, EPODOC

CLASSIFICATION SYMBOLS DEFINING EXTENT OF THE SEARCH:

IPC:

CPC: H04L45/00, H04L43/00

FI/F-TERMS:

KEYWORDS OR OTHER ELEMENTS FEATURING THE INVENTION:

EPO FORM P04A42

PATENT COOPERATION TREATY

From the
INTERNATIONAL SEARCHING AUTHORITY

To:

see form PCT/ISA/220

PCT

WRITTEN OPINION OF THE INTERNATIONAL SEARCHING AUTHORITY
(PCT Rule 43*bis*.1)

Date of mailing
(day/month/year) see form PCT/ISA/210 (second sheet)

Applicant's or agent's file reference
see form PCT/ISA/220

FOR FURTHER ACTION
See paragraph 2 below

International application No.	International filing date *(day/month/year)*	Priority date *(day/month/year)*
PCT/US2015/062691	25.11.2015	10.02.2015

International Patent Classification (IPC) or both national classification and IPC
INV. H04L12/701 H04L12/26

Applicant
CENTRIPETAL NETWORKS, INC.

1. This opinion contains indications relating to the following items:

- ☒ Box No. I Basis of the opinion
- ☐ Box No. II Priority
- ☐ Box No. III Non-establishment of opinion with regard to novelty, inventive step and industrial applicability
- ☐ Box No. IV Lack of unity of invention
- ☒ Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step and industrial applicability; citations and explanations supporting such statement
- ☐ Box No. VI Certain documents cited
- ☒ Box No. VII Certain defects in the international application
- ☒ Box No. VIII Certain observations on the international application

2. **FURTHER ACTION**

If a demand for international preliminary examination is made, this opinion will usually be considered to be a written opinion of the International Preliminary Examining Authority ("IPEA") except that this does not apply where the applicant chooses an Authority other than this one to be the IPEA and the chosen IPEA has notifed the International Bureau under Rule 66.1*bis*(b) that written opinions of this International Searching Authority will not be so considered.

If this opinion is, as provided above, considered to be a written opinion of the IPEA, the applicant is invited to submit to the IPEA a written reply together, where appropriate, with amendments, before the expiration of 3 months from the date of mailing of Form PCT/ISA/220 or before the expiration of 22 months from the priority date, whichever expires later.

For further options, see Form PCT/ISA/220.

Name and mailing address of the ISA:	Date of completion of this opinion	Authorized Officer
European Patent Office D-80298 Munich Tel. +49 89 2399 - 0 Fax: +49 89 2399 - 4465	see form PCT/ISA/210	Pais Gonçalves, A Telephone No. +49 89 2399-0

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 545

Box No. I Basis of the opinion

1. With regard to the **language**, this opinion has been established on the basis of:

   ☒ the international application in the language in which it was filed.

   ☐ a translation of the international application into , which is the language of a translation furnished for the purposes of international search (Rules 12.3(a) and 23.1 (b)).

2. ☐ This opinion has been established taking into account the **rectification of an obvious mistake** authorized by or notified to this Authority under Rule 91 (Rule 43*bis*.1(a))

3. ☐ With regard to any **nucleotide and/or amino acid sequence** disclosed in the international application, this opinion has been established on the basis of a sequence listing:

   a. ☐ forming part of the international application as filed:

      ☐ in the form of an Annex C/ST.25 text file.

      ☐ on paper or in the form of an image file.

   b. ☐ furnished together with the international application under PCT Rule 13*ter*.1(a) for the purposes of international search only in the form of an Annex C/ST.25 text file.

   c. ☐ furnished subsequent to the international filing date for the purposes of international search only:

      ☐ in the form of an Annex C/ST.25 text file (Rule 13*ter*.1(a)).

      ☐ on paper or in the form of an image file (Rule 13*ter*.1(b) and Administrative Instructions, Section 713).

4. ☐ In addition, in the case that more than one version or copy of a sequence listing has been filed or furnished, the required statements that the information in the subsequent or additional copies is identical to that forming part of the application as filed or does not go beyond the application as filed, as appropriate, were furnished.

5. Additional comments:

Box No. V Reasoned statement under Rule 43*bis*.1(a)(i) with regard to novelty, inventive step or industrial applicability; citations and explanations supporting such statement

1. Statement

Novelty (N)	Yes:	Claims	1-25
	No:	Claims	
Inventive step (IS)	Yes:	Claims	
	No:	Claims	1-25
Industrial applicability (IA)	Yes:	Claims	1-25
	No:	Claims	

2. Citations and explanations

see separate sheet

Box No. VII Certain defects in the international application

The following defects in the form or contents of the international application have been noted:

see separate sheet

Box No. VIII Certain observations on the international application

The following observations on the clarity of the claims, description, and drawings or on the question whether the claims are fully supported by the description, are made:

see separate sheet

Reference is made to the following documents:

D1 US 2014/281030 A1 (CUI LIANG [CN] ET AL) 18 September 2014 (2014-09-18)

D2 US 2004/073655 A1 (KAN CHAO [US] ET AL) 15 April 2004 (2004-04-15)

section VIII

1 Claims 1, 24 and 25 do not fulfil the requirements of Article 6 PCT for the following reasons:

1.1 In the claims it is not clear what is meant by the expression "data identifying the host", how such data is determined or the purpose of generating and transmitting said data.

1.2 According to the description, it seems that this data is used to identify which packets are associated with a particular flow or session. However, this is not clear from the wording of the claims.

1.3 In addition, the desired scope of protection of independent claim 25 is not clearly defined due to the wording "one or more".

section V

2 Assuming the interpretation indicated in point 1.2 above, it seems that claim 1 would be directed to a method of monitoring packets received by a network device and packets transmitted by the device and correlating them in order to determined that the packets are associated with a flow.

2.1 Such a method is already known from the disclosure of D1, see e.g. [0002] to [0004] or D2, see e.g. [0037].

2.2 The only apparent distinguishing feature between claim 1 and e.g. D1 could be that the network device and the host are located in different networks.

2.3 However, such distinguishing feature is a simple modification of the basic concept disclosed in D1 lying within the general design competence of a person skilled in the art who would consider such measure to be at least implicitly disclosed in D1, see e.g. figures 7 and 10A.

2.4 In addition, these arguments are also valid for independent Claim 24, which contains the same combination of features as claim 1 in terms of an apparatus claim, i.e. the apparatus (system) for carrying out the method of claim 1.

2.5 Said arguments are also valid for independent claim 25 which is directed to computer-readable media comprising instructions essentially to carry out the method of claim 1.

2.6 Therefore, the subject-matter of Claims 1, 24 and 25 does not involve an **inventive step**, Article 33 (1), (3) PCT.

2.7 The additional features of the dependent claims seem to relate to minor design details and/or implementation measures, which are known or directly derivable from the above-referred documents or constitute normal practice of a person skilled in the art. These claims thus, either alone or in combination, do not seem to add anything new or of inventive significance to any of the previously-referred claims, see passages indicated in the search report.

section VII

3 Having regard to the above objections, it is not clear at present which part of the application could serve as a basis for a new and inventive subject-matter.

4 Should the Applicant nevertheless regard some particular matter of the application as having an inventive contribution over the prior art, then an independent claim relating to such subject-matter should be filed. The Applicant is also requested to indicate in the letter of reply the difference of said subject-matter vis-à-vis the state of the art and the significance thereof.

4.1 The above-cited documents D1 and D2 should be acknowledged and briefly discussed in the opening part of the description, Rule 5.1 (a) (ii) PCT, making clear any inventive contribution of the claimed invention over the prior art.

4.2 The independent claims should be drafted in a two-part form, having in their pre-characterizing part all features known from D1 (refer to point 2 above), Rule 6.3 (b) (ii) PCT.

4.3 The opening part of the description should be modified to bring it into agreement with the amended claim, Rule 5.1(a)(iii) PCT, if required.

4.4 The claims should include reference signs relating to the features referred to therein, Rule 6.2 (b) PCT.

US 20140281030A1

(19) **United States**

(12) **Patent Application Publication**

Cui et al.

(10) **Pub. No.: US 2014/0281030 A1**

(43) **Pub. Date: Sep. 18, 2014**

(54) **VIRTUAL NETWORK FLOW MONITORING**

(71) Applicant: **VMWARE, INC.**, Palo Alto, CA (US)

(72) Inventors: **Liang Cui**, Beijing (CN); **Wei Zhang**, Beijing (CN); **Jingchun Jason Jiang**, Beijing (CN); **Caixia Jiang**, Beijing (CN)

(73) Assignee: **VMWARE, INC.**, Palo Alto, CA (US)

(21) Appl. No.: **13/831,685**

(22) Filed: **Mar. 15, 2013**

Publication Classification

(51) **Int. Cl.**
H04L 12/717 (2006.01)

(52) **U.S. Cl.**
CPC *H04L 45/42* (2013.01)
USPC **709/244**

(57) **ABSTRACT**

Embodiments perform end-to-end virtual network flow monitoring in a virtual datacenter and provide differentiated views to users based on user role. A target flow pattern describing data packets of interest is distributed to a plurality of applications managing VMs in the virtual datacenter, such as hosts, virtual gateways, and other virtual network applications. Each of the applications monitors data packets routed by the application by comparing the data packets to the flow pattern and selectively collecting context data describing the data packets. The context data collected by the applications is aggregated at a remote server for analysis and reporting.

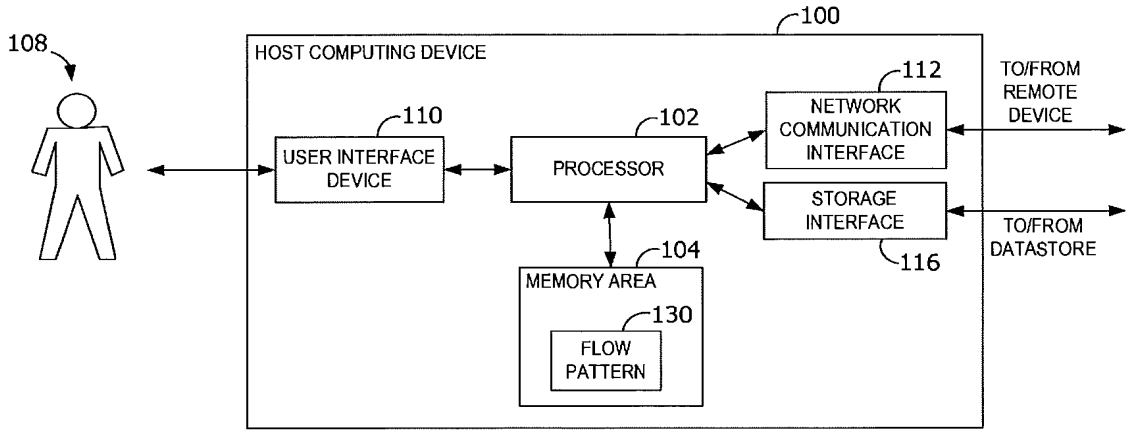

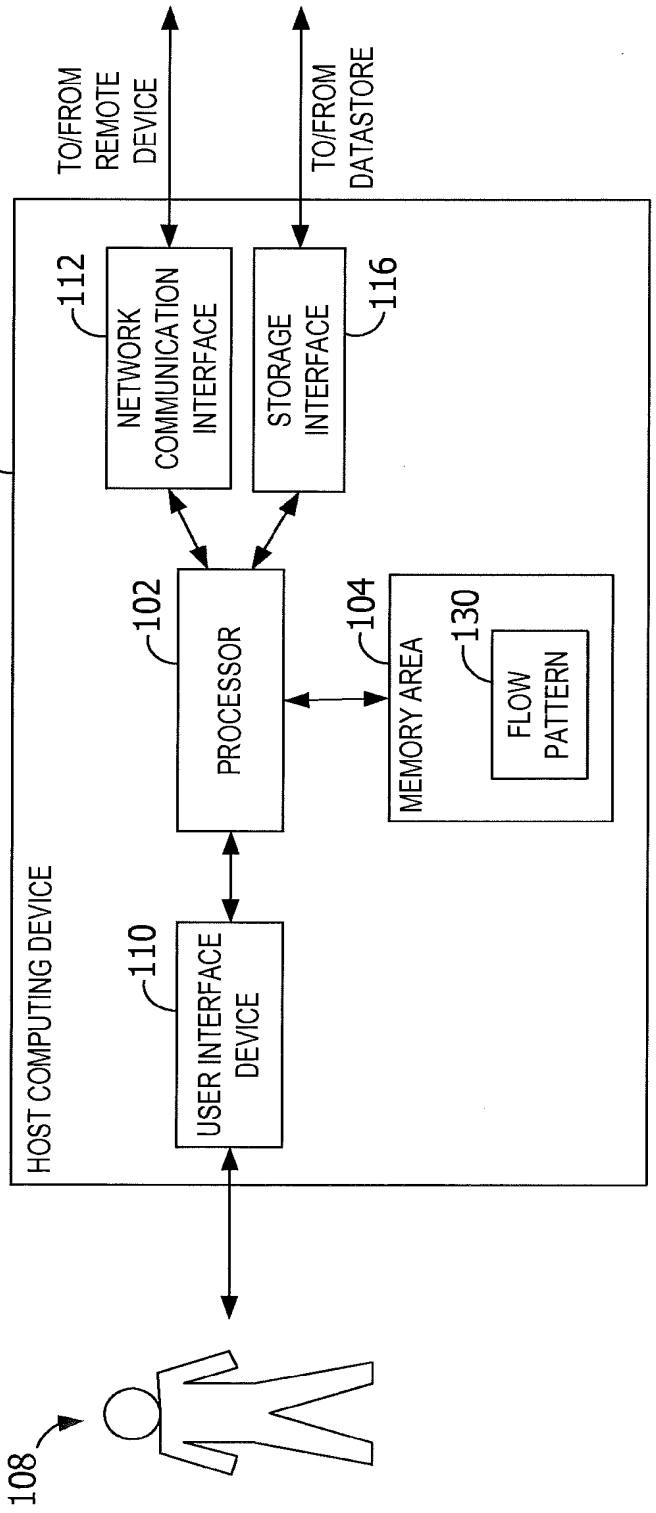
TO/FROM REMOTE DEVICE
TO/FROM DATASTORE
100
112
NETWORK COMMUNICATION INTERFACE
116
STORAGE INTERFACE
102
PROCESSOR
104
MEMORY AREA
130
FLOW PATTERN
110
USER INTERFACE DEVICE
HOST COMPUTING DEVICE
108

FIG. 1

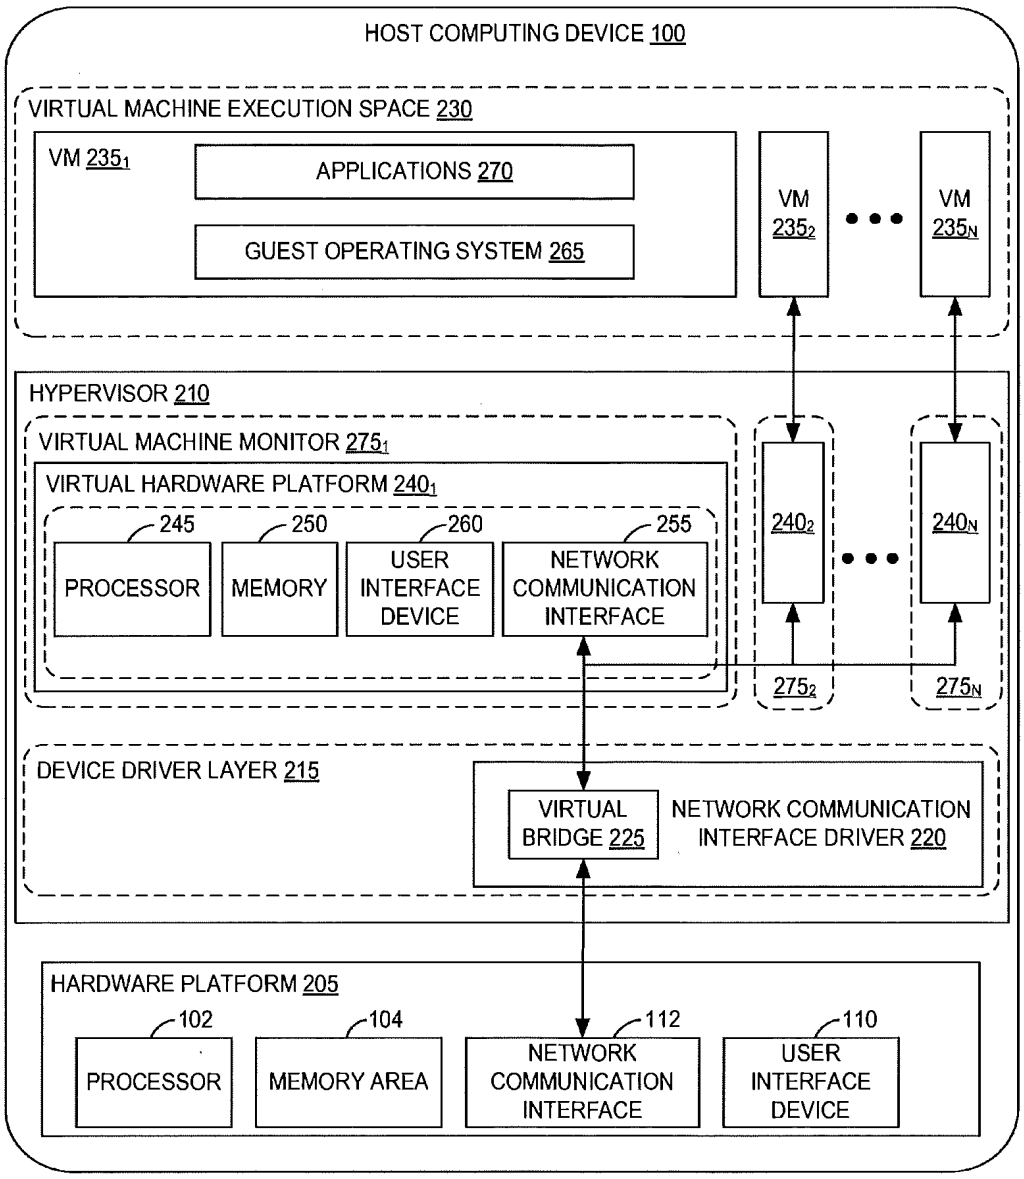
HOST COMPUTING DEVICE 100
VIRTUAL MACHINE EXECUTION SPACE 230
VM 235₁
APPLICATIONS 270
GUEST OPERATING SYSTEM 265
VM 235₂
VM 235N
HYPERVISOR 210
VIRTUAL MACHINE MONITOR 275₁
VIRTUAL HARDWARE PLATFORM 240₁
245
250
260
255
PROCESSOR
MEMORY
USER INTERFACE DEVICE
NETWORK COMMUNICATION INTERFACE
240₂
240N
275₂
275N
DEVICE DRIVER LAYER 215
VIRTUAL BRIDGE 225
NETWORK COMMUNICATION INTERFACE DRIVER 220
HARDWARE PLATFORM 205
102
104
112
110
PROCESSOR
MEMORY AREA
NETWORK COMMUNICATION INTERFACE
USER INTERFACE DEVICE

FIG. 2

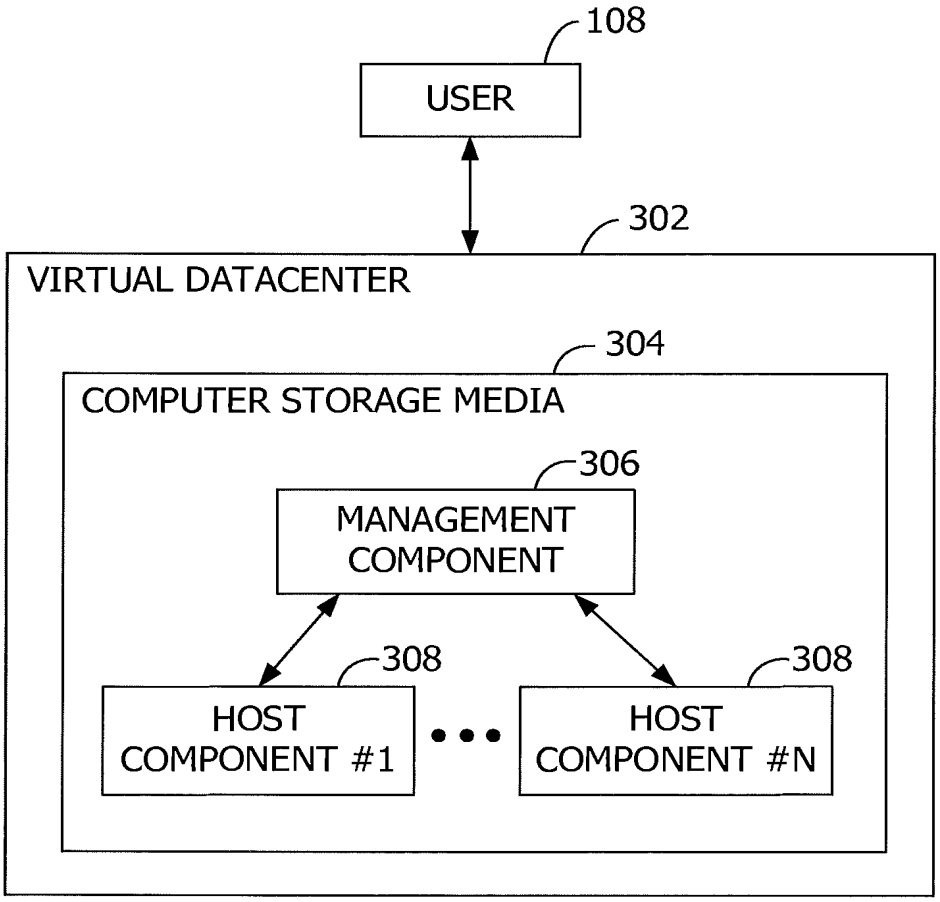
108
USER
302
VIRTUAL DATACENTER
304
COMPUTER STORAGE MEDIA
306
MANAGEMENT COMPONENT
308
HOST COMPONENT #1
308
HOST COMPONENT #N

FIG. 3

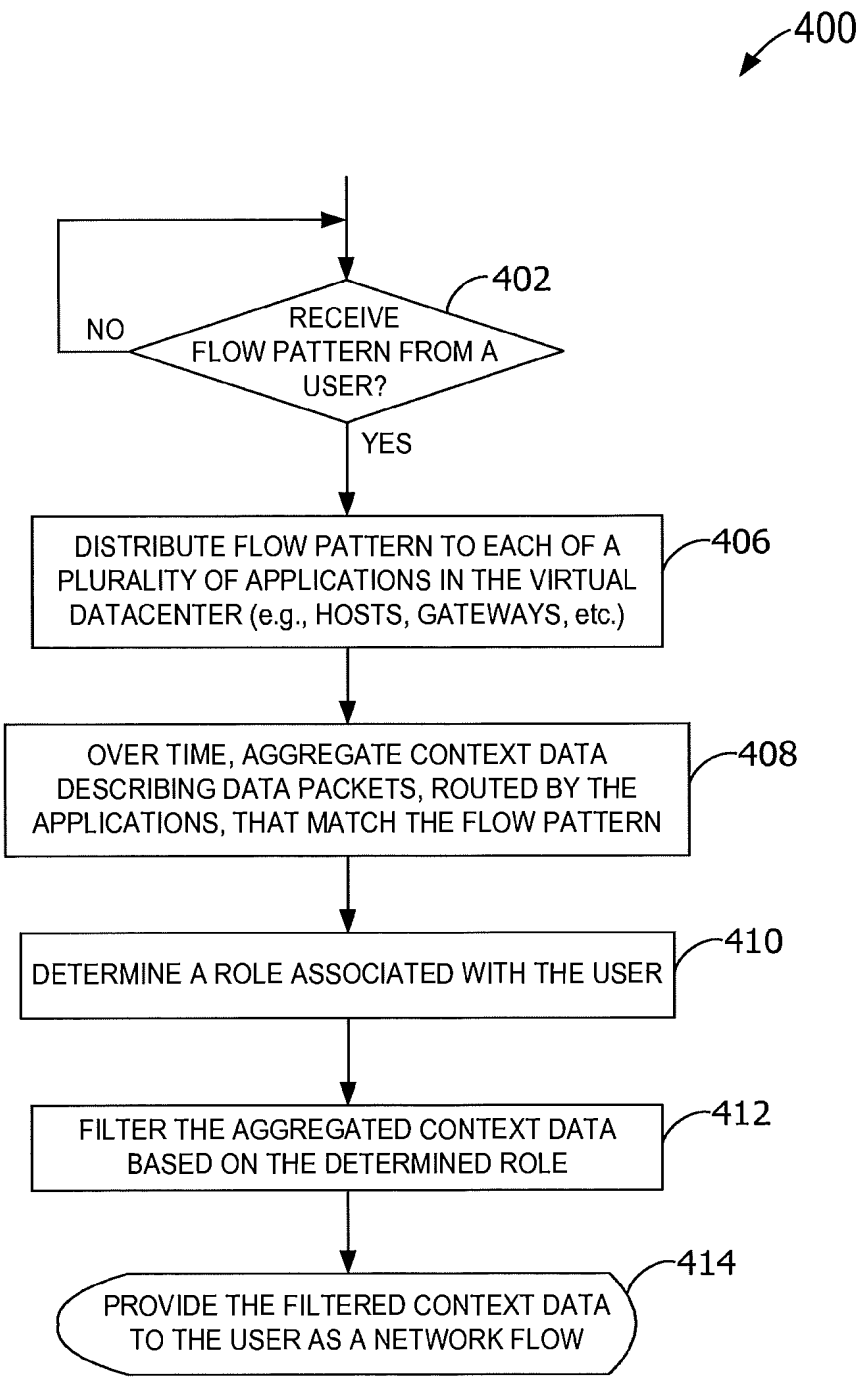
400
402
RECEIVE FLOW PATTERN FROM A USER?
NO
YES
DISTRIBUTE FLOW PATTERN TO EACH OF A PLURALITY OF APPLICATIONS IN THE VIRTUAL DATACENTER (e.g., HOSTS, GATEWAYS, etc.)
406
OVER TIME, AGGREGATE CONTEXT DATA DESCRIBING DATA PACKETS, ROUTED BY THE APPLICATIONS, THAT MATCH THE FLOW PATTERN
408
DETERMINE A ROLE ASSOCIATED WITH THE USER
410
FILTER THE AGGREGATED CONTEXT DATA BASED ON THE DETERMINED ROLE
412
PROVIDE THE FILTERED CONTEXT DATA TO THE USER AS A NETWORK FLOW
414

FIG. 4

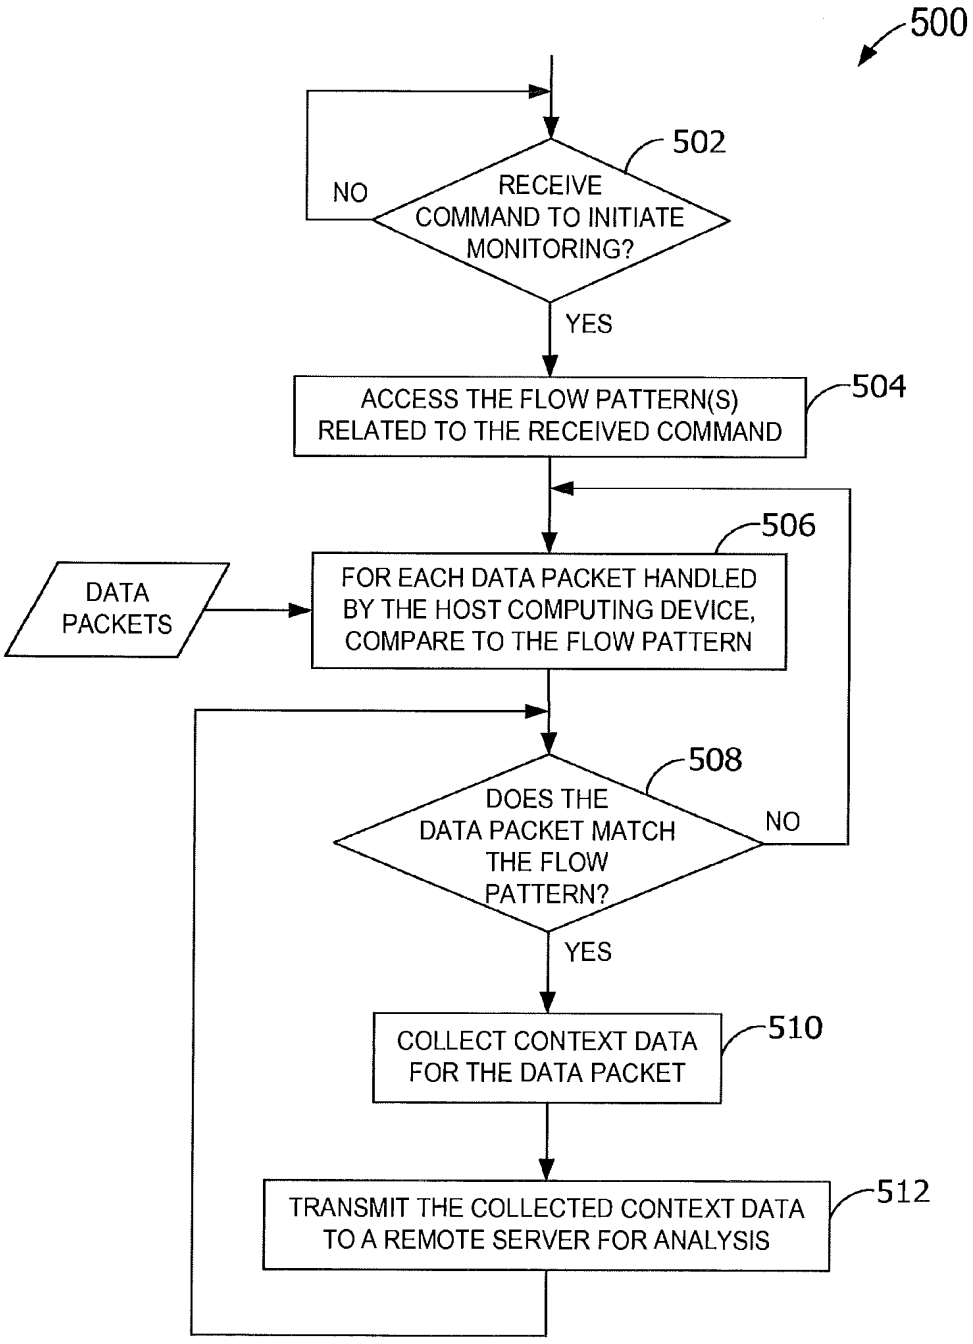
500
502
RECEIVE COMMAND TO INITIATE MONITORING?
NO
YES
504
ACCESS THE FLOW PATTERN(S) RELATED TO THE RECEIVED COMMAND
506
DATA PACKETS
FOR EACH DATA PACKET HANDLED BY THE HOST COMPUTING DEVICE, COMPARE TO THE FLOW PATTERN
508
DOES THE DATA PACKET MATCH THE FLOW PATTERN?
NO
YES
510
COLLECT CONTEXT DATA FOR THE DATA PACKET
512
TRANSMIT THE COLLECTED CONTEXT DATA TO A REMOTE SERVER FOR ANALYSIS

FIG. 5

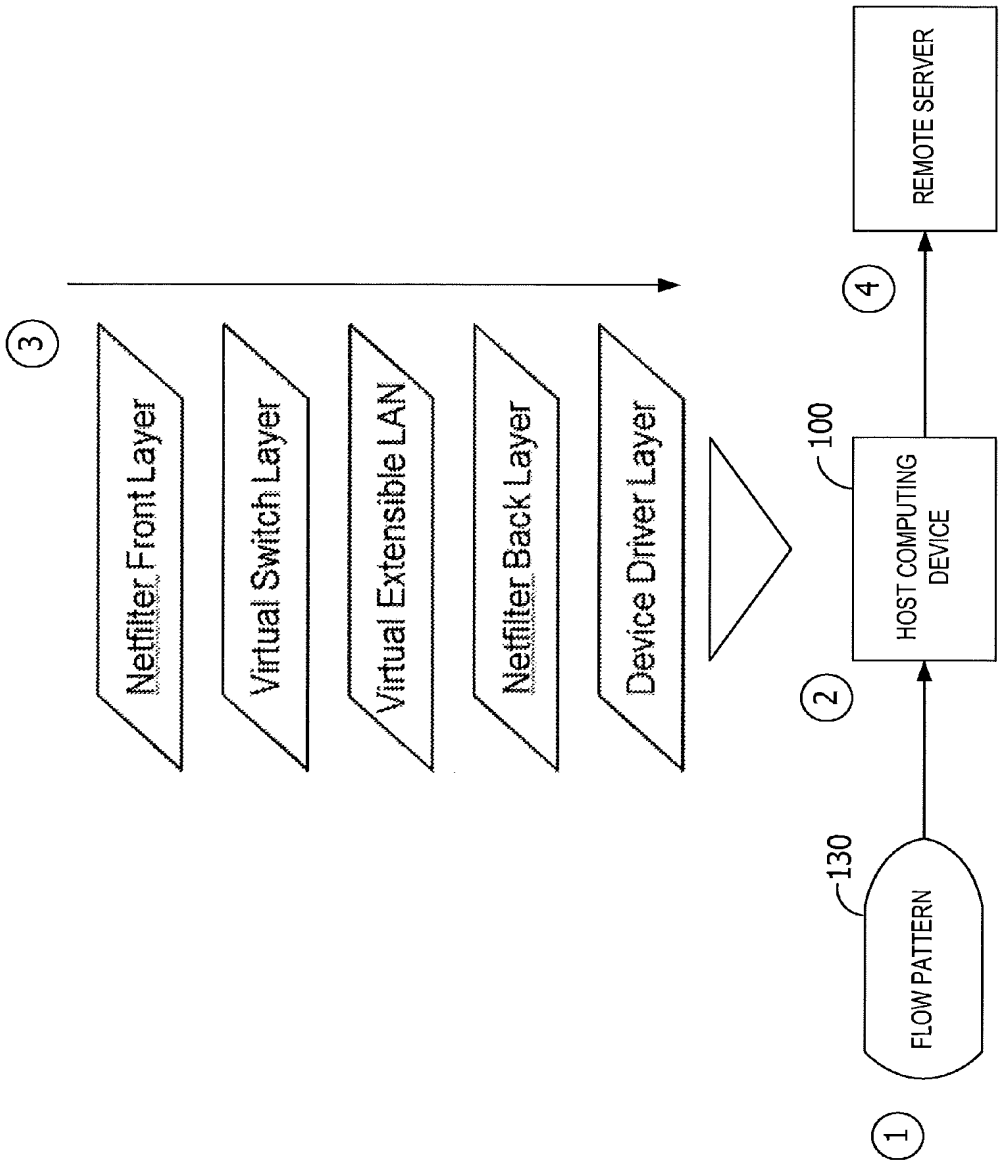
REMOTE SERVER
4
HOST COMPUTING DEVICE
100
2
FLOW PATTERN
130
1
3
Netfilter Front Layer
Virtual Switch Layer
Virtual Extensible LAN
Netfilter Back Layer
Device Driver Layer

FIG. 6

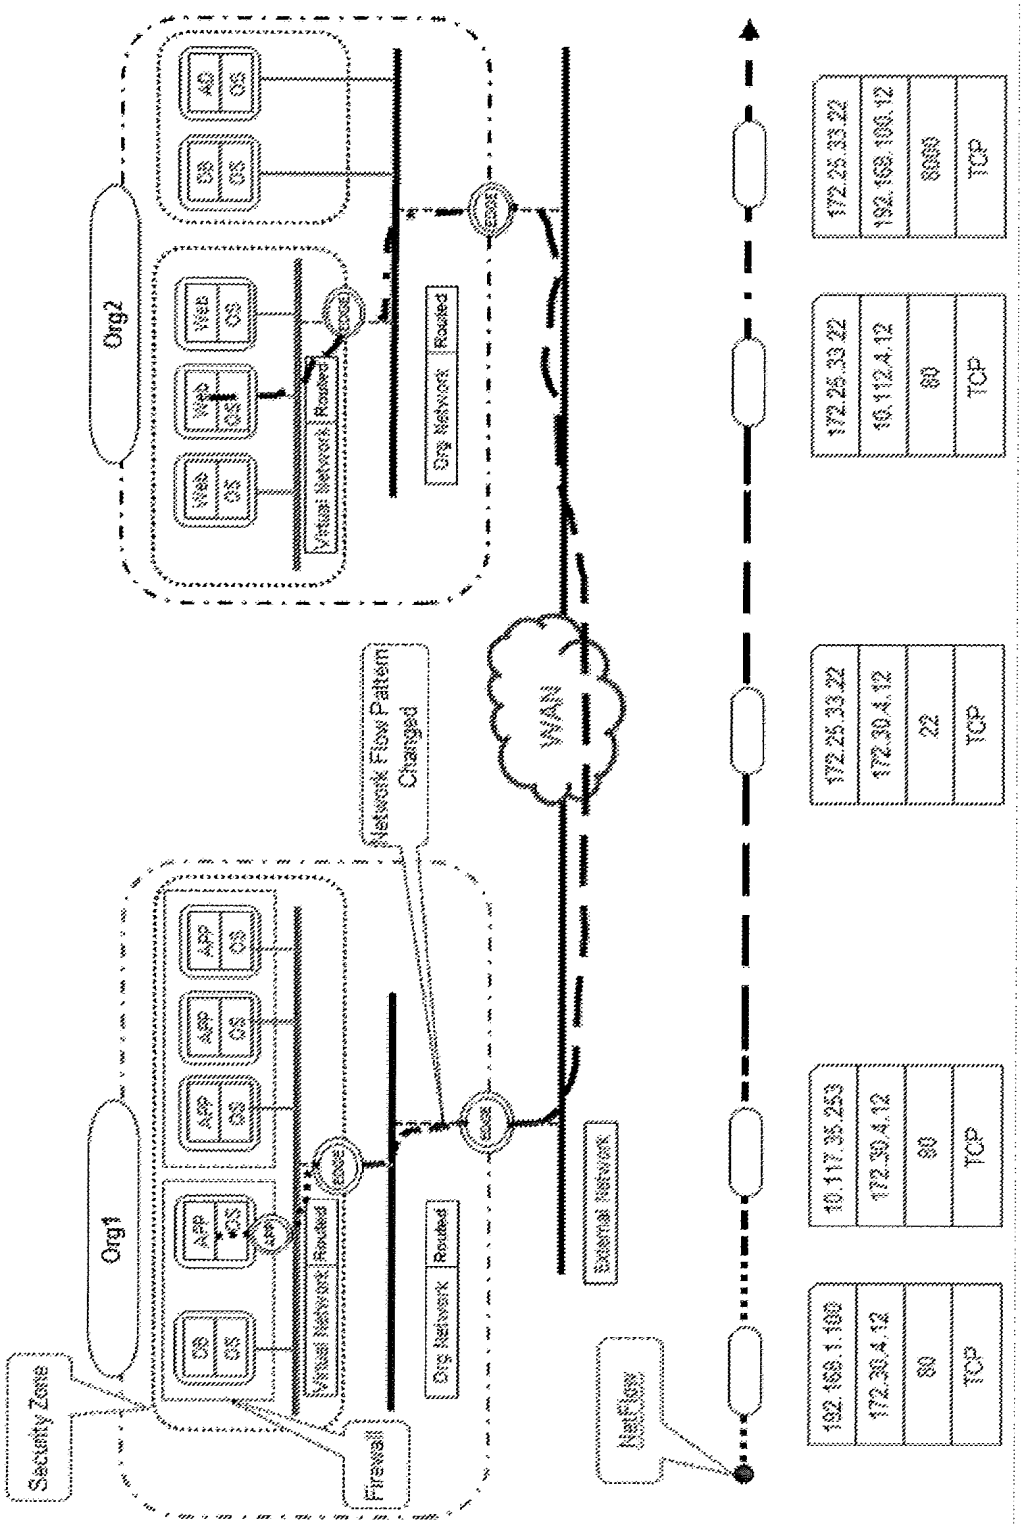
Org2
AD
OS
DB
OS
Web
OS
Web
OS
Web
OS
Edge
Virtual Network
Routed
Org Network
Routed
Edge
Network Flow Pattern Changed
WAN
Org1
APP
OS
APP
OS
APP
OS
APP
OS
DB
OS
Edge
Virtual Network
Routed
Org Network
Routed
Edge
External Network
Security Zone
Firewall
NetFlow
192.168.1.160
172.30.4.12
80
TCP
10.117.35.253
172.30.4.12
80
TCP
172.25.33.22
172.30.4.12
22
TCP
172.25.33.22
10.112.4.12
80
TCP
172.25.33.22
192.168.196.12
8080
TCP

FIG. 7

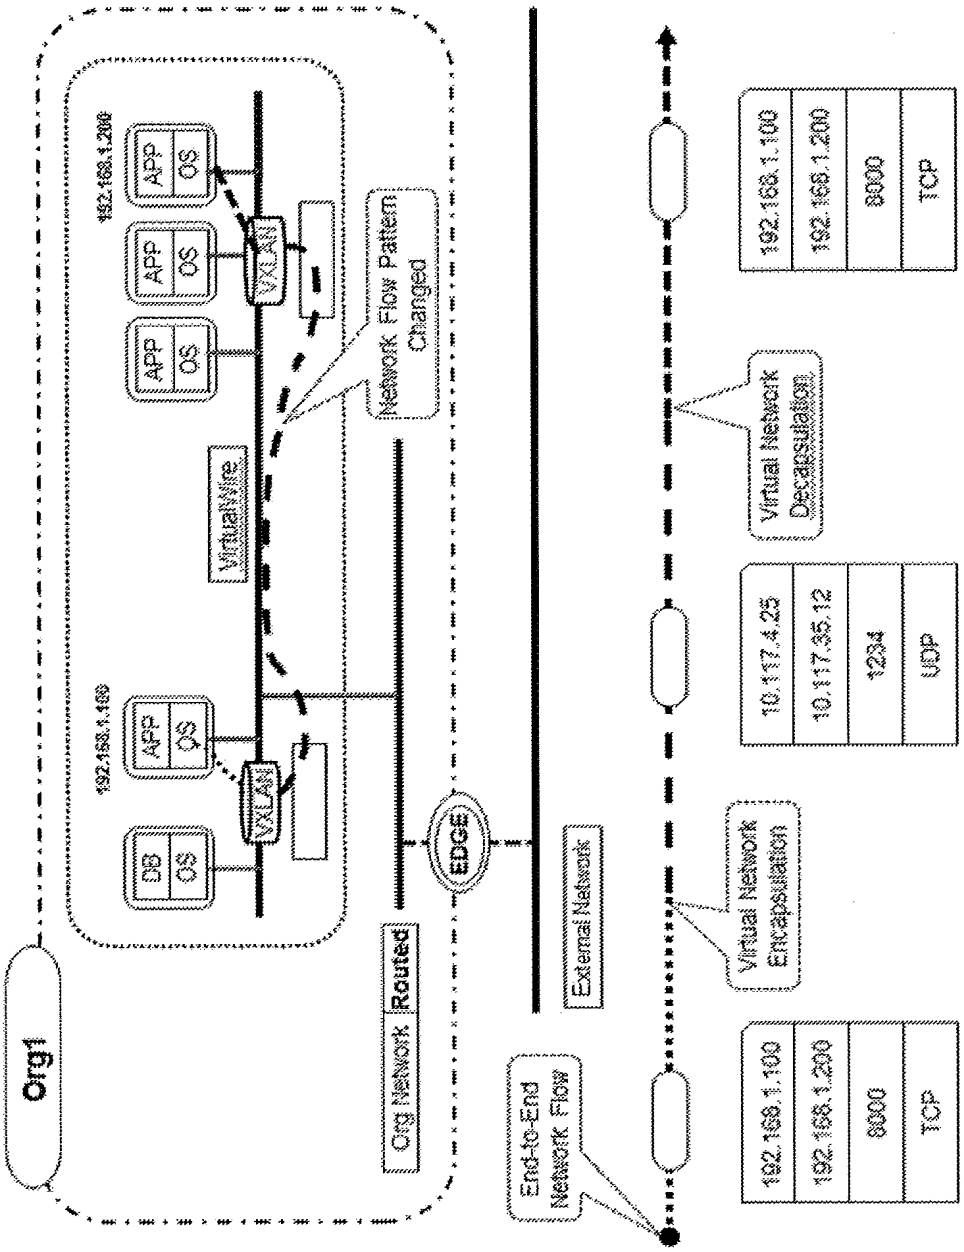
Org1
192.168.1.200
192.168.1.100
APP
OS
DB
VXLAN
VirtualWire
Network Flow Pattern Changed
Org Network
Routed
EDGE
External Network
End-to-End Network Flow
Virtual Network Encapsulation
Virtual Network Decapsulation
192.168.1.100
192.168.1.200
8000
TCP
10.117.4.25
10.117.35.12
1234
UDP

FIG. 8

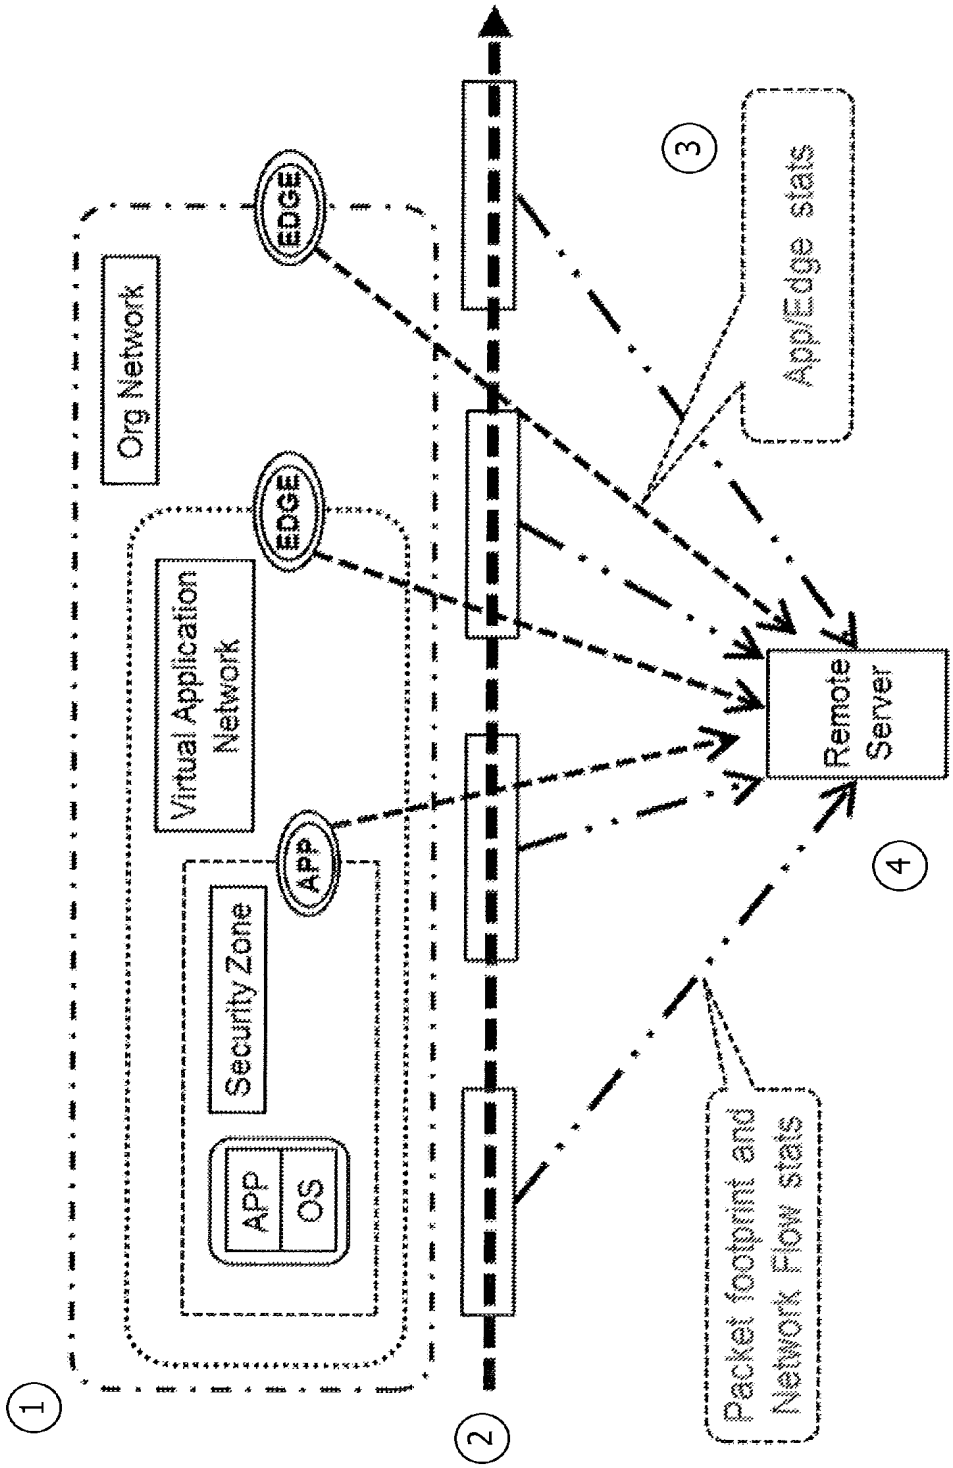
Org Network
EDGE
EDGE
Virtual Application Network
APP
Security Zone
APP
OS
App/Edge stats
Remote Server
Packet footprint and Network Flow stats
1
2
3
4

FIG. 9

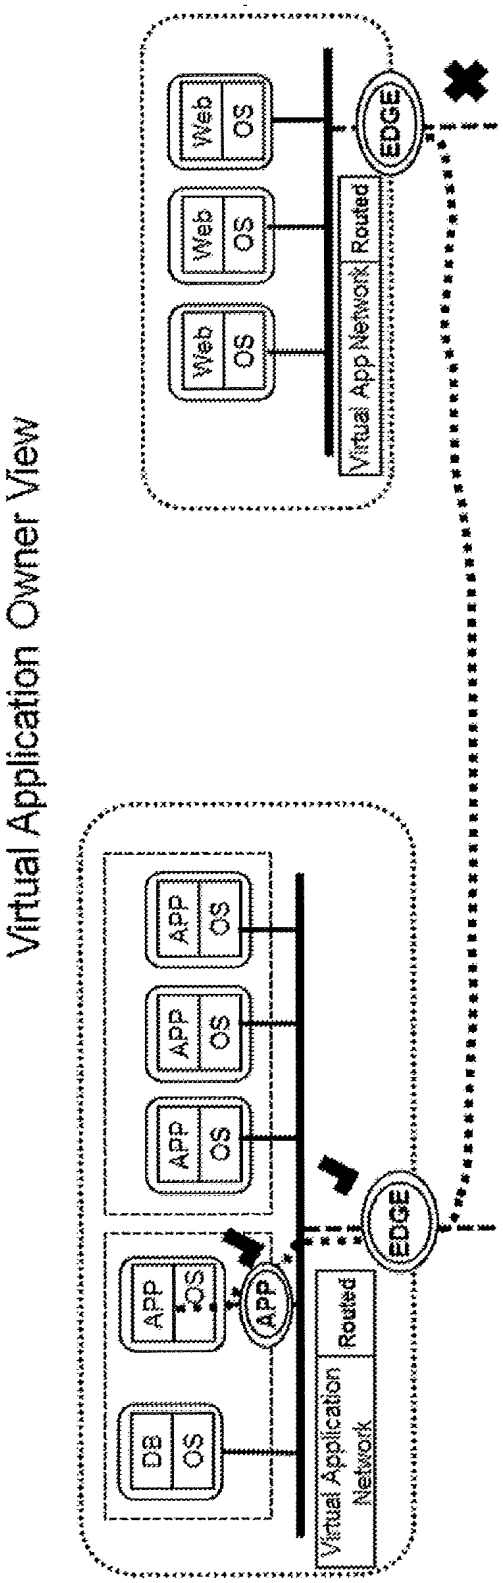
Virtual Application Owner View
Web
OS
Web
OS
Web
OS
Virtual App Network
Routed
EDGE
APP
OS
APP
OS
APP
OS
DB
OS
APP
OS
APP
Virtual Application Network
Routed
EDGE

FIG. 10A

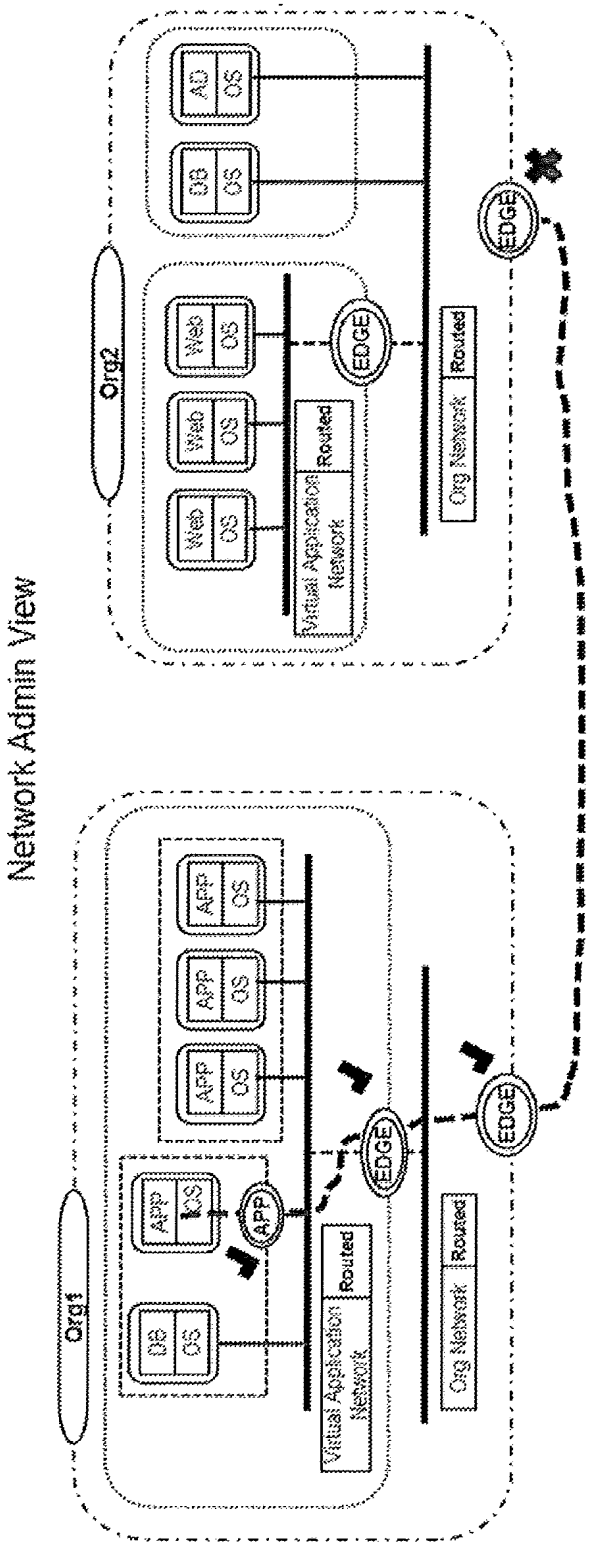
Network Admin View
Org1
Org2
DB
OS
APP
Web
AD
Virtual Application Network
Routed
Org Network
EDGE

FIG. 10B

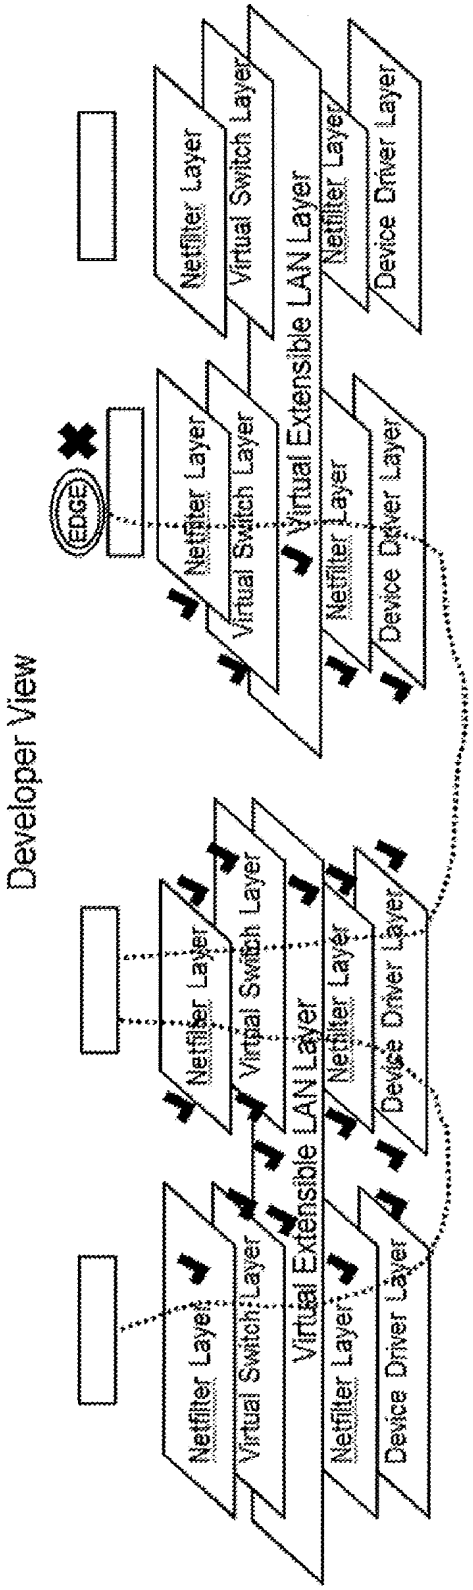
Developer View
Edge
Netfilter Layer
Virtual Switch Layer
Virtual Extensible LAN Layer
Netfilter Layer
Device Driver Layer
Netfilter Layer
Virtual Switch Layer
Virtual Extensible LAN Layer
Netfilter Layer
Device Driver Layer
Netfilter Layer
Virtual Switch Layer
Virtual Extensible LAN Layer
Netfilter Layer
Device Driver Layer

FIG. 10C

VIRTUAL NETWORK FLOW MONITORING

BACKGROUND

[0001] With the increased use of cloud computing and virtualization technologies, virtual datacenter architectures have grown in scale and complexity. In some instances, thousands of virtual machines (VMs) act as hosts, virtual gateways, and other network appliances to implement complex virtual networks. Troubleshooting network problems in such virtualized environments is difficult with existing monitoring systems that are not adapted to operate in cloud computing environments. Further, some virtual datacenters have multiple management and data path layers on top of the physical network infrastructure. For example, a data path may have a network interface card layer, a local area network encapsulation layer, a virtual network appliance layer, an input/output layer, and so on. Further, there may be logical network concepts such as a virtual network, an organization network, zone, endpoint, and so on. Because some of these layers and logical network concepts do not correspond to traditional physical infrastructure, some of the existing monitoring systems are in capable of identifying the source of network problems in virtualized environments.

[0002] Some existing monitoring systems provide packet-based probing by sending probe packets along the data path to identify physical network connectivity issues. Such systems, however, fail to present a high-level logical view to users, fail to provide network flow-based monitoring, and fail to provide detailed flow information within hosts. Further, the probe packets typically terminate at virtual gateways within the virtual datacenter due to security policies. In addition, such systems have a high resource cost (e.g., processing, memory, latency, etc.) by collecting all network traffic going through routers and switches.

[0003] Other existing systems focus on recording packet information within a single host, but these systems fail to provide end-to-end network flow monitoring along the entire network data path.

SUMMARY

[0004] One or more embodiments described herein provide end-to-end virtual network flow monitoring in a virtual datacenter having a plurality of virtual machines (VMs). In some embodiments, the virtual datacenter distributes a flow pattern to a plurality of applications managing the VMs. For example, the applications are associated with host computing devices, virtual gateways, and other network applications. The flow pattern describes data packets of interest to a user. Each of the applications monitors data packets routed by the application by, for example, comparing the data packets to the flow pattern. For each of the routed data packets that match the flow pattern, applications collect context data describing the data packet and transmit the collected context data to a remote server. In this manner, the virtual datacenter aggregates, from the plurality of applications routing data packets, context data for the data packets that match the flow pattern. The virtual datacenter filters the aggregated context data based on a role associated with the user. The filtered context data is presented to the user.

[0005] This summary introduces a selection of concepts that are described in more detail below. This summary is not intended to identify essential features, nor to limit in any way the scope of the claimed subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] FIG. **1** is a block diagram of an exemplary host computing device.

[0007] FIG. **2** is a block diagram of virtual machines that are instantiated on a computing device, such as the host computing device shown in FIG. **1**.

[0008] FIG. **3** is a block diagram of components associated with a virtual datacenter.

[0009] FIG. **4** is a flowchart of an exemplary method performed by the virtual datacenter to initiate flow monitoring in virtual networks in the virtual datacenter.

[0010] FIG. **5** is a flowchart of an exemplary method performed by a host computing device, or other entity managing virtual machines, to monitor data packets and collect context data for selected data packets.

[0011] FIG. **6** is a block diagram illustrating exemplary flow monitoring in a stack in a host computing device.

[0012] FIG. **7** is a block diagram illustrating exemplary end-to-end flow monitoring between two organizations.

[0013] FIG. **8** is a block diagram illustrating exemplary end-to-end flow monitoring with network encapsulation.

[0014] FIG. **9** is a block diagram illustrating exemplary flow monitoring at virtual gateways.

[0015] FIGS. **10**A, **10**B, and **10**C are block diagrams illustrating exemplary filtering of flow information based on user role.

[0016] Corresponding reference characters indicate corresponding parts throughout the drawings.

DETAILED DESCRIPTION

[0017] Embodiments described herein provide a virtual network tracing and monitoring system that identifies and selectively monitors network flows, based on a target flow pattern **130**, end-to-end across a virtual network path in a virtual datacenter **302**. In addition to tracing the network flow within hosts, aspects of the disclosure leverage knowledge of virtual datacenter **302** infrastructure to identify, and enable tracing across, logical network edges or other boundaries (e.g., virtual network edges, organization network edges, zones, etc.) along the virtual network path. In some embodiments, virtual network appliances selectively produce a "footprint" of context data for data packets matching flow pattern **130** (e.g., provided by a user **108**). The context data is aggregated at a remote server for analysis and reporting.

[0018] Aspects of the disclosure further enable different users **108** to troubleshoot and identify the source of network issues in different levels within virtual datacenter **302**. In particular, the aggregated context data is filtered or otherwise mined based on different user roles to generate different virtual network monitoring views for presentation to the users **108**. This enables the different users **108** to identify the source of network issues within virtual datacenter **302** according their goals and knowledge levels.

[0019] Accordingly, aspects of the disclosure trace and selectively collect context data for data packets along the entire virtual network data path, including across multiple processing layers. Further, the tracing may occur periodically, intermittently, and/or on-demand, with a reduced negative impact on performance and traffic flooding relative to existing systems due at least to the selectively monitoring. For example, the network traffic is monitored without generating ping messages. As such, aspects of the disclosure provide a consistent, user-differentiated, one-stop troubleshooting and

BNSDOCID: <US____20140281030A1_I_>

reporting experience to users **108**. For example, different users **108** can troubleshoot network issues by operating on the same management workflow and user interface, without accessing lower level infrastructure information.

[0020] An example of a virtualized environment is next described.

[0021] FIG. **1** is a block diagram of an exemplary host computing device **100**. Host computing device **100** includes a processor **102** for executing instructions. In some embodiments, executable instructions are stored in a memory area **104**. Memory area **104** is any device allowing information, such as executable instructions and/or other data, to be stored and retrieved. For example, memory area **104** may include one or more random access memory (RAM) modules, flash memory modules, hard disks, solid state disks, and/or optical disks.

[0022] As described further herein, memory **104** stores at least one of flow pattern **130** for use in monitoring data packets.

[0023] Host computing device **100** may include a user interface device **110** for receiving data from user **108** and/or for presenting data to user **108**. User **108** may interact indirectly with host computing device **100** via another computing device such as VMware's vCenter Server or other management device. User interface device **110** may include, for example, a keyboard, a pointing device, a mouse, a stylus, a touch sensitive panel (e.g., a touch pad or a touch screen), a gyroscope, an accelerometer, a position detector, and/or an audio input device. In some embodiments, user interface device **110** operates to receive data from user **108**, while another device (e.g., a presentation device) operates to present data to user **108**. In other embodiments, user interface device **110** has a single component, such as a touch screen, that functions to both output data to user **108** and receive data from user **108**. In such embodiments, user interface device **110** operates as a presentation device for presenting information to user **108**. In such embodiments, user interface device **110** represents any component capable of conveying information to user **108**. For example, user interface device **110** may include, without limitation, a display device (e.g., a liquid crystal display (LCD), organic light emitting diode (OLED) display, or "electronic ink" display) and/or an audio output device (e.g., a speaker or headphones). In some embodiments, user interface device **110** includes an output adapter, such as a video adapter and/or an audio adapter. An output adapter is operatively coupled to processor **102** and configured to be operatively coupled to an output device, such as a display device or an audio output device.

[0024] Host computing device **100** also includes a network communication interface **112**, which enables host computing device **100** to communicate with a remote device (e.g., another computing device) via a communication medium, such as a wired or wireless packet network. For example, host computing device **100** may transmit and/or receive data via network communication interface **112**. User interface device **110** and/or network communication interface **112** may be referred to collectively as an input interface and may be configured to receive information from user **108**.

[0025] Host computing device **100** further includes a storage interface **116** that enables host computing device **100** to communicate with one or more datastores, which store virtual disk images, software applications, and/or any other data suitable for use with the methods described herein. In exemplary embodiments, storage interface **116** couples host computing device **100** to a storage area network (SAN) (e.g., a Fibre Channel network) and/or to a network-attached storage (NAS) system (e.g., via a packet network). The storage interface **116** may be integrated with network communication interface **112**.

[0026] FIG. **2** depicts a block diagram of virtual machines 235_1, 235_2 . . . 235_N that are instantiated on host computing device **100**. Host computing device **100** includes a hardware platform **205**, such as an x86 architecture platform. Hardware platform **205** may include processor **102**, memory area **104**, network communication interface **112**, user interface device **110**, and other input/output (I/O) devices, such as a presentation device **106** (shown in FIG. **1**). A virtualization software layer, also referred to hereinafter as a hypervisor **210**, is installed on top of hardware platform **205**.

[0027] The virtualization software layer supports a virtual machine execution space **230** within which multiple virtual machines (VMs 235_1-235_N) may be concurrently instantiated and executed. Hypervisor **210** includes a device driver layer **215**, and maps physical resources of hardware platform **205** (e.g., processor **102**, memory area **104**, network communication interface **112**, and/or user interface device **110**) to "virtual" resources of each of VMs 235_1-235_N such that each of VMs 235_1-235_N has its own virtual hardware platform (e.g., a corresponding one of virtual hardware platforms 240_1-240_N), each virtual hardware platform having its own emulated hardware (such as a processor **245**, a memory **250**, a network communication interface **255**, a user interface device **260** and other emulated I/O devices in VM 235_1). Hypervisor **210** may manage (e.g., monitor, initiate, and/or terminate) execution of VMs 235_1-235_N according to policies associated with hypervisor **210**, such as a policy specifying that VMs 235_1-235_N are to be automatically restarted upon unexpected termination and/or upon initialization of hypervisor **210**. In addition, or alternatively, hypervisor **210** may manage execution VMs 235_1-235_N based on requests received from a device other than host computing device **100**. For example, hypervisor **210** may receive an execution instruction specifying the initiation of execution of first VM 235_1 from a management device via network communication interface **112** and execute the execution instruction to initiate execution of first VM 235_1.

[0028] In some embodiments, memory **250** in first virtual hardware platform 240_1 includes a virtual disk that is associated with or "mapped to" one or more virtual disk images stored on a disk (e.g., a hard disk or solid state disk) of host computing device **100**. The virtual disk image represents a file system (e.g., a hierarchy of directories and files) used by first VM 235_1 in a single file or in a plurality of files, each of which includes a portion of the file system. In addition, or alternatively, virtual disk images may be stored on one or more remote computing devices **100**, such as in a storage area network (SAN) configuration. In such embodiments, any quantity of virtual disk images may be stored by the remote computing devices **100**.

[0029] Device driver layer **215** includes, for example, a communication interface driver **220** that interacts with network communication interface **112** to receive and transmit data from, for example, a local area network (LAN) connected to host computing device **100**. Communication interface driver **220** also includes a virtual bridge **225** that simulates the broadcasting of data packets in a physical network received from one communication interface (e.g., network communication interface **112**) to other communication interfaces (e.g., the virtual communication interfaces of VMs

235_1-235_N). Each virtual communication interface for each VM 235_1-235_N, such as network communication interface **255** for first VM 235_1, may be assigned a unique virtual Media Access Control (MAC) address that enables virtual bridge **225** to simulate the forwarding of incoming data packets from network communication interface **112**. In an embodiment, network communication interface **112** is an Ethernet adapter that is configured in "promiscuous mode" such that all Ethernet packets that it receives (rather than just Ethernet packets addressed to its own physical MAC address) are passed to virtual bridge **225**, which, in turn, is able to further forward the Ethernet packets to VMs 235_1-235_N. This configuration enables an Ethernet packet that has a virtual MAC address as its destination address to properly reach the VM in host computing device **100** with a virtual communication interface that corresponds to such virtual MAC address.

[0030] Virtual hardware platform 240_1 may function as an equivalent of a standard x86 hardware architecture such that any x86-compatible desktop operating system (e.g., Microsoft WINDOWS brand operating system, LINUX brand operating system, SOLARIS brand operating system, NETWARE, or FREEBSD) may be installed as guest operating system (OS) 265 in order to execute applications **270** for an instantiated VM, such as first VM 235_1. Virtual hardware platforms 240_1-240_N may be considered to be part of virtual machine monitors (VMM) 275_1-275_N that implement virtual system support to coordinate operations between hypervisor **210** and corresponding VMs 235_1-235_N. Those with ordinary skill in the art will recognize that the various terms, layers, and categorizations used to describe the virtualization components in FIG. **2** may be referred to differently without departing from their functionality or the spirit or scope of the disclosure. For example, virtual hardware platforms 240_1-240_N may also be considered to be separate from VMMs 275_1-275_N, VMMs 275_1-275_N may be considered to be separate from hypervisor **210**. One example of hypervisor **210** that may be used in an embodiment of the disclosure is included as a component in VMware's ESX brand software, which is commercially available from VMware, Inc.

[0031] In some embodiments, host computing device **100** represents any device executing instructions (e.g., as application programs, operating system functionality, or both) to implement the operations and functionality described herein. For example, one or more of host computing devices **100** execute instructions to implement the operations illustrated in FIG. **4** and/or FIG. **5**. Host computing device **100** may include any computing device or processing unit. For example, host computing device **100** may represent a group of processing units or other computing devices, such as in a cloud computing configuration.

[0032] Processor **102** includes any quantity of processing units, and is programmed to execute computer-executable instructions for implementing aspects of the disclosure. The instructions may be performed by processor **102** or by multiple processors executing within host computing device **100**, or performed by a processor external to host computing device **100** (e.g., another processor available within virtual datacenter **302**). In some embodiments, processor **102** is programmed to execute instructions such as those illustrated in the figures.

[0033] Memory area **104** includes any quantity of computer-readable media associated with or accessible by host computing device **100** or other entity within virtual datacenter **302**. Memory area **104**, or portions thereof, may be internal to host computing device **100**, external to host computing device **100**, or both.

[0034] FIG. **3** is a block diagram of components associated with virtual datacenter **302**. In the example of FIG. **3**, user **108** interacts with virtual datacenter **302**. User **108** may be classified, categorized, or typed based on a role, employment title, and/or interests of user **108**. For example, user **108** includes, but is not limited to, an administrator of virtual datacenter **302**, an owner of a virtual application, a developer, and/or other user involved with virtual datacenter **302**. In some embodiments, the amount and content of flow monitoring and/or tracing data (e.g., collected context data describing data packets) is dependent on the type, category, role, or classification of user **108**.

[0035] As shown in FIG. **3**, virtual datacenter **302** includes one or more computer storage media **304**, such as memory area **104** in some embodiments. Computer storage media **304** store components that, when executed by a processor such as processor **102**, perform various operations. The components illustrated in FIG. **3** include a management component **306** communicating with each of a plurality of host components **308**, such as host component #1 through host component #N.

[0036] The functionality associated with management component **306** is described below with reference to FIG. **4**. In some embodiments, this functionality is implemented by a cloud operating system, a cloud application, a virtual datacenter operating system, a virtual datacenter application, and/or other logic or code supporting an environment having a plurality of VMs executing therein. An example of a cloud operating system includes VMware's vSphere brand software, while an example of a virtual datacenter operation system and/or application includes VMware's vCloud Director software.

[0037] Further, the functionality associated with host components **308** is described below with reference to FIG. **5**. In some embodiments, this functionality is implemented by a virtualization layer such as hypervisor **210**. For example, each of host components **308** manages a plurality of VMs.

[0038] FIG. **4** is a flowchart of an exemplary method **400** performed by virtual datacenter **302** to initiate flow monitoring in virtual networks in virtual datacenter **302**. For example, one or more of the operations illustrated in FIG. **4** may be performed by management component **306** illustrated in FIG. **3**. However, while method **400** is described with reference to execution by virtual datacenter **302** (shown in FIG. **3**), it is contemplated that method **400** may be performed by any computing device, application, and/or other entity. For example, method **400** may be performed by host computing device **100**.

[0039] At **402**, virtual datacenter **302** checks whether flow pattern **130** has been received from user **108**. For example, flow pattern **130** may be received or identified as part of a request from user **108** to initiate virtual network flow monitoring. Flow pattern **130** describes characteristics of data packets of interest to user **108**. The characteristics relate to the content and/or structure of the data packets. In some embodiments, an exemplary flow pattern **130** includes a source address and a destination address. Flow pattern **130** may include other data such as user-defined keys, a protocol type, and/or other data. Exemplary protocol types include, but are not limited to, a transmission control protocol (TCP) and/or a user datagram protocol (UDP). In some embodiments, flow pattern **130** identifies a source Internet Protocol (IP) address,

BNSDOCID: <US____20140281030A1_I_>
CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 565

a destination IP address, a source port for TCP/UDP, a destination port for TCP/UDP, and an IP protocol. While flow pattern **130** may be represented by any data structure, an example of a data structure includes the following 5-tuple: <sourceip_address, destination_ip_address, source port, destination port, protocol_type>.

[0040] Flow pattern **130** may also include rules or other operations or tests to be performed on the data packets when evaluating whether the data packets match flow pattern **130**. For example, as an alternative or addition to analyzing the source address and destination address of each of the data packets, a hash of each data packet may be calculated and compared to a hash included in flow pattern **130**. In this manner, particular data packets may be selected, as further described below.

[0041] Upon receipt of flow pattern **130**, virtual datacenter **302** stores flow pattern **130** in a memory area associated with virtual datacenter **302**. At **406**, virtual datacenter **302** distributes flow pattern **130** to one or more devices, applications, and/or other entities within virtual datacenter **302** that handle (e.g., route) data packets. In some embodiments, one or more of these entities manage a plurality of VMs as part of at least one virtual network. For example, virtual datacenter **302** may distribute flow pattern **130** to one or more of host components **308** illustrated in FIG. **3**.

[0042] In some embodiments, virtual datacenter **302** transmits flow pattern **130** to each of host computing devices **100** and virtual gateways within virtual datacenter **302**. In another example, only particular entities within virtual datacenter **302** are selectively targeted (e.g., a subset of the available host computing devices **100**, applications, and/or virtual gateways). In some embodiments, a plurality of applications operates as, or otherwise implements some functions of, host computing device **100** and/or a virtual gateway.

[0043] At **408**, over time, virtual datacenter **302** aggregates context data describing data packets that match the distributed flow pattern **130** from entities (e.g., host components **308**) within virtual datacenter **302**. For example, the context data is received from the entities in virtual datacenter **302** that handle the data packets, have received flow pattern **130**, and have found data packets that match flow pattern **130**. The context data may be aggregated across virtual networks within virtual datacenter **302** and/or across virtual datacenters **302**. For example, the context data may be aggregated across one or more logical network boundaries. The context data relating to a single packet may be referred to as a footprint record.

[0044] At **410**, virtual datacenter **302** determines a role of user **108**. The role may be specified in, for example, the request received from user **108**. The role may also be derived from past requests from user **108**, feedback from user **108** resulting from past requests, default values, and/or other factors.

[0045] At **412**, virtual datacenter **302** filters the aggregated context data based on the role of user **108**. For example, the context data may be reduced to a level of detail associated with the determined role of user **108**. At **414**, virtual datacenter **302** provides the filtered context data to user **108**. In some embodiments, virtual datacenter **302** displays the filtered context data to user **108**.

[0046] FIG. **5** is a flowchart of an exemplary method **500** performed by host computing device **100**, or other entity managing virtual machines (e.g., an edge appliance such as a virtual gateway at the edge of a virtual network to monitor data packets and collect context data for selected data packets. For example, one or more of the operations illustrated in FIG. **5** may be performed by one or more of host components **308** illustrated in FIG. **3**. However, while method **500** is described with reference to execution by host computing device **100** (shown in FIG. **2**), it is contemplated that method **500** may be performed by any computing device, application, and/or other entity. For example, method **500** may be performed by a virtual gateway, hypervisor **210**, VMs **235**, applications **270**, and/or guest operating system **265**.

[0047] At **502**, host computing device **100** checks whether a command or request to initiate flow monitoring has been received. For example, the command is received from management component **306** or other entity within virtual datacenter **302**. Upon receipt of the command to initiate flow monitoring, host computing device **100** receives, obtains, or otherwise accesses flow pattern **130** associated with the command at **504**. Host computing device **100** stores flow pattern **130** in memory area **104**. For example, host computing device **100** caches flow pattern **130** into a hash table at the kernel level.

[0048] At **506**, for each data packet handled by host computing device **100** (e.g., routed) within one more virtual networks, host computing device **100** compares the data packet to flow pattern **130**. For example, host computing device **100** compares the source and destination addresses of each data packet to the source and destination addresses specified in flow pattern **130**. At **508**, for each of the data packets, host computing device **100** determines whether the data packet matches flow pattern **130**, or otherwise satisfies criteria associate with flow pattern **130**.

[0049] If host computing device **100** determines that the data packet matches flow pattern **130**, host computing device **100** increases a pattern reference counter (e.g., at the kernel level) and collects context data describing the matching data packet at **510** (e.g., at the user level). For example, host computing device **100** collects the context data by identifying at least one of a forwarding port, any filters applied to the data packet, results of any applied filters, a virtual network encapsulating the data packet, and whether the received data packet has been copied, cloned, and/or dropped.

[0050] In some embodiments, host computing device **100** collects the context data for only those data packets matching flow pattern **130**, out of a plurality of data packets handled by host computing device **100**. As such, host computing device **100** does not collect the context data for each and every data packet received, but collects data for only those data packets matching flow pattern **130**. In this manner, the cost of flow monitoring, in terms of latency and/or resource consumption (e.g., processing, memory, etc.), is reduced relative to merely sending a report on every single routed data packet whether or not the data packets are of interest to user **108**.

[0051] At **512**, host computing device **100** transmits the collected context data, for each matching data packet, to a reporting device, server, or other entity. In some embodiments, each of host computing devices **100** within virtual datacenter **302** has access to the reporting device to enable the aggregate of the context data on the reporting device.

[0052] If the data packet does not match flow pattern **130** at **508**, context data is not collected and transmitted to the reporting device. Rather, processing continues with the next handled data packet at **506**.

[0053] In embodiments, the operations illustrated in FIG. **5** may be performed by a virtual gateway application, operating

as a virtual gateway from a first virtual network to a second virtual network. In such embodiments, the virtual gateway application derives flow table information from flow pattern **130**. As such, the virtual gateway application compares each routed data packet to flow pattern **130** by comparing each routed data packet to the derived flow table information.

[0054] FIG. **6** is a block diagram illustrating exemplary flow monitoring or tracing in a stack in host computing device **100**. In the example of FIG. **6**, the stack includes a netfilter front layer (e.g., an IOChain/DVFilter layer for distributed virtual filtering), a virtual switch layer (e.g., a PortGroup layer), a virtual extensible local area network (LAN) layer, a netfilter back layer (e.g., another IOChain/DVFilter layer for distributed virtual filtering), and a device driver layer (e.g., a NicDriver layer). The virtual extensible LAN, such as the VXLAN brand software from VMware, Inc., represents an encapsulation mechanism that runs between virtual switches to enable VMs **235** to be deployed on, and move between, any servers within virtual datacenter **302**.

[0055] Exemplary operations associated with flow monitoring are next described. While the operations are described as being performed by host computing device **100** in some embodiments, the operations may be performed by any module, device, application, etc. that acts as a bi-directional packet filter between a virtual switch and a device driver to filter network traffic entering the virtual switch from VMs **235** or exiting the virtual switch to VMs **235**. In some embodiments, the bi-directional packet filter is embodied in a software Ethernet switch providing network connectivity within a virtual network environment.

[0056] User **108** defines flow pattern **130** at Operation 1 and flow tracing is enabled at Operation 2. The packet footprint is traced according to flow pattern **130** at Operation 3, and the network flow is analyzed via the data packets. For each matching data packet, host computing device **100** collects the context data described in Table 1 below.

TABLE 1

Exemplary Context Data.

Packet Metadata	Layer 3 5-tuple Metadata
portIDs	List of portIDs that the packet accessed
Distributed virtual filters	List of filters that the packet passed (e.g. invoked)
IOChains	List of function addresses that the packet passed (e.g., called)
<func, result>	Function calls and the results therefrom
<clone/copy>	Clone or copy records of the packet

[0057] Each of the elements of the context data shown in Table 1 may include additional data, such as a timestamp.

[0058] At Operation 4, the collected context data is sent to a remote server. For example, the footprint record for the data packet is sent to the remote server when the data packet is about to be freed. Alternatively or in addition, the collected context data is batched for a plurality of matching data packets and flushed to the remote server periodically and/or intermittently. For example, host computing device **100** may send 32 footprints records at a time. Upon collecting enough footprint records (e.g., satisfying a pre-defined threshold), the remote server determines the flow record.

[0059] In an exemplary implementation, host computing device **100** reserves extra memory space in each data packet when creating, copying and/or cloning the packet. The extra memory space is used to capture context data, and may be referred to as the footprint record. When the data packet is copied or cloned, the footprint record is also copied. In some embodiments, host computing device **100** allocates 64 kilobytes for the footprint record, which includes 8 bytes for the functional address, 4 bytes for the result, 4 bytes for the portID, 16 bytes for the timestamp, and 32 bytes for other metadata.

[0060] FIG. **7** is a block diagram illustrating exemplary end-to-end flow monitoring between two organizations. In this example, a user in one organization (Org1) is accessing web services provided by another organization (Org2). The exemplary network path between Org1 and Org2 is illustrated in FIG. **7** and shown to pass through several zones or edges:

[0061] Source→vAppNetwork1→OrgNetwork→ExternalNetwork→OrgNetwork2→vAppNetwork2→Destination

[0062] In the example of FIG. **7**, "edge" represents an edge gateway and "app" represents an application firewall. The virtual gateways between Org1 and Org2 may be implemented using the vShield Edge and/or the vShield App, both from VMware, Inc. Aspects of the disclosure access virtual gateway network profiles for the virtual gateways to understand the virtual network flow patterns. For example, the virtual gateway network profiles may be obtained from a cloud operating system and/or cloud application executed by virtual datacenter **302**. In the example of FIG. **7**, five virtual network flow patterns are illustrated. Flow pattern **130** is given by the application in Org1 with <SRC: 192.168.1.100, DST: 172.30.4.12, DST Port: 80, Protocol: TCP>. Host computing device **100** communicates with the other virtual network appliances shown in FIG. **7** to understand the routing policies on the virtual gateways and then deduce the remaining virtual network flow patterns. The remaining virtual network flow patterns represent the flow patterns (e.g., source address, destination address, port number) along the entire network path. In the example of FIG. **7**, the virtual network flow patterns change several times as the traffic passes through several routers in the network path (e.g., from 192.168.1.100 to 172.25.33.22).

[0063] FIG. **8** is a block diagram illustrating exemplary end-to-end flow monitoring with network encapsulation. In the example of FIG. **8**, network traffic with flow pattern **130** represented by <SRC: 192.168.1.100, DST: 192.168.1.200, DST Port: 8000, Protocol: TCP> is encapsulated and decapsulated within the same virtual network. The virtual extensible LAN application performing the encapsulation and the virtual extensible LAN application performing the decapsulation execute the operations illustrated in FIG. **5** to enable identification and tracing of data packets matching flow pattern **130**. Aspects of the disclosure are operable with any encapsulation/decapsulation operations, such as those associated with the VXLAN brand software from VMware, Inc.

[0064] FIG. **9** is a block diagram illustrating exemplary flow monitoring at virtual gateways. An application implementing the operations illustrated in FIG. **4** (e.g., executing on host computing device **100**) enables flow monitoring by communicating with virtual datacenter **302** (e.g., a cloud operating system) to obtain a description of the infrastructure within virtual datacenter **302** at Operation 1. For example, the application establishes an authenticated channel with the cloud operating system managing virtual datacenter **302** to obtain a network profile describing the virtual network infrastructure. The network profile represents the virtual network

BNSDOCID: <US____20140281030A1_I_>
CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 567

topology and configuration information that identifies, among other data, host computing devices **100** sitting in the network flow path.

[0065] At Operation 2, the application identifies the hosts involved in a target network flow path, and distributes flow pattern **130** to each of the hosts to initiate flow monitoring. At Operation 3, flow pattern **130** is distributed to virtual gateways and other virtual network edge appliances to initiate flow monitoring. At Operation 4, the virtual network appliances (e.g., hosts, virtual gateways, etc.) send the context data for data packets matching flow pattern **130** to a remote server for analysis and reporting.

[0066] FIGS. **10**A, **10**B, and **10**C are block diagrams illustrating exemplary filtering of flow information based on user role. The flow information represents the aggregated context data for data packets matching flow pattern **130**. In this example, the flow information indicates a network issue and is filtered based on three different user roles: a virtual application owner (e.g., a customer), a network administrator, and a developer. FIG. **10**A illustrates the virtual application owner view, FIG. **10**B illustrates the network administrator view, and FIG. **10**C illustrates the developer view. The checkmarks in the figures indicate successful passing of the data packets, whereas the "X" in the figures indicates where the network issue is occurring.

[0067] The virtual application owner view represents a top level view (e.g., high-level logical view) presenting the network flow from a source virtual network to a destination virtual network. The virtual application owner view indicates to the virtual application owner that the network issue is occurring at the edge of the destination virtual network. The network administrator view represents a middle level view (e.g., infrastructure or stack view) presenting the network flow across the overall virtual network infrastructure. The network administrator view indicates to the network administrator that the network issue is occurring at the edge of the destination organization network. The developer view represents a low level view (e.g., component level) presenting the functional level packet footprints and network flow across the hosts. The developer view indicates to the developer that the network issue is occurring because of a destination edge gateway configuration.

[0068] The different views illustrated in FIGS. **10**A, **10**B, and **10**C may be presented to users **108** by an application executing on the remote server, on a device within virtual datacenter **302**, by the cloud operating system managing virtual datacenter **302**, and/or by another entity.

[0069] In this manner, aspects of the disclosure present differentiated troubleshooting views according to user role. For example, the functional packet tracing path is presented to technical support and developers, while the logical network path is presented to network administrators to enable efficient identification of network issues in a top-down fashion.

Additional Examples

[0070] The following scenarios are merely exemplary and not intended to be limiting in any way.

[0071] In one scenario, the remote server aggregating the context data for data packets matching flow pattern **130** allows third party vendors to develop filtering criteria to create customized views of the aggregated context data. For example, rather than presenting the three different views in FIGS. **10**A, **10**B, and **10**C, other views may be created.

Exemplary Operating Environment

[0072] The operations described herein may be performed by a computer or computing device. The computing devices communicate with each other through an exchange of messages and/or stored data. Communication may occur using any protocol or mechanism over any wired or wireless connection. A computing device may transmit a message as a broadcast message (e.g., to an entire network and/or data bus), a multicast message (e.g., addressed to a plurality of other computing devices), and/or as a plurality of unicast messages, each of which is addressed to an individual computing device. Further, in some embodiments, messages are transmitted using a network protocol that does not guarantee delivery, such as User Datagram Protocol (UDP). Accordingly, when transmitting a message, a computing device may transmit multiple copies of the message, enabling the computing device to reduce the risk of non-delivery.

[0073] Exemplary computer readable media include flash memory drives, digital versatile discs (DVDs), compact discs (CDs), floppy disks, and tape cassettes. By way of example and not limitation, computer readable media comprise computer storage media and communication media. Computer storage media include volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media are tangible and are mutually exclusive to communication media. In some embodiments, computer storage media are implemented in hardware. Exemplary computer storage media include hard disks, flash drives, and other solid-state memory. In contrast, communication media typically embody computer readable instructions, data structures, program modules, or other data in a modulated data signal such as a carrier wave or other transport mechanism and include any information delivery media.

[0074] Although described in connection with an exemplary computing system environment, embodiments of the disclosure are operative with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with aspects of the disclosure include, but are not limited to, mobile computing devices, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, gaming consoles, microprocessor-based systems, set top boxes, programmable consumer electronics, mobile telephones, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

[0075] Embodiments of the disclosure may be described in the general context of computer-executable instructions, such as program modules, executed by one or more computers or other devices. The computer-executable instructions may be organized into one or more computer-executable components or modules. Generally, program modules include, but are not limited to, routines, programs, objects, components, and data structures that perform particular tasks or implement particular abstract data types. Aspects of the disclosure may be implemented with any number and organization of such components or modules. For example, aspects of the disclosure are not limited to the specific computer-executable instructions or the specific components or modules illustrated in the figures and described herein. Other embodiments of the disclosure may include different computer-executable instruc-

tions or components having more or less functionality than illustrated and described herein.

[0076] Aspects of the disclosure transform a general-purpose computer into a special-purpose computing device when programmed to execute the instructions described herein.

[0077] The embodiments illustrated and described herein as well as embodiments not specifically described herein but within the scope of aspects of the invention constitute exemplary means for end-to-end tracing of the one or more data packets through virtual datacenter **302** based on a footprint for each of the one or more data packets.

[0078] At least a portion of the functionality of the various elements illustrated in the figures may be performed by other elements in the figures, or an entity (e.g., processor, web service, server, application program, computing device, etc.) not shown in the figures.

[0079] In some embodiments, the operations illustrated in the figures may be implemented as software instructions encoded on a computer readable medium, in hardware programmed or designed to perform the operations, or both. For example, aspects of the disclosure may be implemented as a system on a chip or other circuitry including a plurality of interconnected, electrically conductive elements.

[0080] The order of execution or performance of the operations in embodiments of the disclosure illustrated and described herein is not essential, unless otherwise specified. That is, the operations may be performed in any order, unless otherwise specified, and embodiments of the disclosure may include additional or fewer operations than those disclosed herein. For example, it is contemplated that executing or performing a particular operation before, contemporaneously with, or after another operation is within the scope of aspects of the disclosure.

[0081] When introducing elements of aspects of the disclosure or the embodiments thereof, the articles "a," "an," "the," and "said" are intended to mean that there are one or more of the elements. The terms "comprising," "including," and "having" are intended to be inclusive and mean that there may be additional elements other than the listed elements. The term "exemplary" is intended to mean "an example of"

[0082] Having described aspects of the disclosure in detail, it will be apparent that modifications and variations are possible without departing from the scope of aspects of the disclosure as defined in the appended claims. As various changes could be made in the above constructions, products, and methods without departing from the scope of aspects of the disclosure, it is intended that all matter contained in the above description and shown in the accompanying drawings shall be interpreted as illustrative and not in a limiting sense.

We claim:

1. A system for end-to-end virtual network flow monitoring in a virtual datacenter having a plurality of virtual machines (VMs), said system comprising:

a memory area associated with the virtual datacenter, said memory area storing a flow pattern including at least a source address and a destination address, the flow pattern being received from a user in a request to perform virtual network flow monitoring; and

a processor programmed to:

distribute the flow pattern stored in the memory area to a plurality of applications in the virtual datacenter, each of the plurality of applications managing a plurality of VMs as part of at least one virtual network;

aggregate, by the virtual datacenter from the plurality of applications, context data for one or more data packets routed by the plurality of applications and matching the flow pattern; and

determine a role associated with the user; and

filter, by the virtual datacenter based on the determined role, the aggregated context data for presentation to the user.

2. The system of claim **1**, wherein the plurality of applications operate as at least one of a host computing device or a virtual gateway.

3. The system of claim **1**, wherein the processor is further programmed to display the aggregated context data to the user.

4. The system of claim **1**, wherein the flow pattern further comprises at least one of a transmission control protocol (TCP) or a user datagram protocol (UDP).

5. The system of claim **1**, wherein the user comprises at least one of a virtual application owner, an administrator, or a developer.

6. The system of claim **1**, wherein the processor aggregates the context data across a logical network boundary.

7. The system of claim **1**, further comprising means for end-to-end tracing of the one or more data packets through the virtual datacenter based on a footprint for each of the one or more data packets.

8. A method comprising:

for each of a plurality of computing devices associated with one or more virtual networks in a virtual datacenter, the virtual datacenter comprising a plurality of virtual machines (VMs):

accessing, by the computing device, a flow pattern providing a source address and a destination address;

receiving a data packet from the one or more virtual networks;

comparing, by the computing device, the received data packet to the accessed flow pattern; and

based on the comparison, collecting data describing the received data packet and transmitting the collected data for analysis to a reporting device accessible by each of the plurality of computing devices associated with the one or more virtual networks in the virtual datacenter.

9. The method of claim **8**, wherein collecting the data comprises identifying at least one of a forwarding port, a filter applied to the data packet, a virtual network encapsulating the data packet, whether the received data packet has been copied, or whether the received data packet has been cloned.

10. The method of claim **8**, wherein collecting the data comprises collecting the data on determining that the received data packet matches the accessed flow pattern.

11. The method of claim **8**, wherein the computing device receives a plurality of the data packets, and wherein collecting the data comprises collecting the data for only those data packets, from the plurality of data packets, that match the accessed flow pattern.

12. The method of claim **8**, wherein receiving the data packet comprises receiving, by a host computing device, a data packet.

13. The method of claim **8**, wherein the computing device comprises a virtual gateway application, and further comprising deriving, by the virtual gateway application, flow table information from the accessed flow pattern.

BNSDOCID: <US____20140281030A1_I_>

14. The method of claim **13**, wherein comparing the received data packet to the accessed flow pattern comprises comparing the received data packet to the derived flow table information.

15. The method of claim **8**, wherein the computing device comprises a host, and wherein collecting the data comprises collecting the data across a network stack within the host.

16. One or more computer-readable storage media storing computer-executable components for performing end-to-end virtual network flow monitoring in a virtual datacenter having a plurality of virtual machine (VMs), the computer-executable components comprising:

a plurality of host components that each, when executed, cause at least one processor to:

receive a data packet;

compare the received data packet to a flow pattern defined by a user, the flow pattern including at least a source address and a destination address; and

based on the comparison, collect context data describing the received data packet; and

a management component in communication with each of the plurality of host components, the management component, when executed, causing at least one processor to:

aggregate the collected context data from one or more of the plurality of host components;

determine a role associated with the user; and

filter, based on the determined role, the aggregated, collected context data for presentation to the user.

17. The computer storage media of claim **16**, wherein the management component further receives, from the user, a request to perform the virtual network flow monitoring, the request including the flow pattern.

18. The computer storage media of claim **16**, wherein the management component further distributes the flow pattern to each of the plurality of host components.

19. The computer storage media of claim **16**, wherein each of the plurality of host components manages a plurality of VMs.

20. The computer storage media of claim **16**, wherein at least one of the plurality of host components represents a virtual gateway from a first virtual network to a second virtual network.

* * * * *

US 20040073655A1

(19) United States

(12) Patent Application Publication (10) Pub. No.: US 2004/0073655 A1

Kan et al. (43) Pub. Date: Apr. 15, 2004

(54) PACKET SEQUENCE NUMBER NETWORK MONITORING SYSTEM

(76) Inventors: **Chao Kan**, Plano, TX (US); **Pierrick Guingo**, Dallas, TX (US); **Emanuele Jones**, Ottawa (CA); **Hubert Ogier**, Paris (FR); **Thierry Labbe**, Dallas, TX (US)

Correspondence Address:
ALCATEL USA
INTELLECTUAL PROPERTY DEPARTMENT
3400 W. PLANO PARKWAY, MS LEGL2
PLANO, TX 75075 (US)

(21) Appl. No.: **10/267,813**

(22) Filed: **Oct. 9, 2002**

Publication Classification

(51) **Int. Cl.**[7] **G06F 15/173**

(52) **U.S. Cl.** **709/224**; 709/230

(57) **ABSTRACT**

A network monitoring system **(60)**. The system comprises storage circuitry **(32)** for storing network packet information, wherein the network packet information includes a predicted identifier. The network monitoring system also comprises at least one monitoring circuit **(36)** coupled to a network **(70)** along which network traffic flows in a form of packets. The at least one monitoring circuit programmed to perform the steps **(44)** of receiving a packet communicated along the network and determining whether the received packet is communicated between a source and destination in a first set of network nodes. Each packet in a sequence of communications between the source and the destination comprises a packet identifier that uniquely identifies the packet from all other communications in a flow between the source and the destination. The at least one monitoring circuit programmed to perform the step of, responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the received packet to the predicted identifier to determine an identifier deviation between the packet identifier and the predicted identifier for identifying an irregularity in the network traffic.

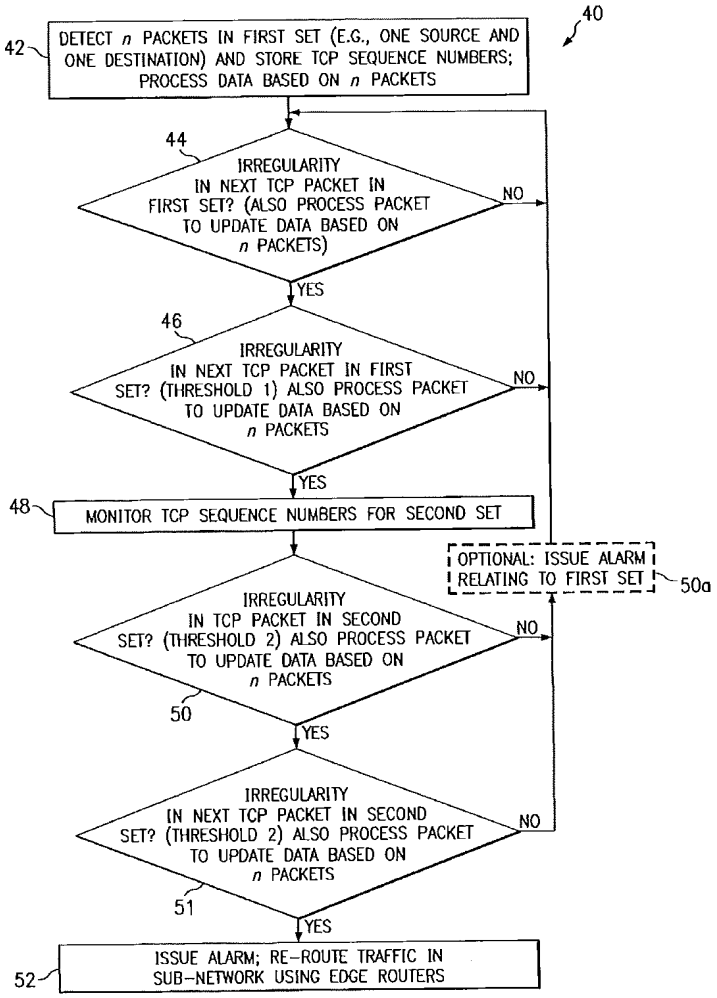

FIG. 1
(PRIOR ART)

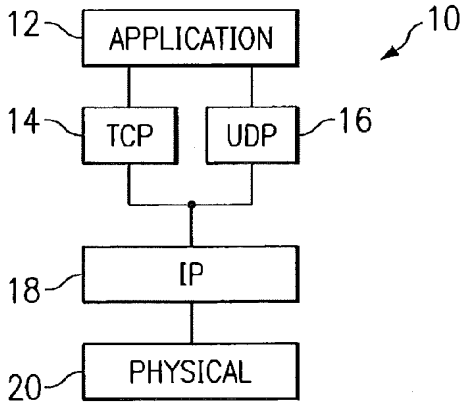

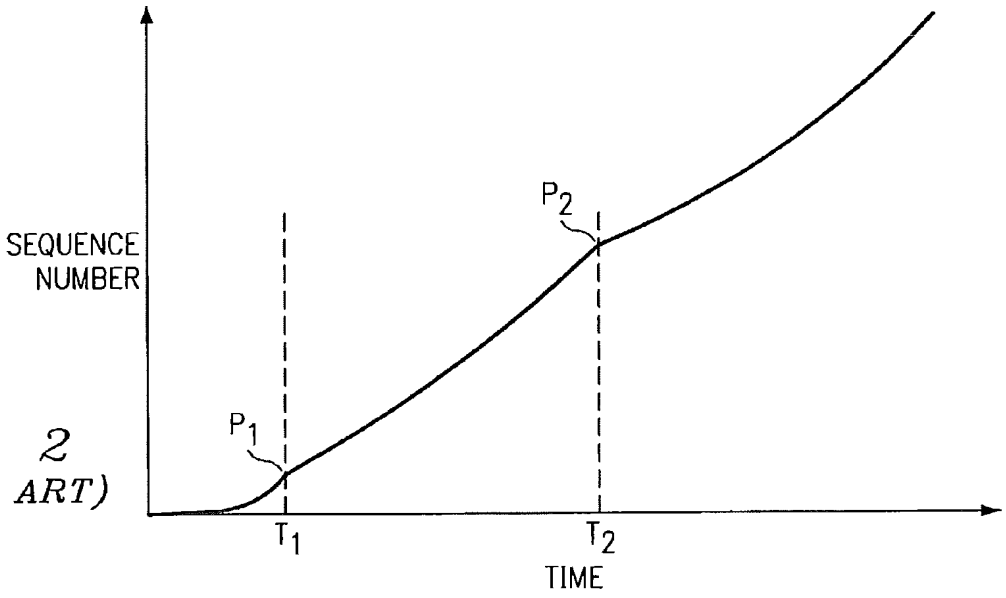

FIG. 2
(PRIOR ART)

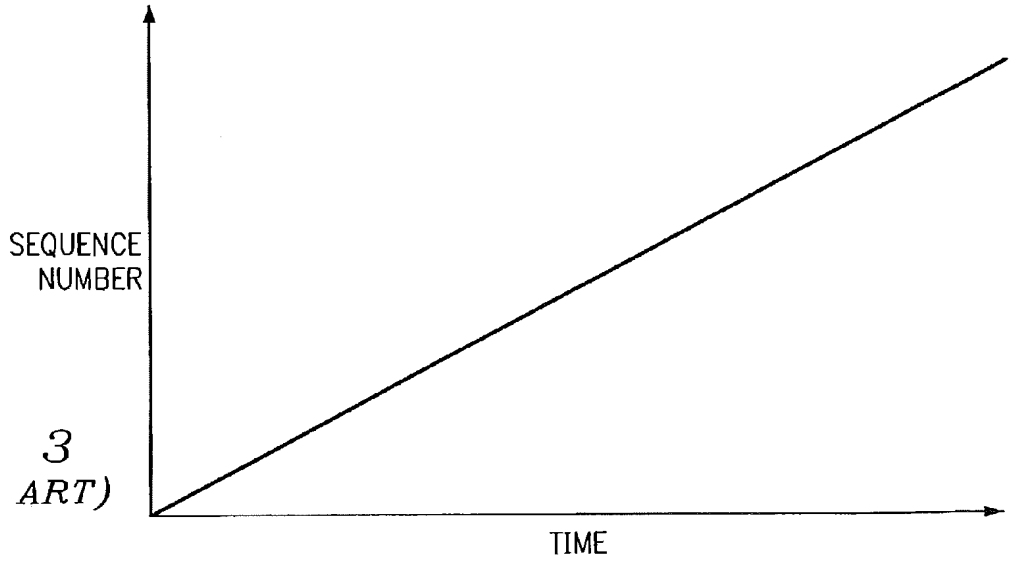

FIG. 3
(PRIOR ART)

FIG. 4
(PRIOR ART)

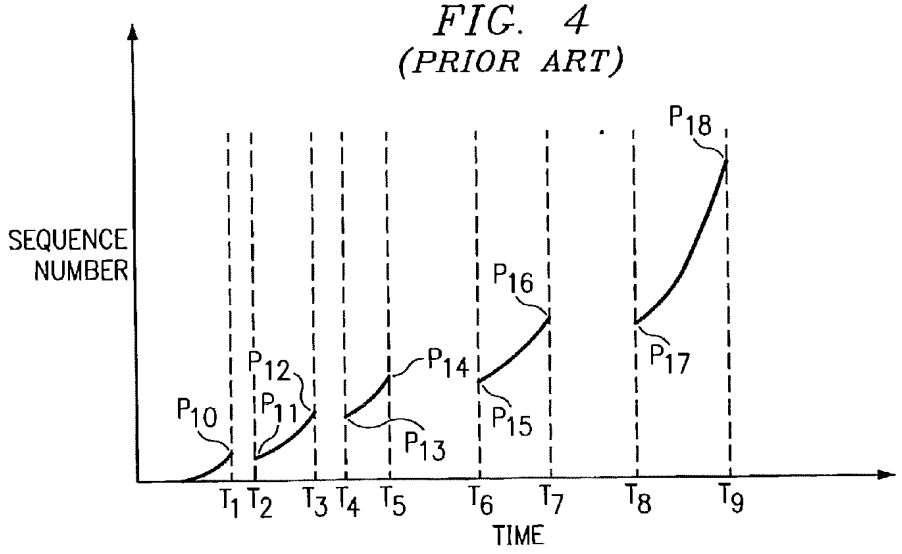

FIG. 5

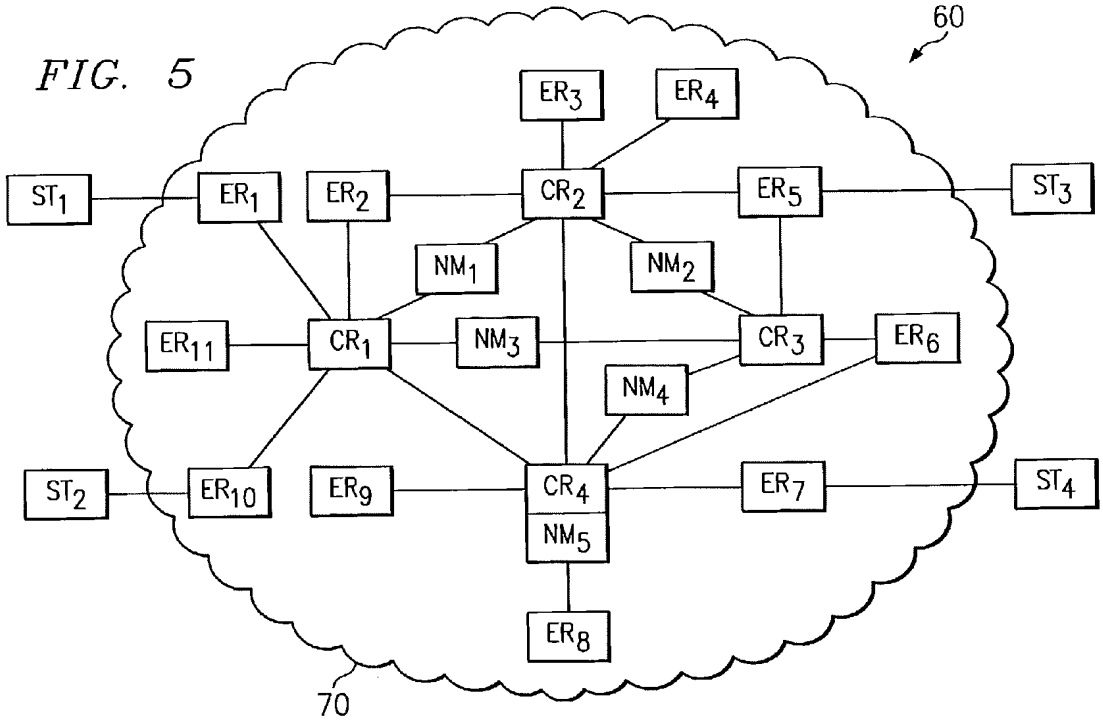

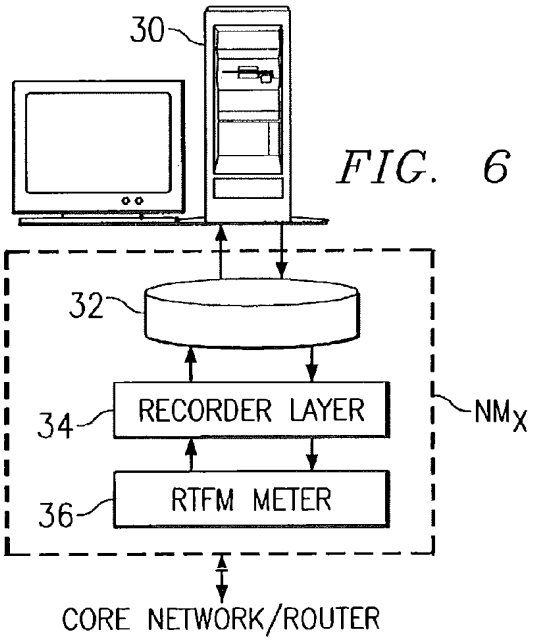
30
32
34
RECORDER LAYER
NMX
36
RTFM METER
CORE NETWORK/ROUTER

FIG. 6

FIG. 8

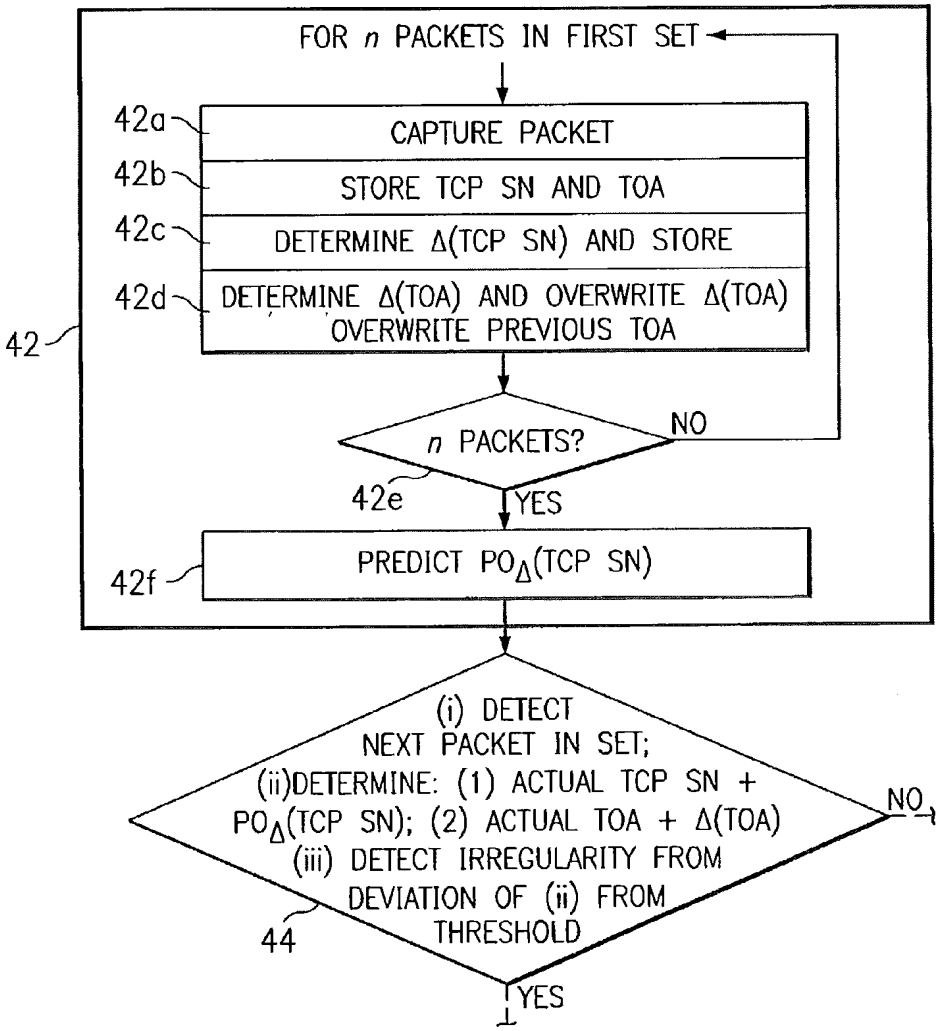
FOR n PACKETS IN FIRST SET
42a
CAPTURE PACKET
42b
STORE TCP SN AND TOA
42c
DETERMINE Δ(TCP SN) AND STORE
42d
DETERMINE Δ(TOA) AND OVERWRITE Δ(TOA) OVERWRITE PREVIOUS TOA
42
n PACKETS?
NO
42e
YES
42f
PREDICT PO$_\Delta$(TCP SN)
(i) DETECT NEXT PACKET IN SET;
(ii)DETERMINE: (1) ACTUAL TCP SN + PO$_\Delta$(TCP SN); (2) ACTUAL TOA + Δ(TOA)
(iii) DETECT IRREGULARITY FROM DEVIATION OF (ii) FROM THRESHOLD
NO
44
YES

FIG. 7

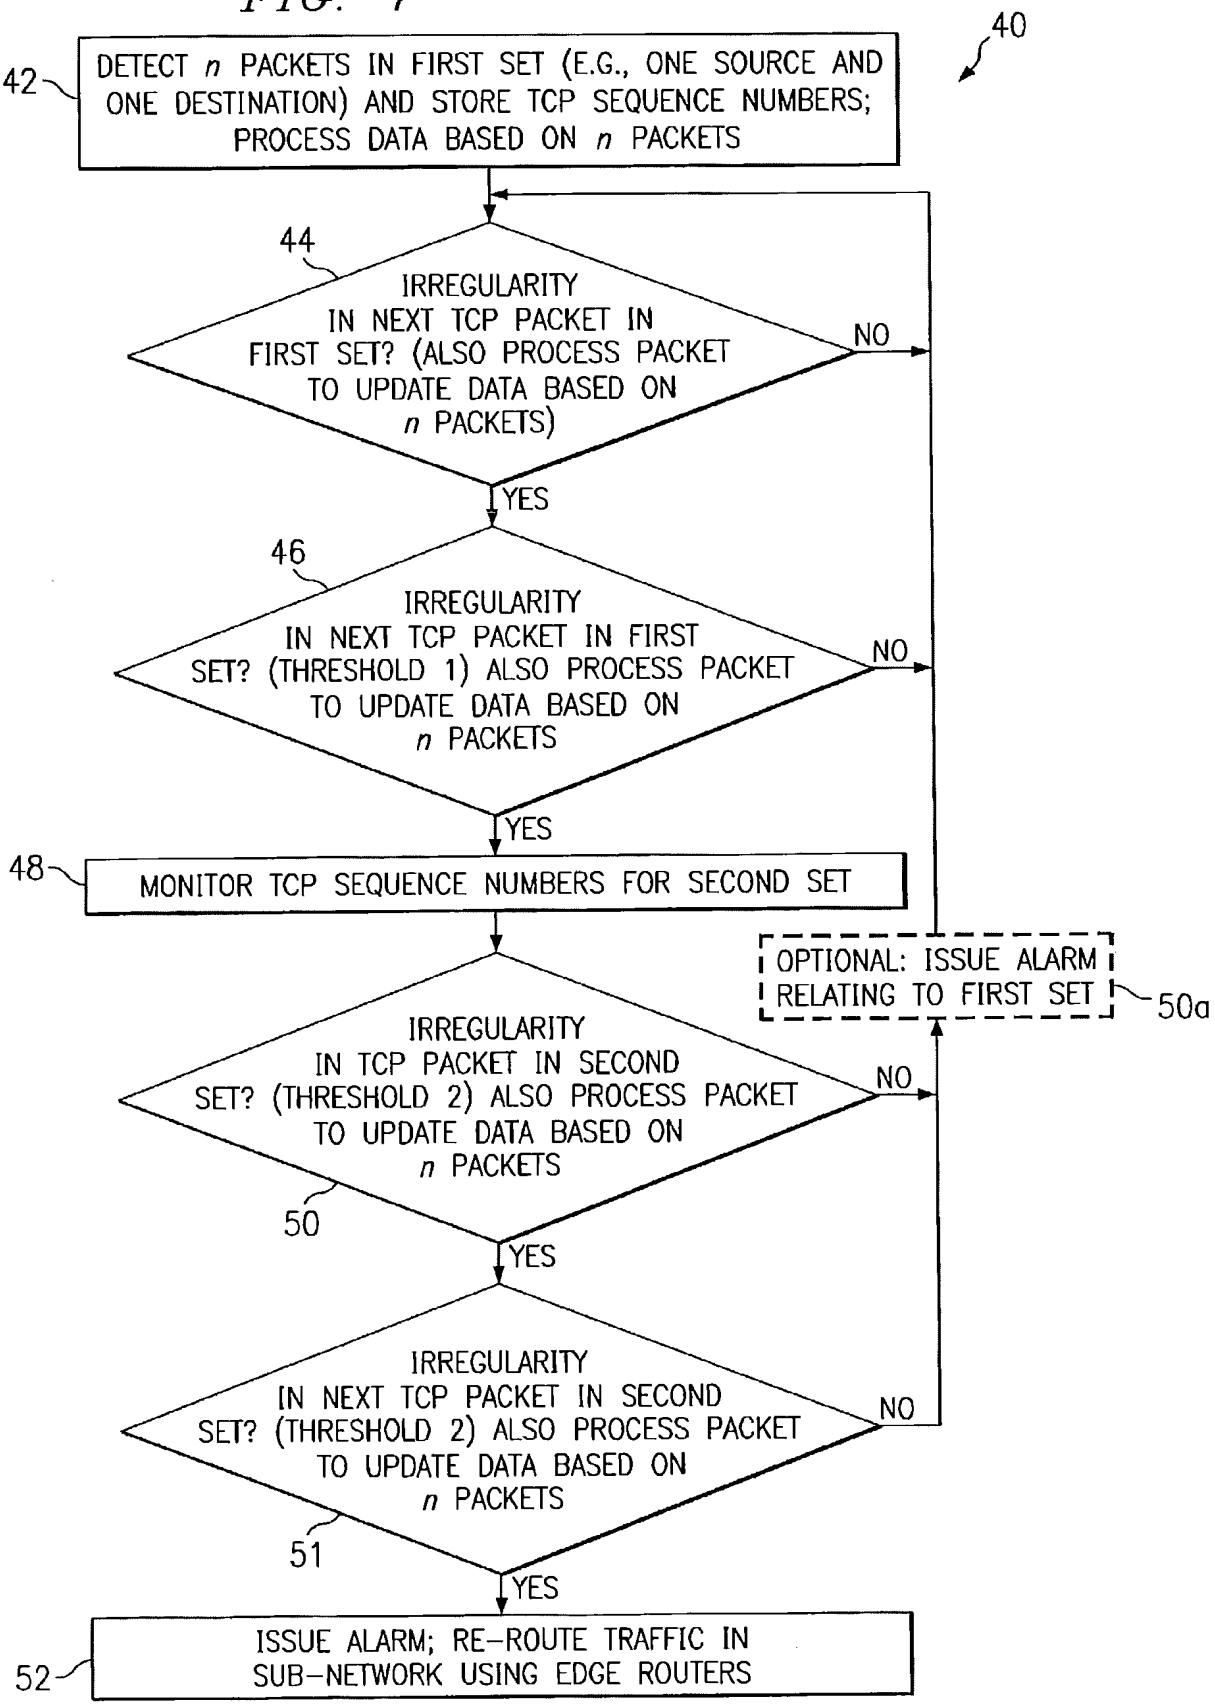

BNSDOCID: <US____20040073655A1_I_>

PACKET SEQUENCE NUMBER NETWORK MONITORING SYSTEM

CROSS-REFERENCES TO RELATED APPLICATIONS

[0001] Not Applicable.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

[0002] Not Applicable.

BACKGROUND OF THE INVENTION

[0003] The present embodiments relate to computer networks and are more particularly directed to a real-time system for monitoring network performance by evaluating packet identifiers such as packet sequence numbers.

[0004] As the number of users and traffic volume continue to grow on the internet and other networks, an essential need has arisen for both the users and network operators to have a set of mechanisms to monitor network performance, detect router availability, and troubleshoot network congestion without introducing additional traffic on the network. This is especially true for Internet Service Providers in order to comply with the Service Level Agreements ("SLAs") that they are now providing to customers. Additionally, such a need is prevalent for the underlying internet protocol ("IP") networks in the global internet. With these mechanisms, the users and network operators also can monitor and query the reliability and availability of the network nodes (e.g., IP routers) and the given internet paths, and the mechanisms also can identify the location of a segment of a given internet path, if any, that has a limitation. Additionally, the internet is evolving towards an advanced architecture that seeks to guarantee the quality of service ("QoS") for real-time applications, such as by putting limits on the upper bound on certain QoS parameters including jitter, throughput, one-way packet delay, and packet loss ratio. Accordingly, tracking QoS performance is also desirable. Given the preceding, various mechanisms for network monitoring have been developed in the art, some of which are described below.

[0005] A first type of network performance analysis mechanism is known in the art as an active mechanism. The active network performance monitoring system sends a series of special test packets or query packets to the underlying networks or routers, and it analyzes the response with a specific purpose. Currently, most tools for active end-to-end QoS monitoring in IP networks are based on the traditional "ping" (i.e., ICMP echo and echo response messages) to measure the roundtrip delay between two hosts. Variations of ping include "Nikhef ping,""fping,""pingplotter,""gnuplotping,""Imeter," and "traceping". Several other tools are variations based on the traditional "traceroute" (which exploits the time-to-live field in the IP packet header), for example, "Nikhef traceroute,""pathchar,""whatroute," and "network probe daemon." A few large-scale research projects are using "ping" to continually and actively monitor the network performance between various points in the global internet. Examples of these types of projects are: (i) the PingER project at Stanford Linear Accelerator Center ("SLAC") uses repeated pings around various Energy Sciences Network ("ESnet") sites; (ii) the Active Measurement Program ("AMP") project by National Laboratory for Applied Network Research ("NLANR") performs pings and traceroutes between various NSF-approved high-performance connection sites; (iii) the Surveyor project attempts one-way packet delay and loss measurements between various sites using the global positioning system ("GPS") for time synchronization; and (iv) the national Internet measurement infrastructure ("NIMI") project measures the performance between various sites using traceroute or transmission control protocol ("TCP") bulk transfer.

[0006] A second type of network performance analysis mechanism is known in the art as a passive mechanism and may be contrasted to the active mechanism discussed above. The passive network performance monitoring system performs its traffic analysis in a non-invasive way with respect to the observed networking environment, namely, the system does not introduce additional test packets to the network. As a result, the system does not affect the performance of the network while doing the measurements and querying. The traditional approach usually involves the following three steps: (i) collection of the entire TCP/IP packet (e.g. traffic sniffers, etc.) or the entire packet header data into a database or databases; (ii) hardware and software for analyzing the collected database(s); and (iii) off-line traffic characterization and modeling. Various examples of passive mechanisms have been implemented by certain entities. For example, the NLANR has been using OCXmon monitors to tap into the light of a fiber interconnection by means of optical splitters, and from the connection the system collects a portion of the packet, namely, the packet header. Specifically, traffic data is collected in an abstract format by extracting and storing the packet header in the database within a preset traffic aggregation period. The packet header traffic data in the database is later analyzed (i.e., an off-line traffic analysis). As another example, Cisco offers a NetFlow capability in its large routers. NetFlow identifies traffic flows based on IP source/destination addresses, protocol ID field, type of service ("TOS") field, and router port. Once identified, statistics can be collected for a traffic flow, and exported via user datagram protocol ("UDP") when the flow expires. Flow statistics may include the flow start/stop times, number of bytes/packets, and all IP header fields. As a final example, Lucent Bell Labs has various research projects in traffic analysis, which are mainly concentrated in collection of TCP/UDP/IP packet headers and off-line traffic analysis, modeling and visualization.

[0007] A third type of network performance analysis mechanism is also a type of passive mechanism in that it does not add traffic to the network, but rather than collecting information directly from a packet it instead collects statistics pertaining to the packet; this system is sometimes referred to as a network or element management system. In this system, instead of separately setting up the packet collection process, IP routers periodically (e.g., in 30 minute intervals) store and collect a set of traffic statistics into a built-in Management Information Base ("MIB") using a simple network management protocol ("SNMP") interface, and those statistics may be retrieved by a network or element manager. Typical traffic statistics might be the number of received/forwarded packets, discarded packets, error packets, port utilization, CPU utilization, buffer utilization, and so forth. As in other passive systems, these statistics are collected for later analysis. When a network congestion or event occurs, the SNMP agent embedded in the IP router

sends a trap message to the SNMP manager, which then indicates an alarm in its graphic user interface.

[0008] While the preceding systems have provided certain levels of information about the behavior of traffic in IP networks, the present inventors have observed that such systems also provide various drawbacks. As an example of a drawback with respect to the active systems, they are required to send special test or PING packets and, thus, these additional packets themselves introduce additional traffic onto the route of the regular data packets. Accordingly, at the time of detecting the network congestion or server availability, the network performance and associated QoS being measured will be negatively impacted by this additional burden. Additionally, the QoS is measured for the test packets instead of the actual data packets, and the QoS for the data packets must be inferred indirectly from the QoS measured for the test packets. As an example of a drawback with respect to the passive system, although many traffic studies have attempted to understand the random behavior or composition of internet traffic, these traffic studies have been off-line analysis of historical data. There are no known prominent research projects attempting traffic analysis and control based on real-time traffic measurement or comprehensive traffic profiling. For example, Lucent's projects reflect the traditional approach of collecting large traffic measurement datasets for lengthy periods of time such as on the order of many days and then performing subsequent off-line statistical analysis. Cisco NetFlow essentially measures the volume and duration of each traffic flow for the purposes of accounting and off-line traffic analysis, but is not intended to be used for real-time network monitoring and querying. The OCXmon tool from NLANR is only for asynchronous transfer mode ("ATM") traffic, and is not for the purpose of traffic monitoring and control. Given the off-line use of these tools which, by definition, is delayed by some time which may be on the order of days, any conclusion that may be drawn from the collected data is stale in the sense that it applies to the network environment as it existed at some previous time rather than having a high likelihood of characterizing the network at the time the analysis is concluded; in other words, the status of the network is likely to have changed between the time these prior art systems collect their data and the time they drawn any inferences or conclusions during their off-line analysis. Thus, by time the results provide meaningful indicators, the real-time status of the network may have changed. Finally, as an example of a drawback of the element management system, the router MIB is usually designed based on the specific structure and implementation of the underlying network or IP router and, therefore, it will not be the same for equipment from different vendors. For example, the Argent Guardian tool from Argent Software Inc. for performance monitoring and proactive problem detection and correction has different versions for different monitored entities. The Argent Guardian for Cisco can only be used for Cisco routers because it uses the Cisco router MIB to retrieve and query the traffic information. In addition, the MIB is commonly embedded in the IP router. As a result of these constraints, there is considerable difficulty in changing the MIB once the design is complete. Accordingly, there are considerable barriers to the use of such a system in the instance where there is a change in the parameters or statistics to be monitored and where the change is not already built into the MIB.

[0009] In view of the above, there arises a need to address the drawbacks of the prior art as is accomplished by the preferred embodiments described below.

BRIEF SUMMARY OF THE INVENTION

[0010] In the preferred embodiment, there is a network monitoring system. The system comprises storage circuitry for storing network packet information, wherein the network packet information includes a predicted identifier. The network monitoring system also comprises at least one monitoring circuit coupled to a network along which network traffic flows in a form of packets. The at least one monitoring circuit is programmed to perform the steps of receiving a packet communicated along the network and determining whether the received packet is communicated between a source and destination in a first set of network nodes. Each packet in a sequence of communications between the source and the destination comprises a packet identifier that uniquely identifies the packet from all other communications in a flow between the source and the destination. The at least one monitoring circuit programmed to perform the step of, responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the received packet to the predicted identifier to determine an identifier deviation between the packet identifier and the predicted identifier for identifying an irregularity in the network traffic.

[0011] Other aspects are also described and claimed.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

[0012] **FIG. 1** illustrates a prior art network layer stack **10** defined to include the transmission control protocol ("TCP").

[0013] **FIG. 2** illustrates a prior art plot of the evolution of TCP sequence numbers over time when a TCP session, including the source, destination, and network, behave normally.

[0014] **FIG. 3** illustrates a prior art plot of the evolution of TCP sequence numbers over time when during a TCP session the network operates normally but either the source or destination is limiting the bandwidth of communication.

[0015] **FIG. 4** illustrates a prior art plot of the evolution of TCP sequence numbers over time when during a TCP session the network is limiting the bandwidth of communication between the source and destination by dropping packets.

[0016] **FIG. 5** illustrates a block diagram of a network system **10** into which the preferred embodiments are implemented.

[0017] **FIG. 6** illustrates a block diagram of each network monitor NM_1 through NM_5 of **FIG. 5**.

[0018] **FIG. 7** illustrates a flow chart of the operation of each network monitor NM_1 through NM_5 of **FIG. 5**.

[0019] **FIG. 8** illustrates a flow chart of the predictor methodology preferably implemented within the flow chart of **FIG. 7**.

DETAILED DESCRIPTION OF THE INVENTION

[0020] By way of introduction to one preferred inventive implementation, **FIG. 1** illustrates a known prior art network layer stack **10** defined to include the internet connection protocol ("IP") as discussed below. At the top of stack **10** is an application layer **12**, which typically represents one or more programs that are commonly used by consumers for various reasons, including personal and business computing. Beneath application layer **12** are located two protocol layers, namely, a transmission control protocol ("TCP") layer **14** and a user datagram protocol ("UDP") layer **16**. TCP layer **14** is a connection-oriented layer, that is, once a connection is established from a source to a destination and through that layer, then traffic through that connection is confirmed so that if a packet is lost before it is confirmed by its destination then the source re-transmits the packet. In contrast, UDP layer **16** is connectionless and, thus, no confirmation is made as between a packet passing from a source to its destination; accordingly, if a packet is lost in a UDP communication then there is no re-transmission of that packet. Comparing then the TCP and UDP layers **14** and **16**, TCP assures greater integrity of communications yet it requires greater overhead than UDP.

[0021] Below layers **14** and **16** is an IP layer **18**, using IP as introduced above. Thus, both TCP and UDP network packets along a given network pass through the IP layer **18** and therefore also must comply with the requirements of the IP protocol. IP layer **18** is also a connectionless layer, although for TCP packets passing through the layer the connection-oriented nature of TCP layer **14** may be relied upon to provide additional data integrity. Due to the connectionless nature of IP layer **18**, there is a trend in the networking art to provide measures of so-called quality of service ("QoS") in order to give priority or otherwise provide guidance to the flow of connectionless IP packets. In any event, and importantly for one example of the preferred embodiment as detailed in this document, each TCP (as opposed to UDP) network packet passing from a single source to a single destination and through IP layer **18** includes a unique TCP sequence number that thereby distinguishes that packet from all other TCP packets communicated from the source to the destination in a given flow session. This TCP sequence number is required as a TCP parameter and may be used for various purposes as known in the art, while the preferred embodiments further monitor the number to evaluate network performance as detailed below. Before reaching that discussion and completing the discussion of **FIG. 1**, located below IP layer **18** is a physical layer **20**. Physical layer **20** is commonly associated with the physical hardware at a source or destination, as typically represented by a network interface having a media access control ("MAC") address.

[0022] As further explored below, the preferred embodiments evaluate network performance, preferably as it is reflected in communications between specified nodes such as in specified sub-networks in the network, in response to a real-time evaluation of the progression of TCP packet sequence numbers. In this regard, in part the evaluation by the preferred embodiments includes consideration of three prior art observations made by the NLANR as relating to TCP sequence numbers. These three observations are explored below in connection with **FIGS. 2 through 4**.

[0023] **FIG. 2** illustrates a plot of the evolution of TCP sequence numbers over time when a TCP session, including the source, destination, and network, behave normally. In **FIG. 2**, the horizontal axis represents the passage of time since the beginning of the TCP session observation, while the vertical axis represents the sequence number of the last byte of each data packet. Note that the sequence numbers are shown for simplicity of illustration to start from zero, whereas in reality they typically commence from a non-zero, and likely random, number, and then increase from that point. In any event, the sequence numbers will increase monotonically unless there are reordered or retransmitted packets. While intuitively one might expect a linear relationship for normal behavior, instead note that the plot provides a number of upward curves each ending with a peak, such as shown by ways of example with a curve prior to a point P_1 at time t_1 and a point P_2 at a time t_2. Note that these curves, followed by respective peaks, occur under ideal network operation because the source transmission rate will tend to increase toward and reach the normal upper limit of the traffic rate for the destination (or the network). In other words, under TCP protocol, transmission occurs in a TCP control window, where the source starts transmitting at a certain rate, and hence with a corresponding steady increase in TCP sequence numbers, followed by an attempt within the window over time to increase the transmission rate and hence with a change in the rate of the ascension of the TCP sequence numbers. The source's periodic attempts during the TCP control window to increase the TCP sequence number frequency continues until the limit of the destination (or network) to accurately communicate packets is reached. As this limit is reached, the source will reduce its transmission rate for a time and thereafter again continue to increase it until the limit is again reached later in time.

[0024] **FIG. 3** illustrates a plot of the evolution of TCP sequence numbers over time, when during a TCP session the network operates normally but either the source or destination is limiting the bandwidth on communication. Here again, the horizontal axis represents the passage of time while the vertical axis represents the increase in sequence numbers. However, the plot of TCP sequence numbers reflects a straight line. Typically such a result is anticipated because it represents that the source is not reaching the maximum bandwidth limitation of the network or destination. Thus, while the TCP sequence numbers increase, the increase of the sequence numbers from one to the next is in a constant fashion across the entire illustrated plot, as opposed to an acceleration of those numbers that increases over time as shown in **FIG. 2** and that also peaks from time to time as a normal limit is reached.

[0025] **FIG. 4** illustrates a plot of the evolution of TCP sequence numbers over time, when during a TCP session the network is limiting the bandwidth of communication between the source and destination and, thus, it is presumed during such time that the source and destination are otherwise capable of communicating at their respective maximum bandwidths. However, the plot of TCP sequence numbers illustrates curves with discontinuities between the curves. For example, at time ti, the TCP sequence numbers are at a point P_{10}, but transmission is interrupted until a time t_2. At time t_2, however, note that a slight drop occurs in the TCP sequence number as shown at point P_{11}. This type of pattern repeats later at time t_3 where a sequence number shown as point P_{12} is reached, but an interruption then occurs until

BNSDOCID: <US____20040073655A1_I_>

communication through to the destination from the source is achieved at time t_4 and having a sequence number at point P_{13}. The behavior illustrated by these discontinuities occurs due to packet losses along the network as resulting from network irregularities. In other words, recalling that TCP communications are connection-oriented, when a network irregularity occurs one or more TCP packets will be lost, and those packets will be required to be re-transmitted by the source. As a result, the source, which receives the indication of the failure at a time when it was prepared to transmit a new packet having a certain TCP sequence number, is instead required to backtrack to re-transmit a previous packet with a lower TCP sequence number, followed by all packets following the re-transmitted packet. The beginning of such a re-transmission, therefore, is shown by way of example at point P_{11} (and various other points in **FIG. 4**), where a drop is shown in TCP sequence number as compared to point P_{10}, where the drop occurs due to the backtrack and re-transmission by the source of a previously-transmitted packet(s).

[0026] **FIG. 5** illustrates a block diagram of a system **60** into which the preferred embodiments are implemented. System **60** generally includes a number of stations ST_1 through ST_4, each coupled to a network **70** via a router, and each operable to send packets as a source or receive packets as a destination. By way of example, network **70** is an internet protocol ("IP") network such as the global internet or other IP-using network, where each station and IP networks in general are well known in the art. However, one skilled in the art should appreciate that the use of the IP protocol is by way of illustration, and many of the various inventive teachings herein may apply to numerous other protocols, including by way of examples asynchronous transfer mode ("ATM"), token ring, Novell, Apple Talk, and still others. In any event, returning to network **70** as in IP network, and also by way of an example, each station ST_x may be constructed and function as one of various different types of computing devices, all capable of communicating according to the IP and to include the TCP protocol. Lastly and also by way of example, only four stations ST_x are shown so as to simplify the illustration and example, where in reality each such station may be proximate other stations (not shown) and at a geography located at a considerable distance from the other illustrated stations. For example, station ST_1 might be located in Montana while station ST_2 is located in California, and station ST_3 might be located in New York while station ST_4 is located in Florida.

[0027] Continuing with **FIG. 5**, along the outer periphery of network **70** are shown a number of edge routers ER_1 through ER_{11}, while within network **70** are shown a number of core routers CR_1 through CR_4. The terms edge router and core router are known in the art and generally relate to the function and relative network location of a router. Typically, edge routers connect to remotely located networks and handle considerably less traffic than core routers. In addition and due in part to the relative amount of traffic handled by core routers, they tend to perform less complex operations on data and instead serve primarily a switching function; in other words, because of the tremendous amount of throughput expected of the core routers, they are typically hardware bound as switching machines and not given the capability to provide operations based on the specific data passing through the router. Indeed, core routers typically do not include much in the way of control mechanism as there could be 10,000 or more connections in a single trunk. Further, typically core routers do not involve their operations with TCP related items and instead deal at the IP level and below. In contrast, edge routers are able to monitor various parameters within data packets encountered by the respective router. In any event, the various routers in **FIG. 5** are shown merely by way of example, where one skilled in the art will recognize that a typical network may include quite a different number of both types of routers. Finally, note that each core router CR_x also generally may be constructed and function according to the art, with the exception that preferably the core routers may include additional functionality for purposes of traffic irregularity detection based on TCP sequence number evaluation as described below. In addition, each edge router ER_x also generally may be constructed and function according to the art, with the exception that preferably the edge routers are responsive to the traffic irregularity detection of the core routers, as also described below.

[0028] Completing the discussion of **FIG. 5**, note that the various stations, edge routers, and core routers therein are shown connected to one another in various fashions and also by way of example. Such connections are intended to illustrate an example for later discussion of the preferred operation and also to reflect a general depiction of how networks are generally established. Thus, each station ST_x is shown connected to a single edge router ER_x, where that edge router ER_x is connected to one or more core routers CR_x. The core routers CR_x, also by way of example, are shown connected to multiple ones of the other core routers CR_x. By way of reference, the following Table 1 identifies each station and router shown in **FIG. 5** as well as the other device(s) to which each is connected.

TABLE 1

station or router	connected nodes
ST_1	ER_1
ST_2	ER_{10}
ST_3	ER_5
ST_4	ER_7
ER_1	ST_1; CR_1
ER_2	CR_3; CR_2
ER_3	CR_2
ER_4	CR_2
ER_5	ST_3; CR2; CR_3
ER_6	CR_3; CR_4
ER_7	ST_4; CR_4
ER_8	CR_4
ER_9	CR_4
ER_{10}	ST_2; CR_1
ER_{11}	CR_4
CR_1	ER_1; ER_{11}; ER_{10}; ER_2; CR_2; CR_3; CR_4
CR_2	ER_2; ER_3; ER_4; CR_1, CR_3; CR_4; ER_5
CR_3	ER_5; ER_6; CR_2, CR_1; CR_4
CR_4	ER_7; ER_8; ER_9, CR_1, CR_2; CR_3; ER_6

[0029] Given the various illustrated connections as also set forth in Table 1, in general IP packets flow along the various illustrated paths of network **70**, and in groups or in their entirety such packets are often referred to as network traffic. In this regard and as developed below, the preferred embodiments operate to identify any irregularities in such network traffic. Finally, note that **FIG. 5** may represent a simplified

version of a network or the internet in that only a few stations and routers are shown, while one skilled in the art will readily appreciate that the inventive concepts in this document may be applied to a larger number of stations, routers, and the network interconnectivity between those devices.

[0030] **FIG. 5** also illustrates a number of network monitors NM_1 through NM_5, where the choice of five such network monitors is only by way of example given the limited amount of other hardware that is shown for network **70**. As detailed below, each network monitor NM_x is operable to sample, and possibly respond to, irregularities in the TCP sequence number in an IP packet that passes along the conductor(s) to which the network monitor is connected. Further, the components of each network monitor NM_x are described below, but at this point the connections of each such monitor are noted in the following Table 2:

TABLE 2

network monitor	connected nodes
NM_1	CR_1; CR_2
NM_2	CR_2; CR_3
NM_3	CR_1; CR_3
NM_4	CR_3; CR_4
NM_5	CR_1; CR_2; CR_3; ER_7; ER_8; ER_9

[0031] **FIG. 5** and Table 2 demonstrate that each of network monitors NM_1 through NM_4 is connected to sample packets passing along the conductor(s) between a pair of core routers. However, network monitor NM_5 is by way of an alternative example incorporated into core router CR_4. As a result, network monitor NM_5 is able to sample packets communicated with any of the nodes to which core router CR_4 is connected, namely, core routers CR_1, CR_2, CR_3, and edge routers ER_7, ER_8, and ER_9. Thus, the contrast of network monitor NM_5 to the other illustrated network monitors NM_1 through NM_4 is shown to demonstrate that in the preferred embodiment each network monitor NM_x may sample packets as a stand alone entity or may be combined with the hardware and software of an existing router; indeed, in the preferred embodiments a network monitor NM_x also may be combined with network or element management systems. In any event and by way of introduction to details provided later, in the preferred embodiments the sampling functionality of each network monitor NM_x permits real-time monitoring of the TCP sequence numbers in network packets, and in response determinations may be made, and actions may be taken, based on detecting network traffic irregularities based on the behavior of the TCP sequence numbers.

[0032] **FIG. 6** illustrates a block diagram of each network monitor NM_1 through NM_4 of **FIG. 1**, with the further understanding that functionally the following description also may be applied to network monitor NM_5 with the addition that certain functionality may be provided by the hardware and software already available from core router CR_4. Turning then to **FIG. 6**, first note that a console **30** is associated with network monitor NM_x, where in the preferred embodiment a single such console **30** communicates with multiple network monitors NM_x. For example, returning briefly to **FIG. 5**, preferably each of network monitors NM_1 through NM_5 communicates with a single console **30**, where such communications also may be by way of packets between console **30** and the network monitors NM_x. Console **30** may be constructed by one skilled in the art using various forms of hardware and software, where the selection is a matter of implementation choice in order to achieve the functionality described in this document. Turning to that functionality, console **30** preferably provides an administration (configuration) function and a characterization (reporting) function. To permit a user to perform these functions, various interfaces may be used such as a graphical Web-based configuration tool or a command line tool, where one preferable approach implements console **30**, by way of example, using the Apache Web server with the PHP server-side scripting language to provide a dynamic interface to a flow store **32**. In any event, the user interface preferably provides different screens for each of the administration and characterization functions. The actual interface and screens, however, may be implemented in various forms, where it is desirable for any form that a user can properly control console **30** to perform its administration and characterization function with respect to its flow store **32**. Preferably, a network administrator or the like uses the administration functionality of console **30** to create a set of rules and to provide those rules along with a predictor methodology to an RTFM meter **36** described later, where those rules cause meter **36** to retrieve TCP sequence numbers and the predictor methodology is later used to evaluate the history of the TCP sequence numbers for a flow, as also discussed later. Also, once the system is configured and left to collect TCP sequence number information from or about the monitored packets, the network administrator can use the characterization screen to query the information so as to generate network status reports based on the collected information and to thereby detect any network irregularities.

[0033] As introduced above, console **30** is connected to a flow store **32**, which preferably represents a storage medium that stores a flow database relating to monitored TCP sequence numbers. In the preferred embodiment, each network monitor NM_x includes its own flow database, although alternative embodiments may be created where more than one network monitor NM_x shares a common flow database. In the preferred embodiment the flow database in flow store **32** is an SQL-compatible database using the PostgreSQL relational database management system, and console **30** communicates with this database through the Web server's PHP link to the SQL database. Thus, any administration and configuration changes made via console **30** are passed directly to flow store **32**. Given the preceding, one skilled in the art should appreciate that access to flow store **32** can be achieved by SQL queries, enabling network administrators to automate the configuration process or integrate report generation. In the preferred embodiment, flow store **32** stores what is referred to in this document as a "rule set" (or "rule sets" when plural), which is initially provided to the flow store **32** from console **30** as part of the administration function and which is also conveyed to meter **36**. Additionally, flow store **32** stores the progression of TCP sequence numbers from a monitored flow between a source and destination node, as further detailed below. Moreover and as also discussed below, flow store **32** may store numerous different progressions of TCP sequence numbers, each corresponding to a monitored flow between sets of corresponding source and corresponding destination nodes. Still fur-

BNSDOCID: <US____20040073655A1_I_>

ther, flow store **32** preferably stores packet time of arrival data, as also detailed later. Lastly, recall that an RTFM meter **36**, described below, stores a predictor methodology. From this predictor methodology, analyses are preferably made on the TCP sequence number progression(s) and times of arrival as stored in flow store **32**, and those analyses may be reported to console **30** and are also used as a basis to respond to network irregularities.

[0034] Continuing with **FIG. 6**, recorder layer **34** provides an interface between flow store **32** and a meter **36** (or meters) in use and coupled to the network, and meter **36** copies TCP sequence numbers from respective selected network packets passing through the meter. Specifically, as each network packet is identified as between a particular same source and destination node, meter **36**, in response to a predictor methodology, provides the TCP sequence number from the packet to recorder layer **34**, which in turn communicates that information to flow store **32**. In the preferred embodiment, the monitored source and destination node can be determined randomly, or it can be identified from one or more rules in a rule set in flow store **32**. Thus, the applications of recorder layer **34** can be separated into two categories: manager applications and meter reading applications. Manager applications configure meter **36**, based on information, including rules in one or more rule sets, in flow store **32**. Meter reading applications permit the TCP sequence number data collected by meter **36** to be provided in a data format usable by flow store **32** and, indeed, recorder layer **34** facilitates the passage of that data into the flow database of flow store **32**. Recorder layer **34** applications may pass information between flow store **32** and the network probes of meter **36** either synchronously or asynchronously. This gives network administrators the flexibility of either using real-time network flow meters (i.e. NeTraMet) or importing data from other network analysis tools that are not able to provide real-time data (e.g. Cisco NetFlow data).

[0035] In the preferred embodiment, meter **36** is a Real-Time Traffic Flow Measurement ("RTFM") meter which is a concept from the Internet Engineering Task Force ("IETF"). As known in the RTFM art, RTFM meters are previously known to be used in systems for determining the service requested by IP packets that are passing through a network for purposes of collecting revenue, where such a service is identifiable by the transport port number specified in each IP packet. For example, RTFM meters are currently being considered for use in systems whereby an internet user is charged based on the service he or she is using on the internet; for example, a different fee may be charged to the user for each different internet service, including mail, video, phone calls, and web browsing. However, as detailed in this document, the preferred embodiment implements the RTFM meter instead to analyze each IP packet and to determine if the packet is a TCP packet between a given source and destination and, if so, to copy its TCP sequence number for storage so that the copied numbers can be interpreted for purposes of monitoring network activity so as to evaluate network performance and, hence, identify any irregularities with such performance. More specifically, meter **36** physically probes the underlying network traffic and each time meter **36** detects an IP packet on the network, it determines whether the packet is a TCP packet and is between a given (i.e., either randomly or specifically identified) source and destination. Accordingly, during the real-time passage of numerous IP packets by the evaluating meter **36**, that meter **36** does not always copy a fixed portion of each such packet into a database such as the entire packet or the entire packet header; instead, each meter **36** evaluates the source and destination address of each TCP packet and, if a given source and destination are matched, then meter **36** communicates the TCP sequence number for the packet to recorder layer **34**. In response, recorder layer **34** inserts the captured TCP sequence number into flow store **32**. Accordingly, that stored information is then available for analysis and reporting to console **30**, as detailed below.

[0036] **FIG. 7** illustrates a flow chart of an overall method **40** of operation for each network monitor NM_x of **FIGS. 5 and 6**, where method **40** is now described and additional preferred details are provided later with respect to **FIG. 8**. Additionally, while method **40** is described below with respect to a single network monitor NM_x, preferably each network monitor NM_x performs a comparable method at its respective location in network **70**. As a result, distributed operations occur at different network locations, where by way of example the results of each individual network monitor NM_x may be accumulated by a common resource, such as at a single console **30**. From this accumulated information, greater accuracy may be achieved in identifying network irregularities and their causes.

[0037] Turning to method **40**, in step **42** a network monitor NM_x according to the preferred embodiment first detects and stores the TCP sequence numbers for each of a number n TCP packets between a first set of source and destination nodes. In the preferred embodiment, the first set of source and destination nodes can be selected randomly or may be specified by a rule in flow store **32** that is conveyed to meter **36**. By way of one example, preferably such a first set includes only a single source and a single destination. Also by way of example, each of the single source and single destination is likely in a different physical sub-network, where each such sub-network is identifiable in that nodes in each sub-network share a same portion of their address as directed to the sub-network, where typically such an address portion is that following the subnet mask in the packet address. By way of an illustrative example, assume that the source is the **FIG. 5** station ST_1 in a first physical sub-network in Montana and the destination is station ST_3 in a second physical sub-network in New York. Thus, for step **42**, the RTFM meter **36** of the given network monitor NM_x evaluates each packet passing through the meter to determine if the packet is a TCP packet from the certain source to the certain destination in the two respective sub-networks. If so, then the TCP sequence number for that packet is stored. Once n such TCP sequence numbers are stored (i.e., corresponding to n packets between the same source and destination along the logical sub-network of interest), then step **42** also processes those sequence numbers so that the regularity of subsequently received packets may be predicted, but a discussion of such processing is deferred until later. In any event, note that at the initial completion of step **42**, a first window of n packets will have been detected and processed; this window is updated to maintain a total of n packets by later processing subsequent selected packets, where for each such subsequent packet its information is processed and added to the overall information while at the same time the information for the then-oldest packet is removed from the processed information. These details are

BNSDOCID: <US____20040073655A1_I_>
CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 581

better understood below. In any event, once step **42** is complete, method **40** continues to step **44**.

[0038] In step **44**, the network monitor NM_x continues to monitor IP packets until it detects another TCP packet between the step **42** source and destination, and then step **44** determines whether an irregularity exists in the TCP sequence number of the additionally-detected packet. Specifically, from the plots of **FIGS. 2 through 4**, once a sufficient number n of TCP sequence numbers are collected in step **42**, a plot of those numbers and the TCP sequence number from the additional packet in step **44** may be compared as against the trends represented by the plots of **FIGS. 2 through 4**, where the case of **FIG. 4** is of particular interest in the preferred embodiments. However, in practice, such an approach may require a considerably large value set for the value n of step **42** or it may involve a considerable delay in either collection or analysis time. Accordingly, in the preferred embodiment, step **44** collects one TCP sequence number from a next TCP packet in the first set. Additionally and as detailed below, through the use of the predictor methodology in meter **36**, as also introduced with respect to the processing of n packets in step **42** so that the regularity of future packets may be predicted, step **44** determines whether the actual TCP sequence number of the newly-collected TCP packet is within a range as predicted from the n previously-collected TCP sequence numbers of step **42**; further, step **44** also preferably determines whether the newly-collected TCP packet arrives within a time as expected based on the time of reception of one or more of the n previously-collected sequence numbers of step **42**. Additional details for this predictor methodology are described below. For sake of continuing with **FIG. 7** at this point, however, step **44** generally determines, such as through the use of the predictor methodology, whether an irregularity in the TCP sequence number of the newly-detected packet exists, such as may be indicated if the TCP sequence number of the newly-detected packet, and/or the time of arrival of the newly-detected packet, deviates beyond a threshold. Such an irregularity may reflect the conditions of **FIG. 4**. If no such irregularity is detected, then the flow returns to step **44** to detect another TCP packet in the first set and to repeat the same analysis with respect to its time of arrival and TCP sequence number. If an irregularity is detected, then flow continues from step **44** to step **46**.

[0039] Before discussing step **46**, note also in step **44** that regardless of whether an irregularity is detected, step **44** also indicates that the newly-detected packet is processed to update the data based on n packets. Specifically, recall that step **42** detects a first group of n packets in the first set and processes the sequence numbers of packets so that the regularity of future packets may be predicted. Looking now to step **44**, the latter indication of processing in that step **44** is intended to demonstrate that step **44** also updates the previous data based on the n packets of the preceding step **42**. In the preferred embodiment, this updating is such that the packet detected in step **44** becomes one of the n packets, whereas the oldest first set packet received by step **42** is now removed from the n packets. In other words, preferably the data relating to n packets is updated in a sliding window type manner, where at any given time after step **42**, the data and the predictions derived from that data as detailed below, relate only to the most recently received n packets in the first set. Once updated in this manner, the flow continues from step **44** either to step **46** or to return to step **44**, as discussed earlier.

[0040] In step **46**, having been reached due to a detected TCP sequence number irregularity between a source and destination in the first set from step **44**, the network monitor NM_x detects a TCP sequence number from one or more additional TCP packets communicated in the first set (e.g., as between the single source and single destination mentioned in the example of step **42**). For each of these additionally-detected packets in the first set, step **46** determines whether the TCP sequence number (and arrival time) also indicates an irregularity as compared to previous TCP sequence numbers and times of arrival of packets between the same source and destination. If an irregularity is not detected, then the preferred embodiment operates under the assumption that the irregularity detected in the preceding step **44** was either an isolated incident, a relatively minor irregularity, or otherwise not reason to make a network response or adjustment, and the flow of method **40** returns from step **46** to step **44**. On the other hand, if an irregularity is detected in step **46**, then the preferred method flow continues to step **48**. As detailed below, preferably an irregularity is considered as detected if the TCP sequence number and/or time of arrival of the TCP sequence packet deviates beyond a threshold, where the deviation relative to the threshold is measured based on an predicted TCP sequence number and time of arrival as provided by the predictor methodology. This threshold may be selected by one skilled in the art, and by way of example may be on the order of five percent from the predicted value.

[0041] Before discussing step **48**, note that step **46** also indicates at its conclusion, and in the same manner as discussed above with respect to step **44**, that regardless of whether an irregularity is detected, step **46** also processes the newly-detected packet to update the data based on n packets. Once more, therefore, this indication is intended to represent the update operation achieved by the sliding window, to again encompass information based on a total of n packets. Accordingly, in step **46** the updating is such that the packet (or packets) detected in step **46** becomes one of the n packets, whereas the oldest first set packet received by step **42** or step **44** is now removed from the n packets. Again, therefore, the data after step **46**, and the predictions derived from that data as detailed below, relate only to the most recently received n packets in the first set. Once updated in this manner, the flow continues from step **46** either to step **48** or to return to step **44**, as discussed earlier.

[0042] In step **48**, the network monitor NM_x continues to monitor IP packets and now collects IP sequence numbers and time of arrival data for packets in a second set of source and destination nodes, where the second set is preferably related in some rational sense to the first set (e.g., the single source and single destination) of steps **42**, **44**, and **46** and the second set is also defined by a rule or rules conveyed from flow store **32** to meter **36**. For example, recall earlier that step **44** was demonstrated by monitoring packets to detect those between source station ST_1 in a first sub-network in Montana and destination station ST_3 in a second sub-network in New York. Further, for step **48**, and since the second set is preferably related to the first set, then as a first example of step **48** consider that the second set includes any source node in the same physical sub-network as station ST_1 in

BNSDOCID: <US____20040073655A1_I_>

Montana and any destination node in the same physical sub-network as station ST_3 in New York. As a second example of step **48**, the second set may include any source node in a different physical Montana sub-network (i.e., not including the sub-network of station ST_1) and any destination node in a different physical New York sub-network (i.e., not including the sub-network of station ST_3). As a third example of step **48**, the second set may include any source node in the same physical sub-network as station ST_1 in Montana, but any destination node in a specified physical sub-network as station ST_4 in Florida; this latter example may prove beneficial where again this second set relates in some sense to the first set, such as may be the case where the destination sub-network of the first set (i.e., in New York) communicates packages received from Montana onward to Florida. In any event, the relationship between the first and second set is preferably stated or implied in the rule set(s) stored in flow store **32**. In this manner, one skilled in the art may establish different rules to thereby encompass different relationships between the first and second monitored steps, and to thereby draw corresponding conclusions based on any irregularities in the traffic flow between the nodes of the two different sets. Also in this regard, in one preferred embodiment the second set may includes the first set, and it also preferably includes a greater number of nodes than the first set. For example, step **44** (and **46**) may be directed to a single source node and a single destination node, while step **48** is directed to multiple source nodes and multiple destination nodes. In any event, therefore, in step **48** and for each packet passing through the network monitor NM_x, if it is a TCP packet of the second set, then the network monitor NM_x preferably copies the TCP sequence number, source address, destination address, and time of arrival for each of n packets between a same source and destination in the second set. In other words, if the second set has three sources, each communicating with three respective destinations (in addition to the first set source and destination), then step **48** endeavors to collect packet data for each these three pairs of source/destination nodes, thereby accumulating data for a total of three times n packets. Thereafter, method **40** continues from step **48** to step **50**.

[0043] In step **50**, the network monitor NM_x generally operates in the same fashion as described above with respect to step **46**, that is, another packet is detected and processed; however, step **50** performs its analysis with respect to the packet information obtained from step **48** in view of a next received second set packet. In other words, while step **46** analyzed packets from the first set, step **50** analyses the TCP sequence numbers and times of arrival for packets from the second set, again determining whether those two parameters deviate beyond a certain percent relative to a predicted value provided by a predictor methodology. In the preferred embodiment, the threshold of step **50** may be the same or different than the threshold used by step **46**. In any event, if the deviation, if any, is less than or equal to the step **50** threshold, then method **40** proceeds under the assumption that the earlier-detected irregularity from steps **44** and **46** were limited to the nodes in the first set, and the method returns to step **44**; optionally, however, as also shown in **FIG. 7**, before returning to step **44**, the flow can pass through a step **50***a* wherein an alarm is issued relating solely to the first set of packets. In other words, if no irregularity is found in the second set packet as examined by step **50**, then there is still the observation from the affirmative findings in steps **44** and **46** that there was an irregularity in at least two first set packets. Thus, if the level of granularity at interest is such that irregularities solely within the first set are of interest, then step **50***a* represents an indication, such as an alarm, relating to the first set so that future examination or action may be taken with respect to the cause of the first set irregularity. Continuing with step **50**, however, if the second set packet deviation exceeds the step **50** threshold, then method **40** continues to step **51**.

[0044] Before discussing step **51**, note that step **50**, in the same manner as steps **44** and **46** discussed above, also indicates at its conclusion the operation of processing the newly-detected packet to update the data based on n packets. Here, note that the packet at issue is a second set packet (as monitored from step **48**). Accordingly, the data based on a sliding window of n packets continues to be updated, but now that data is reflective of packets in the second set, which recall preferably reflects a larger number of source and/or destinations as compared to the first set. Further, the updated data is only that relative to the specific source and destination between which the step **50** packet was communicated. Lastly, as this update occurs for a packet, the data from the oldest packet of the previously-processed n packets is preferably no longer maintained.

[0045] In step **51**, having been reached due to a detected TCP sequence number irregularity between a source and destination in the second set from step **50**, the network monitor NM_x detects one or more additional TCP packets communicated in the second set. For each of these additionally-detected packets in the second set, step **51** determines whether the TCP sequence number (and arrival time) also indicates an irregularity. If an irregularity is not detected, then the preferred embodiment operates under the assumption that the irregularity detected in the preceding step **50** was either an isolated incident, a relatively minor irregularity, or otherwise not reason to make a network response or adjustment, and the flow of method **40** returns from step **51** to step **44**, again with the option as shown in step **50***a* of indicating an alarm relating to the irregularities detected earlier with respect to the first set packets. On the other hand, if an irregularity is detected in step **51**, then the preferred method flow continues to step **52**. Lastly, as with steps **44**, **46**, and **50**, for a given packet between a specific source and destination, step **51** also concludes by updating the data from the sliding window of the most recent n packets in the second set as between that specific source and destination, as updated for each packet (or packets) as detected by step **51**.

[0046] From the preceding, one skilled in the art will appreciate that step **52** is reached when method **40** has detected irregularities in the TCP sequence numbers and/or TCP packet times of arrival in both the first and second sets of source and destination nodes. Thus, irregularities have been detected in communications between a plural number of source and destination nodes having a relationship to one another, where that relationship is implied by the rule(s) that define the first and second sets. Note that the actual irregularity may be for various reasons. Moreover, the specific reason(s) causing the irregularity may be inferred based on the amount of the deviation between the actual and predicted values. For example, if the predicted and real values differ by a figure of approximately 50 percent, then the preferred embodiment concludes that a packet (or more than one packet) was lost. This might occur because the source TCP

control window was halved or the destination may have dropped a packet. Still other examples may be ascertained by one skilled in the art. In any event, in step **52**, method **40** attempts a corrective action to improve traffic flow in view of the irregularity. Preferably, such a corrective action is taken by the network meter issuing an event report, such as via a network packet, to the edge routers ER_x on the logical sub-network having the detected irregularity(ies). The edge router, in response, may take various actions. For example, one or more alarms may be set off to indicate to a user or network manager that an irregularity or irregularities have been detected. As another example, traffic queuing at the edge routers ER_x relating to the logical sub-network at issue may be adjusted. As still another example which is likely a rare case, there may be an observation that 90% of the flows are losing packets independently of their source and destination; in this case there is likely a considerable issue with the link where the abnormal behavior has been observed.

[**0047**] Two additional observations are also noteworthy with respect to method **40**. As a first observation, note that method **40** only generally illustrates the use of two sets of packets, where by way of example the first set is reflected in steps **42** through **46** and in one instance may be packets between only a single source and single destination of two respective physical sub-networks, where the second set is reflected in steps **48** through **51** and in one instance may be all packets between any source and any destination of the same two respective physical sub-networks. However, the present inventive teachings contemplate different divisions for each set, where there may be additional iterations for number, where preferably each additional set encompasses all packets in a lesser analyzed set. In this manner, the analysis of method **40** effectively continues to expand its scope to consider a greater amount of network traffic, each time including the traffic as considered during the lesser scope of analysis. In operating in this fashion, a greater amount of data may be collected to determine the magnitude of a potential irregularity in network traffic. As a second observation, note that each of steps **46** and **51** states that it evaluates one or more additional packets. This is intended to demonstrate that the irregularity detection for a given set can be based on any number of packets; thus, by looking only at **FIG. 7**, steps **44** and **46** may appear to represent the analysis of only two first set packets, but the earlier statement that step **46** may include multiple packets is intended to indicate that the inventive scope includes a variation where more than two first set packets are processed, that is, for each of those packets, they are analyzed to determine if an irregularity exists, and the data from those packets is also used to update the sliding window of n packets. Similarly, from **FIG. 7** alone, steps **50** and **51** may appear to represent the analysis of only two second set packets, but the earlier statement that step **51** may include multiple packets is intended to indicate that the inventive scope includes a variation where more than two second set packets are processed, that is, for each of those packets, they are analyzed to determine if an irregularity exists, and the data from those packets is also used to update the sliding window of n packets corresponding to each source and destination between which each packet is communicated.

[**0048**] **FIG. 8** illustrates a flow chart of the predictor methodology introduced above and which is therefore preferably used in connection with method **40** of **FIG. 7**. Indeed, for sake of cross-reference between **FIGS. 8 and 7**, the predictor steps of **FIG. 8** may be considered to occur as part of various steps from method **40** of **FIG. 7**. Thus, by way of example, **FIG. 8** illustrates steps **40** and **42** from **FIG. 7** and which are shown generally with an elaboration of various sub-steps as described below. Further, while the illustration of **FIG. 8** is directed to steps **42** and **44**, one skilled in the art should appreciate that the discussion below with respect to newly-received and processed as described above, also receives an additional packet(s) and updates the corresponding sliding window of n packet data, where from the following it will become apparent that from each update additional parameters are predicted and those predictions form the basis for detecting irregularities in each of steps **44**, **46**, **50**, and **51**. This processing is now further explored below in the various sub-steps of step **42**.

[**0049**] In sub-step **42***a*, the network monitor NM_x detects a network packet of the first set, that is, a packet having a source and destination address matching the source and destination (or multiple sources and destinations) that are specified, by a rule(s), to be monitored. In sub-step **42***b*, the TCP sequence number ("SN") and the time of arrival ("TOA") of the detected TCP packet are stored, such as in a buffer located in flow store **32**. In sub-step **42***c*, a difference, or offset, is determined and stored for the TCP SN parameter of sub-step **42***b* relative to the immediately-preceding packet that was received for the same source and destination. In other words, a TCP SN offset is determined by subtracting from the TCP SN of the presently-detected packet the TCP SN of the immediately-preceding packet that was received for the same source and destination. For the sake of reference, this offset value is hereafter referred to as Δ(TCP SN), and in the preferred embodiment it is stored in flow store **32**. Further, because sub-step **42***c* is repeated for each of the n packets, then in the preferred embodiment the TCP SN is stored and the value Δ(TCP SN) is computed for each of the n packets (although the first packet in the n packets may not provide a value of Δ(TCP SN) because there is no immediately-preceding packet from which to make the determination of the then-present packet relative to the immediately-preceding packet).

[**0050**] In sub-step **42***d*, a difference, or offset, is determined for the TOA parameter of sub-step **42***b* relative to the immediately-preceding packet that was received for the same source and destination. However, unlike the operation with respect to Δ(TCP SN), preferably there is no storage of multiple offset values of TOA for the multiple ones of the n packets. Instead, in the preferred embodiment and in order to reduce complexity and storage requirements, for each packet received, first the difference between the actual TOA of the received packet and the stored value of the immediately-previous packet (of the same set) is determined, and that difference, hereafter referred to as Δ(TOA), is stored to a single location. Once that value of Δ(TOA) is stored, then the actual TOA of the present packet is also stored to single storage location, thereby overwriting the actual TOA of any previous packet. Thus, when a next packet in the n packets is received, sub-step **42***d* again operates in this manner, thereby determining the Δ(TOA) as between that packet and the immediately-preceding packet, and then overwriting the previous value of Δ(TOA) with a new value of Δ(TOA) and also overwriting the previously-stored actual TOA of the immediately-preceding packet with the actual TOA value of the then-processed packet.

BNSDOCID: <US____20040073655A1_I_>

[0051] Sub-steps **42***a* through **42***d* continue for a total of n packets in the first set, where such a repetition is shown by the condition sub-step **42***e*. In other words, until n packets are processed through sub-steps **42***a* through **42***d*, then those sub-steps are repeated, thereby storing the value Δ(TCP SN) for each of the n packets relative to the packet that preceded it. However, with respect to the TOA, only the TOA of the immediately-preceding packet is saved and it is overwritten once another packet is detected and the value of Δ(TOA) with respect to that packet is determined and stored. Once n packets are processed in this manner, step **42***e* causes the flow to continue to sub-step **42***f*.

[0052] Sub-step **42***f* predicts an offset to the TCP SN for a future packet having the same source and destination as the n packets already detected. In the preferred embodiment, in order to provide the prediction of the offset with respect to the next TCP SN, first step **42***f* reviews the n stored values of Δ(TCP SN) and chooses the most common stored value as the predicted TCP SN offset value for future packets. For sake of reference, the predicted offset value is hereafter referred to as PO_{Δ}(TCP SN). Note also that one skilled in the art may implement alternative manners of processing the n stored values of Δ(TCP SN) in order to provide the predicted value of PO_{Δ}(TCP SN). In any event, once PO_{Δ}(TCP SN) is determined, method **40** continues to step **44**.

[0053] Looking to step **44** in **FIG. 8**, recall first from **FIG. 7** that step **44** detects another TCP packet between the step **42** source and destination, and then step **44** determines whether an irregularity exists in the TCP sequence number of the new packet as compared to past TCP sequence number history. According to the preferred predictor methodology, the determination of an irregularity is made using the predicted offset value PO_{Δ}(TCP SN) and the last-determined offset time of arrival value, Δ(TOA), from sub-step **42***d* by determining predicted values with respect to the newly-detected packet. Specifically, the offset value of PO_{Δ}(TCP SN) is added to the TCP SN of the immediately-preceding received packet in the first set, that is, the previous TCP SN is increased by the offset value it is predicted to change based on step **42***f*. Also, the value of Δ(TOA) is added to the TOA of the immediately-preceding received packet in the first set, that is, the TOA of the immediately-previous first set packet is increased by the amount of time that occurred between that previous packet and the packet that preceded it. These two totals, therefore, represent what are a predicted values for the TCP SN and TOA of the packet detected in step **44**. Given these expected values, the actual detected values of the packet detected in step **44** are compared to the predicted values, and a measure of deviation may be determined between the two. If the deviation, for either TCP SN or TOA, between the actual value and the predicted value is beyond a threshold, then an irregularity is deemed to occur and, as described above with respect to **FIG. 7**, method **40** continues to step **46**. To the contrary, if the deviation is equal to or less than the threshold, then no irregularity is deemed to occur and, as described above with respect to **FIG. 7**, method **40** returns to step **44** to detect a next TCP packet in the first set and to process that step as just described. Note also now that the threshold may be adjusted by one skilled in the art based on various considerations, including the observations made in **FIGS. 2 through 4**.

[0054] Before proceeding recall also that step **44** in **FIG. 7** concludes with an update of the data based on n packets. Given the sub-steps shown in **FIG. 8**, one skilled in the art should now appreciate that this update operates in the same general manner as for the first n packets in step **42**, but for step **44** the packet at issue is the step **44** detected packet. Thus, for step **44**, the offset value Δ(TCP SN) is determined with respect to the step **44** packet and in view of the immediately-preceding packet from step **42**; particularly, Δ(TCP SN) is determined by subtracting from the TCP SN of the presently-detected step **44** packet the TCP SN of the immediately-preceding packet that was received for the same source and destination in step **42**. Also, the step **44** value Δ(TCP SN) is stored while the oldest value of the n stored values of Δ(TCP SN), as corresponding to the oldest received first set packet in a group of n packets, is discarded (e.g., overwritten). As a result, the determination of the predicted offset value, PO_{Δ}(TCP SN), also may be updated given the newest value of Δ(TCP SN) and the removal of the oldest value of Δ(TCP SN). Also as part of the step **44** update, a new value of Δ(TOA) is determined, here by comparing the TOA of the step **44** packet with the last received packet in step **42**. Again, the new value of Δ(TOA) is stored, thereby overwriting the actual TOA of any previous packet. Lastly, note that these same update principles, as applying to Δ(TCP SN) and Δ(TOA), also apply to steps **46**, **50**, and **51**, that is, each indicated update is preferably achieved in the same manner with respect to these two parameters, where in each instance the most recently detected packet, for either the corresponding first or second set, is used to adjust Δ(TCP SN) so that is reflective of the last n received packets as between the same source and destination as the previous packets, and similarly Δ(TOA) is adjusted as a comparison between the presently received packet and its immediate predecessor of the same set.

[0055] Having described the application of the preferred predictor methodology with respect to step **44**, as mentioned above note that a comparable analysis is also preferably performed with respect to steps **46**, **50**, and **51** in **FIG. 7**. For example, recall from FIG. **7** that step **46** detects whether an irregularity occurs for an additional packet in the first set only once step **44** detects a first TCP packet in the first set with an irregular deviation in its throughput (e.g., as detected based on TCP SN and TOA). In the case where step **46** detects multiple packets, then it updates and uses the most current value of PO_{Δ}(TCP SN) as relating to the immediately-preceding n packets in the first set. Thus, for each received packet its actual TCP SN is compared to the sum of the TCP SN of the immediately-preceding packet plus the then-updated value of PO_{Δ}(TCP SN). Additionally, a time of arrival analysis is made for each packet, where the time of arrival of the then-detected packet is compared to the predicted time of arrival based on the Δ(TOA) value from the immediately-preceding packet. Thus, the value of Δ(TOA) is also changed for each successive packet. In other words, each time a first set packet is detected, the TOA of the immediately-preceding first set packet is subtracted from the TOA of the then-received first set packet, and the difference is stored as Δ(TOA); that difference is then added to the TOA of the then-received first set packet to predict the TOA of the next received first set packet, and when that next received first set packet is received, its TOA is compared to the immediately-previous predicted TOA. Again, a deviation beyond a threshold gives rise to some type of corrective/

warning response. Further, the preceding as applied to step **46** will also be readily applied by one skilled in the art to steps **50** and **51**.

[0056] An additional observation is also contemplated with respect to the TOA and Δ(TOA) values included within the present inventive scope. Particularly, the preceding describes the use of only two successive packets to determine those two values with respect to any of steps **42**, **44**, **46**, **50**, and **51**. However, note that the changing of Δ(TOA) immediately after an irregular packet is received may cause difficulty in detecting further TOA irregularities. Specifically, if the TOA of the irregular packet is used to overwrite a previous TOA and also to determine the next Δ(TOA), then that next value of Δ(TOA) may itself be irregular because it is based on an irregular packet. As such, the next packet may arrive at a regular timing, but its timing may be misconstrued as irregular because it will be compared to Δ(TOA) as derived from the previously-received irregular packet. Accordingly, further contemplated within the present inventive scope are other manners of determining Δ(TOA). As one example, rather than using only the TOA from two successive packets, the TOA of a number of packets greater than two may be compared against one another to determine a Δ(TOA), such as via the average of the difference between the TOAs of each of the previously-received packets.

[0057] Having described the application of the preferred predictor methodology with respect to steps **44** and **46**, note that a comparable analysis is also preferably performed with respect to steps **50** and **51**, but as relating to the second set. Therefore, when a new second set packet is detected, the value of PO_{Δ}(TCP SN) comparable to that from sub-step **42***f* is added to the TCP SN of the immediately-preceding second set packet and the result is compared to the actual TCP SN of the new second set packet. Similarly, when the new second set packet is detected, its actual TOA is compared to the sum of the TOA of the immediately-preceding second set packet with the difference in TOA between that immediately-preceding second packet and the second set packet which preceded it. The comparisons provide an indication of a deviation, if any, between the predicted TCP SN and TOA values and the actual values, and for the sake of step **50** or **51** those deviations are compared to the step **50** or **51** threshold as described above, with a response taken based on how the deviations, if any, compare to the step **50** threshold.

[0058] From the above illustrations and description, one skilled in the art should appreciate that the preferred embodiments provide a manner of monitoring the flow of network packets, with a particular benefit having been shown for monitoring TCP packets based on their TCP sequence numbers as well as their times of arrival. By the preferential analysis of flows in terms of TCP sequence numbers and times of arrival, a measure of throughput is realized, which as evidenced in **FIGS. 2 through 4** depicts the change of TCP sequence numbers as it relates to the times of arrival. The embodiments provide numerous benefits over the prior art. As one example of a benefit, unlike the prior art active systems, the preferred embodiments do not send special test packets (e.g., PING) and do not introduce additional test traffic onto the route of the regular data packets. Thus, in this sense the preferred embodiments do not impact the network performance. As another example of a benefit, as compared to the traditional passive measurement mechanisms in which the IP packets (or portions thereof) are routinely stored and then later studied using an off-line analysis of historical data, the preferred embodiments use real-time RTFM meters to retrieve packet information from packets as they are incurred during actual real-time traffic flow and to store portions of those packets; this stored data is also preferably used in a real-time or near real-time manner, such as on the order of less than one second of when the packet sequence numbers and times of arrival are collected, to evaluate the regularity of the underlying network and to take corrective action(s), if necessary. As still another example of a benefit, prior art MIBs provide single point analysis directed to the traffic flow at the location of the MIB. In contrast, the preferred embodiments contemplate analyzing real-time collected packet information from multiple points in the network and they are not constrained to the hardware type or manufacturer of each router. Still further, the preferred embodiments also provide numerous benefits inherited from the following intrinsic advantages of RTFM meter architecture: (1) all data on the IP network is recordable; (2) functional even when the given network elements are incapable of flow monitoring; (3) independence from physical and data layer technologies; and (4) good performance even during fault conditions. In all events, the preceding as well as other benefits should be appreciated by one skilled in the art. As another benefit, while the preferred embodiments are implemented with respect to TCP sequence numbers, the inventive teachings likewise may apply to another selected field(s) in a packet that readily distinguishes that packet from past and future packets in the sequence, where based on past packets the future packets can be monitored given an expectation of the value of the selected field for each of the future packets. As a final benefit, while the preferred embodiments are particularly advantageous in IP networks, they also may be applied to numerous other networks as well. In addition, while the present embodiments have been described in detail, various substitutions, modifications or alterations could be made to the descriptions set forth above without departing from the inventive scope which is defined by the following claims.

1. A network monitoring system, comprising:

storage circuitry for storing network packet information, wherein the network packet information includes a predicted identifier; and

at least one monitoring circuit coupled to a network along which network traffic flows in a form of packets, the at least one monitoring circuit programmed to perform the steps of:

receiving a packet communicated along the network;

determining whether the received packet is communicated between a source and destination in a first set of network nodes, wherein each packet in a sequence of communications between the source and the destination comprises a packet identifier that uniquely identifies the packet from all other communications in a flow between the source and the destination; and

responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the received packet to the predicted identifier to determine an identifier deviation

BNSDOCID: <US____20040073655A1_I_>

between the packet identifier and the predicted identifier for identifying an irregularity in the network traffic.

2. The system of claim 1 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the identifier deviation to a threshold; and

responsive to the identifier deviation exceeding a threshold, indicating an alarm.

3. The system of claim 1 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the identifier deviation to a threshold; and

responsive to the identifier deviation exceeding a threshold, adjusting packet communications between the source and the destination.

4. The system of claim 3:

wherein the at least one monitoring circuit monitors packets between core routers; and

wherein the adjusting step adjusts packet communications at an edge router.

5. The system of claim 1 wherein the packet identifier comprises a sequence number.

6. The system of claim 1:

wherein the packet comprises a TCP packet; and

wherein the packet identifier comprises a TCP sequence number.

7. The system of claim 1 wherein the at least one monitoring circuit is programmed to further perform, prior to the comparing step, the steps of:

receiving a plurality of packets communicated along the network;

for each packet in the plurality of packets:

determining whether the received packet is communicated between a source and destination in the first set of network nodes; and

responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, storing in the storage circuitry the packet identifier of the received packet.

8. The system of claim 7 wherein the at least one monitoring circuit is programmed to further perform the step of determining the predicted identifier in response to a plurality of packet identifiers stored by the storing step corresponding to a plurality of packets in the first set.

9. The system of claim 8 wherein the at least one monitoring circuit is programmed to further perform the step of determining the predicted identifier as a most common one of the plurality of packet identifiers.

10. The system of claim 1 wherein the at least one monitoring circuit is programmed to further perform the step of, responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing a time of arrival of the received packet to a predicted time of arrival to determine a time deviation between the time of arrival and the predicted time of arrival.

11. The system of claim 10 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the time deviation to a threshold; and

responsive to the time deviation exceeding a threshold, indicating an alarm.

12. The system of claim 10 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the time deviation to a threshold; and

responsive to the time deviation exceeding a threshold, adjusting packet communications between the source and the destination.

13. The system of claim 10 wherein the at least one monitoring circuit is programmed to further perform, prior to the step of comparing a time of arrival, the step of determining the predicted time of arrival in response to a difference in time between two successively-received packets between the source and the destination in the first set of network nodes.

14. The system of claim 1 wherein the first set consists of a single source and a single destination node.

15. The system of claim 1 wherein the first set comprises a plurality of source nodes in a first physical sub-network and a plurality of destination nodes in a second physical sub-network that differs from the first physical sub-network.

16. The system of claim 1 wherein the at least one monitoring circuit is programmed to further perform the steps of, responsive to determining that the identifier deviation exceeds a threshold:

receiving an additional packet communicated along the network;

determining whether the additional packet is communicated between the source and the destination in the first set of network nodes; and

responsive to determining the additional packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the additional packet to a predicted identifier to determine an identifier deviation between the packet identifier of the additional packet and a predicted identifier.

17. The system of claim 16 wherein the at least one monitoring circuit is programmed to further perform, responsive to determining that the identifier deviation between the packet identifier of the additional packet to a predicted identifier exceeds a threshold, the steps of:

receiving a plurality of packets communicated along the network; and

determining whether each packet in the plurality of packets is communicated between a source and destination in a second set of network nodes, wherein the second set of network nodes comprises the first set of network nodes.

18. The system of claim 17:

wherein the first set consists of a first number of nodes; and

wherein the second set consists of a second number of nodes greater than the first set of nodes.

19. The system of claim 18 wherein the first set consists of a single source and a single destination node.

20. The system of claim 18:

wherein the single source node corresponds to a first physical sub-network and the single destination node corresponds to a second physical sub-network; and

wherein the second set comprises a plurality of source nodes corresponding to the first physical sub-network and a plurality of destination nodes corresponding to the second physical network.

21. The system of claim 17 wherein the at least one monitoring circuit is programmed to further perform, relative to each packet in the plurality of packets and determined to be communicated between a source and destination in the second set of network nodes, the step of comparing the packet identifier of each respective received packet to a predicted identifier to determine a respective identifier deviation between the packet identifier of the respective received packet and the predicted identifier.

22. The system of claim 21 wherein the at least one monitoring circuit is programmed to further perform, relative to each packet in the plurality of packets and determined to be communicated between a source and destination in a second set of network nodes, the steps of:

comparing the respective identifier deviation to a threshold; and

responsive to the respective identifier deviation exceeding a threshold, adjusting packet communications along the second set of network nodes.

23. The system of claim 22 wherein the step of adjusting packet communications along the second set of network nodes comprises adjusting packet communications at least one an edge router.

24. The system of claim 21 wherein the at least one monitoring circuit is programmed to further perform, relative to each packet in the plurality of packets and determined to be communicated between a source and destination in a second set of network nodes, the steps of:

comparing the respective identifier deviation to a threshold; and

responsive to the respective deviation exceeding a threshold, indicating an alarm.

25. The system of claim 16 wherein the at least one monitoring circuit is programmed to further perform the step of, responsive to determining the additional packet is communicated between a source and destination in the first set of network nodes, comparing a time of arrival of the additional packet to a first predicted time of arrival to determine a time deviation between the time of arrival of the additional packet and the first predicted time of arrival.

26. The system of claim 25 wherein the at least one monitoring circuit is programmed to further perform, responsive to determining that the time deviation exceeds a threshold, the steps of:

receiving a plurality of packets communicated along the network; and

determining whether each packet in the plurality of packets is communicated between a source and destination in a second set of network nodes, wherein the second set of network nodes comprises the first set of network nodes.

27. The system of claim 26 wherein the at least one monitoring circuit is programmed to further perform, relative to each packet in the plurality of packets and determined to be communicated between a source and destination in a second set of network nodes, the step of comparing a time of arrival of the packet to a corresponding predicted time of arrival to determine a corresponding time deviation between the time of arrival of the packet and the corresponding predicted time of arrival.

28. The system of claim 27 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the corresponding time deviation to a threshold; and

responsive to the corresponding time deviation exceeding a threshold, indicating an alarm.

29. The system of claim 27 wherein the at least one monitoring circuit is programmed to further perform the steps of:

comparing the corresponding time deviation to a threshold; and

responsive to the corresponding time deviation exceeding a threshold, adjusting packet communications between at least some of the second set of nodes.

30. The system of claim 1 wherein the network comprises an internet protocol network.

31. The system of claim 1 wherein the internet protocol network comprises the global internet.

32. The system of claim 1 and further comprising a plurality of monitoring circuit, each coupled to the network, and each programmed to perform the steps of:

receiving a packet communicated along the network;

determining whether the received packet is communicated between a source and destination in a first set of network nodes, wherein each packet in a sequence of communications between the source and the destination comprises a packet identifier that uniquely identifies the packet from all other communications in a flow between the source and the destination; and

responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the received packet to the predicted identifier to determine an identifier deviation between the packet identifier and the predicted identifier for identifying an irregularity in the network traffic.

33. A method of monitoring a network along which network traffic flows in a form of packets, the method comprising:

storing network packet information in storing circuitry, wherein the network packet information includes a predicted identifier and wherein the storing circuitry is accessible to the network; and

BNSDOCID: <US____20040073655A1_I_>

operating at least one monitoring circuit coupled to the network along which network traffic flows in a form of packets, the operating step comprising the steps of:

receiving a packet communicated along the network;

determining whether the received packet is communicated between a source and destination in a first set of network nodes, wherein each packet in a sequence of communications between the source and the destination comprises a packet identifier that uniquely identifies the packet from all other communications in a flow between the source and the destination; and

responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing the packet identifier of the received packet to the predicted identifier to determine an identifier deviation between the packet identifier and the predicted identifier for identifying an irregularity in the network traffic.

34. The method of claim 33, wherein the operating step further comprises:

comparing the identifier deviation to a threshold; and

responsive to the identifier deviation exceeding a threshold, one of the steps from the group consisting of indicating an alarm and adjusting packet communications between the source and the destination.

35. The method of claim 33 wherein the packet identifier comprises a sequence number.

36. The method of claim 33, wherein the operating step further comprises, responsive to determining the received packet is communicated between a source and destination in the first set of network nodes, comparing a time of arrival of the received packet to a predicted time of arrival to determine a time deviation between the time of arrival and the predicted time of arrival.

37. The method of claim 36 wherein the operating step further comprises:

comparing the time deviation to a threshold; and

responsive to the time deviation exceeding a threshold, one of the steps from the group consisting of indicating an alarm and adjusting packet communications between the source and the destination.

38. The method of claim 33 wherein the operating step further comprises, after determining that the identifier deviation between the packet identifier of the packet to a predicted identifier exceeds a threshold, the steps of:

receiving a plurality of packets communicated along the network; and

determining whether each packet in the plurality of packets is communicated between a source and destination in a second set of network nodes.

39. The method of claim 38 wherein the second set of network nodes comprises the first set of network nodes and wherein the second set consists of a number of nodes greater a number of nodes in the first set of nodes.

40. The method of claim 38 wherein the operating step further comprises, relative to each packet in the plurality of packets and determined to be communicated between a source and destination in the second set of network nodes, the steps of:

comparing the packet identifier of each respective received packet to a predicted identifier to determine a respective identifier deviation between the packet identifier of the respective received packet and the predicted identifier;

comparing the respective identifier deviation to a threshold; and

responsive to the respective identifier deviation exceeding a threshold, one of the steps from the group consisting of indicating an alarm and adjusting packet communications between the network.

* * * * *

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
1 November 2012 (01.11.2012)

WIPO | PCT

(10) International Publication Number
WO 2012/146265 A1

(51) International Patent Classification:
H04L 29/06 (2006.01) *H04L 12/26* (2006.01)
H04L 12/24 (2006.01)

(21) International Application Number:
PCT/EP2011/002145

(22) International Filing Date:
28 April 2011 (28.04.2011)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant *(for all designated States except US)*: **VOIPFUTURE GMBH** [DE/DE]; Wendenstrasse 4, 20097 Hamburg (DE).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **SCHOLZ, Hendrik** [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE). **WALLBAUM, Michael** [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE). **KRUEGER, Michael** [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE). **BASTIAN, Jan** [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE).

(74) Agent: **BOEHMERT & BOEHMERT**; Felix Hermann, Hollerallee 32, 28209 Bremen (DE).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

(54) **Title**: CORRELATION OF MEDIA PLANE AND SIGNALING PLANE OF MEDIA SERVICES IN A PACKET-SWITCHED NETWORK

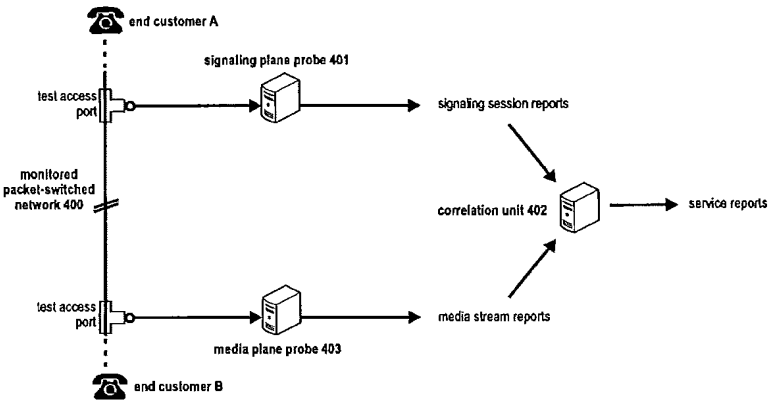

Fig. 4

(57) **Abstract**: This invention relates to methods for correlating media streams and signaling sessions of services, for example, in a passive monitoring system of a packet-switched network. Furthermore, the invention also relates to an implementation of these methods in hardware and software, and provides a signaling plane probe, a media plane probe and a correlation unit. Moreover, a passive monitoring system comprising one or more of these hardware devices is provided. To correlate media streams and signaling sessions of services, the invention proposes to independently generate correlation keys in a media plane probe for monitored media streams and correlation keys for signaling sessions that are monitored by a signaling plane probe in a fashion that matching correlation keys are generated for a respective service. By identifying reports on the media streams and reports on the signaling session that contain matching correlation keys respectively, the media streams and the signaling streams are correlated in a correlation unit.

WO 2012/146265 A1

BNSDOCID: <WO_____2012146265A1_I_>

CISCO EXHIBIT 1002
File History of 9565213 Patent

Correlation of Media Plane and Signaling Plane of Media Services in a Packet-Switched Network

FIELD OF THE INVENTION

This invention relates to methods for correlating media streams and signaling sessions of services, for example, in a passive monitoring system of a packet-switched network. Furthermore, the invention also relates to an implementation of these methods in hardware and software, and provides a signaling plane probe, a media plane probe and a correlation unit. Moreover, a passive monitoring system comprising one or more of these hardware devices is provided.

TECHNICAL BACKGROUND

VoIP Services

VoIP ("Voice over IP" – IP denoting the Internet Protocol) services may be considered to consist of a signaling plane and a media plane. On the signaling plane various protocols describe the session (call) flow in terms of involved parties, intermediary VoIP entities (i.e. VoIP proxies, routers) and the characteristics of the VoIP service (call).The media plane typically carries the media information (i.e. audio and/or video data) between the involved parties. Neither the media plane nor the signaling plane alone is sufficient to carry a VoIP service. On the signaling plane protocols like SIP (see IETF RFC 3261, "SIP: Session Initiation Protocol", available at http://www.ietf.org) or ITU-T recommendation H.323 (see H.323, "Packet-based multimedia communications systems", Edition 7, 2009, available at http://www.itu.int) are commonly used, whereas protocols like RTP (Real-time Transport Protocol, see IETF RFC 3550, "RTP: A Transport Protocol for Real-Time Applications", available at http://www.ietf.org), MSRP (see IETF RFC 4975, "The Message Session Relay Protocol (MSRP)", available at http://www.ietf.org) or ITU recommendation T.38 (see T.38, "Procedures for real-time Group 3 facsimile communication over IP networks", Edition 5 (2007) or Edition 6 (2010), available at http://www.itu.int) may be present on the media plane.

In contrast to the traditional PSTN (Public Switched Telephone Network) network both planes may be on different infrastructure using different protocols and even take different routes through a network.

Correlation Information

The property of disjunct signaling stream(s) and media stream(s) of a VoIP service within packet-switched networks may necessitate a correlation mechanism. This correlation may only be based on the common parameters carried on both planes of the session. Monitoring equipment used in VoIP networks may be limited to either the signaling or media plane due to the network structure. A correlation of signaling stream(s) and media stream(s) is thus needed outside the VoIP network in between the monitoring components.

In RTP-based VoIP networks the data basis is rather limited as defined by the RTP protocol and the underlying transport and Internet layer protocols. An RTP stream which typically carries voice or video information from one party to another of a VoIP service is made up of plural RTP packets. The header of each RTP packet does not contain an identifier which can be used to distinctively associate single RTP packets or the entire RTP stream to a VoIP service (defined by the VoIP signaling). The same way the VoIP signaling protocol does not carry information to relate a signaling session (signaling plane of the VoIP service) to RTP protocol header information.

The parties originating or receiving a VoIP service actively dictate the transport and Internet layer address (IP address and port) on which they want to receive media streams. This configuration is done on the signaling plane. Using SIP as a signaling protocol, the address tuple(s) to be used on the media plane is/are defined in a so-called SDP body (see IETF RFC 3264, "An Offer/Answer Model with the Session Description Protocol (SDP)", available at http://www.ietf.org). The address tuple(s) - usually consisting of an IP address and port – is/are the sole information to correlate media plane of a VoIP service with the signaling plane of the VoIP service, and vice versa.

Challenges in Correlation of Signaling Plane and Media Plane

Beside the limited information available for correlation, the correlation of the signaling plane and media plane imposes also numerous further challenges, especially in passive mid-point monitoring scenarios. While an address tuple (IP address and port) may be enough information to identify a single RTP stream at a given point in time, multiple streams

identified by different address tuples may exist in a regular VoIP call. In fact at least two media streams are needed to allow two-way communication between the caller and callee.

It will be outlined in further detail in the following how multiple media streams may occur in a single VoIP service and how they are communicated between the parties involved (which may not always be the case).

Mid-point monitoring equipment (i.e. the signaling and media probes) is introduced into the network path which active components (e.g. VoIP phones and servers) use to communicate with each other. This passive equipment is not involved in the communication itself and has to derive all information from the network packets it observes on said network path. The correlation challenge increases as the passive equipment lacks information the active components internally have. A VoIP phone may internally assign a media port to a signaling session but never send out this information combined so that a mid-point monitoring probe could benefit from this information.

Time Constraints due to Port Reuse

The IP address and port combination transmitted inside the signaling plane of a VoIP service as well as the IP address and port observed on the media plane of the VoIP service may be reused. For exemplary purposes it can be assumed that each VoIP device has a single IP address. The port range from which the client can draw is limited to roughly 65000 ports by the underlying transport layer protocol UDP (see RFC 768, "User Datagram Protocol", available at http://www.ietf.org). In practice additional specifications and configuration parameters significantly further limit this port range. RFC 3550 which defines the Real-Time Transport Protocol mandates that only even port numbers may be used to transport media streams. The odd ports are reserved for Real-Time Transport Control Protocol (RTCP) packets. This standardization alone effectively halves the number of ports available for VoIP communication.

As a result a VoIP device has to re-use a certain IP address and port combination, thus eliminating a unique identification criterion. Essentially the IP address and port combination can only be assumed to be unique for a limited amount of time. The time-frame depends on the utilization of the VoIP device as well as the number of IP addresses required in an Internet Service Provider's network. If a correlation of the signaling and media plane takes place without a time criteria, it would therefore produce so-called false positives: Multiple media streams would match the requested IP address and port combination, since the signaling plane may have caused re-use of the very same IP address and/or port for another VoIP service.

Media Streams

Multiple streams may exist per direction of each VoIP service. The caller sends data to the callee and vice versa resulting in at least one media stream per direction. Desirable features, such as for example codec changes, may trigger additional streams. On the network multiple streams may use the same IP address and port tuple depending on the media protocol. For example, each RTP stream is additionally identified by a unique identifier - the Sender Synchronization Source (SSRC) – which is used to distinguish between multiple streams.

Fig. 1 shows exemplary situations where multiple media streams may be triggered. The focus is on the media plane traffic and its timing. The x-axis denotes the time basis, starting on the left side. The vertical lines in Fig. 1 denote signaling events and the bold horizontal lines are RTP streams present in the respective exemplified VoIP service example (please note that the different VoIP service examples are separated by broken horizontal lines). Only the RTP stream direction from the callee to the caller is shown per VoIP service example. The RTP stream from the caller to the callee can be assumed to be less complex (e.g. as shown in the "normal" VoIP service example).

The three vertical lines denote important events in the phase of a call. The leftmost line denotes the start timestamp. This is the time at which the caller submitted the call setup request to the network. The connected point is the timestamp at which a callee or his endpoint/handset picked up and moved the service into a fully established state. The service ends when one of the parties sends a termination request.

The most common VoIP services will be the trivial session examples where there is only one RTP stream per direction (see "normal" or "early media"). In a "normal" VoIP service flow both parties only start sending once the session is fully established on the signaling plane. More often the callee would start sending early media, i.e. when sending ring-tones or (pre-call) announcements. In "session forking" VoIP service example there may be multiple concurrent or sequential RTP streams at the beginning of a service e.g. due ring-tones being sent to multiple VoIP devices of the callee, while one is answered at the connect event

During a service media codecs may be changed (see "code change" VOIP service example) which is commonly resulting in a new media stream. Another service scenario may be the VoIP call being put on hold and picked up at a later time essentially generating multiple media streams (see "call on hold" VoIP service example). Silence suppression and other RTP features may cause multiple streams sporadically during a service (see "silence suppression" VoIP service example). Also retransmissions on the signaling plane may

desynchronize the signaling plane and media plane (see "SIP retransmissions" VoIP service examples).

Geographic Distribution or Media Plane and Signaling Plane

As indicated above, signaling plane data and media plane data of VoIP services may take different routes. In examples, where SIP is used for session setup, this is indeed the case as those sessions usually contain signaling-only components (i.e. SIP proxies, REGISTRARs). As a result passive mid-point monitoring solutions may be split geographically along the routes of the VoIP traffic.

An exemplary passive mid-point monitoring scenario for this case is shown in Fig. 2. Fig. 2 shows three geographically distributed POPs (Point of Presence) which are linked over the carrier's internal network. The signaling data flows on the signaling plane between POP A and B via the signaling-only POP C. POP C would typically host a centralized routing entity which reduces the complexity in the POPs A and B. To prevent back-hauling traffic from POP A to POP C before sending it to POP B the VoIP system is configured in a way that would allow the media to flow directly on the media plane between POPs A and B.

Typical mid-point monitoring locations would be inside each POP. Due to the layout of the network no media traffic would be visible in POP C. To allow for successful correlation of signaling and media streams by the correlation mechanism it is therefore desirable to support distributed networks. If the same session traverses multiple POPs it would be desirable for the correlation mechanism to identify this and support correlation of the session legs across multiple sites.

Limitations/Challenges in Correlation of Signaling Plane and Media Plane

The correlation may be solely based on the IP address and port used on the transport and Internet layer. On the media plane each RTP stream is destined to exactly one IP address and port on the transport layer. On the signaling plane one IP address and port is typically exchanged inside the Session Description Body of the SDP protocol during the VoIP service setup. However this IP address and port combination may not be the combination used throughout the VoIP service. In this case the above mentioned correlation mechanism fails: It fails upon searching a media stream for a specific signaling session of the VoIP service. The same way it fails to locate a signaling session of the VoIP service for a specific media stream on the media plane. The reasons for using another IP address port combination than advertised upon session setup are manifold.

Network Address Translation (NAT)

Network Address Translation (NAT, RFC 1631) was developed to solve Internet address space depletion and routing problems. In today's Internet one may distinguish between internal addresses which are often only unique locally (i.e. typically used inside company or home networks), and external addresses of global validity and uniqueness as used on the public Internet.

Fig. 3 shows an exemplary network using internal and external addresses. The internal network on the left side uses IP addresses from the so-called private space, as for example defined in IETF RFC 1918, "Address Allocation for Private Internets", available at http://www.ietf.org. The external network on the right side is the public Internet with globally reachable addresses. A router with embedded Network Address Translation (NAT) functionality is located in between the two networks. The router does not directly route packets in between the two connected networks, but "translates" (i.e. replaces) IP source and destination addresses within the routed IP packets on demand.

Essentially, the NAT functionality of the router hides the entire internal network (10.0.0.0/8 in this example) behind one public IP address (192.0.2.1). Every IP connection from within the internal network to the Internet will seem to originate from the sole public IP address (192.0.2.1) of the entire network.

For each node inside the internal network this imposes some limitations. Although Internet access works in general, a node in the internal network may not be reachable from the outside. This is due to the fact that IP traffic from an outside node (i.e. a node in the external network) would be destined to the public IP address (192.0.2.1), but the router with this IP address does not know at this point in time where to forward the traffic to on the internal network, given that there was no prior communication from an node inside of the internal network and the external node.

Additionally the internal node may not be aware that it is behind a router which performs the NAT. As required by the transport protocols, the internal node will use the internal IP address for communication purposes. On the VoIP signaling and media layer (i.e. session layer or application layer), the application will also use the internal IP address (i.e. 10.0.0.123) instead of the public IP address. VoIP signaling protocols (e.g. SIP or H.323) contain various IP addresses on the application layer. An intermediary NAT device (i.e. the router in Fig. 3) will not perform translation of the addresses on the session layer or application layer, but only translates IP addresses in the network layer. The internal IP addresses inside the VoIP

signaling (i.e. SIP or H.323 messages) impose a problem to the application as internal IP addresses are most likely not reachable from the public Internet.

VoIP devices involved in the processing of a certain call may detect this problem by comparing the IP addresses seen on the Internet and transport layers with the IP addresses set on the application layer. The following message is an exemplary SIP message as seen on the network and transport layer, which means that both the Internet and transport layer information (IP addresses and ports) are visible:

```
U 2010/09/09 13:57:31.925226 192.0.2.99:37682 -> 10.0.0.1:5060
```

The Internet and transport layers show that the UDP packet (U) was received from the external address 192.0.2.99 (port 37682).

The following message represents the same exemplary SIP message as seen on the application layer, which means that both the Internet and transport layer information (IP addresses and ports) as well as the application layer (SIP message contents) are visible:

```
INVITE sip:echo@example.com SIP/2.0
Via: SIP/2.0/UDP 192.168.1.216;branch=z9hG4bK3BB64874;rport=37682;received=192.0.2.99
CSeq: 7726 INVITE
To: <sip:echo@example.com>
Content-Type: application/sdp
From: "John Doe" <sip:jdoe@example.com>;tag=693C7725
Call-ID: 1341723561@192.168.1.216
Subject: sip:jdoe@example.com
Content-Length: 230
User-Agent: kphone/4.2
Contact: "John Doe" <sip:jdoe@192.168.1.216;transport=udp>
v=0
o=username 0 0 IN IP4 192.168.1.216
s=The Funky Flow
c=IN IP4 192.168.1.216
t=0 0
m=audio 60942 RTP/AVP 0 97 8 3
a=rtpmap:0 PCMU/8000
a=rtpmap:3 GSM/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:97 iLBC/8000
a=fmtp:97 mode=30
```

In contrast to the information available at the network and transport layers claiming that the message was received from the external address 192.0.2.99 (port 37682), the application claims that the source was the internal address 192.168.1.216 (port 5060 - the port 5060 is not viewed for the source address as it is the default port).

In general the translation of the network (IP) address also applies to the port numbers which are used on the transport layer. This feature is named Port Address Translation (PAT). NAT and PAT commonly happen at the same time in the same device.

PAT becomes necessary if multiple devices attempt to use the same port. Since different devices have different IP addresses their IP address and port tuple is unique. If packet streams from both devices traverse the same NAT device, the NAT device has to translate the port of one of the two packet streams. Otherwise there would be two distinct packet streams sharing the same external IP address of the NAT device and the same port. While this may not be a problem if the destination IP address and port tuple of the destination device are different for the two packet streams, it adds overhead which the NAT device wants to prevent. By rewriting the source port of one of the packet streams the address tuples of the external IP address and ports become unique again.

However, in networks where NAT (and PAT) is performed to for example hide private networks from the outside Internet, the correlation of signaling sessions of VoIP services and the associated media plane streams is becoming complex.

Multiple Communication Paths

Numerous solutions have been developed to allow traffic of VoIP services to operate in networks with Network Address Translation (NAT). These mechanisms require the involved caller and callee to exchange information, for instance on additional transport protocols or addresses which can be used to reach the remote party. These include – but are not limited to - IPv6 (Internet Protocol Version 6, see IETF RFC 2460, "Internet Protocol, Version 6 (IPv6), Specification"), TURN (Traversal Using Relays around NAT, see IETF RFC 5766, "Traversal Using Relays around NAT (TURN):Relay Extensions to Session Traversal Utilities for NAT (STUN)") addresses or ICE (Interactive Connection Establishment, see IETF RFC 5245, "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols") candidates (all RFCs available at http://www.ietf.org).

These additional tuples of IP addresses and ports communicated in these solutions establish different data paths in media plane and signaling plane of a VoIP service and thereby break

up the previous one-to-one relation of signaling session to media stream to a given VoIP device. Previously one media stream to the given VoIP device could be distinctively identified by exactly one IP address and port. In the new scenario on the signaling plane a list of IP address and port tuples is transported on which the VoIP device may receive media thereby increasing the complexity of correlation between media plane and signaling plane of a VoIP service.

SUMMARY OF THE INVENTION

One object of the invention is to suggest correlation mechanisms that allow overcoming one or more of the limitations and challenges on the correlation of signaling plane and media plane of services.

The object is solved by the subject matter of the independent claims. Advantageous embodiments of the invention are subject to the dependent claims.

One of the various aspects of the invention is to suggest a definition of correlation keys that allow the correlation of media streams of services and the associated signaling session of the respective services based on data records (reports) that provide information on the media plane and signaling plane of the respective services. These correlation keys may be for example provided in form of hash keys. A correlation key may be associated to a given direction of the media data paths of the media streams belonging to a given service (e.g. caller-to-callee or callee-to-caller for a simple bi-directional service, e.g. VoIP call). Hence, assuming that the service comprises at least two media streams of different directions, at least two correlation keys are provided on the signaling plane for later correlation of the media plane and signaling plane of the service. On the media plane one correlation key is provided per media stream.

The correlation keys for use in correlation are generated by a media plane probe and a signaling plane probe independently from each other, but the two probes generate matching correlation keys for a given media service (i.e. the signaling plane probe and the media plane probe independently each generate the same correlation key for a given direction/media data path of the media service, given that there are not any exceptional situations preventing the match of the correlation keys).

In response to the detection of the end of a signaling session, the signaling plane probe generates for this signaling session at least one correlation key for each "direction" of the service to which the signaling session belongs. Each correlation key (e.g. hash key) is generated by the signaling plane probe based on a timestamp indicative of the point in time

at which the end of the respective signaling session has been detected by the signaling plane probe and using information on respective direction of the media plane data flow (e.g. caller-to-callee or callee-to-caller for a simple bi-directional service, e.g. VoIP call). The correlation keys generated for the different directions of the service corresponding to the signaling session the end of which has been detected are then collated in a signaling session report, which may be for example stored in a database or a flat file.

For the media plane, the media plane probe generates a correlation key (e.g. hash key) for each media stream. The correlation keys for the media steams are generated upon the media plane probe detecting the end of the respective media stream. The correlation key for a respective terminated media stream is generated by the media plane probe based on the direction of the respective media stream, e.g. based on the media data path of a media stream, and on a timestamp indicative of the point in time at which the respective media stream has ended. The timestamp at which the respective media stream has ended is for example the time at which the last packet of the media stream has been received by the media plane probe. The media plane probe stores the correlation key (e.g. hash key) for the media stream in a media stream report, which may be for example stored in a database or exported as a flat file.

Another aspect of the invention is the use of the correlation keys comprised in the signaling session reports and the media stream reports to identify the media streams and signaling session belonging to a single media service. As indicated above, the correlation keys generated by the signaling plane probe and the media plane probe should be matching each other (i.e. the correlation keys of a given signaling session report should be found in the media stream reports, respectively, the correlation key of a media stream report should be found in one of the signaling session reports), provided that signaling session and related media streams are terminated nearly at the same point in time, i.e. within a given timeframe. Hence, a correlation unit can identify which media streams belong to which signaling session (or vice versa) and may therefore generate a service record that summarizes a service based on information contained in the media stream reports and information contained in the signaling session report of the respective service.

According to one exemplary embodiment of the invention, a method for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network is provided. It is assumed for exemplary purposes that one service has one signaling session.

In this method a signaling plane probe generates at least one correlation key for each direction of a service based on the information of the signaling session of the respective

service monitored at the signaling plane probe. As noted above the direction may be equivalent to a respective media data path of a media stream belonging to a given service. The correlation keys for a respective service may be generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and using information on the direction of the respective service. Furthermore, the signaling plane probe stores a signaling session report for each signaling session in a database. Each of the signaling session reports comprises at least two correlation keys generated for the respective directions of the service.

Moreover, in this method a media plane probe also generates a correlation key for each media stream monitored by the media plane probe. The media plane probe generates the correlation key using a timestamp indicative of the point in time of the end of the respective media stream as detected by the media plane probe and using information on the media data path of the respective media stream. The media plane probe stores a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

A correlation unit further correlates the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports.

In a more detailed exemplary embodiment of the invention, the generation of the respective correlation keys for the respective media streams by the media plane probe is based on the media data path information of the respective media stream and the timestamp indicative of the point in time at which the respective media stream ended.

According to another exemplary embodiment of the invention, information on the direction of the respective service – based on which the correlation keys are generated by signaling plane probe for a respective service – is information on one or more media data paths of media streams of the service learned from messages of the signaling session of the service. In other words, the direction of the service may be for example given by the media data path which is yielding a given direction. For example in a simple VoIP call between two peer devices, there are two directions: caller-to-callee and callee-to-caller as noted above.

In a further embodiment of the invention, the signaling plane probe monitors signaling messages of the respective signaling sessions of the services, and detects media streams belonging to each service associated to a respective monitored signaling session based on the signaling messages of the respective signaling sessions. The correlation keys for a respective one of the signaling sessions may be generated by the signaling plane probe

based on information on the media data path of the media streams belonging to the service associated to the respective one of the signaling sessions and the timestamp indicative of the point in time at which the end of the respective signaling session has been detected by the signaling plane probe. For example, in one possible implementation information on a media data path of a media stream is a destination IP address and port number of the media stream.

In another embodiment, the timestamps used for the generation of the correlation keys by the signaling plane probe and the media plane probe could be for example rounded (e.g. in a configurable or fixed interval) before the generation of the correlation keys, in order to ensure that timing differences in the events monitored in the media plane and signaling plane of the services do not lead to generation of non-matching correlation keys for a given service. Alternatively, in another exemplary implementation, the timestamps – by default – have a granularity that is higher than one second, e.g. a granularity in the range of one to ten seconds.

In a further embodiment, instead of storing a single correlation key per direction of a signaling session or per media stream, additional correlation keys may be generated and stored. The input for these additional correlation keys are the same IP address and port information, but a different timestamp. For example, the signaling plane probe could calculate an additional correlation key using a timestamp earlier than the timestamp indicative of the end of the respective signaling session.

In case of using a rounded timestamp this earlier timestamp may be the next earlier rounded timestamp before the rounded timestamp indicative of the end of the respective signaling session. In one exemplary implementation, the signaling plane probe could generate two timestamps from a (non-rounded) timestamp indicative of the end of the respective signaling session which are each used together with the same IP address and port information to generate two respective correlation keys for a given direction: One of these two generated timestamps used for calculation of the correlation keys may be the timestamp indicative of the end of the respective signaling session rounded to the nearest lower value according to the rounding interval, while the other of these two generated timestamps may be the timestamp indicative of the end of the respective signaling session rounded to the nearest higher value according to the rounding interval.

All correlation keys that are generated by the signaling plane probe for a given direction may be stored as part of the signaling session records.

In a similar fashion as described above for the signaling plane probe and the directions indicated in the signaling, also the media plane probe may optionally generate more than one correlation key for a given media data path/media stream. In case of using rounded timestamps also the media plane probe could generate two timestamps from a (non-rounded) timestamp indicative of the end of the respective media stream which are each used together with the same IP address and port information of the media stream to generate two respective correlation keys. One of the two timestamps generated by the media plane probe may be the timestamp indicative of the end of the respective media stream rounded to the nearest lower value according to the rounding interval, while the other of these two generated timestamps may be the timestamp indicative of the end of the respective media stream rounded to the nearest higher value according to the rounding interval. All correlation keys that are generated by the media plane probe for a given media stream may be stored as part of the media stream records.

In case of rounding timestamps in the media plane probe and the signaling plane probe, the rounding interval in both probes should be identical. The gain of having both probes generating additional correlation keys and correlating a multitude of correlation keys may not reflect in the increased processing needs.

In a further exemplary embodiment of the invention, the correlation of the media streams and the signaling session of the respective services is performed by a correlation unit periodically or in response to a query.

Another exemplary embodiment of the invention is related to a method for correlating media streams and signaling sessions of services. According to this method a correlation unit correlates the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services. For correlating media streams and the signaling sessions, the correlation unit identifies one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports. Further, the correlation unit generates a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained in the signaling session report of said at least one service.

In one exemplary implementation, these steps are performed by a correlation unit periodically for all signaling session reports and media stream reports generated within a configurable time interval, or in response to a request.

In case the correlation unit is unable to match one or more media streams to said signaling session report, the correlation unit may optionally generate a service record without quality information of the media streams of the service or may optionally generate a corresponding indication of the user/operator of the correlation unit.

A service report generated for a service by the correlation unit may comprise – according to one exemplary implementation – information on the service and the quality information of the media streams of the service.

In one exemplary implementation, the correlation keys provided in a respective signaling session report may have been generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the corresponding service has been detected and using information on the media streams of the service. Similarly, the correlation keys provided in a respective media stream report may have been generated using a timestamp indicative of the point in time at which the respective media stream ended and using information on the media data path of the media stream.

In one further exemplary implementation, the correlation unit first attempts a correlation of media stream reports and signaling session reports using the correlation keys that have been generated based on the timestamps indicative of the end of the signaling session, respectively, the end of the media stream(s). In case the correlation is unsuccessful, the correlation unit may further use optionally available additional correlation keys contained in the signaling session reports and/or media plane reports which are based on other timestamps, e.g. before and after the actual end of the respective signaling session or media stream(s), respectively. This layered approach allows to first conduct the correlation using the correlation keys which are more likely to yield a result and as a fallback conduct the correlation using a wider timeframe.

According to a more detailed exemplary embodiment of the invention, the correlation of the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions by the correlation unit comprises matching the respective correlation keys of the media stream reports of media streams are to the respective correlation keys of the signaling session reports, or vice versa. A match of the correlation keys of a one of the media stream reports and one of the signaling session reports (i.e. the

two correlation keys are identical) yields that the media stream of said one media stream report belongs to the same (media) service to which the signaling session of said one signaling session report belongs.

In another exemplary embodiment of the invention, for correlating the media streams and the signaling sessions of one or more services by the correlation unit, same may obtain from a signaling session report correlation keys identifying the media streams belonging to the same service as the signaling session. Further the correlation unit could then identify those media stream reports which include one of the correlation keys obtained from the signaling session report.

Alternatively, in another exemplary embodiment, the correlation of the media streams and the signaling sessions of one or more services could also be implemented by the correlation unit obtaining from a media stream report a correlation key identifying the media stream, and further the correlation unit identifying a signaling session report among the signaling session reports which includes the same correlation key obtained from the signaling session report.

According to one exemplary embodiment of the invention, a method for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network is provided. It is assumed for exemplary purposes that one service has one signaling session.

Another embodiment of the invention is related to a passive monitoring system for correlating media streams and signaling sessions of services in a packet-switched network. The passive monitoring system comprises a signaling plane probe that generates at least one correlation key for each direction of a service based on the information of the signaling session of the respective service monitored at the signaling plane probe. As noted above, there may be one or more correlation keys for each direction of the service. The correlation keys for a respective service are generated by the signaling plane probe using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and further using information on the direction of the respective service (e.g. information on the media data path of the given direction, such as a destination IP address and UDP port number). The signaling plane probe stores a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service.

Furthermore, the passive monitoring system also comprises a media plane probe for generating a correlation key for each media stream monitored by the media plane probe. The

correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream. The media plane probe stores a media stream report for each media stream in a database. Each media stream report comprises at least the correlation key generated for the respective media.

The passive monitoring system further comprises a correlation unit for correlating the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports. The correlation unit may for example use one of the different correlation methods described herein for this purpose.

Another embodiment of the invention is related to the design of the signaling plane probe. According to this embodiment of the invention, the signaling plane probe comprises a network interface card for receiving packets of signaling sessions of multiple media services and a monitoring unit for monitoring the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of the signaling sessions. The monitoring unit of the signaling plane probe is further capable of detecting the point in time at which a respective signaling session of the respective service terminates.

The signaling plane probe also comprises a processing unit for generating at least one correlation key for each direction of a service based on a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the monitoring unit and based on information on the direction of the respective service. Moreover, the monitoring unit stores a signaling session report for each signaling session in a database, wherein each signaling session report comprises at least two correlation keys generated for the respective directions of the service.

Another embodiment of the invention is related to a media plane probe which comprises a network interface card for receiving packets of media streams of multiple media services, and a monitoring unit for monitoring the media streams based on the packets of the media streams received by the media plane probe. Thereby the media plane probe obtains information on a media data path a respective one of the media stream. Further, the monitoring unit detects the point in time at which a respective media stream of the respective service terminates, and a processing unit of the media plane probe generates a correlation key for each media stream monitored by the monitoring unit. This correlation key may be for example generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream.

The media plane probe's monitoring unit further stores a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

The invention according to another embodiment is related to a correlation unit capable of correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network. The correlation unit comprises a database access unit for accessing a database comprising signaling session reports for signaling sessions of services and a database comprising media stream reports of media streams of the services. The signaling session reports each comprise at least one correlation key for each direction of a respective service, and the media stream reports each comprise one correlation key indicating the direction of the media stream. A processing unit of the correlation unit is used to correlate the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services. The processing unit identifies one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports. The correlation unit may further comprise a service record generation unit (e.g. implemented by a processing unit) for generating a service record for a respective service based on information comprised in the media stream reports and information comprised the signaling session report of said at least one service.

In a further embodiment of the invention, the correlation unit is further adapted, e.g. by comprised respective configured means, to perform the steps of the correlation methods according to one of the different embodiments described herein. Furthermore, also the signaling plane probe and the media plane probe may be adapted, e.g. by comprised respective configured means, to implement the different functions thereof as described herein.

Further embodiments of the invention relate to the implementation of the different functions of media plane probe, signaling plane probe and correlation unit in software. In this context, another embodiment of the invention provides a computer-readable medium storing instructions that, when executed by a processor of a signaling plane probe, cause the signaling plane probe to receive packets of signaling sessions of multiple media services, and to monitor the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of the signaling sessions. The execution of the instructions further cause the signaling plane probe to detect the point in time at which a respective signaling session of the respective

service terminates, and to generate at least one correlation key for each direction of a service based a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected and based on information on the direction of the respective service. Moreover, the executed instructions may further cause the signaling plane probe to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service.

Another exemplary embodiment of the invention is providing a computer-readable medium storing instructions that, when executed by a processor of a media plane probe, cause the media plane probe to receive packets of media streams of multiple media services, and to monitor the media streams based on the packets of the media streams received by the media plane probe to thereby obtain information on a media data path a respective one of the media streams. Moreover, the execution of the instructions cause the media plane probe to detect the point in time at which a respective media stream of the respective service terminates, and to generate a correlation key for each media stream monitored by the monitoring unit, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream. Furthermore, the execution of the instructions may further cause the media plane probe to store a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

Another embodiment of the invention is providing a computer-readable medium storing instructions that, when executed by a correlation unit, cause the correlation unit to correlate media streams and signaling sessions of services in a passive monitoring system of a packet-switched network. The execution of the instructions cause the correlation unit to access a database comprising signaling session reports for signaling sessions of services and a database comprising media stream reports of media streams of the services. The signaling session reports each comprise at least one correlation key for each direction of a respective service, and the media stream reports each comprise one correlation key indication the direction of the media stream. The execution of the instructions further cause the correlation unit to correlate the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services and to identify one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports.

The execution of the instructions also cause the correlation unit to generate a service record for a respective one of said at least one service based on information comprised in the media stream reports and information comprised the signaling session report of said at least one service.

Another embodiment is providing a computer readable medium storing instruction that, when executed by the processor of the correlation unit, cause the correlation unit to perform the steps of the method for correlating media streams and signaling sessions of services according to one of the different embodiments described herein.

BRIEF DESCRIPTION OF THE FIGURES

In the following the invention is described in more detail in reference to the attached figures and drawings. Similar or corresponding details in the figures are marked with the same reference numerals.

Fig. 1 shows exemplary situations where multiple media streams may be triggered in a media service,

Fig. 2 shows three geographically distributed POPs (Point of Presence) which are linked over the carrier's internal network,

Fig. 3 shows an exemplary network using internal and external addresses,

Fig. 4 shows an exemplary overview on a system according to an embodiment of the invention comprising signaling plane probe, a correlation unit and a media plane probe in a packet-switched network,

Fig. 5 shows a flow chart of the operation of the correlation unit according to an exemplary embodiment of the invention,

Fig. 6 shows the structure of a correlation unit according to an exemplary embodiment of the invention,

Figs. 7 & 8 shows two exemplary flow charts of the operation of a signaling plane probe according to different exemplary embodiments of the invention,

Fig. 9 shows the structure of a signaling plane probe according to an embodiment of the invention,

Fig 10 shows the structure of a media plane probe according to an embodiment of the invention,

Fig. 11 shows a flow chart of the operation of a media plane probe according to an embodiment of the invention,

Fig. 12 shows a network in which locations are indicated, where signaling and media plane probes of the invention may be advantageously located according to an embodiment of the invention,

Fig. 13 illustrates the use of rounded timestamps within a signaling plane probe and a media plane probe to highlight potential sources of error in correlation, and

Fig. 14 illustrates the use of two rounded timestamps for generating respective correlation keys according to an exemplary embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

The following paragraphs will describe various embodiments of the invention. For exemplary purposes only, most of the embodiments are outlined in relation mainly to VoIP services as an example for a media service, but the invention is not limited to VoIP services. For example, the invention may also be used for other types of media services provided in packet-switched networks. The invention may be advantageously used in media services which have signaling plane and media plane separated and/or taking different routes through the packet-switched network.

The explanations given in the Technical Background section above are intended to better understand the often VoIP call-related exemplary embodiments described herein and should not be understood as limiting the invention to the described specific implementations of processes and functions in a packet-switched communication network.

As indicated above, the invention generally relates to methods for correlating at least one media data path and a signaling session of a service in a passive monitoring system that is monitoring a packet-switched network. Furthermore, the invention also relates to an implementation of these methods in hardware and software, and provides a signaling plane probe, a media plane probe and a correlation unit, as well as passive monitoring system comprising same. Based on the assumption of disjunctive signaling and media components of a service, the correlation mechanism suggested herein supports correlation across

multiple hosts connected via a network (e.g. a Local Area Network (LAN), Wide Area Network (WAN) or the Internet).

One of the various aspects of the invention is to suggest a definition of correlation keys that allow the correlation of media streams of services and the associated signaling session of the respective services based on data records (reports) that provide information on the media plane and signaling plane of the respective services. In the following these correlation keys are implemented as in form of hash keys for exemplary purposes. The correlation keys may be used to query database(s) to correlate media stream reports for media streams monitored by a media plane probe and a signaling session report on the signaling session provided by a signaling plane probe that belongs to the same service – it should be noted that for the proper functioning of the invention, it is not decisive whether media stream reports and signaling session reports are stored in the same or different databases, or whether there are even multiple database partitions provided. The media stream reports and signaling session reports may even be stored in files.

The correlation keys are generated by one or more media plane probes and one or more signaling plane probes that monitor a packet-switched network independently from each other. This independent calculation of the correlation keys facilitates distributed correlation spanning across multiple geographically distributed components or even across multiple vendors of software components. Furthermore, as will become more apparent from the following, for a given media data path – respectively a corresponding given direction –, the parameters that are used for generating correlation keys at the signaling plane probe(s) and the media plane probe(s) ensure that the identical correlation keys are generated for the given media data path by the signaling plane probe(s) and the media plane probe(s) to facilitate correlation. A respective correlation key may be associated to a given direction of the media data paths of the media streams belonging to a given service. For example, assuming for exemplary purposes a simple bi-directional service, e.g. VoIP call, there are two directions: caller-to-callee and callee-to-caller.

The signaling plane probe may generate a correlation key for each direction of the service. For example, the signaling plane probe generates one correlation key for the caller-to-callee direction and another correlation key for the callee-to-caller direction. However, the signaling plane probe may also generate multiple correlation keys per direction, as will be outlined below in further detail. Consequently, in case it is assumed that a service comprises at least two media streams of different directions, then at least two correlation keys are calculated by the signaling plane probe for later correlation of the media plane and signaling plane of the service. The correlation key for a given direction or media data path may be generated based on information on the direction or media data path, such as for example a tuple comprising

the destination IP address and UDP port for a media stream in the service's media plane corresponding to the given direction or media data path. In addition, the generation of the correlation key is further based on a timestamp. With respect to the generation of the correlation key by the signaling plane probe, the timestamp may for example indicate the point in time at which the signaling plane probe has detected the end of the signaling session (and thus the service). In another exemplary implementation, the timestamp used in the generation of a correlation key for a media data path of a media stream in the media plane of the service could also be the point in time of the termination (end) of the media stream in the media plane detected by the signaling plane probe.

The signaling plane probe may generate the correlation keys for the different endpoints (the endpoint indicate a direction of the related media data flow) of a service upon detection of the end of a signaling session and additionally or alternatively, upon detection of the end of a media stream in the service's media plane. Typically but not limited thereto, there is one endpoint per direction, e.g. the caller and callee. The callee may be represented by multiple endpoints in case of more complex media service scenarios, e.g. when session forking is used (see Fig. 1).

The end of a signaling session may be for example detected in the signaling plane probe by detecting a signaling protocol specific hang-up message, a terminal failure to connect the involved endpoints or a session timeout. Each correlation key (e.g. hash key) is then generated by the signaling plane probe based on a timestamp indicative of the point in time at which the end of the respective signaling session has been detected by the signaling plane probe and using information on respective direction of the media plane data flow (e.g. caller-to-callee or callee-to-caller for a simple bi-directional service, e.g. VoIP call).

The direction may be for example given by a media data path of a media stream of the respective service. A media data path could be for example indicated by a tuple comprising the destination IP address and UDP port number for transmitting the media packets of a media stream of the service. The signaling plane probe may thereby generate only one single correlation key per direction or multiple correlation keys per direction. In one exemplary embodiment of the invention, the signaling plane probe generates one correlation key per direction, where the correlation key for a respective direction is based on a media data path of a media stream. This media stream of which the media data path information is used to generate the correlation key should be existent at the end of the signaling session (according to the session signaling monitored by the signaling plane probe), in order to be able to match the correlation key to a corresponding correlation key in a media stream report generated by a media plane probe. The media data path(s) of a service and optionally their

existence (e.g. start time and termination time) could be for example learned by the signaling plane probe from the messages of the signaling session of the service.

The correlation keys generated for the different directions of the service corresponding to the signaling session the end of which has been detected are then comprised in a signaling session report, which may be for example stored in a database or (flat) file.

Alternatively, the signaling plane probe may also generate a correlation key for each media stream in the media plane of the service. In this case it may be assumed that the signaling plane traffic, i.e. the packets of the signaling protocol of the signaling plane, allow the signaling plane probe to detect the termination (i.e. end) of respective media streams in the media plane. In this case, the detection of the end of a media stream could trigger the generation of a correlation key for the respective media stream based on information on its media data path and a timestamp indicating the point in time at which the signaling plane probe detected the termination of the media stream.

Further, in another alternative implementation, the detection of the end of a signaling session, respectively, service as well as the end of a media stream as detected by a signaling plane probe could be used to trigger the generation of correlation keys in the signaling plane probe. For example, during the ongoing service the detection of the termination of a media stream in the media plane as learned from the signaling plane traffic could trigger the signaling plane probe to generate and include respective correlation keys for the media streams terminated during the ongoing session into the signaling session report of the service, while the detection of the end of the signaling session, and thus the service, could be used as a trigger for generating correlation keys for the respective media streams being "alive" in the media plane of the service at the signaling plane probe and their inclusion into the signaling session report of the service.

In the media plane, the media plane probe generates a correlation key (e.g. hash key) for each media stream. The correlation keys for the media steams are generated upon the media plane probe detecting the end of the respective media stream. The correlation key for a respective terminated media stream is calculated by the media plane probe based on the direction of the respective media stream, e.g. based on the media data path of a media stream, and on a timestamp indicative of the point in time at which the respective media stream ended. This point in time is for example the point in time at which the last packet of the media stream has been received at the media plane probe. In this latter case, the media plane probe may thus maintain a timestamp indicating the time of reception of the last packet of the media stream at the media plane probe.

The timestamp used for the generation of the correlation key may indicate another point in time than the point in time at which the media plane probe detects the end of the media stream. For example, in a packet switched network, there may be interruption in the packet stream of the media stream so that there may be a timer that measures the time since the last reception of a packet of the media stream. In case the time since the last reception of a packet of the media stream is exceeding a threshold (e.g. the timer "expires") – for example after 15 or 30 seconds, or the like – then the media plane probe detects the end of the media stream. The timestamp used in the correlation key generation may nevertheless be the point in time at which the last packet of the media stream has been received. The media plane probe stores the correlation key (e.g. hash key) for the media stream in a media stream report, which may be for example stored in a database or exported as a flat file.

Furthermore the timestamp used for the generation of a correlation key may be rounded. This counters differences in the time basis (local time) of multiple signaling and media probes as slightly different timestamps still yield the same rounded timestamp. Additional correlation keys may be generated with different timestamps, e.g. indicative of the – optionally rounded – end of a respective signaling session and/or media stream or a timeframe before or thereafter. These additional correlation keys may be stored in the respective signaling session reports and/or media session reports.

A media stream may also be interrupted due to codec features such as Voice Activity Detection (VAD). If a sender of a media stream is mute the codec may not encode any data resulting in an interruption of the media stream. This is typically indicated in the packet stream transporting the media and should not be considered by the media plane probe as an indication of the end of the media stream itself. In case a timer mechanism is used to detect terminated media streams, the timer should be configured such that codec features that lead to a temporary interruption of the media stream are not triggering the termination of the media stream (according to the detection mechanism for detecting terminated media streams in the media plane probe).

Furthermore, in another exemplary embodiment, the correlation of media streams of services and the signaling sessions of the services may also be based on a compound keys used as a correlation keys. This exemplary alternative implementation may however require a more complex database structure of the media stream reports and signaling session reports, as will be outlined below.

Another further aspect of the invention is the use of the correlation keys comprised in the signaling session reports and the media stream reports to identify the media streams and signaling session belonging to the respective services. This correlation of media streams and

signaling session could be done for example periodically or on request. The correlation keys generated by the signaling plane probe and the media plane probe should be matching each other (i.e. the correlation keys of a given signaling session report should be found in the media stream reports, respectively, the correlation key of a media stream report should be found in one of the signaling session reports). Hence, a correlation unit can identify which media streams belong to which signaling session (or vice versa) based on the correlation keys contained in the media stream reports and the signaling session reports and may therefore generate a service record for a respective service based on information contained in the media stream reports and information contained the signaling session report of the service.

The correlation methods according to the different embodiments of the invention proposed herein are based on a probe monitoring the signaling plane traffic of services (i.e. the signaling plane probe) identifying new services and associated media data paths of the media plane by detecting their corresponding signaling sessions and determining for each service correlation keys allowing to identify at least the media streams of the media plane existent upon termination of the service. It should be noted that "at least the media streams" means that it is optionally possible to generate a correlation key for each media data path of the media plane learned from the signaling plane traffic of the service, in case it is possible to detect the termination of the media data paths of the media plane in the signaling plane traffic. In this case the correlation key for a respective media stream may be based on the information on the media data path of the respective media stream visible in the signaling plane of the service and the point in time at which the signaling plane probe detects the termination of the media stream in the signaling plane.

Regarding the use of the point in time of the termination of the signaling session (e.g. a signaling message in the signaling plane indicating "hang up") for the generation of correlation keys it should be noted that the termination of the signaling session also means that all existent media data paths, respectively, the corresponding media streams of the media plane are terminated. Hence, in case it is not possible to detect the termination of the media data paths of the media plane in the signaling plane traffic, it is assumed to be possible to detect at least the end of the signaling session and to thereby conclude on the termination of the respective media data paths of the media plane and to thereby facilitate correlation of the media streams that have been "alive" at the end of the service and the corresponding signaling session of the service.

As noted above, another possibility may be to use compound keys as correlation keys. A compound key is defining a set of information comprised e.g. in a given set of fields of a database entry of a signaling session report and media stream report that – together with

additional criteria on the point in time of the termination of a media stream or the service – allows to correlate media streams and signaling sessions of the services. In one exemplary embodiment, the one or more databases that are used to store the signaling session reports and the media stream reports are partitioned into partitions covering only a given amount of time (i.e. time interval) of signaling session reports and/or media stream reports, respectively. The "size", i.e. the time interval covered by the partitions may be for example determined based on the average time interval in which port reuse occurs in the target monitored packet-switched network (e.g. the partition could cover ½ of the average time interval in which port reuse occurs). In the given partitions it may thus be assumed that the media data path information, such as destination IP address and UDP port number are unique. In this case, the signaling session reports and media stream reports may not include separate correlation keys, but the correlation key can be generated on-demand as a compound key from the information on the media data paths contained in the signaling session reports and media stream reports.

Hence, it would be sufficient that the signaling session reports and media stream reports store information on the media data paths of the service (e.g. destination IP address and UDP port number) and a timestamp that indicates the point in time of termination of the respective media data path in the media plane, as detected by the signaling plane probe from the monitored signaling sessions and the media plane probe from the monitored media streams. As noted above, the signaling plane probe can also detect the point in time of termination of the media data path(s) in the media plane by detecting the point in time of termination of the signaling session, which is – in most practical cases – coinciding with the termination of all media streams of the service in the media plane as well.

In this case, the correlation could be realized by first selecting the database partition of the media plane reports according to a timestamp of the termination of a media data path/signaling session within a given signaling session report, and by then determining the media stream reports that comprise media stream information matching those indicated in the given signaling session report.

Fig. 4 shows a packet switched network 400 according to an exemplary embodiment of the invention that comprises a signaling plane probe 401, a correlation unit 402 and a media plane probe 403. Please note that more than one signaling plane probe 401 and media plane probe 403 may be provided. Optionally, there may be multiple correlation units 402 present, working independent of each other and not impacting the performance or correlation accuracy of any other component.

The signaling monitoring probe 401 and the media plane probe 402 are monitoring the traffic on a packet-switched network 400, which may be for example an IP based packet switched network. It is assumed for exemplary purposes that the signaling plane probe 401 and the media plane probe 403 monitor the traffic of the packet-switched network 400 at different locations of the packet-switched network 400, but observe the same network traffic.

In one exemplary implementation, the signaling plane probe 401 and the media plane probe 403 are probes of a minimal non-intrusive, passive monitoring system. The respective monitoring probes 401, 403 are connected to the packet-switched network 400 via respective test access ports (TAPs). A TAP is a passive network device that mirrors network traffic of the packet-switched network without interference of the original network traffic. It provides a copy of every packet sent or received on the network, by duplicating the data packets observed on the physical layer of the network. Instead of a TAP a Switch Port Analyzer (SPAN) port may be used to monitor the traffic of the packet-switched network.

In packet-switched environments, the signaling data of the signaling plane and media data of the media plane is packaged into packets on the transmitter side. Packets are transported over the packet-switched network independent of each other. The receiver of the packets – identified by the destination IP address and port – may have to regroup and sometimes reorder packets in order to compile streams from single packets.

On the signaling plane multiple packets belong to one signaling session. The first packet of a signaling session – usually containing a "call setup" message – starts a new signaling session (respectively a new service). At this point, the signaling plane probe 401 generates a session entry along with auxiliary data, i.e. a unique correlation key identifying the session internally (signaling session identifier). Packets following the initial signaling packet are matched to one of the currently ongoing signaling sessions (respectively services) and signaling sessions managed by the signaling plane probe. Packets which do not start a new signaling session and do not belong to existing sessions may be discarded.

Matching single signaling packets to existing signaling sessions is specified in the standards which define the signaling protocol itself. Transport layer information like IP addresses and protocol-specific information (e.g. SIP Call-IDs and tags) may be used.

A signaling session may be torn down upon reception of a special packet – a "call hang-up" message. This message again is specific to the signaling protocol used. Upon reception of this message the signaling plane probe 401 can conclude on the termination (end) of the signaling session and generates the correlation keys for use in correlation of media plane and signaling plane of the services. Furthermore, the signaling session may be removed

from its internal storage. Additionally the signaling monitoring component may remove sessions from its internal storage, if it identifies stale sessions.

The media plane 403 of the service may for example identify the media plane data in form of RTP streams. RTP streams are – analog to the signaling sessions – made up from individual RTP packets. The transmitter of an RTP stream generates RTP packets by splitting the media data into smaller pieces. In general these packets are relatively small to reduce the delay between sender and receiver. As a result many packets are created and form a single RTP stream – typically 50 packets per second. Single RTP packets can be grouped into an RTP stream by media plane probe 403. When using RTP, the media plane probe 403 has transport layer information as well as the RTP protocol header which contains a so called Synchronization source (SSRC). The SSRC is a 32-bit numeric identifier which is contained in every single packet and deemed mandatory by the RTP. When media plane probe 403 first receives a packet which does not match an existing RTP stream in its internal storage it creates an RTP stream. Consecutive RTP packets can then be matched to the existing streams by comparing the transport layer information and SSRC value. If no packets are received for a single RTP stream for a certain time the stream is considered ended and can be removed from the internal storage. A short interruption in the packet stream should not trigger the media plane probe 403 to mark an RTP stream as ended, thus the time before a media plane probe 403 can consider an RTP stream as ended is multiple seconds and has to take features like voice activity detection into account.

The signaling plane probe 401 and the media plane probe 403 generate information about the signaling session and media streams, respectively. This generated information is comprised in signaling session reports and media stream reports, respectively. In order to track running sessions and media streams the entities 401, 403 store information internally while the signaling sessions and media streams are in progress. When a signaling session or media stream is terminated the respective monitoring component 401, 403 may generate a signaling session report with signaling information, respectively media stream report with information on the media stream. As will be outlined below in further detail, the signaling session report provides a summary of the signaling session (service) and includes the correlation keys (e.g. hash keys) for correlation. Similarly, the media stream reports summarize the media stream and also include a respective correlation key (e.g. hash key) to facilitate correlation.

Each signaling session report for a signaling session may be for example written by the signaling plane probe 401 upon completion of a signaling session and contains data which describes the signaling session it monitored. For example, in one embodiment of the invention, the data stored in the signaling session report and describing the signaling session

include – but are not limited to – a set of caller and callee specific properties which will be available per endpoint (caller or callee). These endpoint-specific properties may for example include at least one of the following:

- endpoint address/name (e.g. SIP URI),
- phone number,
- User-Agent (endpoint sóftware and version),
- capabilities/features,
- supported media types,
- IP addresses and ports for signaling and media streams (from SDP body), and
- hash key of the media path (for example described by the destination IP address and UDP port of the expected media stream)
- signaling routing information.

Additionally, one or more of the following signaling session properties which are not endpoint-specific may also be included in the signaling plane report:

- timestamps: e.g. session start, Post-Dial-Delay (PDD), connect time, session end;
- information on the session flow (e.g. signaling message(s) with timestamps),
- session state (e.g. successful, unsuccessful, timed out, reason for failure), and
- events during session (media on hold, codec renegotiation, session forking, authentication)

In one exemplary embodiment of the invention, the signaling session report comprises all of the above noted endpoint-specific properties, and – optionally – all signaling session properties listed above, as far as they are available from the signaling protocol messages.

Depending on the specific implementation and capabilities of the signaling plane probe, the signaling session reports may contain information about multiple endpoints and/or multiple media streams per endpoint. In this case the session end timestamp may be available per involved endpoint, and not per session.

For example, the most trivial VoIP service is a phone call between two parties (including two media streams, one in each direction). In this case, the signaling session report could store each party as one endpoint with the above mentioned information and the media stream's destination IP address and port information. If the VoIP service would contain a video call (i.e. one audio and video stream per direction) the endpoint information may be extended to contain the destination IP address(es) and port numbers of the two media streams per direction (one IP address & port tuple for audio, one IP address & port tuple pair for video). Considering a VoIP service, the respective tuples of destination IP address and UDP port number of the media streams of the session may be for example learned by the signaling plane probe from the SDP body of SIP signaling.

Likewise, the signaling plane probe may extend the signaling session report to include additional endpoints, if multiple endpoints were part of the VoIP service. Multiple endpoints could be for example involved in a VoIP service, in case a caller tries to reach a callee which has call forwarding configured. The caller may then hear a ring-tone from the callee's landline, followed by a ring-tone from the callee's mobile phone. Finally a voicemail box picks up the session generating a third endpoint on the callee's side.

The media plane probe 403 monitors the media traffic, i.e. the individual media streams that convey the media information of the different services. The media plane probe 403 identifies a new media stream of a media service based on the network traffic it observes on the monitored packet switched network 400. The identification of a new media stream for a given service may happen shortly after this service is set up on the signaling plane. Throughout the duration of the media stream information on the media stream may be stored in an internal data record on the media steam. The media plane probe 403 may note for each stream the last RTP packet received for the respective media stream in an internal data record on the media stream in order to allow detecting the media stream's termination.

The media plane probe 403 may optionally perform a real-time analysis of the media streams monitored on the packet switched network and may generate a fine-grained media stream quality report with media stream related information. Fine-grained media stream quality reports could be for example written in predetermined time intervals, e.g. every five seconds, and could provide a summary of the media stream in the respective interval as well as quality information of the media stream in the respective interval. The information of the fine-grained media stream quality reports of a media stream could be used by the media plane probe 403 to generate a single media stream report on the media stream upon the end of the media stream is detected by the media plane probe 403. The media stream report and optionally also the fine-grained media stream quality reports of a media stream may be stored in individual databases.

The media stream report generated by the media plane probe 403 upon detecting the end of the media stream may contain data which describes the unidirectional media stream it monitored as well as correlation key that can be for example generated based on the destination IP address and UDP port number of the media stream and a timestamp that indicates the point in time at which the media stream ended/terminated. The media plane probe 403 could for example learn the destination IP address and UDP port number from the IP header encapsulating the RTP packets on the network layer. The characteristics of each media stream can be grouped into two categories. The transport and RTP header information characterize the media stream, e.g. for correlation purposes, whereas the quality information enable quality related processing of the media stream. In general the transport information is the same for all media packets processed as part of a single media stream. The stream quality information includes information covering all media packets and may be aggregated or stored on a per-packet basis.

In one exemplary embodiment of the invention, the media stream report generated by the media plane probe 403 for a given media stream may for example include one or more of the following transport media information:

- source and destination IP address,
- source and destination UDP port,
- RTP header Synchronization source (SSRC) value,
- RTP payload type, and
- correlation key(s) (e.g. hash value(s)) calculated based on destination IP address, destination UDP port and a timestamp indicating the point in time at which the of the media stream terminated (e.g. the timestamp of the last packet received for the media stream) or timestamp indicative of the timeframe before or after the termination timestamp.

The stream quality information that may be added to the media stream report by the media plane probe 403 may for example include/indicate one or more of the following:

- timestamp information (e.g. start and/or end time),
- RTP format violations,
- protocol conformance (e.g. sequence errors and/or DSCP errors),

- transport information (e.g. packet interval and jitter statistics, information on packet loss, information on duplicate packets, etc.), and

- MOS values (e.g. per time-slice (e.g. every 5 seconds), minimum/average/maximum value for entire stream duration, etc.),

- events (e.g. comfort noise, silence suppression, DTMF tones, codec changes, marker bits).

In one exemplary embodiment of the invention, the media stream report comprises all of the above noted transport media information and optionally further all stream quality information listed above, as far as they are available to the media plane probe from the packets of the media streams.

In addition to "measured values" (e.g. jitter between multiple media packets) the media stream report may also contain calculated information like MOS (Mean Opinion Score) values and information about impairments experienced by the receiver of the stream.

As further shown in Fig. 4, a correlation unit 402 correlates the media stream reports and signaling session reports generated by the signaling plane probe 401 and the media plane probe 403 in order to identify the media streams and the signaling session that belong to a respective service. It should be noted that each service is assumed to have one signaling session, but may have one or more media streams.

Generation of Correlation Key

The correlation key that is used for correlation may be built from:

- the timestamp that indicates the end of the signaling session, respectively the end of a given media stream, and

- information on a media stream, such as for example the destination IP address and UDP port of the media stream.

The correlation of signaling plane and media plane by the correlation unit 402 is based on this compound correlation key.

As noted earlier, the correlation key may be for example a hash key generated based on the before-mentioned parameters. The hash function to calculate the hash key (sometimes also called hash value) from the given data should be known to all parties involved (i.e. signaling plane probe 401, and media plane probe 403 and optionally also correlation unit 402) so that

the result for a given input is the same for all parties. Therefore, the hash keys are determined by the signaling plane probe 401 and media plane probe 403 using media data path information in addition to a timestamp that indicates point in time when the signaling session, respectively, the media stream ended/terminated. The correlation unit 402 may generate additional correlation keys based on the timestamps indicative of the timeframe before or after the timestamp that indicates the point in time when the respective session ended.

The hash function to determine the hash key may be a simple concatenation of one or more of the input parameters. For example, a concatenation of a destination IP address, destination UDP port number and a timestamp may (optionally including a delimiter character in between the parameters) be used for defining the hash keys. Of course more complex mechanisms may also be used to calculate a hash key representing the signaling session identifiers and/or the media path identifiers. For example the MD5 message digest function (see IETF RFC 1321, "The MD5 Message-Digest Algorithm", 1992, available at http://www.ietf.org) could be used to create a valid hash key. The MD5 message digest function may be fed with a concatenation of a destination IP address, destination UDP port number and the timestamp indicating the time of termination of the signaling session, respectively, the media stream and the resulting string can be used as a correlation key. Another hash function that could be used may be selected from the SHA cryptographic hash function family (see FIPS PUB-180, "Secure Hash Standard (SHS)", 2008 available at http://www.itl.nist.gov/fipspubs/). Furthermore, in another embodiment, a linear hash function may be used.

As outlined with respect to Fig. 1 earlier herein, ephemeral media streams may exist at any point during a service, such as a VoIP service. The ephemeral media streams may be generated by early media streams, service forking and codec changes. As a result the start and end timestamps of media streams may not necessarily concur with the start and end time of the VoIP service itself. Yet the timestamps of the start or end of a signaling session or media stream are an important piece of information for a successful correlation, as it significantly reduces the search window and amount of data to search. In typical scenarios, it can be expected that the start time of signaling and media streams deviate significantly and are thus not ideal candidates for a correlation mechanism.

At the end of a VoIP service both caller and callee have an established signaling session and send each other media on the media plane. Once a party hangs up the session a hang-up signal is send on the signaling plane and at the same time stops sending media packets. Once the remote party receives the hang-up signal it also stops transmitting media. The timestamps of the end of the signaling session of a service and the end of the corresponding

media streams in the media plane of the service are commonly very close to each other compared to the start of a media service in signaling plane and media plane. All these events will be subject to the same network- and application-imposed delays. Still the time-frame in which these happen is typically in the few milliseconds range or in the sub-second-range on long-distance or low bandwidth networks. This makes the timestamps of the end of the signaling session and the end of the media streams a good choice for calculating a correlation key.

Nevertheless, since signaling plane probe 401 and media plane probe 403 independently determine the correlation keys based on timestamps (and even minor deviations of the timestamps in the millisecond-range could lead non-matching correlation keys), the granularity of the timestamps used for the generation of the correlation keys may be chosen to account for those deviations of events in signaling plane and media plane events (i.e. deviations in the end of the signaling session in the signaling plane and the end of the corresponding media streams in the media plane of a service). Therefore the timestamps could have a granularity for example in the range of one or more seconds or at least 250 milliseconds or 500 milliseconds. In one exemplary embodiment, UNIX timestamps (which is indicating the seconds passed since January 1, 1970, 00:00:00) are used for the calculation of the correlation keys. Alternatively, also timestamps with a high resolution (e.g. resolution of one or ten or hundred milliseconds) could be used, but the timestamps may be rounded for the calculation of the correlation keys. It should be noted that the choice of the granularity or rounding interval may also impact the performance of a database or databases storing the media stream reports and signaling session reports for correlation. The larger the granularity, respectively the rounding interval and thus the search window the more data have to be processed to match signaling sessions and media streams.

Rounding timestamps could be for example implemented by rounding the timestamp to the nearest n second interval. Assuming for example that the interval is 5 seconds the following UNIX timestamps would be rounded as exemplified in Table 1.

Unix timestamp value	rounded value (5 sec)	description
1283936131	1283936130	rounded down
1283936133	1283936135	rounded up
1283936135	1283936135	unchanged
1283936136	1283936135	rounded down
1283936139	1283936140	rounded up

Table 1

The rounding interval can be any time interval in seconds or even milliseconds whereas the specific value is a trade-off between speed and accuracy. As noted above, the larger the interval the more data may have to be processed to match signaling sessions and media streams. The smaller the search window is the more likely it is to miss a match due to misaligned timestamps or derivations due to the delays outlined above.

The best rounding interval depends on the network monitored. Reasons for increased time differences are possible derivations from a shared time source (NTP) or differences seen in signaling and media events, i.e. the delay between a call being terminated and the endpoints stopping to send media.

For a smaller search window of three seconds the rounded values are different, as for example shown in Table 2.

Unix timestamp value	rounded value (3 sec)	description
1283936131	1283936130	rounded down
1283936133	1283936133	unchanged
1283936135	1283936136	rounded up
1283936136	1283936136	unchanged
1283936139	1283936139	unchanged

Table 2

Using normalized "session end" timestamp (i.e. timestamps of low granularity or timestamps that have been rounded) may be advantageous to counter the effects (e.g. false positive hash key matches) of the re-used port numbers. Since the end of a service (and/or end of media streams) becomes part of the correlation criteria only the port numbers that are used and valid at this point in time are considered in the correlation.

Another issue that could complicate the correlation processing is that deviations in detecting the end of a media stream, respectively a signaling session at the media plane probe 403 and the signaling plane probe 401 leads to the two entities rounding the timestamps to different values. For example, the signaling plane probe 401 could round the timestamp up, while the media plane probe 403 rounds the timestamp down to the next rounding interval value due to (minor) deviations in the timestamps at the signaling plane probe 401 and media plane probe 403. In another exemplary embodiment, additional correlation keys may be therefore generated by either one or both, the signaling plane probe 401 and media plane probe 403.

The signaling plane probe 401 and/or media plane probe 403 may use timestamps indicative of the rounding intervals before or after the timestamp indicative of the point in time where a

particular signaling session or media steam ended. For example a session ended at a point in time identified by the Unix timestamp value 1283936135. The correlation unit applies the rounding mechanism based on 3 second intervals. The rounded value thus is 1283936136. The correlation unit may generate additional keys with the timeframes before or after the mentioned rounded value. In this case the timestamps used for generation of the additional keys are 1283936133 (1283936136 – 3 seconds) and 1283936139 (1283936136 + 3 seconds). Of course, especially in cases where the rounding interval is short, it may also be desirable to more rounded timestamps rounded to the next and subsequent rounding interval(s) based on which correlation keys are calculated.

As noted above, the inclusion of multiple correlation keys calculated based on different timestamps for a given direction of a signaling session and/or media stream may for example be useful to counter correlation failures due to applying different rounding of timestamps in the signaling plane probe 401 and media plane probe 403 as will be outlined in further detail based on Fig. 13 and Fig. 14 in the following. The end of the signaling session $t_{signaling}$ of a service as observed by the signaling plane probe 401 and the end of the one or more media streams t_{stream} of the service as observed in the media plane probe 403 may not perfectly match each other as shown in Fig. 13. Even in case signaling plane probe 401 and media plane probe 403 calculate correlation keys based on the same IP address and UDP port tuple, the use of different timestamps $t_{signaling}$ and t_{stream} for the calculation of the correlation key will result in the correlation keys generated by signaling plane probe 401 and media plane probe 403 not matching each other. Additionally the use of rounded timestamps may not necessarily resolve this problem, as it cannot always be assured that the timestamps $t_{signaling}$ and t_{stream} are rounded to the same value. As also highlighted in Fig. 13 it could well happen that signaling plane probe 401 rounds the timestamp $t_{signaling}$ to the next lower rounded value T_1 in accordance with the rounding interval R, while the media plane probe 403 rounds the timestamp t_{stream} to the next higher rounded value T_2 in accordance with the rounding interval R. Again, correlation keys calculated on the different rounded timestamps T_1 and T_2 would not match each other.

As noted above, one possible solution to this problem is to include additional correlation keys at least to one of the signaling session reports and media stream reports. For example, in one exemplary embodiment, the media plane probe 403 could include two correlation keys per direction to the signaling session report. For a given media data path of a media stream (e.g. IP address and UDP port tuple) observed by the signaling plane probe, the signaling

plane probe could calculate one correlation key $C_{stream,1}$ on the media data path information and a rounded timestamp T_1, i.e. using the timestamp indicative of the end of the signaling session $t_{signaling}$ rounded to the next lower value according to the rounding interval, and the signaling plane probe could further calculate one correlation key $C_{stream,2}$ on the media data path information and a rounded timestamp T_2, i.e. using the timestamp indicative of the end of the signaling session $t_{signaling}$ rounded to the next higher value according to the rounding interval. Both correlation keys could then be included for the direction/media data path to the signaling session report.

A further implementation may include even further correlation keys, e.g. a third correlation key that is based on the rounded timestamp T_0, indicative of the next lower timestamp to T_1 according to the rounding interval.

The media plane probe 403 may calculate only one correlation key C_{stream} based on a rounded timestamp T_1 or T_2, i.e. based on the timestamp indicative of the end of the media stream (media data path) t_{stream} rounded to the next lower or higher value according to the rounding interval (and rounding rules). This correlation key for the media stream would then be included to the media stream report.

The correlation of the signaling session to the media stream(s) of the service would be successful in this embodiment, since either the correlation key $C_{stream,1}$ or the correlation key $C_{stream,2}$ within the signaling session report for the given direction/media stream would correspond to the correlation key C_{stream} included in the media stream's media stream report.

Optionally, in another embodiment of the invention, also the media plane probe 403 may calculate two or more correlation keys for each media stream, one correlation key based on the rounded timestamp T_1, i.e. based on the timestamp indicative of the end of the media stream (media data path) t_{stream} rounded to the next lower value according to the rounding interval, and the other correlation key based on the rounded timestamp T_2, i.e. based on the timestamp indicative of the end of the media stream (media data path) t_{stream} rounded to the next higher value according to the rounding interval. Both timestamps would then be included in the media stream report for the media stream by the media plane probe. In this exemplary implementation, the correlation may check whether the two correlation keys of the service's signaling session report for the direction/media stream match the two correlation keys within

the media stream report of the service's media stream to confirm the signaling session report and media stream report belonging to the same service.

Please note that instead of or in addition to the media plane probe 403 including two or more correlation keys to the media stream reports, also the signaling plane probe 401 could calculate two or more correlation keys per direction/media data path that is to be indicated in the signaling session report. Including multiple correlation keys per direction/media data path in the signaling session reports may be advantageous as it reduces the overhead information, since there are commonly less signaling session reports than media stream reports generated in the passive monitoring system.

Furthermore, in another exemplary implementation, the database(s) for signaling session reports and media stream reports may be further partitioned to reduce the search space for correlation. For example, the (respective) database could be partitioned in chunks of multiple hours, e.g. 4 or 6 hours or the like. The signaling session reports and media stream reports may further include as a separate element the timestamps of the end of the signaling session, respectively, the end of the media streams, based on which the database(s) can be partitioned and based on which the correct partition could be chosen for correlation later. Hence, in a more advanced implementation, besides the correlation keys a service end timestamp can be used as a search criteria to limit the amount of data that are to be searched in a database significantly.

In another exemplary implementation, the signaling plane probe may also generate a respective correlation key for all media streams of a service, in case the end of media streams is detectable within the session signaling, as noted previously. In this case the signaling session reports could further include a timestamp that indicates the end of each media stream, in order to allow choosing the correct partition of media stream reports in the database for correlation.

Correlation Mechanisms

For correlation of the media streams and the signaling sessions, different possibilities exist how the correlation by the correlation unit can be implemented. For example, according to exemplary embodiments of the invention, the correlation unit may correlate the media streams and the signaling session of the respective services periodically (or proactive) or in response to a query (i.e. on demand).

For example, a proactive correlation process could be utilized by a third party entity which has access to both signaling session reports and media stream reports. Their data may be

present in one or more databases to ease access to the raw data. As noted above, the correlation aims at correlating signaling and media information of the services and at merging them into combined service reports that may be then accessed by a third party, e.g. for evaluation of the services.

The proactive correlation process periodically takes all new signaling session reports which outline – but not fully describe – the media services, e.g. the monitored VoIP sessions. For each signaling session report the correlation process draws the calculated correlation keys from the signaling session report. There are at least two correlation keys in each signaling session report, at least one per direction. Additional correlation keys may optionally exist in the signaling session report, for instance if it describes a video call or a session with multiple endpoints. For each correlation key contained in a respective signaling session report the correlation process searches the media stream report database for the media streams with the same correlation key. If the correlation keys in the signaling session report are generated based on a timestamp indicating the end of the signaling session, every media stream of the service that existed at the end of the service in the media plane should have the same correlation key and will be in the result set of the search. If the correlation keys in the signaling session report also contain correlation keys for media streams that have been terminated during the ongoing service, the search may also further detect those media streams in the media stream report database based on the correlation key.

If a search returned two or more media stream reports, the correlation process encountered one of two scenarios. Either both media streams belong to the signaling session (i.e. service) or one of them was returned due to a correlation key collision. Both cases may be recovered by comparing the contents of the media stream report against the signaling session report. If a media stream in the result set starts before or ends after a signaling session, it was added to the result set due to a correlation key collision. Due to time information it cannot match the signaling session and is deleted from the result set by the correlation unit. If the media stream result set still contains multiple entries after the previous correlation key collision check, all remaining media streams (i.e. their media stream reports) belong to the media service (e.g. the VoIP session).

If a search returned no usable entry, i.e. no signaling session reports and media stream reports comprising matching correlation keys have been found, the correlation unit may optionally use additional correlation keys contained in either the signaling session reports and/or media stream reports to run another search, as explained with respect to Fig. 13 and Fig. 14 above.

The correlation process may optionally select only the most significant media stream for each direction for inclusion to the service report on the media service. The parameters to define the significance of the media streams may be for example configurable by the user/administrator. One parameter for deciding the significance of a media stream in the result set could be the duration of a media stream. For example, the service report generated for a given media service by the correlation unit could include the media stream with the longest duration, in case multiple media stream reports match a correlation key of a signaling session report. Another option would be to include in the service report the last media stream that has been alive in the media service, in case multiple media stream reports match a correlation key of a signaling session report. The correlation process may optionally add information from multiple media streams to the service report, e.g. when the service consisted of an audio and video part.

For a given signaling session (i.e. service), the correlation process may run at least twice, one time for each correlation key in a signaling session report. For example, in case the media service is a VoIP service, the correlation process may be for example performed by the correlation unit once for the caller's media stream and once for the callee's media stream, both identified by their respective correlation keys in the signaling session report.

The correlation unit combines the information contained in the media stream reports that are correlated to a respective signaling session report and the information contained in the signaling session report to provide a service report, providing an outline of the service. The content of this combined service report may be for example the sum or a subset of information of the signaling and media stream reports of a given service. Furthermore, the content of the service reports for the respective services may be configurable and/or depending on the application.

In another exemplary embodiment of the invention, the correlation unit may perform the correlation based on the media stream reports as starting points which are matched to their corresponding signaling session report. Ideally, each media stream of a media stream report should match to one signaling session report based on its correlation key. This implementation may optimize the search time in the correlation processing. The above correlation mechanism starting from the signaling session reports searches the rather large media stream report database at least twice for each signaling session. If the correlation process is based on the correlation keys within the media stream reports entries as starting point, the correlation would periodically take all new media stream reports, obtain the correlation keys thereof and search the signaling session reports for matching signaling sessions. At this point the same steps are taken as outlined above.

Another option may be to correlate media stream reports and signaling session reports on demand. An on-demand correlation employs the same mechanism as the proactive correlation mechanism described above, but instead of running periodically, e.g. when introducing new signaling session reports or media stream reports into the database(s), the on-demand correlation process runs when a user requests particular data.

In an exemplary implementation, the database may contain independent sets of signaling session reports and media stream reports. An operator could request service reports of services linked to some given search criteria, e.g. a specific user, a specific source and/destination IP address, a certain service type, services using a particular software application, codec, etc. The user could be for example specified by providing a username, phone number, the user's SIP URI, etc. or another identifier of the user that can be found in a signaling session report. The correlation unit would then search for the signaling session reports for all signaling sessions matching the particular search criteria. The search criteria may be for example provided to the correlation unit through a web application running on a browser.

For each signaling session report matching the search criteria, the correlation unit subsequently searches for the matching media stream reports as described above for the proactive correlation process. The input for this on-demand correlation is a set of signaling session reports matching the search criteria and not all signaling session reports as in the proactive correlation procedure described above. The on-demand correlation process generates for each signaling session report matching the search criteria a service report, for example by enriching the information of the signaling session report with the matching media stream report information and returns the service report for each matching signaling session report. The returned data may be for example viewed by the web application as individual service reports. Optionally the on-demand correlation process may also store the generated service reports in a database. This would save the on-demand correlation processing time in case an operator requests the same information multiple times.

Correlation Unit

Next, the correlation unit 402 as shown in Fig. 4 is discussed in further detail with respect to Fig. 5 and Fig. 6. Fig. 6 shows the structure of a correlation unit 402 according to an exemplary embodiment of the invention.

The correlation unit 402 may be reachable via the packet-switched network, e.g. through a web based application. For this purpose, the correlation unit 402 may comprise one or more network interface cards 601. The correlation unit 402 has a receiving unit 602 and a

transmitter unit 603, for example for receiving correlation related messages, such as for example on-demand correlation requests, and for responding thereto. Furthermore, the receiving unit 602 and a transmitter unit 603 may transmit and receive messages for accessing remote database(s) that contain(s) the signaling session reports and the media stream reports. The receiving unit 602 and transmitter unit 603 may receive and transmit the correlation related messages from its network interface card(s) 601.

The correlation unit 402 may also have a processing unit 604. The processing unit 604 may be used to determine correlation respective signaling sessions to their corresponding media streams. Furthermore, the processing unit 604 may be used to generate services reports. The service reports may for example comprise information on the respective services and the quality information of the media streams of the respective services. This information may be obtained from the signaling session report and media stream reports belonging to a given service. For correlation, it may be necessary to query one or more databases that comprise the signaling session reports and the media stream reports as provided by one or more signaling plane probes 401 and media plane probes 403. The correlation unit 402 may therefore comprise a database access unit 605 which provides functionality to access these one or more databases. In one exemplary implementation, the database(s) may be located remote from the correlation unit 403.

The correlation unit 402 may also have a storage unit 606, e.g. a RAM, for (temporarily) storing information and data related to the correlation processing. For example, the result set of media stream reports for correlation process, the data of the signaling session reports, and the service report generated for a respective service may be stored in the storage unit 606.

Fig. 5 shows a flow chart of the operation of the correlation unit 402 according to an exemplary embodiment of the invention. Exemplarily, the correlation processing for determining the media streams belonging to one given signaling session report is illustrated in Fig. 5. In case a correlation should be performed for more than one signaling session (report), the procedure of Fig. 5 is repeated for each signaling session report. For example, when using a proactive correlation mechanism, the correlation unit 402 could obtain all signaling session reports that have been generated since the previous correlation process from a signaling session database and process each of the signaling session reports according to Fig. 5. In another embodiment related to an on-demand correlation process, the correlation unit 402 could also receive a correlation request that is indicating one or more search criteria, and may first – not shown in Fig. 5 – query a signaling session database for all signaling session reports that match the search criteria. The signaling session reports matching the search one or more criteria may then be each processed according to Fig. 5 and the service report(s) are returned to the querying party.

The correlation unit 402 obtains 501 a signaling session report to be processed and generates 502 a service report for the service corresponding to the given signaling session. This service report may include information contained in the signaling session report and will be further supplemented with additional information as described below. Furthermore, the correlation unit 402 obtains 503 a correlation key contained in the signaling session report – as noted previously there should be at least two correlation keys contained in the signaling session report. Using the obtained correlation key, the correlation unit 402 queries 504 a media report database to receive all media stream reports that match the correlation key i.e. that contain the same correlation key.

As noted above, it may happen that the query results in the database returning more than one media stream report for the given correlation key, for example due to a hash collision of the correlation key. Therefore, the correlation unit 402 checks 505 whether the returned result set of the database query comprises more than one media stream report. If not (no), the correlation unit 402 can immediately enrich 508 the service report with information on the media stream as contained in the one media stream report returned from the media stream database.

In case the result set returned by the database comprises multiple media stream reports, the correlation unit 402 may try to detect 506 and resolve a hash collision. For example, the hash collision could be resolved by checking further information in the media stream reports of the result set for parameters that identify the media stream (report) that belongs to the correlation key that has been queried. For example, in case the media stream reports contain timestamps of the beginning and/or end of the media stream, same could be compared to respective timestamps contained in the signaling session report, which yield the beginning and/or end of the signaling session. The media stream report(s) of media stream(s) that start before or end after the signaling session may be obtained from the database due to a hash collision. Accordingly, the correlation unit 402 may select 507 only the most relevant media stream report, i.e. the media stream report that is most likely corresponding to the service of the signaling session report from the result set and may add 508 at least a part of its information to the service report. Further, the correlation unit checks 509 whether there is a further correlation key contained in the signaling session report and, if so (yes), repeats steps 503 to 508. These steps are repeated for each further correlation key(s) comprised in the signaling session report.

Once all correlation keys contained in the signaling session report have been processed, the service report for the signaling session is complete, and may be optionally stored 510 by the correlation unit 402 in a service report database.

Depending on the correlation type or application, the procedure of Fig. 5 could be supplemented with additional steps. For example, in case of a proactive/periodic correlation, the correlation unit 402 may check whether there is another signaling session report to be matched to its corresponding media stream reports, and if so the procedure of Fig. 5 would be executed for each signaling session report to be processed. In case of performing a correlation on demand, the correlation unit 402 may likewise check whether there is another signaling session report to be matched to its corresponding media stream reports, and if so the procedure of Fig. 5 would be executed for each signaling session report to be processed. Once all signaling session reports have been processed the correlation unit 402 may generate a correlation response message comprising the service reports generated by the correlation unit 402 in response to the search criteria provided in a correlation request, and returns the correlation response message to the querying entity, e.g. another device/computer.

Signaling Plane Probe

Fig. 9 shows an exemplary implementation of a signaling plane probe 401 and Figs. 7 and 8 show two exemplary flow charts of the operation of same according to different embodiments of the invention. In general, the signaling plane probe 401 observes one or more network links to obtain a copy of all packets passing these network links/interfaces. For this purpose, the signaling plane probe 401 comprises a monitoring unit 901 that receives 701 packets of a service on the packet-switched network 400. For example, this monitoring unit 901 may receive the packets from a one or more specialized packet capture cards 906 (network interface card - NIC) which is capable of receiving and processing packets on the physical interface and to provide them to the application layer, nearly without requiring CPU processing time and operation system functionality. An optional filter in each packet capture card of the signaling plane probe 401 may be used to prevent unwanted packets from being copied from the network interfaces to the processing processes. These filters may be for example implemented as stateless or stateful Internet Protocol (IP) or payload filters. For example, this filter may check 702, whether the currently processed packet is a signaling packet or not. If the packet does not belong to a signaling session, i.e. is not packet of a signaling related protocol, the next packet may be processed.

Each signaling packet received at the signaling plane probe 401 will be processed by a detection unit 902 which analyzes the contents of the packets provided by the monitoring unit 901 so as to identify and process packets as signaling packets. Specifically, the detection unit 902 detects signaling packets and matches these to signaling sessions of the media service, and is capable of identifying one or more media data paths of a media plane of the

service based on signaling information of the service comprised in the received packets. Optionally, the detection unit 902 may be provided in one or more specialized packet capture cards.

For example, the detection unit 902 identifies packets containing signaling information (signaling packets) in a parsing process based on their header information. If parsed successfully as a signaling packet, the signaling packet is forwarded to a state engine. The state engine is designed to identify 703 and form signaling sessions from individual signaling packets.

The state engine (which may be considered an internal database with correlation related information maintained in the storage unit 903) of the detection unit 902 has a processing facility and a state table which holds information about signaling sessions (signaling session records). Based on the protocol-specific identifiers in the signaling packets the detection unit 902 can identify the start and end of a signaling session, respectively media service, or "append" a respective processed signaling packet to an existing signaling session. Signaling packets may also contain information on the start and termination of media streams in the media plane of the media service, which are events that may be tracked by the detection unit 902. Depending on whether the identified signaling session is already registered in the signaling plane probe 401 (step 704), the detection unit 902 registers 705 the new signaling session in the state table, including the generation of a new signaling session record for the signaling session.

The signaling session record may be supplemented with information on the signaling session while same is running to account – for example – for endpoint address/name (e.g. SIP URI); the phone number of the user; the user-agent (endpoint software and version); capabilities/features; supported media types; IP addresses and ports for signaling and media streams (from SDP body); signaling routing information; timestamps: e.g. session start, Post-Dial-Delay (PDD), connect time, session end; information on the session flow (e.g. signaling message(s) with timestamps); session state (e.g. successful, unsuccessful, timed out, reason for failure), and/or events during session (media on hold, codec renegotiation, session forking, authentication), etc.

Furthermore, by analyzing 706 the contents of the signaling packets, e.g. a SDP body of a SIP messages comprised in signaling packets, the detection unit 902 may further learn one or more media data paths of the media plane of the service belonging to the media service that is associated to a given signaling session. The signaling plane probe 401 may further comprise a storage unit 903, e.g. a volatile memory, to store the information describing the state of the respective signaling sessions in the state table. In case the signaling packet

indicates the start of a new media stream (step 706: yes), the detection unit 902 may add 707 information on the media data path (e.g. IP address and port number of the destination, and optionally the source of the media stream) of this media stream to the signaling session record. Likewise, in case the signaling packet indicates the end/termination of a new media stream (step 708: yes), the detection unit 902 may remove the information on the respective media stream from the signaling session record maintained in storage unit 903. In case the signaling packet indicates the termination of the signaling session (step 710: yes), the detection unit 902 may either itself or trigger the processing unit 904 to generate 711 a respective correlation key for each media data path indicated in the signaling session record. The correlation key for a respective media data path of a media stream may be generated for example based on information on the directivity of the media stream (e.g. as given by the tuple of destination IP address and UDP port number) and a timestamp indicating the point in time at which the signaling session was terminated. For example, this timestamp may indicate the point in time of reception of the signaling packet indicating the event of the termination of the signaling session at the signaling plane probe 401. For example, the network interface card(s) 906 of the signaling plane probe could generate a timestamp indicating the time of reception of the respective signaling plane packet and could make same available to the other components of the signaling plane probe 401. Furthermore, the detection unit 902 or the processing unit 904 may further generate 912 a signaling session report based on the information on the signaling session gathered in the signaling session record. Further, the signaling session report also included the correlation keys for the media data paths of the service that exist at the end of the service according to the signaling plane information. The signaling session report may then be stored in a signaling session report database, for example using the database access unit 907 of the signaling plane probe 401. Moreover, the signaling session is deregistered 713 from the state table.

Once a given signaling packet has been processed, it may be discarded. It should be also noted that the order of the checks in steps 706, 708 and 710 and the respective actions to be taken is only exemplary. It is only important that the information on the media streams of the service that exist in the media plane are tracked by the signaling plane probe 401, and that the signaling plane probe 401 is able to detect the termination of the signaling session.

Optionally (not shown in Fig. 7), in another exemplary embodiment, the detection unit 902 of the signaling plane probe 401 periodically (e.g. every 5, 10 or 15 seconds) traverses all entries in the state table in order to find and clean up stale signaling sessions. For this purpose, the signaling plane probe 401 may add a timestamp to each signaling session record in the state table that is updated each time a new signaling packet/message of the respective signaling session is received. Based on this timestamp allows to identify stale

signaling session by simply checking based on the signaling sessions' timestamps whether the last signaling packet/message for a given signaling session has been received longer than a threshold time period ago. For example, stale sessions might be those which did not receive new signaling packets/messages for 86400 seconds or signaling sessions may be defined stale and are terminated after a certain time. This clean-up removes unnecessary information from the state.

Fig. 8 shows another flow chart of the operation of a signaling plane probe 401 according to another embodiment of the invention. In the example discussed with respect to Fig. 7, the signaling plane probe 401 keeps track of the media streams currently existing in the media plane of the service, eliminating those media streams from the signaling session record that are terminated during the ongoing signaling session/service. In the flow chart of Fig. 8, which is corresponding to Fig. 7 to a large extent, the termination of one or more media streams in the media plane, as detected in step 708, causes the detection unit 902 or the processing unit 904 (triggered by the detection unit 902) to generate 801 a correlation key for each terminated media stream. Similar to the other correlation keys, the correlation key for a terminated media stream may be generated based on its media data path information (e.g. destination IP address and UDP port number) and a timestamp that indicates the point in time the media stream was terminated. This timestamp may indicate the point in time of reception of the signaling packet indicating the event of the termination of the media stream at the signaling plane probe 401. The correlation key generated is then added 802 to the signaling session record of the signaling session by the detection unit 902 or the processing unit 904 and is later then also stored in the signaling session report in step 712 together with the other correlation keys either generated in step 801 or step 711.

Media Plane Probe

Next, the media plane probe 403 according to an exemplary embodiment of the invention is described with respect to Fig. 10 and Fig. 11. Fig 10 shows an exemplary embodiment of a media plane probe 403, while Fig. 11 shows a flow chart of its exemplary operation according to an embodiment of the invention.

Similar to the signaling plane probe 401, the media plane probe 403 observes one or more network links to obtain a copy of all packets passing these interfaces. The media plane probe 403 may observe one or more network links to obtain a copy of all packets passing these interfaces. The media plane probe 403 comprises a monitoring unit 1001 that receives 1101 packets of a service on the packet-switched network 401. For example, this monitoring unit 1001 received the packets from one or more specialized packet capture cards 1006 (network interface card - NIC) which are capable of receiving and processing packets on the physical

interface and to provide them to the application layer, nearly without requiring CPU processing time and operation system functionality. An optional filter in the packet capture card of the media plane probe 403 may be used to prevent unwanted packets from being copied from the network interfaces to the processing processes. This filter may be for example implemented as stateless or stateful Internet Protocol (IP) or payload filter. For example, this filter may check 1102, whether the currently processed packet is a media plane packet (media packet) or is some other packet. If the packet does not belong to the media plane, i.e. is not packet of a media related protocol, such as for example RTP, the next packet may be processed.

Each media packet received will be processed by the media plane probe 403. The detection unit 1002 of the media plane probe 403 is designed to detect and form media streams from single media packets. The media plane probe 403 may further comprise a processing unit 1002 that may hold a state table which stores information about the media streams monitored by the media plane probe 403 in a local storage unit 1003, e.g. a volatile memory. Based on the protocol-specific identifiers in the media streams the media plane probe 403 (e.g. based on header information in the RTP packets and/or lower layer headers) the detection unit 1002 identifies 1103 the media data path, i.e. the media stream, the media packet belongs to. If the media data path has not been yet registered (step 1104: no), i.e. the processed media packet starts a new media stream (with the given media data path) the detection unit 1002 generates a new media data record in the state table to account for the new media data path observed by the media plane probe 403. This media stream record may contain information on the media stream such as for example the destination (and optionally the source) IP address and port combination of the media stream, the SSRC in the packet header (in case RTP is used for media transport), etc. If the media data path is already known (step 1104: yes), the media stream record already exists. In any case, the detection unit 1002 updates a timestamp 1106 in the media stream record to indicate the point in time at which the last media packet of the media stream has been received. This timestamp may be used for calculating the correlation key for the media stream. Furthermore, the detection unit 1002 may (re)start a timer 1108 for the media stream record, which when timing out is indicating the termination of the media stream as will be outlined below. Alternatively, no timer is used, but the media plane probe 403 periodically checks for media streams the last media packet of which was received more than a threshold time period ago, in order to determine terminated media streams.

The media stream records of the media streams may for example indicate the respective media data paths thereof by an IP address and port tuple of the destination device, i.e. the receiving device of the media stream. The media plane probe's detection unit 1002 stores

information describing the state of the detected media streams in the media stream records maintained in the state table. Optionally, the individual media streams (i.e. their media data paths) may be identified by a media path identifier (e.g. stream hash key) as unique key. Further optionally, upon detecting a new media stream, respectively media path (step 1104: no), the processing unit 1004 could generate 1105 a media path identifier (e.g. a media path identifier (e.g. stream hash key)) for the media stream, respectively media path and stores 1106 the media stream record in the state table as outlined above.

Optionally, the processing unit 1004 could analyze every received packet for quality information 1107 and – further optionally – store them in the storage unit 1003 for a given time period. Periodically the processing unit 1004 may aggregate or compress the quality information to form fine-grained media stream quality reports with media stream related information. Fine-grained media stream quality reports could be for example written in predetermined time intervals, e.g. every five seconds, and could provide a summary of the media stream in the respective interval and the aggregated quality information of the media stream in the respective interval. Upon generation of a respective fine-grained media stream quality report, the packets used for the generation of this fine-grained media stream quality report can be removed from the storage unit 1003.

The processing unit 1004 of the media plane probe may for example periodically check 1109 (e.g. every 5, 10 or 15 seconds); whether a timer set of a media stream (record) in step 1108 has expired. Alternatively, as noted above, the media plane probe 403 may also check timestamps indicating the last update (i.e. last reception of a media packet) of the respective media stream reports in the state table. A media stream may be for example considered terminated, if no new media packets have been received by the monitoring unit 1001 for the existing media stream for a given amount of time (e.g. 60 seconds). Accordingly, if using a timer to monitor the termination, the timer value could be set to this given amount of time. Further alternatively, instead of periodically checking the media stream records or expiry of their timers, the expiry of a timer for a media stream (record) could cause an interrupt invoking the processing unit 1004 to perform steps 1110 to 1112 as outlined below.

Upon detecting the termination of a media stream, the processing unit 1004 processes the respective information in the media stream record of this media stream the state table. The processing unit 1004 generates 1110 a correlation key for the media stream based on its media data path information (e.g. the destination IP address and UDP port number) and the timestamp in the media stream record that is indicating the point in time of receiving the last media packet of the media stream (see step 1106). Furthermore, the processing unit 1004 generates 1111 a media stream report for the media data path, respectively media stream, which includes the generated correlation key and further information on the media stream.

This further information may be for example transport media information and optionally stream quality information. For example, in case RTP is used for media transport, the transport media information could include the source and destination IP address of the RTP stream, the source and destination UDP port of the RTP stream, the RTP header Synchronization source (SSRC) value and the RTP payload type. Furthermore, the processing unit 1004 may also generate information on the media stream quality that is included to the media stream report as stream quality information. This stream quality information could for example include timestamp information (e.g. start and/or end time of the media stream), an indication of RTP format violations (e.g. the number thereof), information on the protocol conformance (e.g. sequence errors and/or DSCP errors), and transport information (e.g. packet interval and jitter statistics, information on packet loss, information on duplicate packets, etc.), MOS values (e.g. per time-slice (e.g. every 5 seconds), minimum/average/maximum value for entire stream duration, etc.), and/or events (e.g. comfort noise, silence suppression, DTMF tones, codec changes, marker bits). The contents of the media stream reports may be configurable by an operator and may be adapted to customers' needs.

The media stream record may be further stored in a media stream report database using the database access unit 1007 of the media plane probe 403. Finally, the processing unit 1004 of the media plane probe 403 deregisters 1112 the media data path from the state table, e.g. by deleting the corresponding media stream record.

Steps 1110 to 1112 are performed for every media stream that is considered terminated by the media plane probe.

Optimizations may be used to speed up processing of entries in the state tables inside the signaling plane probe 401, the correlation unit 402 and the media plane probe 403. A simple list would be sufficient to implement the features as outlined above. Add, modify, delete and search operations on simple lists may however be too slow. These operations may be implemented utilizing more complex but faster data structures, e.g. hash tables or data structures heavily dependent on CPU and Operating System specific behavior and functionality. The Intel Thread Building Block (TBB) library for instance can be used to implement in-memory databases used for the state tables in the signaling plane probe 401, the correlation unit 402 and the media plane probe 403. The Intel Thread Building Block (TBB) library may also be used to aid implementation of the processing queues inside and in between the components outlined above.

Please note that in case the signaling plane probe 401 and the media plane probe 403 should perform passive non-intrusive monitoring of the signaling plane traffic, respectively

media plane traffic, a separate network interface card is employed for the transmission of correlation related messages (i.e. the network interface card(s) for monitoring the traffic are only employed for packet reception but not for packet transmission).

Please also note that in the procedures described above the detection unit 902, 1002 and the processing unit 904, 1004 may use the database access unit 907, 1007 to access the database/state table in the storage unit 903, 1003.

Application

The correlation mechanism for signaling plane and media plane of a service may be used in manifold applications. For example, the correlation mechanism is useful in small networks but is more important in larger scale networks where the various signaling and media components are decomposed and present on different, sometimes geographically split, network entities. This is the case in IP Multimedia Subsystem (IMS) or Next Generation Network (NGN) setups.

Fig. 12 shows a simplified network in which locations are indicated where correlation may be desired and can be implemented using the correlation mechanism according to one of the various embodiments described herein. Possible monitoring locations are for example (and not limited thereto):

- The Access Session Border Controller (SBC) network interfaces with end customers (shown of the left side of Fig. 12),
- a decomposed gateways system, e.g. to the PSTN, is shown as Media Gateway (MGW) and a Media Gateway Controller (MGC) – see In the top right section of Fig. 12, and
- an Interconnection SBC interfaces to one or more interconnection or transit partners (see the lower right side of Fig. 12.

It should be noted that the location of the monitoring probes for signaling plane and media plane may also be subject to requirements on the monitoring solution.

In some cases, signaling plane and media plane may originate and/or terminate at the same endpoint (device). In case of an end customer this endpoint most likely uses the same IP address on the signaling and media plane. These streams may be for example observed, if the signaling plane probes 401 and the media plane probes 403 are placed close to the end customer devices. The Access Session Border Controller may be considered to be close to the carrier internal network's "border" towards end customer devices and to therefore receive

traffic from the end customers on one side and transmits this traffic on the other side towards the carrier internal network. Most commonly these Session Border Controllers (SBC) are used to separate the devices outside the own trust domain from the inside network. Although the traffic is in one geographic location the SBC may use different IP addresses to load balance traffic. In this case a simple IP address based correlation mechanism not using the principles and ideas outlined herein would already fail whereas the correlation mechanism described herein would succeed. The signaling plane probes 401 and media plane probes 403 may be for example placed in a way that they observe the signaling streams and media streams respectively.

In the upper right corner of the diagram shown in Fig. 12, a decomposed gateway is shown which uses different physical entities as media and signaling components (MGC and MGW, respectively). In between those devices a third party protocol may be used to synchronize the devices. For the sake of a minimum size network diagram only one of each component types is shown but in real life scenarios multiple media components are controlled by one signaling component (as it is the case in IMS, NGN). The media components may be co-located in the same location with the signaling component or spread across the internal network of the carrier. In either case the correlation task is rather complex as it involves multiple IP addresses on different physical network links and thus may require multiple signaling plane probes 401 as well as media plane probes 403.

In a third exemplary scenario the correlation has to correlate signaling and media streams from different locations. A media plane probe 403 component on the internal side of the interconnection point (lower right corner of the diagram – Interconnection SBC) observes a media stream from the Access SBC (middle left). The corresponding signaling session is observed next to the aforementioned Access SBC by a separate signaling plane probe 401.

The location of the correlation unit 402 does not impact what signaling sessions or media streams can be correlated. The correlation unit 402 should have access to the signaling session reports and media stream reports of the signaling plane probes 401 and media plane probes 403 and can be implemented as distinct physical machine or co-located on an existing signaling plane probe 401 or media plane probe 403.

The terms probe and correlation unit used herein can be considered to refer to dedicated hardware, a software program or a combination thereof to implement the desired functionality of the respective probe, respectively, correlation unit. Accordingly, another embodiment of the invention relates to the implementation of the above described various embodiments using hardware and software. It is recognized that probes and correlation unit of the invention may be implemented by having computing devices (processors) execute a software

program that causes the desired functionality being executed. A computing device or processor may for example be general purpose processors, digital signal processors (DSP), application specific integrated circuits (ASIC), field programmable gate arrays (FPGA) or other programmable logic devices, etc. Some of the computing devices may store the software program internally.

The various embodiments of the invention may also be performed or embodied by a combination of computing devices and software programs providing the desired functionality stored on any kind of computer readable storage media, for example RAM, EPROM, EEPROM, flash memory, registers, hard disks, CD-ROM, DVD, etc.

It should be further noted that the individual features of the different embodiments of the invention may individually or in arbitrary combination be subject matter to another invention.

It would be appreciated by a person skilled in the art that numerous variations and/or modifications may be made to the present invention as shown in the specific embodiments without departing from the spirit or scope of the invention as broadly described. The present embodiments are, therefore, to be considered in all respects to be illustrative and not restrictive.

CLAIMS

1. A method for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network, the method comprising:

a) generating at least one correlation key for each direction of a service based on the information of the signaling session of the respective service monitored at the signaling plane probe, wherein the correlation keys for a respective service are generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and using information on the direction of the respective service,

b) storing, by the signaling plane probe, a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service,

c) generating a correlation key for each media stream monitored by the media plane probe, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,

d) storing, by the media plane probe, a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media, and

e) correlating the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports.

2. The method according to claim 1, wherein the generation of the respective correlation keys for the respective media streams by the media plane probe is based on the media data path information of the respective media stream and the timestamp indicative of the point in time at which the respective media stream ended.

3. The method according to claim 1 or 2, wherein information on the direction of the respective service, based on which the correlation keys are generated by

signaling plane probe for a respective service, is information on one or more media data paths of media streams of the service learned from messages of the signaling session of the service.

4. The method according to one of claims 1 to 3, further comprising the steps of:

f) monitoring, by the signaling plane probe, signaling messages of the respective signaling sessions of the services, and

g) detecting, by the signaling plane probe, media streams belonging to each service associated to a respective monitored signaling session based on the signaling messages of the respective signaling sessions,

wherein the correlation keys for a respective one of the signaling sessions are generated based on information on the media data path of the media streams belonging to the service associated to the respective one of the signaling sessions and the timestamp indicative of the point in time at which the end of the respective signaling session has been detected by the signaling plane probe.

5. The method according to one of claims 1 to 4, wherein information on a media data path of a media stream is a destination IP address and port number of the media stream.

6. The method according to one of claims 1 to 5, wherein the timestamps used for the generation of the correlation keys by the signaling plane probe and the media plane probe are rounded before the generation of the correlation keys or have a granularity in the range of one to ten seconds.

7. The method according to one of claims 1 to 6, wherein the signaling plane probe generates for each direction of the service at least one correlation key using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and being rounded to the next lower timestamp value according to a rounding interval, and at least another correlation key using the timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and being rounded to the next higher timestamp value according to the rounding interval.

8. The method according to one of claims 1 to 7, wherein the media plane probe generates for each media stream at least one correlation key using a timestamp indicative of the point in time at which the respective media ended and being rounded to the next lower timestamp value according to a rounding interval, and at least another correlation key using the timestamp indicative of the point in time at which the respective media stream ended and being rounded to the next higher timestamp value according to the rounding interval.

9. The method according to one of claims 1 to 8, wherein the correlation of the media streams and the signaling session of the respective services is performed by a correlation unit periodically or in response to a query.

10. A method for correlating media streams and signaling sessions of services, the method comprising the steps of:

a) correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services, wherein correlating the media streams and the signaling sessions comprises identifying one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys contained in the media stream reports and signaling session reports, and

b) generating a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained the signaling session report of said at least one service.

11. The method according to claim 10, wherein correlation keys provided in a respective signaling session report have been generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the corresponding service has been detected and using information on the media streams of the service.

12. The method according to claim 10 or 11, wherein correlation key provided in a respective media stream report has been generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the media stream.

13. The method according to one of claims 10 to 12, wherein correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions comprises:

matching the respective correlation keys of the media stream reports of media streams are to the respective correlation keys of the signaling session reports, or vice versa,

wherein a match of the correlation keys of a one of the media stream reports and one of the signaling session reports yields the media stream of said one media stream report belonging to the same service to which the signaling session of said one signaling session report belongs.

14. The method according to one of claims 10 to 13, wherein the service report comprises information on the service and the quality information of the media streams of the service.

15. The method according to one of claims 10 to 13, wherein correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions comprises:

obtaining from a signaling session report correlation keys identifying the media streams belonging to the same service as the signaling session, and

identifying those media stream reports which include one of the correlation keys obtained from the signaling session report.

16. The method according to one of claims 10 to 13, wherein correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions comprises:

obtaining from a media stream report a correlation key identifying the media stream, and

identifying a signaling session report among the signaling session reports which includes the same correlation key obtained from the signaling session report.

17. The method according to one of claims 10 to 16, wherein steps a) and b) are performed by a correlation unit periodically for all signaling session reports and media stream reports generated within a configurable time interval, or in response to a request.

18. The method according to one of claims 10 to 17, further comprising steps a) to d) of claim 1 or the steps of the method according to one of claims 2 to 9.

19. A passive monitoring system for correlating media streams and signaling sessions of services in a packet-switched network, the passive monitoring system comprising:

a signaling plane probe for generating at least one correlation key for each direction of a service based on the information of the signaling session of the respective service monitored at the signaling plane probe, wherein the correlation keys for a respective service are generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and using information on the direction of the respective service,

wherein the signaling plane probe is further adapted to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service,

a media plane probe for generating a correlation key for each media stream monitored by the media plane probe, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,

wherein the media plane probe is adapted to store a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media, and

a correlation unit for correlating the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports.

20. A signaling plane probe comprising:

a network interface card for receiving packets of signaling sessions of multiple media services,

monitoring unit for monitoring the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of the signaling sessions, wherein the monitoring unit is further adapted to detect the point in time at which a respective signaling session of the respective service terminates, and

a processing unit for generating at least one correlation key for each direction of a service based on a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the monitoring unit and based on information on the direction of the respective service,

wherein the monitoring unit is adapted to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises at least two correlation keys generated for the respective directions of the service.

21. A media plane probe comprising:

a network interface card for receiving packets of media streams of multiple media services,

monitoring unit for monitoring the media streams based on the packets of the media streams received by the media plane probe to thereby obtain information on a media data path a respective one of the media streams, wherein the monitoring unit is further adapted to detect the point in time at which a respective media stream of the respective service terminates, and

a processing unit for generating a correlation key for each media stream monitored by the monitoring unit, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,

wherein the monitoring unit is adapted to store a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

22. A correlation unit for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network, the correlation unit comprising:

a database access unit for accessing a database containing signaling session reports for signaling sessions of services and a database comprising media stream reports of media streams of the services, wherein the signaling session reports each contain at least one correlation key for each direction of a respective service, and the media stream reports each contain one correlation key indication the direction of the media stream,

a processing unit for correlating the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services, wherein the processing unit is adapted to identify one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys contained in the media stream reports and signaling session reports, and

a service record generation unit for generating a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained the signaling session report of said at least one service.

23. The correlation unit according to claim 22, further comprising means to perform the steps of the method according to one of claims 10 to 17.

24. A computer-readable medium storing instructions that, when executed by a processor of a signaling plane probe, cause the signaling plane probe:

to receive packets of signaling sessions of multiple media services,

to monitor the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of the signaling sessions,

to detect the point in time at which a respective signaling session of the respective service terminates,

to generate at least one correlation key for each direction of a service based a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected and based on information on the direction of the respective service, and

to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service.

25. A computer-readable medium storing instructions that, when executed by a processor of a media plane probe, cause the media plane probe:

to receive packets of media streams of multiple media services,

to monitor the media streams based on the packets of the media streams received by the signaling plane probe to thereby obtain information on a media data path a respective one of the media streams,

to detect the point in time at which a respective media stream of the respective service terminates, and

to generate a correlation key for each media stream monitored by the monitoring unit, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,

to store a media stream report for each media stream in a database, wherein each media stream report contains at least the correlation key generated for the respective media.

26. A computer-readable medium storing instructions that, when executed by a correlation unit, cause the correlation unit to correlate media streams and signaling sessions of services in a passive monitoring system of a packet-switched network, by:

accessing a database containing signaling session reports for signaling sessions of services and a database containing media stream reports of media streams of the services, wherein the signaling session reports each contain at least one correlation key for each direction of a respective service, and the media stream reports each contain one correlation key indication the direction of the media stream,

correlating the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services,

identifying one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports, and

generating a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained the signaling session report of said at least one service.

27. The computer-readable medium according to claim 26, further storing instruction that, when executed by the processor of the correlation unit, cause the correlation unit to perform the steps of the method according to one of claims 10 to 17.

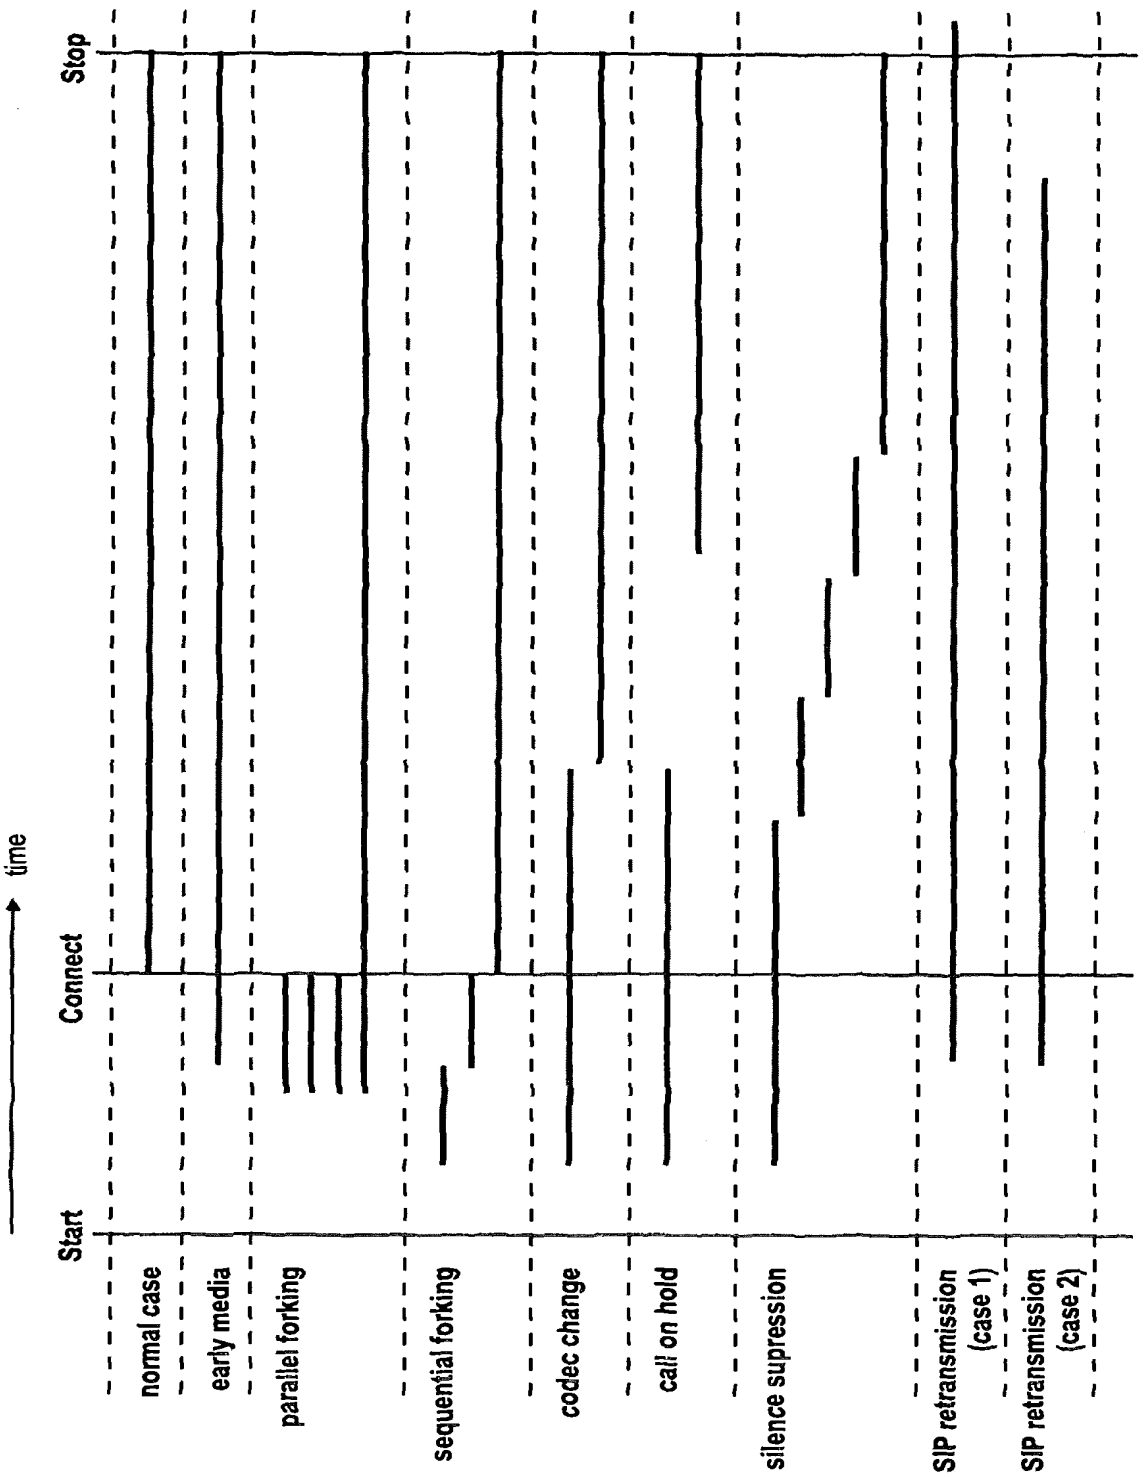
Start
Connect
Stop
time
normal case
early media
parallel forking
sequential forking
codec change
call on hold
silence supression
SIP retransmission (case 1)
SIP retransmission (case 2)

Fig. 1

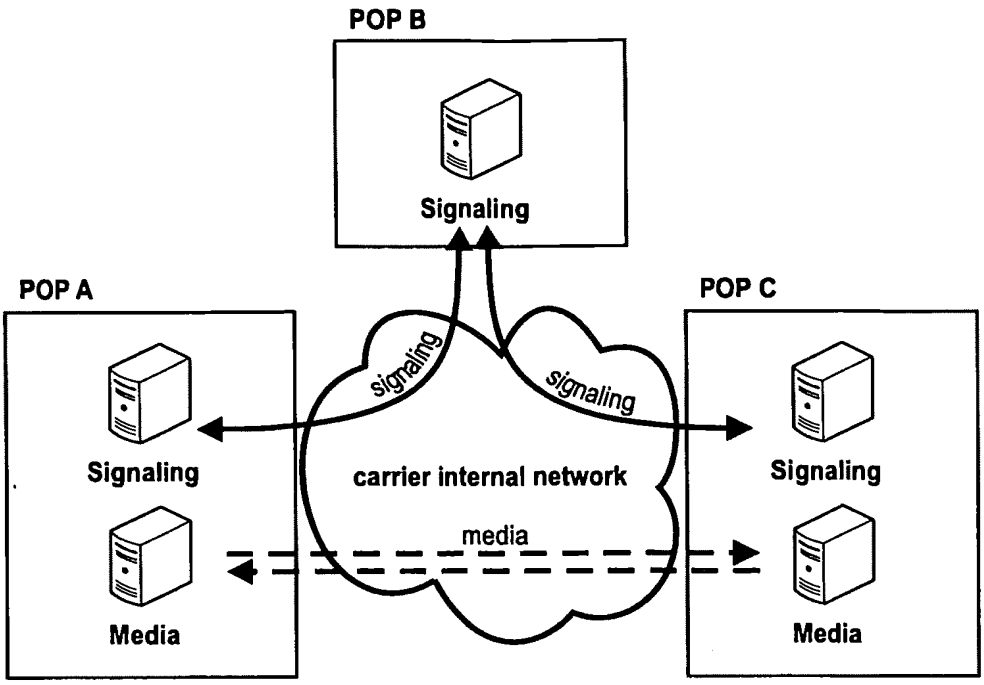
POP B
Signaling
POP A
POP C
signaling
signaling
Signaling
carrier internal network
Signaling
media
Media
Media

Fig. 2

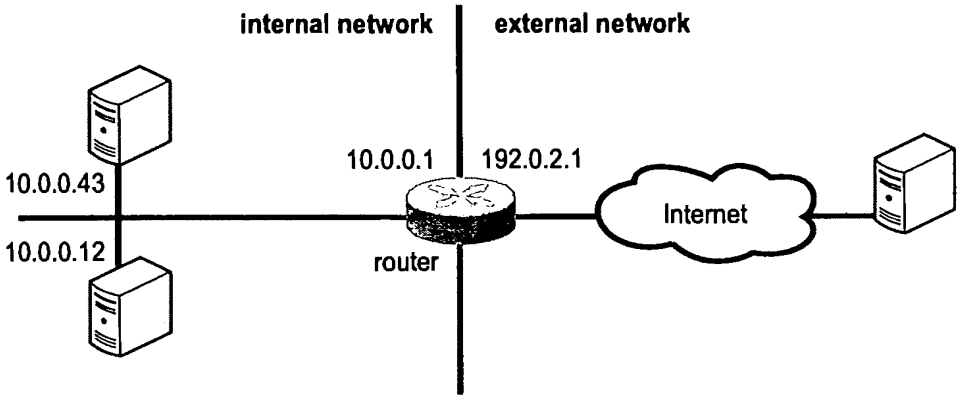
internal network
external network
10.0.0.1
192.0.2.1
10.0.0.43
Internet
10.0.0.12
router

Fig. 3

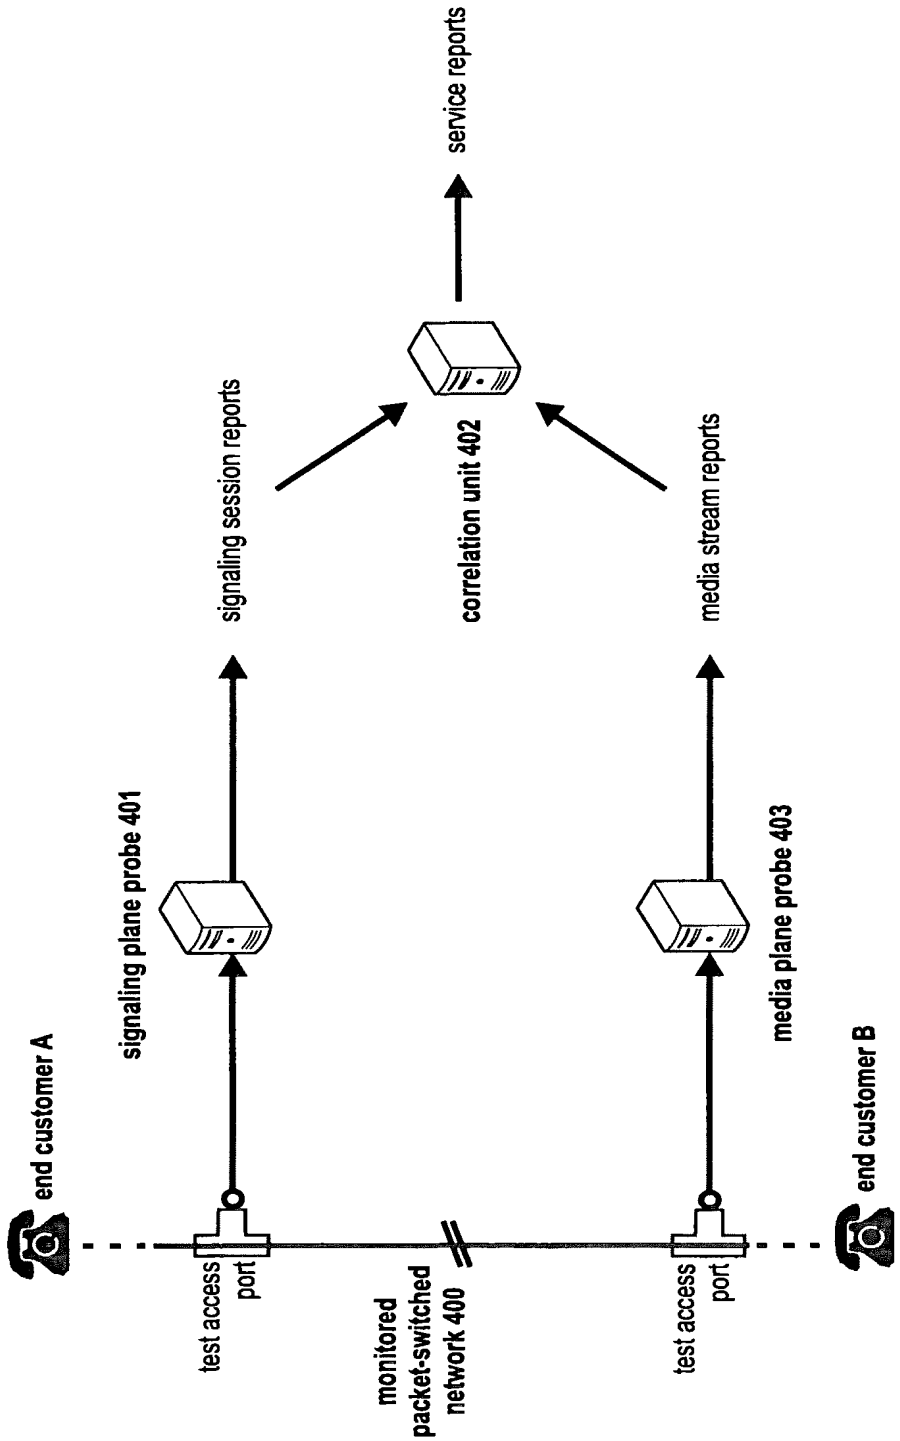
service reports
signaling session reports
correlation unit 402
media stream reports
signaling plane probe 401
media plane probe 403
end customer A
end customer B
test access port
test access port
monitored packet-switched network 400

Fig. 4

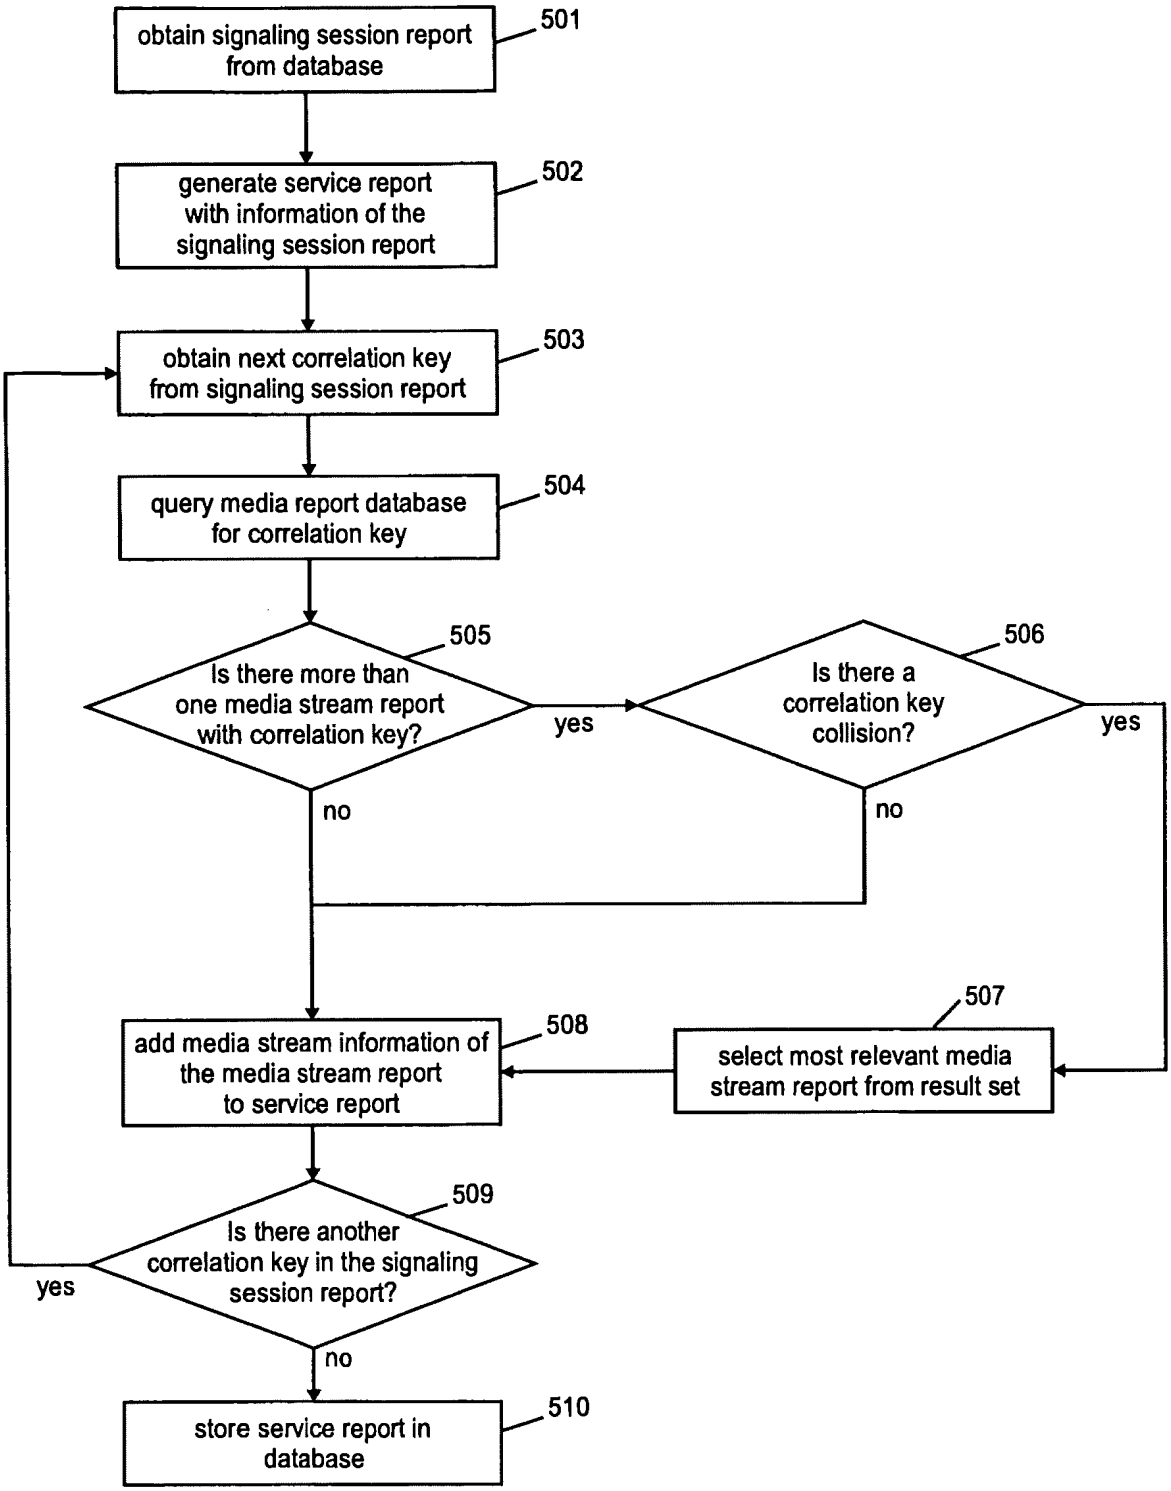
obtain signaling session report from database
501
generate service report with information of the signaling session report
502
obtain next correlation key from signaling session report
503
query media report database for correlation key
504
Is there more than one media stream report with correlation key?
505
yes
Is there a correlation key collision?
506
yes
no
no
add media stream information of the media stream report to service report
508
select most relevant media stream report from result set
507
Is there another correlation key in the signaling session report?
509
yes
no
store service report in database
510

Fig. 5

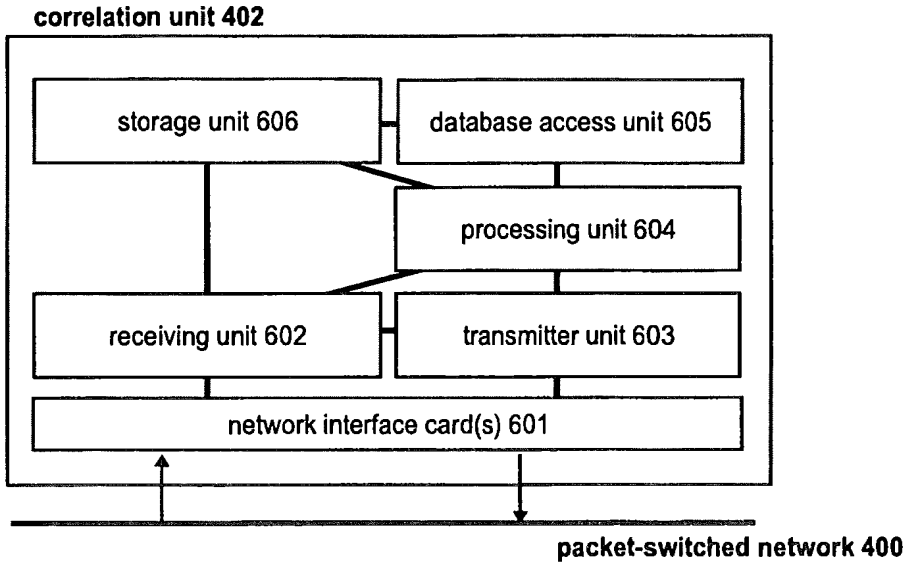
correlation unit 402
storage unit 606
database access unit 605
processing unit 604
receiving unit 602
transmitter unit 603
network interface card(s) 601
packet-switched network 400

Fig. 6

Fig. 7

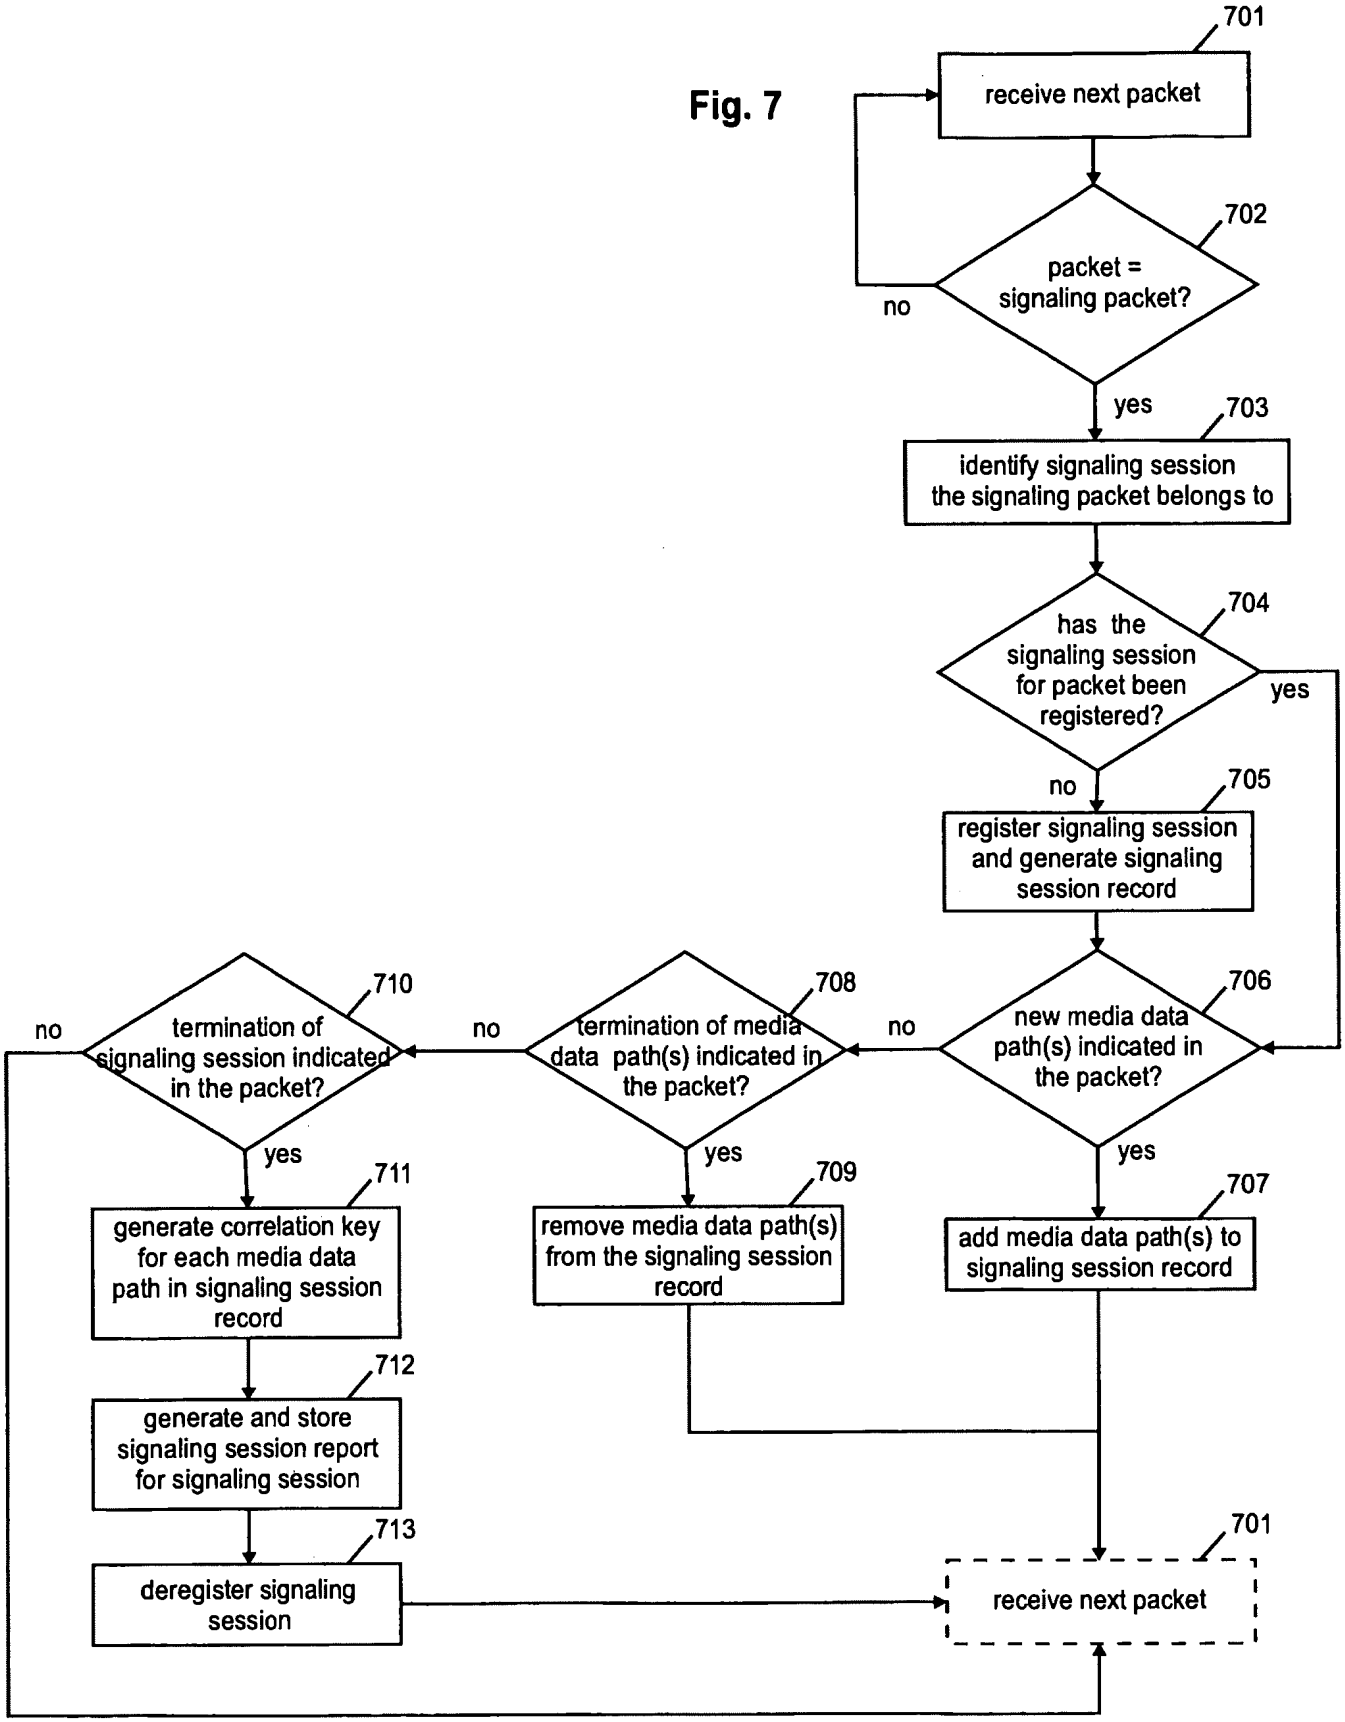
701
receive next packet
702
packet = signaling packet?
no
yes
703
identify signaling session the signaling packet belongs to
704
has the signaling session for packet been registered?
yes
no
705
register signaling session and generate signaling session record
706
new media data path(s) indicated in the packet?
no
yes
707
add media data path(s) to signaling session record
708
termination of media data path(s) indicated in the packet?
no
yes
709
remove media data path(s) from the signaling session record
710
termination of signaling session indicated in the packet?
no
yes
711
generate correlation key for each media data path in signaling session record
712
generate and store signaling session report for signaling session
713
deregister signaling session
701
receive next packet

Fig. 8

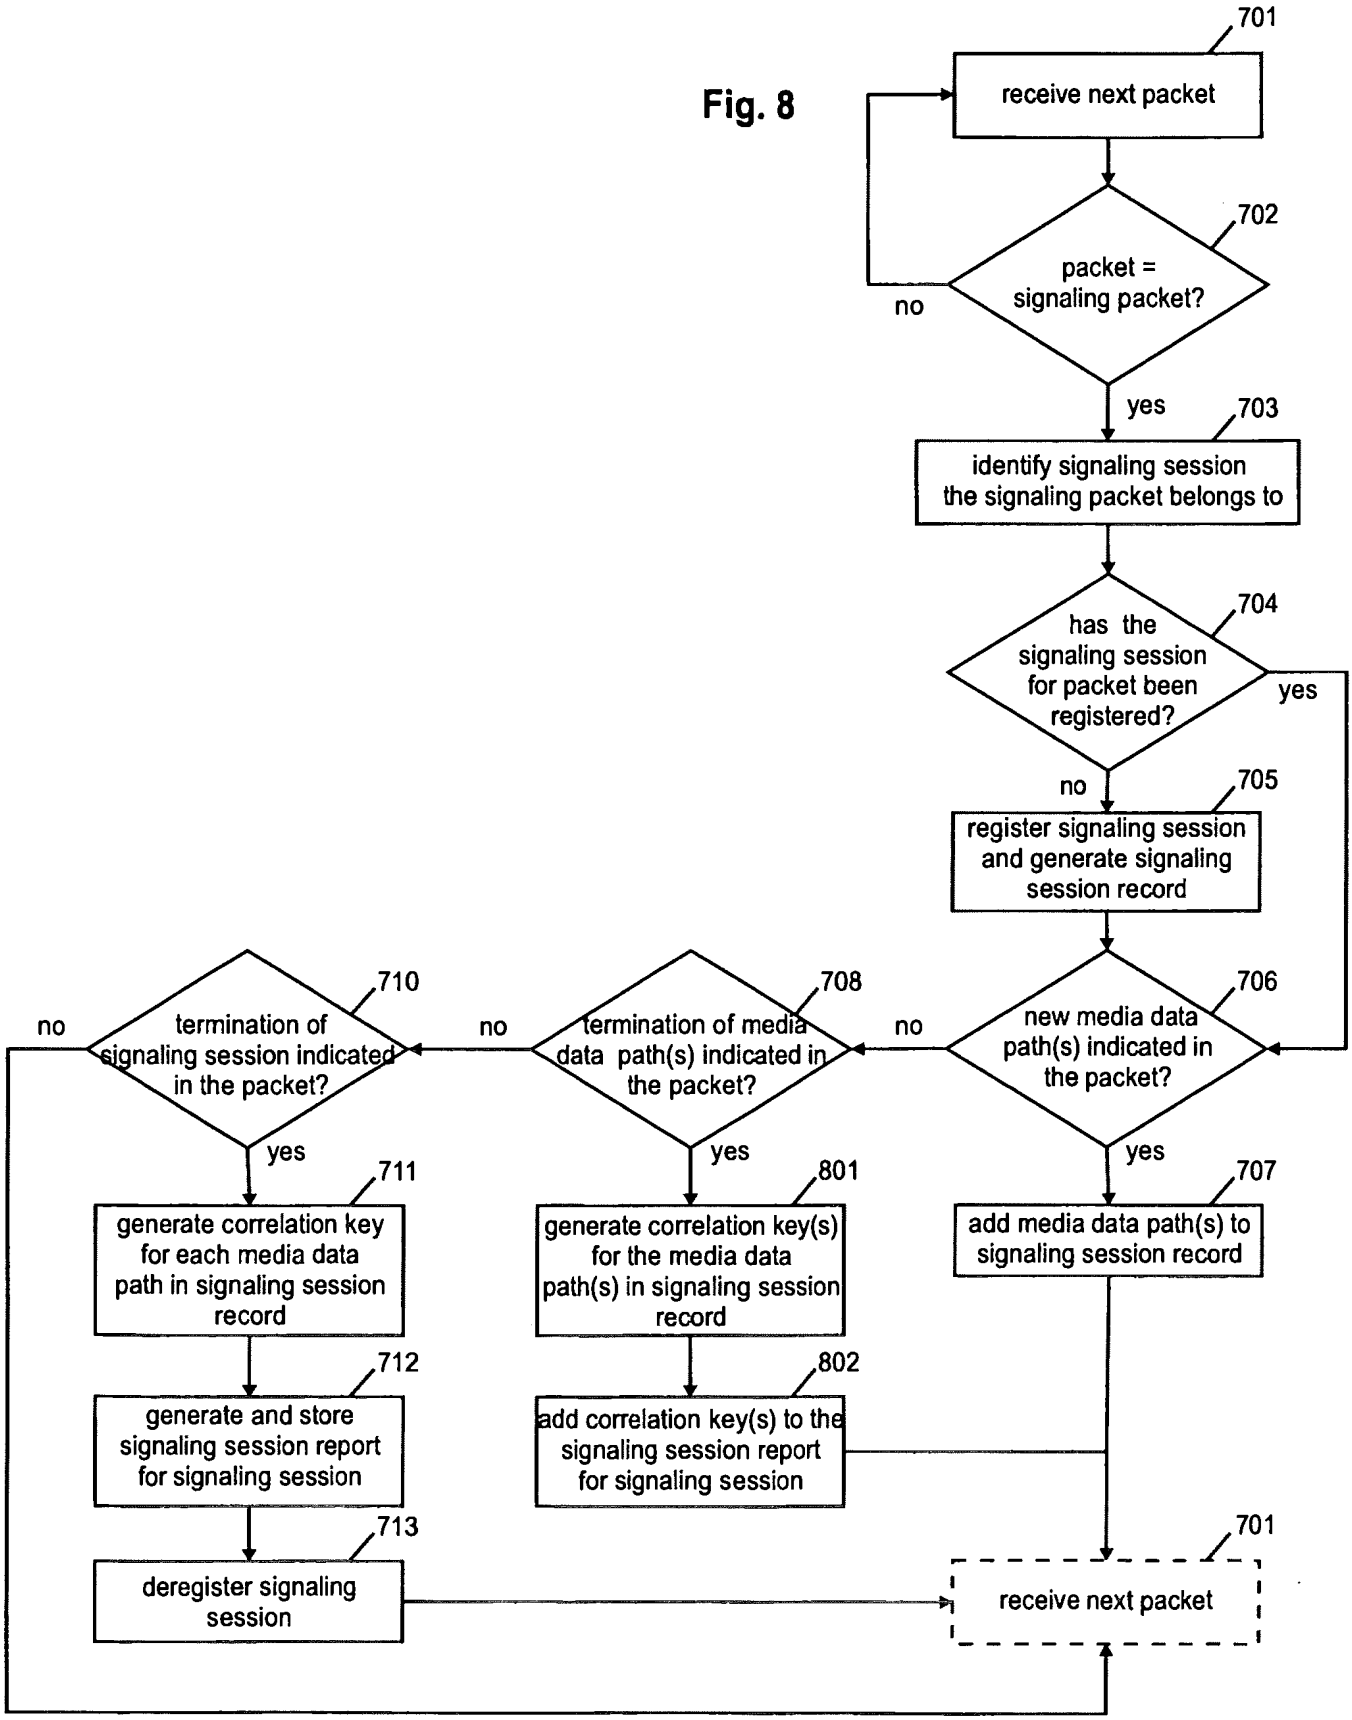
701
receive next packet
702
packet = signaling packet?
no
yes
703
identify signaling session the signaling packet belongs to
704
has the signaling session for packet been registered?
yes
no
705
register signaling session and generate signaling session record
706
new media data path(s) indicated in the packet?
no
yes
708
termination of media data path(s) indicated in the packet?
no
yes
710
termination of signaling session indicated in the packet?
no
yes
707
add media data path(s) to signaling session record
801
generate correlation key(s) for the media data path(s) in signaling session record
802
add correlation key(s) to the signaling session report for signaling session
711
generate correlation key for each media data path in signaling session record
712
generate and store signaling session report for signaling session
713
deregister signaling session
701
receive next packet

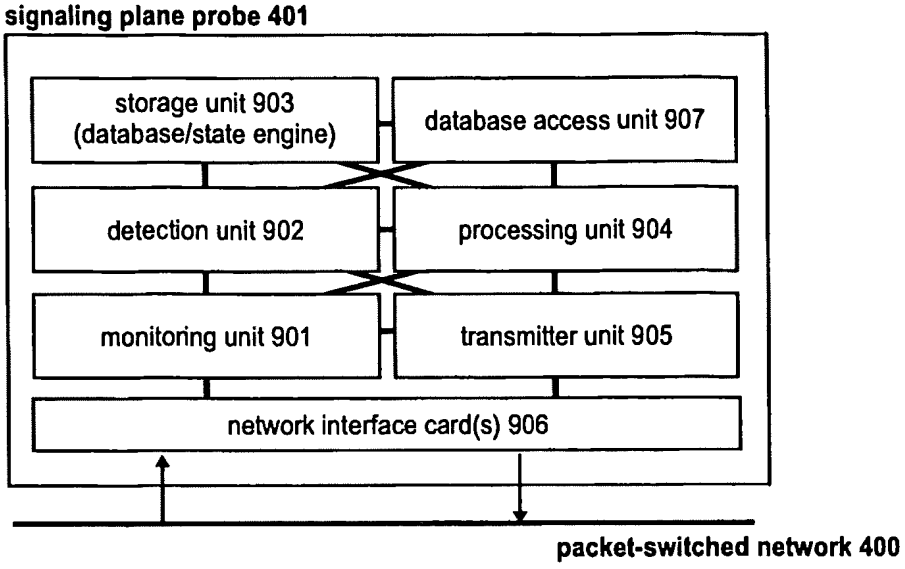
signaling plane probe 401
storage unit 903 (database/state engine)
database access unit 907
detection unit 902
processing unit 904
monitoring unit 901
transmitter unit 905
network interface card(s) 906
packet-switched network 400

Fig. 9

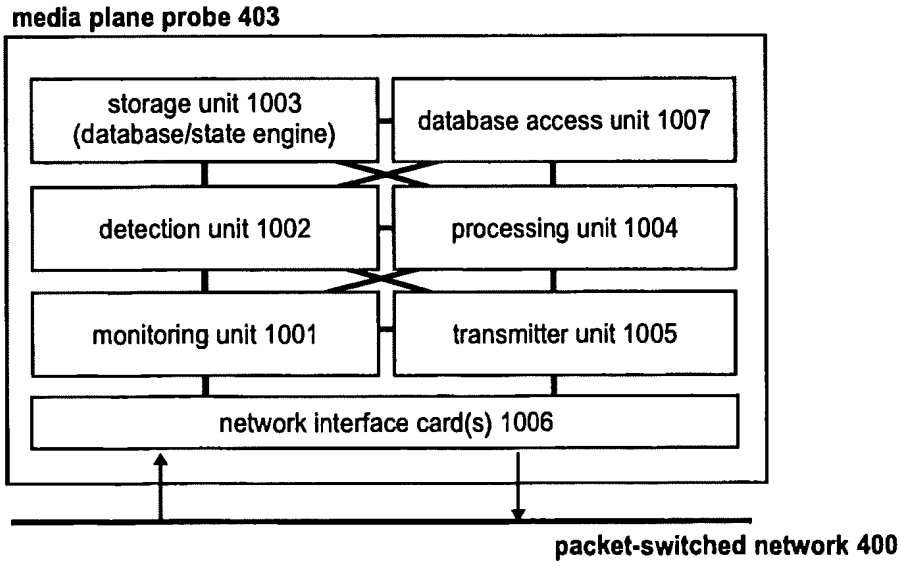
media plane probe 403
storage unit 1003 (database/state engine)
database access unit 1007
detection unit 1002
processing unit 1004
monitoring unit 1001
transmitter unit 1005
network interface card(s) 1006
packet-switched network 400

Fig. 10

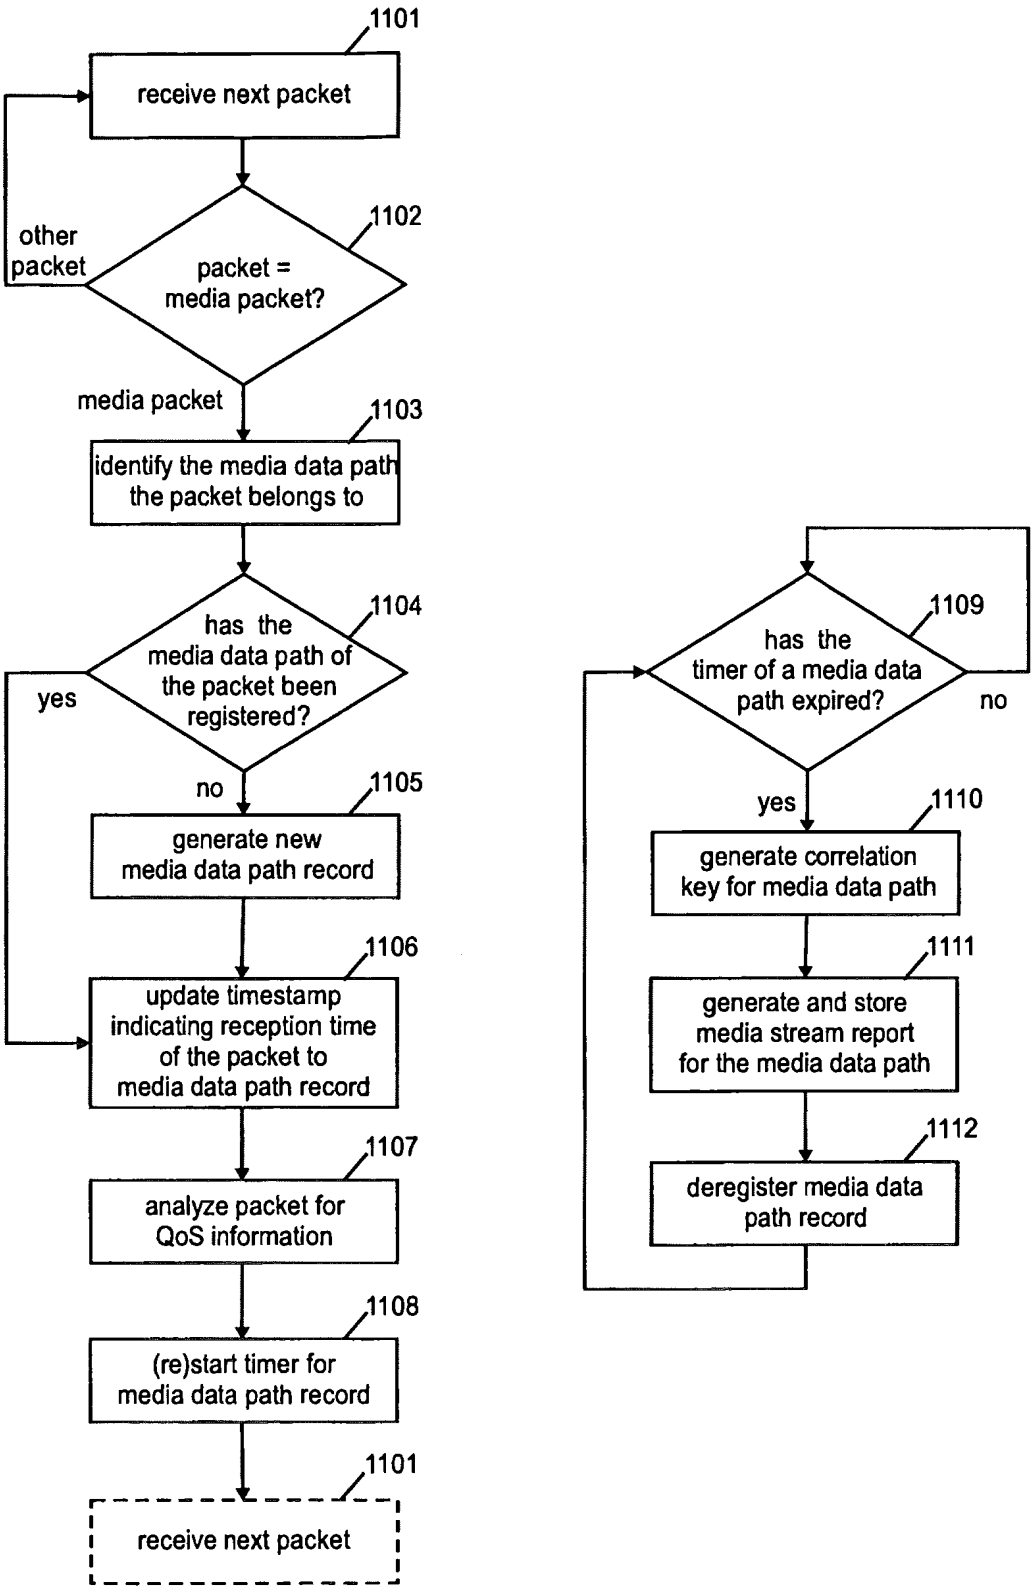

Fig. 11

BNSDOCID: <WO_____2012146265A1_I_>

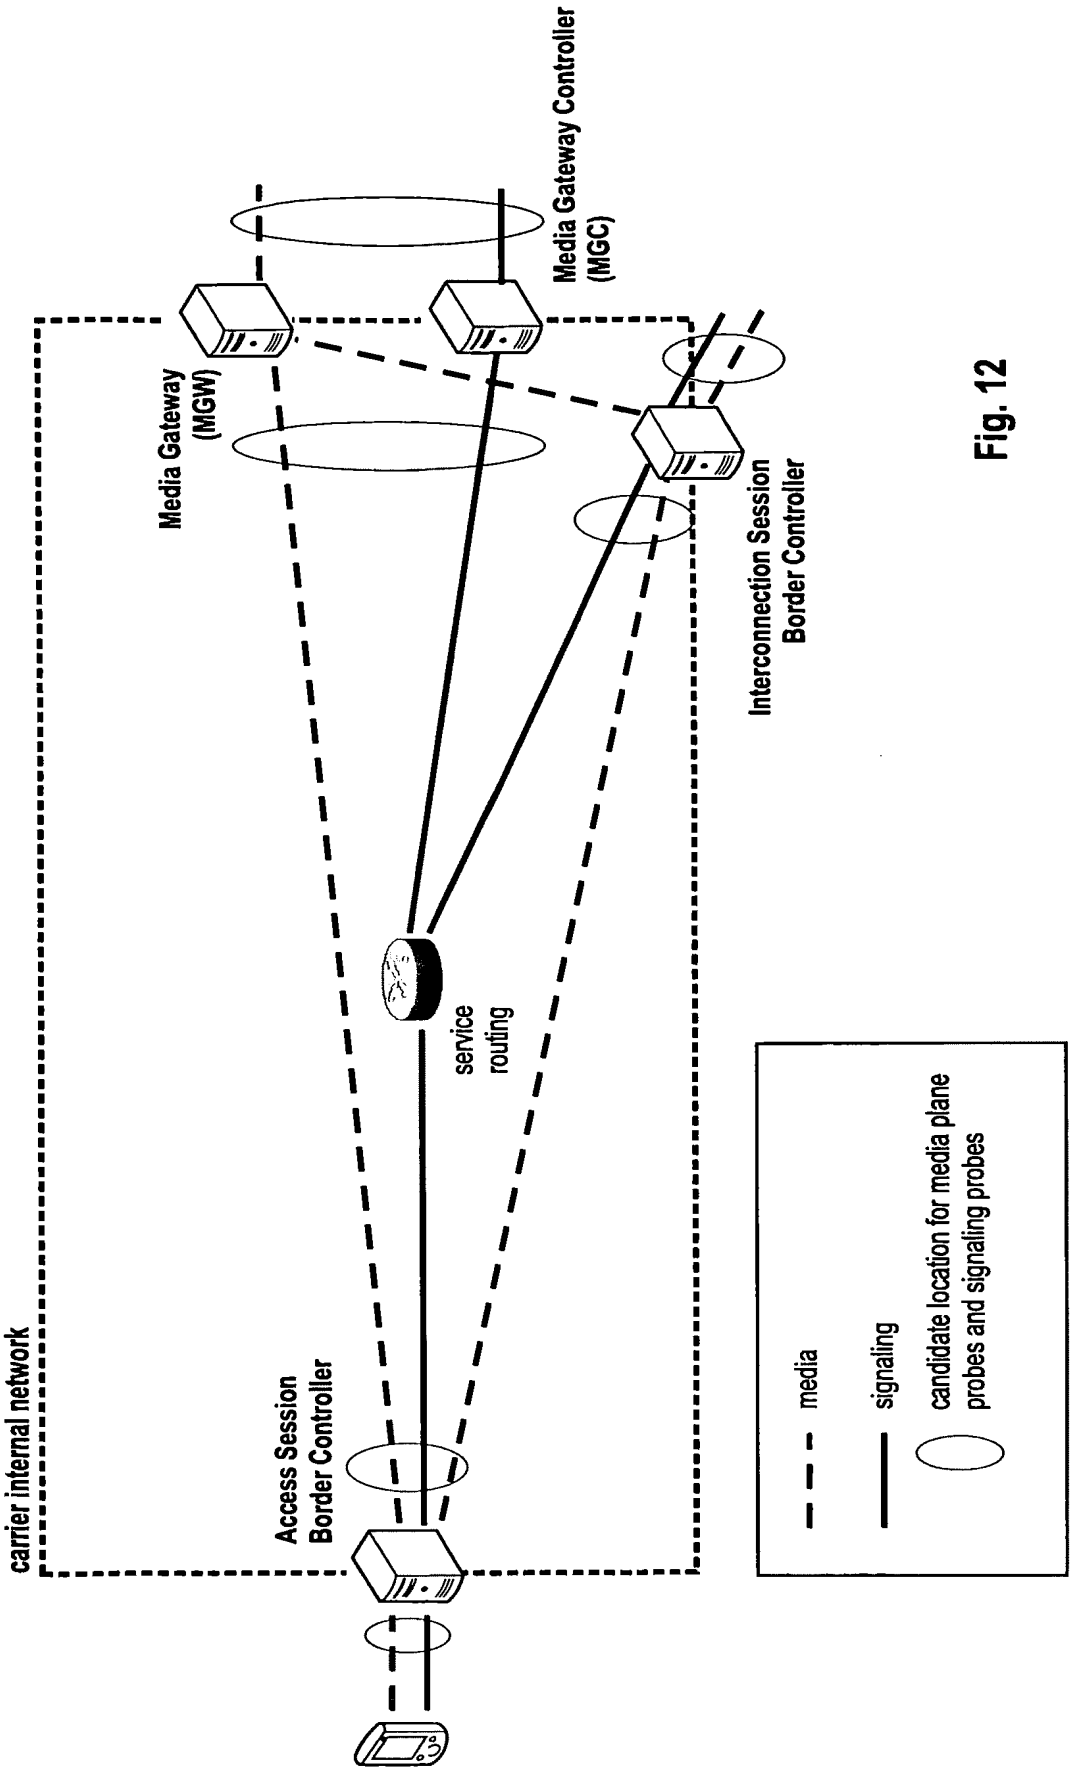
Media Gateway Controller (MGC)
Media Gateway (MGW)
Interconnection Session Border Controller
service routing
Access Session Border Controller
carrier internal network
media
signaling
candidate location for media plane probes and signaling probes

Fig. 12

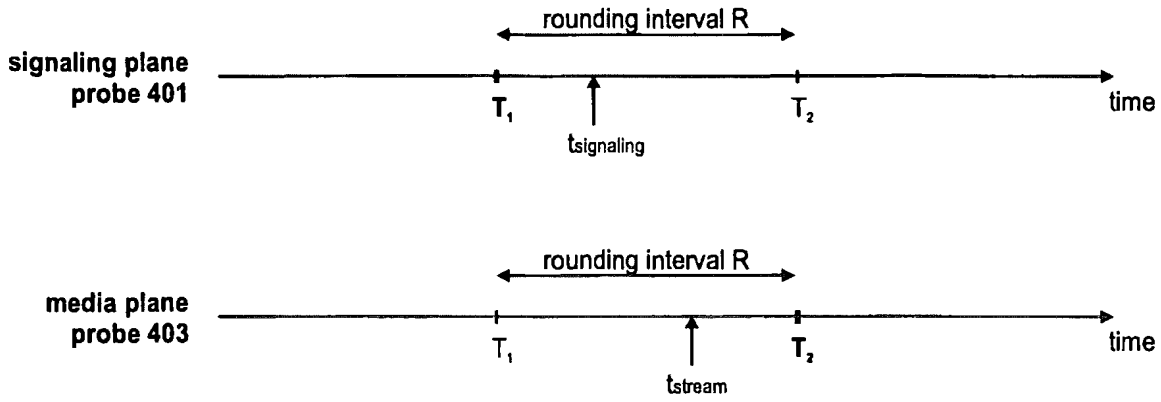
rounding interval R
signaling plane
probe 401
T1
tsignaling
T2
time
rounding interval R
media plane
probe 403
T1
tstream
T2
time

Fig. 13

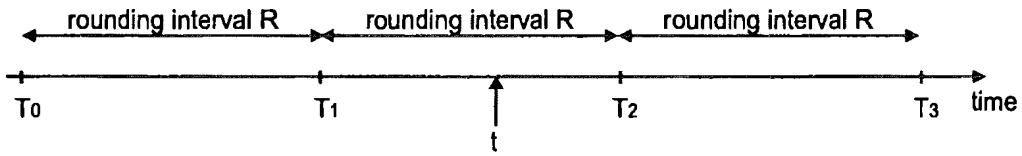
rounding interval R
rounding interval R
rounding interval R
T0
T1
t
T2
T3
time

Fig. 14

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2011/002145

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L29/06 H04L12/24
ADD. H04L12/26

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal, COMPENDEX, INSPEC, IBM-TDB, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	ATSUSHI KOBAYASHI ET AL: "VoIP Measurement Architecture Using Data Mediation", 29 October 2009 (2009-10-29), IP OPERATIONS AND MANAGEMENT, SPRINGER BERLIN HEIDELBERG, BERLIN, HEIDELBERG, PAGE(S) 15 - 26, XP019133039, ISBN: 978-3-642-04967-5 page 18, line 5 - page 20, line 35	1-27
A	US 2008/151918 A1 (FOTI GEORGE [CA]) 26 June 2008 (2008-06-26) the whole document	1-27
A	WO 2010/106075 A1 (KONINKL KPN NV [NL]; TNO [NL]; STOKKING HANS MAARTEN [NL]; WALRAVEN FA) 23 September 2010 (2010-09-23) page 4, line 16 - page 8, line 26	1-27

[ ] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
10 January 2012	17/01/2012

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Aura Marcos, F

3

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/EP2011/002145

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2008151918 A1	26-06-2008	US 2008151918 A1	26-06-2008
		WO 2008078209 A2	03-07-2008
WO 2010106075 A1	23-09-2010	EP 2409432 A1	25-01-2012
		KR 20110125674 A	21-11-2011
		WO 2010106075 A1	23-09-2010

Espacenet

Bibliographic data: EP1006701 (A2) — 2000-06-07

Adaptive re-ordering of data packet filter rules

Inventor(s): KRISHNAN P [US]; RAZ DANNY [US]; SUGLA BINAY [US] ± (KRISHNAN, P, ; RAZ, DANNY, ; SUGLA, BINAY)

Applicant(s): LUCENT TECHNOLOGIES INC [US] ± (LUCENT TECHNOLOGIES INC)

Classification: **- international:** ***G06F13/00; G06F17/30; H04L29/06;*** (IPC1-7): H04L29/06
- cooperative: **H04L63/0227**; **H04L63/0263**

Application number: EP19990309593 19991130

Priority number (s): US19980204464 19981203

Also published as: EP1006701 (A3) EP1006701 (B1) CA2287689 (A1) CA2287689 (C) US2003051165 (A1) more

Abstract of EP1006701 (A2)

A packet data filter which stores ordered rules and sequentially applies the rules to received data packets to determine the disposition of the data packet. The packet filter maintains a match count in memory which indicates the number of times each rule matched an incoming data packet. Periodically, at the initiation of a user, or based on operating parameters of the filter, the rules are automatically re-ordered based on the match count. As a result of the re-ordering, rules with higher match counts are moved earlier in the sequential evaluation order and rules with lower match counts are moved later in the sequential evaluation order. As such, rules which are more likely to match incoming data packets are evaluated earlier, thus avoiding the evaluation of later rules.; In order to prevent a re-ordering which would change the overall security policy of the packet filter, pairs of rules are compared to determine if they conflict (i.e., the swapping of the two rules would result in a change in the overall security policy). During re-ordering, the swapping of conflicting rules is prevented.

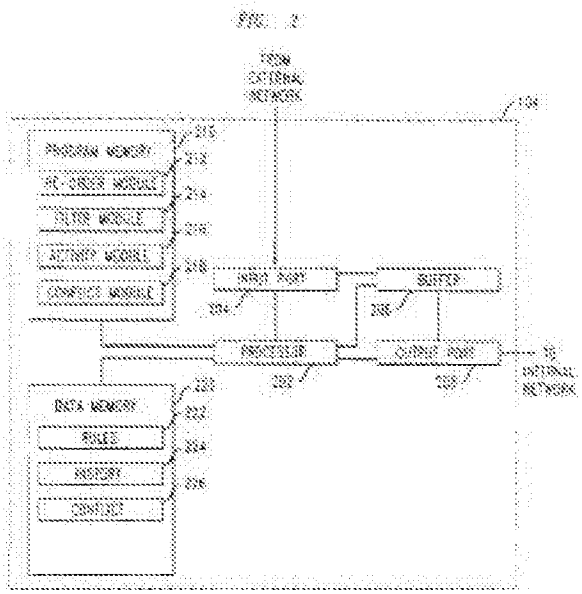

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 006 701 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
07.06.2000 Bulletin 2000/23

(51) Int Cl.[7]: **H04L 29/06**

(21) Application number: **99309593.4**

(22) Date of filing: **30.11.1999**

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **03.12.1998 US 204464**

(71) Applicant: **LUCENT TECHNOLOGIES INC.**
Murray Hill, New Jersey 07974-0636 (US)

(72) Inventors:
- **Krishnan, P.**
  Matawan, New Jersey 07747 (US)
- **Raz, Danny**
  Aberdeen, New Jersey 07747 (US)
- **Sugla, Binay**
  Aberdeen, New Jersey 07747 (US)

(74) Representative:
Buckley, Christopher Simon Thirsk et al
Lucent Technologies (UK) Ltd,
5 Mornington Road
Woodford Green, Essex IG8 0TU (GB)

(54) Adaptive re-ordering of data packet filter rules

(57) A packet data filter which stores ordered rules and sequentially applies the rules to received data packets to determine the disposition of the data packet. The packet filter maintains a match count in memory which indicates the number of times each rule matched an incoming data packet. Periodically, at the initiation of a user, or based on operating parameters of the filter, the rules are automatically re-ordered based on the match count. As a result of the re-ordering, rules with higher match counts are moved earlier in the sequential evaluation order and rules with lower match counts are moved later in the sequential evaluation order. As such, rules which are more likely to match incoming data packets are evaluated earlier, thus avoiding the evaluation of later rules. In order to prevent a re-ordering which would change the overall security policy of the packet filter, pairs of rules are compared to determine if they conflict (i.e., the swapping of the two rules would result in a change in the overall security policy). During re-ordering, the swapping of conflicting rules is prevented.

FIG. 2

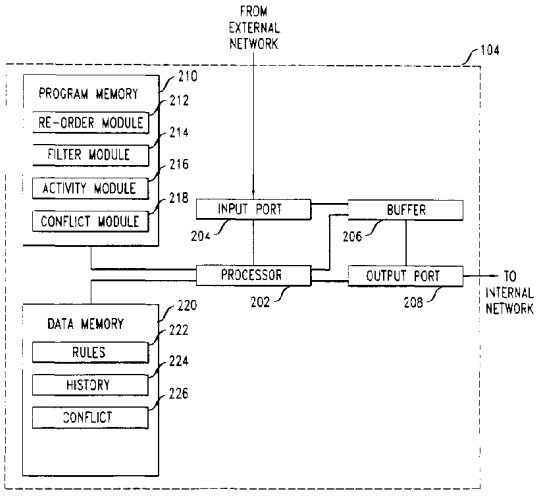

EP 1 006 701 A2

Printed by Jouve, 75001 PARIS (FR)

CISCO EXHIBIT 1002
File History of 9565213 Patent

Description

Field of the Invention

[0001] The present invention relates to data packet filters. More particularly, the present invention relates to adaptive re-ordering of data packet filter rules to improve the performance of the filter while maintaining a security policy.

Background of the Invention

[0002] Data networks are becoming increasingly prevalent. For example, many companies have internal data networks to allow for intra-company communication. Such communication includes email, documents, voice, video, and multimedia. Further, these internal data networks are generally connected to an external data network (e.g. Internet), to allow for the exchange of information between the internal and external networks. However, as a result of the interconnection of data networks, security has become a major concern. The unauthorized access to a company's data network can result in the loss of valuable proprietary information. Further, unauthorized access can also result in attacks on the company's data network and computing system, resulting in loss of data and computer system crashes.

[0003] A computer programmed to implement a security policy for a data network is called a firewall. Typically, firewalls are located at strategic points in the network such that all incoming and outgoing data traffic must pass through a firewall. At these points, all network traffic is trapped and examined by the firewall to determine if the traffic should be allowed to pass through the firewall. There are various techniques for implementing a firewall. One such implementation is a computer programmed to implement a data packet filter. A data packet filter examines all received data packets, and either denies or allows passage of the data packet based on the contents of the data packet. In one implementation, this decision is based on rules stored in the data packet filter which define which data packets are allowed to pass and which data packets will be blocked. Each stored rule defines certain parameters of data packets (e.g. source and destination) and also includes the disposition of data packets which match these parameters. For example, a rule may indicate that any data packet from source X will be blocked, while another rule may indicate that any data packet from source Y to destination Z will be allowed to pass. The rules are stored in sequential order in the data packet filter and every data packet received by the firewall is tested against the rules in sequential order. The first rule whose parameters match the received packet is applied to the packet and the packet is treated as indicated by the rule.

[0004] In such a rules based data packet filter, a system administrator must first define a security policy to be implemented by the firewall, and then must program the data packet filter with rules which will implement that security policy. One major problem with these types of firewalls is that sophisticated security policies require a large number of rules in order to implement the policy. Since every data packet which passes through the firewall must be checked against these rules, the performance of the firewall degrades as the number of rules increases.

[0005] One known solution to the performance problem in a data packet filter firewall is to use a memory cache. In such a system, when a data packet arrives, the relevant parameters (e.g. source and destination) of the data packet are stored in a cache. In addition, after the packet filter rules have been applied to the received data packet, the disposition (e.g. allow or deny) is also stored in the cache associated with the relevant parameters of the received data packet. Thereafter, if a data packet is received with parameters which are the same as parameters previously stored in the cache, the firewall can apply the associated disposition without applying all the rules to the data packet. This enhances performance in view of the fact that for certain applications, ongoing communications will occur between two computers, and there is no need to check every data packet exchanged between the computers during the communication session (i.e., connection). Thus, while this technique improves performance for data packets exchanged during connections, the technique does not improve performance for new connections.

[0006] Packet filters may also be implemented on a client computer, for example a personal computer running World Wide Web (WWW) browsing software (e.g., Microsoft Explorer or Netscape Navigator). Many WWW sites contain material which parents may deem unsuitable for children. As such, several companies offer filtering software which implements packet filters for blocking unsuitable material. These packet filters implement a security policy through the use of rules as described above. As the number of rules increases, these client based packet filters suffer from the same performance problems as described above.

[0007] What is needed is a technique for improving the performance of a data packet filter as the number of rules required to implement a security policy increases.

Summary of the Invention

[0008] As described above, a data packet filter stores a plurality of ordered rules which are sequentially applied to received data packets to implement a security policy. In accordance with the invention, the rules are automatically re-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 669

ordered to improve the performance of the packet filter. Rules which match incoming data packets more frequently are moved earlier in the ordering, and rules which match incoming data packets less frequently are moved later in the ordering. Since the first rule that matches a received data packet controls the disposition of the packet, once a rule matches a packet, the remaining rules need not be evaluated. By re-ordering the rules, data packets are matched against rules more quickly, and the performance of the data packet filter is improved.

[0009] In one embodiment of the invention, the packet filter dynamically maintains, during operation of the packet filter, a count of the number of times each rule matches a received data packet. The packets are then re-ordered based on these counts with the rules having higher match counts being moved earlier in the rules order. Such re-ordering is made on the premise that past history of the packet filter provides a reliable indication of future operation. Thus, by dynamically maintaining match counts during operation, the packet filter can adaptively re-order its rules based on past operating parameters. Such adaptive re-ordering is beneficial to protect against so-called denial of service attacks, as described in further detail below.

[0010] In one embodiment of the invention, rules are re-ordered by swapping pairs of rules in the rules order. However, the sequence of the rules plays a part in defining the security policy of the packet filter. As such, the indiscriminate swapping of rules may result in changing the security policy of the packet filter. Such a result is undesirable and must be prevented. As such, rules are evaluated to determine re-orderings which are not allowed because they result in a change in the security policy. In one embodiment, such an evaluation is made by comparing pairs of rules to determine which pairs of rules conflict with each other. Two rules are said to conflict with each other if swapping the rules would result in changing the security policy of the packet filter. The result of this conflict evaluation is stored in the memory of the packet filter and subsequent re-ordering refers to this information in order to prevent the swapping of conflicting rules.

[0011] The re-ordering in accordance with the invention may be performed periodically, at the initiation of a user, or based on the some operating parameters of the packet filter (e.g. when the performance of the packet filter drops below a threshold).

[0012] The principles of the present invention may be applied to various types of computing devices. For example, and without limitation, the invention may be implemented in a dedicated network firewall which filters data packets for many computers and networks, or in a stand-alone personal computer such that the packet filter only filters data packets for the personal computer.

[0013] These and other advantages of the invention will be apparent to those of ordinary skill in the art by reference to the following detailed description and the accompanying drawings.

Brief Description of the Drawings

[0014]

Fig. 1 shows a network of the type in which the present invention may be implemented;
Fig. 2 shows a computer configured to implemented a firewall in accordance with the present invention;
Fig. 3 shows the format of an exemplary data packet;
Fig. 4 shows an example of a rule set;
Fig. 5 shows a conflict table; and
Fig. 6 shows a personal computer in which the principles of the present invention may be implemented.

Detailed Description

[0015] A network of the type in which the present invention may be implemented is shown in Fig. 1. An internal data network, such as a corporate intranet, is shown as network 102. An external data network, such as the Internet, is shown as network 106. Firewall 104 is shown positioned such that all data traffic exchanged between internal network 102 and external network 106 must pass through the firewall 104. These data networks communicate with each other using the well known Internet Protocol (IP) for data packet communication. Assume that firewall 104 is administered by the same entity which administers internal network 102. The purpose of firewall 104 is to implement the security policy of the internal network 102.

[0016] The firewall 104 may be implemented by a properly configured computer, further details of which are shown in Fig. 2. Fig. 2 shows firewall 104 comprising a processor 202 which controls the overall operation of the firewall 104. Processor 202 is connected to input port 204, buffer 206, output port 208, program memory 210, and data memory 220. Program memory 210 and data memory 220 may be any type of well known computer storage device. In operation, processor 202 controls the operation of the firewall 104 by executing computer program instructions stored in program memory 210. Program memory 210 includes the following program modules: re-order module 212, filter module 214, activity module 216, and conflict module 218. These modules contain computer program code which, when executed

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 670

by processor 202, implements a packet filter in accordance with the present invention. It is to be understood that the program modules shown in program memory 210 are high level representations of functional computer program modules. In various implementations, these modules could include program source code, object code, or a combination of source and object code. These modules will be discussed in further detail below. Of course, a computer implementing a firewall would have other well known components as well. Such other components are not shown in Fig. 2 for clarity.

[0017] The operation of firewall 104 is as follows. When a data packet is received at firewall 104 from external network 106, the data packet is received at input port 204. The format of an exemplary data packet 302 is shown in Fig. 3. Data packet 302 contains a header portion 304 and data portion 306. The header portion 304 contains the following fields: protocol 308, source address 310, destination address 312, source port 314, and destination port 316. The protocol field 308 indicates the protocol with which the current packet is being transmitted. For example, the protocol field 308 may indicate that the data packet 302 is being transmitted using the well known TCP/IP protocol. The source address 310 indicates the IP address of the source of the data packet 302. IP addressing is well known in data networking. IP addresses are 32 bits, and are most commonly represented in a dotted notation as 4 integers, each between 0 and 255, and each separated by a dot (e.g. 125.180.104.4). The destination address 312 indicates the IP address of the destination of the data packet 302. The source port 314 indicates the port number of the source computer. A port number is an integer between 0 and 65,535, and generally is associated with an application that runs on a particular machine. For example, web servers are typically associated with port number 80. The destination port 316 indicates the port number of the destination computer. The data portion 306 of the packet 302 contains the payload of the packet which will be used by an application executing on the destination computer.

[0018] Upon receipt of the data packet at input port 204, the data packet 302 is stored in buffer 206 while the processor 202 executes the filter module 214 in order to apply the filter rules 222 stored in data memory 220 against the received data packet. The filter module 214 applies the rules to the received data packet in sequential order. The first rule that matches the data packet is the rule which will decide whether the data packet is allowed or denied. Once a rule matches the data packet, there is no need to apply the remaining rules.

[0019] An example of a rule set which would be stored as rules 222 is shown in Fig. 4. As shown in Fig. 4, each rule has a sequence number, which indicates the order in which the rules will be applied against an incoming data packet. Each rule also has fields indicating protocol, source address, destination address, source port, destination port, and an action. Each rule indicates the action to be taken when a data packet is received which matches the information in the rule's fields. A packet is said to match the rule if the header information in the packet falls within the defined ranges of values in the above described fields. For example, rule number 1 indicates that any packet received using the TCP protocol with a source address of 125.30.20.200 is to be denied (i.e. blocked), regardless of the contents of the other fields of the packet. The "*" in a rule field indicates a wildcard which will match any value. Rule number 2 indicates that any packet received using the TCP protocol with a source address in the range 135.180.104.4 to 135.180.104.225, a destination address in the range 135.180.204.4 to 135.180.204.255, a source port in the range 1024-65535, and a destination port of 80 is to be allowed (i.e. passed to its destination). Rule number 3 indicates that any packet received using the TCP protocol from a source address in the range 125.180.100.0 - 125.180.100.255 is to be allowed, regardless of the values in the other fields. Finally, rule number 4 indicates that all data packets are to be denied. Since rules are applied in sequential order, rule number 4 will only be applied if none of the rules 1-3 match the incoming data packet. The inclusion of a rule such as rule number 4 is common practice in firewalls and implements a security policy that all packets are denied unless specifically allowed. As would readily be recognized by one skilled in the art, a wide variety of security policies may be implemented by storing different rule sets as rules 222 in data memory 220. It is also noted that the action indicated in a rule may be something other than allow or deny. For example, it is possible that an action indicates that packets matching the rule be forwarded to some other location for further processing.

[0020] Thus, while the received data packet is stored in buffer 206, processor 202 sequentially applies rules 222 to the data packet until a rule matches the data packet. If the matched rule indicates that the data packet is to be allowed, then the packet is transferred to output port 208 and is allowed to pass to the internal network 102. If the matched rule indicates that the data packet is to be denied, then the packet is deleted from the buffer 206 and is thus not allowed to pass to the internal network 102. It is noted that although the rules are applied in sequential order, they will not necessarily be stored internally in data memory 220 in that sequential order. For various reasons, the internal representation of the rules may be different from the logical representation shown in fig. 4. However, regardless of the internal representation, each rule will be associated with a sequence number or other indication of order of evaluation (e.g. priority), such that the filter module 214 will apply the rules to a received data packet in a given sequence.

[0021] As the number of stored rules 222 increases, the performance of the firewall 104 generally degrades because it takes longer to apply the rules to each incoming data packet. The present invention solves this performance problem by re-ordering the stored rules 222 such that rules which are more likely to match incoming data packets are placed earlier in the sequential order (i.e. assigned a lower sequence number). This improves the performance of the firewall because once a rule is matched, the remaining rules are not applied, thus avoiding the processing time required to apply those remaining rules. If rules which are more likely to match data packets can be moved earlier in the rules

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 671

sequence, performance of the firewall will improve.

[0022] However, prior to re-ordering the rules, it must be determined which rules may be re-ordered without having an affect on the overall security policy of the firewall 104. Due to the nature of the rules and the sequential application of the rules to received data packets, certain re-orderings may result in a change in the overall security policy. Such re-orderings must be prevented. For example, referring to the example rule set of fig. 4, rule number 4 could not be swapped with rule 1, because such a swap will result in rule number 4 being applied first. Since rule number 4 indicates that all data packets are to be denied, placing it first in the rules sequence would result in all data packets being denied. None of the other rules would be reached to allow the data packets to pass. Therefore, it is necessary to determine which rules may be re-ordered while still maintaining the security policy of the firewall.

[0023] In order to determine which rules may be re-ordered, conflict module 218 is executed by processor 202. The conflict module analyzes the rule set 222 to determine whether pairs of rules are in conflict such that the swapping of the pair would result in a change in the security policy. Generally, two rules are said to conflict with each other if at least one data packet could match either one of the rules and the actions of the rules are not the same. In the embodiment described herein, the conflict algorithm to determine whether a Rule-A conflicts with a Rule-B is as follows:

```
If ACTION of Rule-A is not the same as ACTION of Rule-B then

    if the PROTOCOL of Rule-A does not intersect with the
    PROTOCOL OF Rule-B then
```


CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 672

```
                    if the SOURCE ADDRESS of Rule-A does not intersect
                    with the SOURCE ADDRESS of Rule-B then

                          if the DESTINATION ADDRESS of Rule-A does not
                          intersect with the DESTINATION ADDRESS of Rule-
                          B then

                                if the SOURCE PORT of Rule-A does not
                                intersect with the SOURCE PORT of Rule-B
                                then

                                      if the DESTINATION PORT of Rule-A
                                      does not intersect with the
                                      DESTINATION PORT of Rule-B then

                                            no conflict between Rule-A
                                            and Rule-B
                                      else
                                otherwise
                          else
                    else
              else
                    conflict between Rule-A and Rule-B.
        else

              no conflict.
```

[0024] In accordance with the algorithm, a particular of a rule is said to *intersect* with a particular field of another rule if there is any intersection in the range of values of the two fields. For example, with respect to fields which contains IP addresses, the fields will intersect if there is any overlap between the addresses. For example, an IP address field containing the address range 123.30.20.100 intersects with an IP address field containing the address range 123.30.20.70 - 123.30.20.150 in the address range 123.30.20.70 - 123.30.20.100. For example, the IP address 123.30.20.85 would fall within both ranges. Thus, in accordance with the algorithm, if the actions of two rules are different, and if the there is an intersection in any of the fields of the two rules, then the rules conflict. Otherwise, the rules do not conflict. Of course, one skilled in the art could modify the above algorithm in order to detect the conflict between rules with a different set of fields. Such an algorithm would analyze each field of each rule to determine if there were an intersection. If there is any intersection, and if the action of the two rules is different, then the rules conflict.

[0025] It is noted that the conflict module 218 is executed once when the rule set 222 changes (i.e. a rule is added or deleted). The results of the execution of the conflict module 218 are stored in data memory 220 as conflict data 226 which indicates, for every possible pair of rules, whether that pair conflicts. In one embodiment, a two-dimensional conflict table such as conflict table 500 shown in Fig. 5 could be stored as conflict data 226 as a result of the execution of conflict module 218. For each pair of rules, the intersection point in the conflict table 500 indicates whether the rules are in conflict (C) or not in conflict (NC). For example, the intersection of row 3 and column 2 in conflict table 500

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 673

contains NC indicating that rule 3 and rule 2 do not conflict. If a rule is added, only the interaction of the new rule with the other rules needs to be checked with the conflict module 218. The interaction of the other rules will not change. It is noted that the conflict table 500 is symmetrical in that the intersection of column x and row y will contain the same value as the intersection of column y and row x. As such, only half of the table needs to be stored in conflict data 226 in data memory 220. However, the entire table is shown in Fig. 5 for purposes of this description.

[0026] Returning now to the re-ordering of rules, as was stated above, in accordance with the invention, the performance of a packet data filter is improved by re-ordering the stored rules 222 such that rules which are more likely to match incoming data packets are placed earlier in the sequential order. In accordance with one aspect of the invention, the decision as to which rules are more likely to match incoming data packets is made with reference to past history of the firewall 104. In one embodiment, this past history is a record of the number of times each rule has matched an incoming data packet. In accordance with this embodiment, during operation of the firewall 104, an activity module 216 dynamically maintains a running count of the number of times each rule matched an incoming packet. This count is stored as a table in the history section 224 of data memory 220. The table associates a count with each rule. Thus, rules associated with higher counts matched incoming packets more often than rules associated with lower counts. It is noted that the counts can periodically be reinitialized to zero.

[0027] The re-ordering of the rule set 222 may be performed periodically (e.g. every 30 minutes), may be initiated by a system administrator, or may be initiated by the firewall 104 based on certain operating parameters of the firewall 104. For example, the firewall 104 may initiate a re-ordering when the performance of the firewall 104 drops below a certain predetermined threshold. The re-ordering is performed by the re-order module 212 by swapping rules having a lower count in history section 224 with rules having a higher count in history section 224. This results in rules which matched data packets more often being placed earlier in the sequential ordering.

[0028] One algorithm for swapping rules is as follows.

```
For i = 1 to N

    J = N

    While J > 0

        If match count for Rule J > match count for Rule J-1
        then

            if Rule J does not conflict with Rule J-1 then

                Swap Rule J with Rule J-1.

        J = J - 1.
```

This algorithm attempts to move the rules with a higher match count earlier in the sequence of rules. One constraint of the re-ordering of rules is that rules are only swapped if doing so will not change the overall security policy of the rule set (i.e. the rules do not conflict). While this algorithm uses a well known bubble sort type of algorithm, one skilled in the art could readily implement various types of re-ordering algorithms.

[0029] One benefit of the present invention is that the filter rules are re-ordered based on the activity of the firewall. As such, the firewall is dynamically adaptable to changing operating conditions. As an example, one type of attack on a computer system is known as a denial of service attack, which is aimed at reducing the availability of network elements. One technique for mounting such an attack is to continuously transmit packets to a network, which packets are known to be blocked by the network's firewall. If the rule which rejects such packets is late in the sequential ordering of rules, then every time such a packet is received by the firewall, the firewall must apply many rules before rejecting the packet. For example, if the rule which will eventually reject the data packet is rule number 100, then each time

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 674

such a packet is received, the firewall must apply 100 rules prior to rejecting the packet. By sending many such packets, an attacker may seriously impair the performance of the firewall.

[0030] The present invention can improve the performance during a denial of service attack as follows. If an attacker is continuously sending packets which are rejected by rule number 100, then the match count for rule number 100 will increase quickly. Further, since one of the parameters for initiating a re-ordering of rules is the performance of the firewall, then a denial of service attack will substantially degrade the performance such that a re-ordering will be initiated. At such time, since the match count of rule 100 will be high, the re-ordering module will attempt to move rule 100 earlier in the sequential ordering. For example, if the rule were moved to rule 5, then further receipt of the attacking packets will result in the packets being blocked after only 5 rules, rather than 100, an improvement of 95%.

[0031] It is to be understood that although the invention has been described in the context of a firewall implemented in a network as shown in Fig. 1, the principles of the present invention may also be implemented in a personal computer configured to filter data packets received from a data network. In one well known configuration, a personal computer is connected to the Internet using a modem and a dial-up connection to the Internet via the telephone network. Since it is often desirable to limit the Internet sites which a user of a personal computer may access, there are various filter programs available which filter incoming data packets based on filtering rules. The principles of the present invention may be implemented in such a personal computer. Such an embodiment of the invention is shown in Fig. 6 as personal computer 600. Personal computer 600 includes a processor 610 connected to a modem 612 which is used to establish a connection to a data network, such as the Internet. In one embodiment, the computer 600 executes a browser program 604 which allows the computer 600 to communicate with World Wide Web (WWW) sites and to receive information from such sites. In accordance with the invention, the computer 600 also includes a packet filter program module 606, which contains a re-order module, filter module, activity module, and intersection module, which operate to implement a packet filter as described above. Also, the computer 600 contains data memory 608 including rules, history, and conflict data, also as described above. A user of the computer 600 interacts with the computer 600 via keyboard 614, graphical input device 618 (e.g. mouse), and display 616 in a conventional manner. In operation, a user of computer 600 may request information from the Internet, and the packet filter module 606 will operate as described above in order to filter data packets in accordance with stored filter rules. The filter rules may be re-ordered to improve the performance of the packet filter, as long as such re-ordering does not change the overall security policy of computer 600.

[0032] It is noted that network routers may optionally be configured to filter data packets in a manner similar to that described above in connection with firewalls. Thus, the principles of the present invention may be implemented in a network router which is configured for data packet filtering.

[0033] The foregoing Detailed Description is to be understood as being in every respect illustrative and exemplary, but not restrictive, and the scope of the invention disclosed herein is not to be determined from the Detailed Description, but rather from the claims as interpreted according to the full breadth permitted by the patent laws. It is to be understood that the embodiments shown and described herein are only illustrative of the principles of the present invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention.

Claims

1. A method for operation of a data packet filter which implements a security policy by filtering data packets by sequentially evaluating a plurality of ordered rules, said method comprising the step of:
   automatically re-ordering said ordered rules.

2. The method of claim 1 wherein said step of re-ordering is based on prior activity of said data packet filter.

3. The method of claim 2 further comprising the step of:

   dynamically maintaining, during operation of said data packet filter, match data indicating the number of times each of said rules is matched by a data packet; and
   wherein said step of re-ordering is based on said match data.

4. The method of claim 3 wherein said rules are re-ordered such that rules which are matched more often are evaluated earlier.

5. The method of claim 1 wherein said step of automatically re-ordering is performed periodically.

6. The method of claim 1 wherein said step of automatically re-ordering is performed when the performance of said

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 675

data packet filter falls below a threshold.

7. The method of claim 1 wherein said step of re-ordering is performed by swapping pairs of rules, said method further comprising the step of:
determining pairs of rules which can be swapped without changing said security policy.

8. The method of claim 7 wherein said step of re-ordering is performed such that said security policy is not changed.

9. The method of claim 7 wherein said step of determining further comprises the step of determining the intersection of fields of pairs of said rules.

10. A data packet filter for implementing a security policy comprising:

a memory for storing a plurality of ordered rules;
a filter module for sequentially applying said rules to received data packets; and
a re-ordering module for re-ordering said ordered rules.

11. The data packet filter of claim 10 wherein said memory further stores information representing past activity of said data packet filter and wherein said re-ordering is based on said past activity.

12. The data packet filter of claim 11 wherein said information representing past activity comprises match data indicating the number of times each of said rules is matched by a data packet.

13. The data packet filter of claim 12 further comprising:
an activity module for dynamically maintaining said match data.

14. The data packet filter of claim 10 further comprising:
a conflict module for determining pairs of said rules which can be swapped without altering said security policy.

15. The data packet filter of claim 14 wherein said re-ordering module only swaps pairs of rules which can be swapped without altering said security

16. The data packet filter of claim 14 wherein said conflict module determines pairs of said rules which can be swapped without altering said security policy by determining the intersection of fields of pairs of rules.

17. A computer readable medium storing computer program instructions which are executable on a computer system implementing a data packet filter, said computer program instructions comprising instructions defining the steps of:

filtering received data packets by sequentially applying a plurality of ordered rules, said rules defining a security policy; and
automatically re-ordering said ordered rules.

18. The computer readable medium of claim 17 wherein said computer program instructions further comprise instructions defining the steps of:
re-ordering said rules based on prior activity of said computer system.

19. The computer readable medium of claim 18 wherein said computer program instructions further comprise instructions defining the steps of:

dynamically maintaining, during operation of said computer system, match data indicating the number of times each of said rules is matched by a data packet; and
re-ordering said rules based on said match data.

20. The computer readable medium of claim 19 wherein said computer program instructions further comprise instructions defining the steps of:
re-ordering said rules such that rules which are matched more often are applied earlier.

21. The computer readable medium of claim 17 wherein said computer program instructions further comprise instruc-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 676

tions defining the step of:
periodically automatically re-ordering said rules.

22. The computer readable medium of claim 17 wherein said computer program instructions further comprise instructions defining the step of:
automatically re-ordering said rules when the performance of said computer system falls below a threshold.

23. The computer readable medium of claim 17 wherein said computer program instructions further comprise instructions defining the step of:
determining pairs of rules which can be swapped without changing said security policy.

24. The computer readable medium of claim 23 wherein said computer program instructions further comprise instructions defining the step of:
re-ordering said rules by swapping only those rules which may be swapped without changing said security policy.

25. The computer readable medium of claim 23 wherein said computer program instructions further comprise instructions defining the step of:
determining pairs of said rules which can be swapped without changing said security policy by determining the intersection of fields of pairs of rules.

26. A firewall for filtering data traffic between a first network and a second network comprising:

an input port for receiving data packets;
a memory storing a plurality of ordered rules defining a security policy of said firewall;
a filter module for sequentially applying said ordered rules to received data packets; and
a re-order module for automatically re-ordering said plurality of rules.

27. The firewall of claim 26 further comprising:
a conflict module for determining whether pairs of rules can be swapped without altering said security policy.

28. The firewall of claim 27 wherein said conflict module determines whether pairs of rules can be swapped without altering said security policy by determining the intersection of fields of pairs of rules.

29. The firewall of claim 27 wherein said re-order module only swaps pairs of rules which can be swapped without altering said security policy.

30. The firewall of claim 26 further comprising an activity module for dynamically maintaining activity information about the operation of said firewall.

31. The firewall of claim 30 wherein said activity information comprises match data indicating the number of times each of said rules is matched by a data packet.

32. The firewall of claim 30 wherein said re-order module re-orders said plurality of rules based on said activity information.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 677

FIG. 1

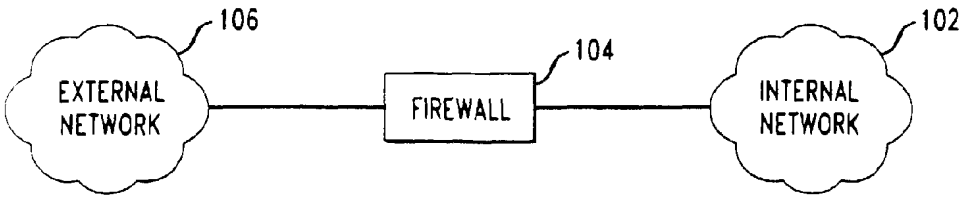

FIG. 2

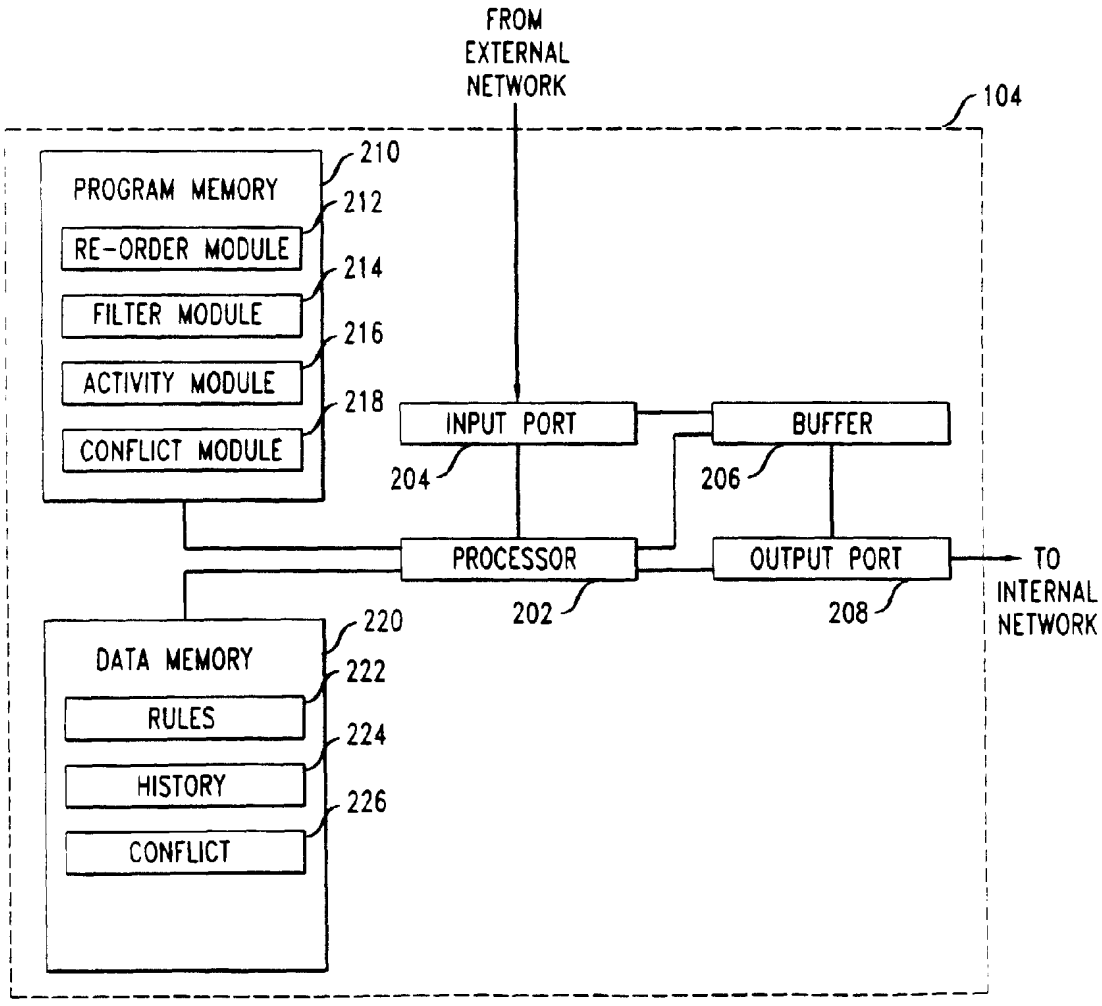

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 678

FIG. 3

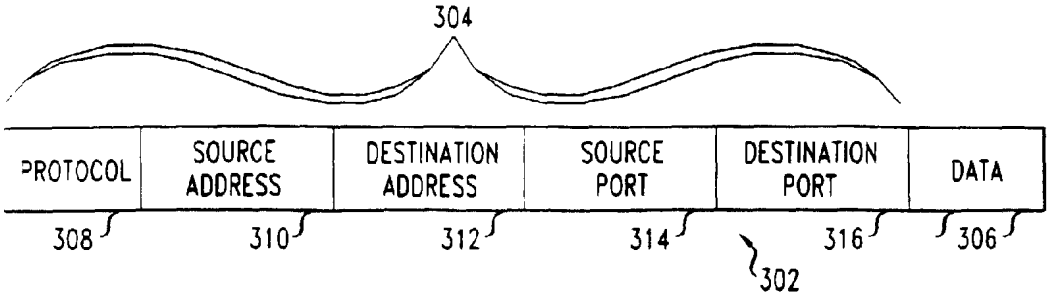

FIG. 4

SEQ. NO.	PROTOCOL	SOURCE ADDRESS	DESTINATION ADDRESS	SOURCE PORT	DESTINATION PORT	ACTION
1	TCP	125.30.20.200	*	*	*	DENY
2	TCP	135.180.104.4-135.180.104.255	135.180.204.4-135.180.204.255	1024-65535	80	ALLOW
3	TCP	125.180.100.*	*	*	*	ALLOW
4	*	*	*	*	*	DENY

FIG. 5

	1	2	3	4	5	
1	NC	C	C	NC	C	
2	C	NC	NC	C	NC	
3	C	NC	NC	NC	C	
4	NC	C	NC	NC	C	
5	C	NC	C	C	NC	

500

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 679

FIG. 6

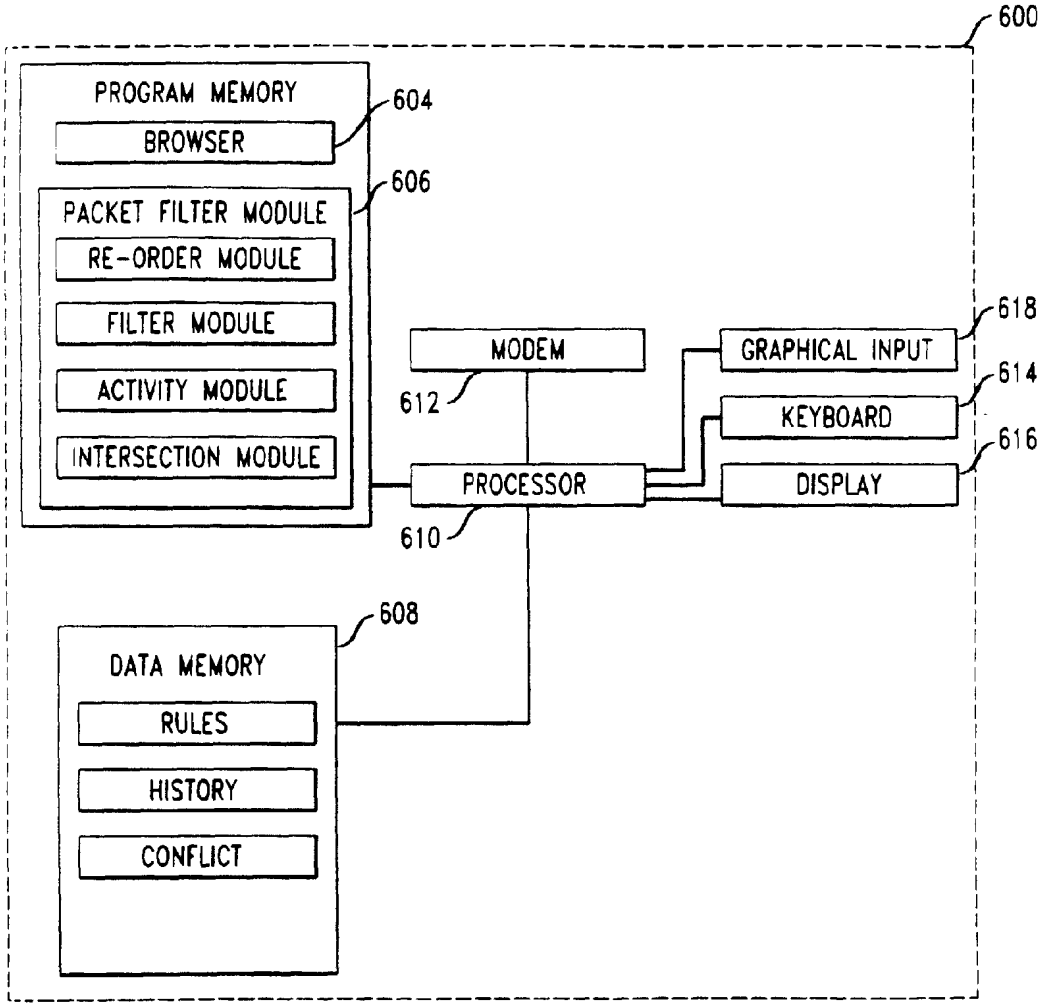

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 680

Espacenet

Bibliographic data: WO2006105093 (A2) — 2006-10-05

METHODS, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR NETWORK FIREWALL POLICY OPTIMIZATION

Inventor(s): FULP ERRIN W [US]; TARSA STEPHEN J [US] ± (FULP, ERRIN, W, ; TARSA, STEPHEN, J)

Applicant(s): UNIV WAKE FOREST [US]; FULP ERRIN W [US]; TARSA STEPHEN J [US] ± (WAKE FOREST UNIVERSITY, ; FULP, ERRIN, W, ; TARSA, STEPHEN, J)

Classification: **- international:** ***G06F15/16***
- cooperative: **H04L63/0263**; **H04L63/20**

Application number: WO2006US11291 20060328

Priority number (s): US20050665664P 20050328

Also published as: WO2006105093 (A3) US2006248580 (A1) US8042167 (B2) EP1864226 (A2) EP1864226 (A4) more

Abstract of WO2006105093 (A2)

Methods, systems, and computer program products for firewall policy optimization are disclosed. According to one method, a firewall policy including an ordered list of firewall rules is defined. For each rule, a probability indicating a likelihood of receiving a packet matching the rule is determined. The rules are sorted in order of non-increasing probability in a manner that preserves the firewall policy.

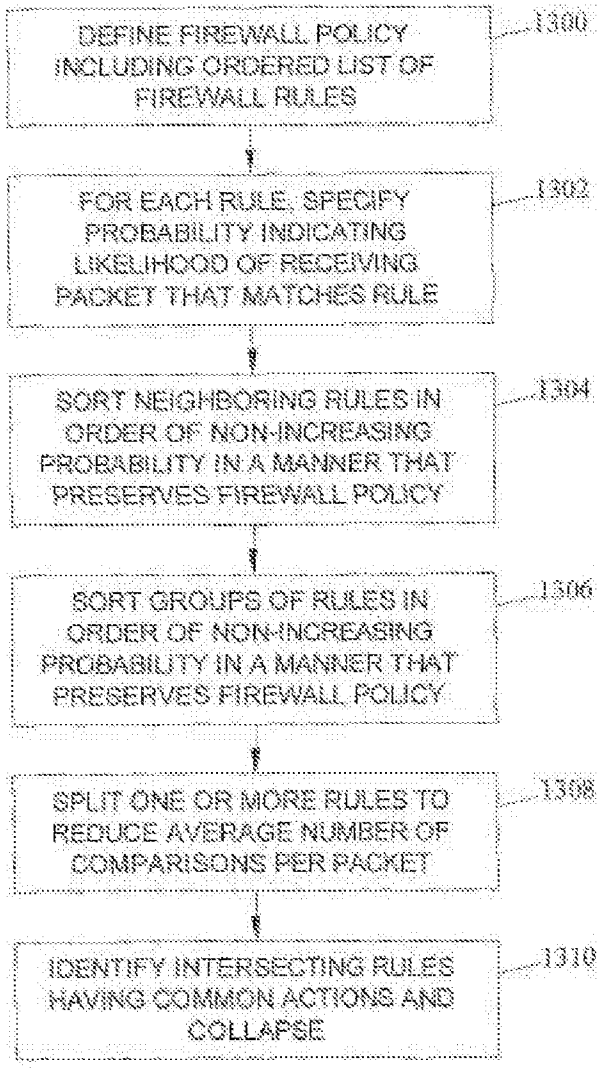
DEFINE FIREWALL POLICY INCLUDING ORDERED LIST OF FIREWALL RULES
1300
FOR EACH RULE, SPECIFY PROBABILITY INDICATING LIKELIHOOD OF RECEIVING PACKET THAT MATCHES RULE
1302
SORT NEIGHBORING RULES IN ORDER OF NON-INCREASING PROBABILITY IN A MANNER THAT PRESERVES FIREWALL POLICY
1304
SORT GROUPS OF RULES IN ORDER OF NON-INCREASING PROBABILITY IN A MANNER THAT PRESERVES FIREWALL POLICY
1306
SPLIT ONE OR MORE RULES TO REDUCE AVERAGE NUMBER OF COMPARISONS PER PACKET
1308
IDENTIFY INTERSECTING RULES HAVING COMMON ACTIONS AND COLLAPSE
1310

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
5 October 2006 (05.10.2006)

PCT

(10) International Publication Number
WO 2006/105093 A2

(51) International Patent Classification:
G06F 15/16 (2006.01)

(21) International Application Number:
PCT/US2006/011291

(22) International Filing Date: 28 March 2006 (28.03.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/665,664 28 March 2005 (28.03.2005) US

(71) Applicant *(for all designated States except US)*: **WAKE FOREST UNIVERSITY** [US/US]; OFFICE OF TECHNOLOGY ASSET MANAGEMENT, WAKE FOREST UNIVERSITY HEALTH SCIECNES, 200 East First Street, Suite 101, Winston Salem, North Carolina 27101 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **FULP, Errin, W.** [US/US]; 308 N. Hawthorne Road, Winston-Salem, North Carolina 27104 (US). **TARSA, Stephen, J.** [US/US]; 135 King Phillips Path, Duxbury, Massachusetts 02332-3503 (US).

(74) Agent: **HUNT, Gregory, A.**; JENKINS, WILSON, TAYLOR & HUNT, P.A., Suite 1200 University Tower, 3100 Tower Boulevard, Durham, North Carolina 27707 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),

[Continued on next page]

(54) Title: METHODS, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR NETWORK FIREWALL POLICY OPTIMIZATION

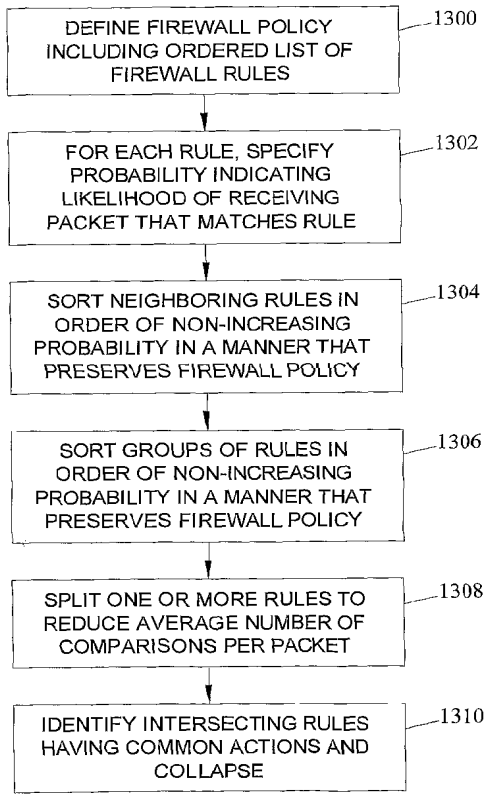

(57) Abstract: Methods, systems, and computer program products for firewall policy optimization are disclosed. According to one method, a firewall policy including an ordered list of firewall rules is defined. For each rule, a probability indicating a likelihood of receiving a packet matching the rule is determined. The rules are sorted in order of non-increasing probability in a manner that preserves the firewall policy.

WO 2006/105093 A2

European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

DESCRIPTION

METHODS, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR NETWORK FIREWALL POLICY OPTIMIZATION

RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Patent Application Serial No. 60/665,664, filed March 28, 2005; the disclosure of which is incorporated herein by reference in its entirety.

GOVERNMENT INTEREST

This invention was made with Government support under Grant No. DE-FG02-03ER25581 awarded by U.S. Department of Energy, Mathematical and Computational Information Sciences Division. The Government has certain rights in the invention.

TECHNICAL FIELD

The subject matter described herein relates to optimizing polices for network firewalls. More particularly, the subject matter described herein relates to methods, systems, and computer program products for optimizing performance of network firewalls.

BACKGROUND ART

Network firewalls remain the forefront defense for most computer systems. Guided by a security policy, these devices provide access control, auditing, and traffic control [3, 30, 31]. A security policy is a list of ordered rules, as seen in Table 1, that defines the action to perform on matching packets. For example, an accept action passes the packet into or from the secure network, while deny causes the packet to be discarded. In many implementations, the rule set is stored internally as a linked list. A packet is sequentially compared to the rules, starting with the first, until a match is found; otherwise, a default action is performed [30, 31]. This is referred to as a first-match policy and is used in many firewall systems including the Linux firewall implementation iptables [25].

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 685

No.	Proto.	Source IP	Source Port	Destination IP	Destination Port	Action	Prob.
1	UDP	1.1.*	*	*	80	deny	0.01
2	TCP	2.*	*	1.*	90	accept	0.02
3	UDP	*	*	1.*	*	accept	0.10
4	TCP	2.*	*	1.*	20	accept	0.17
5	UDP	1.*	*	*	*	accept	0.20
6	*	*	*	*	*	deny	0.50

Table 1: Example Security Policy Consisting of Multiple Ordered Rules

Traditional firewall implementations consist of a single, dedicated machine, similar to a router, that sequentially applies rules to each arriving packet. However, packet filtering represents a significantly higher processing load than routing decisions [24, 29, 31]. For example, a firewall that interconnects two 100 Mbps networks would have to process over 300,000 packets per second [30]. Successfully handling these traffic loads becomes more difficult as rule sets become more complex [4, 22, 31]. Furthermore, firewalls must be capable of processing even more packets as interface speeds increase. In a high-speed environment (e.g. Gigabit Ethernet), a single firewall can easily become a bottleneck and is susceptible to DoS attacks [4, 9, 13, 14]. An attacker could simply inundate the firewall with traffic, delaying or preventing legitimate packets from being processed.

One approach to increase firewall performance focuses on improving hardware design. Current research is investigating different distributed firewall designs to reduce processing delay [4, 9, 22], and possibly provide service differentiation [11]. Another approach focuses on improving performance via better firewall software [6, 7, 12, 16, 17, 24]. Similar to approaches that address the longest matching prefix problem for packet classification [8, 10, 25, 28], solutions typically represent the firewall rule set in different fashions (e.g. tree structures) to improve performance. While both approaches, together or separate, show great promise, each requires radical changes to the firewall system, and therefore are not amenable to current or legacy systems.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 686

Accordingly, there exists a need for improved methods, systems, and computer program products for network firewall policy optimization.

SUMMARY

The subject matter described herein includes methods, systems, and computer program products for network firewall policy optimization. According to one method, a firewall policy including an ordered list of firewall rules is defined. For each rule, a probability indicating a likelihood of receiving a packet matching the rule is defined. Neighboring rules are sorted in order of non-increasing probability in a manner that preserves the firewall policy.

As used herein, the term "firewall" refers to a logical entity that is adapted to filter packets at the ingress and/or egress portions of a network based on a policy. The term "firewall" is intended to include systems that block packets from entering or leaving a network and systems that perform intrusion detection and intrusion protection functions for packets entering or leaving a network. Thus, the methods and systems described herein for firewall policy optimization can be used to optimize policies for intrusion detection and intrusion protection firewalls.

The subject matter described herein includes a method to improve firewall performance and lower packet delay that can be applied to both legacy and current systems. In one implementation, a firewall rule set may be re-ordered to minimize the average number of rule comparisons to determine the action, while maintaining the integrity of the original policy. Integrity is preserved if the reordered and original rules always arrive at the same result. To maintain integrity, a firewall rule set may be modeled as a Directed Acyclical Graph (DAG), where vertices are firewall rules and edges indicate precedence relationships. Given this representation, any linear arrangement of the DAG is proven to maintain the original policy integrity. Unfortunately, determining the optimal rule order from all the possible linear arrangements is proven to be *NP*-complete, since it is equivalent to sequencing jobs with precedence constraints for a single machine [15]. Although determining the optimal order is *NP*-complete, a heuristic will be described below to order firewall rules that reduces the average number of comparisons while maintaining integrity. Simulation

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 687

results show the proposed reordering method yields rule orders that are comparable to optimal; thus, provides a simple means to significantly improve firewall performance and lower packet delay.

The subject matter described herein for network firewall policy optimization may be implemented using any combination of hardware, software, or firmware. For example, the subject matter described herein may be implemented using a computer program product comprising computer executable instructions embodied in a computer readable medium. Exemplary computer readable media suitable for implementing the subject matter described herein include disk memory devices, chip memory devices, programmable logic devices, application specific integrated circuits, and downloadable electrical signals. In addition, a computer program product that implements the subject matter described herein may be implemented using a single device or computing platform or may be distributed across multiple devices or computing platforms.

BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the subject matter described herein will now be explained with reference to the accompanying drawings of which:

Figure 1A is a block diagram of a rule list Directed Acyclical Graph (DAG) for the firewall rules in Table 1;

Figure 1B is a linear arrangement of firewall rules corresponding to the original rule order;

Figure 2A is a policy DAG representing a reordered version of the firewall rule set in Figure 1 after sorting;

Figure 2B is a policy DAG for the rule set illustrated in Table 1 illustrating an optimal rule order;

Figure 3A is a graph illustrating average numbers of packet comparisons versus intersection percentage for different orderings of a firewall policy;

Figure 3B is a graph illustrating the percent difference in average numbers of packet comparisons between sorted and optimal sortings of a firewall policy;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 688

Figure 4A is a graph of percent difference comparisons between original and sorted versions of a firewall policy;

Figure 4B is a graph illustrating the performance impact of considering rule intersection and actions in performing firewall rule sorting;

Figure 5A is a table illustrating a firewall policy before sorting;

Figure 5B is a table illustrating a sorted version of the firewall policy of Figure 5A;

Figure 6A is a policy trie representation of the firewall policy of Figure 5B;

Figure 6B is a table representing the firewall policy of Figure 5A after TCP and UDP sub-trie rotation;

Figure 7A is a policy trie illustrating the splitting of rule r_5 into rules r_{5a} and r_{5b};

Figure 7B is a table illustrating the firewall policy of Figure 5A after splitting rule r_5 into rules r_{5a} and r_{5b};

Figure 8A is a graph illustrating firewall performance with and without rule splitting;

Figure 8B is a graph illustrating locations and probabilities where splitting is better;

Figure 9A is a policy DAG of a firewall rule set;

Figure 9B is a policy DAG illustrating compression of rules r_1 and r_5;

Figure 10A is a policy DAG of a firewall rule list;

Figure 10B is an original policy trie before sorting;

Figure 10C is a policy trie illustrating the firewall rules after rotating sub-tries T_1 and T_2;

Figure 11A is a policy DAG illustrating a firewall policy;

Figure 11B is a policy trie illustrating firewall rules of the policy before sorting;

Figure 11C is a policy trie illustrating the firewall policy after exchanging nodes r_1 and r_2 and nodes r_7 and r_8;

Figure 11D is a policy trie illustrating the firewall policy after exchanging sub-tries T_3 and T_4;

Figure 11E is a policy trie illustrating the firewall policy after exchanging sub-tries T_1 and T_2;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 689

Figure 12A is a policy DAG illustrating an original order for a firewall policy;

Figure 12B is a policy DAG illustrating a sorted version of the firewall policy of Figure 12A after exchanging neighbors;

Figure 12C is a policy DAG illustrating a sorted version of the firewall policy after exchanging lists L_1 including rules r_1 and r_3 and lists L_2 including rules r_2 and r_4;

Figure 13 is a flow chart illustrating exemplary steps for network firewall policy optimization according to an embodiment of the subject matter described herein;

Figure 14 is a block diagram illustrating exemplary components of a system for firewall policy optimization according to an embodiment of the subject matter described herein; and

Figure 15 is a block diagram illustrating a plurality of firewall nodes for implementing a firewall policy according to an embodiment of the subject matter described herein.

DETAILED DESCRIPTION OF THE INVENTION

Modeling Firewall Security Policies

As described above, a firewall rule set, also known as a firewall policy, is traditionally an ordered list of firewall rules. Firewall policy models have been the subject of recent research [1, 2, 16]; however, the primary purpose is anomaly detection and policy verification. In contrast, the policy model described in this herein is designed for firewall performance optimization and integrity. Firewall performance refers to reducing the average number of comparisons required to determine an action, while integrity refers to maintaining the original policy intent. Although improving the worst-case performance is important, it is not possible without changing the list-based representation [16, 24].

Firewall Rule and Policy Models

In the examples described herein, a rule r is modeled as an ordered tuple of sets, $r = (r[1], r[2], ..., r[k])$. Order is necessary among the tuples since

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 690

comparing rules and packets requires the comparison of corresponding tuples. Each tuple $r[l]$ is a set that can be fully specified, specify a range, or contain wildcards '*' in standard prefix format. For the Internet, security rules are commonly represented as a 5-tuple consisting of: protocol type, IP source address, source port number, IP destination address, and destination port number [30, 31]. Given this model, the ordered tuples can be supersets and subsets of each other, which forms the basis of precedence relationships. In addition to the prefixes, each filter rule has an action, which is to accept or deny. However, the action will not be considered when comparing packets and rules. Similar to a rule, a packet (IP datagram) d can be viewed as an ordered k-tuple $d = (d[1], d[2], ..., d[k])$; however, ranges and wildcards are not possible for any packet tuple.

Using the previous rule definition, a standard security policy can be modeled as an ordered set (list) of n rules, denoted as $R = \{r_1, r_2, ..., r_n\}$. A packet d is sequentially compared against each rule r_i starting with the first, until a match is found ($d \Rightarrow r_i$) then the associated action is performed. A match is found between a packet and rule when every tuple of the packet is a subset of the corresponding tuple in the rule.

Definition Packet d matches r_i if

$$d \Rightarrow r_i \quad \text{iff} \quad d[l] \subseteq r_i[l], \quad l = 1, ..., k$$

The rule list R is comprehensive if for every possible legal packet d a match is found using R. Furthermore, two rule lists R and R' are equivalent if for every possible legal packet d the same action is performed by the two rule lists. If R and R' are different (e.g. a reorder) and the lists are equivalent, then the policy integrity is maintained.

As previously mentioned, a rule list has an implied precedence relationship where certain rules must appear before others if the integrity of the policy is to be maintained. For example consider the rule list in Table 1. Rule r_1 must appear before rules r_3 and r_5, likewise rule r_6 must be the last rule in the policy. If for example, rule r_3 was moved to the beginning of the policy, then it will shadow [2] the original rule r_1. However, there is no precedence relationship between rules r_2 and r_4 given in Table 1. Therefore, the relative ordering of these two rules will not impact the policy integrity and can be

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 691

changed to improve performance. The present example assumes that the original policy is free from any anomalies. Likewise, when a policy is reordered to improve performance it should not introduce any anomalies, which will occur if precedence relationships are not maintained. As a result, a model is needed to effectively represent precedence relationships.

Modeling Rule List Precedence Relationships

The precedence relationship between rules in a policy will be modeled as a Directed Acyclical Graph (DAG) [23, 18]. Such graphs have been successfully used to represent the relative order of individual tasks that must take place to complete a job (referred to as a task graph model). Since certain rules must appear before others to maintain policy integrity, this structure is well suited for modeling the precedence of firewall rules. Let $G = (R, E)$ be a rule list DAG for a rule list R, where vertices are rules and edges E are the precedence relationships (constraint). A precedence relationship, or edge, exists between rules r_i and r_j, if $i < j$, the actions for each rule are different, and the rules intersect.

Definition The intersection of rule r_i and r_j, denoted as $r_i \cap r_j$ is

$$r_i \cap r_j = (r_i[l] \cap r_j[l]), \qquad l = 1, \ldots, k$$

Therefore, the intersection of two rules results in an ordered set of tuples that collectively describes the packets that match both rules. The rules r_i and r_j intersect if every tuple of the resulting operation is non-empty. In contrast, the rules r_i and r_j do not intersect, denoted as $r_i \not\cap r_j$, if at least one tuple is the empty set. Note the intersection operation is symmetric; therefore, if r_i intersects r_j, then r_j will intersect r_i. The same is true for rules that do not intersect.

For example consider the rules given in Table 1, the intersection of r_1 and r_3 yields (UDP, 1.1.*, *, 1.*, 80). Again, the rule actions are not considered in the intersection or match operation. Since these two rules intersect, a packet can match both rules for example d = (UDP, 1.1.1.1, 80, 1.1.1.1, 80). Furthermore, the actions of the two rules are different. Therefore, the relative order must be maintained between these two rules and an edge drawn from r_1 to r_3 must be present in the rule list DAG, as seen in Figures 1A and 1B. More

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 692

particularly, Figure 1A illustrates a rule list DAG for the rules in Table 1. Figure 1B illustrates a linear arrangement of the rules in Figure 1A. In Figures 1A and 1B, the vertices represent rules, the circles represents an accept rules, and the squares represents deny rules. Edges that connect the vertices represent precedence requirements. As can be seen from Figure 1B, because of the edge between rules r_1 and r_3, precedence between rules r_1 and r_3 must be maintained. In contrast consider the intersection of rules r_3 and r_5. These two rules intersect, indicating packets belonging to the set (UDP, 1.*, *, 1.*, *) would match both rules. However, it does not matter which of the two rules a packet matches first, since the action is the same for both rules. Therefore, an edge does not exist between rules r_3 and r_5 in the diagram. Similarly, rules r_2 and r_4 do not intersect due to the fifth tuple (destination port). A packet cannot match both rules indicating the relative order can change; therefore, an edge will not exist between them.

The match operation can be used to identify precedence relationships, but it cannot do so in every case. Consider a partial-match example [1], where r_a = (UDP, *, 80, 10.*, 90, accept) and r_b = (UDP, 10.*, 80, *, 90, deny). The intersection of r_a and r_b is (UDP, 10.*, 80, 10.*, 90); therefore a packet, such as d = (UDP, 10.10.10.10, 80, 10.10.10.10, 90), can match both rules. If r_a appears before r_b then the packet d is accepted, but if r_b occurs before r_a then d is rejected. As a result, the order of r_a and r_b in the original policy must be maintained. However, the match operation is unable to identify the precedence in this example. A partial match exists in between rules r_3 and r_5 in Table 1, but as previously discussed an edge does not exist between the rules since the actions are the same.

Using the rule list DAG representation a linear arrangement is sought that improves the firewall performance. As depicted in Figure 1B, a linear arrangement (permutation or topological sort) is a list of DAG vertices where all the successors of a vertex appear in sequence after that vertex [23]. Therefore it follows that a linear arrangement of a rule list DAG represents a rule order, if the vertices are read from left to right. Furthermore, it is proven in the following theorem that any linear arrangement of a rule list DAG maintains integrity.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 693

Theorem Any linear arrangement of a rule list DAG maintains integrity.

Proof Assume a rule list DAG *G* is constructed from the security policy *R* that is free of anomalies. Consider any two rules r_i and r_j in the policy, where $i < j$. If an edge between r_i and r_j in *G* does not exist, then a linear arrangement of *G* can interchange the order of the two rules. An edge will not exist if the rules do not intersect; however, a reorder will not affect integrity since a packet cannot match both rules. Shadowing is not introduced due to the reorder since the intersection operation is symmetric. An edge will not exist if the two rules intersect but have the same action; however, a reorder will not affect integrity since the same action will occur regardless of which rule is matched first. If an edge does exist between the rules, then their relative order will be maintained in every linear arrangement of *G*; thus maintaining precedence and integrity.

Rule List Optimization

As described above, it is important to inspect packets as quickly as possible given increasing network speeds and QoS requirements. Using the rule list DAG to maintain policy integrity, a linear arrangement is sought that minimizes the average number of comparisons required. However, this will require information not present in the firewall rule list. Certain firewall rules have a higher probability of matching a packet than others. As a result, it is possible to develop a policy profile over time that indicates frequency of rule matches (similar to cache hit ratio). Let $P = \{p_1, p_2, ..., p_n\}$ be the policy profile, where p_i is the probability that a packet will match rule *i* (first match in the policy). Furthermore, assume a packet will always find a match, $\sum_{i=1}^{n} p_i = 1$; therefore *R* is comprehensive. Using this information, the average number of rule comparisons required is

$$E[n] = \sum_{i=1}^{n} i \cdot p_i \qquad (1)$$

For example, the average number of comparisons required for the rule set in Table 1 is 5.03.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 694

Given a rule list DAG $G = (R, E)$ and policy profile $P = \{p_1, p_2, ..., p_n\}$ a linear arrangement π of G is sought that minimizes Equation 1. In the absence of precedence relationships, the average number of comparisons is minimized if the rules are sorted in non-increasing order according to the probabilities [26], which is also referred to as Smith's algorithm [27]. Precedence constraints cause the problem to be more realistic; however, such constraints also make determining the optimal permutation more problematic.

Determining the optimal rule list permutation can be viewed as job scheduling for a single machine with precedence constraints [15, 21]. The notation for such scheduling problems is $\alpha|\beta|\gamma|\delta$, where α is the number of machines, β is the precedence (or absence of) which can be represented as a DAG, γ is a restriction on processing time, and δ is the optimality criterion [15]. Determining the optimal rule order is similar to the $1|\beta|1|\sum \omega C_i$ scheduling problem, or optimality criterion, where ω is a weight associated with a job (for example, importance) and C_i is the completion time. As previously noted, the $1|||\sum \omega C_i$ problem can be solved in linear time the using Smith's algorithm [27], which orders jobs according to non-decreasing $\frac{t_i}{\omega}$ ratio, where t_i is the processing time of job i. In this case set $t_i = 1$ and $\omega = p_i \ \forall i$. However, Lawler [19] and Lenstra et al. [21] proved $1|\beta|1|\sum \omega C_i$ to be *NP*-complete via the linear arrangement problem, which implies determining the optimal firewall rule order is also *NP*-complete. Note, determining the number of possible permutations has been proven to be *NP*-hard [5].

Theorem $1|\beta|1|\sum \omega C_i \ \alpha$ *Determining the optimal order of a firewall rule list*

Proof Consider the $1|\beta|1|\sum \omega C_i$ problem. Each of n jobs J_i, $i \in I$, has to be processed without preemption on a single machine that can handle at most one job at a time. For each $i \in I$, let ω be the associated weight. Furthermore, let $G = (V, E)$ be a DAG that represents the precedence order of the jobs J_i. Assume the processing time of each job equals 1 time unit, the weights to be $0 \leq \omega \leq 1$ such that $\sum \omega = 1$, and β, which is G, to be a rule list DAG. In this case, the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 695

optimization criterion $\sum \omega \cdot C_i$ is the same as $\sum p_i \cdot i$, which is given in equation 1. Clearly, the optimal firewall rule ordering problem has a solution if and only if $1|\beta|1|\sum \omega C_i$ has a solution. Therefore, determining the optimal permutation of firewall rules is *NP*-complete.

Exemplary Rule Sorting Algorithm

Although determining the optimal rule permutation was proven to be *NP*-complete, reducing the average number of comparisons required to process a packet remains an important objective. As previously discussed, a sorting algorithm must maintain the precedence relationships among rules. Of course an exhaustive search is possible if the number of rules is small (generate and test all possible topological orders); however as proven in the previous section, this is not feasible with a realistic rule list.

One exemplary algorithm starts with the original the rule set, then sorts neighboring rules based on nonincreasing probabilities. However, an exchange of neighbors should never occur if the rules intersect and have different actions. This test preserves any precedence relationships in the policy. For example, the following sorting algorithm uses such a comparison to determine if neighboring rules should be exchanged. Note a_i denotes the action associated with the i^{th} rule.

```
done = false
while(!done)
    done = true
    for(i = 1; i <n; i++)
        if(p_i < p_i+1 AND (r_i ∩ r_i+1 OR a_i == a_i+1))then
            exchange rules, actions, and probabilities
            done = false
        endif
    endfor
endwhile
```

Other sorting algorithms are possible if the completion time of the sort is an issue; however, the method presented is easy to implement and only requires a simple neighbor comparison. Assume the match probabilities for the rule list given in Table 1. Applying the sorting algorithm to this rule list results in the ordering depicted in Figure 2A, which has 10% fewer comparisons on

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 696

average. However when using the algorithm, it is possible that one rule can prevent another rule from being reordered. For example, rule r_1 prevents rule r_5 from being placed closer to the beginning of the rule set. However, if both rules r_1 and r_5 are placed closer to the beginning of the policy while maintaining their relative order, the average number of comparison will be further reduced, 15% fewer. This is the optimal order for the 6 rule set, which is depicted in Figure 2B. Although this simple sorting algorithm is unable to move groups of rules, it can still improve the performance of the firewall system. Its effectiveness is measured experimentally in the next section.

Experimental Results

In this section, the average number of rule comparisons required after sorting the rules is compared with average number of comparisons required using the original order and the optimal order. Note the optimal rule ordering was determined via an exhaustive search. As a result, it was not feasible to determine the optimal ordering once the number of rules equaled 15 given the number of permutations to consider. Two different variations of the rule sorting method described in the preceding section were implemented. The difference between the methods was the comparison used to determine if neighboring rules should be exchanged. The first sorting algorithm exchanged neighboring rules if they were out of order and did not intersect. The if-condition was as follows.

if ($p_i < p_{i+1}$ **AND** $r_i \cap r_{i+1}$) **then**

Therefore, rule actions were not considered and the method will be referred to as *non-action sort*. The second sorting method did consider the rule action (as described in the preceding section) and will be referred to as *action sort*. The comparison between the two sorting algorithms will indicate the importance of considering rule actions when ordering rules.

In the first experiment, lists of 10 firewall rules were generated with random precedence relationships. The match probability of each rule was given by a Zipf distribution [20], which assigns probabilities according to the rank of an item. For this simulation, the last rule had the highest probability

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 697

which is consistent with most policies (last rules are more general). As a result, the original order yields the worst average number of comparisons. The intersection percentage measured the percentage of rules that intersect in the policy. This metric gives a value for the dependency level in the policy; however note, rule action is not considered when calculating this metric. Rules were equally likely to have an accept or deny action and the results (average number of rule comparisons) for a particular intersection percentage were averaged over 1000 simulations.

Results of the first experiment are given in Figures 3A and 3B. More particularly, Figure 3A is a graph illustrating the average number of packet comparisons versus intersection percentage for the original rule sorting, the non-action sort, the action sort, and the optimal sort. Figure 3B is a graph illustrating the percent difference in the average number of comparisons required between the sorted and optimal firewall policy configurations. The average number of comparisons required was lower for the sorted and the optimal lists when the intersection percentage was low, as seen in Figure 3A. This was expected since there is a large number of possible rule permutations (few edges in the DAG). When the intersection percentage approached 100%, the values converged to the number required for the original list. This is due to the limited number of possible rule orders, only one order in the extreme case. The percent difference between the sorted and optimal order is shown in Figure 3B. At zero intersection percentage the sorted and optimal orders are equal, since any ordering is possible. Similarly at an intersection percentage of 100, the sorted and optimal orders are equal since only one order is possible. Between these two extremes, the action sorting algorithm remains close to the optimal value. In contrast, the non-action sort had a maximum difference of 33%.

The benefit of rule sorting on larger policies is also of interest; however as previously described, it was not possible to determine the optimal ordering once the number of rules approaches 15. Therefore, the second experiment used larger rule sets, but only compared the original order and the sorted orders. The number of rules ranged from 10 to 1000, while the matching probabilities and intersection percentages were the same as the previous

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 698

experiment. The results of this experiment are depicted in Figure 4A. The sorted rule sets always performed equal to or better than the original order, while the action sort consistently performed better than the non-action sort. As noted in the previous experiment, the percent difference is very large given few intersections (e.g. 80% decrease for 1000 rules with 0% dependency), but approaches zero as all the rules intersect. As the number of rules increases, sorting is increasingly beneficial, although only at low intersection percentages. The benefit of sorting drops more significantly as the intersection percentage increases with larger rule sets. This is primarily due to the low matching probabilities of each rule, which requires a complete reordering to have a significant impact on the average number of comparisons. As a result, large rule sets can benefit from sorting if the intersection percentage is low.

The impact of considering rule actions when sorting is illustrated in Figure 4B. In this experiment lists of 1000 rules were generated, where the intersection percentages ranged from 0% to 100% and the percentage of rules with the same action varied from 50% to 100%. The performance of the action sort to the original list was then compared. As the percentage of rules with the same action increased, the percent difference (reduction) in the average number of comparisons also increased. This is due to the increased number of permutations possible when the rule actions are increasingly the same (fewer edges in the rule list DAG to consider). This is depicted in Figure 4B where a policy with a 100% intersection percentage can significantly reduce the number of comparisons (80% reduction) if all the rules have the same action. This performance increase is not possible with the non-action sort. Therefore, considering the rule action increases the number of possible rule orders, thereby providing more possibilities to improve firewall performance.

Conclusions

Network firewalls enforce a security policy by comparing arriving packets to a list of rules. An action, such as accept or deny, is then performed on the packet based on the matching rule. Unfortunately packet filtering can impose significant delays on traffic due to the complexity and size of rule sets.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 699

Therefore, improving firewall performance is important, given network Quality of Service (QoS) requirements and increasing network traffic loads.

The sections above describe rule ordering methods to improve the performance of network firewalls. Assuming each rule has a probability of a packet matching, firewall rules should be sorted such that the matching probabilities are non-increasing. This reduces the average number of comparisons and the delay across the firewall. However, a simple sort is not possible given precedence relations across rules. It is common in a security policy that two rules may match the same packet yet have different actions. It is this precedence relationship between rules that must be maintained to preserve integrity. The method described above uses Directed Acyclical Graphs (DAG) to represent the precedence order rules must maintain. Given this representation, a topological sort can be used to determine the optimal order (minimum average number of comparisons); however, the examples above prove this problem to be *NP*-complete (similar to job scheduling for a single nonpreemptive machine with precedence constraints). As an alternative, a simple sorting method was introduced that maintained the precedence order of the rules. Simulation results indicate this method can significantly reduce the average number of comparisons required and is comparable to optimal ordering.

Several areas exist for future research in optimizing firewall rule lists. The sorting method proposed above is based on a simple algorithm. Although it can offer an improvement over the original rule order, algorithms that can move groups of rules could provide a larger reduction in the average number of comparisons. The effect of stateful firewalls should also be addressed in future research. Security can be enhanced with connection state and packet audit information. For example, a table can be used to record the state of each connection, which is useful for preventing certain types of attacks (e.g., TCP SYN flood) [30, 31]. The impact of such rules on the firewall needs to be investigated and whether sorting can be done on-line to reflect temporal changes. In addition, more research is needed to determine more accurate probability distributions for packet matching and dependency percentages.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 700

Given this information, better algorithms can be designed and more realistic simulations can be performed.

Trie-Based Network Firewall Optimization Methods and Rule Splitting

Overview

As described above, while assuring firewall policy integrity is critical, performance is equally important given increasing network speeds and traffic volumes. Firewall processing delay can be reduced via hardware as well as policy optimization, where rules are manipulated to reduce the number of comparisons. For a given firewall installation, it can be determined that certain security rules have a higher probability of matching a packet than others. As a result, it is possible to develop a *policy profile* over time that indicates frequency of rule matches (similar to cache hit ratio). Given this information, trie-based methods are described herein to minimize the average number of comparisons required per packet. Although most firewall systems still utilize an ordered list representation, the proposed enhancements still are applicable to current firewall systems since a policy trie can be converted into an ordered list using an inorder traversal.

Rule Sorting

As described above, rule sorting of neighboring rules based on ascending probabilities in a manner that considers rules intersection and rule actions can be used to improve firewall performance by reducing the number of packet comparisons. The methods described above use Directed Acyclic Goal Graphs (DAGs) to represent firewalls. The methods described in this section will use the rule sorting algorithms in combination with policy trie representations to perform sorting of groups of rules in a manner that improves firewall performance.

In this example, it is assumed that a security policy contains a list of n ordered rules, $\{r_1, r_2, ..., r_n\}$. In addition, let p_i represent the probability that a packet will match rule i (first match in the policy). Therefore, the policy is comprehensive, which means every packet will have a match, if $\sum_{i=1}^{n} p_i = 1$.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 701

Using this information, the average number of rule comparisons required can be calculated as

$$E[n] = \sum_{i=1}^{n} i \cdot p_i$$

The average number of comparisons is minimized if the rules are sorted in non-ascending order according to the probabilities. However, this can only be achieved if the sorting algorithm considers the probabilities and whether one rule is a superset of another. For example, sorting algorithms must use the following comparison to determine if neighboring rules should be interchanged Note, a_i denotes the action associated with the i^{th} rule.

if ($p_i < p_{i+1}$ **AND** $r_i \not\subset r_{i+1}$ **OR** $a_i == a_{i+1}$)) **then**

Sorting rules in this fashion can have positive impact on the average number of comparisons required. Figure 5A is a table illustrating a firewall policy prior to rule sorting. The expected number of comparisons for the policy of Figure 5A is 4.26 comparisons per packet. Figure 5B is a table illustrating the firewall policy of Figure 5A after sorting using the above-described algorithm. In Figure 5B, the expected number of comparisons per packet has been reduced to 3.94. Thus, using the method described above, the number of comparisons per packet can be reduced by 8%.

However, it can be seen from Figures 5A and 5B that one rule can prevent another rule from being reordered. For example in Figure 5B, rule r_1 prevents rule r_4 from being placed closer to the beginning of the rule set. However, if both rules r_1 and r_4 are placed closer to the beginning of the policy while maintaining their relative order, the average number of comparison will be reduced.

To solve this problem, rule sets can be sorted a using policy tries. First, the rule list is converted into an equivalent policy trie. Each sub-trie will have an associated probability $\overline{p}$ that is average probability (hit-ratio) of the rules comprising the sub-trie. Sub-tries can be rotated around their parent node to increase performance, using the method described earlier. Since the policy trie (or equivalent rule set) is tested from left to right, rotation should occur if the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 702

probability of the right sub-trie is greater than the left sub-trie. As a result, the average number of comparisons required will be reduced and the policy integrity is maintained. Figure 6A is a trie representation of the rules in Figure 5B. Figure 6B is a table representing resulting rule set after rotating the TCP and UDP sub-tries. The TCP sub-trie in Figure 6A has an expected packet matching probability of 0.095 while the UDP sub-trie has an expected packet matching probability of 0.155. Thus, the UDP sub-trie and the TCP sub-trie should be rotated so that the UDP sub-trie becomes the left sub-trie and the TCP sub-trie becomes the right sub-trie. In the table illustrated in Figure 6B, which lists the ordering of the rules after rotating the sub-tries, the average number of comparisons per packets has been reduced to 3.70.

Rule Splitting

Rule splitting takes a general rule and creates more specific rules that collectively perform the same action over the same set of packets. Here, rule splitting is used to reduce the average number of comparisons. For example in Figure 7A, rule r_5 is split into two separate rules, r_{5a} for UDP and r_{5b} for TCP. Figure 7B is a table corresponding to the policy trie of Figure 7A. Once the rules are positioned based on their probabilities and their relation to other rules, the average number of rule comparisons is reduced to 2.98 (after sort, trie rotation, and splitting) which is a 30% less.

It many not be advantageous to split a general rule since it adds another rule to the policy. For example, assume a policy contains 20 rules where the first 19 rules have the same probability. Assume the last rule can be split and the new specific rule has a probability that is $\frac{3}{4}$ of the last rule. The impact of the probability of the last rule and the location of the new split rule, *m*, is depicted in Figures 8A and 8B. More particularly, Figure 8A is a graph illustrating the average number of comparisons for firewalls with and without rule splitting. Figure 8B is a graph illustrating average number of comparisons versus different locations of the new rule and probabilities. The average number of comparisons is reduced as the split rule is located closer to the first rule (surface decreases as *m* approaches one). Furthermore, splitting yields

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 703

better results when the general rule has a high probability. However as illustrated in Figure 8B, the closer the specific rule is to the location of the original rule the average number of comparisons increases, which is the penalty of adding one more rule to the policy.

The effect of splitting a single rule can be described mathematically as follows. Consider n rules where rule r_n can be split into rules $r_{n,1}$ and $r_{n,r}$ (whose union is the original rule). In addition, assume the split rule $r_{n,1}$ will be located at the m^{th} position ($1 \leq m < n$), while rule $r_{n,r}$ will remain at the n^{th} location. The goal is to determine the best location m, which yields a lower average number of comparisons as compared to the original rule set. This can be defined mathematically in the following formula.

$$\sum_{i=1}^{n} i \cdot p_i > \sum_{i=1}^{m-1} i \cdot p_i + m \cdot p_{n,l} + \sum_{i=m}^{n-1} (i+1) \cdot p_i + (n+1) \cdot p_{n,r}$$

The left side of the inequality is the average number of comparisons for the original rule set. The right hand side of the inequality is the average number of comparisons with the specific rule at location m. If it is assumed that the rules located between m and n have an equal probability (denoted as p) the previous equation can be solved for m.

$$m < \frac{n \cdot p - n \cdot p_n + (n+1) \cdot p_{n,r}}{p - p_{n,l}}$$

The new rule must be located between the first and m^{th}; however, its final location will depend on the relationship with the other rules (cannot be placed before any rule for which it is a superset). This result can be applied iteratively to multiple rules and repeatedly to the same rule.

Rule Optimization for QoS

In the preceding sections, optimizing the entire security policy has been discussed. However, it may be desirable to optimize the rules for a certain type of traffic, in order to reduce the average number of comparisons this traffic encounters. This can be done by organizing the policy trie via traffic

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 704

classification, then optimizing the sub-tries. The result is an ordering where the rules for the most important traffic are tested first.

Network Firewall Optimization Using Rule Compression

Another method for network firewall optimization is referred to as rule compression. Given a security policy R, one objective in optimizing firewall performance is to reduce number of rules in the policy while maintaining the original integrity. A method for reducing the number of rules called *rule compression*, which combines and replaces two rules in the policy will now be described. Although compression does reduce the number of rules, it can only be done if the integrity of the original policy is maintained.

Consider two rules r_i and r_j in policy R, where r_i appears before r_j. The two rules can be compressed if and only if the following conditions are true

1. The two rules intersect, $r_i \cap r_j \neq \emptyset$ (an edge between r_i and r_j exists in the policy DAG) (Tuples that are proper subsets and supersets are considered. Compression may still occur if the rules do not intersect or are adjacent in the set space, the resulting rule may contain ranges.)
2. The action of r_i and r_j is the same.
3. Rule r_i is not dependent on any rule r_k that is directly or indirectly dependent on r_j (only one path from r_i to r_j exists in the policy DAG).

The result of compression is a new rule, $r_{i,j}$ whose tuples are the union of the corresponding tuples of rules r_i and r_j. The new rule $r_{i,j}$ replaces both r_i and r_j in R; however, the location of $r_{i,j}$ in R may require the relocation of other rules. Let D_j be the ordered set of rules that appear after r_i that directly or indirectly depend upon r_j. The new rule is placed at the original location of r_i in R and the rules D_j are placed before $r_{i,j}$.

For example, consider the policy given in Table 2. It is possible to compress rules r_1 and r_5, creating the new rule $r_{1,5}$ =[TCP, *, *, 2.*, *, accept]. The ordered set D_j in this example consists of the rules r_3 and r_4. As a result, the new rule $r_{1,5}$ is placed at the original location of r_1 while the rules in D_j are placed before the new rule. Figures 9A and 9B respectively illustrate the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 705

original policy DAG for Table 2 and the policy DAG after rules r_1 and r_5 are compressed.

No.	Proto.	Source IP	Source Port	Destination IP	Destination Port	Action
1	TCP	1.1.*	*	2.2.*	*	accept
2	TCP	1.*	*	*	80	deny
3	TCP	3.3.*	*	*	80	accept
4	TCP	3.*	*	*	*	deny
5	TCP	*	*	2.*	90	accept

Table 2: Example Policy Consisting of Multiple Ordered Rules

Theorem *Compressing rules r_i and r_j of policy R will maintain integrity if the three conditions are met and method described is used.*

Proof Before r_i and r_j are compressed, relocate the rules D_j immediately before r_i and D_i. This relocation will maintain integrity due to the third requirement for compression (no rule in D_i can be dependent on any rule in D_j). Now place r_j directly after r_i. This relocation will not affect integrity since no rule in D_i can have an edge to r_j (again, due to the third condition). Compressing the neighboring rules r_i and r_j creates $r_{i,j}$. This does not affect integrity since the resulting rule only matches the packets that match r_i or r_j in the original policy. Therefore, the result is the compression of rule r_i and r_j , and the integrity of the policy is maintained.

Determining if Compression Should Occur

Although compression may be possible, it may not improve the performance of the policy due to the relocation of rules. Compression should only occur if the average number of comparisons required for the new policy R_c is less than the original policy R; therefore, $E[R_c] < E[R]$.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 706

Firewall Rule Sub-Trie and Sub-List Ordering

Policy Sub-Trie Ordering

Additional methods for firewall optimization are referred to as firewall sub-trie and sub-list ordering. Given a policy trie T, it may be desirable to order all sub-tries such that the average number of tuple comparisons is reduced and policy integrity is maintained. Consider a policy trie T that contains sub-tries T_i and T_j having the same parent node, as seen in Figures 10A-10C. Let P_i denote the sum of the probabilities of the rules contained in sub-trie i, while C_i denotes the number of comparisons required to completely traverse sub-trie i (which is equal to the number of branches). In order to reduce the average number of tuple comparisons, sub-tries that share the same parent node should be ordered such that the following two conditions are observed.

1. Sub-tries that share a parent node are ordered such that the P_i values are non-ascending (higher match probabilities occur first, from left to right).
2. If sub-tries that share a parent node have the same probability (P_i equals P_j), then order the sub-tries such that the C_i values are non-descending (sub-tries consisting of fewer comparisons occur first, from left to right).

These conditions are recursively applied throughout the policy trie as seen in Figures 11A-11D. However, these conditions do not necessarily maintain policy integrity.

Consider a group of n sub-tries that share the same parent node that are numbered sequentially from left to right. Consider any two non-intersecting sub-tries T_i and T_{i+k} in this group that are out of order. If two sub-tries do not intersect, denoted as $T_i \pitchfork T_j$, then no rule in one sub-trie will intersect with any rule in the other sub-trie, $r_k \pitchfork r_l$, $\forall r_k \in T_i$, $\forall r_l \in T_j$. The following *integrity condition* must be observed when reordering the sub-tries. The two sub-tries can be exchanged (rotated about the parent) if and only the sub-tries T_i through T_{i+k-1} do not intersect with T_{i+k} and the sub-tries T_{i+1} through T_{i+k} do not intersect with T_i. Similar to finding the linear sequence of a policy DAGs, these conditions maintain the policy trie integrity. This is observed in Figure 10A,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 707

where no edge exists between any rule in T_1 and T_2. More particularly, Figure 10A is a policy DAG of a six-rule firewall list. Figure 10B is a policy trie representation of the firewall rule list represented by Figure 10A. In Figure 10B, the average number of tuple compares required per packet is 5.5. Figure 10C illustrates the policy trie after rotating sub-tries T_1 and T_2. In Figure 10C, the average number of tuple compares required is 4.5. In Figures 10A-10C, the following probabilities are assumed:

$P_1 = 0.5$

$C_1 = 4$

$P_2 = 0.5$; and

$C_2 = 2$.

The average number of tuple comparisons is an estimate based on the longest trie path.

Using the conditions for comparing and ordering, the following sorting algorithm sorts neighboring sub-tries that share the same parent node.

```
m = a parent node in T
done = false
while(!done)
    done = true
    for each sub-trie Ti having parent m AND a right neighboring sub-trie
        if((Pi < Pi+1 AND (Pi == Pi+1 AND Ci > Ci+1)) AND Ti ⋔ Ti+1) then
        interchange Ti and Ti+1
        done = false
        endif
    endfor
endwhile
```

Figures 11A-11D illustrate policy trie rotation for an eight-rule firewall policy. More particularly, Figure 11A illustrates the policy DAG representation of an eight-rule firewall policy. Figure 11B is a policy trie representation of the original firewall policy. In this example, ordering occurs in three stages, starting at the bottom level and moving towards the top level of the trie. In Figure 11C, the leaves are ordered at the lowest level. More particularly, nodes r_1 and r_2 are exchanged with nodes r_7 and r_8. The average number of tuple compares for the ordering represented by Figure 11C is 10.45. In Figure 11D, the sub-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 708

tries on the second lowest level (T_3, T_4, T_5, and T_6) are ordered. More particularly, Figure 11D illustrates the policy trie after exchanging sub-tries T_3 and T_4. The average number of tuple compares for the representation in Figure 11D is 10.275. In Figure 11E, the sub-tries on the second highest level, (T_1 and T_2) are ordered. Figure 11E illustrates the final ordering after exchanging T_1 and T_2. In the policy trie of Figure 11E, the average number of tuple compares required is 6.775. The average number of tuple comparisons is an estimate based on the longest trie path.

Policy Sub-List Ordering

The preceding section describes conditions for sorting policy sub-tries, which allow the exchange of groups of rules (sub-tries). The same conditions can be applied to list based policies, where sub-lists are ordered to improve the average number of rule comparisons.

In the sections above relating to rule sorting, a method is described to exchange neighboring rules in a list-based security policy, as seen in Figure 12B. However, it was observed that when ordering rules in this fashion it is possible that one rule can block the exchange of another. If groups of rules were allowed to be exchanged the average number of rule comparisons could be further reduced.

Given a policy list L, it may be desirable to order all sub-lists such that the average number of rule comparisons is reduced and policy integrity is maintained. Consider a policy list L that contains sub-tries L_i and L_j. Let P_i denote the sum of the probabilities of the rules contained in sub-list i, while C_i denotes the number of rules in sub-list i. In order to reduce the average number of rule comparisons, sub-lists should be ordered such that the following two conditions are observed.

1. Sub-lists are ordered such that the P_i values are non-ascending (higher match probabilities occur first, from left to right).
2. If sub-lists have the same probability (P_i equals P_j), then order the sub-lists such that the C_i values are non-descending (sub-lists consisting of fewer comparisons occur first, from left to right).

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 709

These conditions are applied throughout the policy list; however as with policy tries, these conditions do not necessarily maintain policy integrity.

Consider a group of n sub-lists that are numbered sequentially from left to right. Consider any two non-intersecting sub-lists L_i and L_{i+k} in this group that are out of order. If two sub-lists do not intersect, denoted as $L_i \pitchfork L_j$, then no rule in one sub-list will intersect with any rule in the other sub-list, $r_k \pitchfork r_l$, $\forall r_k \in L_i$, $\forall r_l \in L_j$. The following *integrity condition* must be observed when reordering the sub-lists. The two sub-lists can be exchanged if and only if the sub-lists L_i through L_{i+k-1} do not intersect with any rule in L_{i+k} and sub-lists L_{i+1} through L_{i+k} do not intersect with L_i. Similar to finding the linear sequence of a policy DAGs, these conditions maintain the policy trie integrity. For example, consider the sub-lists $L_1 = \{r_1, r_3\}$ and $L_2 = \{r_2, r_4\}$ given in Figure 12B. No edge exists between any rule in L_1 and L_2. The two sub-lists are out of order with respect to P and should be exchanged. The resulting list is given in Figure 12C. Figure 12A illustrates the original policy order and results in an average number of rule comparisons of 3.4. In the sorted version in Figure 12B, the average number of rule comparisons is 3.3. In Figure 12C, the average number of rule comparisons has been reduced to 2.7. Using the conditions for comparing and ordering, the following sorting algorithm sorts neighboring sub-lists.

```
done = false
while(!done)
    done = true
    for each sub-list Li having a right neighboring sub-list
        if ((Pi < Pi+1 OR (Pi == Pi+1 AND Ci > Ci+1)) AND Li ⋔ Li+1)then
            interchange rules and probabilities
            done = false
        endif
    endfor
endwhile
```


CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 710

Summary of Firewall Policy Optimization Techniques and Methods

As stated above, the methods and systems described herein for a network firewall policy optimization can be implemented using any combination of hardware, software, and/or firmware. Alternatively, the steps for firewall policy optimization can be performed manually with the number of rules being optimized is small enough. Figure 13 is a flow chart illustrating an exemplary process for firewall policy optimization according to an embodiment of the subject matter described herein. Referring to Figure 13, in block **1300**, a firewall policy including an ordered list of firewall rules is defined. In block **1302**, for each rule, a probability indicating a likelihood of receiving a packet that matches each rule is specified. In block **1304**, neighboring rules are sorted in order of non-increasing probability in a manner that preserves firewall policy. In block **1306**, groups of rules are sorted in order of non-increasing probability in a manner that preserves firewall policy. For example, any of the policy-trie-based methods described above may be used. In block **1308**, one or more rules are split to reduce the number of average packet comparisons. In block **1310**, intersecting rules having common actions are identified and collapsed into single rules.

As described above, the subject matter described herein may be implemented in hardware, software, and/or firmware. In one implementation, the subject matter described herein for firewall policy optimization may be implemented on a general purpose computing platform. Figure 14 is a block diagram of a general purpose computing platform including software for firewall policy optimization according to an embodiment of the subject matter described herein. Referring to Figure 14, computing platform **1400** may be a general purpose computer, such as a personal computer, that includes a microprocessor **1402**, memory **1404**, and I/O interfaces **1406**. Microprocessor **1400** may be any suitable microprocessor, such as any of the Intel or AMD families of microprocessors. Memory **1404** may include volatile memory for running programs and persistent memory, such as one or more disk storage devices. I/O interface **1406** may include interfaces with I/O devices, such as user input devices and output devices.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 711

In the illustrated example, software that may be resident in memory **1404** includes a firewall policy editor **1408** and a firewall policy optimizer **1410**. Firewall policy editor **1408** may allow a user to define a firewall policy. For example, firewall policy editor **1408** may allow a user to define a firewall policy by specifying different rules. The rules may be specified in a graphical manner, for example using policy DAGs as described above. Alternatively, firewall policy editor **1408** may allow a user to input rules in a tabular manner, as illustrated in any of the tables described herein. Firewall policy optimizer **1410** may implement any or all of the firewall policy optimization techniques to order rules entered via firewall policy editor **1408** in a manner that preserves policy integrity and that enhances firewall performance.

A firewall rule set that is optimized using the subject matter described herein may be implemented on any firewall system that includes one or more firewalls. For example, an optimized firewall rule set according to embodiments of the subject matter described herein may be implemented using any of the hierarchical, multi-node firewall systems described in commonly assigned, co-pending U.S. patent application no. 11/316,331, filed December 22, 2005, the disclosure of which is incorporated herein by reference in its entirety. Figure 15 illustrates an example of a multi-node firewall system suitable for implementing the firewall rules according to an embodiment of the subject matter described herein. Referring to Figure 15, a plurality of firewall nodes **1500**, **1502**, **1504**,and **1506** may collectively implement any of the optimized firewall policies described herein. For example, the firewall nodes may collectively implement different portions of a firewall policy data structure including an ordered list of firewall rules. A firewall policy engine **1508** resident on each firewall node may filter packets using its local portion **1510** of the firewall policy data structure. A control node **1512** may include a firewall policy optimizer **1514** that measures the average number of comparisons per packets for the current rule configuration and may dynamically reorder the rules to reduce the average number of comparisons per packet. For example, firewall policy optimizer **1514** may utilize any of the methods described herein to rearrange rules and improve firewall performance. Rules may be tested for rearrangement at user-specified

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 712

intervals or when average number of packet comparisons increases by a user-specified amount.

References

The disclosure of each of the following references is hereby incorporated herein by reference in its entirety.

[1] E. Al-Shaer and H. Hamed. Firewall Policy Management Advisor for Anomaly Detection and Rule Editing. In *Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management*, 2003.

[2] E. Al-Shaer and H. Hamed. Modeling and Management of Firewall Policies. *IEEE Transactions on Network and Service Management*, 1(1), 2004.

[3] S. M. Bellovin and W. Cheswick. Network Firewalls. *IEEE Communications Magazine*, pages 50–57, Sept. 1994.

[4] C. Benecke. A Parallel Packet Screen for High Speed Networks. In *Proceedings of the 15th Annual Computer Security Applications Conference*, 1999.

[5] G. Brightwell and P. Winkler. Counting Linear Extensions is #P-Complete. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*, 1991.

[6] M. Christiansen and E. Fleury. Using Interval Decision Diagrams for Packet Filtering. Technical report, BRICS, 2002.

[7] D. Decasper, Z. Dittia, G. Parulkar, and B. Plattner. Router Plugins: A Software Architecture for Next-Generation Routers. *IEEE/ACM Transactions on Networking*, 8(1), February 2000.

[8] M. Degermark, A. Brodnik, S. Carlsson, and S. Pink. Small Forwarding Tables for Fast Routing Lookups. In *Proceedings of ACM SIGCOMM*, pages 4 – 13, 1997.

[9] U. Ellermann and C. Benecke. Firewalls for ATM Networks. In Proceedings of *INFOSEC'COM*, 1998.

[10] A. Feldmann and S. Muthukrishnan. Tradeoffs for Packet Classification. In Proceedings of the *IEEE INFOCOM*, pages 397 – 413, 2000.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 713

[11] E. W. Fulp. Firewall Architectures for High Speed Networks. Technical Report 20026, Wake Forest University Computer Science Department, 2002.

[12] E. W. Fulp. Firewall Policy Models Using Ordered-Sets. Technical report, Wake Forest University Computer Science Department, 2004.

[13] R. Funke, A. Grote, and H.-U. Heiss. Performance Evaluation of Firewalls in Gigabit-Networks. In *Proceedings of the Symposium on Performance Evaluation of Computer and Telecommunication Systems*, 1999.

[14] S. Goddard, R. Kieckhafer, and Y. Zhang. An Unavailability Analysis of Firewall Sandwich Configurations. In *Proceedings of the 6th IEEE Symposium on High Assurance Systems Engineering*, 2001.

[15] R. L. Graham, E. L. Lawler, J. K. Lenstra, and A. H. G. R. Kan. Optimizing and Approximation in Determinstic Sequencing and Scheduling: A Survey. *Annals of Discrete Mathematics*, 5:287 – 326, 1979.

[16] A. Hari, S. Suri, and G. Parulkar. Detecting and Resolving Packet Filter Conflicts. In *Proceedings of IEEE INFOCOM*, pages 1203–1212, 2000.

[17] HiPAC. High Performance Packet Classification. http://www.hipac.org.

[18] E. Horowitz, S. Sahni, and D. Mehta. *Fundamentals of Data Structures in C++*. Compuer Science Press, 1995.

[19] E. L. Lawler. Sequencing Jobs to Minimize Total Weighted Completion Time Subject to Precedence Constraints. *Annals of Discrete Mathematics*, 2:75 – 90, 1978.

[20] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the Self-Similar Nature of Ethernet Traffic. *IEEE Transactions on Networking*, 2:1 – 15, 1994.

[21] J. K. Lenstra and A. H. G. R. Kan. Complexity of Scheduling under Precedence Constraints. *Operations Research*, 26(1):22 – 35, 1978.

[22] O. Paul and M. Laurent. A Full Bandwidth ATM Firewall. In Proceedings of the 6th European Symposium on *Research in Computer Security ESORICS'2000*, 2000.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 714

[23] B. R. Preiss. *Data Structures and Algorithms with Object-Oriented Design Patterns in C++*. John Wiley & Sons, 1999.

[24] L. Qui, G. Varghese, and S. Suri. Fast Firewall Implementations for Software and Hardware-Based Routers. In *Proceedings of ACM SIGMETRICS*, June 2001.

[25] V. P. Ranganath and D. Andresen. A Set-Based Approach to Packet Classification. In *Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems*, pages 889–894, 2003.

[26] R. Rivest. On Self-Organizing Sequential Search Heuristics. *Communications of the ACM*, 19(2), 1976.

[27] W. E. Smith. Various Optimizers for Single-Stage Production. *Naval Research Logistics Quarterly*, 3:59 – 66, 1956.

[28] V. Srinivasan, G. Varghese, S. Suri, and M. Waldvogel. Fast and Scalable Layer Four Switching. In *Proceedings of ACM SIGCOMM*, pages 191–202, 1998.

[29] S. Suri and G. Varghese. Packet Filtering in High Speed Networks. In *Proceedings of the Symposium on Discrete Algorithms*, pages 969 – 970, 1999.

[30] R. L. Ziegler. *Linux Firewalls*. New Riders, second edition, 2002.

[31] E. D. Zwicky, S. Cooper, and D. B. Chapman. *Building Internet Firewalls*. O'Reilly, 2000.

It will be understood that various details of the invention may be changed without departing from the scope of the invention. Furthermore, the foregoing description is for the purpose of illustration only, and not for the purpose of limitation.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 715

CLAIMS

What is claimed is:

1. A method for producing a performance optimized firewall policy, the method comprising:
   (a) defining a firewall policy including an ordered list of firewall rules;
   (b) for each rule, defining a probability indicating the likelihood of receiving a packet matching each rule; and
   (c) sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy.
2. The method of claim 1 wherein sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy includes, for each pair of neighboring rules, rearranging the rules if a first rule of the pair has a lower probability than a second rule of the pair and the first rule of the pair is not a subset of the second rule of the pair with a different action.
3. The method of claim 1 comprising sorting groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy.
4. The method of claim 3 wherein sorting groups of two or more rules in a manner that preserves firewall policy includes assigning each group of rules to a sub-trie in a policy trie, computing a sum of packet matching probabilities for rules in each sub-trie, and, in response to determining that a first sub-trie has a lower sum than a second sub-trie, rotating the first and second sub-tries about a common parent so that rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy .
5. The method of claim 4 comprising, in response to determining that the first and second sub-tries have equal sums, determining whether the second sub-trie has an lower number of branches than the first sub-trie, and, in response to determining that the second sub-trie has a lower number of branches than the first sub-trie, rotating the first and second sub-tries about a common parent so that the rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 716

6. The method of claim 3 wherein sorting groups of two or more rules in a manner that preserves the firewall policy includes:

(a) identifying first and second sub-lists of rules in the ordered list of firewall rules, the first sub-list preceding the second sub-list in the ordered list of firewall rules;

(b) determining a sum of packet matching probabilities for each of the first and second sub-lists;

(c) determining whether the first and second sub-lists intersect; and

(d) in response to determining that the sum of the packet matching probabilities for the second sub-list is greater than the sum of the packet matching probabilities for the first sub-list and that the first and second sub-lists do not intersect, switching the order of the first and second sub-lists in the ordered list of firewall rules.

7. The method of claim 6 comprising, in response to determining that the sums of the packet matching probabilities for the first and sub-lists are equal, determining whether the second sub-list has a lower number of rules than the first sub-list, and, in response to determining that the second sub-list has a lower number of rules than the first sub-list, switching the order of the first and second sub-lists in the ordered list of firewall rules.

8. The method of claim 1 comprising splitting at least one of the firewall rules to reduce an average number of comparisons required per received packet.

9. The method of claim 1 comprising identifying intersecting rules having a common action and collapsing the rules into a single rule representing a union of the intersecting rules.

10. The method of claim 1 wherein sorting neighboring rules includes sorting rules such that comparisons for a particular class of packets are reduced.

11. The method of claim 1 comprising assigning the rules to vertices in a directed acyclical graph (DAG), determining precedence relationships between the rules, assigning the precedence relationships to edges that connect the vertices in the DAG, and wherein sorting the neighboring

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 717

rules in order of non-increasing probability in a manner that preserves the firewall policy includes sorting the neighboring rules in a manner that preserves the precedence relationships represented by the DAG.

12. The method of claim 1 wherein the rules comprise intrusion detection rules.

13. The method of claim 1 wherein the rules comprise intrusion protection rules.

14. A system for firewall policy optimization, the system comprising:

(a) a firewall policy editor for allowing a user to define a firewall policy including an ordered list of firewall rules and for receiving input regarding probabilities indicating the likelihood of receiving a packet matching each rule; and

(b) a firewall policy optimizer for sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy.

15. The system of claim 14 wherein the firewall policy optimizer is adapted to, for each pair of neighboring rules, rearrange the rules if a first rule of the pair has a lower probability than a second rule of the pair and the first rule of the pair is not a subset of the second rule of the pair with a different action.

16. The system of claim 14 wherein the firewall policy optimizer is adapted to sort groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy.

17. The system of claim 16 wherein the firewall policy optimizer is adapted to assign each group of rules to a sub-trie in a policy trie, to compute a sum of packet matching probabilities for rules in each sub-trie, and, in response to determining that a first sub-trie has a lower sum than a second sub-trie, to rotate the first and second sub-tries about a common parent so that rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy.

18. The system of claim 17 wherein the firewall policy optimizer is adapted to, in response to determining that the first and second sub-tries have equal sums, determine whether the second sub-trie has an lower

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 718

number of branches than the first sub-trie, and, in response to determining that the second sub-trie has a lower number of branches than the first sub-trie, rotate the first and second sub-tries about a common parent so that the rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy.

19. The system of claim 16 wherein the firewall policy optimizer is adapted to:

(a) identify first and second sub-lists of rules in the ordered list of firewall rules, the first sub-list preceding the second sub-list in the ordered list of firewall rules;

(b) determine a sum of packet matching probabilities for each of the first and second sub-lists;

(c) determine whether the first and second sub-lists intersect; and

(d) in response to determining that the sum of the packet matching probabilities for the second sub-list is greater than the sum of the packet matching probabilities for the first sub-list and that the first and second sub-lists do not intersect, switch the order of the first and second sub-lists in the ordered list of firewall rules.

20. The system of claim 19 wherein the firewall policy optimizer is adapted to, in response to determining that the sums of the packet matching probabilities for the first and sub-lists are equal, determine whether the second sub-list has a lower number of rules than the first sub-list, and, in response to determining that the second sub-list has a lower number of rules than the first sub-list, switch the order of the first and second sub-lists in the ordered list of firewall rules.

21. The system of claim 14 wherein the firewall policy optimizer is adapted to split at least one of the firewall rules to reduce an average number of comparisons required per received packet.

22. The system of claim 14 wherein the firewall policy optimizer is adapted to identify intersecting rules having a common action and collapse the rules into a single rule representing a union of the intersecting rules.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 719

23. The system of claim 14 wherein the firewall policy optimizer is adapted to sort neighboring rules such that comparisons for a particular class of packets are reduced.

24. The system of claim 14 wherein the firewall policy optimizer is adapted to assign the rules to vertices in a directed acyclical graph (DAG), determine precedence relationships between the rules, assign the precedence relationships to edges that connect the vertices in the DAG, and to sort the neighboring rules in a manner that preserves the precedence relationships represented by the DAG.

25. The system of claim 14 wherein the rules comprise intrusion detection rules.

26. The system of claim 14 wherein the rules comprise intrusion protection rules.

27. A network firewall comprising:

(a) a firewall policy data structure for storing an ordered list of firewall rules;

(b) a firewall policy engine for implementing a firewall policy by comparing packets to the rules in the order specified by the firewall policy data structure; and

(c) a firewall policy optimizer for optimizing performance of the network firewall by reordering the rules in a manner that reduces an average number of rule comparisons per packet and hat preserves the firewall policy.

28. The network firewall of claim 27 wherein the firewall policy optimizer is adapted to measure average number of comparisons per packet for a rule ordering specified by the firewall policy data structure and to dynamically re-order to rules to reduce the average number of comparisons per packet.

29. A computer program product comprising computer executable instructions embodied in a computer readable medium for performing steps comprising:

(a) defining a firewall policy including an ordered list of firewall rules;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 720

(b) for each rule, defining a probability indicating the likelihood of receiving a packet matching each rule; and

(c) sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy.

30. The computer program product of claim 29 wherein sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy includes, for each pair of neighboring rules, rearranging the rules if a first rule of the pair has a lower probability than a second rule of the pair and the first rule of the pair is not a subset of the second rule of the pair with a different action.

31. The computer program product of claim 29 comprising sorting groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy.

32. The computer program product of claim 31 wherein sorting groups of two or more rules in a manner that preserves firewall policy includes assigning each group of rules to a sub-trie in a policy trie, computing a sum of packet matching probabilities for rules in each sub-trie, and, in response to determining that a first sub-trie has a lower sum than a second sub-trie, rotating the first and second sub-tries about a common parent so that rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy .

33. The computer program product of claim 32 comprising, in response to determining that the first and second sub-tries have equal sums, determining whether the second sub-trie has an lower number of branches than the first sub-trie, and, in response to determining that the second sub-trie has a lower number of branches than the first sub-trie, rotating the first and second sub-tries about a common parent so that the rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy.

34. The computer program product of claim 31 wherein sorting groups of two or more rules in a manner that preserves the firewall policy includes:

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 721

(a) identifying first and second sub-lists of rules in the ordered list of firewall rules, the first sub-list preceding the second sub-list in the ordered list of firewall rules;

(b) determining a sum of packet matching probabilities for each of the first and second sub-lists;

(c) determining whether the first and second sub-lists intersect; and

(d) in response to determining that the sum of the packet matching probabilities for the second sub-list is greater than the sum of the packet matching probabilities for the first sub-list and that the first and second sub-lists do not intersect, switching the order of the first and second sub-lists in the ordered list of firewall rules.

35. The computer program product of claim 34 comprising, in response to determining that the sums of the packet matching probabilities for the first and sub-lists are equal, determining whether the second sub-list has a lower number of rules than the first sub-list, and, in response to determining that the second sub-list has a lower number of rules than the first sub-list, switching the order of the first and second sub-lists in the ordered list of firewall rules.

36. The computer program product of claim 29 comprising splitting at least one of the firewall rules to reduce an average number of comparisons required per received packet.

37. The computer program product of claim 29 comprising identifying intersecting rules having a common action and collapsing the rules into a single rule representing a union of the intersecting rules.

38. The computer program product of claim 29 wherein sorting neighboring rules includes sorting rules such that comparisons for a particular class of packets are reduced.

39. The computer program product of claim 29 comprising assigning the rules to vertices in a directed acyclical graph (DAG), determining precedence relationships between the rules, assigning the precedence relationships to edges that connect the vertices in the DAG, and wherein sorting the neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy includes sorting the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 722

neighboring rules in a manner that preserves the precedence relationships represented by the DAG.

40. The computer program product of claim 29 wherein the rules comprise intrusion detection rules.

41. The computer program product of claim 29 wherein the rules comprise intrusion protection rules.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 723

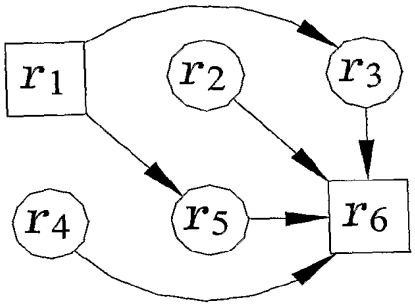

FIG. 1A

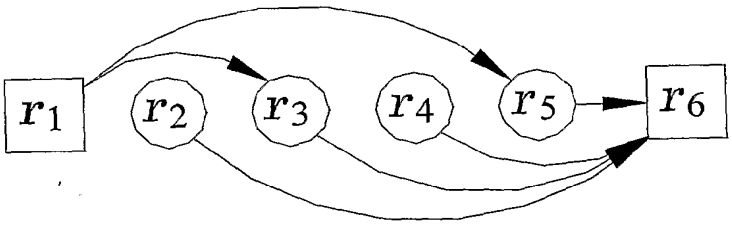

FIG. 1B

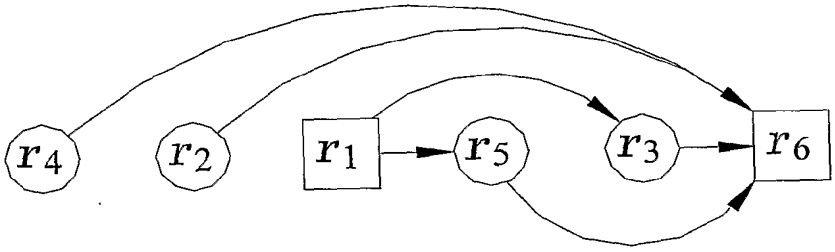

FIG. 2A

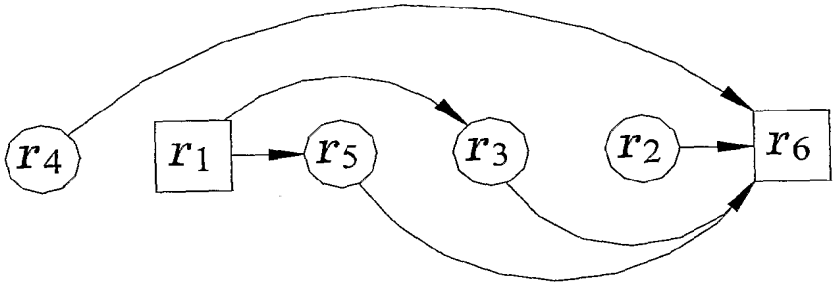

FIG. 2B

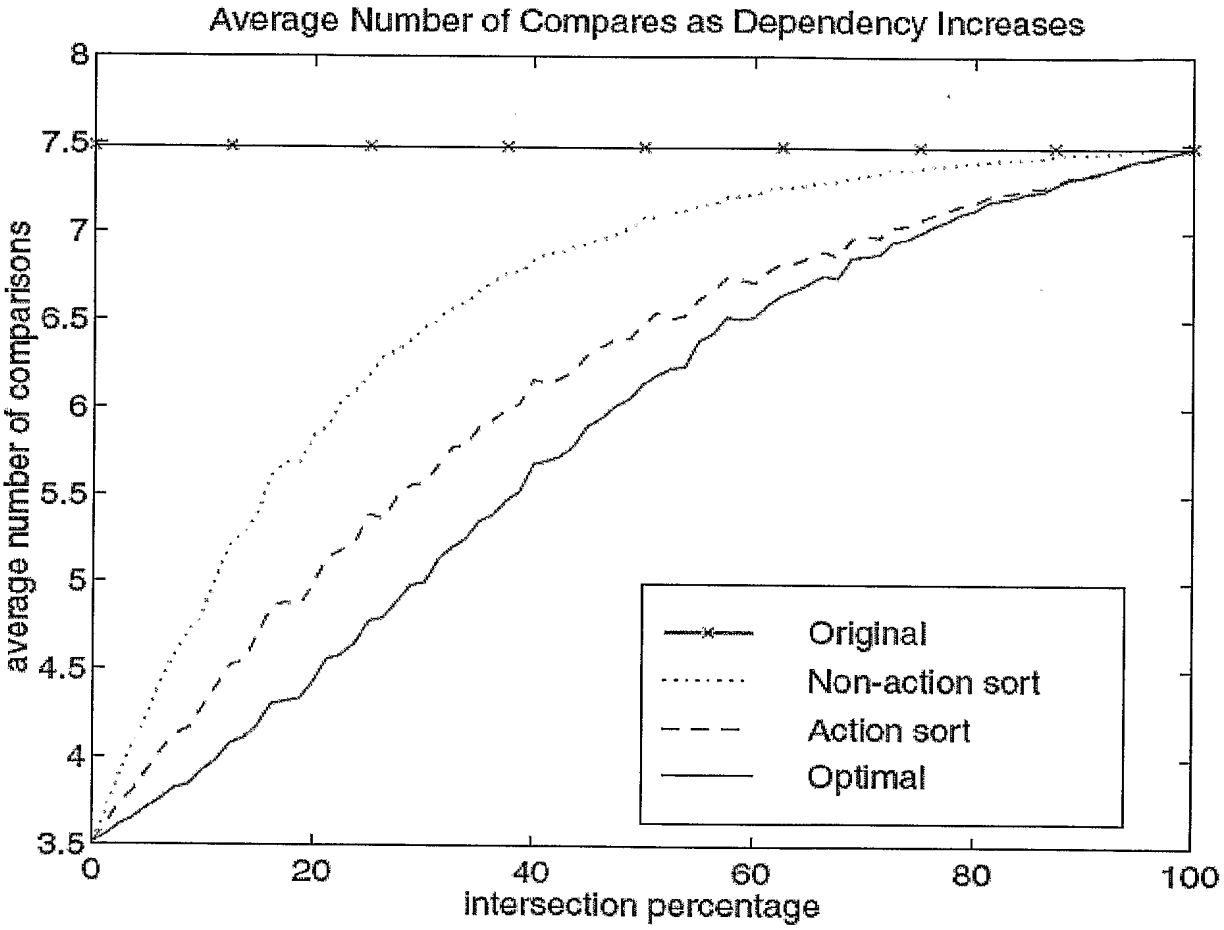
Average Number of Compares as Dependency Increases
average number of comparisons
intersection percentage
8
7.5
7
6.5
6
5.5
5
4.5
4
3.5
0
20
40
60
80
100
Original
Non-action sort
Action sort
Optimal

FIG. 3A

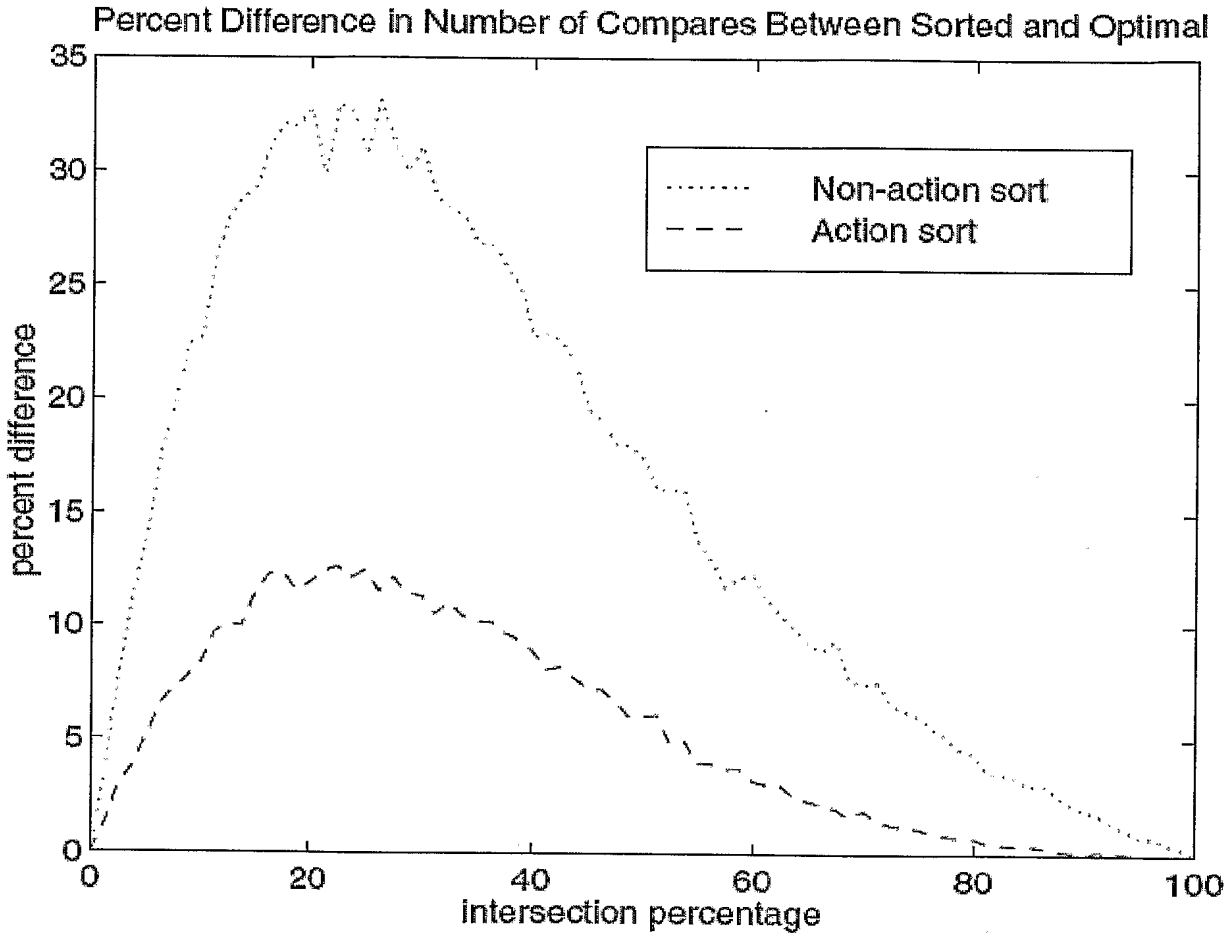
Percent Difference in Number of Compares Between Sorted and Optimal
Non-action sort
Action sort
percent difference
intersection percentage
0
5
10
15
20
25
30
35
0
20
40
60
80
100

FIG. 3B

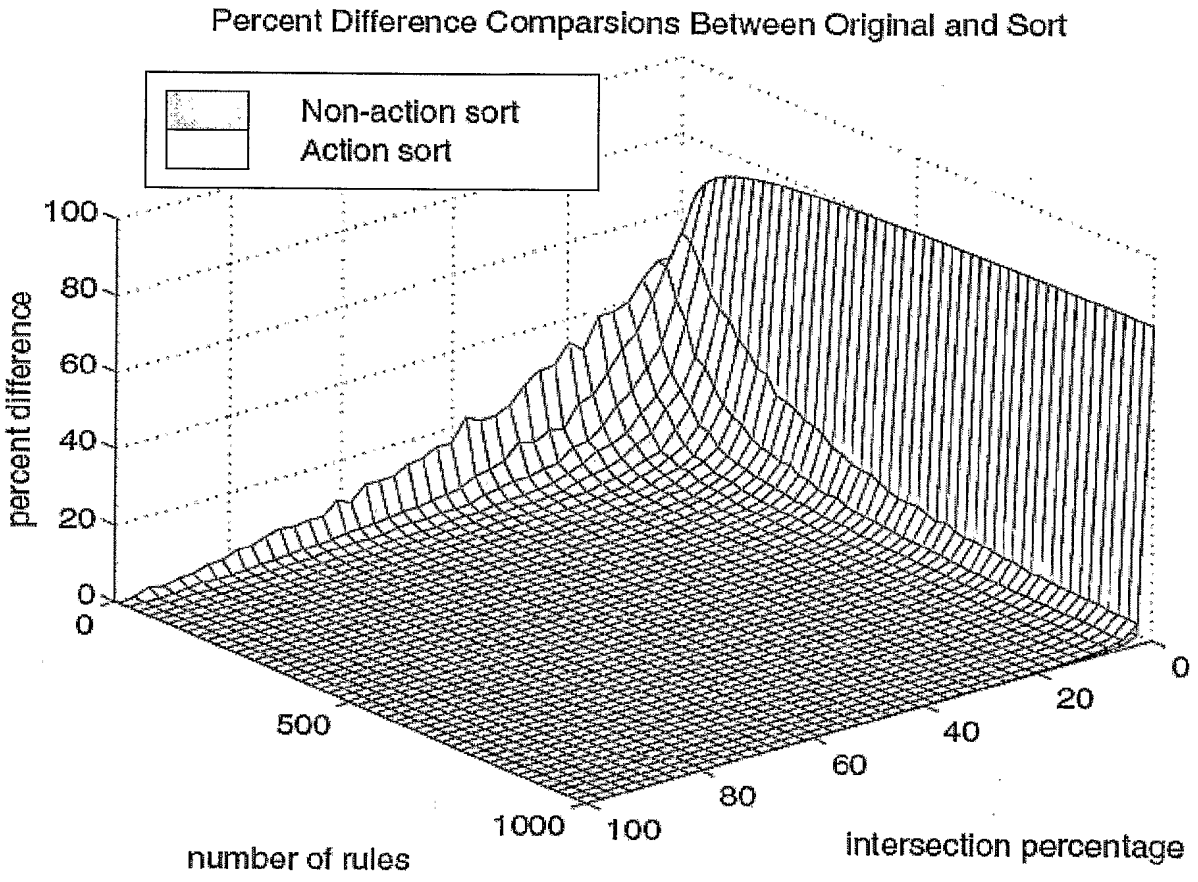
Percent Difference Comparsions Between Original and Sort
Non-action sort
Action sort
percent difference
100
80
60
40
20
0
0
500
1000
number of rules
0
20
40
60
80
100
intersection percentage

FIG. 4A

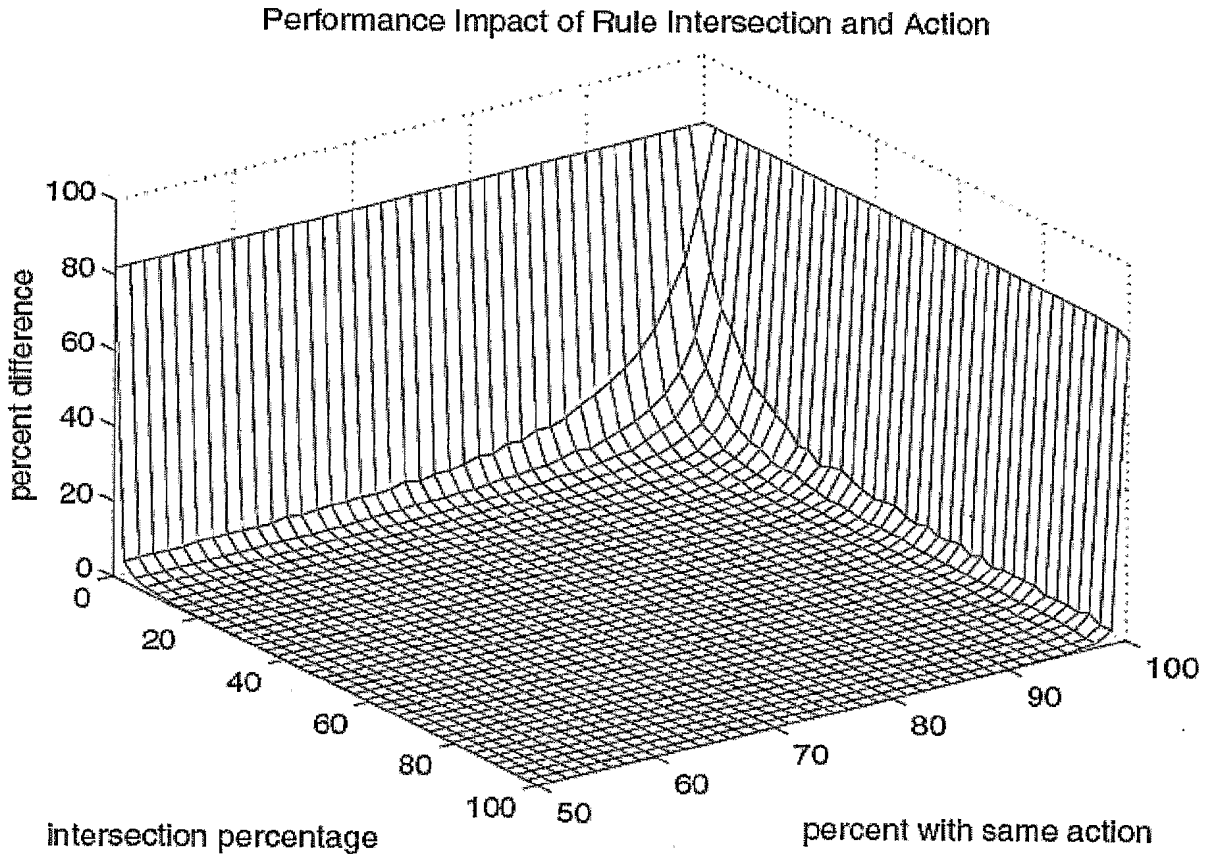
Performance Impact of Rule Intersection and Action
percent difference
100
80
60
40
20
0
0
20
40
60
80
100
intersection percentage
50
60
70
80
90
100
percent with same action

FIG. 4B

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 5A

r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 5B

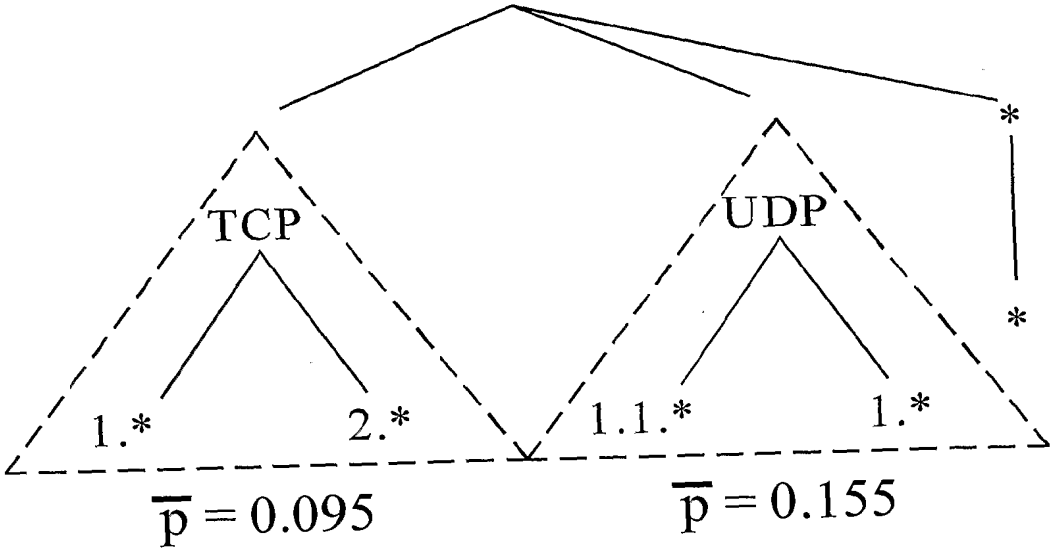

FIG. 6A

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 6B

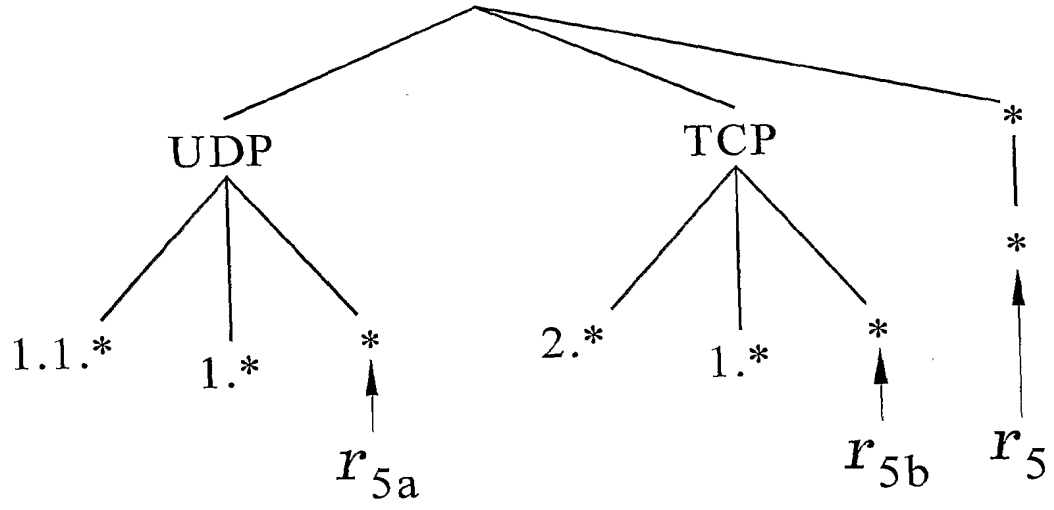

FIG. 7A

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	ACCEPT)	$p_4 = 0.30$
r_{5a}	(UDP,	*,	DENY)	$p_{5a} = 0.47$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_{5b}	(TCP,	*,	DENY)	$p_{5b} = 0.03$

FIG. 7B

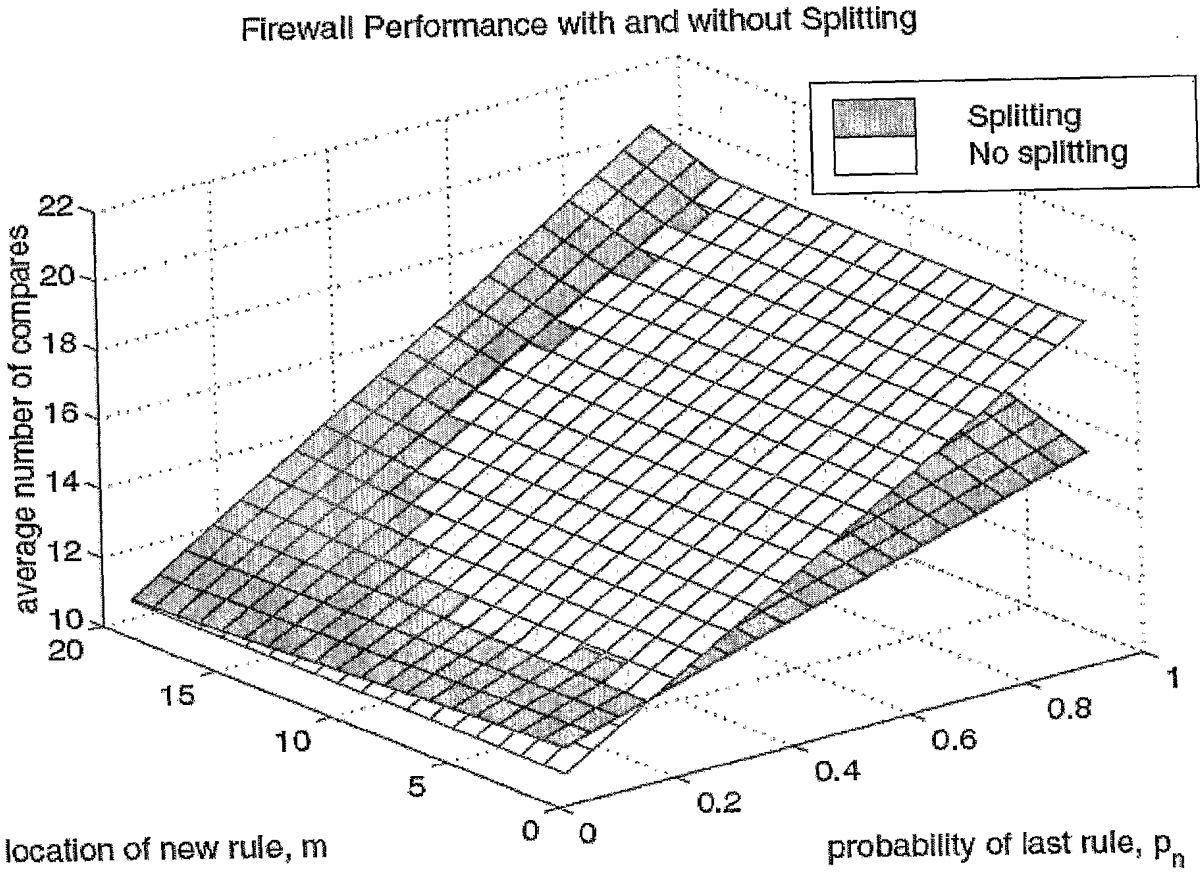
Firewall Performance with and without Splitting
Splitting
No splitting
average number of compares
22
20
18
16
14
12
10
20
15
10
5
0
0
0.2
0.4
0.6
0.8
1
location of new rule, m
probability of last rule, p_n

FIG. 8A

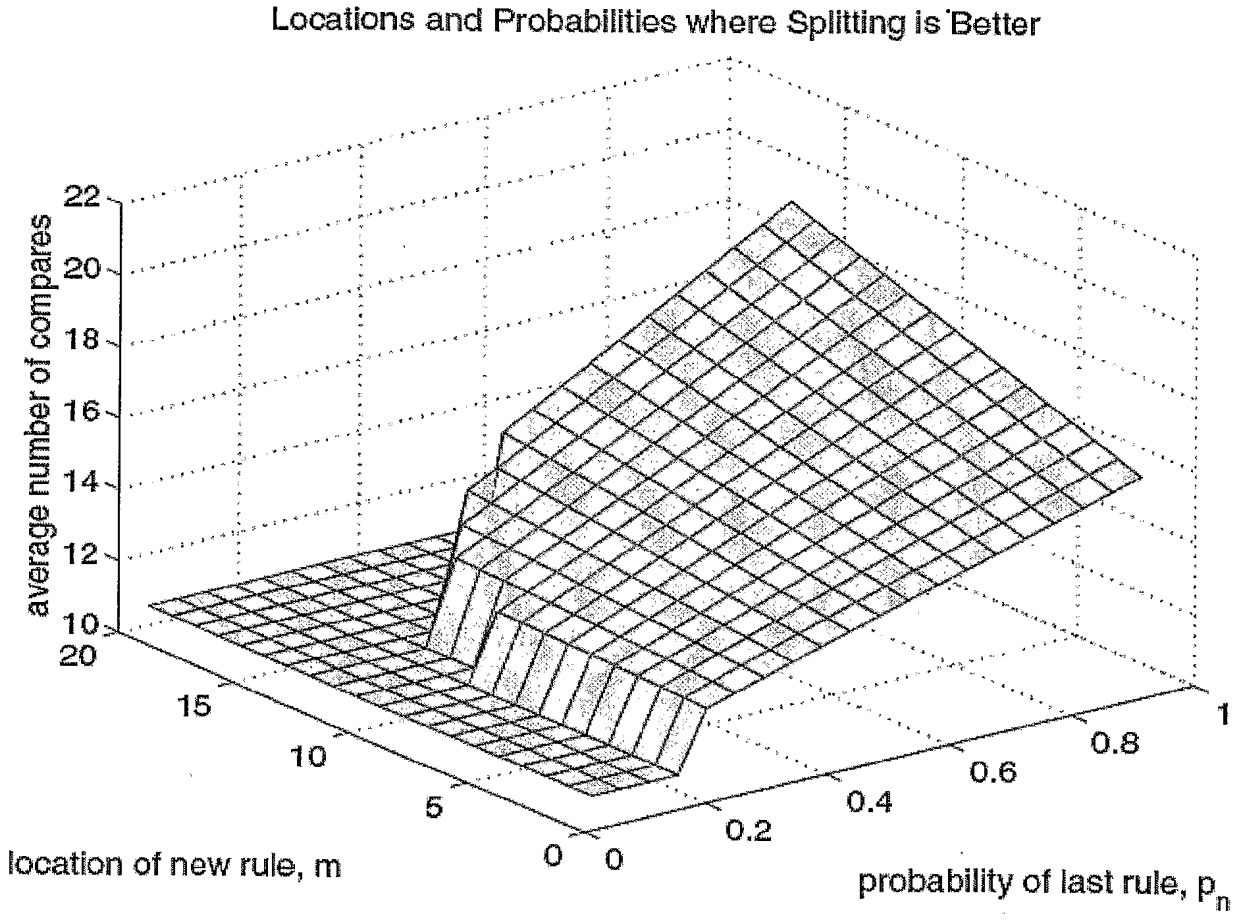
Locations and Probabilities where Splitting is Better
average number of compares
22
20
18
16
14
12
10
20
15
10
5
0
location of new rule, m
0
0.2
0.4
0.6
0.8
1
probability of last rule, p_n

FIG. 8B

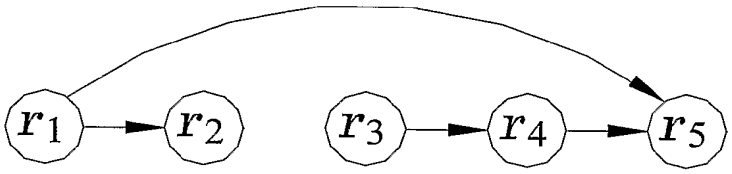

FIG. 9A

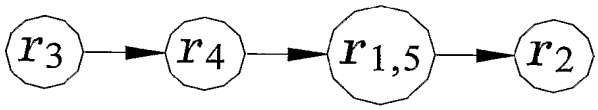

FIG. 9B

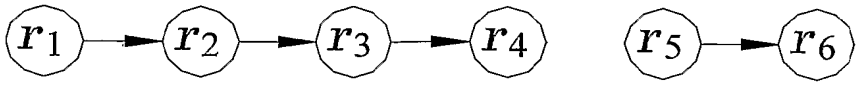

FIG. 10A

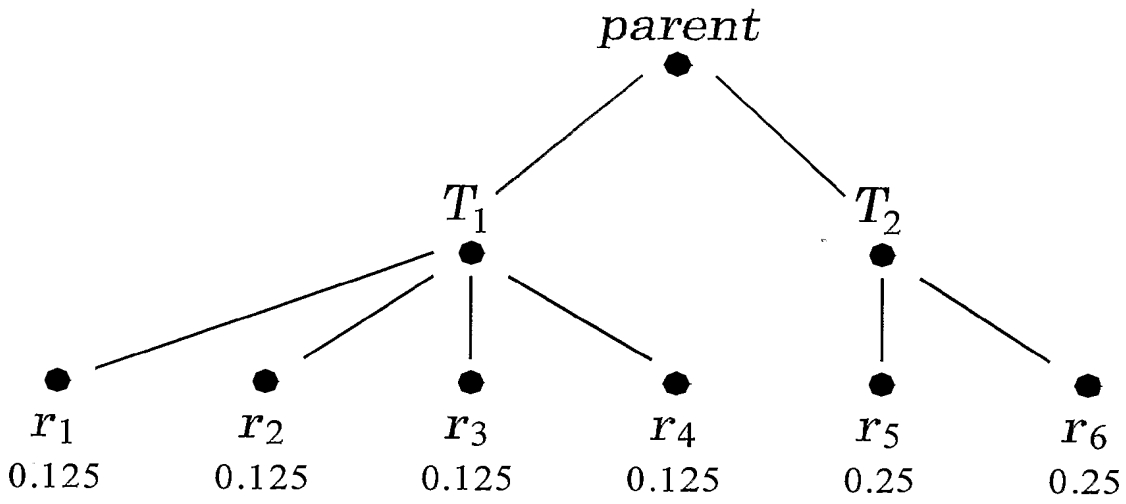

FIG. 10B

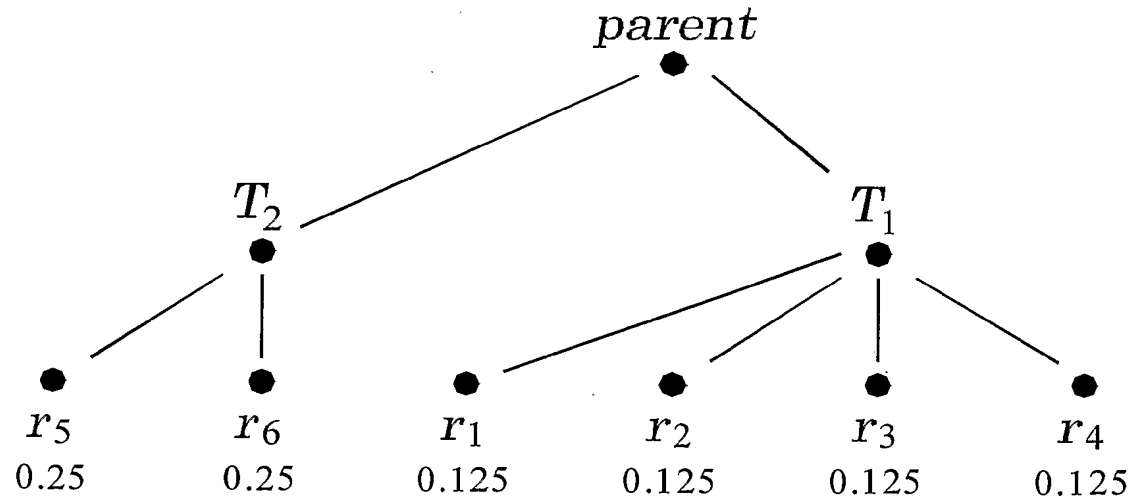

FIG. 10C

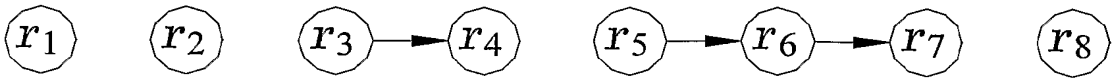

FIG. 11A

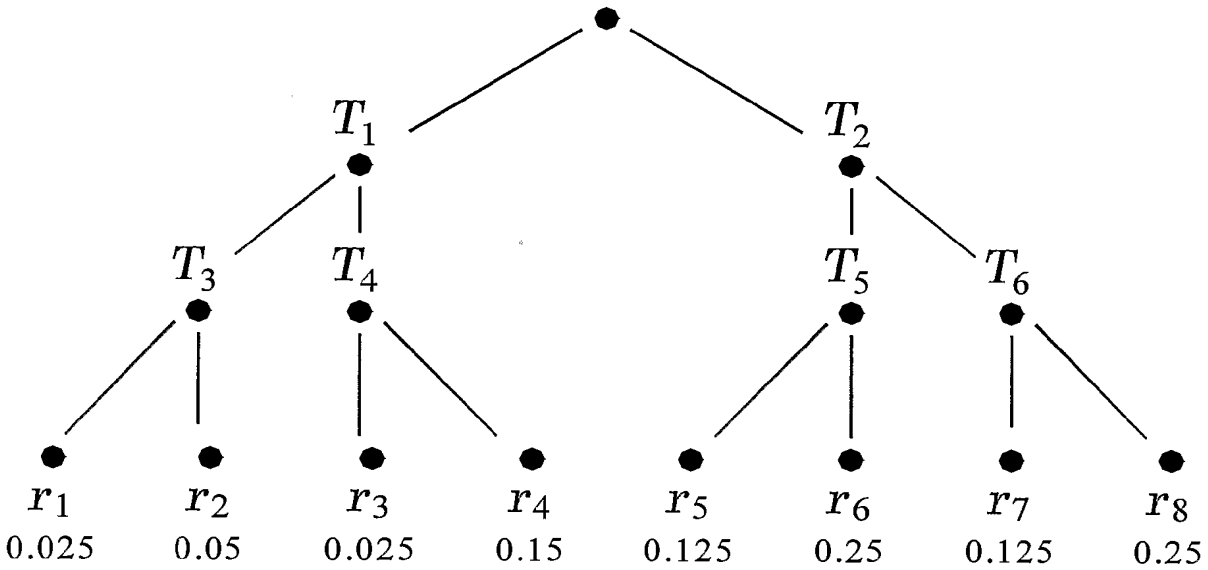

FIG. 11B

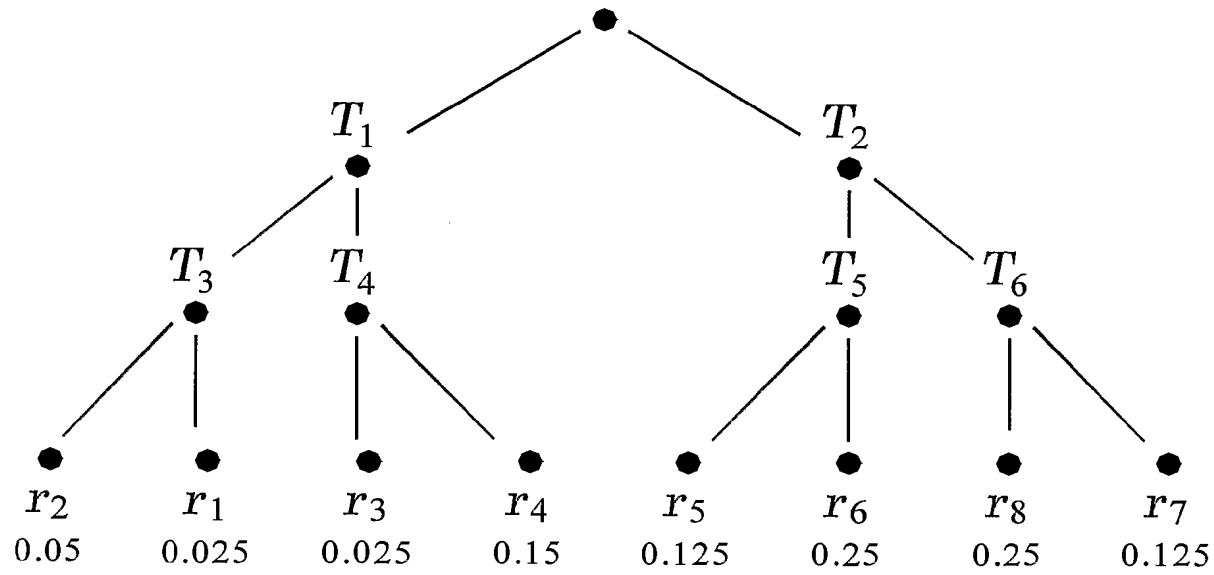

FIG. 11C

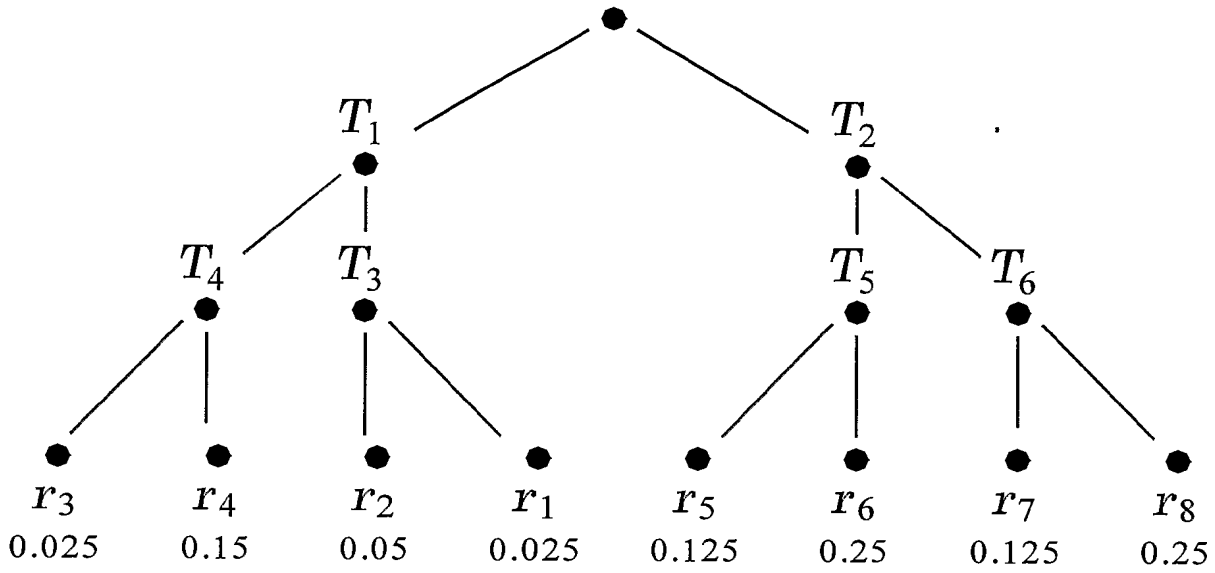

FIG. 11D

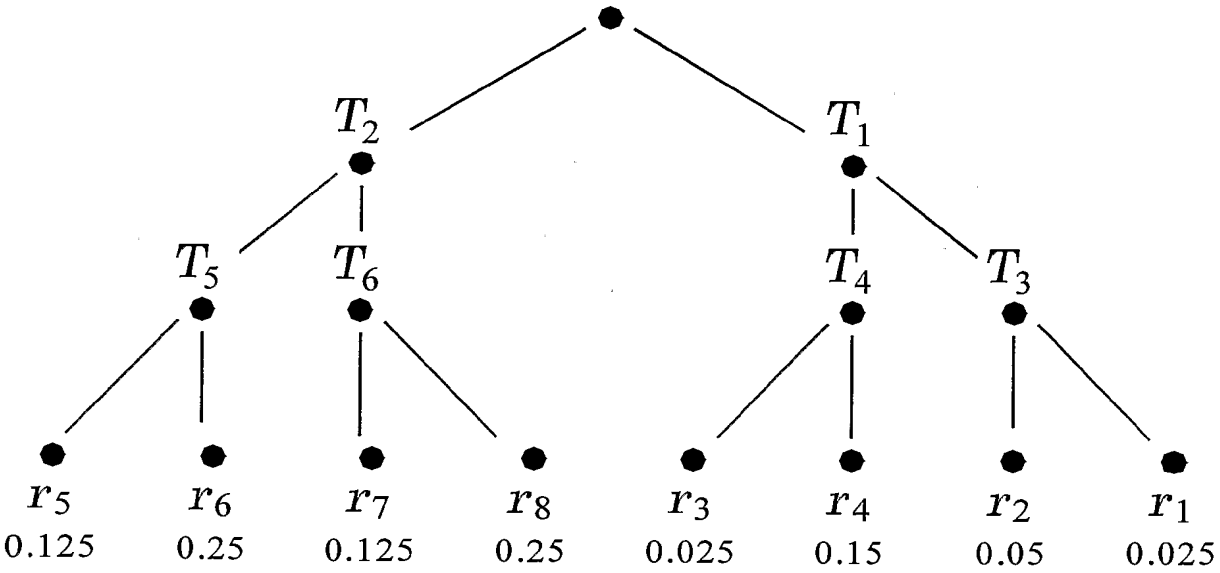

FIG. 11E

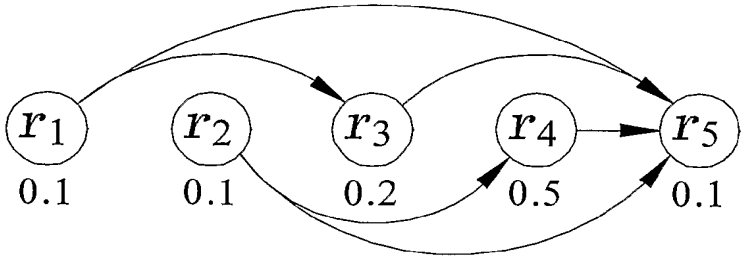

FIG. 12A

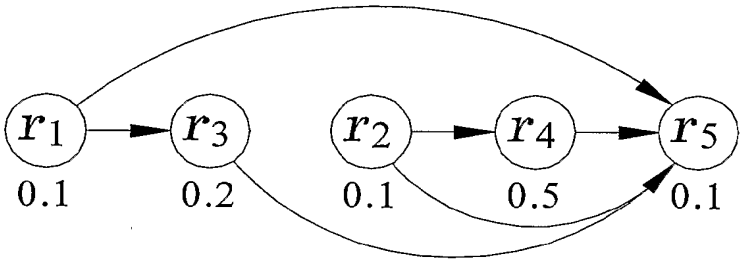

FIG. 12B

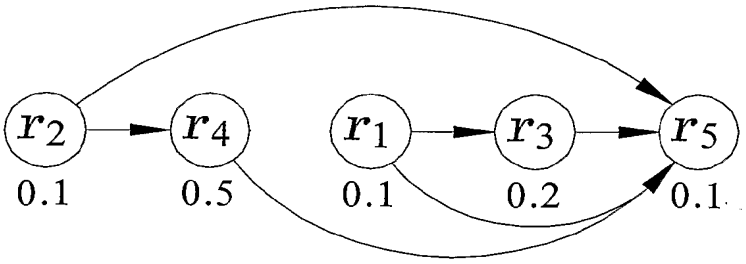

FIG. 12C

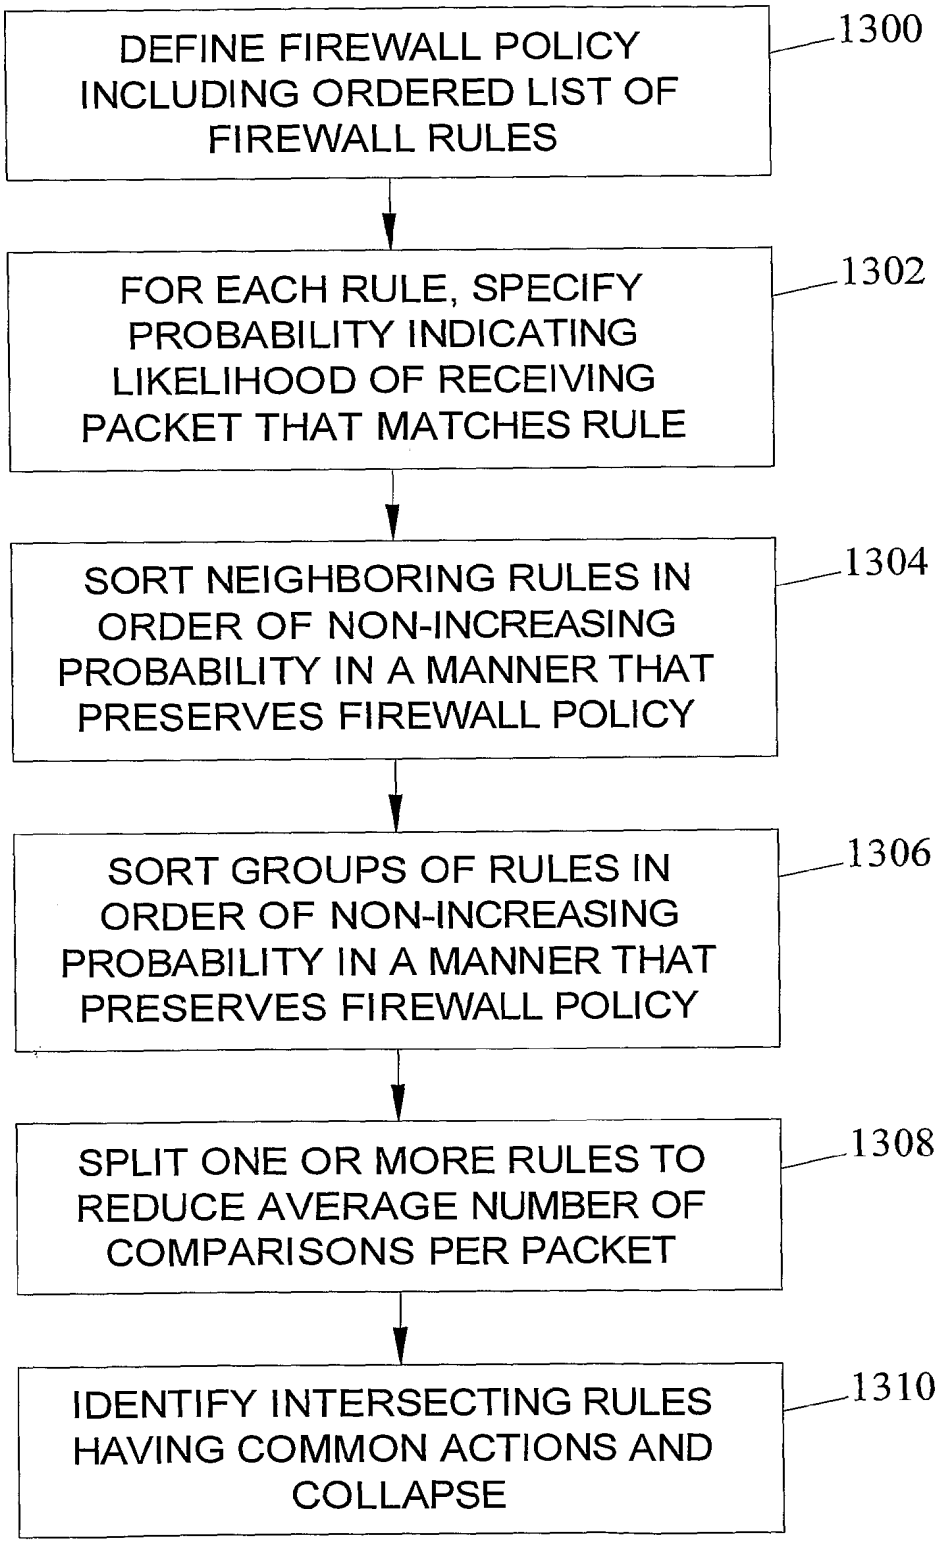
DEFINE FIREWALL POLICY INCLUDING ORDERED LIST OF FIREWALL RULES
1300
FOR EACH RULE, SPECIFY PROBABILITY INDICATING LIKELIHOOD OF RECEIVING PACKET THAT MATCHES RULE
1302
SORT NEIGHBORING RULES IN ORDER OF NON-INCREASING PROBABILITY IN A MANNER THAT PRESERVES FIREWALL POLICY
1304
SORT GROUPS OF RULES IN ORDER OF NON-INCREASING PROBABILITY IN A MANNER THAT PRESERVES FIREWALL POLICY
1306
SPLIT ONE OR MORE RULES TO REDUCE AVERAGE NUMBER OF COMPARISONS PER PACKET
1308
IDENTIFY INTERSECTING RULES HAVING COMMON ACTIONS AND COLLAPSE
1310

FIG. 13

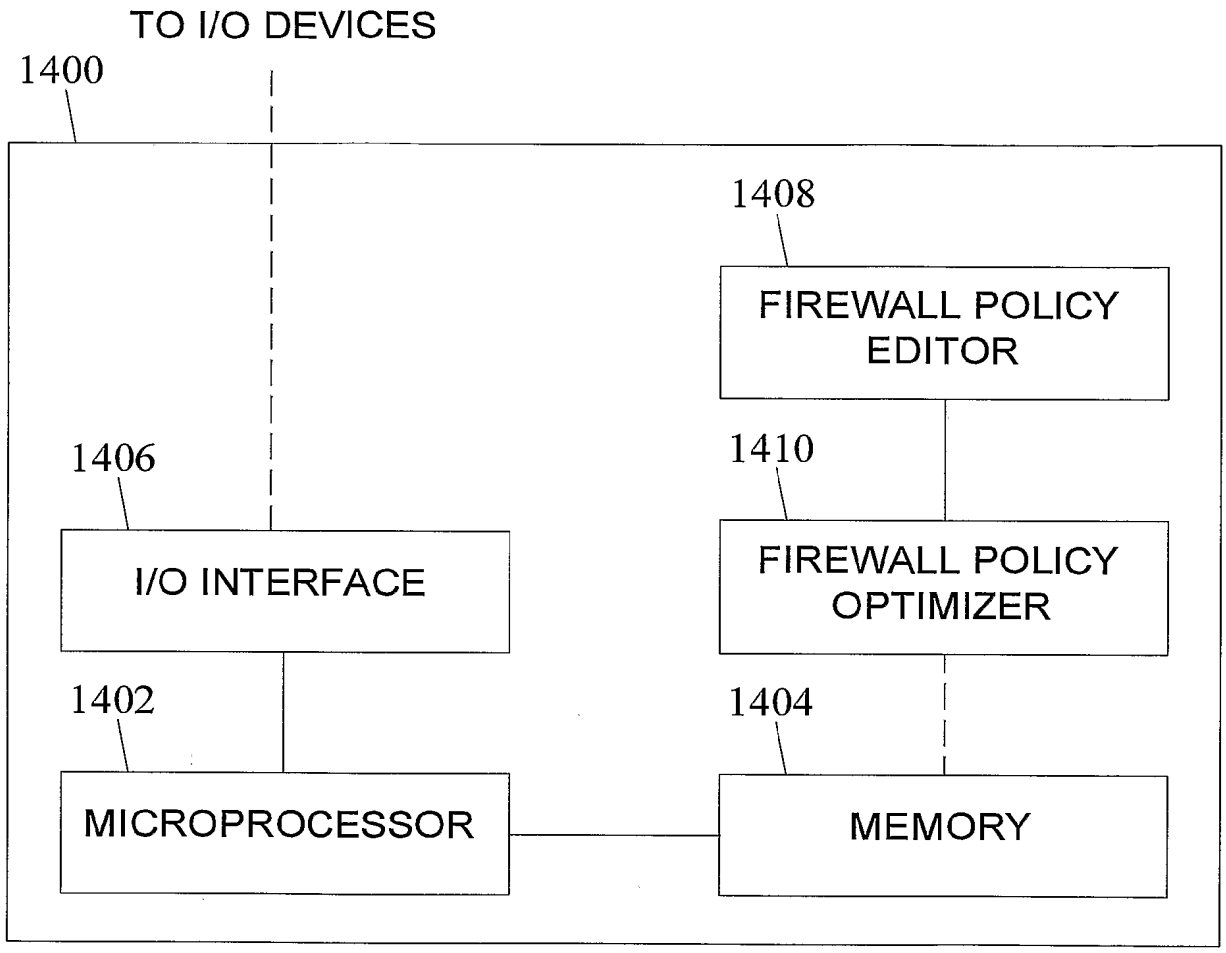
TO I/O DEVICES
1400
1408
FIREWALL POLICY EDITOR
1406
1410
I/O INTERFACE
FIREWALL POLICY OPTIMIZER
1402
1404
MICROPROCESSOR
MEMORY

FIG. 14

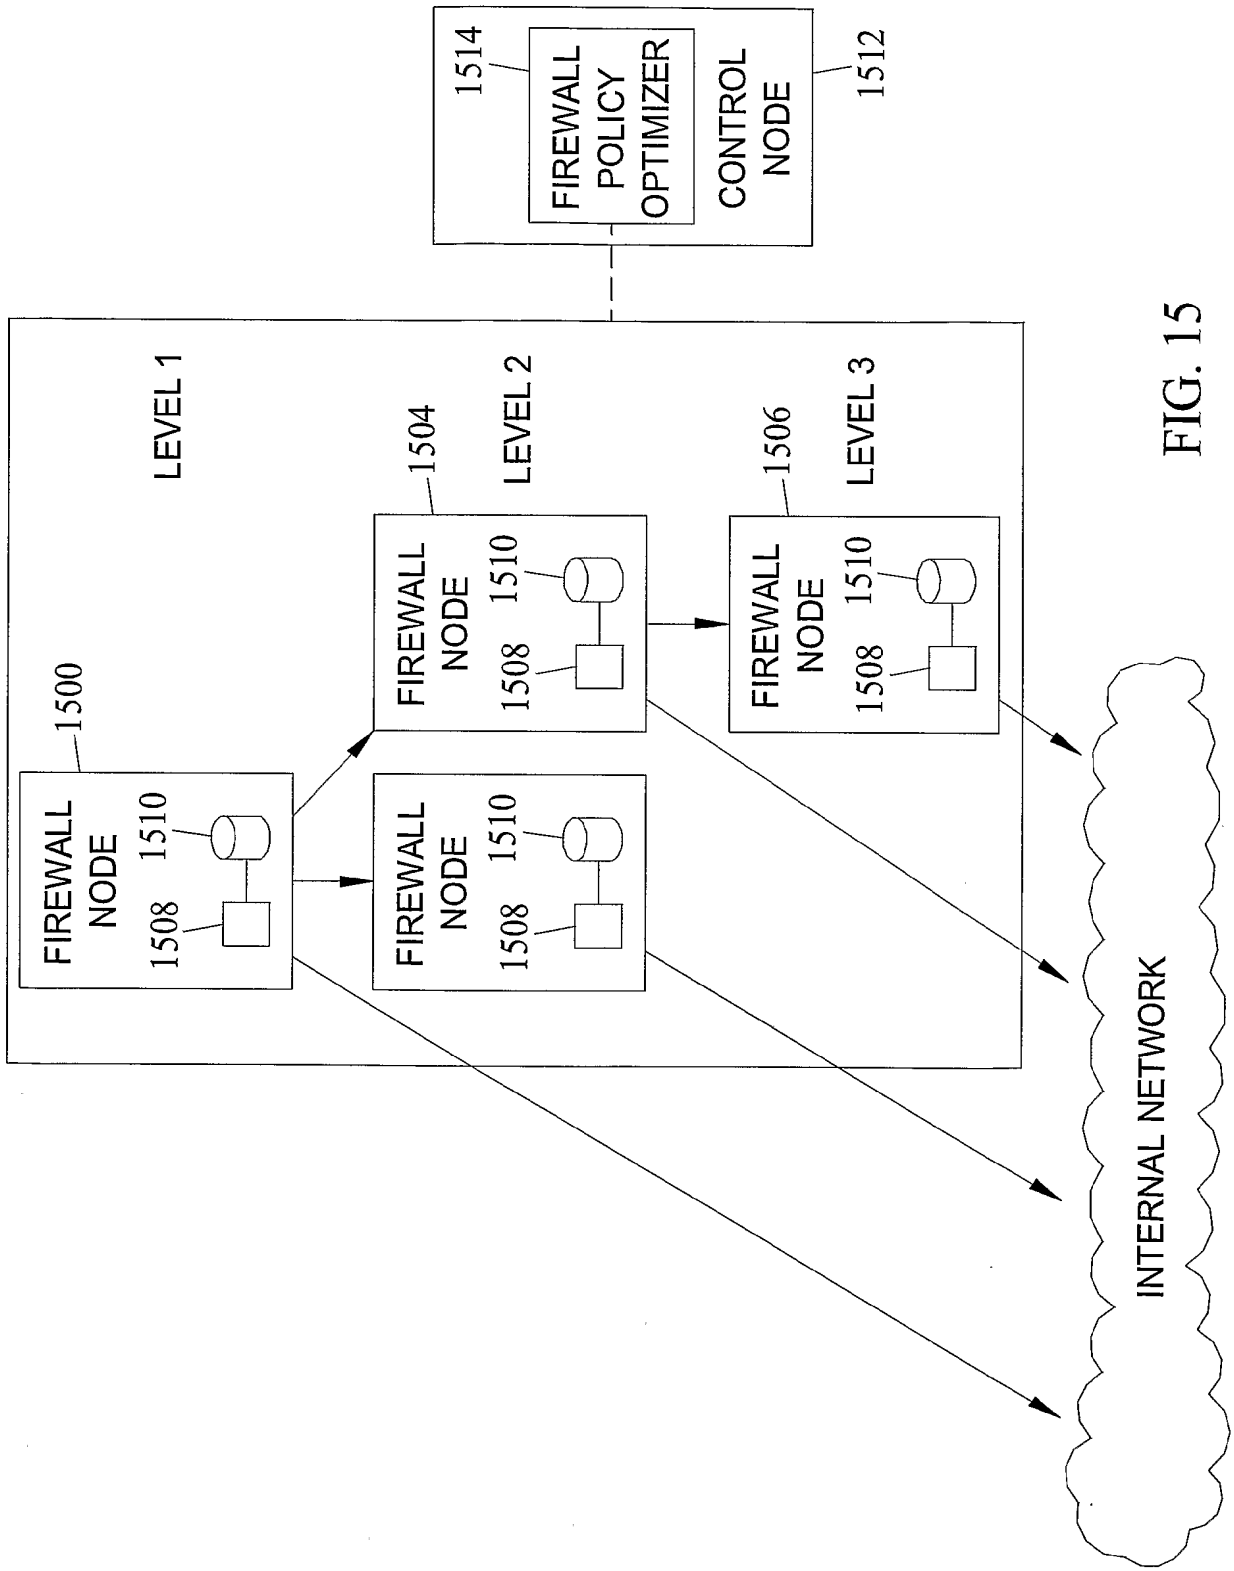
1514
FIREWALL POLICY OPTIMIZER
CONTROL NODE
1512
LEVEL 1
LEVEL 2
LEVEL 3
1500
1504
1506
FIREWALL NODE
1510
1508
INTERNAL NETWORK

FIG. 15

Espacenet

Bibliographic data: WO2006093557 (A2) — 2006-09-08

METHOD, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR IMPLEMENTING FUNCTION-PARALLEL NETWORK FIREWALL

Inventor(s): FULP ERRIN W [US]; FARLEY RYAN J [US] ± (FULP, ERRIN, W, ; FARLEY, RYAN, J)

Applicant(s): UNIV WAKE FOREST [US]; FULP ERRIN W [US]; FARLEY RYAN J [US] ± (WAKE FOREST UNIVERSITY, ; FULP, ERRIN, W, ; FARLEY, RYAN, J)

Classification: **- international:**
- cooperative: **H04L63/0218**; **H04L63/0263**

Application number: WO2005US47008 20051222

Priority number (s): US20040638436P 20041222

Also published as: WO2006093557 (A3) US2006195896 (A1) US8037517 (B2) EP1839188 (A2) EP1839188 (A4) more

Abstract of WO2006093557 (A2)

Methods, systems, and computer program products for providing function-paralell firewalls are disclosed. According to one aspect, a function-parallel firewall includes a first firewall node for filtering received packets using a first portion of a rule set including a plurality of rules. The first portion includes less than all of the rules in the rule set. At least one second firewall node filters packets using a second portion of the rule set. The second portion includes at least one rule in the rule set that is not present in the first portion. The first and second portions together include all of the rules in the rule set.

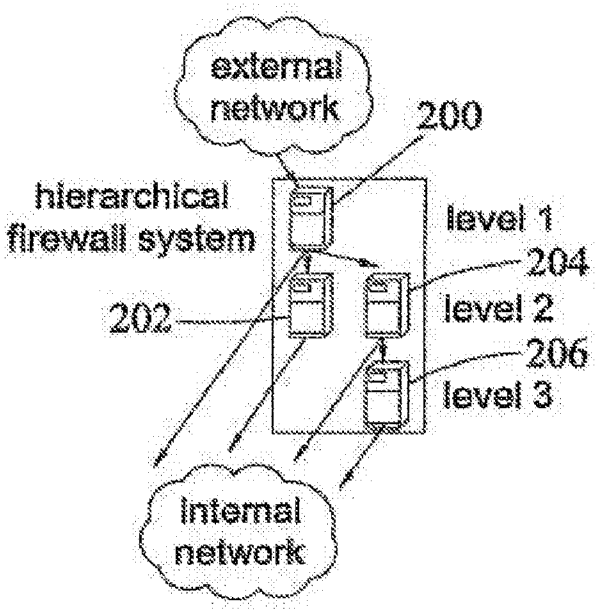

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
8 September 2006 (08.09.2006)

PCT

(10) International Publication Number
WO 2006/093557 A2

(51) International Patent Classification: Not classified

(21) International Application Number:
PCT/US2005/047008

(22) International Filing Date:
22 December 2005 (22.12.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/638,436 22 December 2004 (22.12.2004) US

(71) Applicant *(for all designated States except US)*: **WAKE FOREST UNIVERSITY** [US/US]; OFFICE OF TECHNOLOGY ASSET MANAGEMENT, WAKE FOREST UNIVERSITY HEALTH SCIENCES, 200 East First Street, Suite 101, Winston Salem, North Carolina 27101 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **FULP, Errin, W.** [US/US]; 308 N. Hawthorne Road, Winston-salem, North Carolina 27104 (US). **FARLEY, Ryan, J.** [US/US]; 212 Radford Street, Winston-salem, North Carolina 27106 (US).

(74) Agent: **HUNT, Gregory, A.**; JENKINS, WILSON & TAYLOR, P.A., SUITE 1400 UNIVERSITY TOWER, 3100 Tower Boulevard, Durham, North Carolina 27707 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR IMPLEMENTING FUNCTION-PARALLEL NETWORK FIREWALL

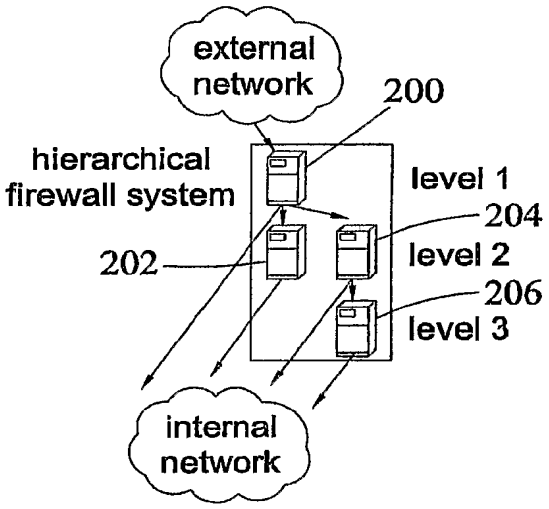

(57) Abstract: Methods, systems, and computer program products for providing function-paralell firewalls are disclosed. According to one aspect, a function-parallel firewall includes a first firewall node for filtering received packets using a first portion of a rule set including a plurality of rules. The first portion includes less than all of the rules in the rule set. At least one second firewall node filters packets using a second portion of the rule set. The second portion includes at least one rule in the rule set that is not present in the first portion. The first and second portions together include all of the rules in the rule set.

WO 2006/093557 A2

DESCRIPTION

METHOD, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR IMPLEMENTING FUNCTION-PARALLEL NETWORK FIREWALL

RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Patent Application Serial No. 60/638,436, filed December 22, 2004; the disclosure of which is incorporated herein by reference in its entirety.

GOVERNMENT INTEREST

This invention was made with Government support under Grant No. DE-FG02-03ER25581 awarded by U.S. Department of Energy, Mathematical and Computational Information Sciences Division. The U.S. Government has certain rights in the invention.

TECHNICAL FIELD

The subject matter described herein relates to implementing network firewalls. More particularly, the subject matter described herein relates to methods, systems, and computer program products for implementing a function-parallel network firewall.

BACKGROUND ART

Firewalls remain the frontier defense for securing networks that are vital to private industry, government agencies, and the military. They enforce a security policy by inspecting and filtering traffic arriving at or departing from a secure network [3, 21, 22]. Meaningful inspections often involve a complex process, and if firewalls are to remain effective, they must adapt to constantly changing network demands, technology, and security threats.

As firewall technology evolves to meet new demands and threats, it must continue to act transparently to legitimate users, with little or no effect on the perceived network performance. The networks between two communicating machines should be invisible, especially if traffic requires specific network Quality of Service (QoS), such as bounds on the packet delay, jitter, and throughput. The firewall should process the legitimate traffic quickly, efficiently,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 745

and never violate the desired QoS. Unfortunately, the firewall can easily become a bottleneck, given increasing traffic loads and network speeds [12, 14, 18, 20, 21]. Packets must be inspected and compared against complex rule sets and tables, which is a time-consuming process. In addition, audit files must be updated with current connection information. As a result, current firewalls have difficulty maintaining QoS guarantees. Since a firewall can be overwhelmed with traffic, it is also susceptible to Denial of Service (DoS) attacks [21]. Such attacks merely overload/saturate the firewall with illegitimate traffic. Audit files and state tables quickly fill available storage space and legitimate traffic queued behind suspicious traffic encounters long delays. Legitimate users notice poor network performance, and, in the worst case, the secure network is disconnected from the outside world.

Inspecting traffic sent between networks, a firewall provides access control, auditing, and traffic control based on a security policy [3, 21, 22]. The security policy is a list of ordered rules, as seen in Table 1, that defines the action to perform on matching packets. A rule r can be viewed as a k-tuple $(r[1], r[2], r[3], ..., r[k])$, where $r[i]$ is specified as a variable length prefix.

No.	Proto.	Source		Destination		Action
		IP	Port	IP	Port	
1	TCP	140.*	*	*	80	accept
2	TCP	150.*	*	120	80	accept
3	TCP	140.*	*	130	20	accept
4	UDP	150.*	*	*	3030	accept
5	*	*	*	*	*	deny

Table 1: Example Security Policy Consisting of Multiple Ordered Rules

In the examples described herein, rules will initially be represented as the 5-tuple (protocol, IP source address, source port number, IP destination address, destination port number). Fields can be fully specified or contain wildcards '*' in standard prefix format. For example the prefix 192.* would represent any IP address that has 192 as the first dotted-decimal number. In addition to the prefixes, each filter rule has an action, which is to accept or deny. An accept action passes the packet into or from the secure network,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 746

while deny causes the packet to be discarded. Rules are applied in order to every arriving packet until a match is found; otherwise, a default action is performed [21, 22].

Unfortunately, security policies may contain anomalies, such as a packet that matches multiple rules, where the matching rules specify conflicting actions. Security policy anomaly detection and correction is the subject of continued research [13, 21] and is not the focus of the present subject matter. Given the type of security policy required, rule sets can become quite large and complex. Security can be further enhanced with connection state and packet audit information. For example, a table can be used to record the state of each connection, which is useful for preventing certain types of attacks (e.g., TCP SYN flood) [21, 22].

Traditional firewall implementations consist of a single dedicated machine, similar to a router, that sequentially applies the rule set to each arriving packet. However, packet filtering represents a significantly higher processing load than routing decisions [15, 18, 22]. For example, a firewall that interconnects two 100 Mbps networks would have to process over 300,000 packets per second [21]. Successfully handling this high traffic becomes more difficult as rule sets become more complex [4, 14, 22]. Furthermore, firewalls must be capable of processing even more packets as interface speeds increase. In a high-speed environment (e.g. Gigabit Ethernet), a single firewall can easily become a bottleneck and is susceptible to DoS attacks [4, 7, 10, 12]. An attacker could simply inundate the firewall with traffic, delaying or preventing legitimate packets from being processed. Building a faster single firewall is possible [8, 15, 18, 19, 20]; however, the benefits are temporary (traffic loads and interface speeds are increasing); it is not scalable; it is a single point of failure; and it is generally not cost-effective for all installations.

A data-parallel firewall (or firewall sandwich) where each firewall node implements an entire firewall rule set is another approach to increase the speed of processing traffic [4, 12, 14, 22]. As shown in Figure 1, the system consists of multiple identical firewalls **100** connected in parallel. Each machine implements the complete security policy and arriving packets are distributed across the machines for processing in parallel [4]. A packet distributor **102**

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 747

implements a load balancing algorithm to distribute different packets to different firewall nodes. How the load-balancing algorithm distributes packets is vital to the system and is typically implemented as a high-speed switch in commercial products [11, 12]. Although data-parallel firewalls achieve higher throughput than traditional firewalls [4] and have a redundant design, they have difficulty maintaining QoS across networks. For example, legitimate traffic can encounter delays if it is queued behind traffic that requires more processing. Under these circumstances, users notice poor network performance, which is growing concern as more network applications require QoS assurances. Furthermore, stateful inspection requires all traffic from a certain connection or exchange to traverse the same parallel machine, which is difficult to perform at high speeds [14]. Therefore, new firewall architectures are needed to meet the demands of future networks and the challenges of increasing security threats.

SUMMARY

According to one aspect, the subject matter described herein includes a function-parallel firewall. The function-parallel firewall includes a first firewall node for filtering received packets using a portion of a rule set. The portion of the rule set may be less than all of the rules in the rule set. At least one second firewall node filters packets using at least one rule in the rule set that is not present in the portion implemented by the first firewall node. The first and second potions together include all of the rules in the rule set.

The term "firewall node," as used herein refers to a physical entity that implements one or more packet filtering rules. For example, a firewall node may be a single-processor computing device that implements the packet filtering rules. Alternatively, a firewall node may be a processor within a multi-processor computing device that implements one or more packet filtering rules. Thus, a multi-processor device may include multiple firewall nodes.

According to another aspect, the subject matter described herein includes a stateful function-parallel firewall system. The system includes a first function-parallel firewall subsystem including a plurality of firewall nodes implementing a first rule set for filtering packets arriving at a network. The first function-parallel firewall subsystem includes at least first and second firewall nodes respectively implementing first and second portions of the first rule set.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 748

The first portion includes at least one rule that is not present in the second portion. A second function-parallel firewall subsystem includes a second set of firewall nodes. The second set of firewall nodes implements a firewall rule set for filtering packets departing from the network. The second function-parallel firewall subsystem includes at least third and fourth firewall nodes. The third and fourth firewall nodes implement at least one rule that is not present in the second portion of the second rule set. The function-parallel firewall subsystems share state information regarding connections established through the first and second sets of firewall nodes.

According to another aspect, the subject matter described herein includes a firewall grid. The firewall grid includes a plurality of firewall nodes that are physically connected to each other via a network for filtering packets. A controller controls logical connections between the firewall nodes. The controller is adapted to configure the firewall nodes in a logical firewall hierarchy without changing physical connections between the firewall nodes.

According to yet another aspect, the subject matter described herein includes a method for distributing rules in a function-parallel firewall. The method includes defining a rule set for a function-parallel firewall. The method further includes assigning rules in the rule set to nodes and branches in a trie data structure. Each node in the trie data structure corresponds to a data field and each branch represents a value for each data field to be compared to be values in corresponding fields of received packets. The trie data structure is pruned in a manner that preserves ordering of the rules in the rule set.

According to yet another aspect, the subject matter described herein includes a method for assigning rules to a plurality of firewall nodes. The method includes assigning rules to nodes in a directed acyclical graph (DAG). Relationships between rules are represented by edges in the DAG. Rules are distributed among a plurality of firewall nodes using the DAG in a function-parallel manner and such that relationships between the rules specified in the DAG are preserved.

As used herein, the term "firewall" refers to a logical entity that is adapted to filter packets at the ingress and/or egress portions of a network based on a policy. The term "firewall" is intended to include systems that block packets from entering or leaving a network and systems that perform intrusion

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 749

detection and intrusion protection functions for packets entering or leaving a network. Thus, the term "function-parallel firewall" is also intended to include an intrusion detection system or an intrusion protection system that operates in a function-parallel manner as described herein.

The subject matter described herein can be implemented using any combination of hardware, software, or firmware. In one exemplary implementation, the subject matter described herein may be implemented using a computer program product comprising computer executable instructions embodied in a computer readable medium. Exemplary computer readable media suitable for implementing the subject matter described herein include disk memory devices, chip memory devices, programmable logic devices, application specific integrated circuits, and downloadable electrical signals. In addition, a computer program product that implements the subject matter described herein may be located on a single device or computing platform or may be distributed across multiple devices or computing platforms.

One objective of the subject matter described herein is to provide a new security architecture, referred to herein as a hierarchical firewall, that provides the functionality of a traditional firewall (enforcing a security policy) and includes the following improvements:

1. **Suitable for high speed networks and scalable for large traffic loads**. Firewalls must continue to perform security tasks quickly and efficiently, appearing transparent to legitimate users under varying traffic loads.

2. **Maintains QoS requirements and traffic differentiation**. A growing number of network applications require QoS for their operation. Firewalls must not violate these guarantees while performing their security tasks.

3. **Greatly lessens the impact of DoS attacks**. Firewalls are candidates for DoS attacks, since they serve as the point of entry between two networks. New architectures must be capable of handling these new threats.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 750

4. **Highly available.** Firewalls introduce a single point of failure between two networks. New firewall architectures must provide redundancy that prevents segmenting the secure network from the external network.

According to one aspect, the subject matter described herein includes a set of hierarchically configured computers that collectively enforce a desired security policy. Given this distributed configuration, traffic is quickly divided among the machines based on perceived threat. In one exemplary implementation, traffic that is considered safe may be promptly forwarded into the secure network, while the remaining traffic may be forwarded to different machines in the hierarchy for further scrutiny. Hence, traffic is segregated and queued based on security threat, yielding minimal delays for legitimate traffic. The system is also robust and highly available, since multiple machines are employed for the architecture and low cost, since Linux-installed PCs and open source software can be used for its implementation [21]. The proposed architecture is believed to be applicable to a wide variety of agencies in the public and private sectors.

The firewall architecture described herein will provide a scalable solution to varying traffic loads and QoS requirements, which are increasingly important as computing becomes more ubiquitous and applications become more diverse. The architecture is of interest to the networking community, since the architecture will result in new security policy models and optimization methods. The architecture is also relevant to the military and business community, since it provides a new, robust and cost-effective solution to security threats.

Some of the objects of the subject matter described herein having been stated hereinabove, and which are addressed in whole or in part by the subject matter described herein, other objects will become evident as the description proceeds when taken in connection with the accompanying drawings as best described hereinbelow.

BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the subject matter described herein will now be explained with reference to the accompanying drawings of which:

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 751

Figure 1 is a block diagram of a conventional data-parallel firewall;

Figure 2 is a block diagram illustrating an exemplary function-parallel firewall architecture according to an embodiment of the subject matter described herein;

Figure 3A is graph comparing simulation results for average packet delay for different firewall architectures;

Figure 3B is a graph comparing simulation results for maximum packet delay for different firewall architectures;

Figure 4A is a graph illustrating average packet delay for different packet types for a data-parallel firewall;

Figure 4B is a graph illustrating simulation results for packet delay for different packet types for a function-parallel hierarchical firewall according to an embodiment of the subject matter described herein;

Figure 5 is a data structure diagram illustrating a full trie representation of the security policy illustrated in Table 1;

Figure 6 is a data structure diagram illustrating a pruned trie representation of the security policy given in Table 1 where each packet is processed within four field comparisons;

Figure 7A is a data structure diagram illustrating a trie representing two security rules assuming that the default action per node is deny;

Figure 7B is a data structure diagram illustrating a reordered version of the trie data structure illustrated in Figure 7A;

Figure 8A is a diagram illustrating a Jackson model for a function-parallel hierarchical firewall according to an embodiment of the subject matter described herein;

Figure 8B is a diagram illustrating a Jackson model for a data-parallel firewall;

Figure 9 is a graph illustrating simulation results for average delay for different firewall architectures;

Figure 10 is a block diagram illustrating a stateful hierarchical firewall according to an embodiment of the subject matter described herein;

Figure 11A is a block diagram illustrating physical connections of a firewall grid according to an embodiment of the subject matter described herein;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 752

Figure 11B is a block diagram illustrating logical connections of the firewall grid illustrated in Figure 11A according to an embodiment of the subject matter described herein;

Figure 12 is a policy directed acyclical graph (DAG) representation of the firewall rules given in Table 2;

Figure 13A is a block diagram illustrating rule distribution among a data-parallel firewall;

Figure 13B is a block diagram illustrating a function-parallel firewall and rule distribution among nodes in the function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 14A is a block diagram illustrating one exemplary rule distribution among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 14B illustrates another distribution of rules among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 14C illustrates an illegal rule distribution among nodes in a function-parallel firewall;

Figure 15A is a block diagram illustrating non-redundant rule distribution among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 15B is a block diagram illustrating redundant rule distribution among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 16A illustrates a policy DAG where deny rules are indicated by a square node;

Figure 16B is a block diagram illustrating initial gateless rule distribution among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 16C is a block diagram illustrating final gateless rule distribution among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 17A is a graph illustrating simulation results for average packet delay as packet arrival rate increases for different firewall designs;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 753

Figure 17B is a graph illustrating simulation results for packet delay as the arrival rate and number of rules increase for different firewall designs;

Figure 18A is a graph illustrating simulation results for average packet delay for different firewall architectures;

Figure 18B is a graph illustrating simulation results for maximum packet delay for different firewall architectures; and

Figure 19 is a flow chart illustrating exemplary steps for filtering packets using a rule set distributed among nodes in a function-parallel manner according to an embodiment of the subject matter described herein.

DETAILED DESCRIPTION OF THE INVENTION

According to one aspect of the subject matter described herein may include multiple firewall machines configured in a hierarchy, as shown in Figure 2. In the example illustrated in Figure 2, a plurality of firewall nodes **200**, **202**, **204**, and **206** collectively implement a firewall policy. Each firewall node or set of nodes implements a different level of the firewall policy in a hierarchy. For example, firewall node **200** may implement a first level of a firewall policy that includes rules that are not present in subsequent levels. If a packet passes the first level of screening, the packet may be forwarded immediately to the internal network without impacting other firewall nodes. If the packet is identified as requiring further scrutiny, the packet may be forwarded to the next level where it is screened by firewall node **202** or **204** using rules that were not applied by firewall node **200**. If the packet passes level 2 screening, the packet may be forwarded directly to the internal network. If the packet is indicated as requiring further screening, the packet will be passed to level 3 where level 3 screening is applied by firewall node **206**.

At each level in the firewall hierarchy illustrated in Figure 2, firewall nodes may implement different portions of a packet filtering rule set. An architecture in which different firewall nodes implement different portions of a firewall rule set is referred to herein as a function-parallel architecture. A function-parallel firewall architecture in which packet filtering rules are distributed among firewall nodes such that if a packet passes or fails at a

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 754

certain level, the packet may bypass later levels is referred to herein as a hierarchical firewall architecture.

In the architecture illustrated in Figure 2, at each level, machines may be configured to perform a specific security task quickly and efficiently. In one implementation, the firewall rules may be distributed across the different hierarchy levels such that increased scrutiny is placed at lower levels. Since the rules are distributed among the levels, no machine in the hierarchy is required to implement the entire rule set. In contrast, a data-parallel firewall architecture requires all machines to implement the complete rule set.

As described above, in the exemplary architecture illustrated in Figure 2, packets start at the highest level and progress downward through the hierarchy. At each level, a packet can be forwarded to the internal network if deemed safe, forwarded to the next lower level for more processing, or dropped. The benefit of this divide-and-conquer approach is that traffic will incur delays consistent with its perceived security threat; therefore, traffic that is deemed safe will not be delayed by traffic that requires more processing. As a result, the hierarchical firewall can provide differentiated service and support QoS guarantees, unlike current designs. Furthermore, traffic that is forwarded to the internal network can be shaped at the particular firewall level, constraining the amount of traffic entering the network and lessening the effect of DoS attacks.

For example, in the hierarchical firewall depicted in Figure 2, packets are forwarded from the external network to the first level, where initial security checks are performed, such as a quick validity check and comparisons of IP addresses. Packets with trusted IP addresses are immediately forwarded to the internal network, while the remaining packets are sent to the next level. At the second level, packets are separated based on transport protocols and/or port numbers. Again, if a packet is deemed safe it is immediately forwarded to the secure network; otherwise, they are sent to the next level for auditing. Log files may be used at one or more of the firewall nodes indicate which connections are open and have been requested. Each level may also drop packets based on perceived threat.

The preceding example illustrates a simple hierarchical firewall; however, the size and complexity can vary. In general, the hierarchy attempts to quickly divide traffic and provide differentiated service. More levels can be used for

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 755

stronger security policies and to distribute traffic more efficiently. Determining the appropriate number of levels, the number of machines per level, and how to distribute the rule set are implementation details that may be used to optimize a function-parallel firewall according to an embodiment of the subject matter described herein.

Preliminary Experimental Results

In this section, simulation is used to compare a hierarchical firewall with a traditional single-machine firewall and a data-parallel firewall [4, 22]. Each architecture implemented the same 325-rule security policy, where each rule required 1 unit of time to complete. The single firewall architecture consisted of one machine that processed each packet received from the external network. The data-parallel architecture consisted of four machines, as shown in Figure 1, each implementing the same 325-rule set. The hierarchical firewall consisted of four machines configured in three levels, as shown in Figure 2. Rules were distributed across the levels such that the first level implemented 25 rules; the second level implemented 50 rules each; and the third level implemented 75 rules. The group of rules performed at a certain machine will be referred to as a security task. The hierarchical firewall incurred an additional delay of one time unit when a packet was transmitted from machine to machine. Packets randomly arrived to the system using a Poisson distribution with a mean of 30 units and would randomly match one of the 325 rules based on a uniform distribution. For classification purposes, each packet was associated with the security task it matched (which group of rules). In this experiment, the performance parameters analyzed are the overall average packet delay, the maximum packet delay, and the average delay experienced by each packet type.

Experimental results for the different architectures are shown in Figures 3 and 4. More particularly, Figure 3A illustrates average packet delay for a single firewall design, for the data-parallel firewall design illustrated in Figure 1, and for the hierarchical firewall design illustrated in Figure 2. Figure 3B illustrates maximum packet delay for the single, data-parallel, and hierarchical firewall designs. Figure 4A illustrates average packet delay for the data-parallel firewall design for different packet types. Figure 4B illustrates average packet

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 756

delay for the hierarchical firewall design for different packet types. Packet types refer to packets of a classified based on which security task was satisfied. Unlike the data-parallel firewall, the hierarchical design consistently yields different delays for each packet type.

For both delay measurements, the hierarchical architecture performed better than the single firewall and the data-parallel firewall. The single firewall had an average delay 31 times greater than the hierarchical architecture, which demonstrates how a single firewall machine can quickly become a bottleneck. Although the data-parallel firewall performed better than the single firewall, it had an average and maximum delay 6 times longer than the hierarchical architecture. In the parallel configuration, traffic that needed little processing often queued behind traffic that required more processing, which resulted in longer delays. In contrast, the hierarchical system consistently yielded better results, since traffic was distributed among the levels based on processing requirements.

The average delay experienced by each packet type is also important and is illustrated in Figures 4A and 4B. Delays in the hierarchical firewall were proportional to the amount of processing required; for example, security task 1 traffic had the shortest delay, security task 2 traffic had the second shortest delay, security task 3 traffic had the third shortest delay; and security task 4 traffic had the highest delay. Even though security tasks 2 and 3 were on the same level, security task 3 traffic had a delay slightly higher that security task 2 traffic, since it was queued with security task 4 traffic. Therefore, the hierarchical firewall is capable of providing differentiated service to multiple packet types/classes. In contrast, the delays varied for each packet type in the data-parallel firewall. Packets were constantly delayed by lower type/class traffic (e.g., security task 1 packets delayed by security task 4 packets); thus, no differentiation was achieved. Although these results are preliminary, they indicate the hierarchical architecture has merit, offering significantly better performance than current firewall systems.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 757

Refinements

The previous section described an initial hierarchical firewall architecture for employing a security policy. Although preliminary results indicate the architecture has merit, refinements can be made so that the design can achieve its full potential. One refinement relates to the hierarchical architecture. Important questions include determining the appropriate number of levels and how to effectively distribute the security policy rules. Another refinement may include providing stateful packet inspection (maintaining connection tables) and increasing redundancy. Finally, another refinement may include examining other high-speed firewall architectures, such as a dynamic firewall grid. Analytical models and simulation will be used for design and analysis; in addition, a prototype will be implemented along with analysis tools and software.

Architecture Design and Rule Set Distribution

Important questions about the firewall architecture include the number of levels, the number of machines per level, and the number of security tasks per machine. The security policy is central for addressing these questions and, for this reason, will be formally defined and modeled. The model will be optimized to reduce the physical size, minimize the traversal time, and provide differentiated service. Afterward, the optimized policy model will be used to implement the actual hierarchical firewall.

One method for designing the hierarchy is to represent the security policy graphically as a tree data structure. This representation will aid the translation to the actual firewall, since it has a hierarchical form. Although a variety of tree data structures exist, a trie will be initially investigated [1]. A trie is a special type of data structure that has the advantages of fast retrievals and quick unsuccessful search determination, which are important for enforcing a security policy. In a security policy trie, each path from the root node to a leaf corresponds to one rule. The nodes of the trie correspond to the rule fields, while branches are possible values. Figure 5 is a trie data structure for the rules in Table 1. For example, in Figure 5, the root node is the protocol field. The next level represents the source IP address, while the leaf nodes represent

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 758

the rule action. Rules that have the same field value at a node will share a branch for that value.

The trie in Figure 5 referred to as full trie, since every noncycle path from the root to a leaf has the same length, and every leaf lies at the same depth [6]. Unlike other policy models, the rule order must be preserved if the trie is to be successfully translated to a hierarchical firewall. The action to apply to a packet is determined by a preorder traversal. Starting at the root node, branches are compared to the fields of the packet. Since multiple branches may match a packet, they are traversed in order from left to right. If a packet field does not match any possible branches, it moves back to the previous node (parent), where the next branch is tested. Once the policy trie is developed, it can be translated into a hierarchical firewall. The nodes of the trie represent the machines in the hierarchy, while the branches represent the links between nodes. The value associated with each branch indicates the packet field value required to traverse it. Although a full trie representation of a security policy can be created, it is not necessarily optimal. Further trie manipulation is required create an effective hierarchical firewall, such as preventing routing to parents and limiting trie traversal time.

Pruning yields a trie that never requires routing to parent nodes, which provides faster packet processing [1, 15]. Assuming a full trie consists of n nodes and k levels (where k <= n), the worst case packet processing would be O(n*k) tuple-compares. In contrast, a pruned trie requires only O(k + n) tuple-compares in the worst case which is a processing reduction of 1/k.. Pruning is accomplished by replacing general rules with descendants. Rule r_1 is a descendant of rule r_2, if $r_2[i]$ is a prefix of $r_1[i]$ for $i = 1, \ldots, k$. Therefore, pruning occurs by pushing a rule down to all its descendents, then deleting the rule. An example of a pruned version of the trie illustrated in Figure 5 is given in Figure 6, where [*, *, *, deny] is a general rule that can be pushed down to every other rule. Although the original branch has been removed and added to the other branches, the order is preserved, if branches are tested from left to right. Current pruning algorithms do not consider rule order, since tries were primarily developed for storing data [6]. New pruning techniques that preserve order may be developed and used to distribute rules among nodes in a functional-parallel firewall according to an embodiment of the subject matter described

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 759

herein. Pruning may also cause the number of nodes in the hierarchy to increase, since a general rule is duplicated at different levels. For this reason, other techniques are needed to reduce the trie size.

Reducing the trie size is important, since delay increases each time a packet is forwarded from one node to another. The trie can be reduced by compressing nodes, which is achieved by removing single branching paths or consolidating sub-ties into one node [2, 6, 15]. As a result, a node may compare multiple rule fields or perform multiple security tasks per packet. For example, in Figure 5, the rule [UDP, 150.*, *, 3030, accept] is considered a leaf chain [6] and can be compressed into one node (a single machine). Two other methods commonly used to reduce the size of a trie are Directed Acyclic Graphs (DAG) and level reordering [1, 6, 15]. If the same sub-trie appears in multiple places, then the trie structure can be changed to a DAG. The possible use of DAG is seen in Figure 6, where nodes labeled $Port^a$ and $Port^b$ are equivalent sub-tries. Parent nodes of both sub-tries could point to just one, eliminating the need for the other.

Another method for reducing the trie size is level reordering. As seen in Figures 7A and 7B, the order of the levels can significantly affect the number of nodes in the trie. More particularly, in Figure 7A, an original trie representing the security rule [TCP, 120.*, 210.*, 90, accept] and [TCP, 130.*, 210.*, 90, accept] is illustrated. Both Figures 7A and 7B assume that the default action per node is deny. Since some of the rules in the original trie in Figure 7A include common elements, the rules can be reordered so that the number of branches in the subtrie are reduced. For example, in Figure 7A, the destination IP address and port for each branch of the subtrie are the same. Accordingly, if these common portions of the rule are implemented at higher levels, the duplicate nodes in the trie can be reduced, as illustrated by Figure 7B. More particularly, in Figure 7B, the nodes for destination IP and port are placed before the node for source IP address. As can be seen from Figure 7B, reordering the levels produces a smaller trie. Using a DAG can also reduce the number of nodes.

The problem of finding an ordering that guarantees a minimum size trie has been shown to be NP-complete; however, several heuristics have been developed [2, 6]. It should be noted that DAGs and reordering only reduce the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 760

number of nodes in the hierarchy and do not reduce the depth of the trie, which influences the traversal time.

Regardless of the applicability of these different trie modifications, they must be performed in concert to achieve an optimal hierarchical design. For example, the architecture utilizes multiple nodes to distribute traffic load and provide differentiated service; however, the trie path length (depth) impacts the packet traversal time. Therefore, compression and ordering must be performed together to provide differentiated service and optimize packet delay.

To achieve these goals, information not explicitly given in the security policy (e.g., QoS, level of trust, and processing time) may be used. For example, it is advantageous to place security tasks that require the most time at the lowest levels. This prevents packets from queueing behind traffic that requires more processing time; thus, lowering delay. A weight w_i can be assigned to each security task s_i (a single rule or multiple rules) based on the amount of time it requires. For example, an auditing security task requires more time (updating system files) than those that only perform field comparisons. Using these weights, the trie can be ordered to minimize the average traversal time. For example assuming i and j are hierarchy levels, where $i < j$, the trie can be ordered such that the average weight per level increases from the highest to the lowest level $\overline{w}_i < \overline{w}_j, \forall i \neq j$. In addition, the trie can be designed to separate traffic based on QoS or level of trust. Traffic would traverse different sub-tries, based on the type of service required, providing class isolation and different delays. The algorithms described in this section can then be applied to these sub-tries independently to reduce size and traversal time. These ideas were used to design the hierarchical firewall simulated in the preceding section, which showed excellent results. New algorithms for trie compression based on traversal time with constraints on QoS will be developed for optimizing function-parallel firewall design.

Packet traversal time remains a key performance metric for firewall systems. For this reason, queueing models will be used to analyze the delay of different firewall designs. The node delay encountered by a packet consists of the processing time, queueing delay, and transmission time. The processing delay is the time required to perform the security task. Assume node i must

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 761

process R_i rules per packet and can perform xi rules per unit time, then the service rate is x_i/R_i packets per unit time. Queueing is the delay associated with waiting for processing and depends on the system state. The transmission delay is the time required to transmit the packet, which depends on the packet's length and the line speed. The maximum delay experienced by a packet traversing the firewall is the sum of the node delays along the path from the root to a leaf.

If it is assumed the packet arrival and processing (transmission and security task) rates are exponentially distributed, the node can be modeled as an M/M/1 queue [5, 17]. The firewall can then be considered an open network of M/M/1 queues, where probabilities are assigned to each link to indicate the likelihood of moving to the next node. This is a Jackson network model, which has a product-form solution for the average delay [5, 17]. The average end-to-end delay for N cascading nodes (traversal path) can be computed as

$$E(T) = \sum_{i=1}^{N} \frac{1}{\mu_i - \lambda_i} \tag{1}$$

where $1/\mu_i$ is the service time (processing and transmission) and λ_i is the arrival rate to node i. Figures 8A and 8B respectively illustrate Jackson models for hierarchical and data-parallel firewalls. In Figures 8A and 8B, it is assumed that the total arrival rate to each system is λ packets per unit time, each system implements the same R rule security policy, and each node can perform x rules per unit time. For the hierarchical firewall, shown in Figure 8A, the policy rules are evenly distributed to each node; therefore, the service rate is $3x/R$ per node. The traffic arrival rate to node A is λ packets per unit time, half of the packets are forwarded to node B, while the other half is forwarded to node C. Although not considered in this example, some traffic may be forwarded into the secure network or dropped by node A; therefore in this case, the sum of the rates to nodes B and C would be less than λ. Using these values, the average end-to-end delay for either path in the hierarchical firewall is

$$E_h(T) = \frac{1}{\frac{3\chi}{R} - \lambda} + \frac{1}{\frac{3\chi}{R} - \frac{\lambda}{2}} \tag{2}$$

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 762

For the data-parallel configuration, illustrated in Figure 8B, traffic arrives at the packet distributor, which evenly distributes the traffic to each node. Each node implements the complete rule set *R*; therefore, the service rate is *x/R*. The end-to-end delay across any node in the data-parallel firewall is

$$E_p(T) = \frac{1}{\frac{\chi}{R} - \frac{\lambda}{3}} \quad (3)$$

Using these models, the effect of increasing traffic load λ on the average end-to-end delay, for the two architectures is shown in Figure 9. For this example, the arrival rate λ ranged from 1 to 50 packets per unit time. The rule set *R* consisted of 240 rules, while *x* equaled 5000 rules per unit time. Each of the hierarchical nodes incurred an additional delay of five time units to compensate for inter-node transmission. As expected, the average delay increases as the network load increases for both architectures; however, the average delay is consistently lower for the hierarchical firewall. This is due to the decreased rule set per node and the reduced arrival rates for nodes B and C. The average is an upper bound for the hierarchical firewall. Allowing node A to forward packets to the secure network or to drop packets would reduce the load on the second level and decrease delays. The arrival rate is less per node in the data-parallel firewall; however, the system has a higher end-to-end delay because each node must implement the complete rule set. The theoretical results are consistent with the experimental results described above. These models provide an important tool for analyzing the effect of load and rule distribution on end-to-end delay. In addition to analytical results, simulation will be used to measure the performance of the hierarchical designs under various conditions.

Stateful Firewalls

Another aspect of the subject matter described herein includes a stateful hierarchical firewall. The preceding design is considered stateless, or static, since the security action taken is based solely on the packet currently inspected. A stateful firewall provides better security by tracking and controlling network flows (packets belonging to a particular network connection or

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 763

exchange) [16, 21, 22]. In one exemplary implementation, a stateful firewall can be achieved by storing state information concerning each flow in a table, and then according it with the security policy. For example, each time a TCP connection is established from the internal network to the external network, it is logged in a table. The table may contain the source and destination addresses, port numbers, sequencing information, and connection flags. Afterwards, packets arriving from the external network claiming to be part of an existing TCP connection can be checked against the table. If a packet belongs to an existing connection or exchange, it can be accepted. Otherwise, it is dropped. Table entries expire, if packets belonging to an exchange are not observed by the firewall over a period of time [21].

Therefore, a stateful firewall is dynamic and limits the hackers' opportunity to exploit certain services. The increased security does require additional memory and time for table manipulation, which is problematic for high-speed networks. Furthermore, the firewall must process all traffic arriving and departing the secure network to adequately maintain the connection tables 21]. To solve these problems, the proposed firewall system will be divided into two hierarchies, one for arriving traffic and another for departing traffic, as illustrated in Figure 10. In Figure 10, a stateful firewall system includes first and second hierarchical firewall subsystems **1000** and **1002** and a control node **1004**. First and second hierarchical firewall subsystems **1000** and **1002** may be function-parallel firewall systems including a plurality of firewall nodes with different rules distributed among different nodes as described above. In one exemplary implementation, hierarchical firewall subsystem **1000** may include at least two firewall nodes that implement different portions of a first firewall rule set. Similarly, hierarchical firewall subsystem **1002** may include at least two firewall nodes that implement different portions of a second firewall rule set. Firewall subsystem **1000** may filter packets entering the network. Firewall subsystem **1002** may filter packets departing from the network. Firewall subsystems **1000** and **1002** may share state information regarding connections established through the nodes within each system. Control node **1004** may store the state information and may be accessible by the nodes in hierarchical firewall subsystems **1000** and **1002**. Rules can be distributed to the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 764

hierarchies and the designs can be optimized using the methods described in preceding section.

However, a new problem with the stateful model is the need to distribute state information across nodes. For example, nodes that scrutinize arriving traffic must be aware of any new connections permitted by nodes that inspect departing traffic. Furthermore, nodes must be informed of table entries that expire over time. To communicate state information, a central database machine, such as control node **1004**, may collect and distribute the information (table entries) to the appropriate nodes. In one exemplary implementation, control node **1004** may not perform table queries on behalf of nodes. Individual nodes will only query their local tables, which prevents the central machine from becoming a bottleneck. In addition, the central machine will only send information (table entries) to nodes that might require it. As a result of this policy, the table searched by an individual node is reduced (contains only relevant information) as well as the search time. In an alternate implementation, control node **1004** may store the connection table and may perform queries for packets received by firewall subsystems **1000** and **1002**.

Survivability and Firewall Grids

Another aspect of the subject matter described herein focuses on increasing the survivability of the hierarchical design. Although the proposed design has the benefit of increased packet throughput and minimal delay, the reliability of the system is decreased, since more machines (nodes) are deployed. For example, if the highest level node fails (is unable to inspect more traffic or goes off-line), then the entire hierarchy is unable to process packets. For this reason, the hierarchy is considered an n of n system, where all *n* machines must be functioning for the system to operate correctly [4, 12].

Survivability of the system can be improved with the use of redundant (parallel) nodes at different levels. For example, the first level may include multiple machines in parallel, where each machine implements the same security task (e.g., an IP check). An arriving packet would be sent to only one of the parallel machines, where it would be accepted, dropped, or forwarded to the next level for further processing. If one of the parallel nodes fails, then the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 765

remaining nodes can process the traffic. An additional benefit of parallel processing is the equal distribution of traffic among nodes, which can result in higher throughput and lower delays. Details such as where to locate parallel nodes and the number of machines required for each parallel configuration may depend on the desired degree of reliability versus cost of deploying redundant machines.

Yet another aspect of the subject matter described herein includes a grid of configurable firewalls (nodes) that can be used to efficiently process traffic, while being highly available. In one exemplary implementation, the grid may include a network of highly interconnected nodes (possibly via a high-speed switch or router) that can be logically deployed in various configurations (e.g., parallel, hierarchical, or hybrid). The advantage of the firewall grid is its ability to quickly adapt to new demands. When security needs change (e.g., due to policy changes, increased traffic volume, or node failure), these machines can be deployed or configured to maintain throughput and QoS requirements. For example, in a logical hierarchy, if a higher level encounters an increased traffic load, less frequently used redundant machines at lower levels can be reconfigured to process packets at the higher level. Unlike the hierarchical model, physical changes are not required, since all the machines are interconnected. Only the logical connection between nodes must be updated, which can be performed quickly. Like the stateful hierarchical firewall, the initial grid model will use a central machine to coordinate node configuration and communication. This central machine will play an important role in the system (balancing processing demand per node) and may constrain the efficiency of the firewall grid. Future versions of the firewall grid will attempt to place more intelligence in the nodes and allow them to interact independently and adapt dynamically under the guidance of a general security policy.

Figures 11A and 11B illustrate an example of a firewall grid according to an embodiment of the subject matter described herein. Referring to Figure 11A, a plurality of firewall nodes **1100-1114** are connected via a control node **1116**. Each firewall node **1100-1114** may be configured to implement different portions of a firewall rule set, as controlled by control node **1116**. Each of the firewall nodes is physically connected to the other firewall nodes via physical

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 766

connections **1118**. Physical connections **1118** may be local area network connections, such as Ethernet connections. As such, control node **1116** may be an Ethernet switch.

In Figure 11B, firewall nodes **1100**-**1114** are connected via logical connections **1120**. Logical connections **1120** may be implemented by control node **1116** without changing physical connections **1118**. For example, control node **1116** may configure the output port to which firewall **1106** is connected to receive frames from firewalls **1100**, **1102**, and **1104**. Similar internal configurations can be implemented for the remaining firewalls illustrated in Figure 11B. In Figure 11B, logical connections **1120** between the firewalls implement a hierarchical, function-parallel firewall design, such as that illustrated in Figure 1. In addition, the design illustrated in Figure 11B is also data-parallel at each level. For example, each firewall node at each level may implement a common rule or rule set. For example, the rule set implemented by the firewall nodes in level 1 is illustrated in Figure 11B as R1. The firewall nodes in level 2 implement a rule set represented by R2, and the firewall nodes in level 3 implement a rule set represented by R3. If any of the firewall nodes in a given level fails, the remaining firewall nodes in that level may take over operation for the failed firewall. Alternatively, as described above, if a failure or a bottleneck occurs at one level, a less-used firewall at another level may be logically configured to operate at the level where the failure or congestion occurred. Packet distributors **1122** may be configured to distribute the firewall rules between the firewalls at each level. Thus, as illustrated by Figures 11A and 11B, logical connections between firewall nodes and a firewall grid may be dynamically altered (i.e., changed on the fly without ceasing packet filtering operation) by a controller node without changing the physical connections between the firewalls.

Firewall Implementation

In addition to the analytical models and simulations described above, an implementation of a hierarchical firewall to measure performance will now be described. One exemplary implementation may utilize Linux-based PCs and Ethernet connections. Each PC will contain multiple Ethernet cards and be

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 767

loaded with the Linux 2.4 kernel. The 2.4 kernel has support for various packet queueing disciplines, traffic control methods, and packet processing (e.g., iproute2, tc, and iptable). Collectively, these functions are required for advanced packet filtering. For example, iproute2 and iptable can route packets based on information other than the destination address, while tc can mark and shape traffic. PCs will serve as hierarchical nodes, sources and receivers of network traffic, and hierarchy controllers.

Results from an initial implementation of a three-node hierarchical firewall indicate that these open-source components are suitable for high-speed networks. The Computer Science Department at Wake Forest University has 10 PCs available for this project, where each has a 600 Mhz processor and multiple 100 Mbps Ethernet cards. These machines will serve as an initial testbed for system development; however, faster machines and high-speed network connections are required for actual system testing. For this reason, this project requests the purchase of 6 additional PCs (2.4 Ghz processors), each with multiple Gigabit Ethernet cards (3 cards per machine). These new machines will be configured as the firewall, while the current machines will serve as traffic generators and receivers. For comparison purposes, parallel and single-machine firewalls will also be implemented using this equipment. Any software required will be developed and provided by the Computer Science Department.

Firewall Rule Distribution for Function-Parallel Firewall Design

As stated above, one aspect with creating an optimal function-parallel firewall design includes distributing rules among firewalls in the firewall rule set. Firewall rules are preferably distributed among firewall nodes in a manner that preserves firewall policy integrity. The following sections describe an exemplary firewall rule set and different options for distributing the firewall rule set among firewall nodes in a function-parallel firewall design.

Firewall Policy

In this section and the following sections, a rule r is modeled as an ordered tuple of sets, $r = (r[1], r[2], ..., r[k])$. Order is necessary among the tuples since comparing rules and packets requires the comparison of

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 768

corresponding tuples. Each tuple $r[l]$ is a set that can be fully specified, specify a range, or contain wildcards '*' in standard prefix format. For the Internet, security rules are commonly represented as a 5-tuple consisting of: protocol type, IP source address, source port number, IP destination address, and destination port number [21, 22]. For example, the prefix 192.* would represent any IP address that has 192 as the first dotted-decimal number. Given this model, the ordered tuples can be supersets and subsets of each other, which forms the basis of precedence relationships. In addition to the prefixes, each filter rule has an action, which is to accept or deny. However, the action will not be considered when comparing packets and rules. Similar to a rule, a packet (IP datagram) d can be viewed as an ordered k-tuple $d = (d[1], d[2], ..., d[k])$; however, ranges and wildcards are not possible for any packet tuple.

Using the previous rule definition, a standard security policy can be modeled as an ordered set (list) of n rules, denoted as $R = \{r_1, r_2, ..., r_n\}$. A packet d is sequentially compared against each rule r_i starting with the first, until a match is found ($d \Rightarrow r_i$) then the associated action is performed. A match is found between a packet and rule when every tuple of the packet is a subset of the corresponding tuple in the rule.

Definition Packet d matches r_i if

$$d \Rightarrow r_i \text{ iff } d[l] \subseteq r_i[l],\ l = 1, ..., k$$

The rule list R is comprehensive if for every possible legal packet d a match is found using R. Furthermore, two rule lists R and R' are equivalent if for every possible legal packet d the same action is performed by the two rule lists. If R and R' are different (e.g. a reorder) and the lists are equivalent, then the policy integrity is maintained.

As described above, a rule list has an implied precedence relationship where certain rules must appear before others if the integrity of the policy is to be maintained. Table 2 shown below illustrates an exemplary rule set that will be used herein to illustrate precedence relationships and rule assignment according to an embodiment of the subject matter described herein.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 769

No.	Proto.	Source		Destination		Action	Prob.
		IP	Port	IP	Port		
1	UDP	1.1.*	*	*	80	deny	0.01
2	TCP	1.*	*	1.*	90	accept	0.02
3	TCP	2.*	*	2.*	20	accept	0.25
4	UDP	1.*	*	*	*	accept	0.22
5	*	*	*	*	*	deny	0.50

Table 2: Example Security Policy Consisting of Multiple Ordered Rules

In Table 2, rule r_1 must appear before rule r_4, likewise rule r_5 must be the last rule in the policy. If for example, rule r_4 was moved to the beginning of the policy, then it will shadow [23] the original rule r_1. Shadowing is an anomaly that occurs when a rule r_j matches a preceding rule r_i, where $i < j$. As a result of the relative order r_j will never be utilized. However, there is no precedence relationship between rules r_1, r_2, or r_3 given in Table 2. Therefore, the relative order of these three rules will not impact the policy integrity and can be changed to improve firewall performance.

As described in [28], the precedence relationship between rules in a policy can be modeled as a Policy Directed Acyclical Graph (DAG). Let $G = (R,E)$ be a *policy DAG* for a rule list R, where vertices are rules and edges E are the precedence relationships (constraint). A precedence relationship, or edge, exists between rules r_i and r_j, if $i < j$ and the rules intersect [28].

Definition The intersection of rule r_i and r_j, denoted as $r_i \cap r_j$ is

$$r_i \cap r_j = (r_i[l] \cap r_j[l]),\ l = 1, \ldots, k$$

Therefore, the intersection of two rules results in an ordered set of tuples that collectively describes the packets that match both rules. The rules r_i and r_j intersect if every tuple of the resulting operation is non-empty. In contrast, the rules r_i and r_j do not intersect, denoted as $r_i \not\cap r_j$, if at least one tuple is the empty set. It should be noted that the intersection operation is symmetric; therefore, if r_i intersects r_j, then r_j will intersect r_i. The same is true for rules that do not intersect.

For example, for the rules given in Table 2, the intersection of r_1 and r_4 yields (UDP, 1.1.*, *, *, 80). Again, the rule actions are not considered in the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 770

intersection or match operation. Since these two rules intersect, a packet can match both rules for example d = (UDP, 1.1.1.1, 80, 2.2.2.2, 80). The relative order must be maintained between these two rules and an edge drawn from r_1 to r_4 must be present in the DAG, as seen in Figure 12. In Figure 12, the policy DAG represents the rules in Table 2. Vertices **1200** indicate rules while edges **1202** indicate precedence relationships. Rules r_2 and r_3 in Table 2 provide an example of rules that can be re-ordered without changing the firewall policy. These two rules do not intersect due to the fifth tuple (destination port). A packet cannot match both rules indicating the relative order can change; therefore, an edge will not exist between them.

As described above, traditional firewall implementations consist of a single, dedicated machine, similar to a router, that applies the rule set to each arriving packet. However, packet filtering represents a significantly higher processing load than routing decisions [15, 18, 22]. For example, a firewall that interconnects two 100 Mbps networks would have to process over 300,000 packets per second [21]. Successfully handling these traffic loads becomes more difficult as rule sets become more complex [4, 14, 22]. Furthermore, firewalls must be capable of processing even more packets as interface speeds increase. In a high-speed environment (e.g. Gigabit Ethernet), a single firewall can easily become a bottleneck and is susceptible to DoS attacks [4, 7, 10, 11]. Building a faster single firewall is possible [8, 15, 18, 19, 20]; however, the benefits are temporary (traffic loads and interface speeds are increasing); it is not scalable; it is a single point of failure; and it is generally not cost-effective for all installations.

Parallelization has been demonstrated to offer a scalable approach for improving the performance of network firewalls. As seen in Figure 13A, parallel designs consist of multiple firewalls, referred to as firewall nodes, connected in an array. In Figure 13A, firewall nodes **1300** each implement a common rule set. A packet distributor **1302** load shares packets among firewall nodes **1300**. In Figure 13B, a function-parallel design is illustrated where firewall nodes **1304** implement different portions of a firewall rule set. A packet duplicator **1306** duplicates each packet to each firewall node. A gate **1308** controls packets entering the network based on control information received from firewall nodes

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 771

1304. The operation of the system illustrated in Figure 13B will be described in more detail below.

The two designs depicted in Figures 13A and 13B are different based on what is distributed: packets or rules. For example, the design developed by Benecke et al consisted of multiple identical firewall nodes connected in parallel, as shown in Figure 13A. Each firewall node implements the complete security policy *R*, and arriving packets are distributed across the firewall nodes for processing (one packet is sent to one firewall node) [4]; therefore different packets are processed in parallel. Using terminology developed for parallel computing, this design is considered data-parallel since the data (packets) is distributed across the firewall nodes [25]. How the load-balancing algorithm distributes the packets is vital to the data-parallel system and typically implemented as a specialized high-speed switch [11, 12]. Although data-parallel firewalls have been shown to achieve higher throughput than traditional firewalls and have a redundant design [4], they suffer from two major disadvantages. First, these designs have difficulty maintaining QoS across networks. For example, legitimate traffic can encounter delays if it is queued behind traffic that requires more processing. Under these circumstances, users notice poor network performance, which is a growing concern as more network applications require QoS assurances. Second, stateful inspection requires all traffic from a certain connection or exchange to traverse the same firewall node, which is difficult to perform at high speeds using the data-parallel approach [4, 14].

Rule Distribution and Function-Parallel Firewall

As described above, new scalable firewall designs are needed to cope with increasing traffic loads, network speeds, and security threats. To address these issues, one implementation of the subject matter described herein includes a firewall design that consists of multiple firewalls connected in parallel and a gate machine, as shown in Figure 13B. Again in this design, the individual firewalls will be referred to as firewall nodes. However unlike the data-parallel design that distributes the packets, the function-parallel design distributes the policy rules across the firewall nodes. The distribution is done such that policy integrity is maintained. Therefore, a single traditional firewall

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 772

and a function-parallel firewall will always give the same result for any packet. The rules distributed to a firewall node will be referred to as the local rule set or local policy. Using parallel computing terminology this design will be referred to as function-parallel, since the rules are distributed across the firewall nodes [25].

The general operation of one exemplary function-parallel system can be described as follows. When a packet arrives to the function-parallel system, the packet is forwarded to every firewall node. Each firewall node processes the packet using its local rule list. The firewall node then signals the gate indicating either no match was found or provides the rule number and action if a match was found. Therefore, no match is a valid response and is required for this basic function-parallel design. The gate keeps track of the results and determines the final action to perform on the packet. For example, a first match policy can be implemented by applying the action of the lowest numbered matching rule to the packet.

The performance of this basic design can be improved by allowing the firewall nodes to pipeline packets [25]. For example, once a firewall node has completed processing a packet, it can immediately begin processing the next packet. This allows the firewall nodes to work independently, thus increasing the overall system performance. However, pipelining requires the gate to keep track of multiple packets simultaneously. Performance can also be increased if the gate can signal the firewall nodes that further processing of a certain packet is no longer needed (similar to short-circuit evaluation). This occurs when the appropriate match has been found by a firewall node before the other firewall nodes have completed processing the same packet. Short-circuit evaluation requires the gate to know how the rules are distributed as well as the dependencies, which is discussed in detail below.

The function-parallel design has several significant advantages over traditional and data-parallel firewalls. First, the function-parallel design results in the faster processing since every firewall node is utilized to process a single packet. Reducing the processing time, instead of the arrival rate (the data-parallel paradigm), yields better performance that will be proven theoretically and demonstrated experimentally. Second, unlike the data-parallel design, function-parallel can maintain state information about existing connections.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 773

Maintaining state can be viewed as the addition of a new rule corresponding to a requested connection. Unlike the data-parallel design, the new rule can be placed in any firewall node since a packet will be processed by every firewall node. Third, the function-parallel design has the unique ability to provide service differentiation between classes of traffic; therefore, the design can provide QoS requirements. Finally, the function-parallel system can be implemented using any current firewall system (as discussed below, the gate can be removed), without any specialized hardware. The disadvantage of the system is the limited scalability, since the system cannot have more firewall nodes than rules, unless data-redundancy is provided. However, given the size of most firewall policies range in the thousands to hundred of thousands of rules, which is expected to increase, the scalability limit is not an important concern.

The operation and performance of the function-parallel design depends on the correct distribution of policy rules. Rule distribution techniques that maintain integrity, increase processing speed, allow redundancy and robustness, and eliminate the gate are described in the following sections.

Rule Distribution and Policy Integrity

As described above, the function-parallel design distributes the security policy rules across an array of firewall nodes. Using this design, all the firewall nodes process an arriving packet in parallel using its local rule set. The rules must be distributed such that the integrity of the policy is maintained. Again, it will be assumed in the examples described herein that a first-match policy is desired. Therefore, the distribution of the rules must not introduce any anomalies into the original rule list, such as shadowing.

Let m be the number of firewall nodes in the system and j be a firewall node $j = 1...m$. Each firewall node has a local rule set that is a portion of the security policy. Denote the local rule list for firewall node j as R_j that has n_j rules. The firewall node will process an arriving packet using the local rule list in a top-down fashion to find the first match. It should be noted that a software implementation of the firewall node according to the subject matter described herein is not required to have a rule list. The function-parallel design is independent of software implementation. In this scenario integrity is maintained

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 774

under two conditions: every rule is assigned to at least one firewall node, and a policy DAG edge never traverses upward within a firewall node. The first condition assures every rule in the original policy exists in the function-parallel system, while the second condition prevents rule shadowing within a firewall node. Consider a two-firewall node system and the five-rule policy given in Table 2. Possible rule distributions are depicted in Figures 14A-14C. The distribution illustrated in Figure 14A will be referred to as vertical, while the distribution illustrated in Figure 14B will be referred to as horizontal. In both cases the five rules are present and all policy DAG edges are downward; thus integrity is maintained. In contrast, Figure 14C is an illegal distribution because rule r_4 shadows r_1 (upward policy DAG edge).

The gate can also utilize the policy DAG to determine the first match, and if a short-circuit evaluation can occur. For example consider the distribution illustrated in Figure 14A. Assume a packet matches r_3 in the second firewall node and this result is reported to the gate. Using the policy DAG, the gate can determine this is the first match, since it is impossible for a packet that matches r_3 to also match any rules that preceded it in the original policy (rules r_1 or r_2). The gate can inform the remaining firewall nodes to stop processing the packet.

Optimal Rule Distribution

The previous section described the distribution requirements for maintaining policy integrity. However it was proven in [27] that any linear arrangement of a policy DAG maintains integrity, which also holds for a local policy DAG. Furthermore as evident in the previous section, many different distributions are possible. Given the possible rule distributions for the function-parallel design, a distribution is sought that minimizes the average number of comparisons required per packet [27].

Achieving this goal will require information not present in the firewall rule list. Certain firewall rules have a higher probability of matching a packet than others. As a result, it is possible to develop a policy profile over time that indicates frequency of rule matches (similar to cache hit ratio). Let $P = \{p_1, p_2, \ldots, p_n\}$ be the policy profile, where p_i be the probability that a packet will match rule r_i (first match in the policy). Furthermore, assume a packet will always find

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 775

a match, $\sum_{i=1}^{n} p_i = 1$; therefore R is comprehensive. Using this information, the average number of rule comparisons required for the policy R using a standard single machine firewall is

$$E(R) = \sum_{i=1}^{n} i \cdot p_i \quad (4)$$

This equation also applies to firewall nodes. However, the rules of a local policy include rules of the original policy; therefore, a mapping between the local and original policy rules is needed. Let $r_{i,j}$ represent the i^{th} rule of firewall node j, that corresponds a certain rule in the original rule list R. Consider the local rule list for the second firewall node depicted in Figure 14A. The local policy is $R_2 = \{r_4, r_5\}$, where $r_{1,2} = r_4$ and $r_{2,2} = r_5$. Similarly, let $p_{i,j}$ be the probability of matching the i^{th} rule of firewall node j. Therefore the average number of comparisons required for firewall node j is

$$E(R_j) = \sum_{i=1}^{n_j} i \cdot p_{i,j} \quad (5)$$

Using the previous equation, the average number of comparisons required for the function-parallel system is

$$E_f(R) = \sum_{j=1}^{m} E(R_j) = \sum_{j=1}^{m} \sum_{i=1}^{n_j} i \cdot p_{i,j} \quad (6)$$

For example, the average number of comparisons required for the distribution given in Figure 14A is 2.24, while the average number of comparisons for the distribution given in Figure 14B is 2.45.

Given a policy DAG $G = (R,E)$ and policy profile P, a distribution is sought that minimizes Equation 6. In the absence of precedence relationships, the average number of comparisons is minimized if the rules are sorted in non-increasing order according to the probabilities [28], which is also referred to as Smith's algorithm [29]. The sorted rules are then distributed in a horizontal fashion, as depicted in Figure 14B. Precedence constraints cause the problem

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 776

to be more realistic; however, precedence constraints also make determining the optimal permutation more problematic. In [27] determining the optimal rule order for a single machine firewall is proven to be *NP*-complete; therefore, so is determining the optimal function-parallel rule distribution. However, [27] presents a simple sorting algorithm to improve the optimal order, that can be modified to apply to function-parallel rule distribution. This modified algorithm sorts the policy rules using the method described in [27], then distributes the rules in a horizontal fashion.

Redundancy

Similar to the traditional and data-parallel firewalls, a disadvantage of the current function-parallel design is its inability to withstand a single firewall node failure. If a single firewall node fails in the function-parallel system, then part of the policy is lost and integrity is not maintained. A simple solution commonly used for traditional firewalls is to duplicate the entire system; however, this solution is cost prohibitive, not efficient, and difficult to manage.

Redundancy can be provided in the function-parallel design by replicating rules, instead of firewall nodes. Consider the function-parallel system depicted in Figure 15A, that includes four firewall nodes and an eight-rule policy distributed in a horizontal manner. If the local rule list of the right neighbor is added to the end of every firewall node, as illustrated in Figure 15B, then any non-continuous firewall nodes can fail and integrity will be maintained. It should be noted that the criteria for maintaining integrity still apply: every rule must be implemented at a node and no upward edges in the local policy DAG. Furthermore, the redundant design requires short-circuit evaluations to prevent multiple firewall nodes from evaluating the same rule.

Gateless Design

The gate plays an important role in the function-parallel design. It accepts information from the firewall nodes to determine the correct action to apply to a packet. Unfortunately, the gate is a special device that must coordinate with the individual firewall nodes in the system, which may be problematic to implement. However, it is possible to eliminate the gate and allow the firewall nodes to operate independently without any coordination. The

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 777

significant advantage is that the gateless function-parallel design can be implemented with any firewall system (from open source to proprietary).

Eliminating the gate requires a rule distribution such that for any given packet, only one machine will ever accept and the remaining machines will always deny. As described above, rule actions are either accept or deny. This can be seen in the policy DAG by denoting deny rules with a square and accept rules with a circle. To assure only one machine will accept a packet, an accept rule in a firewall node cannot intersect with another accept rule in any other firewall node. As described above, the intersection of two rules describes the set of packets that can match both rules. If an edge spans the local rules of firewall nodes, then the set of packets would be accepted by both firewall nodes; thus, duplicate packets would be accepted. In addition, every drop rule must be present in each firewall node to assure the remaining firewall nodes will drop the packet. Given these two requirements a policy DAG edge will not span firewall nodes, as illustrated in Figures 16B and 16C.

For example, consider a two-firewall node function-parallel system that implements the five-rule policy given in Table 2. The policy DAG is given in Figure 16A, where the first and last rules are deny. An initial gateless rule distribution is given in Figure 16B which has no edges spanning the firewall nodes. Given the local policy DAG, deny rules can be removed if: no inbound edges exist for the deny rule; and all out-going edges for the deny rule connect to a superset deny rule. Using these constraints, the first rule in the second firewall node depicted in Figure 16B can be removed (it is a subset of r_5 and has no other edges). This gives the final gateless rule distribution shown in Figure 16C. As a result, the first firewall node in this distribution will only process UDP traffic, while the second machine only processes TCP. The rules can be divided with respect to other tuples, and can be done to provide QoS guarantees.

Theoretical Results

Assuming arrivals and service times are exponentially distributed, a firewall system can be considered an open network of M/M/1 queues (Jackson network) [5, 17]. Probabilities can be assigned to each link to indicate the likelihood of moving to the next node, which are given from the policy profile (hit

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 778

ratio) described above. The average end-to-end delay for q cascading machines (traversal path) can be computed as

$$E(T) = \sum_{i=1}^{q} \frac{1}{\mu_i - \lambda_i} \tag{7}$$

where $1/\mu_i$ is the service time (processing and transmission) and λi is the arrival rate to node i. As a result, we have a theoretical model for the average delay across a firewall system. Consider the data-parallel and functional-parallel firewall designs given in Figures 13A and 13B. Assume each system consists of m machines and implements the same n rule security policy. Let the total arrival rate to each system be λ packets per unit time and each node performs x rules per unit time.

For the data-parallel firewall, traffic arrives at the packet distributor, which evenly distributes the traffic. As a result the arrival rate to each machine is λ m. The service rate is x n, since each node implements the complete rule set. The end-to-end delay across any machine in the parallel firewall is

$$E_d(T) = \frac{1}{\frac{x}{n} - \frac{\lambda}{m}} \tag{8}$$

In contrast, rules are distributed across each machine, in the function-parallel firewall. Furthermore, all traffic arriving to the system is forwarded to each machine. If we assume the rules in each machine are independent (no policy DAG edges spanning machines), then the end-to-end delay across any machine in the functional-parallel firewall is

$$E_f(T) = \frac{1}{\frac{m \cdot x}{n} - \lambda} \tag{9}$$

The reduction in delay compared to a data-parallel firewall system is 1 m; therefore, the functional-parallel system has the potential to be m times faster than a (data) parallel system. However, the scalability of the functional-parallel system is dependent on the rule set.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 779

Figures 17A and 17B compare the performance of function-parallel, data-parallel, and single firewall designs. More particularly, Figure 17A compares the average among the different firewall designs as the arrival rate λ increases. As can be seen from Figure 17A, the function-parallel design achieves lower average delay as the arrival rate increases than the data-parallel design. Figure 17B illustrates the average delay as arrival rate and the number of rules increase. In Figure 17B, the function-parallel design again achieves better simulated performance than the data-parallel design.

Experimental Results

The previous section described models that provide insight to the average delay across different firewall designs. This information is beneficial in guiding initial design; however, we are also interested in the maximum delay encountered. For this reason, a discrete event simulator has been developed to model the different systems.

The data-parallel firewall and function-parallel firewall designs were simulated, where both systems implemented a 50-rule policy and each consisted of 5 firewall nodes. Three classes of traffic were modeled, first, second, and background. The first and second classes represent legitimate traffic that was eventually accepted by the firewall system, while background traffic was typically dropped. The two designs were simulated under normal conditions, where each traffic type had the same volume. Figures 18A and 18B shows the simulation results. As seen in Figure 18A and predicted by the theoretical models, the average delay encountered by packets in the functional-parallel design was lower than the data-parallel designs. The function-parallel design also provided lower maximum delays than the data-parallel, as seen in Figure 18B. In summary, the theoretical models and simulation results indicate function-parallel is a better design than a traditional single firewall system or a data-parallel system; yet, function-parallel requires no more computing resources than data-parallel.

Summary

Thus, as described above, the subject matter described herein includes methods, systems, and computer program products for providing a function-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 780

parallel firewall where different portions of a firewall rule set are distributed among different firewall nodes. Figure 19 is a flow chart illustrating exemplary steps of a process for filtering packets using a rule set distributed among firewall nodes in a function-parallel manner. Referring to Figure 19, in step **1900**, packet filtering rules of a rule set are distributed among nodes in a function-parallel manner so that different nodes implement different portions of the rule set. In step **1902**, packets to be filtered are replicated to each of the nodes. In step **1904**, the packets to be filtered are replicated to each of the nodes. In a hierarchical design, step **1904** may be implemented at each level in the hierarchy. If a packet passes or fails at one level, the packet may be forwarded into a network or dropped, depending on the distribution of the rule set. In step **1904**, the rules are applied to filter the packets.

The steps illustrated in Figure 19 and in all of the examples described herein may be applied to firewall systems that filter packets based on packet header information, intrusion detection systems that filter packets based on packet payload information, or intrusion protection systems that filter packets based on a combination of packet header and packet payload information.

The following references correspond to the citations in this document. The disclosure of each of these references is hereby incorporated herein by reference it its entirety.

References

[1] A. V. Aho, J.E. Hopcroft, and J. D. Ullman. *Data Structures and Algorithms*. Addison-Wesley, 1987.

[2] M. Al-Suwaiyel and E. Horowitz. Algorithms for Trie Compaction. *ACM Transactions on Database Systems*, 9(2):243–263, June 1984.

[3] S. M. Bellovin and W. Cheswick. Network Firewalls. *IEEE Communications Magazine*, pages 50–57, Sept. 1994.

[4] C. Benecke. A Parallel Packet Screen for High Speed Networks. In *Proceedings of the 15th Annual Computer Security Applications Conference*, 1999.

[5] G. Bolch, S. Greiner, H. de Meer, and K. S. Trivedi. *Queueing Networks and Markov Chains*. John Wiley and Sons, Inc., 1998.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 781

[6] D. Comer. Analysis of a Heuristic for Full Trie Minimization. *ACM Transactions on Database Systems*, 6(3):513–537, Sept. 1981.

[7] U. Ellermann and C. Benecke. Firewalls for ATM Networks. In *Proceedings of INFOSEC'COM*, 1998.

[8] D. Eppstein and S. Muthukrishnan. Internet Packet Filter Management and Rectangle Geometry. In *Proceedings of the Symposium on Discrete Algorithms*, pages 827–835, 2001.

[9] E. Fulp, Z. Fu, D. S. Reeves, S. F. Wu, and X. Zhang. Preventing Denial of Service Attacks on Quality of Service. In *Proceedings of the 2001 DARPA Information Survivability Conference and Exposition II*, 2001.

[10] R. Funke, A. Grote, and H.-U. Heiss. Performance Evaluation of Firewalls in Gigabit-Networks. In *Proceedings of the Symposium on Performance Evaluation of Computer and Telecommunication Systems*, 1999.

[11] X. Gan, T. Schroeder, S. Goddard, and B. Ramamurthy. LSMAC vs. SLNAT: Scalable ClusterbasedWeb Servers. *Journal of Networks, Software Tools, and Applications*, 3(3):175–185, 2000.

[12] S. Goddard, R. Kieckhafer, and Y. Zhang. An Unavailability Analysis of Firewall Sandwich Configurations. In *Proceedings of the 6th IEEE Symposium on High Assurance Systems Engineering*, 2001.

[13] A. Hari, S. Suri, and G. Parulkar. Detecting and Resolving Packet Filter Conflicts. In *Proceedings of IEEE INFOCOM*, pages 1203–1212, 2000.

[14] O. Paul and M. Laurent. A Full Bandwidth ATM Firewall. In *Proceedings of the 6th European Symposium on Research in Computer Security ESORICS'2000*, 2000.

[15] L. Qui, G. Varghese, and S. Suri. Fast Firewall Implementations for Software and Hardware-Based Routers. In *Proceedings of ACM SIGMETRICS*, June 2001.

[16] G. V. Rooij. Real Stateful TCP Packet Filtering in IP Filter. In *Proceedings of the 10th USENIX Security Symposium*, 2001.

[17] M. Schwartz. *Telecommunication Networks: Protocols, Modeling, and Analysis*. Addison-Wesley, 1987.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 782

[18] S. Suri and G. Varghese. Packet Filtering in High Speed Networks. In *Proceedings of the Symposium on Discrete Algorithms*, pages 969 – 970, 1999.

[19] P. Warkhede, S. Suri, and G. Varghese. Fast Packet Classification for Two-Dimensional Conflict-Free Filters. In *Proceedings of IEEE INFOCOM*, pages 1434–1443, 2001.

[20] J. Xu and M. Singhal. Design and Evaluation of a High-Performance ATM Firewall Switch and Its Applications. *IEEE Journal on Selected Areas in Communications*, 17(6):1190–1200, June 1999.

[21] R. L. Ziegler. *Linux Firewalls*. New Riders, second edition, 2002.

[22] E. D. Zwicky, S. Cooper, and D. B. Chapman. *Building Internet Firewalls*. O'Reilly, 2000.

[23] E. Al-Shaer and H. Hamed. Modeling and Management of Firewall Policies. *IEEE Transactions on Network and Service Management*, 1(1), 2004.

[24] S. M. Bellovin and W. Cheswick. Network Firewalls. *IEEE Communications Magazine*, pages 50–57, Sept. 1994.

[25] D. E. Culler and J. P. Singh. *Parallel Computer Architecture: A Hardware/Software Approach*. Morgan Kaufman, 1999.

[26] E. W. Fulp. Firewall Architectures for High Speed Networks. Technical Report 20026, Wake Forest University Computer Science Department, 2002.

[27] E. W. Fulp. Optimization of Firewall Policies Using Directed Acyclical Graphs. Submitted for publication, 2004.

[28] R. Rivest. On Self-Organizing Sequential Search Heuristics. *Communications of the ACM*, 19(2), 1976.

[29] W. E. Smith. Various Optimizers for Single-Stage Production. *Naval Research Logistics Quarterly*, 3:59 – 66, 1956.

It will be understood that various details of the present subject matter may be changed without departing from the scope of the present subject matter. Furthermore, the foregoing description is for the purpose of illustration only, and not for the purpose of limitation.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 783

CLAIMS

What is claimed is:

1. A function-parallel firewall comprising:
   (a) a first firewall node for filtering received packets using a first portion of a rule set including a plurality of rules, the first portion including less than all of the rules in the rule set; and
   (b) at least one second firewall node for filtering packets using a second portion of the rule set, the second portion including at least one rule in the rule set that is not present in the first portion, wherein the first and second portions together include all of the rules in the rule set.
2. The function-parallel firewall of claim 1 wherein the first and second portions respectively include rules present in first and second levels of a firewall hierarchy.
3. The function-parallel firewall of claim 2 wherein the first and second firewall nodes are connected to each other such that the rules in the first level are applied to received packets before the rules in the second level.
4. The function-parallel firewall of claim 3 wherein the first firewall node is adapted to forward packets that pass one of the rules in the first portion to an internal network in a manner that bypasses the second firewall node.
5. The function-parallel firewall of claim 3 wherein the first firewall node is adapted to discard packets that fail one of the rules in the first portion.
6. The function-parallel firewall of claim 3 wherein the rules in the second portion are more processor-intensive than the rules in the first portion.
7. The function-parallel firewall of claim 1 comprising:
   (a) a packet distributor for replicating each packet to the first and second firewall nodes; and
   (b) a gate machine for collecting results of applying the first and second portions of the rule set to each packet from the first and second firewall nodes and for making a forwarding decision based on the results.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 784

8. The function-parallel firewall of claim 7 wherein each of the first and second firewall nodes is adapted to process the packets in a pipelined manner and wherein the gate machine is adapted to maintain state information for multiple packets simultaneously.

9. The function-parallel firewall of claim 7 wherein the gate machine is adapted to signal the second firewall node when processing of a packet is no longer needed due to matching of the packet with one of the rules used by the first firewall node.

10. The function-parallel firewall of claim 1 wherein the first and second firewall nodes implement rules of the same hierarchical level.

11. The function-parallel firewall of claim 1 wherein at least one of the rules in the first portion is present in the second portion to provide redundancy.

12. The function-parallel firewall of claim 1 wherein the first and second firewall nodes are configured to implement a gateless design wherein the rules are distributed among the first and second firewall nodes such that for any given packet, one of the first and second firewall nodes will accept the packet and the other of the first and second firewall nodes will deny the packet.

13. The function-parallel firewall of claim 1 wherein the first and second firewall nodes are implemented on different physical machines.

14. The function-parallel firewall of claim 1 wherein the first and second firewall nodes are implemented on different processors in the same machine.

15. The function-parallel firewall of claim 1 wherein the first and second firewall nodes are adapted to simultaneously process copies of the same packets.

16. The function-parallel firewall of claim 1 wherein the rule set includes firewall rules for filtering packets based on packet headers.

17. The function-parallel firewall of claim 1 wherein the rule set includes intrusion detection rules for filtering packets based on packet payloads.

18. The function-parallel firewall of claim 1 wherein the rule set includes intrusion protection rules for filtering packets based on a combination of packet headers and packet payloads.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 785

19. A stateful function-parallel firewall system comprising:

(a) a first function-parallel firewall subsystem including a plurality of firewall nodes implementing a first rule set for filtering packets arriving in a network, wherein the first function-parallel firewall subsystem includes at least first and second firewall nodes respectively implementing first and second portions of the first rule set, the first portion including at least one rule that is not present in the second portion; and

(b) a second function-parallel firewall sub-system including a second set of firewall nodes implementing a second rule set for filtering packets departing from the network, wherein the second function-parallel firewall subsystem includes at least third and fourth firewall nodes, the third and fourth firewall nodes implementing first and second portions of the second rule set, the first portion of the second rule set including at least one rule that is not present in the second portion of the second rule set, wherein the first and second function-parallel firewall subsystems share state information regarding connections established through the first and second sets of firewall nodes.

20. The stateful function-parallel firewall system of claim 19 comprising a database node coupled to the first and second function-parallel firewall subsystems, wherein the database node stores the state information, wherein the first set of firewall nodes obtains state information regarding connections established through the second set of firewall nodes from the database node, and wherein the second set of firewall nodes obtains state information regarding connections established through the first set of firewall nodes from the database node.

21. The stateful function-parallel firewall system of claim 20 wherein the first and second sets of firewall nodes store copies of the state information maintained by the database node.

22. The stateful function-parallel firewall system of claim 19 wherein the first and second function-parallel firewall subsystems each include a gateless design wherein rules are distributed among individual firewall nodes

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 786

such that a received packet will be accepted by a single firewall node and will be denied by remaining firewall nodes.

23. The stateful function-parallel firewall system of claim 19 wherein the firewall nodes of the first subsystem are adapted to simultaneously process copies of the same packets arriving in the network and wherein the firewall nodes in the second subsystem are adapted to simultaneously process copies of packets departing from the network.

24. The stateful function-parallel firewall system of claim 19 wherein the first and second rule sets include firewall rules for filtering packets based on packet header information.

25. The stateful function-parallel firewall system of claim 19 wherein the first and second rule sets include intrusion detection rules for filtering packets based on payloads of the packets.

26. The stateful function-parallel firewall system of claim 19 wherein the first and second rule sets include intrusion protection rules for filtering packets based on a combination of packet header and packet payload information.

27. A firewall grid comprising:

(a) a plurality of firewall nodes being physically connected to each other via a network for filtering packets; and

(b) a controller for controlling logical connections between the firewall nodes, wherein the controller is adapted to configure the firewall nodes in a logical firewall hierarchy without changing physical connections between the firewall nodes.

28. The firewall grid of claim 27 wherein the firewall nodes are adapted to simultaneously process copies of the same packets.

29. The firewall grid of claim 27 wherein the logical firewall hierarchy comprises a function-parallel hierarchy where individual firewall nodes implement different portions of a rule set.

30. The firewall grid of claim 27 wherein the controller is adapted to detect congestion at a first level in the firewall hierarchy and is adapted to reconfigure the logical connections between the firewall nodes to add a firewall node from another level in the firewall hierarchy to the level at which congestion was detected.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 787

31. A system for providing network access control based on a function-parallel policy, the system comprising:
    (a) a first node for filtering received packets using a first portion of a rule set including a plurality of rules, the first portion including less than all of the rules in the rule set; and
    (b) at least one second node for filtering packets using a second portion of the rule, the second portion including at least one rule in the rule set that is not present in the first portion, wherein the first and second portions together include all of the rules in the rule set.
32. The system of claim 31 wherein the rule set includes firewall rules for filtering packets based on packet header information.
33. The system of claim 31 wherein the rule set includes intrusion detection rules for filtering packets based on packet payload information.
34. The system of claim 31 wherein the rule set includes intrusion protection rules for filtering packets based on a combination of packet header and packet payload information.
35. A method for controlling access to a network based on a set of packet filtering rules distributed in a function-parallel manner, the method comprising:
    (a) distributing packet filtering rules of a rule set among a plurality of different nodes in a function-parallel manner so that at least some of the different nodes implement different portions of the rule set;
    (b) replicating packets to each of the nodes; and
    (c) applying the rules to filter the packets.
36. The method of claim 35 wherein the packet filtering rules comprise firewall rules for filtering packets based on packet header information.
37. The method of claim 35 wherein the packet filtering rules comprise intrusion detection rules for filtering packets based on packet payload information.
38. The method of claim 35 wherein the packet filtering rules comprise intrusion protection rules for filtering packets based on a combination of packet header and packet payload information.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 788

39. The method of claim 35 comprising, receiving a first packet at a first node, determining whether the packet matches a first rule implemented by the first node, and, in response to determining that a match occurs, notifying the other nodes of the match.

40. The method of claim 36 comprising, at the other nodes, ceasing processing of the first packet.

41. A method for distributing rules in a function-parallel firewall, the method comprising:

(a) defining a rule set for a function-parallel firewall;

(b) assigning rules in the rule set to nodes and branches in a trie data structure, wherein each node in the trie data structure corresponds to a data field and each branch represents a value for each data field to be compared to values in corresponding fields in received packets; and

(c) pruning the trie data structure in a manner that preserves ordering of the rules in the rule set.

42. The method of claim 41 wherein pruning the trie data structure includes pushing a rule to its descendants and deleting the original rule.

43. The method of claim 41 comprising reducing the depth of the trie using level reordering.

44. The method of claim 41 wherein assigning rules in the rule set to nodes and branches in the trie data structure includes assigning rules requiring longer processing time to lower levels in the trie.

45. A method for assigning rules to a plurality of firewall nodes, the method comprising:

(a) assigning packet filtering rules to nodes in a directed acyclical graph (DAG);

(b) representing relationships between rules by edges in the DAG; and

(c) distributing the rules among a plurality of firewall nodes using the DAG such that different firewall nodes implement different rules and such that relationships between the rules specified in the DAG are preserved.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 789

46. The method of claim 45 wherein distributing the rules among a plurality of firewall nodes using the DAG includes distributing the rules such that intersecting accept rules are not present in more than one firewall node.

47. The method of claim 45 wherein distributing the rules among a plurality of firewall nodes using the DAG includes distributing the rules such that the edges to not span multiple firewall nodes.

48. The method of claim 45 wherein distributing the rules among a plurality of firewall nodes using the DAG includes replicating deny rules in each of the firewall nodes.

49. The method of claim 48 comprising eliminating a deny rule implemented by a first firewall node in response to determining that the deny rule for the first firewall node includes no inbound edges and outbound edges from the first deny rule terminate only in a superset deny rule.

50. A computer program product comprising computer executable instructions embodied in a computer readable medium for performing steps comprising:

(a) distributing packet filtering rules of a rule set among a plurality of different nodes in a function-parallel manner so that at least some of the different nodes implement different portions of the rule set;

(b) replicating packets to each of the nodes; and

(c) applying the rules to filter the packets.

51. A computer program product comprising computer executable instructions embodied in a computer readable medium for performing steps comprising:

(a) defining a rule set for a function-parallel firewall;

(b) assigning rules in the rule set to nodes and branches in a trie data structure, wherein each node in the trie data structure corresponds to a data field and each branch represents a value for each data field to be compared to values in corresponding fields in received packets; and

(c) pruning the trie data structure in a manner that preserves ordering of the rules in the rule set.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 790

52. A computer program product comprising computer executable instructions embodied in a computer readable medium for performing steps comprising:

(a) assigning rules to nodes in a directed acyclical graph (DAG);

(b) representing relationships between rules by edges in the DAG; and

(c) distributing the rules among a plurality of firewall nodes using the DAG such that the relationships between the rules specified in the DAG are preserved.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 791

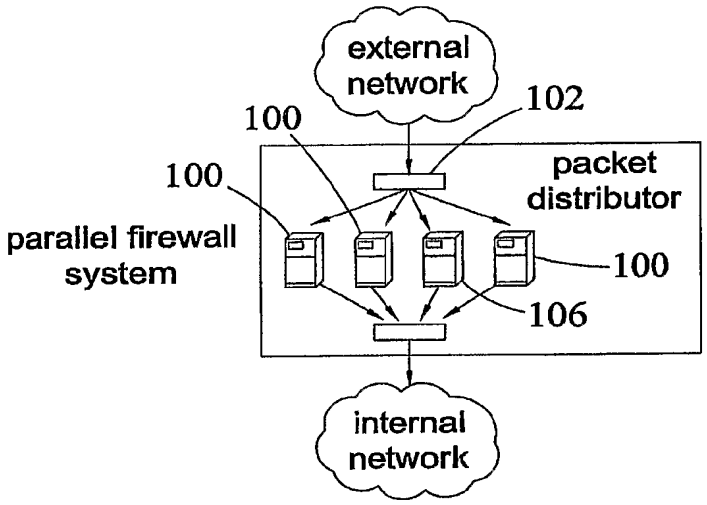

FIG. 1
(PRIOR ART)

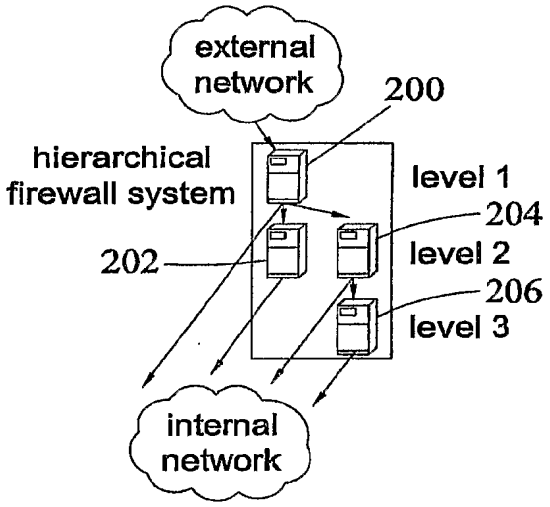
external network
200
hierarchical firewall system
level 1
204
level 2
202
206
level 3
internal network

FIG. 2

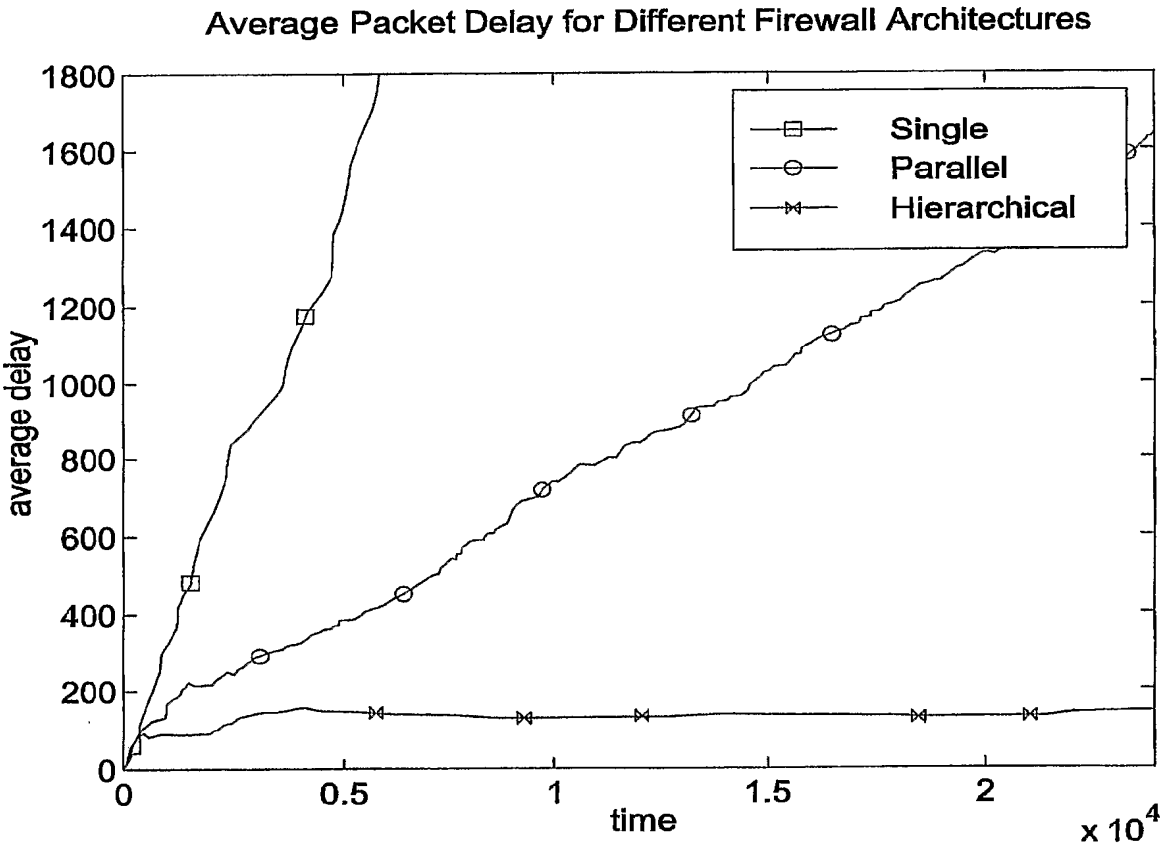
Average Packet Delay for Different Firewall Architectures
average delay
time
x 10^4
Single
Parallel
Hierarchical
0
200
400
600
800
1000
1200
1400
1600
1800
0
0.5
1
1.5
2

FIG. 3A

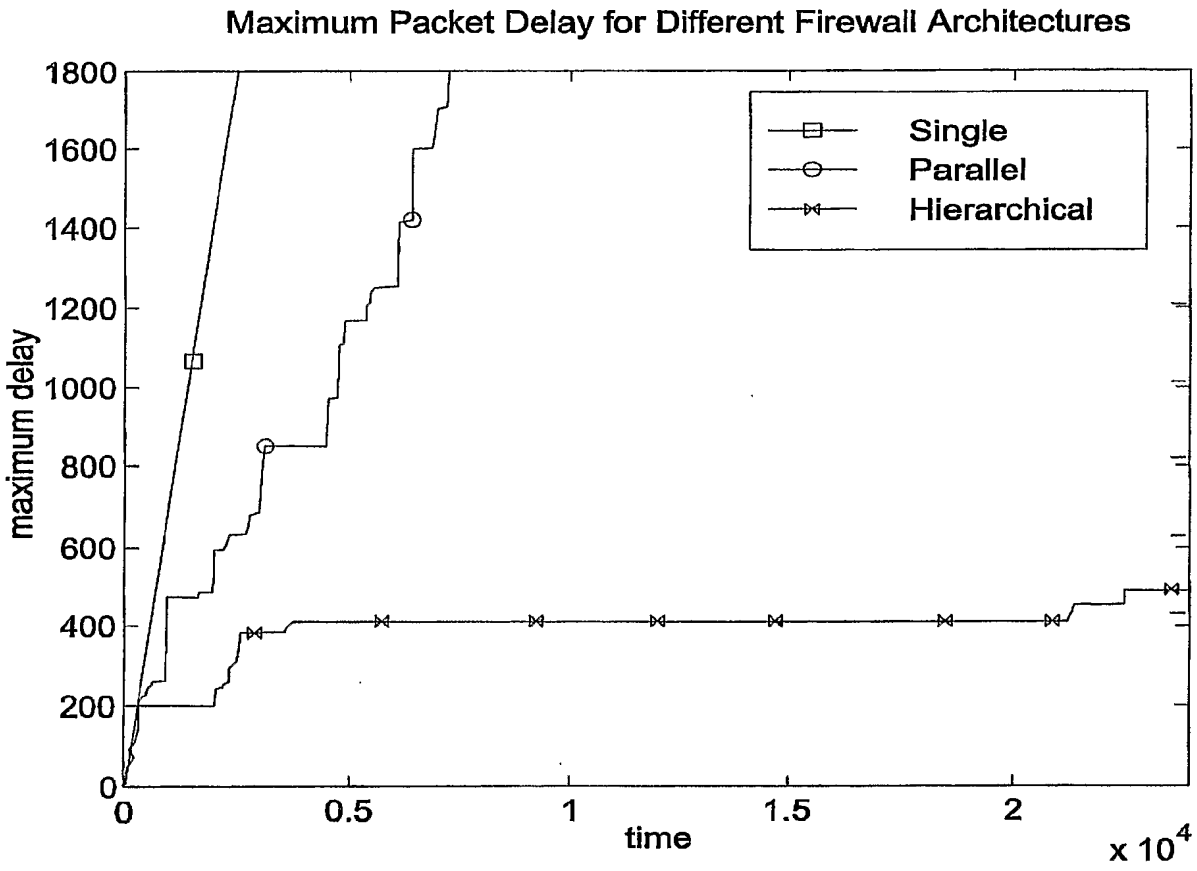
Maximum Packet Delay for Different Firewall Architectures
maximum delay
time
x 10^4
Single
Parallel
Hierarchical
0
200
400
600
800
1000
1200
1400
1600
1800
0
0.5
1
1.5
2

FIG. 3B

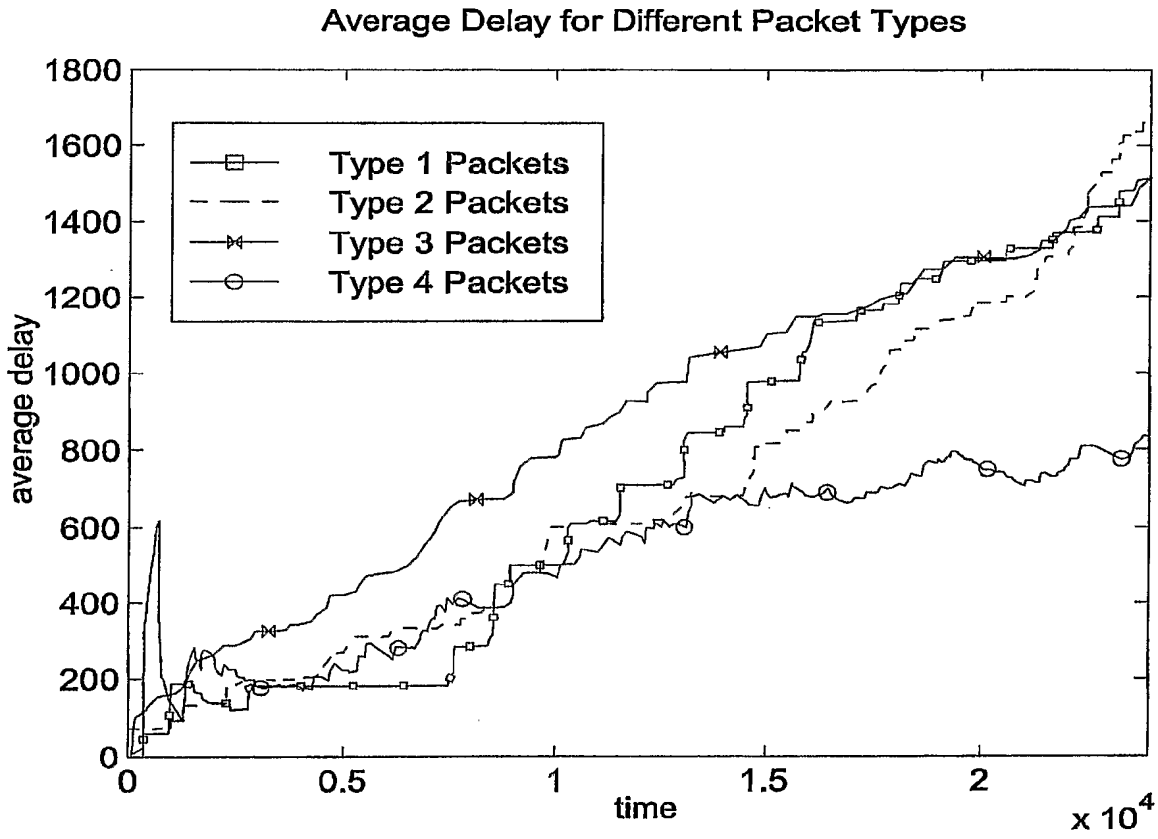
Average Delay for Different Packet Types
Type 1 Packets
Type 2 Packets
Type 3 Packets
Type 4 Packets
average delay
time
x 10^4
0
200
400
600
800
1000
1200
1400
1600
1800
0.5
1
1.5
2

FIG. 4A

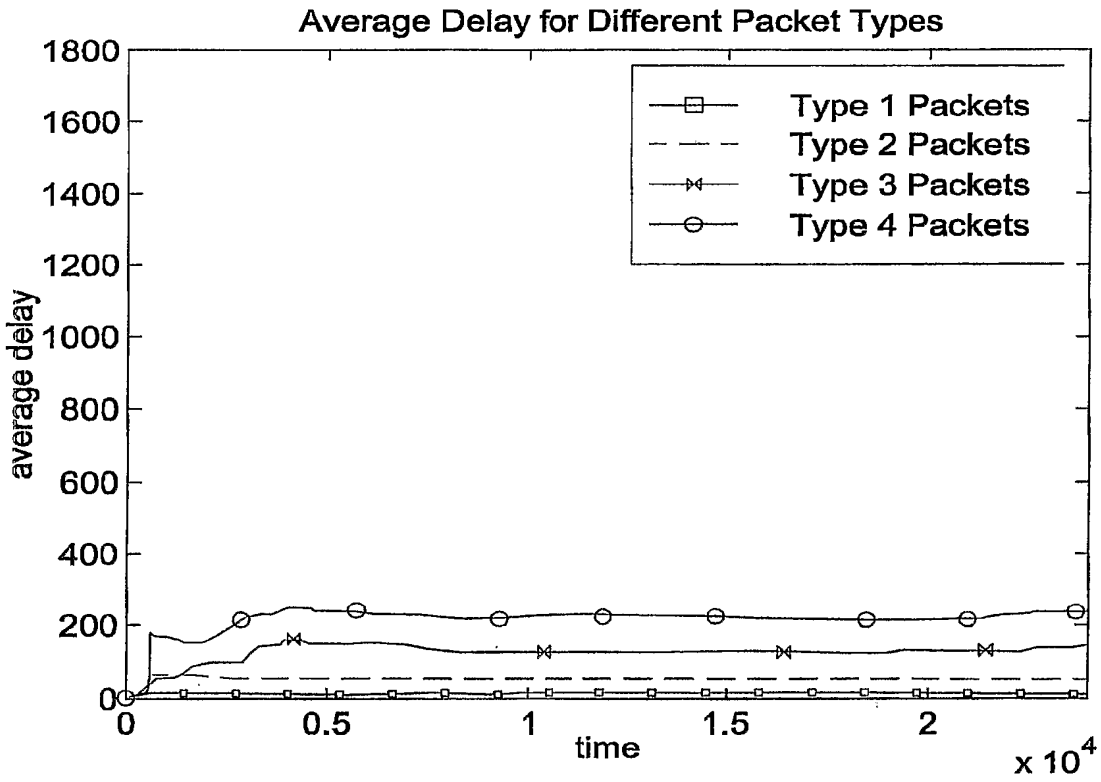
Average Delay for Different Packet Types
Type 1 Packets
Type 2 Packets
Type 3 Packets
Type 4 Packets
average delay
1800
1600
1400
1200
1000
800
600
400
200
0
0
0.5
1
1.5
2
time
x 10^4

FIG. 4B

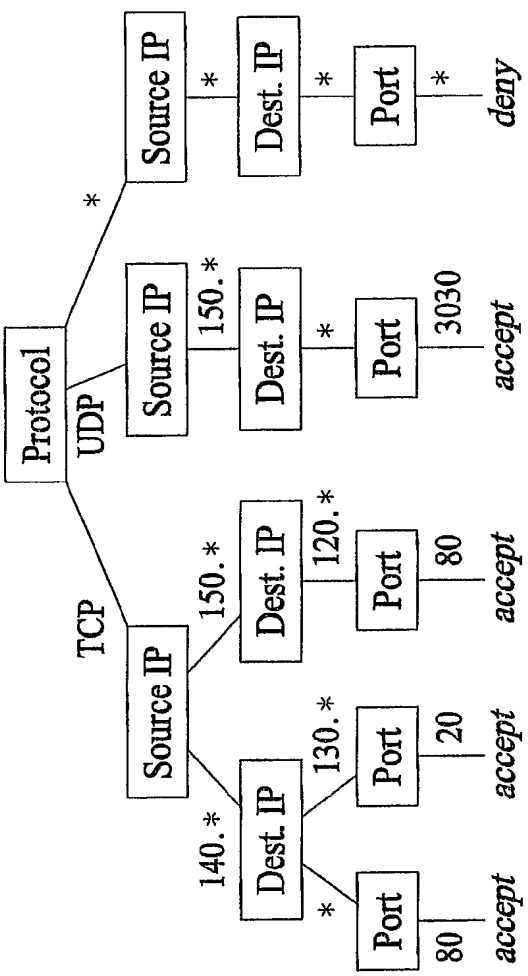
Protocol
TCP
UDP
Source IP
140.*
150.*
*
Dest. IP
130.*
120.*
150.*
*
Port
20
80
80
3030
*
accept
deny

FIG. 5

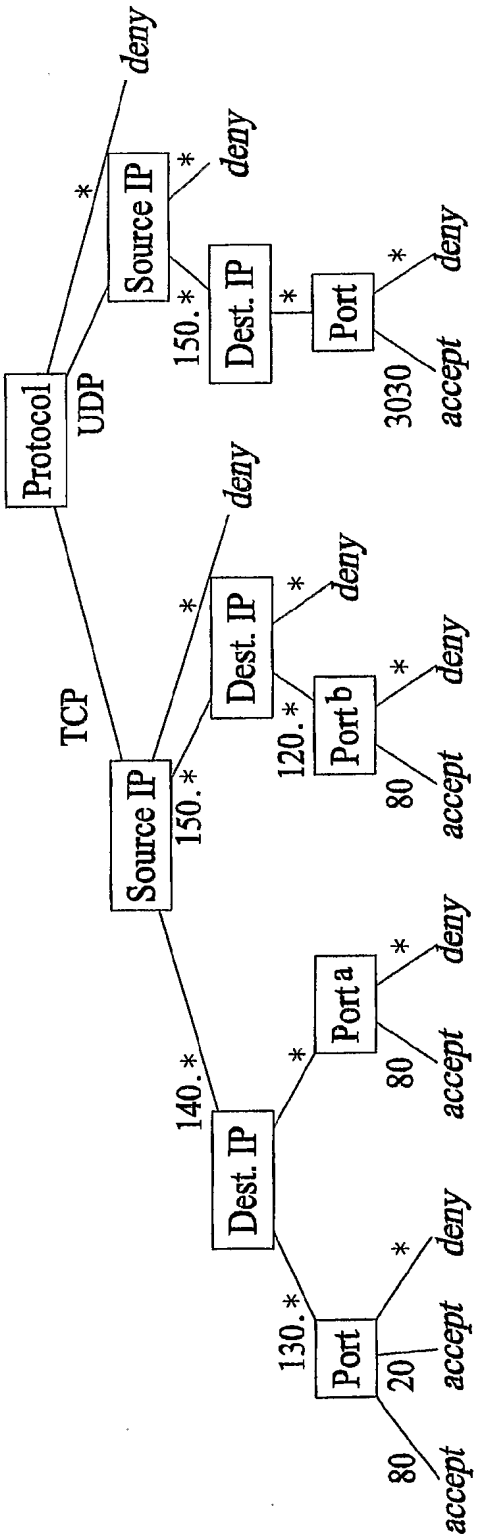
Protocol
TCP
UDP
Source IP
140.*
150.*
*
Dest. IP
130.*
120.*
150.*
*
Port a
Port b
Port
20
80
80
3030
*
accept
deny

FIG. 6

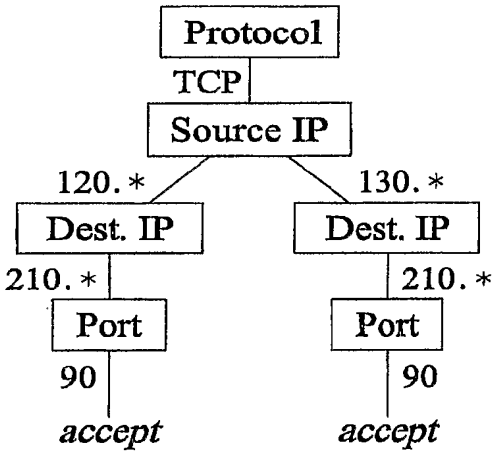
Protocol
TCP
Source IP
120.*
130.*
Dest. IP
Dest. IP
210.*
210.*
Port
Port
90
90
accept
accept

FIG. 7A

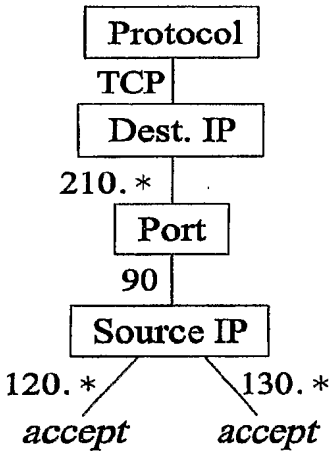
Protocol
TCP
Dest. IP
210.*
Port
90
Source IP
120.*
130.*
accept
accept

FIG. 7B

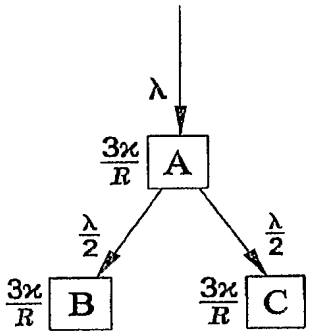
λ
3κ/R
A
λ/2
λ/2
3κ/R
B
3κ/R
C

FIG. 8A

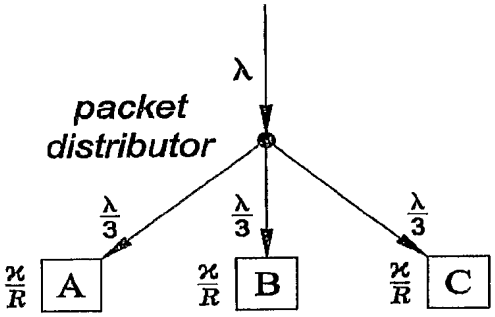
λ
packet
distributor
λ/3
λ/3
λ/3
κ/R
A
κ/R
B
κ/R
C

FIG.8B

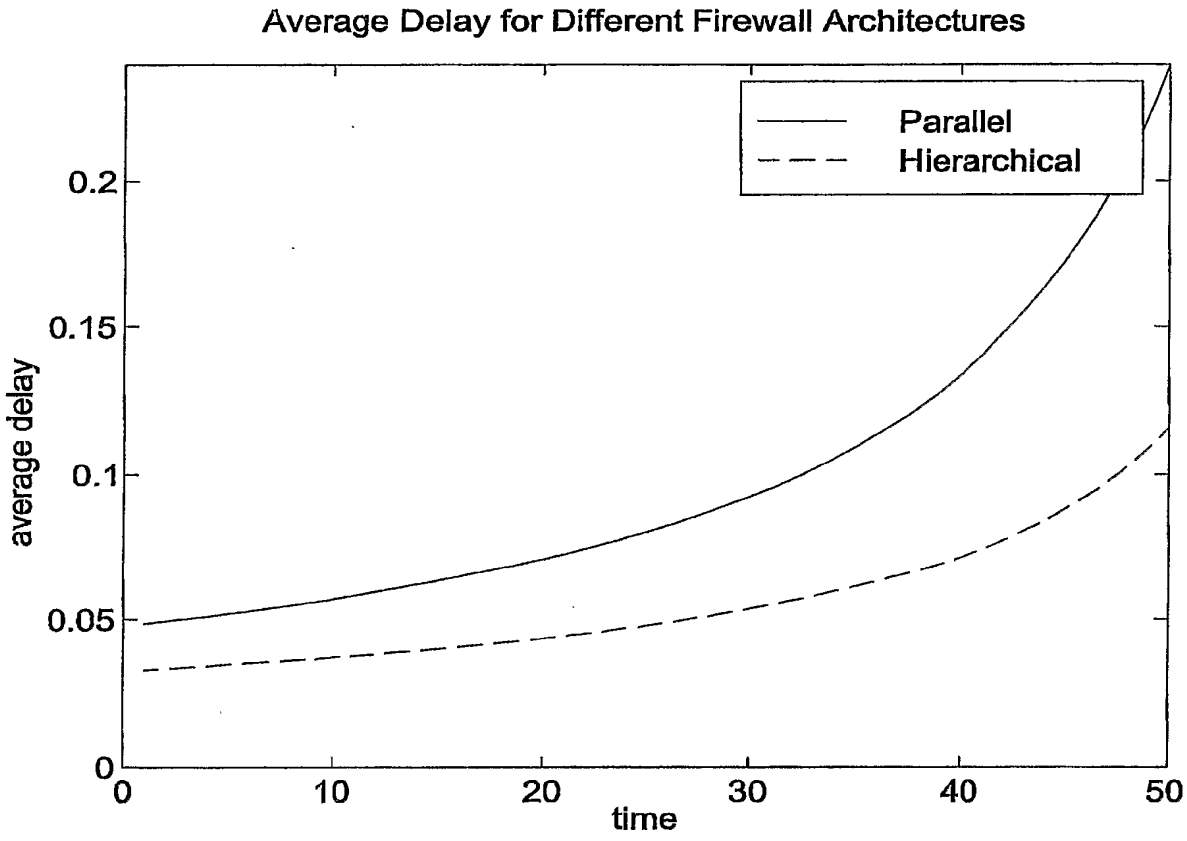
Average Delay for Different Firewall Architectures
Parallel
Hierarchical
average delay
time
0
0.05
0.1
0.15
0.2
0
10
20
30
40
50

FIG. 9

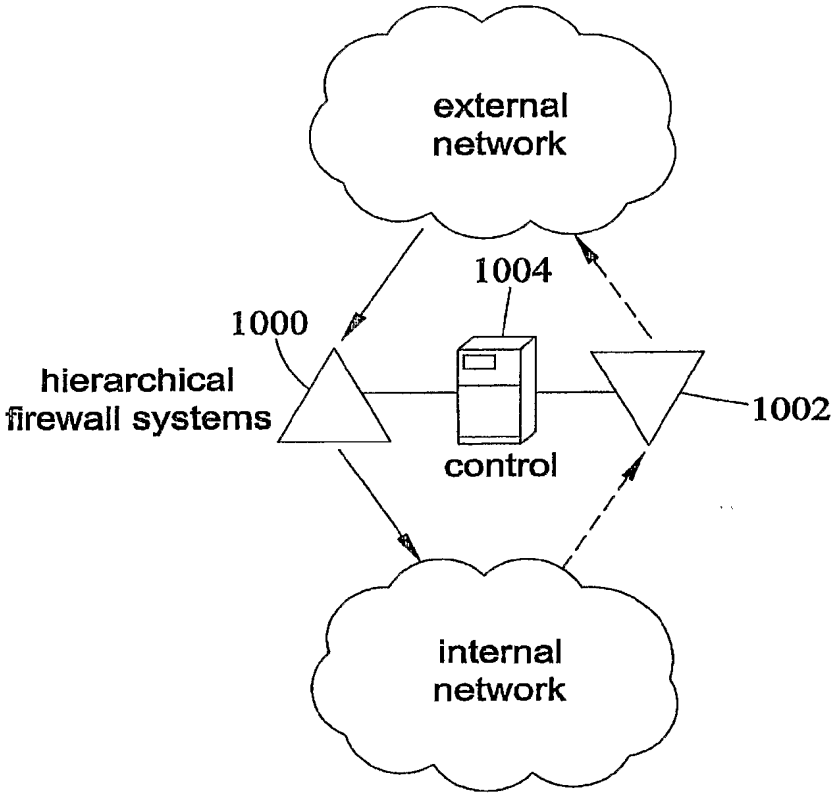
external network
1004
1000
hierarchical firewall systems
1002
control
internal network

FIG. 10

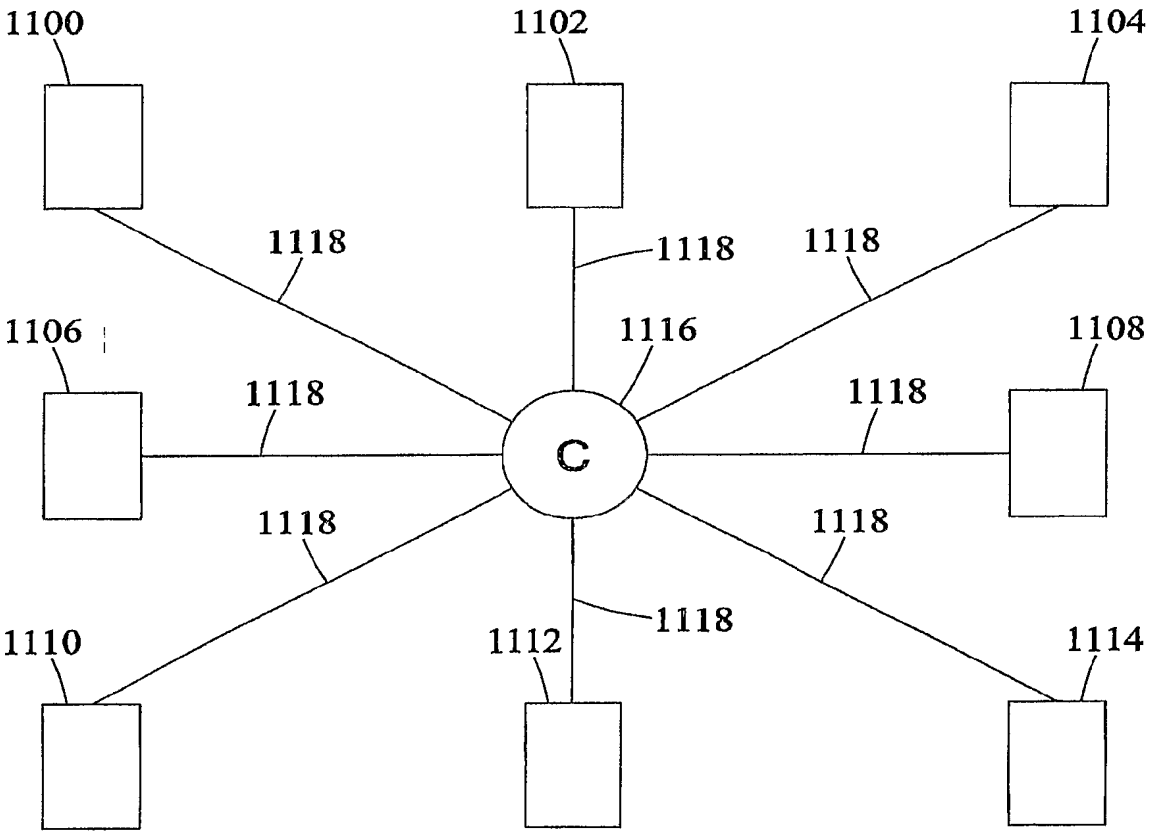
1100
1102
1104
1118
1118
1118
1106
1116
1108
1118
1118
C
1118
1118
1118
1110
1112
1114

FIG. 11A

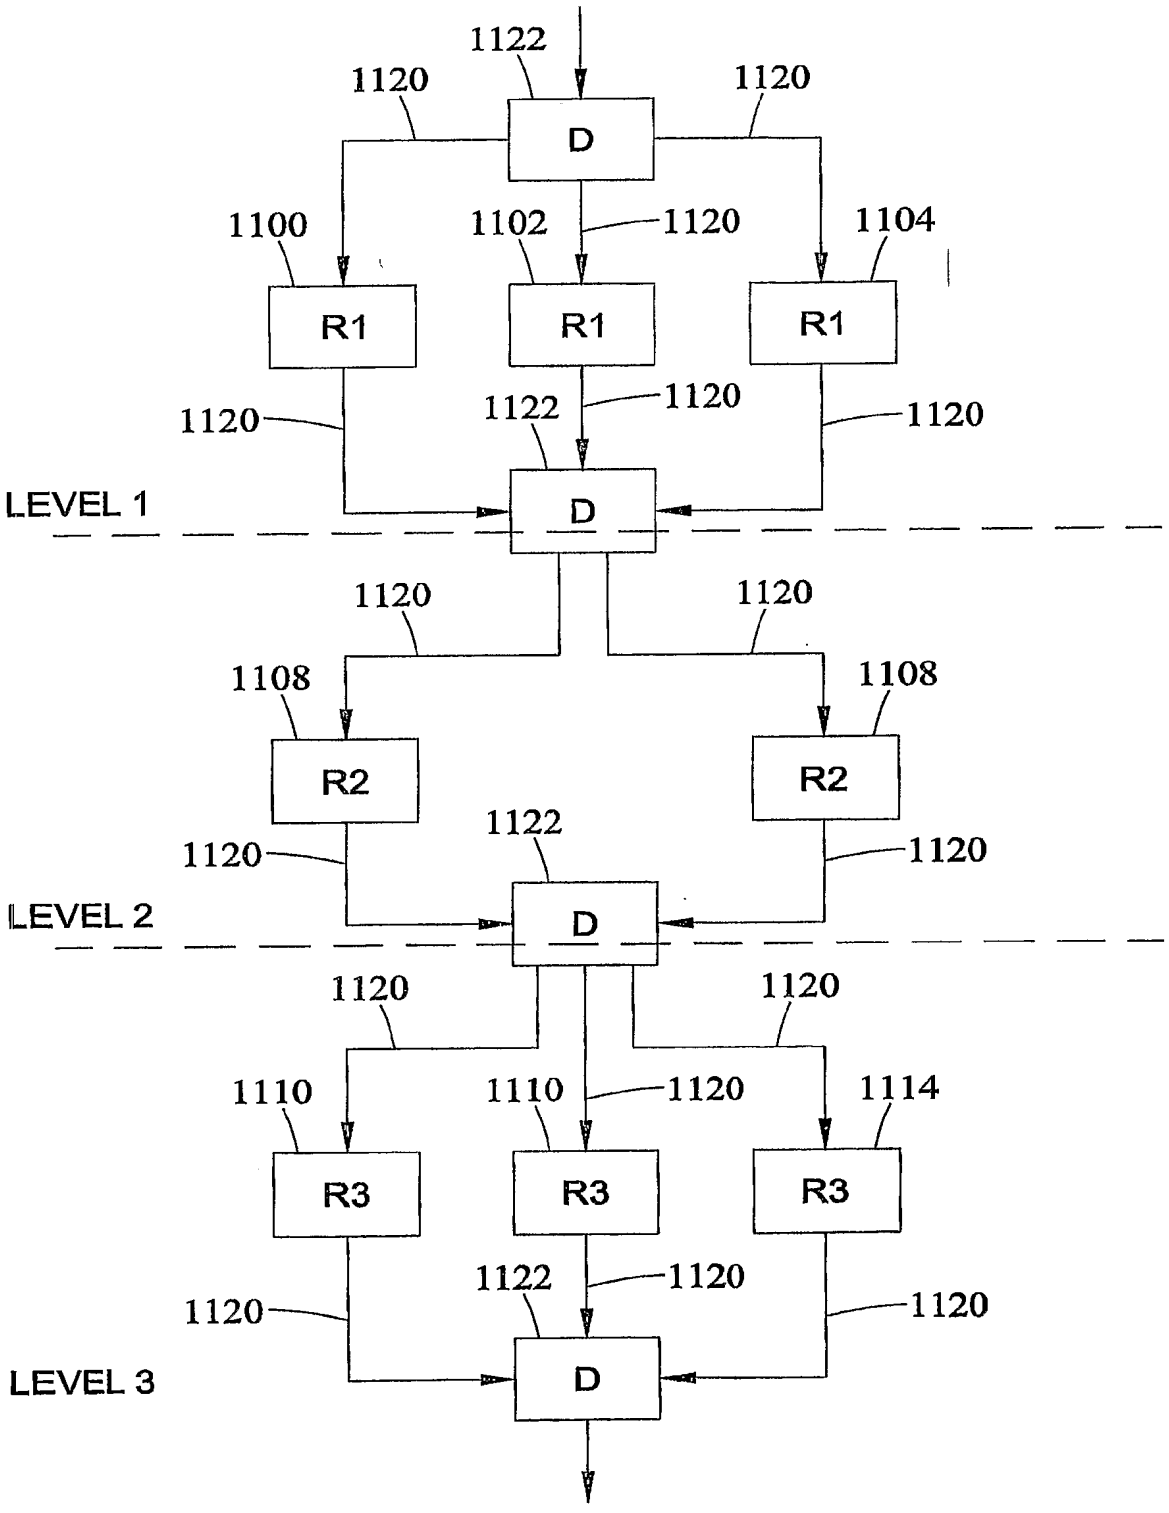
1122
1120
1120
D
1100
1102
1120
1104
R1
R1
R1
1120
1122
1120
1120
LEVEL 1
D
1120
1120
1108
1108
R2
R2
1122
1120
1120
LEVEL 2
D
1120
1120
1110
1110
1120
1114
R3
R3
R3
1122
1120
1120
1120
LEVEL 3
D

FIG. 11B

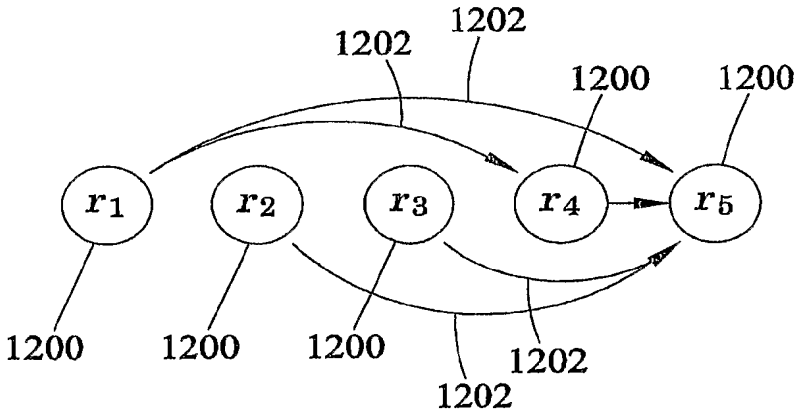

FIG. 12

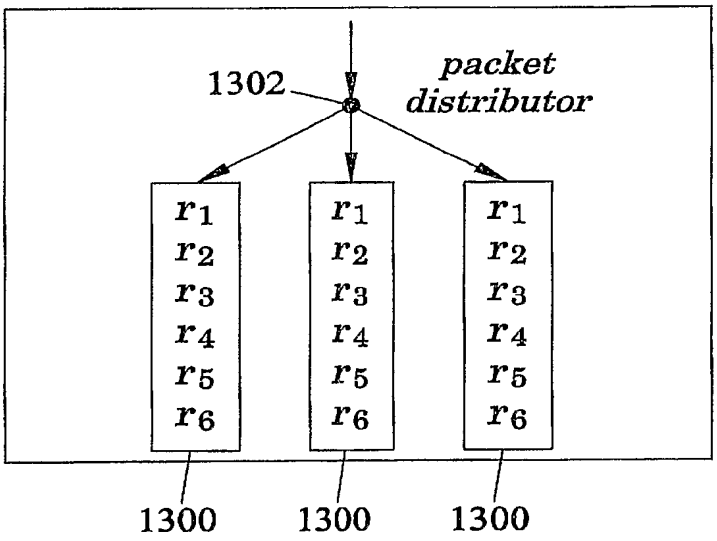

FIG. 13A

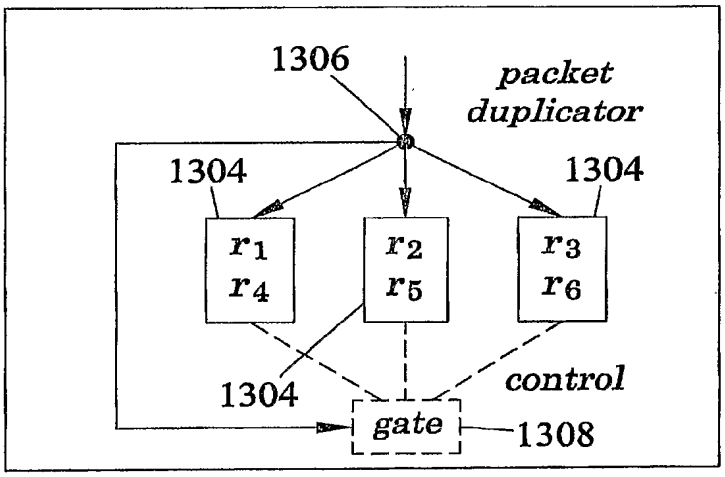

FIG. 13B

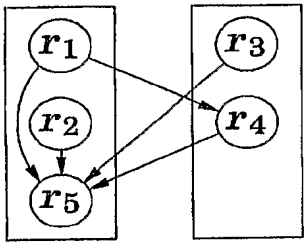

FIG. 14A

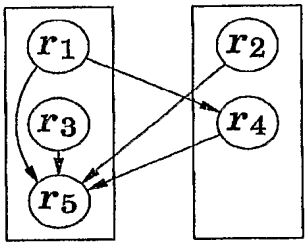

FIG. 14B

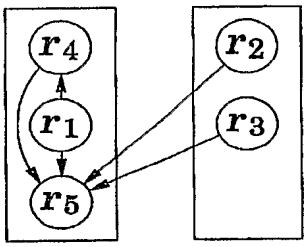

FIG. 14C

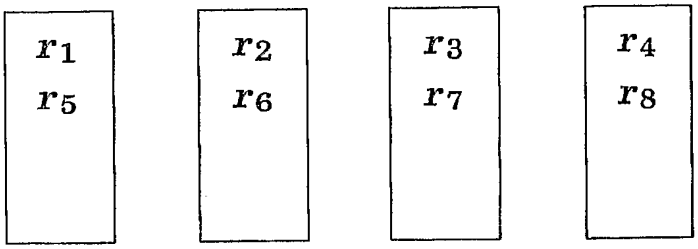

FIG. 15A

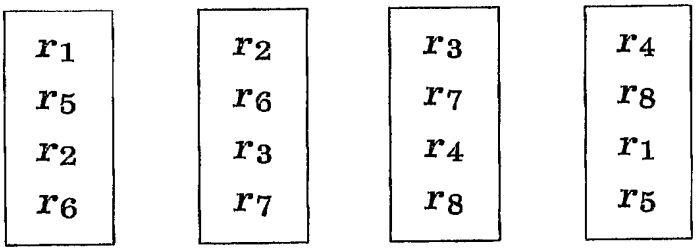

FIG. 15B

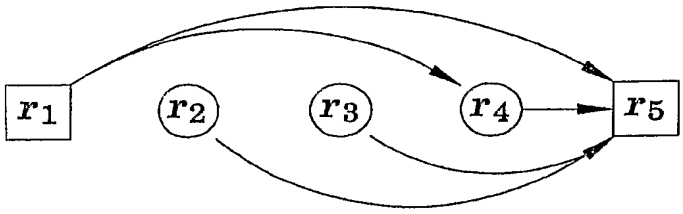

FIG. 16A

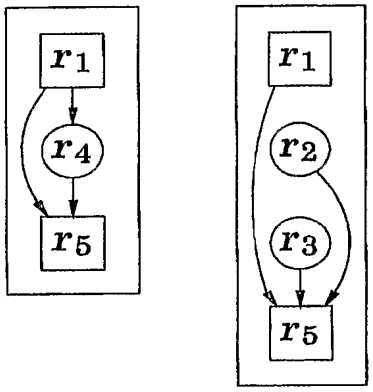

FIG. 16B

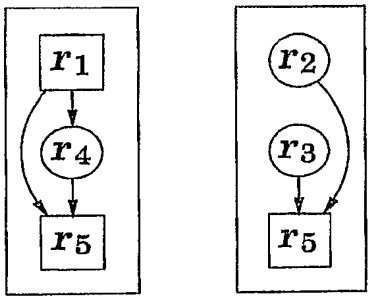

FIG. 16C

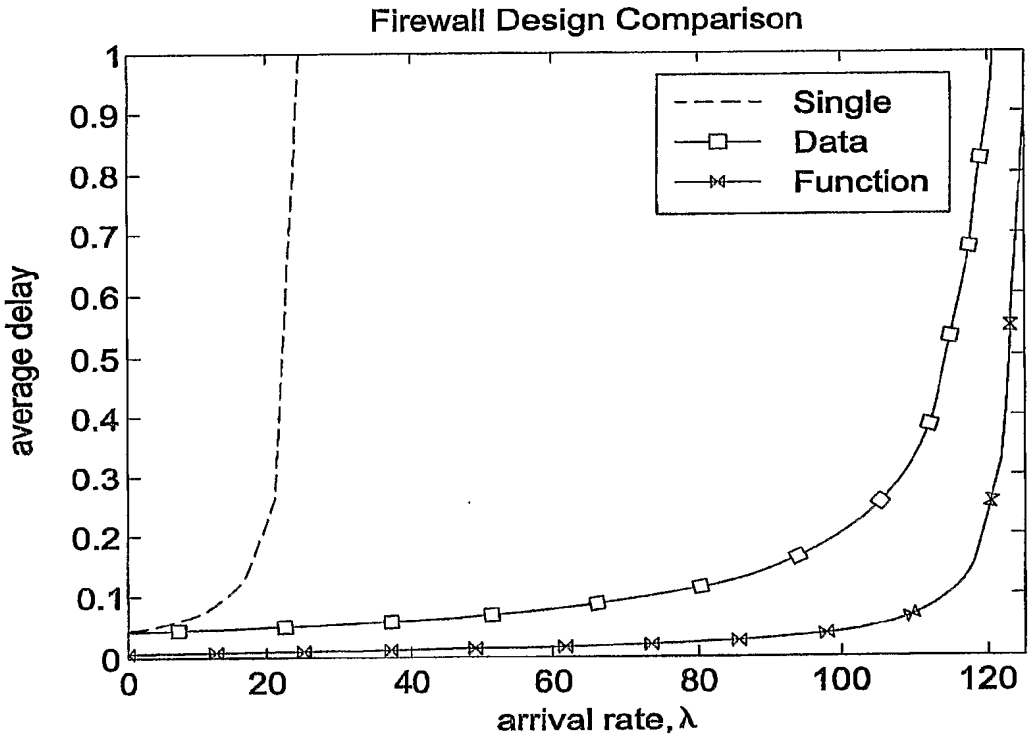
Firewall Design Comparison
average delay
arrival rate, λ
Single
Data
Function
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
20
40
60
80
100
120

FIG. 17A

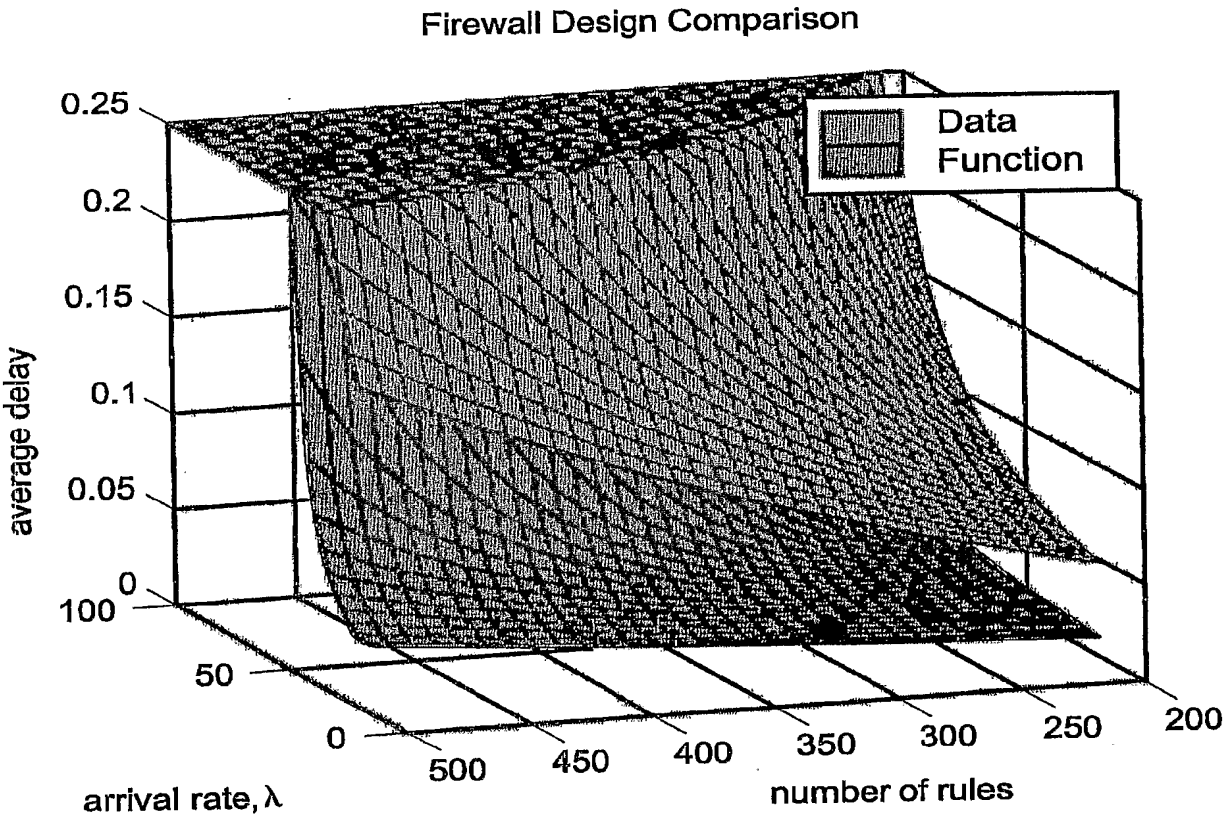
Firewall Design Comparison
Data
Function
0.25
0.2
0.15
0.1
0.05
0
average delay
100
50
0
arrival rate, λ
500
450
400
350
300
250
200
number of rules

FIG. 17B

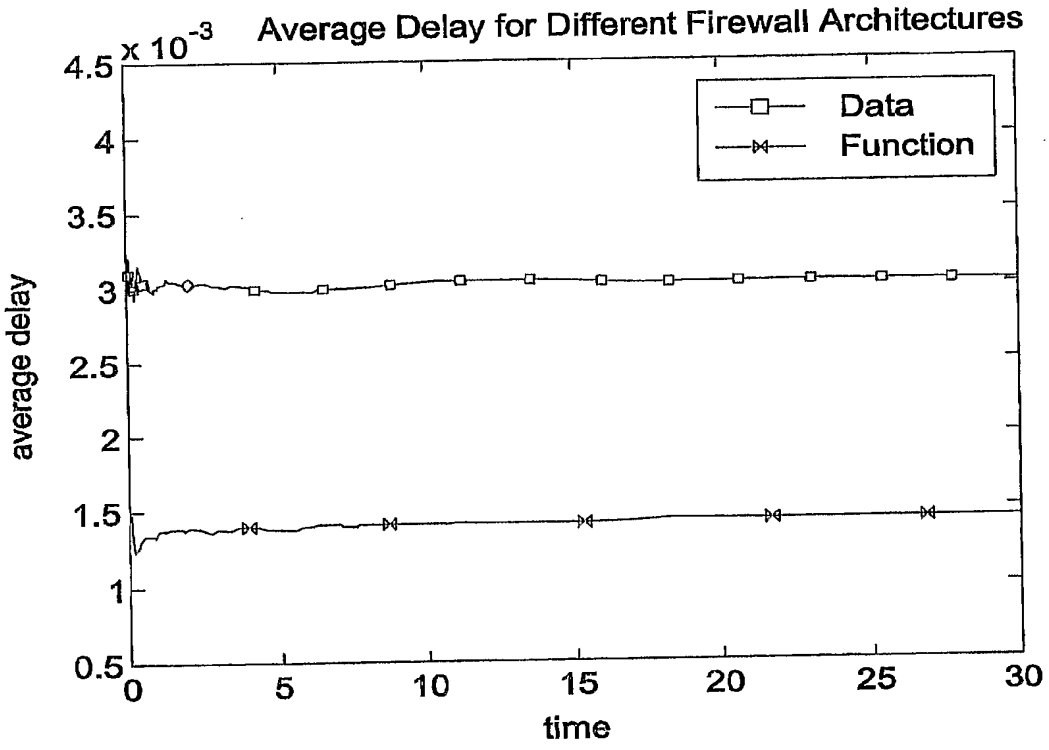
Average Delay for Different Firewall Architectures
x 10^-3
4.5
4
3.5
3
2.5
2
1.5
1
0.5
average delay
0
5
10
15
20
25
30
time
Data
Function

FIG. 18A

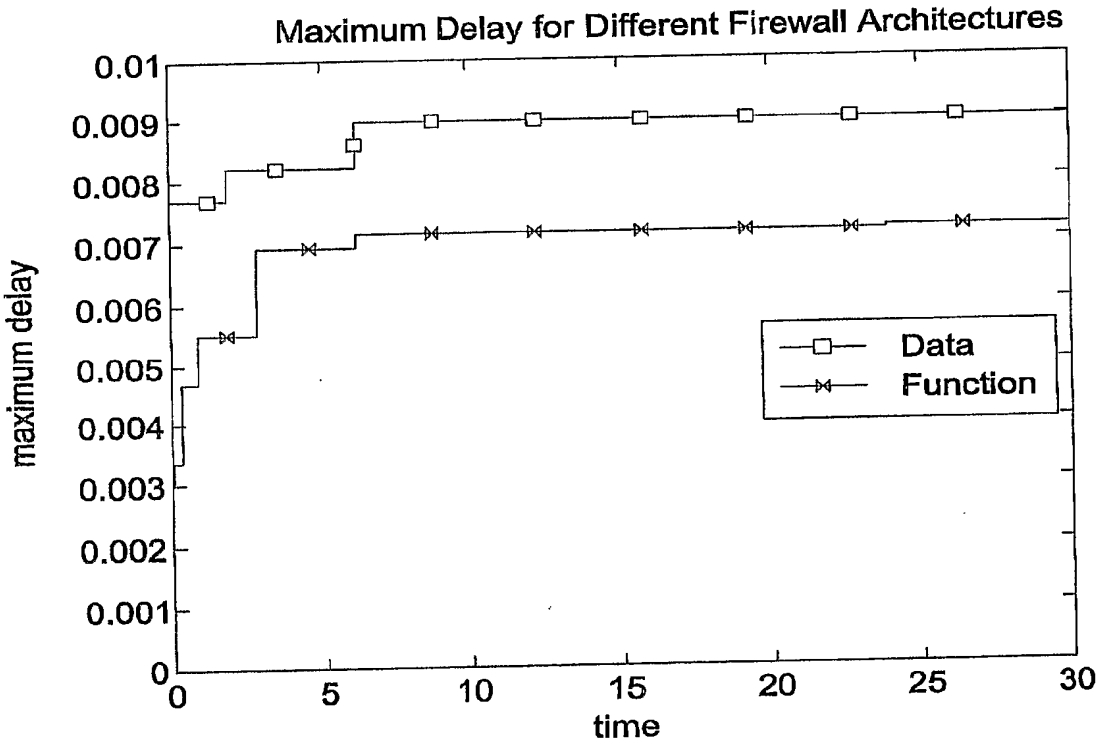
Maximum Delay for Different Firewall Architectures
maximum delay
time
Data
Function
0.01
0.009
0.008
0.007
0.006
0.005
0.004
0.003
0.002
0.001
0
0
5
10
15
20
25
30

FIG. 18B

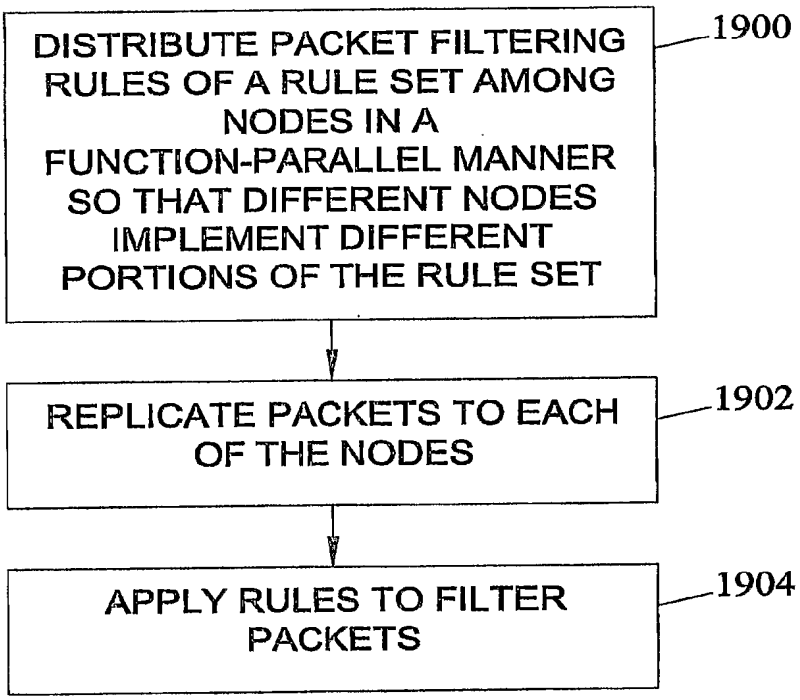
DISTRIBUTE PACKET FILTERING RULES OF A RULE SET AMONG NODES IN A FUNCTION-PARALLEL MANNER SO THAT DIFFERENT NODES IMPLEMENT DIFFERENT PORTIONS OF THE RULE SET
1900
REPLICATE PACKETS TO EACH OF THE NODES
1902
APPLY RULES TO FILTER PACKETS
1904

FIG. 19

Australian Government

IP Australia

LETTERS PATENT

STANDARD PATENT

2005328336

I, Victor Portelli, the Commissioner of Patents, grant a Standard Patent with the following particulars:

Name and Address of Patentee(s):

Wake Forest University
Office of Technology Asset Management, Wake Forest University Health Services, 200 East First Street, Suite 101, Winston Salem, North Carolina, 27101, United States of America

Name of Actual Inventor(s):

Fulp, Errin W. and Farley, Ryan J..

Title of Invention:

Method, systems, and computer program products for implementing function-parallel network firewall

Term of Letters Patent:

Twenty years from 22 December 2005

Priority Details :

Number	*Date*	*Filed with*
60/638,436	22 December 2004	US

COMMONWEALTH OF AUSTRALIA
AUSTRALIA
PATENT OFFICE

Dated this 5th day of January 2012

Victor Portelli
Commissioner of Patents

PATENTS ACT 1990

(12) STANDARD PATENT (11) Application No. **AU 2005328336 B2**
(19) AUSTRALIAN PATENT OFFICE

(54) Title
Method, systems, and computer program products for implementing function-parallel network firewall

(51) International Patent Classification(s)
G06F 17/00 (2006.01)

(21) Application No: **2005328336** (22) Date of Filing: **2005.12.22**

(87) WIPO No: **WO06/093557**

(30) Priority Data

(31) Number	(32) Date	(33) Country
60/638,436	**2004.12.22**	**US**

(43) Publication Date: **2006.09.08**
(44) Accepted Journal Date: **2011.09.15**

(71) Applicant(s)
Wake Forest University

(72) Inventor(s)
Fulp, Errin W.;Farley, Ryan J.

(74) Agent / Attorney
Griffith Hack, Level 10 167 Eagle Street, Brisbane, QLD, 4001

(56) Related Art
US 2002/0038339 A1

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
8 September 2006 (08.09.2006)

PCT

(10) International Publication Number
WO 2006/093557 A3

(51) International Patent Classification:
G06F 17/00 (2006.01)

(21) International Application Number:
PCT/US2005/047008

(22) International Filing Date:
22 December 2005 (22.12.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/638,436 22 December 2004 (22.12.2004) US

(71) Applicant *(for all designated States except US)*: **WAKE FOREST UNIVERSITY** [US/US]; OFFICE OF TECHNOLOGY ASSET MANAGEMENT, WAKE FOREST UNIVERSITY HEALTH SCIENCES, 200 East First Street, Suite 101, Winston Salem, North Carolina 27101 (US).

(72) Inventors; and
(75) Inventors/Applicants *(for US only)*: **FULP, Errin, W.** [US/US]; 308 N. Hawthorne Road, Winston-salem, North Carolina 27104 (US). **FARLEY, Ryan, J.** [US/US]; 212 Radford Street, Winston-salem, North Carolina 27106 (US).

(74) Agent: **HUNT, Gregory, A.**; JENKINS, WILSON & TAYLOR, P.A., SUITE 1400 UNIVERSITY TOWER, 3100 Tower Boulevard, Durham, North Carolina 27707 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*

(88) Date of publication of the international search report:
18 January 2007

[Continued on next page]

(54) Title: METHOD, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR IMPLEMENTING FUNCTION-PARALLEL NETWORK FIREWALL

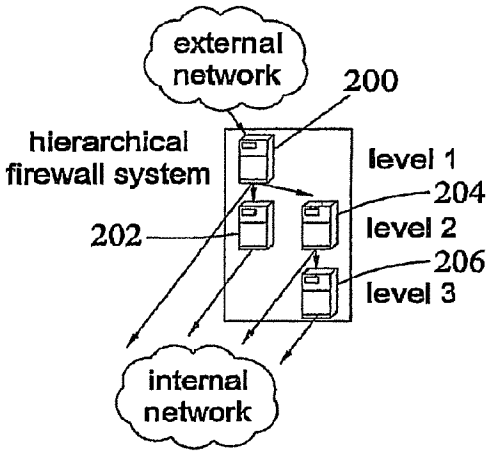

(57) Abstract: Methods, systems, and computer program products for providing function-paralell firewalls are disclosed. According to one aspect, a function-parallel firewall includes a first firewall node for filtering received packets using a first portion of a rule set including a plurality of rules. The first portion includes less than all of the rules in the rule set. At least one second firewall node filters packets using a second portion of the rule set. The second portion includes at least one rule in the rule set that is not present in the first portion. The first and second portions together include all of the rules in the rule set.

WO 2006/093557 A3

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

DESCRIPTION

METHOD, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR IMPLEMENTING FUNCTION-PARALLEL NETWORK FIREWALL

RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Patent Application Serial No. 60/638,436, filed December 22, 2004; the disclosure of which is incorporated herein by reference in its entirety.

GOVERNMENT INTEREST

This invention was made with Government support under Grant No. DE-FG02-03ER25581 awarded by U.S. Department of Energy, Mathematical and Computational Information Sciences Division. The U.S. Government has certain rights in the invention.

TECHNICAL FIELD

The subject matter described herein relates to implementing network firewalls. More particularly, the subject matter described herein relates to methods, systems, and computer program products for implementing a function-parallel network firewall.

BACKGROUND ART

Firewalls remain the frontier defense for securing networks that are vital to private industry, government agencies, and the military. They enforce a security policy by inspecting and filtering traffic arriving at or departing from a secure network [3, 21, 22]. Meaningful inspections often involve a complex process, and if firewalls are to remain effective, they must adapt to constantly changing network demands, technology, and security threats.

As firewall technology evolves to meet new demands and threats, it must continue to act transparently to legitimate users, with little or no effect on the perceived network performance. The networks between two communicating machines should be invisible, especially if traffic requires specific network Quality of Service (QoS), such as bounds on the packet delay, jitter, and throughput. The firewall should process the legitimate traffic quickly, efficiently,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 819

and never violate the desired QoS. Unfortunately, the firewall can easily become a bottleneck, given increasing traffic loads and network speeds [12, 14, 18, 20, 21]. Packets must be inspected and compared against complex rule sets and tables, which is a time-consuming process. In addition, audit files must be updated with current connection information. As a result, current firewalls have difficulty maintaining QoS guarantees. Since a firewall can be overwhelmed with traffic, it is also susceptible to Denial of Service (DoS) attacks [21]. Such attacks merely overload/saturate the firewall with illegitimate traffic. Audit files and state tables quickly fill available storage space and legitimate traffic queued behind suspicious traffic encounters long delays. Legitimate users notice poor network performance, and, in the worst case, the secure network is disconnected from the outside world.

Inspecting traffic sent between networks, a firewall provides access control, auditing, and traffic control based on a security policy [3, 21, 22]. The security policy is a list of ordered rules, as seen in Table 1, that defines the action to perform on matching packets. A rule r can be viewed as a k-tuple $(r[1], r[2], r[3], ..., r[k])$, where $r[i]$ is specified as a variable length prefix.

No.	Proto.	Source		Destination		Action
		IP	Port	IP	Port	
1	TCP	140.*	*	*	80	accept
2	TCP	150.*	*	120	80	accept
3	TCP	140.*	*	130	20	accept
4	UDP	150.*	*	*	3030	accept
5	*	*	*	*	*	deny

Table 1: Example Security Policy Consisting of Multiple Ordered Rules

In the examples described herein, rules will initially be represented as the 5-tuple (protocol, IP source address, source port number, IP destination address, destination port number). Fields can be fully specified or contain wildcards '*' in standard prefix format. For example the prefix 192.* would represent any IP address that has 192 as the first dotted-decimal number. In addition to the prefixes, each filter rule has an action, which is to accept or deny. An accept action passes the packet into or from the secure network,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 820

while deny causes the packet to be discarded. Rules are applied in order to every arriving packet until a match is found; otherwise, a default action is performed [21, 22].

Unfortunately, security policies may contain anomalies, such as a packet that matches multiple rules, where the matching rules specify conflicting actions. Security policy anomaly detection and correction is the subject of continued research [13, 21] and is not the focus of the present subject matter. Given the type of security policy required, rule sets can become quite large and complex. Security can be further enhanced with connection state and packet audit information. For example, a table can be used to record the state of each connection, which is useful for preventing certain types of attacks (e.g., TCP SYN flood) [21, 22].

Traditional firewall implementations consist of a single dedicated machine, similar to a router, that sequentially applies the rule set to each arriving packet. However, packet filtering represents a significantly higher processing load than routing decisions [15, 18, 22]. For example, a firewall that interconnects two 100 Mbps networks would have to process over 300,000 packets per second [21]. Successfully handling this high traffic becomes more difficult as rule sets become more complex [4, 14, 22]. Furthermore, firewalls must be capable of processing even more packets as interface speeds increase. In a high-speed environment (e.g. Gigabit Ethernet), a single firewall can easily become a bottleneck and is susceptible to DoS attacks [4, 7, 10, 12]. An attacker could simply inundate the firewall with traffic, delaying or preventing legitimate packets from being processed. Building a faster single firewall is possible [8, 15, 18, 19, 20]; however, the benefits are temporary (traffic loads and interface speeds are increasing); it is not scalable; it is a single point of failure; and it is generally not cost-effective for all installations.

A data-parallel firewall (or firewall sandwich) where each firewall node implements an entire firewall rule set is another approach to increase the speed of processing traffic [4, 12, 14, 22]. As shown in Figure 1, the system consists of multiple identical firewalls **100** connected in parallel. Each machine implements the complete security policy and arriving packets are distributed across the machines for processing in parallel [4]. A packet distributor **102**

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 821

2005328336 22 Aug 2011

implements a load balancing algorithm to distribute different packets to different firewall nodes. How the load-balancing algorithm distributes packets is vital to the system and is typically implemented as a high-speed switch in commercial products [11,12]. Although data-parallel firewalls achieve higher throughput than traditional firewalls [4] and have a redundant design, they have difficulty maintaining QoS across networks. For example, legitimate traffic can encounter delays if it is queued behind traffic that requires more processing. Under these circumstances, users notice poor network performance, which is growing concern as more network applications require QoS assurances. Furthermore, stateful inspection requires all traffic from a certain connection or exchange to traverse the same parallel machine, which is difficult to perform at high speeds [14]. Therefore, new firewall architectures are needed to meet the demands of future networks and the challenges of increasing security threats.

SUMMARY

In accordance with the present invention, there is provided a function-parallel firewall system comprising (a) a first firewall node for filtering received packets using a first portion of a rule set including a plurality of rules, the first portion including less than all of the rules in the rule set; and (b) at least one second firewall node for filtering packets using a second portion of the rule set, the second portion including at least one rule in the rule set that is not present in the first portion, wherein the first and second portions together include all of the rules in the rule set, wherein the first and second firewall nodes are configured to implement a gateless design wherein the rules are distributed among the first and second firewall nodes such that for a given packet, one of the first and second firewall nodes will accept the packet and the other of the first and second firewall nodes will deny the packet.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 822

2005328336 22 Aug 2011

In accordance with the present invention, there is provided a method for controlling access to a network based on a set of packet filtering rules distributed in a function-parallel manner, the method comprising (a) distributing packet filtering rules of a rule set among a plurality of different nodes in a function-parallel manner so that at least some of the different nodes implement different portions of the rule set; (b) replicating packets to each of the nodes; and (c) applying the rules to filter the packet using a gateless design wherein the rules are distributed among first and second firewall nodes such that for a given packet, one of the first and second firewall nodes will accept the packet and the other of the first and second firewall nodes will deny the packet.

In accordance with the present invention, there is provided a computer program product comprising computer executable instructions embodied in a computer readable medium for performing steps comprising (a) a first firewall node for filtering received packets using a first portion of a rule set including a plurality of rules, the first portion including less than all of the rules in the rule set; and (b) at least one second firewall node for filtering packets using a second portion of the rule set, the second portion including at least one rule in the rule set that is not present in the first portion, wherein the first and second portions together include all of the rules in the rule set, wherein the first and second firewall nodes are configured to implement a gateless design wherein the rules are distributed among the first and second firewall nodes such that for any given packet, one of the first and second firewall nodes will accept the packet and the other of the first and second firewall nodes will deny the packet.

2005328336 22 Aug 2011

According to one aspect, the subject matter described herein includes a function-parallel firewall. The function-parallel firewall includes a first firewall node for filtering received packets using a portion of a rule set. The portion of the rule set may be less than all of the rules in the rule set. At least one second firewall node filters packets using at least one rule in the rule set that is not present in the portion implemented by the first firewall node. The first and second potions together include all of the rules in the rule set.

The term "firewall node," as used herein refers to a physical entity that implements one or more packet filtering rules. For example, a firewall node may be a single-processor computing device that implements the packet filtering rules. Alternatively, a firewall node may be a processor within a multi-processor computing device that implements one or more packet filtering rules. Thus, a multi-processor device may include multiple firewall nodes.

According to another aspect, the subject matter described herein includes a stateful function-parallel firewall system. The system includes a first function-parallel firewall subsystem including a plurality of firewall nodes implementing a first rule set for filtering packets arriving at a network. The first function-parallel firewall subsystem includes at least first and second firewall nodes respectively implementing first and second portions of the first rule set.

The first portion includes at least one rule that is not present in the second portion. A second function-parallel firewall subsystem includes a second set of firewall nodes. The second set of firewall nodes implements a firewall rule set for filtering packets departing from the network. The second function-parallel firewall subsystem includes at least third and fourth firewall nodes. The third and fourth firewall nodes implement at least one rule that is not present in the second portion of the second rule set. The function-parallel firewall subsystems share state information regarding connections established through the first and second sets of firewall nodes.

According to another aspect, the subject matter described herein includes a firewall grid. The firewall grid includes a plurality of firewall nodes that are physically connected to each other via a network for filtering packets. A controller controls logical connections between the firewall nodes. The controller is adapted to configure the firewall nodes in a logical firewall hierarchy without changing physical connections between the firewall nodes.

According to yet another aspect, the subject matter described herein includes a method for distributing rules in a function-parallel firewall. The method includes defining a rule set for a function-parallel firewall. The method further includes assigning rules in the rule set to nodes and branches in a trie data structure. Each node in the trie data structure corresponds to a data field and each branch represents a value for each data field to be compared to be values in corresponding fields of received packets. The trie data structure is pruned in a manner that preserves ordering of the rules in the rule set.

According to yet another aspect, the subject matter described herein includes a method for assigning rules to a plurality of firewall nodes. The method includes assigning rules to nodes in a directed acyclical graph (DAG). Relationships between rules are represented by edges in the DAG. Rules are distributed among a plurality of firewall nodes using the DAG in a function-parallel manner and such that relationships between the rules specified in the DAG are preserved.

As used herein, the term "firewall" refers to a logical entity that is adapted to filter packets at the ingress and/or egress portions of a network based on a policy. The term "firewall" is intended to include systems that block packets from entering or leaving a network and systems that perform intrusion

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 825

detection and intrusion protection functions for packets entering or leaving a network. Thus, the term "function-parallel firewall" is also intended to include an intrusion detection system or an intrusion protection system that operates in a function-parallel manner as described herein.

The subject matter described herein can be implemented using any combination of hardware, software, or firmware. In one exemplary implementation, the subject matter described herein may be implemented using a computer program product comprising computer executable instructions embodied in a computer readable medium. Exemplary computer readable media suitable for implementing the subject matter described herein include disk memory devices, chip memory devices, programmable logic devices, application specific integrated circuits, and downloadable electrical signals. In addition, a computer program product that implements the subject matter described herein may be located on a single device or computing platform or may be distributed across multiple devices or computing platforms.

One objective of the subject matter described herein is to provide a new security architecture, referred to herein as a hierarchical firewall, that provides the functionality of a traditional firewall (enforcing a security policy) and includes the following improvements:

1. **Suitable for high speed networks and scalable for large traffic loads.** Firewalls must continue to perform security tasks quickly and efficiently, appearing transparent to legitimate users under varying traffic loads.

2. **Maintains QoS requirements and traffic differentiation.** A growing number of network applications require QoS for their operation. Firewalls must not violate these guarantees while performing their security tasks.

3. **Greatly lessens the impact of DoS attacks.** Firewalls are candidates for DoS attacks, since they serve as the point of entry between two networks. New architectures must be capable of handling these new threats.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 826

4. **Highly available.** Firewalls introduce a single point of failure between two networks. New firewall architectures must provide redundancy that prevents segmenting the secure network from the external network.

According to one aspect, the subject matter described herein includes a set of hierarchically configured computers that collectively enforce a desired security policy. Given this distributed configuration, traffic is quickly divided among the machines based on perceived threat. In one exemplary implementation, traffic that is considered safe may be promptly forwarded into the secure network, while the remaining traffic may be forwarded to different machines in the hierarchy for further scrutiny. Hence, traffic is segregated and queued based on security threat, yielding minimal delays for legitimate traffic. The system is also robust and highly available, since multiple machines are employed for the architecture and low cost, since Linux-installed PCs and open source software can be used for its implementation [21]. The proposed architecture is believed to be applicable to a wide variety of agencies in the public and private sectors.

The firewall architecture described herein will provide a scalable solution to varying traffic loads and QoS requirements, which are increasingly important as computing becomes more ubiquitous and applications become more diverse. The architecture is of interest to the networking community, since the architecture will result in new security policy models and optimization methods. The architecture is also relevant to the military and business community, since it provides a new, robust and cost-effective solution to security threats.

Some of the objects of the subject matter described herein having been stated hereinabove, and which are addressed in whole or in part by the subject matter described herein, other objects will become evident as the description proceeds when taken in connection with the accompanying drawings as best described hereinbelow.

BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the subject matter described herein will now be explained with reference to the accompanying drawings of which:

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 827

Figure 1 is a block diagram of a conventional data-parallel firewall;

Figure 2 is a block diagram illustrating an exemplary function-parallel firewall architecture according to an embodiment of the subject matter described herein;

Figure 3A is graph comparing simulation results for average packet delay for different firewall architectures;

Figure 3B is a graph comparing simulation results for maximum packet delay for different firewall architectures;

Figure 4A is a graph illustrating average packet delay for different packet types for a data-parallel firewall;

Figure 4B is a graph illustrating simulation results for packet delay for different packet types for a function-parallel hierarchical firewall according to an embodiment of the subject matter described herein;

Figure 5 is a data structure diagram illustrating a full trie representation of the security policy illustrated in Table 1;

Figure 6 is a data structure diagram illustrating a pruned trie representation of the security policy given in Table 1 where each packet is processed within four field comparisons;

Figure 7A is a data structure diagram illustrating a trie representing two security rules assuming that the default action per node is deny;

Figure 7B is a data structure diagram illustrating a reordered version of the trie data structure illustrated in Figure 7A;

Figure 8A is a diagram illustrating a Jackson model for a function-parallel hierarchical firewall according to an embodiment of the subject matter described herein;

Figure 8B is a diagram illustrating a Jackson model for a data-parallel firewall;

Figure 9 is a graph illustrating simulation results for average delay for different firewall architectures;

Figure 10 is a block diagram illustrating a stateful hierarchical firewall according to an embodiment of the subject matter described herein;

Figure 11A is a block diagram illustrating physical connections of a firewall grid according to an embodiment of the subject matter described herein;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 828

Figure 11B is a block diagram illustrating logical connections of the firewall grid illustrated in Figure 11A according to an embodiment of the subject matter described herein;

Figure 12 is a policy directed acyclical graph (DAG) representation of the firewall rules given in Table 2;

Figure 13A is a block diagram illustrating rule distribution among a data-parallel firewall;

Figure 13B is a block diagram illustrating a function-parallel firewall and rule distribution among nodes in the function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 14A is a block diagram illustrating one exemplary rule distribution among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 14B illustrates another distribution of rules among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 14C illustrates an illegal rule distribution among nodes in a function-parallel firewall;

Figure 15A is a block diagram illustrating non-redundant rule distribution among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 15B is a block diagram illustrating redundant rule distribution among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 16A illustrates a policy DAG where deny rules are indicated by a square node;

Figure 16B is a block diagram illustrating initial gateless rule distribution among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 16C is a block diagram illustrating final gateless rule distribution among nodes in a function-parallel firewall according to an embodiment of the subject matter described herein;

Figure 17A is a graph illustrating simulation results for average packet delay as packet arrival rate increases for different firewall designs;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 829

Figure 17B is a graph illustrating simulation results for packet delay as the arrival rate and number of rules increase for different firewall designs;

Figure 18A is a graph illustrating simulation results for average packet delay for different firewall architectures;

Figure 18B is a graph illustrating simulation results for maximum packet delay for different firewall architectures; and

Figure 19 is a flow chart illustrating exemplary steps for filtering packets using a rule set distributed among nodes in a function-parallel manner according to an embodiment of the subject matter described herein.

DETAILED DESCRIPTION OF THE INVENTION

According to one aspect of the subject matter described herein may include multiple firewall machines configured in a hierarchy, as shown in Figure 2. In the example illustrated in Figure 2, a plurality of firewall nodes **200**, **202**, **204**, and **206** collectively implement a firewall policy. Each firewall node or set of nodes implements a different level of the firewall policy in a hierarchy. For example, firewall node **200** may implement a first level of a firewall policy that includes rules that are not present in subsequent levels. If a packet passes the first level of screening, the packet may be forwarded immediately to the internal network without impacting other firewall nodes. If the packet is identified as requiring further scrutiny, the packet may be forwarded to the next level where it is screened by firewall node **202** or **204** using rules that were not applied by firewall node **200**. If the packet passes level 2 screening, the packet may be forwarded directly to the internal network. If the packet is indicated as requiring further screening, the packet will be passed to level 3 where level 3 screening is applied by firewall node **206**.

At each level in the firewall hierarchy illustrated in Figure 2, firewall nodes may implement different portions of a packet filtering rule set. An architecture in which different firewall nodes implement different portions of a firewall rule set is referred to herein as a function-parallel architecture. A function-parallel firewall architecture in which packet filtering rules are distributed among firewall nodes such that if a packet passes or fails at a

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 830

certain level, the packet may bypass later levels is referred to herein as a hierarchical firewall architecture.

In the architecture illustrated in Figure 2, at each level, machines may be configured to perform a specific security task quickly and efficiently. In one implementation, the firewall rules may be distributed across the different hierarchy levels such that increased scrutiny is placed at lower levels. Since the rules are distributed among the levels, no machine in the hierarchy is required to implement the entire rule set. In contrast, a data-parallel firewall architecture requires all machines to implement the complete rule set.

As described above, in the exemplary architecture illustrated in Figure 2, packets start at the highest level and progress downward through the hierarchy. At each level, a packet can be forwarded to the internal network if deemed safe, forwarded to the next lower level for more processing, or dropped. The benefit of this divide-and-conquer approach is that traffic will incur delays consistent with its perceived security threat; therefore, traffic that is deemed safe will not be delayed by traffic that requires more processing. As a result, the hierarchical firewall can provide differentiated service and support QoS guarantees, unlike current designs. Furthermore, traffic that is forwarded to the internal network can be shaped at the particular firewall level, constraining the amount of traffic entering the network and lessening the effect of DoS attacks.

For example, in the hierarchical firewall depicted in Figure 2, packets are forwarded from the external network to the first level, where initial security checks are performed, such as a quick validity check and comparisons of IP addresses. Packets with trusted IP addresses are immediately forwarded to the internal network, while the remaining packets are sent to the next level. At the second level, packets are separated based on transport protocols and/or port numbers. Again, if a packet is deemed safe it is immediately forwarded to the secure network; otherwise, they are sent to the next level for auditing. Log files may be used at one or more of the firewall nodes indicate which connections are open and have been requested. Each level may also drop packets based on perceived threat.

The preceding example illustrates a simple hierarchical firewall; however, the size and complexity can vary. In general, the hierarchy attempts to quickly divide traffic and provide differentiated service. More levels can be used for

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 831

stronger security policies and to distribute traffic more efficiently. Determining the appropriate number of levels, the number of machines per level, and how to distribute the rule set are implementation details that may be used to optimize a function-parallel firewall according to an embodiment of the subject matter described herein.

Preliminary Experimental Results

In this section, simulation is used to compare a hierarchical firewall with a traditional single-machine firewall and a data-parallel firewall [4, 22]. Each architecture implemented the same 325-rule security policy, where each rule required 1 unit of time to complete. The single firewall architecture consisted of one machine that processed each packet received from the external network. The data-parallel architecture consisted of four machines, as shown in Figure 1, each implementing the same 325-rule set. The hierarchical firewall consisted of four machines configured in three levels, as shown in Figure 2. Rules were distributed across the levels such that the first level implemented 25 rules; the second level implemented 50 rules each; and the third level implemented 75 rules. The group of rules performed at a certain machine will be referred to as a security task. The hierarchical firewall incurred an additional delay of one time unit when a packet was transmitted from machine to machine. Packets randomly arrived to the system using a Poisson distribution with a mean of 30 units and would randomly match one of the 325 rules based on a uniform distribution. For classification purposes, each packet was associated with the security task it matched (which group of rules). In this experiment, the performance parameters analyzed are the overall average packet delay, the maximum packet delay, and the average delay experienced by each packet type.

Experimental results for the different architectures are shown in Figures 3 and 4. More particularly, Figure 3A illustrates average packet delay for a single firewall design, for the data-parallel firewall design illustrated in Figure 1, and for the hierarchical firewall design illustrated in Figure 2. Figure 3B illustrates maximum packet delay for the single, data-parallel, and hierarchical firewall designs. Figure 4A illustrates average packet delay for the data-parallel firewall design for different packet types. Figure 4B illustrates average packet

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 832

delay for the hierarchical firewall design for different packet types. Packet types refer to packets of a classified based on which security task was satisfied. Unlike the data-parallel firewall, the hierarchical design consistently yields different delays for each packet type.

For both delay measurements, the hierarchical architecture performed better than the single firewall and the data-parallel firewall. The single firewall had an average delay 31 times greater than the hierarchical architecture, which demonstrates how a single firewall machine can quickly become a bottleneck. Although the data-parallel firewall performed better than the single firewall, it had an average and maximum delay 6 times longer than the hierarchical architecture. In the parallel configuration, traffic that needed little processing often queued behind traffic that required more processing, which resulted in longer delays. In contrast, the hierarchical system consistently yielded better results, since traffic was distributed among the levels based on processing requirements.

The average delay experienced by each packet type is also important and is illustrated in Figures 4A and 4B. Delays in the hierarchical firewall were proportional to the amount of processing required; for example, security task 1 traffic had the shortest delay, security task 2 traffic had the second shortest delay, security task 3 traffic had the third shortest delay; and security task 4 traffic had the highest delay. Even though security tasks 2 and 3 were on the same level, security task 3 traffic had a delay slightly higher that security task 2 traffic, since it was queued with security task 4 traffic. Therefore, the hierarchical firewall is capable of providing differentiated service to multiple packet types/classes. In contrast, the delays varied for each packet type in the data-parallel firewall. Packets were constantly delayed by lower type/class traffic (e.g., security task 1 packets delayed by security task 4 packets); thus, no differentiation was achieved. Although these results are preliminary, they indicate the hierarchical architecture has merit, offering significantly better performance than current firewall systems.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 833

Refinements

The previous section described an initial hierarchical firewall architecture for employing a security policy. Although preliminary results indicate the architecture has merit, refinements can be made so that the design can achieve its full potential. One refinement relates to the hierarchical architecture. Important questions include determining the appropriate number of levels and how to effectively distribute the security policy rules. Another refinement may include providing stateful packet inspection (maintaining connection tables) and increasing redundancy. Finally, another refinement may include examining other high-speed firewall architectures, such as a dynamic firewall grid. Analytical models and simulation will be used for design and analysis; in addition, a prototype will be implemented along with analysis tools and software.

Architecture Design and Rule Set Distribution

Important questions about the firewall architecture include the number of levels, the number of machines per level, and the number of security tasks per machine. The security policy is central for addressing these questions and, for this reason, will be formally defined and modeled. The model will be optimized to reduce the physical size, minimize the traversal time, and provide differentiated service. Afterward, the optimized policy model will be used to implement the actual hierarchical firewall.

One method for designing the hierarchy is to represent the security policy graphically as a tree data structure. This representation will aid the translation to the actual firewall, since it has a hierarchical form. Although a variety of tree data structures exist, a trie will be initially investigated [1]. A trie is a special type of data structure that has the advantages of fast retrievals and quick unsuccessful search determination, which are important for enforcing a security policy. In a security policy trie, each path from the root node to a leaf corresponds to one rule. The nodes of the trie correspond to the rule fields, while branches are possible values. Figure 5 is a trie data structure for the rules in Table 1. For example, in Figure 5, the root node is the protocol field. The next level represents the source IP address, while the leaf nodes represent

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 834

the rule action. Rules that have the same field value at a node will share a branch for that value.

The trie in Figure 5 referred to as full trie, since every noncycle path from the root to a leaf has the same length, and every leaf lies at the same depth [6]. Unlike other policy models, the rule order must be preserved if the trie is to be successfully translated to a hierarchical firewall. The action to apply to a packet is determined by a preorder traversal. Starting at the root node, branches are compared to the fields of the packet. Since multiple branches may match a packet, they are traversed in order from left to right. If a packet field does not match any possible branches, it moves back to the previous node (parent), where the next branch is tested. Once the policy trie is developed, it can be translated into a hierarchical firewall. The nodes of the trie represent the machines in the hierarchy, while the branches represent the links between nodes. The value associated with each branch indicates the packet field value required to traverse it. Although a full trie representation of a security policy can be created, it is not necessarily optimal. Further trie manipulation is required create an effective hierarchical firewall, such as preventing routing to parents and limiting trie traversal time.

Pruning yields a trie that never requires routing to parent nodes, which provides faster packet processing [1, 15]. Assuming a full trie consists of n nodes and k levels (where k <= n), the worst case packet processing would be O(n*k) tuple-compares. In contrast, a pruned trie requires only O(k + n) tuple-compares in the worst case which is a processing reduction of 1/k.. Pruning is accomplished by replacing general rules with descendants. Rule r_1 is a descendant of rule r_2, if $r_2[i]$ is a prefix of $r_1[i]$ for $i = 1, ..., k$. Therefore, pruning occurs by pushing a rule down to all its descendents, then deleting the rule. An example of a pruned version of the trie illustrated in Figure 5 is given in Figure 6, where [*, *, *, deny] is a general rule that can be pushed down to every other rule. Although the original branch has been removed and added to the other branches, the order is preserved, if branches are tested from left to right. Current pruning algorithms do not consider rule order, since tries were primarily developed for storing data [6]. New pruning techniques that preserve order may be developed and used to distribute rules among nodes in a functional-parallel firewall according to an embodiment of the subject matter described

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 835

herein. Pruning may also cause the number of nodes in the hierarchy to increase, since a general rule is duplicated at different levels. For this reason, other techniques are needed to reduce the trie size.

Reducing the trie size is important, since delay increases each time a packet is forwarded from one node to another. The trie can be reduced by compressing nodes, which is achieved by removing single branching paths or consolidating sub-ties into one node [2, 6, 15]. As a result, a node may compare multiple rule fields or perform multiple security tasks per packet. For example, in Figure 5, the rule [UDP, 150.*, *, 3030, accept] is considered a leaf chain [6] and can be compressed into one node (a single machine). Two other methods commonly used to reduce the size of a trie are Directed Acyclic Graphs (DAG) and level reordering [1, 6, 15]. If the same sub-trie appears in multiple places, then the trie structure can be changed to a DAG. The possible use of DAG is seen in Figure 6, where nodes labeled $Port^a$ and $Port^b$ are equivalent sub-tries. Parent nodes of both sub-tries could point to just one, eliminating the need for the other.

Another method for reducing the trie size is level reordering. As seen in Figures 7A and 7B, the order of the levels can significantly affect the number of nodes in the trie. More particularly, in Figure 7A, an original trie representing the security rule [TCP, 120.*, 210.*, 90, accept] and [TCP, 130.*, 210.*, 90, accept] is illustrated. Both Figures 7A and 7B assume that the default action per node is deny. Since some of the rules in the original trie in Figure 7A include common elements, the rules can be reordered so that the number of branches in the subtrie are reduced. For example, in Figure 7A, the destination IP address and port for each branch of the subtrie are the same. Accordingly, if these common portions of the rule are implemented at higher levels, the duplicate nodes in the trie can be reduced, as illustrated by Figure 7B. More particularly, in Figure 7B, the nodes for destination IP and port are placed before the node for source IP address. As can be seen from Figure 7B, reordering the levels produces a smaller trie. Using a DAG can also reduce the number of nodes.

The problem of finding an ordering that guarantees a minimum size trie has been shown to be NP-complete; however, several heuristics have been developed [2, 6]. It should be noted that DAGs and reordering only reduce the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 836

number of nodes in the hierarchy and do not reduce the depth of the trie, which influences the traversal time.

Regardless of the applicability of these different trie modifications, they must be performed in concert to achieve an optimal hierarchical design. For example, the architecture utilizes multiple nodes to distribute traffic load and provide differentiated service; however, the trie path length (depth) impacts the packet traversal time. Therefore, compression and ordering must be performed together to provide differentiated service and optimize packet delay.

To achieve these goals, information not explicitly given in the security policy (e.g., QoS, level of trust, and processing time) may be used. For example, it is advantageous to place security tasks that require the most time at the lowest levels. This prevents packets from queueing behind traffic that requires more processing time; thus, lowering delay. A weight w_i can be assigned to each security task s_i (a single rule or multiple rules) based on the amount of time it requires. For example, an auditing security task requires more time (updating system files) than those that only perform field comparisons. Using these weights, the trie can be ordered to minimize the average traversal time. For example assuming i and j are hierarchy levels, where $i < j$, the trie can be ordered such that the average weight per level increases from the highest to the lowest level $\overline{w}_i < \overline{w}_j, \forall i \neq j$. In addition, the trie can be designed to separate traffic based on QoS or level of trust. Traffic would traverse different sub-tries, based on the type of service required, providing class isolation and different delays. The algorithms described in this section can then be applied to these sub-tries independently to reduce size and traversal time. These ideas were used to design the hierarchical firewall simulated in the preceding section, which showed excellent results. New algorithms for trie compression based on traversal time with constraints on QoS will be developed for optimizing function-parallel firewall design.

Packet traversal time remains a key performance metric for firewall systems. For this reason, queueing models will be used to analyze the delay of different firewall designs. The node delay encountered by a packet consists of the processing time, queueing delay, and transmission time. The processing delay is the time required to perform the security task. Assume node i must

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 837

process R_i rules per packet and can perform xi rules per unit time, then the service rate is x_i/R_i packets per unit time. Queueing is the delay associated with waiting for processing and depends on the system state. The transmission delay is the time required to transmit the packet, which depends on the packet's length and the line speed. The maximum delay experienced by a packet traversing the firewall is the sum of the node delays along the path from the root to a leaf.

If it is assumed the packet arrival and processing (transmission and security task) rates are exponentially distributed, the node can be modeled as an M/M/1 queue [5, 17]. The firewall can then be considered an open network of M/M/1 queues, where probabilities are assigned to each link to indicate the likelihood of moving to the next node. This is a Jackson network model, which has a product-form solution for the average delay [5, 17]. The average end-to-end delay for N cascading nodes (traversal path) can be computed as

$$E(T) = \sum_{i=1}^{N} \frac{1}{\mu_i - \lambda_i} \tag{1}$$

where $1/\mu_i$ is the service time (processing and transmission) and λ_i is the arrival rate to node i. Figures 8A and 8B respectively illustrate Jackson models for hierarchical and data-parallel firewalls. In Figures 8A and 8B, it is assumed that the total arrival rate to each system is λ packets per unit time, each system implements the same R rule security policy, and each node can perform x rules per unit time. For the hierarchical firewall, shown in Figure 8A, the policy rules are evenly distributed to each node; therefore, the service rate is $3x/R$ per node. The traffic arrival rate to node A is λ packets per unit time, half of the packets are forwarded to node B, while the other half is forwarded to node C. Although not considered in this example, some traffic may be forwarded into the secure network or dropped by node A; therefore in this case, the sum of the rates to nodes B and C would be less than λ. Using these values, the average end-to-end delay for either path in the hierarchical firewall is

$$E_h(T) = \frac{1}{\frac{3x}{R} - \lambda} + \frac{1}{\frac{3x}{R} - \frac{\lambda}{2}} \tag{2}$$

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 838

For the data-parallel configuration, illustrated in Figure 8B, traffic arrives at the packet distributor, which evenly distributes the traffic to each node. Each node implements the complete rule set R; therefore, the service rate is x/R. The end-to-end delay across any node in the data-parallel firewall is

$$E_p(T) = \frac{1}{\frac{x}{R} - \frac{\lambda}{3}} \quad (3)$$

Using these models, the effect of increasing traffic load λ on the average end-to-end delay, for the two architectures is shown in Figure 9. For this example, the arrival rate λ ranged from 1 to 50 packets per unit time. The rule set R consisted of 240 rules, while x equaled 5000 rules per unit time. Each of the hierarchical nodes incurred an additional delay of five time units to compensate for inter-node transmission. As expected, the average delay increases as the network load increases for both architectures; however, the average delay is consistently lower for the hierarchical firewall. This is due to the decreased rule set per node and the reduced arrival rates for nodes B and C. The average is an upper bound for the hierarchical firewall. Allowing node A to forward packets to the secure network or to drop packets would reduce the load on the second level and decrease delays. The arrival rate is less per node in the data-parallel firewall; however, the system has a higher end-to-end delay because each node must implement the complete rule set. The theoretical results are consistent with the experimental results described above. These models provide an important tool for analyzing the effect of load and rule distribution on end-to-end delay. In addition to analytical results, simulation will be used to measure the performance of the hierarchical designs under various conditions.

Stateful Firewalls

Another aspect of the subject matter described herein includes a stateful hierarchical firewall. The preceding design is considered stateless, or static, since the security action taken is based solely on the packet currently inspected. A stateful firewall provides better security by tracking and controlling network flows (packets belonging to a particular network connection or

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 839

exchange) [16, 21, 22]. In one exemplary implementation, a stateful firewall can be achieved by storing state information concerning each flow in a table, and then according it with the security policy. For example, each time a TCP connection is established from the internal network to the external network, it is logged in a table. The table may contain the source and destination addresses, port numbers, sequencing information, and connection flags. Afterwards, packets arriving from the external network claiming to be part of an existing TCP connection can be checked against the table. If a packet belongs to an existing connection or exchange, it can be accepted. Otherwise, it is dropped. Table entries expire, if packets belonging to an exchange are not observed by the firewall over a period of time [21].

Therefore, a stateful firewall is dynamic and limits the hackers' opportunity to exploit certain services. The increased security does require additional memory and time for table manipulation, which is problematic for high-speed networks. Furthermore, the firewall must process all traffic arriving and departing the secure network to adequately maintain the connection tables 21]. To solve these problems, the proposed firewall system will be divided into two hierarchies, one for arriving traffic and another for departing traffic, as illustrated in Figure 10. In Figure 10, a stateful firewall system includes first and second hierarchical firewall subsystems **1000** and **1002** and a control node **1004**. First and second hierarchical firewall subsystems **1000** and **1002** may be function-parallel firewall systems including a plurality of firewall nodes with different rules distributed among different nodes as described above. In one exemplary implementation, hierarchical firewall subsystem **1000** may include at least two firewall nodes that implement different portions of a first firewall rule set. Similarly, hierarchical firewall subsystem **1002** may include at least two firewall nodes that implement different portions of a second firewall rule set. Firewall subsystem **1000** may filter packets entering the network. Firewall subsystem **1002** may filter packets departing from the network. Firewall subsystems **1000** and **1002** may share state information regarding connections established through the nodes within each system. Control node **1004** may store the state information and may be accessible by the nodes in hierarchical firewall subsystems **1000** and **1002.** Rules can be distributed to the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 840

hierarchies and the designs can be optimized using the methods described in preceding section.

However, a new problem with the stateful model is the need to distribute state information across nodes. For example, nodes that scrutinize arriving traffic must be aware of any new connections permitted by nodes that inspect departing traffic. Furthermore, nodes must be informed of table entries that expire over time. To communicate state information, a central database machine, such as control node **1004**, may collect and distribute the information (table entries) to the appropriate nodes. In one exemplary implementation, control node **1004** may not perform table queries on behalf of nodes. Individual nodes will only query their local tables, which prevents the central machine from becoming a bottleneck. In addition, the central machine will only send information (table entries) to nodes that might require it. As a result of this policy, the table searched by an individual node is reduced (contains only relevant information) as well as the search time. In an alternate implementation, control node **1004** may store the connection table and may perform queries for packets received by firewall subsystems **1000** and **1002**.

Survivability and Firewall Grids

Another aspect of the subject matter described herein focuses on increasing the survivability of the hierarchical design. Although the proposed design has the benefit of increased packet throughput and minimal delay, the reliability of the system is decreased, since more machines (nodes) are deployed. For example, if the highest level node fails (is unable to inspect more traffic or goes off-line), then the entire hierarchy is unable to process packets. For this reason, the hierarchy is considered an n of n system, where all *n* machines must be functioning for the system to operate correctly [4, 12].

Survivability of the system can be improved with the use of redundant (parallel) nodes at different levels. For example, the first level may include multiple machines in parallel, where each machine implements the same security task (e.g., an IP check). An arriving packet would be sent to only one of the parallel machines, where it would be accepted, dropped, or forwarded to the next level for further processing. If one of the parallel nodes fails, then the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 841

remaining nodes can process the traffic. An additional benefit of parallel processing is the equal distribution of traffic among nodes, which can result in higher throughput and lower delays. Details such as where to locate parallel nodes and the number of machines required for each parallel configuration may depend on the desired degree of reliability versus cost of deploying redundant machines.

Yet another aspect of the subject matter described herein includes a grid of configurable firewalls (nodes) that can be used to efficiently process traffic, while being highly available. In one exemplary implementation, the grid may include a network of highly interconnected nodes (possibly via a high-speed switch or router) that can be logically deployed in various configurations (e.g., parallel, hierarchical, or hybrid). The advantage of the firewall grid is its ability to quickly adapt to new demands. When security needs change (e.g., due to policy changes, increased traffic volume, or node failure), these machines can be deployed or configured to maintain throughput and QoS requirements. For example, in a logical hierarchy, if a higher level encounters an increased traffic load, less frequently used redundant machines at lower levels can be reconfigured to process packets at the higher level. Unlike the hierarchical model, physical changes are not required, since all the machines are interconnected. Only the logical connection between nodes must be updated, which can be performed quickly. Like the stateful hierarchical firewall, the initial grid model will use a central machine to coordinate node configuration and communication. This central machine will play an important role in the system (balancing processing demand per node) and may constrain the efficiency of the firewall grid. Future versions of the firewall grid will attempt to place more intelligence in the nodes and allow them to interact independently and adapt dynamically under the guidance of a general security policy.

Figures 11A and 11B illustrate an example of a firewall grid according to an embodiment of the subject matter described herein. Referring to Figure 11A, a plurality of firewall nodes **1100-1114** are connected via a control node **1116**. Each firewall node **1100-1114** may be configured to implement different portions of a firewall rule set, as controlled by control node **1116**. Each of the firewall nodes is physically connected to the other firewall nodes via physical

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 842

connections **1118**. Physical connections **1118** may be local area network connections, such as Ethernet connections. As such, control node **1116** may be an Ethernet switch.

In Figure 11B, firewall nodes **1100-1114** are connected via logical connections **1120**. Logical connections **1120** may be implemented by control node **1116** without changing physical connections **1118**. For example, control node **1116** may configure the output port to which firewall **1106** is connected to receive frames from firewalls **1100**, **1102**, and **1104**. Similar internal configurations can be implemented for the remaining firewalls illustrated in Figure 11B. In Figure 11B, logical connections **1120** between the firewalls implement a hierarchical, function-parallel firewall design, such as that illustrated in Figure 1. In addition, the design illustrated in Figure 11B is also data-parallel at each level. For example, each firewall node at each level may implement a common rule or rule set. For example, the rule set implemented by the firewall nodes in level 1 is illustrated in Figure 11B as R1. The firewall nodes in level 2 implement a rule set represented by R2, and the firewall nodes in level 3 implement a rule set represented by R3. If any of the firewall nodes in a given level fails, the remaining firewall nodes in that level may take over operation for the failed firewall. Alternatively, as described above, if a failure or a bottleneck occurs at one level, a less-used firewall at another level may be logically configured to operate at the level where the failure or congestion occurred. Packet distributors **1122** may be configured to distribute the firewall rules between the firewalls at each level. Thus, as illustrated by Figures 11A and 11B, logical connections between firewall nodes and a firewall grid may be dynamically altered (i.e., changed on the fly without ceasing packet filtering operation) by a controller node without changing the physical connections between the firewalls.

Firewall Implementation

In addition to the analytical models and simulations described above, an implementation of a hierarchical firewall to measure performance will now be described. One exemplary implementation may utilize Linux-based PCs and Ethernet connections. Each PC will contain multiple Ethernet cards and be

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 843

loaded with the Linux 2.4 kernel. The 2.4 kernel has support for various packet queueing disciplines, traffic control methods, and packet processing (e.g., iproute2, tc, and iptable). Collectively, these functions are required for advanced packet filtering. For example, iproute2 and iptable can route packets based on information other than the destination address, while tc can mark and shape traffic. PCs will serve as hierarchical nodes, sources and receivers of network traffic, and hierarchy controllers.

Results from an initial implementation of a three-node hierarchical firewall indicate that these open-source components are suitable for high-speed networks. The Computer Science Department at Wake Forest University has 10 PCs available for this project, where each has a 600 Mhz processor and multiple 100 Mbps Ethernet cards. These machines will serve as an initial testbed for system development; however, faster machines and high-speed network connections are required for actual system testing. For this reason, this project requests the purchase of 6 additional PCs (2.4 Ghz processors), each with multiple Gigabit Ethernet cards (3 cards per machine). These new machines will be configured as the firewall, while the current machines will serve as traffic generators and receivers. For comparison purposes, parallel and single-machine firewalls will also be implemented using this equipment. Any software required will be developed and provided by the Computer Science Department.

Firewall Rule Distribution for Function-Parallel Firewall Design

As stated above, one aspect with creating an optimal function-parallel firewall design includes distributing rules among firewalls in the firewall rule set. Firewall rules are preferably distributed among firewall nodes in a manner that preserves firewall policy integrity. The following sections describe an exemplary firewall rule set and different options for distributing the firewall rule set among firewall nodes in a function-parallel firewall design.

Firewall Policy

In this section and the following sections, a rule r is modeled as an ordered tuple of sets, $r = (r[1], r[2], ..., r[k])$. Order is necessary among the tuples since comparing rules and packets requires the comparison of

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 844

corresponding tuples. Each tuple $r[l]$ is a set that can be fully specified, specify a range, or contain wildcards '*' in standard prefix format. For the Internet, security rules are commonly represented as a 5-tuple consisting of: protocol type, IP source address, source port number, IP destination address, and destination port number [21, 22]. For example, the prefix 192.* would represent any IP address that has 192 as the first dotted-decimal number. Given this model, the ordered tuples can be supersets and subsets of each other, which forms the basis of precedence relationships. In addition to the prefixes, each filter rule has an action, which is to accept or deny. However, the action will not be considered when comparing packets and rules. Similar to a rule, a packet (IP datagram) d can be viewed as an ordered k-tuple $d = (d[1], d[2], ..., d[k])$; however, ranges and wildcards are not possible for any packet tuple.

Using the previous rule definition, a standard security policy can be modeled as an ordered set (list) of n rules, denoted as $R = \{r_1, r_2, ..., r_n\}$. A packet d is sequentially compared against each rule r_i starting with the first, until a match is found ($d \Rightarrow r_i$) then the associated action is performed. A match is found between a packet and rule when every tuple of the packet is a subset of the corresponding tuple in the rule.

Definition Packet d matches r_i if

$$d \Rightarrow r_i \text{ iff } d[l] \subseteq r_i[l],\ l = 1, ..., k$$

The rule list R is comprehensive if for every possible legal packet d a match is found using R. Furthermore, two rule lists R and R' are equivalent if for every possible legal packet d the same action is performed by the two rule lists. If R and R' are different (e.g. a reorder) and the lists are equivalent, then the policy integrity is maintained.

As described above, a rule list has an implied precedence relationship where certain rules must appear before others if the integrity of the policy is to be maintained. Table 2 shown below illustrates an exemplary rule set that will be used herein to illustrate precedence relationships and rule assignment according to an embodiment of the subject matter described herein.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 845

No.	Proto.	Source IP	Source Port	Destination IP	Destination Port	Action	Prob.
1	UDP	1.1.*	*	*	80	deny	0.01
2	TCP	1.*	*	1.*	90	accept	0.02
3	TCP	2.*	*	2.*	20	accept	0.25
4	UDP	1.*	*	*	*	accept	0.22
5	*	*	*	*	*	deny	0.50

Table 2: Example Security Policy Consisting of Multiple Ordered Rules

In Table 2, rule r_1 must appear before rule r_4, likewise rule r_5 must be the last rule in the policy. If for example, rule r_4 was moved to the beginning of the policy, then it will shadow [23] the original rule r_1. Shadowing is an anomaly that occurs when a rule r_j matches a preceding rule r_i, where $i < j$. As a result of the relative order r_j will never be utilized. However, there is no precedence relationship between rules r_1, r_2, or r_3 given in Table 2. Therefore, the relative order of these three rules will not impact the policy integrity and can be changed to improve firewall performance.

As described in [28], the precedence relationship between rules in a policy can be modeled as a Policy Directed Acyclical Graph (DAG). Let $G = (R,E)$ be a *policy DAG* for a rule list R, where vertices are rules and edges E are the precedence relationships (constraint). A precedence relationship, or edge, exists between rules r_i and r_j, if $i < j$ and the rules intersect [28].

Definition The intersection of rule r_i and r_j, denoted as $r_i \cap r_j$ is

$$r_i \cap r_j = (r_i[l] \cap r_j[l]),\ l = 1, \ldots, k$$

Therefore, the intersection of two rules results in an ordered set of tuples that collectively describes the packets that match both rules. The rules r_i and r_j intersect if every tuple of the resulting operation is non-empty. In contrast, the rules r_i and r_j do not intersect, denoted as $r_i \not\cap r_j$, if at least one tuple is the empty set. It should be noted that the intersection operation is symmetric; therefore, if r_i intersects r_j, then r_j will intersect r_i. The same is true for rules that do not intersect.

For example, for the rules given in Table 2, the intersection of r_1 and r_4 yields (UDP, 1.1.*, *, *, 80). Again, the rule actions are not considered in the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 846

intersection or match operation. Since these two rules intersect, a packet can match both rules for example d = (UDP, 1.1.1.1, 80, 2.2.2.2, 80). The relative order must be maintained between these two rules and an edge drawn from r_1 to r_4 must be present in the DAG, as seen in Figure 12. In Figure 12, the policy DAG represents the rules in Table 2. Vertices **1200** indicate rules while edges **1202** indicate precedence relationships. Rules r_2 and r_3 in Table 2 provide an example of rules that can be re-ordered without changing the firewall policy. These two rules do not intersect due to the fifth tuple (destination port). A packet cannot match both rules indicating the relative order can change; therefore, an edge will not exist between them.

As described above, traditional firewall implementations consist of a single, dedicated machine, similar to a router, that applies the rule set to each arriving packet. However, packet filtering represents a significantly higher processing load than routing decisions [15, 18, 22]. For example, a firewall that interconnects two 100 Mbps networks would have to process over 300,000 packets per second [21]. Successfully handling these traffic loads becomes more difficult as rule sets become more complex [4, 14, 22]. Furthermore, firewalls must be capable of processing even more packets as interface speeds increase. In a high-speed environment (e.g. Gigabit Ethernet), a single firewall can easily become a bottleneck and is susceptible to DoS attacks [4, 7, 10, 11]. Building a faster single firewall is possible [8, 15, 18, 19, 20]; however, the benefits are temporary (traffic loads and interface speeds are increasing); it is not scalable; it is a single point of failure; and it is generally not cost-effective for all installations.

Parallelization has been demonstrated to offer a scalable approach for improving the performance of network firewalls. As seen in Figure 13A, parallel designs consist of multiple firewalls, referred to as firewall nodes, connected in an array. In Figure 13A, firewall nodes **1300** each implement a common rule set. A packet distributor **1302** load shares packets among firewall nodes **1300**. In Figure 13B, a function-parallel design is illustrated where firewall nodes **1304** implement different portions of a firewall rule set. A packet duplicator **1306** duplicates each packet to each firewall node. A gate **1308** controls packets entering the network based on control information received from firewall nodes

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 847

1304. The operation of the system illustrated in Figure 13B will be described in more detail below.

The two designs depicted in Figures 13A and 13B are different based on what is distributed: packets or rules. For example, the design developed by Benecke et al consisted of multiple identical firewall nodes connected in parallel, as shown in Figure 13A. Each firewall node implements the complete security policy *R*, and arriving packets are distributed across the firewall nodes for processing (one packet is sent to one firewall node) [4]; therefore different packets are processed in parallel. Using terminology developed for parallel computing, this design is considered data-parallel since the data (packets) is distributed across the firewall nodes [25]. How the load-balancing algorithm distributes the packets is vital to the data-parallel system and typically implemented as a specialized high-speed switch [11, 12]. Although data-parallel firewalls have been shown to achieve higher throughput than traditional firewalls and have a redundant design [4], they suffer from two major disadvantages. First, these designs have difficulty maintaining QoS across networks. For example, legitimate traffic can encounter delays if it is queued behind traffic that requires more processing. Under these circumstances, users notice poor network performance, which is a growing concern as more network applications require QoS assurances. Second, stateful inspection requires all traffic from a certain connection or exchange to traverse the same firewall node, which is difficult to perform at high speeds using the data-parallel approach [4, 14].

Rule Distribution and Function-Parallel Firewall

As described above, new scalable firewall designs are needed to cope with increasing traffic loads, network speeds, and security threats. To address these issues, one implementation of the subject matter described herein includes a firewall design that consists of multiple firewalls connected in parallel and a gate machine, as shown in Figure 13B. Again in this design, the individual firewalls will be referred to as firewall nodes. However unlike the data-parallel design that distributes the packets, the function-parallel design distributes the policy rules across the firewall nodes. The distribution is done such that policy integrity is maintained. Therefore, a single traditional firewall

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 848

and a function-parallel firewall will always give the same result for any packet. The rules distributed to a firewall node will be referred to as the local rule set or local policy. Using parallel computing terminology this design will be referred to as function-parallel, since the rules are distributed across the firewall nodes [25].

The general operation of one exemplary function-parallel system can be described as follows. When a packet arrives to the function-parallel system, the packet is forwarded to every firewall node. Each firewall node processes the packet using its local rule list. The firewall node then signals the gate indicating either no match was found or provides the rule number and action if a match was found. Therefore, no match is a valid response and is required for this basic function-parallel design. The gate keeps track of the results and determines the final action to perform on the packet. For example, a first match policy can be implemented by applying the action of the lowest numbered matching rule to the packet.

The performance of this basic design can be improved by allowing the firewall nodes to pipeline packets [25]. For example, once a firewall node has completed processing a packet, it can immediately begin processing the next packet. This allows the firewall nodes to work independently, thus increasing the overall system performance. However, pipelining requires the gate to keep track of multiple packets simultaneously. Performance can also be increased if the gate can signal the firewall nodes that further processing of a certain packet is no longer needed (similar to short-circuit evaluation). This occurs when the appropriate match has been found by a firewall node before the other firewall nodes have completed processing the same packet. Short-circuit evaluation requires the gate to know how the rules are distributed as well as the dependencies, which is discussed in detail below.

The function-parallel design has several significant advantages over traditional and data-parallel firewalls. First, the function-parallel design results in the faster processing since every firewall node is utilized to process a single packet. Reducing the processing time, instead of the arrival rate (the data-parallel paradigm), yields better performance that will be proven theoretically and demonstrated experimentally. Second, unlike the data-parallel design, function-parallel can maintain state information about existing connections.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 849

Maintaining state can be viewed as the addition of a new rule corresponding to a requested connection. Unlike the data-parallel design, the new rule can be placed in any firewall node since a packet will be processed by every firewall node. Third, the function-parallel design has the unique ability to provide service differentiation between classes of traffic; therefore, the design can provide QoS requirements. Finally, the function-parallel system can be implemented using any current firewall system (as discussed below, the gate can be removed), without any specialized hardware. The disadvantage of the system is the limited scalability, since the system cannot have more firewall nodes than rules, unless data-redundancy is provided. However, given the size of most firewall policies range in the thousands to hundred of thousands of rules, which is expected to increase, the scalability limit is not an important concern.

The operation and performance of the function-parallel design depends on the correct distribution of policy rules. Rule distribution techniques that maintain integrity, increase processing speed, allow redundancy and robustness, and eliminate the gate are described in the following sections.

Rule Distribution and Policy Integrity

As described above, the function-parallel design distributes the security policy rules across an array of firewall nodes. Using this design, all the firewall nodes process an arriving packet in parallel using its local rule set. The rules must be distributed such that the integrity of the policy is maintained. Again, it will be assumed in the examples described herein that a first-match policy is desired. Therefore, the distribution of the rules must not introduce any anomalies into the original rule list, such as shadowing.

Let m be the number of firewall nodes in the system and j be a firewall node $j = 1...m$. Each firewall node has a local rule set that is a portion of the security policy. Denote the local rule list for firewall node j as R_j that has n_j rules. The firewall node will process an arriving packet using the local rule list in a top-down fashion to find the first match. It should be noted that a software implementation of the firewall node according to the subject matter described herein is not required to have a rule list. The function-parallel design is independent of software implementation. In this scenario integrity is maintained

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 850

under two conditions: every rule is assigned to at least one firewall node, and a policy DAG edge never traverses upward within a firewall node. The first condition assures every rule in the original policy exists in the function-parallel system, while the second condition prevents rule shadowing within a firewall node. Consider a two-firewall node system and the five-rule policy given in Table 2. Possible rule distributions are depicted in Figures 14A-14C. The distribution illustrated in Figure 14A will be referred to as vertical, while the distribution illustrated in Figure 14B will be referred to as horizontal. In both cases the five rules are present and all policy DAG edges are downward; thus integrity is maintained. In contrast, Figure 14C is an illegal distribution because rule r_4 shadows r_1 (upward policy DAG edge).

The gate can also utilize the policy DAG to determine the first match, and if a short-circuit evaluation can occur. For example consider the distribution illustrated in Figure 14A. Assume a packet matches r_3 in the second firewall node and this result is reported to the gate. Using the policy DAG, the gate can determine this is the first match, since it is impossible for a packet that matches r_3 to also match any rules that preceded it in the original policy (rules r_1 or r_2). The gate can inform the remaining firewall nodes to stop processing the packet.

Optimal Rule Distribution

The previous section described the distribution requirements for maintaining policy integrity. However it was proven in [27] that any linear arrangement of a policy DAG maintains integrity, which also holds for a local policy DAG. Furthermore as evident in the previous section, many different distributions are possible. Given the possible rule distributions for the function-parallel design, a distribution is sought that minimizes the average number of comparisons required per packet [27].

Achieving this goal will require information not present in the firewall rule list. Certain firewall rules have a higher probability of matching a packet than others. As a result, it is possible to develop a policy profile over time that indicates frequency of rule matches (similar to cache hit ratio). Let $P = \{p_1, p_2, \ldots, p_n\}$ be the policy profile, where p_i be the probability that a packet will match rule r_i (first match in the policy). Furthermore, assume a packet will always find

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 851

a match, $\sum_{i=1}^{n} p_i = 1$; therefore R is comprehensive. Using this information, the average number of rule comparisons required for the policy R using a standard single machine firewall is

$$E(R) = \sum_{i=1}^{n} i \cdot p_i \quad (4)$$

This equation also applies to firewall nodes. However, the rules of a local policy include rules of the original policy; therefore, a mapping between the local and original policy rules is needed. Let $r_{i,j}$ represent the i^{th} rule of firewall node j, that corresponds a certain rule in the original rule list R. Consider the local rule list for the second firewall node depicted in Figure 14A. The local policy is $R_2 = \{r_4, r_5\}$, where $r_{1,2} = r_4$ and $r_{2,2} = r_5$. Similarly, let $p_{i,j}$ be the probability of matching the i^{th} rule of firewall node j. Therefore the average number of comparisons required for firewall node j is

$$E(R_j) = \sum_{i=1}^{n_j} i \cdot p_{i,j} \quad (5)$$

Using the previous equation, the average number of comparisons required for the function-parallel system is

$$E_f(R) = \sum_{j=1}^{m} E(R_j) = \sum_{j=1}^{m} \sum_{i=1}^{n_j} i \cdot p_{i,j} \quad (6)$$

For example, the average number of comparisons required for the distribution given in Figure 14A is 2.24, while the average number of comparisons for the distribution given in Figure 14B is 2.45.

Given a policy DAG $G = (R,E)$ and policy profile P, a distribution is sought that minimizes Equation 6. In the absence of precedence relationships, the average number of comparisons is minimized if the rules are sorted in non-increasing order according to the probabilities [28], which is also referred to as Smith's algorithm [29]. The sorted rules are then distributed in a horizontal fashion, as depicted in Figure 14B. Precedence constraints cause the problem

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 852

to be more realistic; however, precedence constraints also make determining the optimal permutation more problematic. In [27] determining the optimal rule order for a single machine firewall is proven to be *NP*-complete; therefore, so is determining the optimal function-parallel rule distribution. However, [27] presents a simple sorting algorithm to improve the optimal order, that can be modified to apply to function-parallel rule distribution. This modified algorithm sorts the policy rules using the method described in [27], then distributes the rules in a horizontal fashion.

Redundancy

Similar to the traditional and data-parallel firewalls, a disadvantage of the current function-parallel design is its inability to withstand a single firewall node failure. If a single firewall node fails in the function-parallel system, then part of the policy is lost and integrity is not maintained. A simple solution commonly used for traditional firewalls is to duplicate the entire system; however, this solution is cost prohibitive, not efficient, and difficult to manage.

Redundancy can be provided in the function-parallel design by replicating rules, instead of firewall nodes. Consider the function-parallel system depicted in Figure 15A, that includes four firewall nodes and an eight-rule policy distributed in a horizontal manner. If the local rule list of the right neighbor is added to the end of every firewall node, as illustrated in Figure 15B, then any non-continuous firewall nodes can fail and integrity will be maintained. It should be noted that the criteria for maintaining integrity still apply: every rule must be implemented at a node and no upward edges in the local policy DAG. Furthermore, the redundant design requires short-circuit evaluations to prevent multiple firewall nodes from evaluating the same rule.

Gateless Design

The gate plays an important role in the function-parallel design. It accepts information from the firewall nodes to determine the correct action to apply to a packet. Unfortunately, the gate is a special device that must coordinate with the individual firewall nodes in the system, which may be problematic to implement. However, it is possible to eliminate the gate and allow the firewall nodes to operate independently without any coordination. The

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 853

significant advantage is that the gateless function-parallel design can be implemented with any firewall system (from open source to proprietary).

Eliminating the gate requires a rule distribution such that for any given packet, only one machine will ever accept and the remaining machines will always deny. As described above, rule actions are either accept or deny. This can be seen in the policy DAG by denoting deny rules with a square and accept rules with a circle. To assure only one machine will accept a packet, an accept rule in a firewall node cannot intersect with another accept rule in any other firewall node. As described above, the intersection of two rules describes the set of packets that can match both rules. If an edge spans the local rules of firewall nodes, then the set of packets would be accepted by both firewall nodes; thus, duplicate packets would be accepted. In addition, every drop rule must be present in each firewall node to assure the remaining firewall nodes will drop the packet. Given these two requirements a policy DAG edge will not span firewall nodes, as illustrated in Figures 16B and 16C.

For example, consider a two-firewall node function-parallel system that implements the five-rule policy given in Table 2. The policy DAG is given in Figure 16A, where the first and last rules are deny. An initial gateless rule distribution is given in Figure 16B which has no edges spanning the firewall nodes. Given the local policy DAG, deny rules can be removed if: no inbound edges exist for the deny rule; and all out-going edges for the deny rule connect to a superset deny rule. Using these constraints, the first rule in the second firewall node depicted in Figure 16B can be removed (it is a subset of r_5 and has no other edges). This gives the final gateless rule distribution shown in Figure 16C. As a result, the first firewall node in this distribution will only process UDP traffic, while the second machine only processes TCP. The rules can be divided with respect to other tuples, and can be done to provide QoS guarantees.

Theoretical Results

Assuming arrivals and service times are exponentially distributed, a firewall system can be considered an open network of M/M/1 queues (Jackson network) [5, 17]. Probabilities can be assigned to each link to indicate the likelihood of moving to the next node, which are given from the policy profile (hit

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 854

ratio) described above. The average end-to-end delay for q cascading machines (traversal path) can be computed as

$$E(T) = \sum_{i=1}^{q} \frac{1}{\mu_i - \lambda_i} \tag{7}$$

where $1/\mu_i$ is the service time (processing and transmission) and λi is the arrival rate to node i. As a result, we have a theoretical model for the average delay across a firewall system. Consider the data-parallel and functional-parallel firewall designs given in Figures 13A and 13B. Assume each system consists of m machines and implements the same n rule security policy. Let the total arrival rate to each system be λ packets per unit time and each node performs x rules per unit time.

For the data-parallel firewall, traffic arrives at the packet distributor, which evenly distributes the traffic. As a result the arrival rate to each machine is λ m. The service rate is x n, since each node implements the complete rule set. The end-to-end delay across any machine in the parallel firewall is

$$E_d(T) = \frac{1}{\frac{x}{n} - \frac{\lambda}{m}} \tag{8}$$

In contrast, rules are distributed across each machine, in the function-parallel firewall. Furthermore, all traffic arriving to the system is forwarded to each machine. If we assume the rules in each machine are independent (no policy DAG edges spanning machines), then the end-to-end delay across any machine in the functional-parallel firewall is

$$E_f(T) = \frac{1}{\frac{m \cdot x}{n} - \lambda} \tag{9}$$

The reduction in delay compared to a data-parallel firewall system is 1 m; therefore, the functional-parallel system has the potential to be m times faster than a (data) parallel system. However, the scalability of the functional-parallel system is dependent on the rule set.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 855

Figures 17A and 17B compare the performance of function-parallel, data-parallel, and single firewall designs. More particularly, Figure 17A compares the average among the different firewall designs as the arrival rate λ increases. As can be seen from Figure 17A, the function-parallel design achieves lower average delay as the arrival rate increases than the data-parallel design. Figure 17B illustrates the average delay as arrival rate and the number of rules increase. In Figure 17B, the function-parallel design again achieves better simulated performance than the data-parallel design.

Experimental Results

The previous section described models that provide insight to the average delay across different firewall designs. This information is beneficial in guiding initial design; however, we are also interested in the maximum delay encountered. For this reason, a discrete event simulator has been developed to model the different systems.

The data-parallel firewall and function-parallel firewall designs were simulated, where both systems implemented a 50-rule policy and each consisted of 5 firewall nodes. Three classes of traffic were modeled, first, second, and background. The first and second classes represent legitimate traffic that was eventually accepted by the firewall system, while background traffic was typically dropped. The two designs were simulated under normal conditions, where each traffic type had the same volume. Figures 18A and 18B shows the simulation results. As seen in Figure 18A and predicted by the theoretical models, the average delay encountered by packets in the functional-parallel design was lower than the data-parallel designs. The function-parallel design also provided lower maximum delays than the data-parallel, as seen in Figure 18B. In summary, the theoretical models and simulation results indicate function-parallel is a better design than a traditional single firewall system or a data-parallel system; yet, function-parallel requires no more computing resources than data-parallel.

Summary

Thus, as described above, the subject matter described herein includes methods, systems, and computer program products for providing a function-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 856

2005328336 22 Aug 2011

parallel firewall where different portions of a firewall rule set are distributed among different firewall nodes. Figure 19 is a flow chart illustrating exemplary steps of a process for filtering packets using a rule set distributed among firewall nodes in a function-parallel manner. Referring to Figure 19, in step **1900**, packet filtering rules of a rule set are distributed among nodes in a function-parallel manner so that different nodes implement different portions of the rule set. In step **1902**, packets to be filtered are replicated to each of the nodes. In step **1904**, the packets to be filtered are replicated to each of the nodes. In a hierarchical design, step **1904** may be implemented at each level in the hierarchy. If a packet passes or fails at one level, the packet may be forwarded into a network or dropped, depending on the distribution of the rule set. In step **1904**, the rules are applied to filter the packets.

The steps illustrated in Figure 19 and in all of the examples described herein may be applied to firewall systems that filter packets based on packet header information, intrusion detection systems that filter packets based on packet payload information, or intrusion protection systems that filter packets based on a combination of packet header and packet payload information.

In the claims which follow and in the preceding description of the invention, except where the context requires otherwise due to express language or necessary implication, the word "comprise" or variations such as "comprises" or "comprising" is used in an inclusive sense, i.e. to specify the presence of the stated features but not to preclude the presence or addition of further features in various embodiments of the invention.

It is to be understood that, if any prior art publication is referred to herein, such reference does not constitute an admission that the publication forms a part of the common general knowledge in the art, in Australia or any other country.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 857

2005328336 22 Aug 2011

The following references correspond to the citations in this document. The disclosure of each of these references is hereby incorporated herein by reference it its entirety.

References

[1] A. V. Aho, J.E. Hopcroft, and J. D. Ullman. *Data Structures and Algorithms*. Addison-Wesley, 1987.

[2] M. Al-Suwaiyel and E. Horowitz. Algorithms for Trie Compaction. *ACM Transactions on Database Systems*, 9(2):243–263, June 1984.

[3] S. M. Bellovin and W. Cheswick. Network Firewalls. *IEEE Communications Magazine*, pages 50–57, Sept. 1994.

[4] C. Benecke. A Parallel Packet Screen for High Speed Networks. In *Proceedings of the 15th Annual Computer Security Applications Conference*, 1999.

[5] G. Bolch, S. Greiner, H. de Meer, and K. S. Trivedi. *Queueing Networks and Markov Chains*. John Wiley and Sons, Inc., 1998.

[6] D. Comer. Analysis of a Heuristic for Full Trie Minimization. *ACM Transactions on Database Systems*, 6(3):513–537, Sept. 1981.

[7] U. Ellermann and C. Benecke. Firewalls for ATM Networks. In *Proceedings of INFOSEC'COM*, 1998.

[8] D. Eppstein and S. Muthukrishnan. Internet Packet Filter Management and Rectangle Geometry. In *Proceedings of the Symposium on Discrete Algorithms*, pages 827–835, 2001.

[9] E. Fulp, Z. Fu, D. S. Reeves, S. F. Wu, and X. Zhang. Preventing Denial of Service Attacks on Quality of Service. In *Proceedings of the 2001 DARPA Information Survivability Conference and Exposition II*, 2001.

[10] R. Funke, A. Grote, and H.-U. Heiss. Performance Evaluation of Firewalls in Gigabit-Networks. In *Proceedings of the Symposium on Performance Evaluation of Computer and Telecommunication Systems*, 1999.

[11] X. Gan, T. Schroeder, S. Goddard, and B. Ramamurthy. LSMAC vs. SLNAT: Scalable ClusterbasedWeb Servers. *Journal of Networks, Software Tools, and Applications*, 3(3):175–185, 2000.

[12] S. Goddard, R. Kieckhafer, and Y. Zhang. An Unavailability Analysis of Firewall Sandwich Configurations. In *Proceedings of the 6th IEEE Symposium on High Assurance Systems Engineering*, 2001.

[13] A. Hari, S. Suri, and G. Parulkar. Detecting and Resolving Packet Filter Conflicts. In *Proceedings of IEEE INFOCOM*, pages 1203–1212, 2000.

[14] O. Paul and M. Laurent. A Full Bandwidth ATM Firewall. In *Proceedings of the 6th European Symposium on Research in Computer Security ESORICS'2000*, 2000.

[15] L. Qui, G. Varghese, and S. Suri. Fast Firewall Implementations for Software and Hardware-Based Routers. In *Proceedings of ACM SIGMETRICS*, June 2001.

[16] G. V. Rooij. Real Stateful TCP Packet Filtering in IP Filter. In *Proceedings of the 10th USENIX Security Symposium*, 2001.

[17] M. Schwartz. *Telecommunication Networks: Protocols, Modeling, and Analysis*. Addison-Wesley, 1987.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 859

[18] S. Suri and G. Varghese. Packet Filtering in High Speed Networks. In *Proceedings of the Symposium on Discrete Algorithms*, pages 969 – 970, 1999.

[19] P. Warkhede, S. Suri, and G. Varghese. Fast Packet Classification for Two-Dimensional Conflict-Free Filters. In *Proceedings of IEEE INFOCOM*, pages 1434–1443, 2001.

[20] J. Xu and M. Singhal. Design and Evaluation of a High-Performance ATM Firewall Switch and Its Applications. *IEEE Journal on Selected Areas in Communications*, 17(6):1190–1200, June 1999.

[21] R. L. Ziegler. *Linux Firewalls*. New Riders, second edition, 2002.

[22] E. D. Zwicky, S. Cooper, and D. B. Chapman. *Building Internet Firewalls*. O'Reilly, 2000.

[23] E. Al-Shaer and H. Hamed. Modeling and Management of Firewall Policies. *IEEE Transactions on Network and Service Management*, 1(1), 2004.

[24] S. M. Bellovin and W. Cheswick. Network Firewalls. *IEEE Communications Magazine*, pages 50–57, Sept. 1994.

[25] D. E. Culler and J. P. Singh. *Parallel Computer Architecture: A Hardware/Software Approach*. Morgan Kaufman, 1999.

[26] E. W. Fulp. Firewall Architectures for High Speed Networks. Technical Report 20026, Wake Forest University Computer Science Department, 2002.

[27] E. W. Fulp. Optimization of Firewall Policies Using Directed Acyclical Graphs. Submitted for publication, 2004.

[28] R. Rivest. On Self-Organizing Sequential Search Heuristics. *Communications of the ACM*, 19(2), 1976.

[29] W. E. Smith. Various Optimizers for Single-Stage Production. *Naval Research Logistics Quarterly*, 3:59 – 66, 1956.

It will be understood that various details of the present subject matter may be changed without departing from the scope of the present subject matter. Furthermore, the foregoing description is for the purpose of illustration only, and not for the purpose of limitation.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 860

2005328336 29 Jun 2007

CLAIMS

What is claimed is:

1. A function-parallel firewall system comprising:

(a) a first firewall node for filtering received packets using a first portion of a rule set including a plurality of rules, the first portion including less than all of the rules in the rule set; and

(b) at least one second firewall node for filtering packets using a second portion of the rule set, the second portion including at least one rule in the rule set that is not present in the first portion, wherein the first and second portions together include all of the rules in the rule set, wherein the first and second firewall nodes are configured to implement a gateless design wherein the rules are distributed among the first and second firewall nodes such that for a given packet, one of the first and second firewall nodes will accept the packet and the other of the first and second firewall nodes will deny the packet.

2. The function-parallel firewall system of claim 1 wherein the first and second portions respectively include rules present in first and second levels of a firewall hierarchy.

3. The function-parallel firewall system of claim 2 wherein the first and second firewall nodes are connected to each other such that the rules in the first level are applied to received packets before the rules in the second level.

4. The function-parallel firewall system of claim 3 wherein the first firewall node is adapted to forward packets that pass one of the rules in the first portion to an internal network in a manner that bypasses the second firewall node.

5. The function-parallel firewall system of claim 3 wherein the first firewall node is adapted to discard packets that fail one of the rules in the first portion.

6. The function-parallel firewall system of claim 3 wherein the rules in the second portion are more processor-intensive than the rules in the first portion.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 861

7. The function-parallel firewall system of claim 1 wherein the first and second firewall nodes implement rules of the same hierarchical level.

8. The function-parallel firewall system of claim 1 wherein the first and second firewall nodes are adapted to simultaneously process copies of the same packets.

9. The function-parallel firewall system of claim 1 wherein the rule set includes intrusion detection rules for filtering packets based on packet payloads.

10. The function-parallel firewall system of claim 1 wherein the rule set includes intrusion protection rules for filtering packets based on a combination of packet headers and packet payloads.

11. The function-parallel firewall system of claim 1 wherein the first and second firewall nodes are components of a first firewall subsystem for filtering packets entering a network and wherein the system further comprises a second function-parallel firewall sub-system including a second set of firewall nodes implementing a second rule set for filtering packets departing from the network, wherein the second function-parallel firewall subsystem includes at least third and fourth firewall nodes, the third and fourth firewall nodes implementing first and second portions of the second rule set, the first portion of the second rule set including at least one rule that is not present in the second portion of the second rule set, wherein the first and second function-parallel firewall subsystems share state information regarding connections established through the first and second sets of firewall nodes.

12. The function-parallel firewall system of claim 1 wherein the first and second firewall nodes are components of a firewall grid with logically configurable interconnections between the firewall nodes.

13. A method for controlling access to a network based on a set of packet filtering rules distributed in a function-parallel manner, the method comprising:

(a) distributing packet filtering rules of a rule set among a plurality of different nodes in a function-parallel manner so that at least some of the different nodes implement different portions of the rule set;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 862

2005328336 22 Aug 2011

(b) replicating packets to each of the nodes; and

(c) applying the rules to filter the packets using a gateless design wherein the rules are distributed among first and second firewall nodes such that for a given packet, one of the first and second firewall nodes will accept the packet and the other of the first and second firewall nodes will deny the packet.

14. The method of claim 13 wherein the packet filtering rules comprise intrusion detection rules for filtering packets based on packet payload information.

15. The method of claim 13 wherein the packet filtering rules comprise intrusion protection rules for filtering packets based on a combination of packet header and packet payload information.

16. The method of claim 13 comprising, receiving a first packet at a first node, determining whether the packet matches a first rule implemented by the first node, and, in response to determining that a match occurs, notifying the other nodes of the match.

17. The method of claim 16 comprising, at the other nodes, ceasing processing of the first packet.

18. A computer program product comprising computer executable instructions embodied in a computer readable medium for performing steps comprising:

(a) a first firewall node for filtering received packets using a first portion of a rule set including a plurality of rules, the first portion including less than all of the rules in the rule set; and

(b) at least one second firewall node for filtering packets using a second portion of the rule set, the second portion including at least one rule in the rule set that is not present in the first portion, wherein the first and second portions together include all of the rules in the rule set, wherein the first and second firewall nodes are configured to implement a gateless design

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 863

2005328336 22 Aug 2011

wherein the rules are distributed among the first and second firewall nodes such that for any given packet, one of the first and second firewall nodes will accept the packet and the other of the first and second firewall nodes will deny the packet

19. The method of claim 13, or computer program product of claim 18, wherein distributing the rules in the rule set among a plurality of firewall nodes comprises:

(a) assigning rules in the rule set to nodes and branches in a trie data structure, wherein each node in the trie data structure corresponds to a data field and each branch represents a value for each data field to be compared to values in corresponding fields in received packets; and

(b) pruning the trie data structure in a manner that preserves ordering of the rules in the rule set.

20. The method or product of claim 19 wherein pruning the trie data structure includes pushing a rule to its descendants and deleting the original rule.

21. The method or product of claim 19 comprising reducing the depth of the trie using level reordering.

22. The method or product of claim 19 wherein assigning rules in the rule set to nodes and branches in the trie data structure includes assigning rules requiring longer processing time to lower levels in the trie.

23. The method of claim 13 or computer program product of claim 18, wherein distributing the rules in the rule set among a plurality of firewall nodes comprises:

(a) assigning packet filtering rules to nodes in a directed acyclical graph (DAG);

(b) representing relationships between rules by edges in the DAG; and

(c) distributing the rules among the plurality of firewall nodes using the DAG such that different firewall nodes implement different rules and such that relationships between the rules specified in the DAG are

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 864

2005328336 22 Aug 2011

preserved.

24. The method or product of claim 23, wherein distributing the rules among the plurality of firewall nodes using the DAG includes distributing the rules such that intersecting accept rules are not present in more than one firewall node.

25. The method or product of claim 23, wherein distributing the rules among the plurality of firewall nodes using the DAG includes distributing the rules such that the edges to not span multiple firewall nodes.

26. The method or product of claim 23, wherein distributing the rules among the plurality of firewall nodes using the DAG includes replicating deny rules in each of the firewall nodes.

27. The method or product of claim 26, comprising eliminating a deny rule implemented by a first firewall node in response to determining that the deny rule for the first firewall node includes no inbound edges and outbound edges from the first deny rule terminate only in a superset deny rule.

28. A function-parallel firewall system, a method for controlling access to a network, or a computer program product, substantially as herein described with reference to Figures 2 to 19.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 865

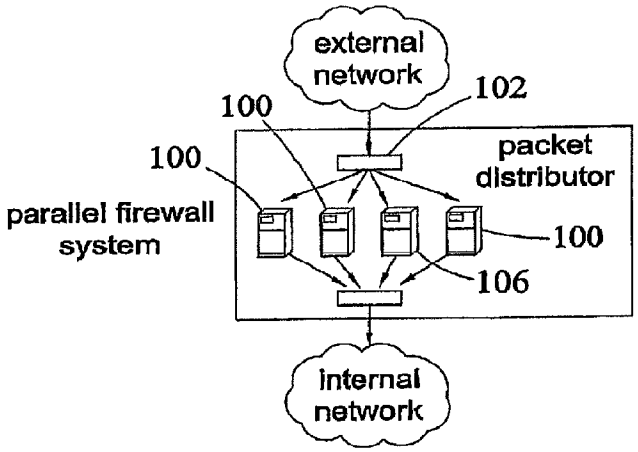

FIG. 1
(PRIOR ART)

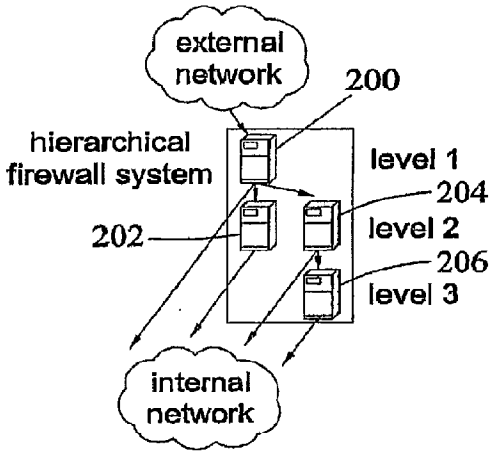
external network
200
hierarchical firewall system
level 1
204
level 2
202
206
level 3
internal network

FIG. 2

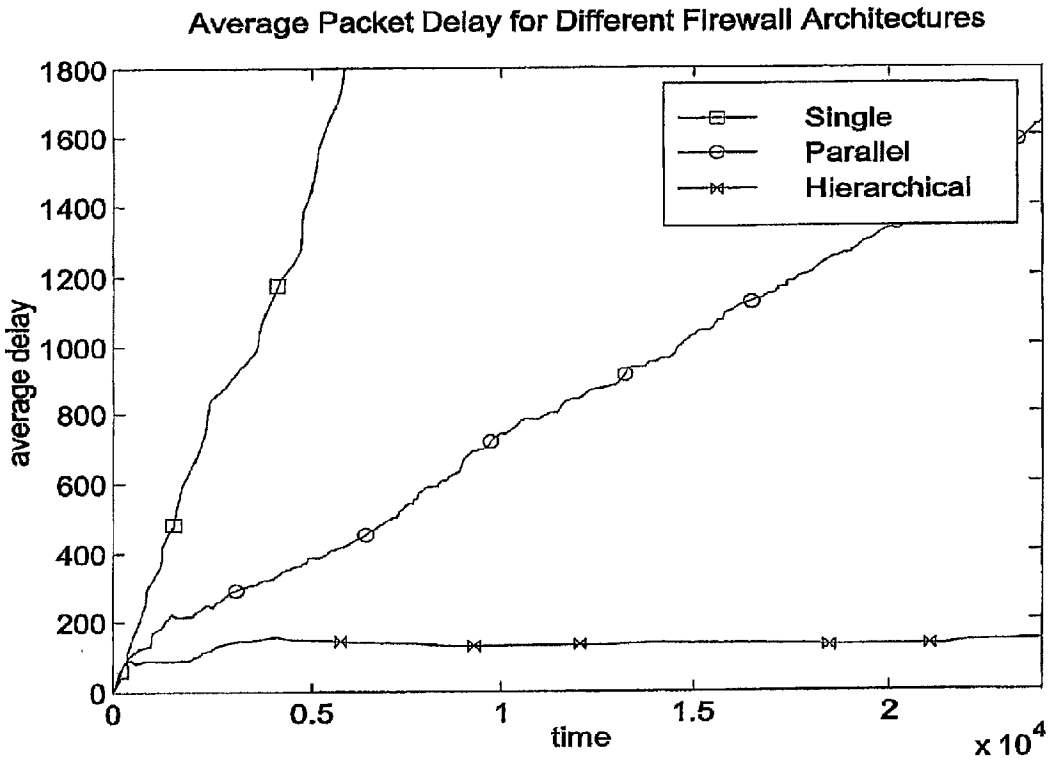
Average Packet Delay for Different Firewall Architectures
average delay
time
x 10^4
Single
Parallel
Hierarchical
0
200
400
600
800
1000
1200
1400
1600
1800
0
0.5
1
1.5
2

FIG. 3A

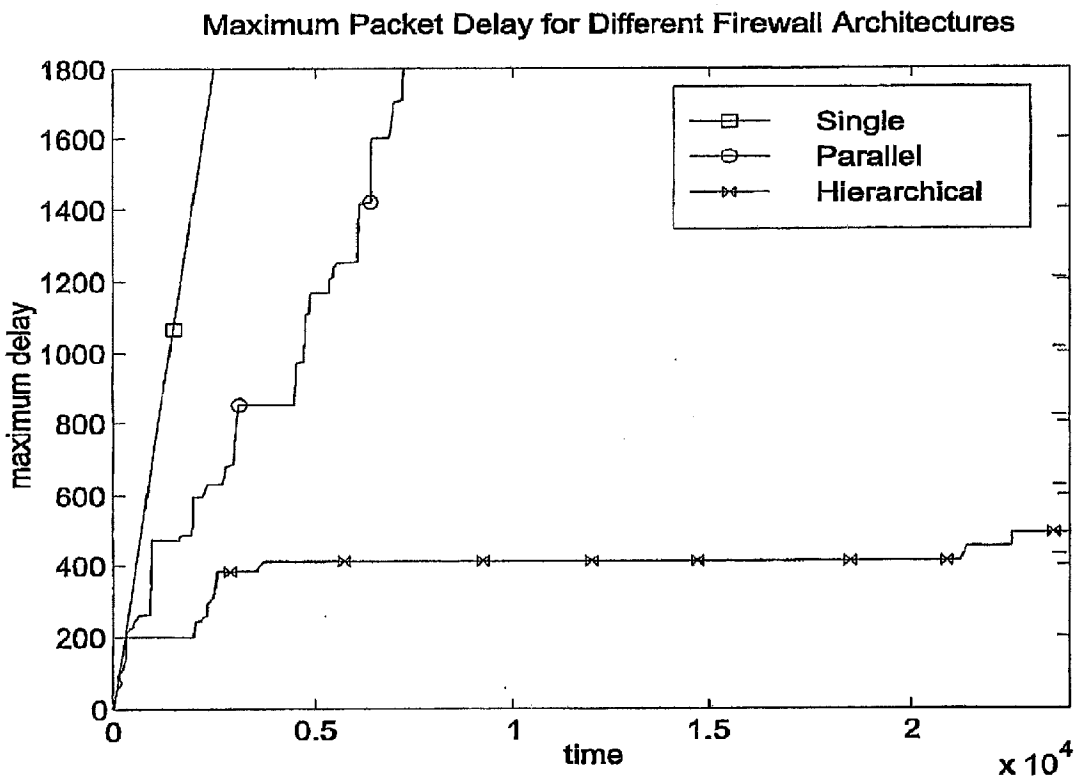
Maximum Packet Delay for Different Firewall Architectures
maximum delay
time
x 10^4
Single
Parallel
Hierarchical
0
200
400
600
800
1000
1200
1400
1600
1800
0
0.5
1
1.5
2

FIG. 3B

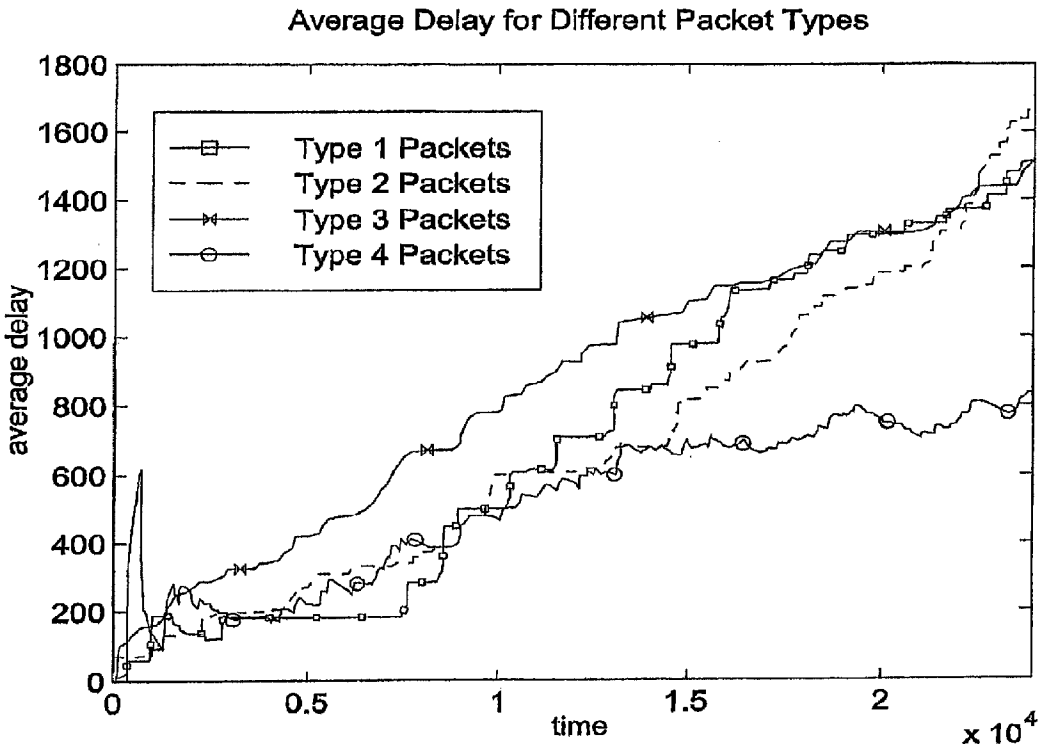
Average Delay for Different Packet Types
Type 1 Packets
Type 2 Packets
Type 3 Packets
Type 4 Packets
average delay
time
x 10^4
0
200
400
600
800
1000
1200
1400
1600
1800
0.5
1
1.5
2

FIG. 4A

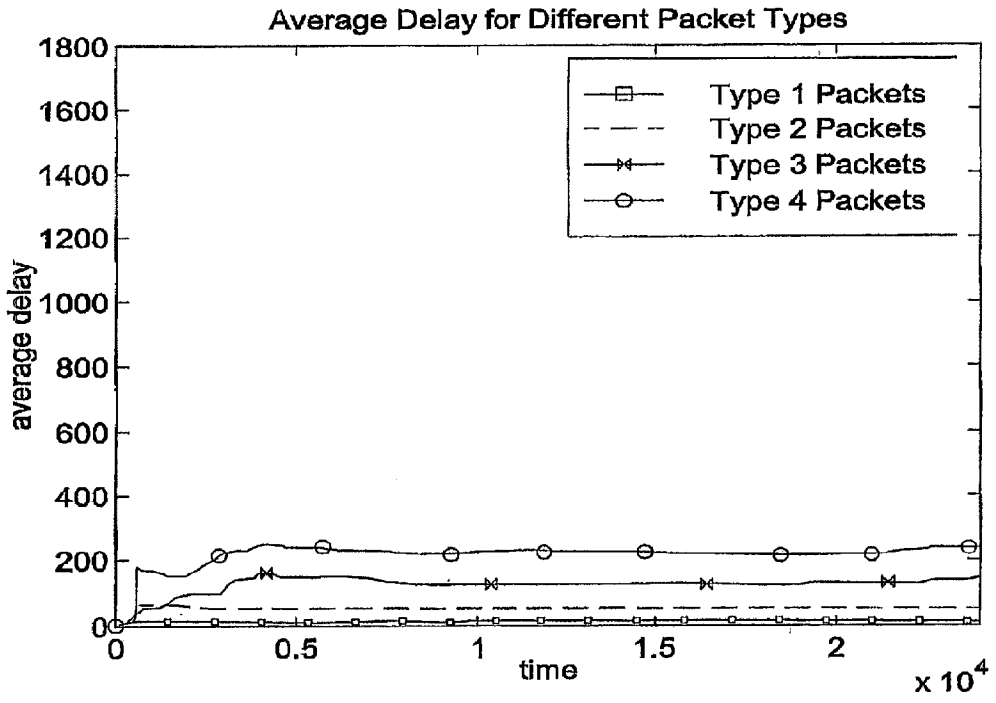
Average Delay for Different Packet Types
1800
1600
1400
1200
1000
800
600
400
200
0
average delay
0
0.5
1
1.5
2
time
x 10^4
Type 1 Packets
Type 2 Packets
Type 3 Packets
Type 4 Packets

FIG. 4B

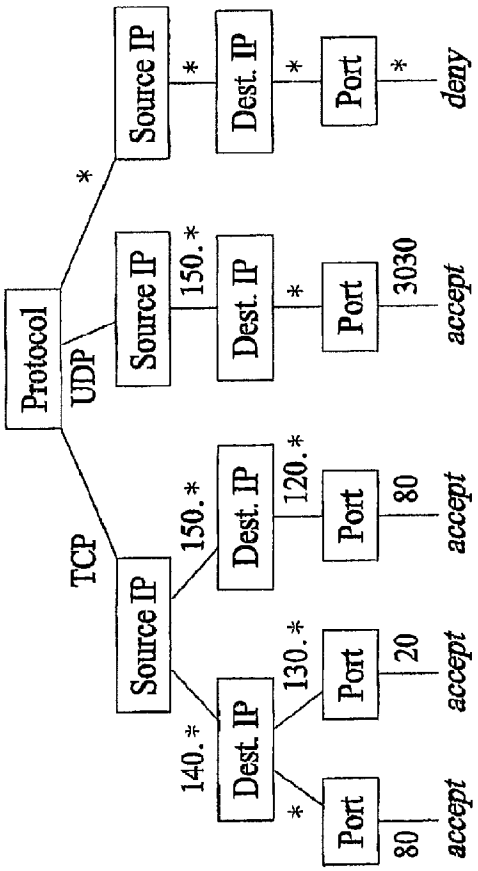
Protocol
TCP
UDP
*
Source IP
140.*
150.*
*
Dest. IP
130.*
120.*
*
Port
20
80
*
accept
deny
150.*
3030

FIG. 5

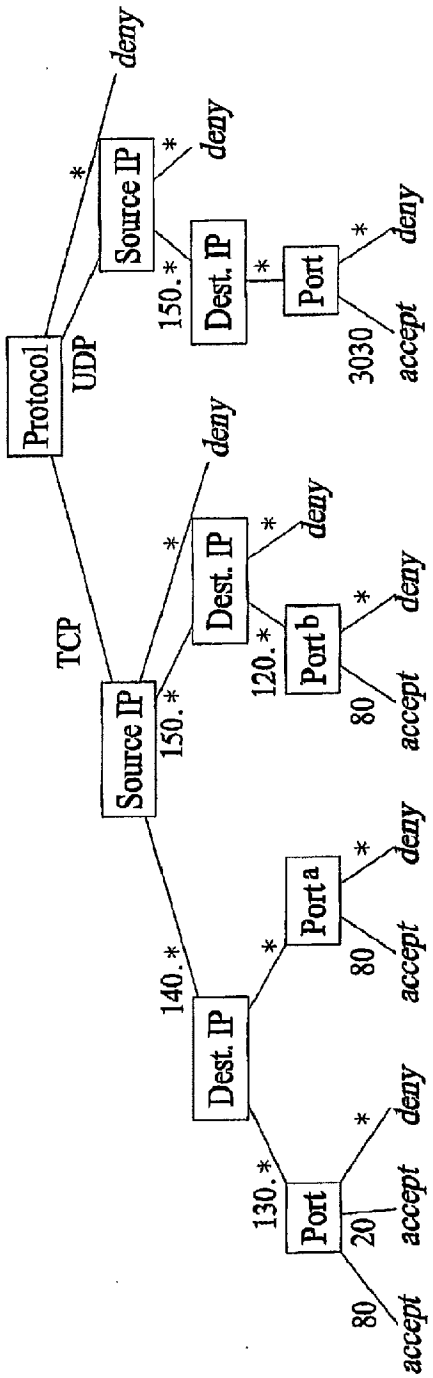
Protocol
TCP
UDP
*
Source IP
140.*
150.*
*
Dest. IP
130.*
120.*
*
Port
Port a
Port b
20
80
*
accept
deny
150.*
3030

FIG. 6

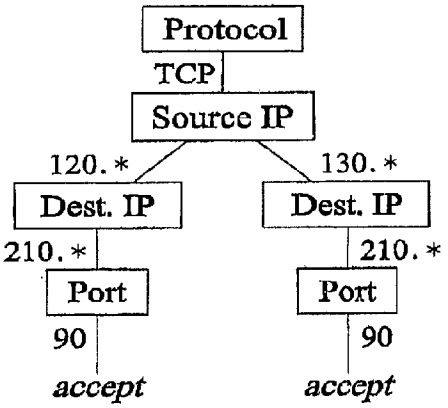
Protocol
TCP
Source IP
120.*
130.*
Dest. IP
Dest. IP
210.*
210.*
Port
Port
90
90
accept
accept

FIG. 7A

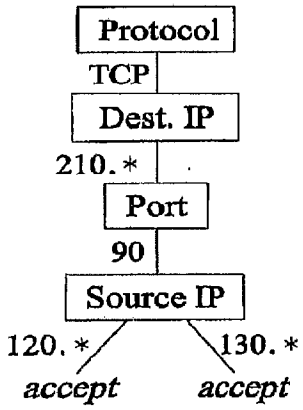
Protocol
TCP
Dest. IP
210.*
Port
90
Source IP
120.*
130.*
accept
accept

FIG. 7B

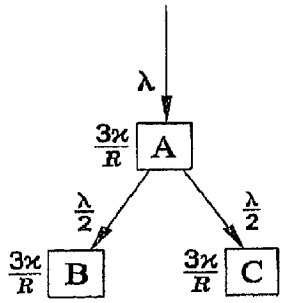
λ
3κ/R
A
λ/2
λ/2
3κ/R
B
3κ/R
C

FIG. 8A

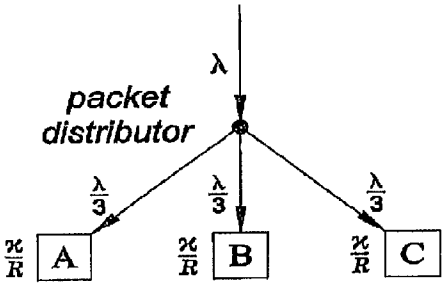
λ
packet distributor
λ/3
λ/3
λ/3
κ/R
A
κ/R
B
κ/R
C

FIG.8B

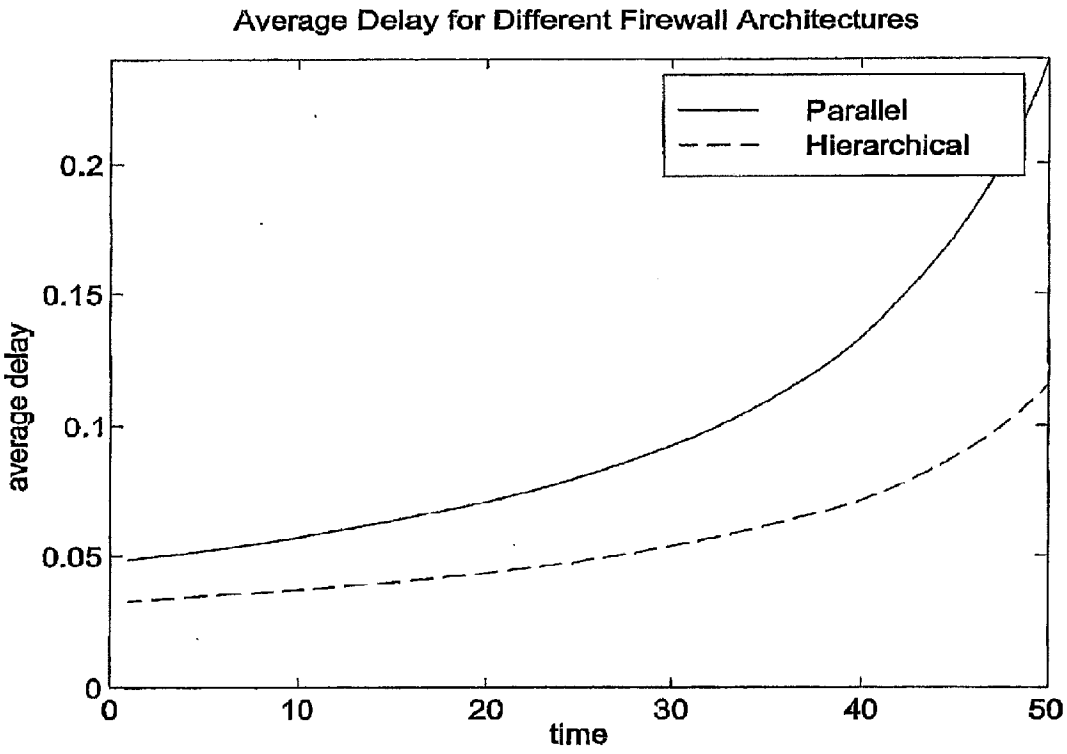
Average Delay for Different Firewall Architectures
Parallel
Hierarchical
average delay
time
0
0.05
0.1
0.15
0.2
0
10
20
30
40
50

FIG. 9

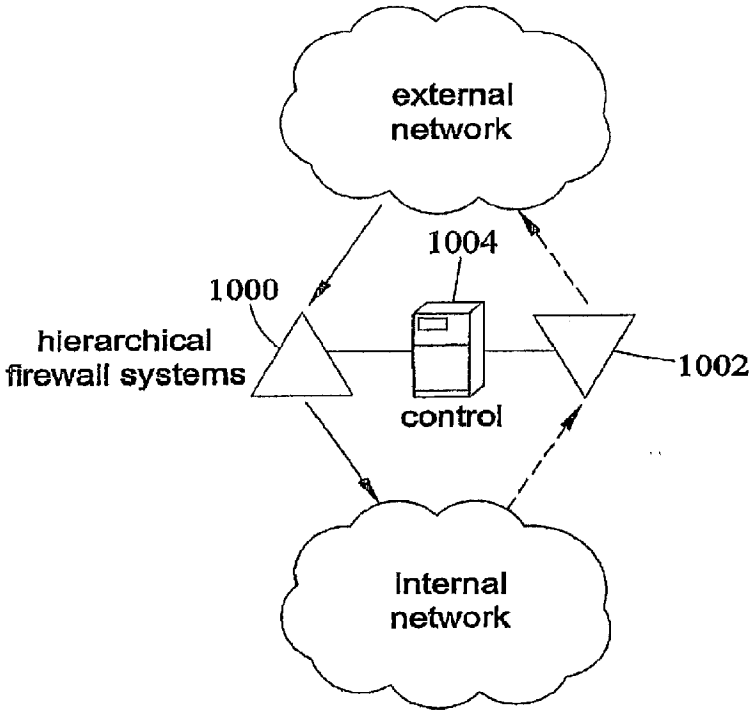
external network
1004
1000
hierarchical firewall systems
1002
control
internal network

FIG. 10

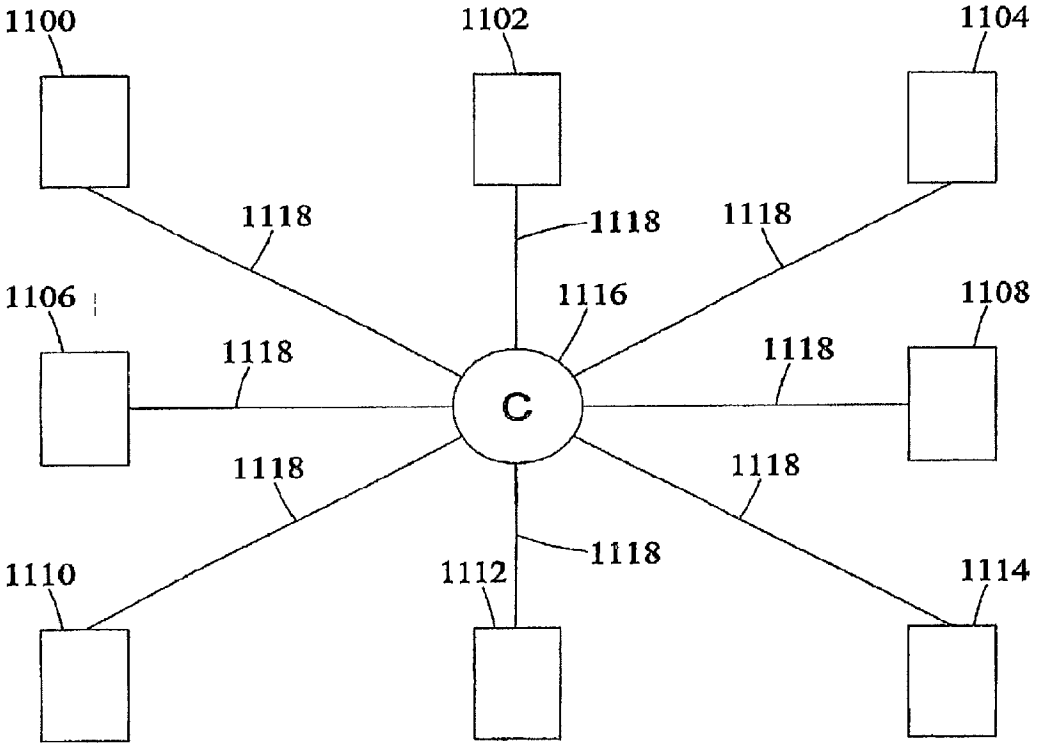
1100
1102
1104
1118
1118
1118
1106
1116
1108
1118
1118
C
1118
1118
1110
1112
1118
1114

FIG. 11A

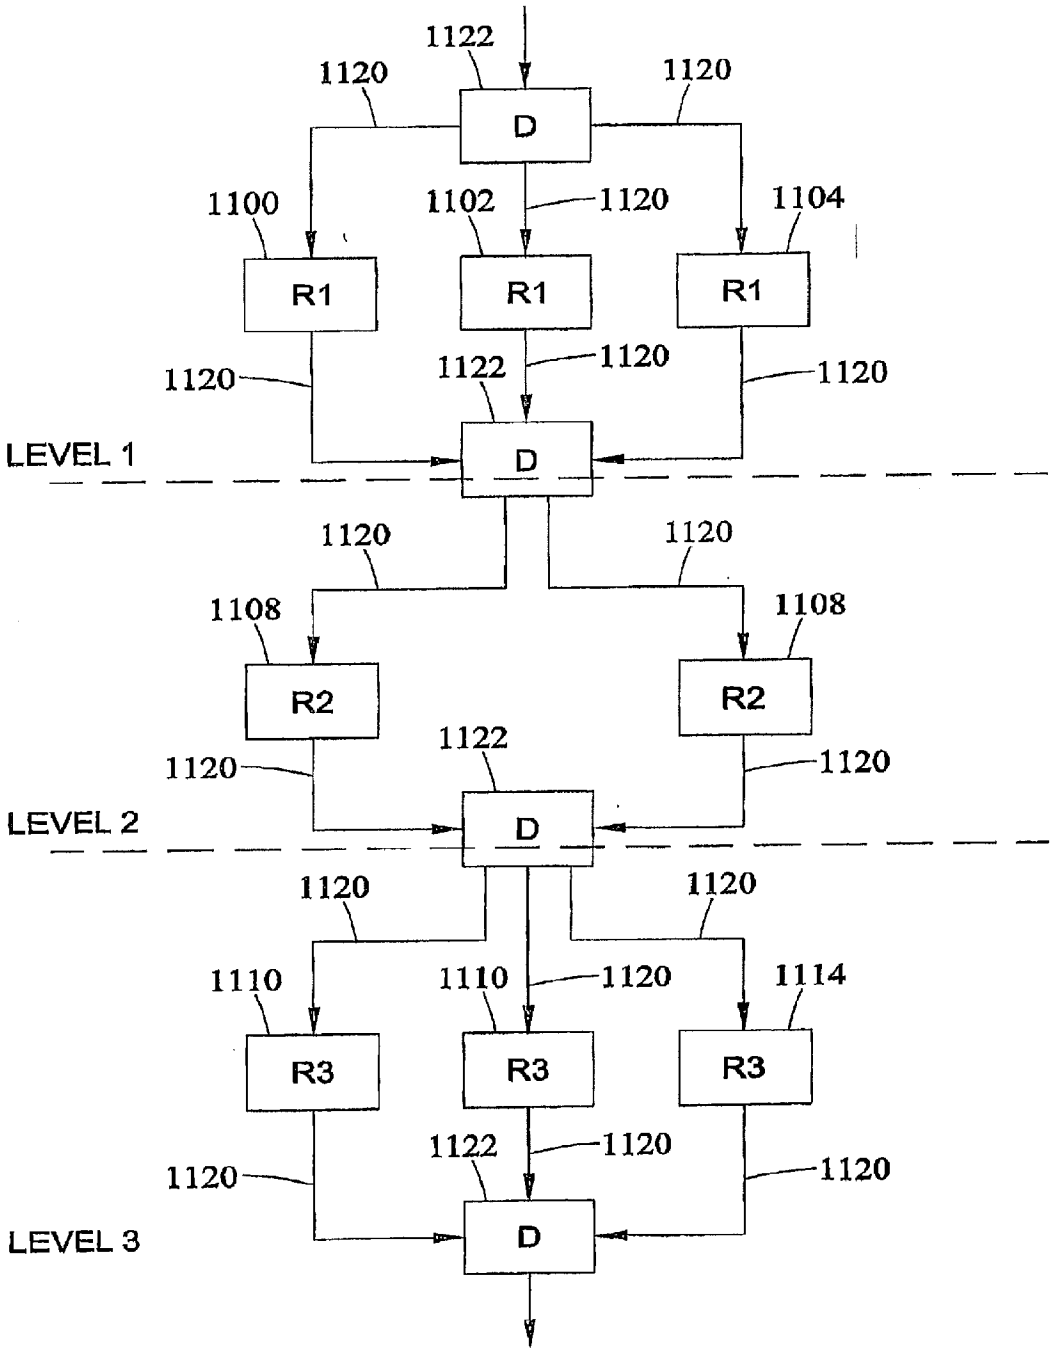
1122
1120
1120
D
1100
1102
1120
1104
R1
R1
R1
1120
1122
1120
1120
LEVEL 1
D
1120
1120
1108
1108
R2
R2
1122
1120
1120
LEVEL 2
D
1120
1120
1110
1110
1120
1114
R3
R3
R3
1122
1120
1120
1120
LEVEL 3
D

FIG. 11B

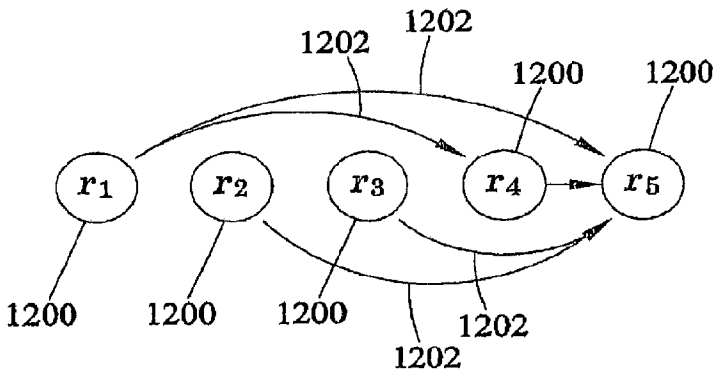

FIG. 12

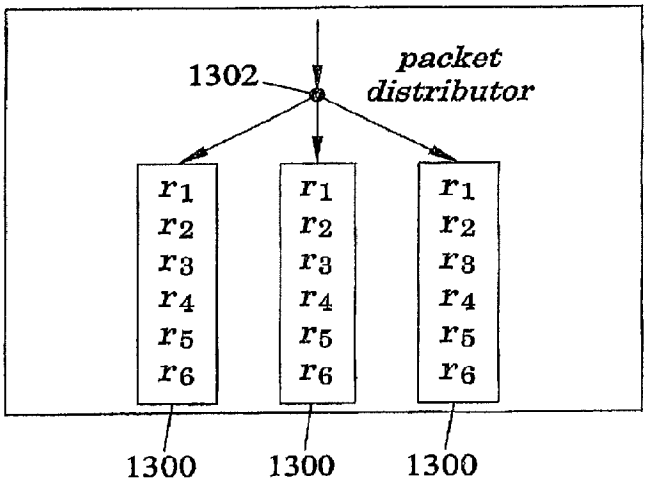

FIG. 13A

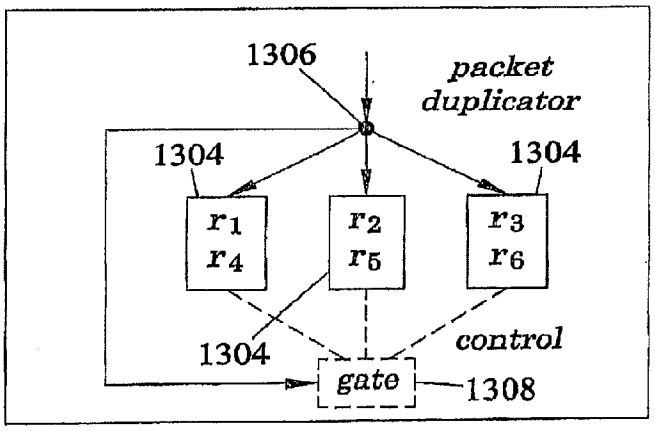

FIG. 13B

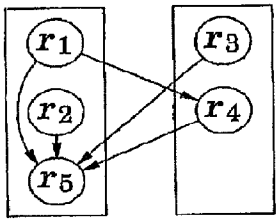

FIG. 14A

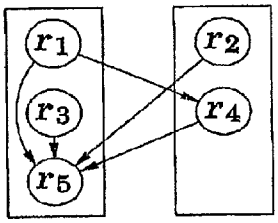

FIG. 14B

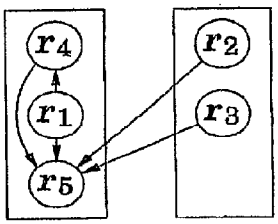

FIG. 14C

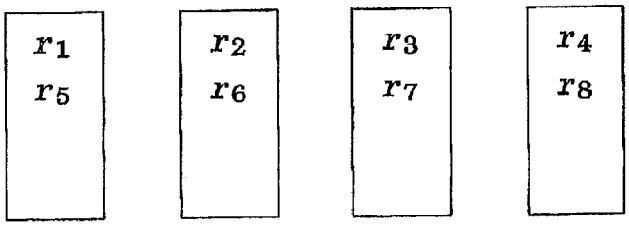

FIG. 15A

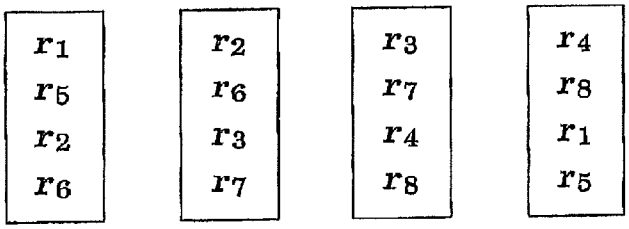

FIG. 15B

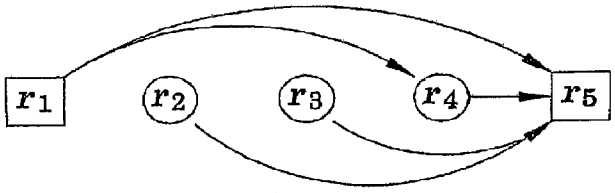

FIG. 16A

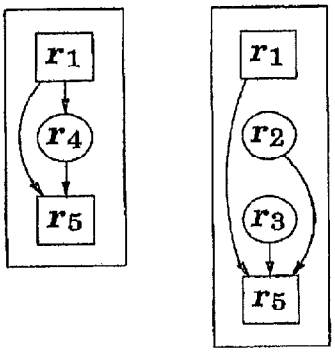

FIG. 16B

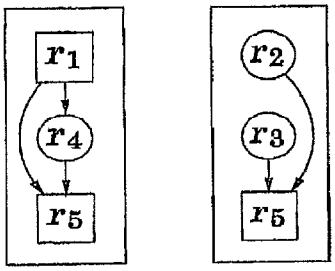

FIG. 16C

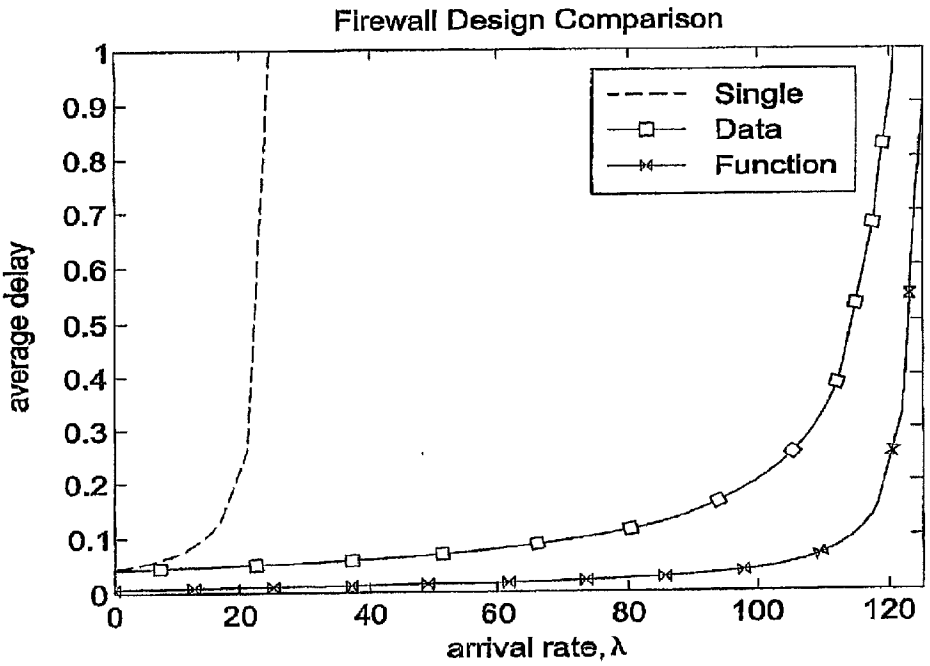
Firewall Design Comparison
Single
Data
Function
average delay
arrival rate, λ
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
20
40
60
80
100
120

FIG. 17A

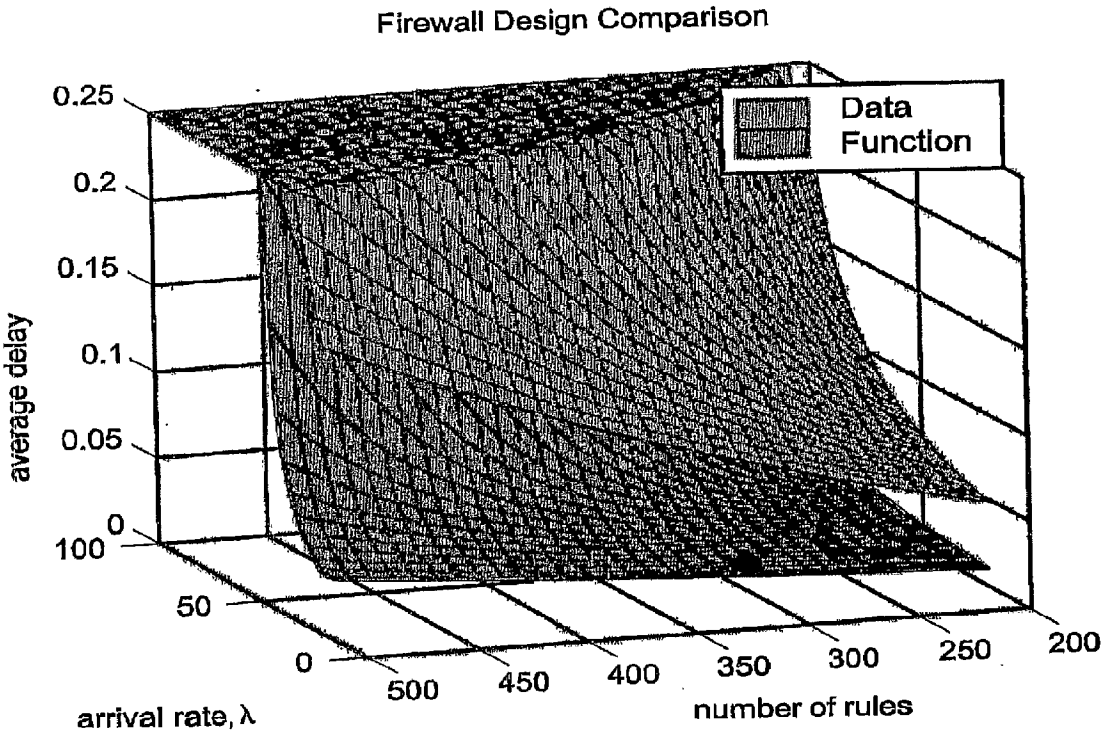
Firewall Design Comparison
Data
Function
average delay
0.25
0.2
0.15
0.1
0.05
0
100
50
0
arrival rate, λ
500
450
400
350
300
250
200
number of rules

FIG. 17B

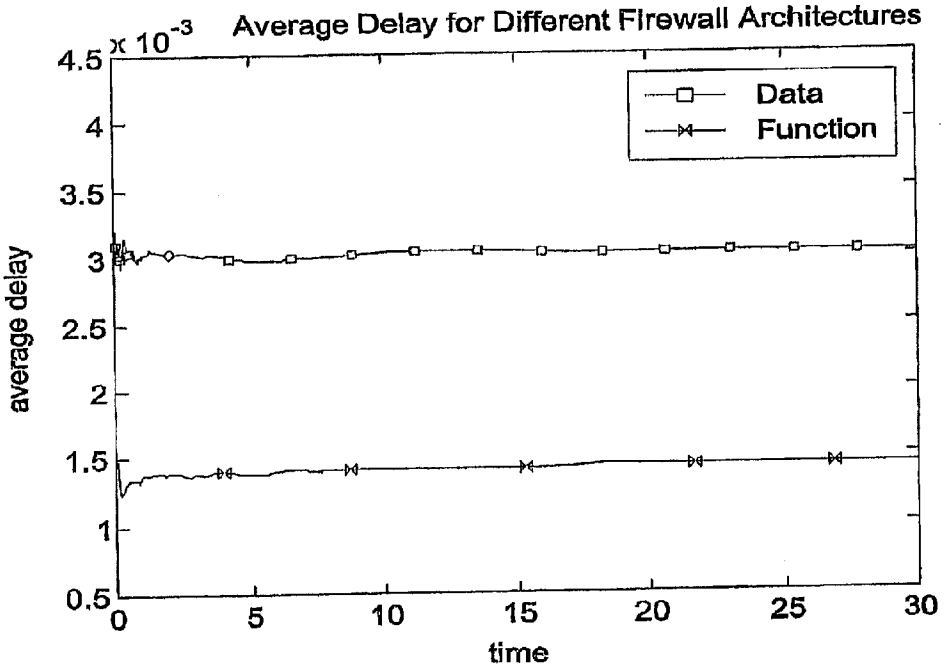
Average Delay for Different Firewall Architectures
x 10^-3
4.5
4
3.5
3
2.5
2
1.5
1
0.5
average delay
0
5
10
15
20
25
30
time
Data
Function

FIG. 18A

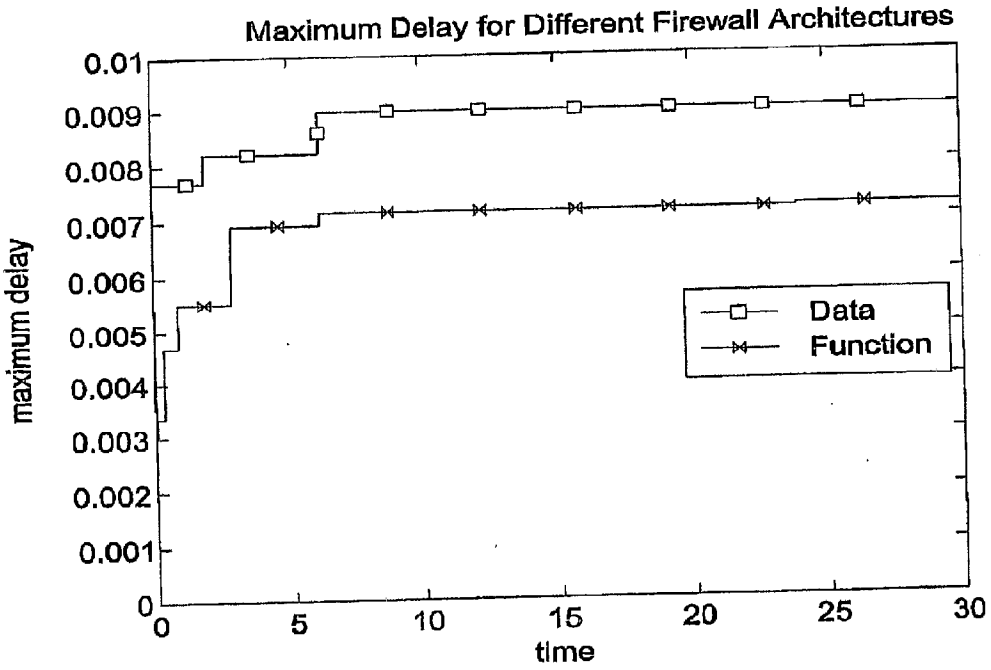
Maximum Delay for Different Firewall Architectures
maximum delay
time
Data
Function
0.01
0.009
0.008
0.007
0.006
0.005
0.004
0.003
0.002
0.001
0
0
5
10
15
20
25
30

FIG. 18B

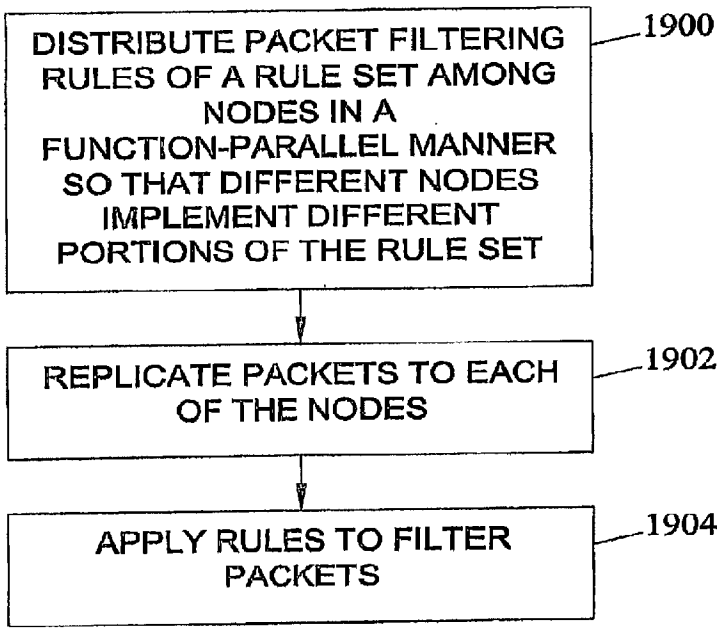
DISTRIBUTE PACKET FILTERING RULES OF A RULE SET AMONG NODES IN A FUNCTION-PARALLEL MANNER SO THAT DIFFERENT NODES IMPLEMENT DIFFERENT PORTIONS OF THE RULE SET
1900
REPLICATE PACKETS TO EACH OF THE NODES
1902
APPLY RULES TO FILTER PACKETS
1904

FIG. 19

Espacenet

Bibliographic data: WO2011038420 (A2) — 2011-03-31

METHODS, SYSTEMS, AND COMPUTER READABLE MEDIA FOR ADAPTIVE PACKET FILTERING

Inventor(s): AHN DAVID K [US] ± (AHN, DAVID, K)

Applicant(s): GREAT WALL SYSTEMS [US]; AHN DAVID K [US] ± (GREAT WALL SYSTEMS, ; AHN, DAVID, K)

Classification: **- international:** ***H04L12/22; H04L12/26; H04L12/56; H04L29/06***
- cooperative: **G06F21/562**; **H04L63/0227**; **H04L63/0263**; H04L63/1416

Application number: WO2010US54520 20101028

Priority number (s): US20090237974P 20090828 ; US20100871806 20100830

Also published as: WO2011038420 (A3) US2011055916 (A1) US8495725 (B2) US2014007216 (A1) EP2471218 (A2) more

Abstract of WO2011038420 (A2)

The subject matter described herein includes methods, systems, and computer readable media for adaptive packet filtering. One method includes identifying at least one subset of rules and an ordered set of firewall packet filtering rules that defines a firewall policy such that the subset contains disjoint rules. Disjoint rules are defined as rules whose order can be changed without changing integrity of the firewall policy. Rules in the subset are sorted to statistically decrease the number of comparisons that will be applied to each packet that a firewall encounters.; Packets are filtered at the firewall using the sorted rules in the subset by comparing each packet to each of the sorted rules in the subset until the packet is allowed or denied and ceasing the comparing for the packet in response to the packet being allowed or denied and

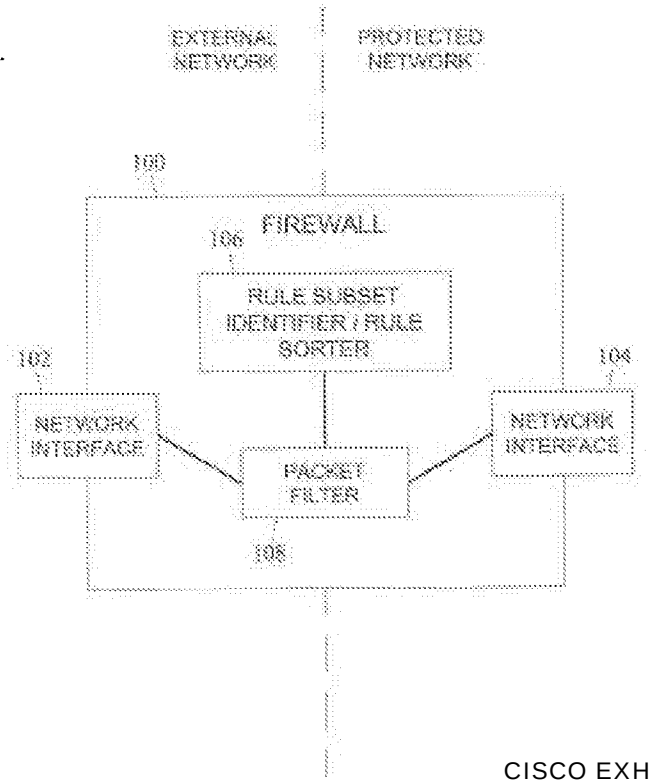

thereby achieving sub-linear searching for packets filtered using the sorted rules in the subset.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
31 March 2011 (31.03.2011)

PCT

(10) International Publication Number
WO 2011/038420 A2

(51) International Patent Classification:
H04L 12/22 (2006.01) *H04L 12/56* (2006.01)
H04L 12/26 (2006.01) *H04L 29/06* (2006.01)

(21) International Application Number:
PCT/US2010/054520

(22) International Filing Date:
28 October 2010 (28.10.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/237,974 28 August 2009 (28.08.2009) US
12/871,806 30 August 2010 (30.08.2010) US

(71) Applicant *(for all designated States except US)*: **GREAT WALL SYSTEMS** [US/US]; 635 W 4th St., Ste 200, Winston-Salem, NC 27101-2740 (US).

(72) Inventor; and

(75) Inventor/Applicant *(for US only)*: **AHN, David, K.** [US/US]; 740 N. Pine Valley Road, Winston-salem, NC 27106-5508 (US).

(74) Agent: **HUNT, Gregory, A.**; Jenkins, Wilson, Taylor & Hunt, P.A., Suite 1200, University Tower, 3100 Tower Boulevard, Durham, NC 27707 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,

[Continued on next page]

(54) Title: METHODS, SYSTEMS, AND COMPUTER READABLE MEDIA FOR ADAPTIVE PACKET FILTERING

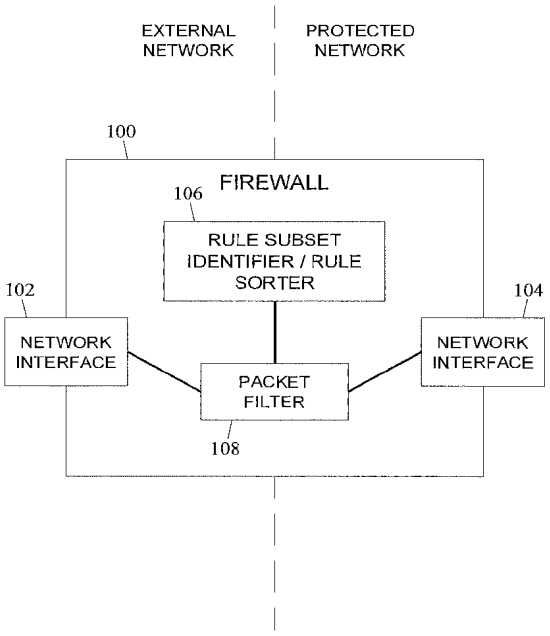

FIG. 1

(57) Abstract: The subject matter described herein includes methods, systems, and computer readable media for adaptive packet filtering. One method includes identifying at least one subset of rules and an ordered set of firewall packet filtering rules that defines a firewall policy such that the subset contains disjoint rules. Disjoint rules are defined as rules whose order can be changed without changing integrity of the firewall policy. Rules in the subset are sorted to statistically decrease the number of comparisons that will be applied to each packet that a firewall encounters. Packets are filtered at the firewall using the sorted rules in the subset by comparing each packet to each of the sorted rules in the subset until the packet is allowed or denied and ceasing the comparing for the packet in response to the packet being allowed or denied and thereby achieving sub-linear searching for packets filtered using the sorted rules in the subset.

WO 2011/038420 A2

SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

— *with information concerning request for restoration of the right of priority in respect of one or more priority claims (Rules 26bis.3 and 48.2(b)(vii))*

DESCRIPTION

METHODS, SYSTEMS, AND COMPUTER READABLE MEDIA FOR ADAPTIVE PACKET FILTERING

PRIORITY CLAIM

This application claims the benefit of U.S. Provisional Patent Application Serial No. 61/237,974, filed August 28, 2009 and U.S. Patent Application Serial No. 12/871,806, filed August 30, 2010; the disclosures of which are incorporated herein by reference in their entireties.

TECHNICAL FIELD

The subject matter described herein relates to network firewall filtering. More particularly, the subject matter described herein relates to methods, systems, and computer readable media for adaptive packet filtering.

BACKGROUND

A firewall generally processes a packet against a list of ordered rules to find the first rule match. The list of ordered rules represents an aggregate security policy, and arbitrarily changing the order of the rules can result in a violation of the aggregate security policy. The Wake Forest University (WFU) techniques described in U.S. patent application publication nos. 2006/0248580 and 2006/0195896 provide the methods to optimally reorder the list while preserving the aggregate security policy, thereby improving the performance of the firewall. The WFU techniques also include methods to break apart rules into functionally independent lists containing (groups of) dependent rules such that a function parallel firewall can simultaneously process one packet against multiple lists, which can substantially improve the performance of the firewall. However, these improvements provided by WFU techniques can be dwarfed by the performance degradation as the number of rules in the list becomes very large.

A key reason for the lack of scalability of most firewall implementations is due to the common use of linear search algorithms for comparing packets against a list of rules. In the worst case, a packet is matched at the last N^{th} rule in the list, so it must also be compared against all *N-1* prior rules for a total of *N*

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 893

comparisons. This poses a computational resource problem when the size of *N* is very large on a single processing node (including when such nodes are arranged in a data-, function-, hierarchical- or hybrid-parallel system), where the time required for processing each packet can quickly increase latency and reduce throughput to unacceptable levels. In fact, the WFU techniques provide good results in part because the reordering of, or the reduction in size of, rules on each processing node allows for a larger percentage of the total rules to reside in each processor's cache(s), which then substantially increases their performance relative to when only a small portion of those rules are cached.

The problem of searching firewall rule sets is well understood and highly researched, and there are some published techniques for sub-linear (substantially faster than linear) techniques applicable to firewall rules. However, these sub-linear techniques generally involve changing the underlying representation of rules. Examples of such an approach might be to use a graph, trie- or tree-like structure instead of a list to represent a set of rules, which would allow a match to be determined using tree search algorithms by traversing down the graph, trie or tree (see E. Fulp, Trie-Based Policy Representations for Network Firewalls, Proceedings of the IEEE International Symposium on Computer Communications, 2005 and Al-Shaer et al., Modeling and Management of Firewall Policies, IEEE Transactions on Network and Service Management, 2004). These approaches have potential but can add complexity or limitations that may reduce their practical usefulness in a commercial high performance firewall product.

SUMMARY

Adaptive packet filtering (APF), a set of techniques for processing firewall rules and packets, is described herein. APF offers improved processing performance compared to the WFU techniques in most cases, and can be combined with the WFU techniques or other parallel, pipelining and optimization techniques to achieve even greater performance.

The subject matter described herein includes methods, systems, and computer readable media for adaptive packet filtering. One method includes identifying at least one subset of rules in an ordered set of firewall packet

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 894

filtering rules that defines a firewall policy such that the subset contains disjoint rules. Disjoint rules are defined as rules whose order can be changed without changing the integrity of the firewall policy. Rules in the subset are sorted to statistically decrease the number of comparisons that will be applied to each packet that a firewall encounters. Packets are filtered at the firewall using the sorted rules in the subset by using binary search, interpolated search, informed search, or hash lookup search algorithms to compare each packet to the sorted rules in the subset until the packet is allowed or denied and ceasing the comparing for the packet in response to the packet being allowed or denied and thereby achieving sub-linear searching for packets filtered using the sorted rules in the subset.

The subject matter described herein for adaptive packet filtering can be implemented in a non-transitory computer readable medium having stored thereon executable instructions that when executed by the processor of a computer control the computer to perform steps. Exemplary computer readable media suitable for implementing the subject matter described herein include chip memory devices, disk memory devices, programmable logic devices and application specific integrated circuits. In addition, a computer readable medium that implements a subject matter described herein may be located on a single device or computing platform or may be distributed across multiple devices or computing platforms.

Further, the subject matter described herein for adaptive packet filtering can be implemented on a particular machine, such as a network firewall including one or more network interfaces for receiving packets and packet filtering hardware and software for optimizing rules as described herein and for filtering packets using the optimized arrangement of rules.

BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the subject matter described herein will now be described with reference to the accompanying drawings of which:

Figure 1 is a block diagram of a system for adaptive packet filtering according to an embodiment of the subject matter described herein;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 895

Figure 2 is a block diagram of application of the present subject matter to a pipelined processing approach according to an embodiment of the subject matter described herein;

Figure 3 is a block diagram illustrating application of the present subject matter to a combination of pipelined and data parallel processing approaches according to an embodiment of the subject matter described herein; and

Figure 4 is a block diagram illustrating application of the present subject matter to a short-circuiting pipelined processing approach according to an embodiment of the subject matter described herein; and

Figure 5 is a block diagram illustrating application of the present subject matter to a combination of pipelined and function parallel processing approaches according to an embodiment of the subject matter described herein.

DETAILED DESCRIPTION

Methods, systems, and computer readable media for adaptive packet filtering are disclosed. Figure 1 is a block diagram illustrating an exemplary system for adaptive packet filtering according to an embodiment of the subject matter described herein. Referring to Figure 1, a firewall **100** may function at the boundary between an external network and a protected network. Firewall **100** may include one or more network interfaces **102** for receiving packets from the external network. Firewall **100** may also include one or more network interfaces **104** for transmitting allowed packets to the protective network. In one implementation, firewall **100** may filter Internet protocol (IP) packets based on a combination of source and destination addresses in the IP headers of the packets. However, the subject matter described herein is not limited to filtering any particular protocol. Any packet network protocol with parameters for which firewall filtering rules can be defined is intended to be within the scope of the subject matter described herein.

As used herein, the term "firewall" includes any network security device or system of devices that inspects network traffic data that originates, terminates, or traverses the device system in any capacity and compares that

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 896

traffic data (headers, payload, raw bits, etc.) to a set of one or more rules, signatures, or conditions, either inline (i.e., in real time) or offline (i.e., capture and replay of the traffic data). The term "firewall" is also intended to include an intrusion detection device that analyzes network traffic in real time or historically to detect the presence of intrusion events in a network. The term "firewall" is also intended to include a deep packet inspection device that analyzes network traffic in real time or historically to detect the presence of certain packet content in a network.

Firewall **100** includes a firewall rule subset identifier/rule sorter **106** for identifying at least one subset of rules in an ordered set of firewall packet filtering rules that defines a firewall policy such that the subset contains disjoint rules, where disjoint rules are defined as rules whose order can be changed without changing the firewall policy. Firewall rule subset identifier/rule sorter **106** may sort the rules in the subset or subsets to statistically decrease the number of comparison that will be applied to each packet that the firewall encounters. Exemplary methods for grouping and sorting rules will be described below.

Although in the example illustrated in Figure 1 rule subset identifier/rule sorter **106** is illustrated as a component of firewall **100**, the subject matter described herein is not limited to such an implementation. Rule subset identifier/rule sorter **106** can be implemented on any computing platform capable of sorting firewall rules using the methods described herein, and the sorted rule set can be provided to firewall **100** through any suitable means, such as communication over a network. In one implementation, rule subset identifier/rule sorter **106** may be implemented on a management platform separate from firewall **100**.

Firewall **100** further includes a packet filter **108** for filtering packets at the firewall using the rules in the subset by using binary search, interpolated search, informed search, hash lookup search algorithms, or other sub-linear algorithms to compare each packet to each of the sorted rules in the subset until the packet is allowed or denied and ceasing the comparing for the packet in response to the packet being allowed or denied and thereby achieving

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 897

sub-linear searching for the packets filtered using the sorted rules in the subset.

Once the subsets of disjoint rules have been identified by rule subset identifier/rule sorter **106**, the rule subsets can be distributed across plural firewall processors in order to improve packet filtering efficiency. Figure 2 is a block diagram illustrating an exemplary pipelined approach where rules and different subsets are distributed across plural firewall processors for processing packets in a pipelined manner. Referring to Figure 2, firewalls **200** and **202** each include separate processors **204** and **206** for executing packet filters **108**. In this example, rule subset identifier/rule sorter **106** identifies two rule subsets, subset A **208** and subset B **210**. The rules within each subset **208** and **210** are disjoint and sorted to statistically decrease the number of comparisons that will be applied to each packet using the methods described herein. However, the rules in subset B **210** are dependent on the rules in rule subset A **208**. Accordingly, rule subset identifier/rule sorter **106** distributes the rule across firewall processors **204** such that the rules in rule subset A **208** are applied before the rules in rule subset B **210**. Because the rules in different subsets are distributed across plural processors in a pipeline manner, packet filtering efficiency is improved over a single-processor approach because the different processors can simultaneously apply rules to different packets. In the example illustrated in Figure 2, packets that pass the filtering of rule subset A **208** are processed by processor **206**, which applies rule subset B **210**, at the same time that processor **204** applies rule subset A to new incoming packets.

In yet another embodiment, rule subset identifier/rule sorter **106** may distribute the grouped, sorted rules across firewall processors such that a combination of pipelined and data parallel processing techniques are used. Figure 3 illustrates an example where firewalls **300**, **302**, **304**, and **306** each include separate processors **308**, **310**, **312**, and **314** for applying their respective packet filters. In the illustrated example, rule subset identifier/rule sorter **106** distributes rule subset A **316** to firewall **300**, rule subset B **318** to firewalls **302** and **304**, and rule subset C **320** to firewall **306**. The rules within each subset A, B and C are disjoint. The rules in subset B are dependent

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 898

upon the rules in subset A. The rules in subset C are dependent upon the rules in subsets A and B.

In operation, packets entering firewall **300** are filtered using rule subset A **316**. The packets that are allowed by rule subset A **316** are divided between firewalls **302** and **304** such that the application of the rules in rule subset B **318** to different packets is performed in parallel. This is referred to as a data parallel approach. The packets that pass the filtering by rule subset B **318** are passed to firewall **306** for application of the rules in rule subset C **320**. Accordingly, Figure 3 illustrates an example where the rule subsets that are identified and sorted by rule subset identifier/rule sorter **106** are distributed across the firewall processors for a combination of pipelined and data parallel processing.

In yet another embodiment, the rules subsets that are identified and in which the rules are sorted using rule subset identifier/rule sorter **106** may be distributed across firewall processors in a short-circuiting pipelined manner. Figure 4 is an example of short-circuiting pipelined filtering using rule subsets that are identified and sorted by rule subset identifier/rule sorter **106**. Referring to Figure 4, a first firewall **400** and a second firewall **402,** respectively including processors **404** and **406,** filter packets using packet filters **108**. In the illustrated example, packet filter **108** uses rule subset A **408** and packet filter **108** uses rule subset B **410**. Rule subsets A and B **408** and **410** may respectively implement different levels of a firewall hierarchy such that packets that pass the filtering by rule subset A **408** are allowed into the protected network. Packets that are identified by rule subset A **408** is requiring further filtering are distributed to rule subset B **410** for that filtering. Thus, rule subset identifier/rule sorter **106** can also be used with short-circuiting pipelined firewall techniques without departing from the scope of the subject matter described herein.

In yet another embodiment, rule subset identifier/rule sorter **106** may distribute the grouped, sorted rules across firewall processors such that a combination of pipelined and function parallel processing techniques are used. Figure 5 illustrates an example where firewalls **500**, **502**, **504**, and **506** each include separate processors **508**, **510**, **512**, and **514** for applying their

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 899

respective packet filters. In the illustrated example, rule subset identifier/rule sorter **106** distributes rule subset A **516** to firewall **500**, rule subset B **518** to firewalls **502**, and rule subset C **520** to firewall **504**, and rule subset D **522** to firewall **506**. The rules within each subset A, B, C and D are disjoint. The rules in subset B and C are dependent upon the rules in subset A. The rules in subset D are dependent upon the rules in subsets A, B and C.

In operation, packets entering firewall **500** are filtered using rule subset A **516**. The packets that are allowed by rule subset A **516** are copied to both firewalls **502** and **504** such that the application of the rules in rule subsets B **518** and C **520** to the packets is performed in parallel. This is referred to as a function parallel approach. The packets that pass the filtering by rule subsets B **518** and C **520** are passed to firewall **506** for application of the rules in rule subset D **522**. Accordingly, Figure 5 illustrates an example where the rule subsets that are identified and sorted by rule subset identifier/rule sorter **106** are distributed across the firewall processors for a combination of pipelined and function parallel processing.

Technique

APF analyzes and orders the list of firewall rules in-place to contain functionally dependent groups, where each group contains a subset of rules that are disjoint, dependent or both, without substantially changing the underlying representation of rules and while preserving the aggregate security policy. APF then uses varying criteria to sort each group containing disjoint rules, then uses sub-linear search algorithms when comparing packets against the rules within that group. APF uses linear search algorithms when comparing packets within a group containing dependent rules or when otherwise appropriate. A detailed computational complexity analysis of APF would need to be completed. However, on average, it is hypothesized that only $O(log(N))$ comparisons would be needed to process a rule list of size N. In the theoretical best case when all rules are disjoint, this translates to about 20 comparisons (instead of 1,000,000) for a list of N = 1,000,000 rules and about 30 comparisons for a list of N = 1,000,000,000 rules. In the worst case when all rules are dependent, APF performs the same as linear search firewall

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 900

cores. In practice, APF should process a packet against a very large list of rules (*N* = millions) in the same amount of time that other techniques can process against a very small list (*N* = hundreds, or thousands). APF does not inherently use parallel techniques; therefore, it can be combined with WFU techniques or other parallel/pipelining techniques to increase performance.

The following table shows preliminary results comparing a single linear search firewall) with a single APF core as the number of rules is increased.

	linear firewall core		APF core	
N = # rules in the core	PPS = packets per second at max throughout with 0 loss	Latency in microseconds at max throughput with 0 loss	PPS = packets per second at max throughout with 0 loss	Latency in microseconds at max throughput with 0 loss
1	844,595	12.3	811,688	12.3
10	730,994	11.9	766,871	12.3
100	314,861	13.9	718,391	12.0
1,000	29,357	44.6	683,060	12.0
10,000	930	1,112.4	464,684	12.0
100,000	Fail	Fail	292,740	17.4
1,000,000			252,525	14.3
10,000,000			132,556	18.6

Detailed Technique

This section describes an exemplary algorithm for implementing the subject matter described herein.

A firewall *rule* is defined as an *n*-tuple criteria and an associated action for matching packets. For example, a 5-tuple rule that matches Internet Protocol version 4 (IPv4) packets might consist of 5 IPv4 header fields (source address, source port, destination address, destination port and protocol) and an action (allow, deny), and might specify the rule ***R1*** as:

Rule	Source Addr	Source Port	Dest Addr	Dest Port	Protocol	Action
R1	192.168.1.1	12345	10.1.1.1	80	TCP	DENY

A firewall *rule set* is defined as an ordered list of *n* rules *R1,R2,R3,...,Rn* where the *i* in *Ri* is the index of the rule in the list. Packets that traverse the firewall are checked against each rule in the rule set until the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 901

first matching rule is found and its associated action is applied. An example rule set is ***S1*** which contains:

Rule	Source Addr	Source Port	Dest Addr	Dest Port	Protocol	Action
R1	192.168.1.1	12345	10.1.1.1	80	TCP	DENY
R2	192.168.2.2	ANY	10.2.2.2	25	TCP	ALLOW
R3	192.168.3.3	ANY	10.3.3.3	53	UDP	ALLOW
R4	ANY	ANY	10.1.1.1	ANY	ANY	ALLOW
R5	ANY	ANY	ANY	ANY	ANY	DENY

An example TCP packet from source 192.168.4.4 port 54321 to destination 10.1.1.1 port 80 would be checked against but not match *R1*, *R2* and *R3*; would be checked against and match *R4* and be allowed; and, would not be checked against *R5* because *R4* was the first matching rule.

A firewall *security policy* is defined as the set of all possible packets that can traverse the firewall along with their specified outcomes as defined by the rule set. Changing the rules in a rule set usually results in a change of its security policy.

Within a rule set, a firewall rule is *dependent* on another rule if swapping the order of the two rules results in a change in the security policy of the rule set. Otherwise, the two rules are *disjoint* if swapping the order does **not** result in a change the security policy. For example, in rule set ***S1*** above, rules ***R1*** and ***R4*** are dependent because placing ***R4*** ahead of ***R1*** would render ***R1*** ineffective, thereby changing the security policy. Rules ***R1*** and ***R2*** are disjoint because placing ***R2*** ahead of ***R1*** does not change the security policy.

A *permutation* of a rule set is defined as a new rule set which contains the same rules as the original rule set, but which lists a different ordering of the rules from the original rule set without changing the original security policy. For example, in the rule set ***S1*** above, swapping the order of the disjoint rules ***R1*** and ***R2*** would result in a permutation rule set ***S1'***:

Rule	Source Addr	Source Port	Dest Addr	Dest Port	Protocol	Action
R2	192.168.2.2	ANY	10.2.2.2	25	TCP	ALLOW
R1	192.168.1.1	12345	10.1.1.1	80	TCP	DENY
R3	192.168.3.3	ANY	10.3.3.3	53	UDP	ALLOW

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 902

R4	ANY	ANY	10.1.1.1	ANY	ANY	ALLOW
R5	ANY	ANY	ANY	ANY	ANY	DENY

Two rules are *spatially disjoint* if they are disjoint and their corresponding tuples are either identical or do not overlap. For example, in the rule set ***S1*** above, rules ***R1*** and ***R2*** are disjoint but not spatially disjoint because the source ports 12345 and ANY overlap. However, rules ***R2*** and ***R3*** are both disjoint and spatially disjoint because the source ports ANY and ANY are identical, and the other 4 tuples do not overlap. (Other examples follow.)

A *transform function* is an algorithm that can be applied to a rule to create a *sortable key* for that rule, which can then be used to sort the rules by their keys using a *key comparison function*. For example, the transform function ***Tfn*** could concatenate the tuples of a rule into a bit array that is interpreted as a large integer, and a corresponding comparison function ***Cfn*** could be a simple integer comparison function. (Other examples follow.)

A *rule subset* is defined as an ordered grouping of one or more rules within a rule set. For example, in rule set ***S1*** above, the rule subsets might be:

Subset	Rules
T1	*R1, R2, R3*
T2	*R4*
T3	*R5*

A *rule group* is defined as a rule subset with a *group type* (dependent, disjoint), *transform function*, *comparison function*, and a *search algorithm hint* (linear, sub-linear). The *group type* can be *dependent* if the group contains dependent rules, or can be *disjoint* if the group strictly contains disjoint rules. For example, the rule set ***S1*** above might contain the following disjoint rule group:

Group	Group Type	Rules	Transform Fn	Comparison Fn	Hint
G1	Disjoint	*R1, R2, R3*	*Tfn*	*Cfn*	Sub-linear

A rule set may be partitioned into a list of ordered rule groups such that the security policy of the rule set is not changed when each rule group is decomposed in the listed order. This partitioning is accomplished by applying

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 903

a rule subset identification method to a given rule set. An example of such a method is:

1. For a given rule set *S* containing *n* rules *R1,R2,...,Rn*:
    a. Create a new empty disjoint rule group *Gj* (initially G1) in *S*.
    b. Place the first ungrouped rule *Ri* (initially R1) into *Gj*.
    c. For each remaining ungrouped rule *Ri* in *S*:
        i. If *Ri* is disjoint from rules in *Gj* and placing *Ri* into *Gj* does not modify the security policy of *S*, then place *Ri* into *Gj*.
        ii. Otherwise, leave *Ri* ungrouped.
    d. If *S* contains ungrouped rules, then go to step 1.a.
2. The rule set *S* now contains *m* disjoint rule groups *G1,G2,G3,...,Gm* which group together the *n* (possibly reordered) rules *R1,R2,...,Rn*.

Applying the above method to rule set ***S1*** might result in its partitioning into the following list of disjoint ordered rule groups:

Group	Group Type	Rules	Transform Fn	Comparison Fn	Hint
G1	Disjoint	*R1, R2, R3*			Linear
G2	Disjoint	*R4*			Linear
G3	Disjoint	*R5*			Linear

Decomposing the disjoint rule groups would result in ***G1,G2,G3*** = [***R1,R2,R3***], [***R4***], [***R5***] = ***R1,R2,R3,R4,R5*** = ***S1***.

A partitioned rule set containing disjoint rule groups may then be sorted by applying a transform function to each rule within each disjoint group to derive a sortable key for each rule. Then, the rules may be reordered within their disjoint groups using their sortable keys. The resulting sorted groups may be searched using sub-linear searching algorithms. An example of the sorting method is:

1. For each disjoint group *Gj* in rule set *S*:
    a. For each rule *Ri* in *Gj*:
        i. Apply Transform *Tfn* to *Ri* to derive a sortable key *Ki*.
    b. Sort the rules in *Gj* using the comparison function *Cfn* on the sortable keys *K*.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 904

3. The rule set *S* now contains *m* sorted disjoint rule groups *G1,G2,G3,...,Gm*.

Applying the above method to the partitioned rule set ***S1*** might result in a new permutation rule set ***S1'*** that contains the following list of sorted rule groups where the ordering of rules within ***G1*** might change from ***R1,R2,R3*** to ***R3,R1,R2***:

Group	Group Type	Rules	Transform Fn	Comparison Fn	Hint
G1	Disjoint	*R3, R1, R2*	*Tfn*	*Cfn*	Sub-linear
G2	Disjoint	*R4*	*Tfn*	*Cfn*	Sub-linear
G3	Disjoint	*R5*	*Tfn*	*Cfn*	Sub-linear

The permutated rule set containing disjoint rule groups may be consolidated to reduce the number of groups that contain a rule count at or below a certain threshold, such as 1 rule, by merging two or more consecutive disjoint groups into a larger dependent group that may be searched using linear searching algorithms. An example of the consolidation method is:

1. For each group *Gj* in rule set S:
    a. If the sum of the number of rules in *Gj* and its subsequent group *Gj+1* is less than or equal to a specified threshold (e.g. 1), then the two groups are merged, their rules are concatenated, and the group type is set to *dependent*.
2. The rule set *S* now contains *m* or fewer rule groups of both dependent and disjoint types.

Applying the above method to the permuted rule set ***S1'*** might merge the disjoint group ***G3*** into ***G2***:

Group	Group Type	Rules	Transform Fn	Comparison Fn	Hint
G1	Disjoint	*R3, R1, R2*	*Tfn*	*Cfn*	Sub-linear
G2	Dependent	*R4, R5*	(none)	(none)	Linear

The APF packet filtering method matches packets against a given rule set by sequentially iterating over each of the ordered rule groups, then performing the specified sub-linear or linear search within each group. An example of a rule filtering method is:

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 905

1. For each packet that traverses the firewall:
    b. For each group *Gj* in rule set *S*:
        i. If *Gj* is a dependent group, then perform linear search within that group until there is a first rule match.
        ii. If *Gj* is a disjoint group:
            1. Apply transform *T* to the packet to derive a lookup key *K*.
            2. Use the comparison function *C* and sub-linear search within the group until there is a first rule match on key *K*.
        iii. If there is a first rule match, then process the packet according to its specified action. Otherwise, continue to the next group *Gj+1*.

Additional Notes

- The primary purpose of the partitioning of rules is to create rule groups that, in aggregate, enable the fastest possible searching of each packet against the rules in the rule set. In most cases, the optimal partitioning should be the grouping of maximal subsets of spatially disjoint rules. However, in some cases where these disjoint groups are small (e.g. less than 10 rules), consolidating the disjoint groups into a single larger group containing rules that are ordered using other criteria (such as hit probabilities or hardware cache friendliness) and employing linear or interpolated search algorithms may improve performance. A critical concept is the flexibility to organize the rules within the rule set in different ways that enable the use of the most efficient and applicable search algorithm that is available that accounts for the hardware capabilities, which is the motivation behind the term "adaptive" in Adaptive Packet Filtering.
- Rule sets may be partitioned, sorted and consolidated in-place.
- When partitioning a rule set into rule groups, each disjoint group should generally contain the maximal subsets of disjoint rules in order to reduce the number of disjoint groups in the rule set.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 906

- When partitioning a rule set into rule groups and/or sorting those groups, the transform and comparison functions may be different for each rule group.
- When partitioning a rule set into rule groups and/or sorting those groups, the algorithms may account for the hit probabilities of each rule and the aggregate hit probabilities of each group.
- When filtering packets, the search performed within any given rule group may employ the fastest available search algorithm applicable to that group even if it may be different from the specified search algorithm hint.
- When filtering packets, a sub-linear binary search within a disjoint rule group may account for hit probabilities at each pivot so that each recursion could maximize the probability of a rule match.
- When filtering packets, a constant-time search within a disjoint rule group is possible by defining a hashing function as the transform function such that the hash values for all rules within a group are unique within that group.

Examples

This section provides examples of the following items described in the algorithm in the previous section.

1) Example of the rule breakup

2) Example of transform function

3) Definition of rule representation

4) Difference between sub-linear and linear

5) Explanation of disjoint rules

1) Examples of the rule breakup

Define

```
S1 = (R1, R2, R3, R4, R5)
```

and

```
R1 = from 1.2.3.4 to 3.4.5.6 deny
R2 = from 2.3.4.5 to 4.5.6.7 allow
```


CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 907

```
R3 = from 3.4.5.6 to 5.6.7.8 allow
R4 = from *.*.*.* to 3.4.5.6 allow
R5 = from *.*.*.* to *.*.*.* deny
```

Then S1' contains 3 groups of disjoint rules:

```
S1' = (G1, G2, G3)
G1 = (R1, R2, R3)
G2 = (R4)
G3 = (R5)
```

Or, S1' can contain 2 groups of disjoint and dependent rules:

```
S1' = (G1, G2)
G1 = (R1, R2, R3), disjoint, sub-linear
G2 = (R4, R5), dependent, linear
```

or

Define

```
S2 = (R1, R2, R3, R4, R5, R6, R7, R8, R9)
```

and

```
R1 = from 1.2.3.4 to 3.4.5.6 deny
R2 = from 2.3.4.5 to 4.5.6.7 allow
R3 = from 1.2.3.4 to 4.5.6.7 deny
R4 = from 1.*.*.* to *.*.*.* allow
R5 = from 3.4.5.6 to 5.6.7.8 deny
R6 = from 2.*.*.* to *.*.*.* deny
R7 = from *.*.*.* to 3.4.5.6 allow
R8 - from *.*.*.* to 5.6.7.8 allow
R9 = from *.*.*.* to *.*.*.* deny
```

Then S2' contains 4 groups of "spatially disjoint" rules:

```
S2' = (G1, G2, G3, G4)
G1 = (R1, R2, R3, R5)
G2 = (R4, R6)
G3 = (R7, R8)
G4 = (R9)
```

Or, S2' can contain 2 groups of disjoint and dependent rules:

```
S2' = (G1, G2)
G1 = (R1, R2, R3, R5), disjoint, sub-linear
G2 = (R4, R6, R7, R8, R9), dependent, linear
```

---

2) Example of transform function

A rule R usually consists of an N-tuple, most basically a 3-tuple such as "from 1.2.3.4 to 3.4.5.6 deny"

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 908

* tuple a = source IP address, e.g. "1.2.3.4"
* tuple b = destination IP address, e.g. "3.4.5.6"
* tuple c = action, e.g. "deny"

Each of these tuples have underlying scalar integer/bit vector representation, so in the above example:
* source IP address = "1.2.3.4" = 32 bit integer 0016909060
* destination IP address = "3.4.5.6" = 32 bit integer 0050595078
* action = "deny" = 8 bit integer 000

One possible transform function is a transform to scalar key which concatenates the digits of each of the tuples into a large integer value:
* T(a, b, c) = abc
* T(1.2.3.4, 3.4.5.6, deny) = 0016909060 0050595078 000 = 001690906000050595078000

Another possible transform function is a transform to scalar key which concatenates the bits of each of the tuples into a large integer/bit vector:
* a = 1.2.3.4 = 00000001000000100000001100000100 (32 bits)
* b = 3.4.5.6 = 00000011000001000000010100000110 (32 bits)
* c = 0 = 00000000 (8 bits)

T(a, b, c) = abc =
000000010000001000000011000001000000001100000100000001010000011000000000

Another possible transform function is an identity function (i.e. transformation function that does not do anything), then defining a multi-dimensional comparison function for sorting purposes. An example of this is a comparison function that is radix-based for each tuple, which would essentially result in a rule set that is radix sorted by each tuple.

Note that the transform function must convert the rule into a sortable key, which does not necessarily have to be a scalar key (i.e. it can be a multi-dimensional key that uses a multi-dimensional comparison function for sorting).

---

3) Definition of rule representation.
Rule representation is the way a rule and a rule set are conceptually represented in software. The most common representation of a rule is as an N-tuple object or structure that simply holds all the tuples together:

```
struct rule {
      unsigned char ip_proto;
      unsigned int ip_src_addr;
      unsigned int ip_dst_addr;
      unsigned short ip_src_port;
```


CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 909

```
    unsigned short ip_dst_port;
    ...
    unsigned char action;
};
```

The most common representation of a rule set is an array or linked list that hold the rules in a fixed order, and allows for iteration forwards and backwards in the array or list.

Example of array:

memory location	0	1	2	3	4	5	6	7	8
value	R1	R2	R3	R4	R5	R6	R7	R8	R9

Example of linked list:

memory location	0	1	2	3	4	5	6	7	8
value	R1	R9	R2	R8	R3	R5	R4	R6	R7
linked list prev	-	0	2	4	6	5	7	8	3
rule	R1	R2	R3	R4	R5	R6	R7	R8	R9
next	2	4	6	5	7	8	3	1	-

An alternate rule representation is to hold the rule and rule set in a trie or other graph structure. An example of this is described in "Balancing Trie-Based Policy Representations for Network Firewalls."Stephen J. Tarsa and Errin W. Fulp. Proceedings of the IEEE International Symposium on Computer Communications, 2006

Another alternate rule set representation is a hierarchical one as described in OPTWALL, described in Acharya et al., "OPTWALL: A Hierarchical Traffic-Ware Firewall," available at:

http://www.isoc.org/isoc/conferences/ndss/07/papers/OPTWALL.pdf,

where rule sets are broken down into mutually exclusive rule subsets which are arranged in a hierarchical order. Despite some similarities in terminology, OPTWALL and APF are different. For example, APF does not change the underlying rule or rule set representation. It simply reorders the rules in place and keeps track of the beginning and ending rules in each subset T externally from the rule or rule set. For example, say that a given rule is a standard structure and the rule set is an array form so that the rule set contains the following in the computer's memory:

Rule set S:

Location	0	1	2	3	4	5	6	7	8
value	R1	R2	R3	R4	R5	R6	R7	R8	R9

Suppose that the above rule set S can be divided into subsets containing disjoint rules:

T1 = R2, R9, R1, R3

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 910

```
T2 = R5, R7
T3 = R2, R4, R8
```

so that

```
S' = T1, T2, T3 = R2, R9, R1, R3, R5, R7, R2, R4, R8
```

Then the order of the rules in memory can be changed in place:

Rule set S':

```
location  0  1  2  3  4  5  6  7  8
value     R2 R9 R1 R3 R5 R7 R2 R4 R8
          |---------| |---| |------|
               T1       T2     T3
```

The memory ranges for T1 (0-3), T2 (4-5) and T3 (6-8) are stored outside of the rule and rule set data structures.

---

4) Difference between linear and sub-linear.

A linear algorithm is one whose computational time increases linearly as the size of the set is increased. The best example of this is when looking up a word in a dictionary. If the dictionary is unsorted, then the order of the words would be arbitrary. Therefore, when looking up the word "zebra," one could start from the beginning and search until the end to find it. If the dictionary contains 1,000 entries, you would need to examine all 1,000 words in the worst case.

A sub-linear algorithm is one whose computational time increases sub-linearly as the size of the set is increased. In the above example, if the dictionary were sorted alphabetically, then one could still use a linear search by starting from the beginning and searching until the end to find "zebra." However, one could also use a sub-linear binary search algorithm by looking in the middle of the dictionary, then seeing if the middle entry comes alphabetically before or after (or is equal), then recursively selecting the middle of the appropriate half again and again to find the word "zebra." Since at each recursion 1/2 of the remaining words are eliminated, it would take about $\log_2(1{,}000)$ or about 10 examinations to find the entry in the worst case.

Another example of a sub-linear algorithm is hashing. Suppose that the above dictionary contains only 5 letter words. If we define a hash function that sums the alphabet order of each letter in the word (z = 26, e = 05, b = 02, r = 18, a = 01), then hash("zebra") = 26 + 05 + 02 + 18 + 01 = 52. The computer could have an array containing all the words in the dictionary where each word's position in the array is the hash value of the word (subject to collisions). In the above example, the array's 52nd position would have the word "zebra," so it would take only 1 comparison to determine a match without collision. This hash technique can be selectively used in APF.

---

5) Explanation of disjoint rules.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 911

A rule R1 is "disjoint" from another rule R2 if their positions in the rule set S can be exchanged without altering the overall security policy. An example of this is rule set S containing:

```
R1 = from 1.*.*.* to 2.3.4.5 deny
R2 = from 1.2.3.4 to 1.2.3.4 allow
```

has the same security policy as rule set S' containing:

```
R2 = from 1.2.3.4 to 1.2.3.4 allow
R1 = from 1.*.*.* to 2.3.4.5 deny
```

because P = P'; therefore, R1 and R2 are disjoint.

The technique set forth above does not explain the concept of "spatially disjoint" rules. This is important if the transform function T cannot account for overlapping tuple values, which can be very common in practical settings.

A rule R1 is "spatially disjoint" from another rule R2 if they are "disjoint" and their corresponding tuples do not unevenly overlap (must be exactly equal, or do not overlap at all). In the above example, R1 and R2 are disjoint but not spatially disjoint because the first tuple of R1 (1.*.*.*) and R2 (1.2.3.4) are not equal but do overlap, i.e. the value of 1.2.3.4 would match the first tuple of both R1 and R2. An example of spatially disjoint rules are:

```
R3 = from 1.2.3.4 to 3.4.5.6 deny
R4 = from 2.3.4.5 to 3.4.5.6 allow
```

Here, the first tuple of R1 (1.2.3.4) and the first tuple of R2 (2.3.4.5) do not overlap, the second tuple of R1 (3.4.5.6) and R2 (3.4.5.6) are equal, and the >> third tuple of R1 (deny) and R2 (allow) do not overlap.

The importance of "spatially disjoint" rules is dependent upon the definition of a transform function T, so it may be possible to define T such that rules need not be "spatially disjoint" so long as rules are "disjoint."

The disclosure of each of the publications referenced herein is hereby incorporated by reference in its entirety.

It will be understood that various details of the presently disclosed subject matter may be changed without departing from the scope of the presently disclosed subject matter. Furthermore, the foregoing description is for the purpose of illustration only, and not for the purpose of limitation.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 912

CLAIMS

What is claimed is:

1. A method for adaptive packet filtering, the method comprising:

identifying at least one subset of rules in an ordered set of firewall packet filtering rules that defines a firewall policy such that the at least one subset contains disjoint rules, where disjoint rules are defined as rules whose order can be changed without changing the integrity of the firewall policy;

sorting the rules in the at least one subset to statistically decrease the number of comparisons that will be applied to each packet that a firewall encounters; and

filtering packets at the firewall using the sorted rules in the at least one subset by comparing each packet to each of the sorted rules in the at least one subset until the packet is allowed or denied and ceasing the comparing for the packet in response to the packet being allowed or denied and thereby achieving sub-linear searching for the packets filtered using the sorted rules in at least one subset.

2. The method of claim 1 comprising, for the disjoint rules within the at least one subset, identifying spatially disjoint rules and grouping the spatially disjoint rules together within the at least one subset and wherein sorting the rules includes sorting the rules within each of the spatially disjoint groupings.

3. The method of claim 1 comprising identifying at least another subset of rules in the ordered set of firewall rules that are non-disjoint and refraining from sorting the rules in the at least another subset and wherein filtering the packets includes filtering the packets using the sorted, disjoint rules and the non-sorted, non-disjoint rules, thereby achieving sub-linear searching for the sorted, disjoint rules and achieving linear searching for the non-sorted, non-disjoint rules.

4. The method of claim 1 wherein sorting the rules includes applying a transform function to each rule in the at least one subset to compute a sortable key for each of the rules and sorting the rules according to the keys.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 913

5. The method of claim 1 wherein sorting the rules includes sorting the rules according to hit probabilities.

6. The method of claim 1 wherein sorting the rules includes sorting the rules in a binary search tree.

7. The method of claim 1 wherein filtering packets by comparing each packet to the sorted rules includes comparing each packet to the sorted rules using a binary search algorithm.

8. The method of claim 1 wherein filtering packets by comparing each packet to the sorted rules includes comparing each packet to the sorted rules using an interpolated search algorithm.

9. The method of claim 1 wherein filtering packets by comparing each packet to the sorted rules includes comparing each packet to the sorted rules using an informed search algorithm.

10. The method of claim 1 wherein filtering packets by comparing each packet to the sorted rules includes comparing each packet to the sorted rules using a hash lookup algorithm.

11. The method of claim 1 wherein each rule includes an IP source address and an IP destination address.

12. The method of claim 1 wherein the firewall includes a plurality of processors, the at least one subset comprises a plurality of subsets and further comprising distributing the subsets across the plurality of processors to implement at least one of pipelined processing, data parallel processing, and function parallel processing.

13. The method of claim 12 wherein the firewall implements a combination of pipelined and function parallel processing.

14. The method of claim 12 wherein the firewall implements a combination of pipelined and data parallel processing.

15. The method of claim 12 wherein the firewall implements a combination of pipelined, function parallel, and data parallel processing.

16. A system for adaptive packet filtering, the system comprising:

a firewall rule subset identifier/rule sorter for identifying at least one subset of rules in an ordered set of firewall packet filtering rules that defines a firewall policy such that the at least one subset contains

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 914

disjoint rules, where disjoint rules are defined as rules whose order can be changed without changing integrity of the firewall policy, and for sorting the rules in the at least one subset to statistically decrease the number of comparisons that will be applied to each packet that a firewall encounters; and

a packet filter for filtering packets at the firewall using the sorted rules in the at least one subset by comparing each packet to each of the sorted rules in the at least one subset until the packet is allowed or denied and ceasing the comparing for the packet in response to the packet being allowed or denied and thereby achieving sub-linear searching for the packets filtered using the sorted rules in at least one subset.

17. The system of claim 16 comprising, for the disjoint rules within the at least one subset, the firewall rule subset identifier/rule sorter identifies spatially disjoint rules and groups the spatially disjoint rules together within the at least one subset and wherein the firewall rule subset identifier/rule sorter the rules within each of the spatially disjoint groupings.

18. The system of claim 16 wherein the firewall rule subset identifier/rule sorter identifies at least another subset of rules in the ordered set of firewall rules that are non-disjoint and refrains from sorting the rules in the at least another subset and wherein the firewall packet filter filters the packets using the sorted, disjoint rules and the non-sorted, non-disjoint rules, thereby achieving sub-linear searching for the sorted, disjoint rules and achieving linear searching for the non-sorted, non-disjoint rules.

19. The system of claim 16 wherein the firewall rule subset identifier/rule sorter applies a transform function to each rule in the at least one subset to compute a sortable key for each of the rules and sorting the rules according to the keys.

20. The system of claim 16 wherein the firewall rule subset identifier/rule sorter sorts the rules includes sorting the rules according to hit probabilities.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 915

21. The system of claim 16 wherein the firewall rule subset identifier/rule sorter sorts the rules includes sorting the rules in a binary search tree.

22. The system of claim 16 wherein the packet filter compares each packet to the sorted rules using a binary search algorithm.

23. The system of claim 16 wherein the packet filter compares each packet to the sorted rules using an interpolated search algorithm.

24. The system of claim 16 wherein the packet filter compares each packet to the sorted rules using an informed search algorithm.

25. The system of claim 16 wherein the packet filter compares each packet to the sorted rules using a hash lookup algorithm.

26. The system of claim 16 wherein each rule includes an IP source address and an IP destination address.

27. The system of claim 16 wherein the firewall includes a plurality of processors, the at least one subset comprises a plurality of subsets and further comprising distributing the subsets across the plurality of processors to implement at least one of pipelined processing, data parallel processing, and function parallel processing.

28. The system of claim 27 wherein the firewall implements a combination of pipelined and function parallel processing.

29. The system of claim 27 wherein the firewall implements a combination of pipelined and data parallel processing.

30. The system of claim 27 wherein the firewall implements a combination of pipelined, data parallel, and function parallel processing.

31. A non-transitory computer readable medium having stored thereon on executable instructions that when implemented by a processor of a computer controls the computer to perform steps comprising:

identifying at least one subset of rules in an ordered set of firewall packet filtering rules that defines a firewall policy such that the at least one subset contains disjoint rules, where disjoint rules are defined as rules whose order can be changed without changing integrity of the firewall policy;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 916

sorting the rules in the at least one subset to statistically decrease the number of comparisons that will be applied to each packet that a firewall encounters; and

filtering packets at the firewall using the sorted rules in the at least one subset by comparing each packet to each of the sorted rules in the at least one subset until the packet is allowed or denied and ceasing the comparing for the packet in response to the packet being allowed or denied and thereby achieving sub-linear searching for the packets filtered using the sorted rules in at least one subset.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 917

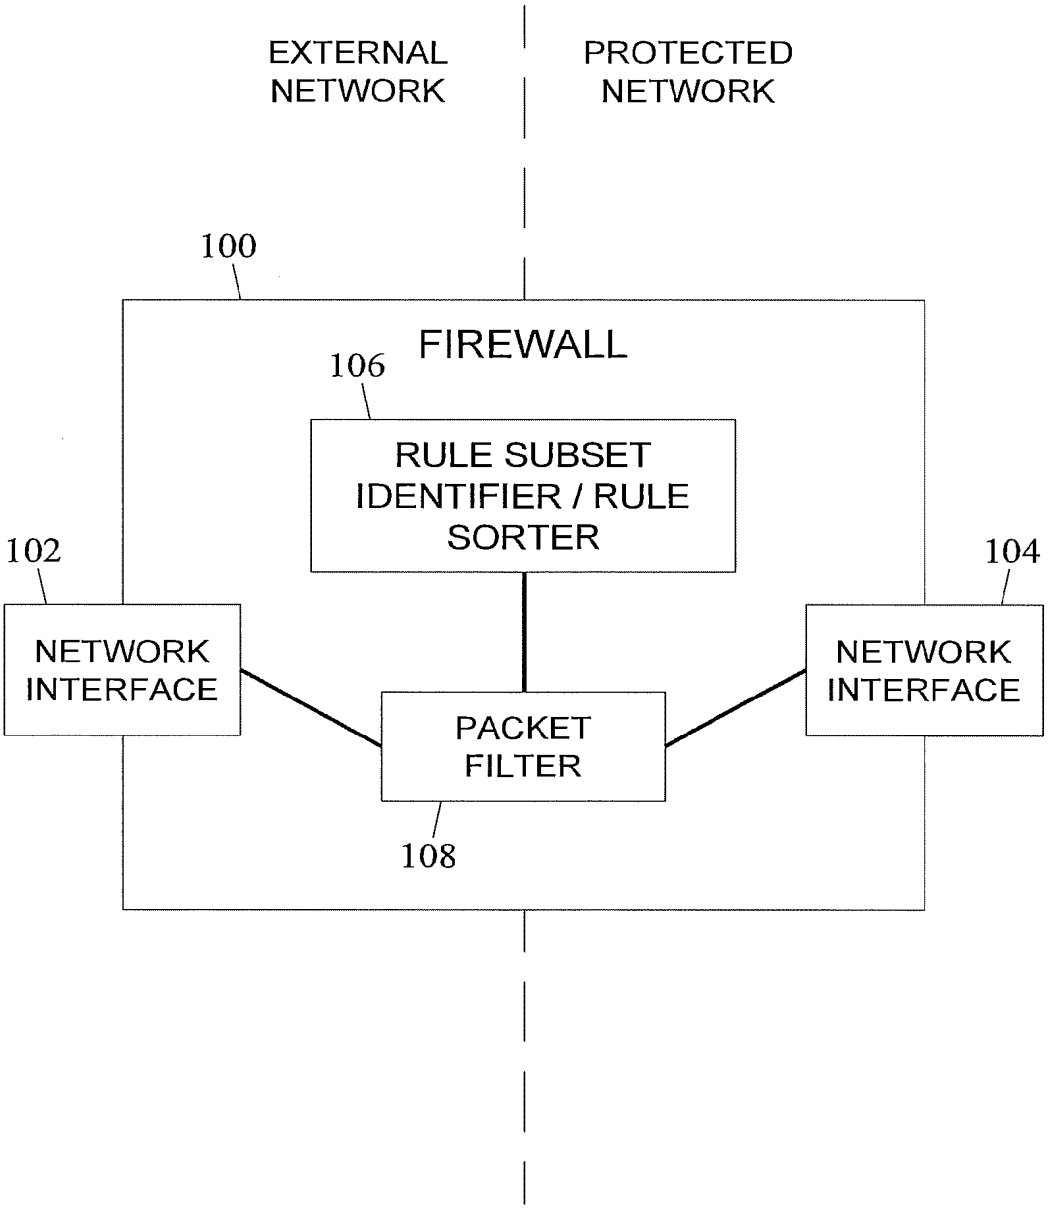
EXTERNAL NETWORK
PROTECTED NETWORK
100
FIREWALL
106
RULE SUBSET IDENTIFIER / RULE SORTER
102
104
NETWORK INTERFACE
NETWORK INTERFACE
PACKET FILTER
108

FIG. 1

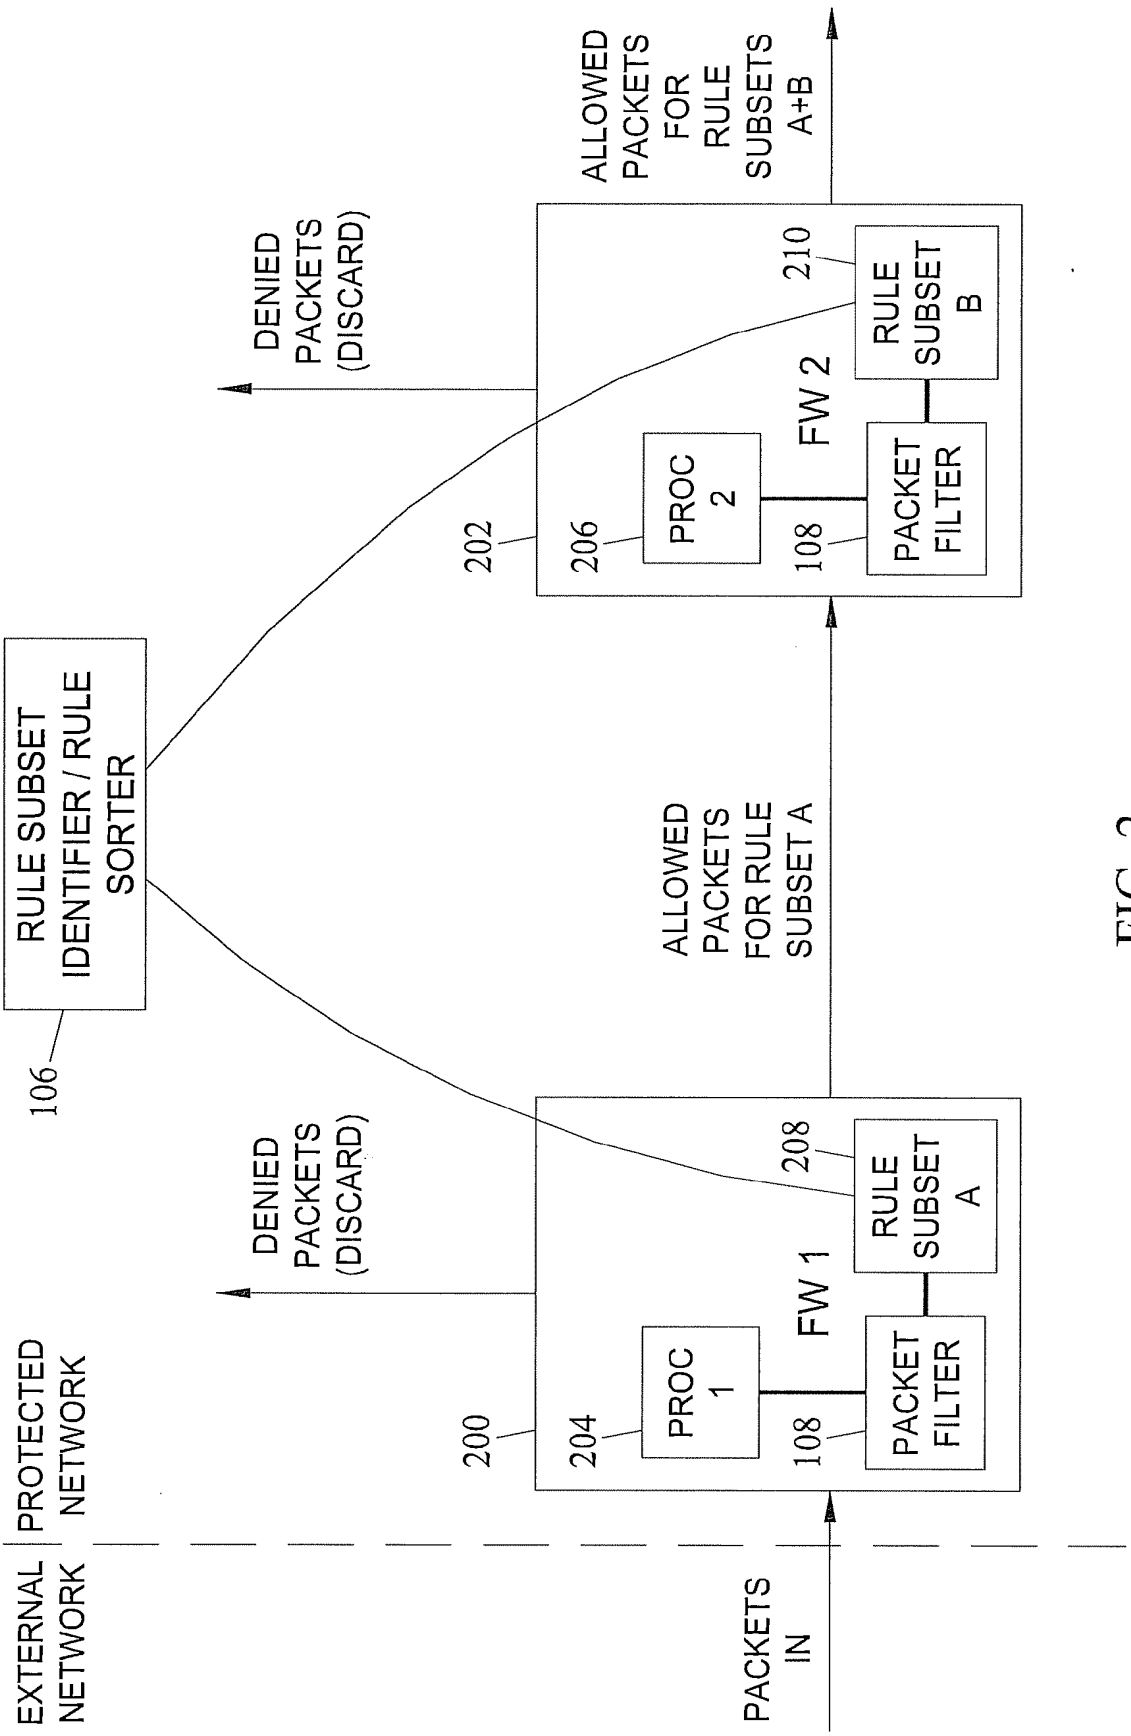
EXTERNAL NETWORK
PROTECTED NETWORK
PACKETS IN
RULE SUBSET IDENTIFIER / RULE SORTER
106
DENIED PACKETS (DISCARD)
DENIED PACKETS (DISCARD)
ALLOWED PACKETS FOR RULE SUBSET A
ALLOWED PACKETS FOR RULE SUBSETS A+B
200
202
204
206
208
210
PROC 1
PROC 2
FW 1
FW 2
PACKET FILTER
PACKET FILTER
108
108
RULE SUBSET A
RULE SUBSET B

FIG. 2

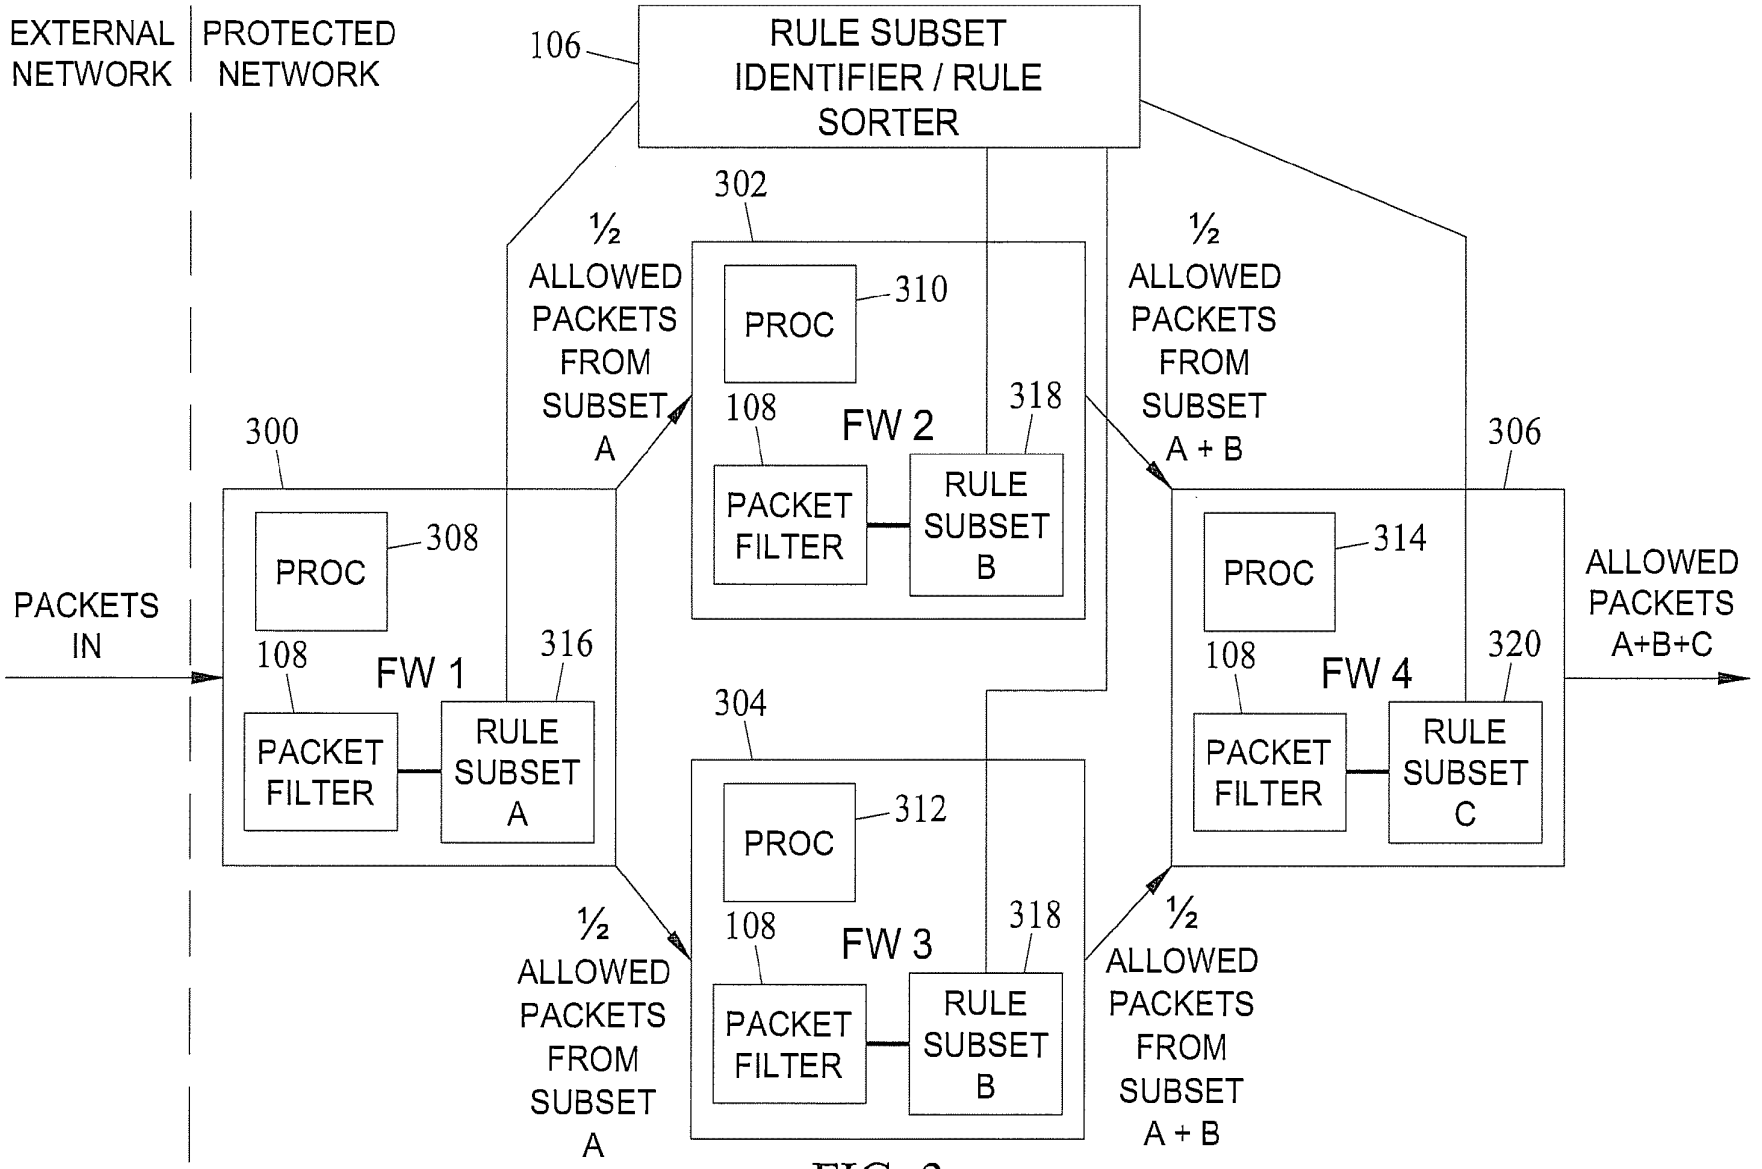

FIG. 3

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 920

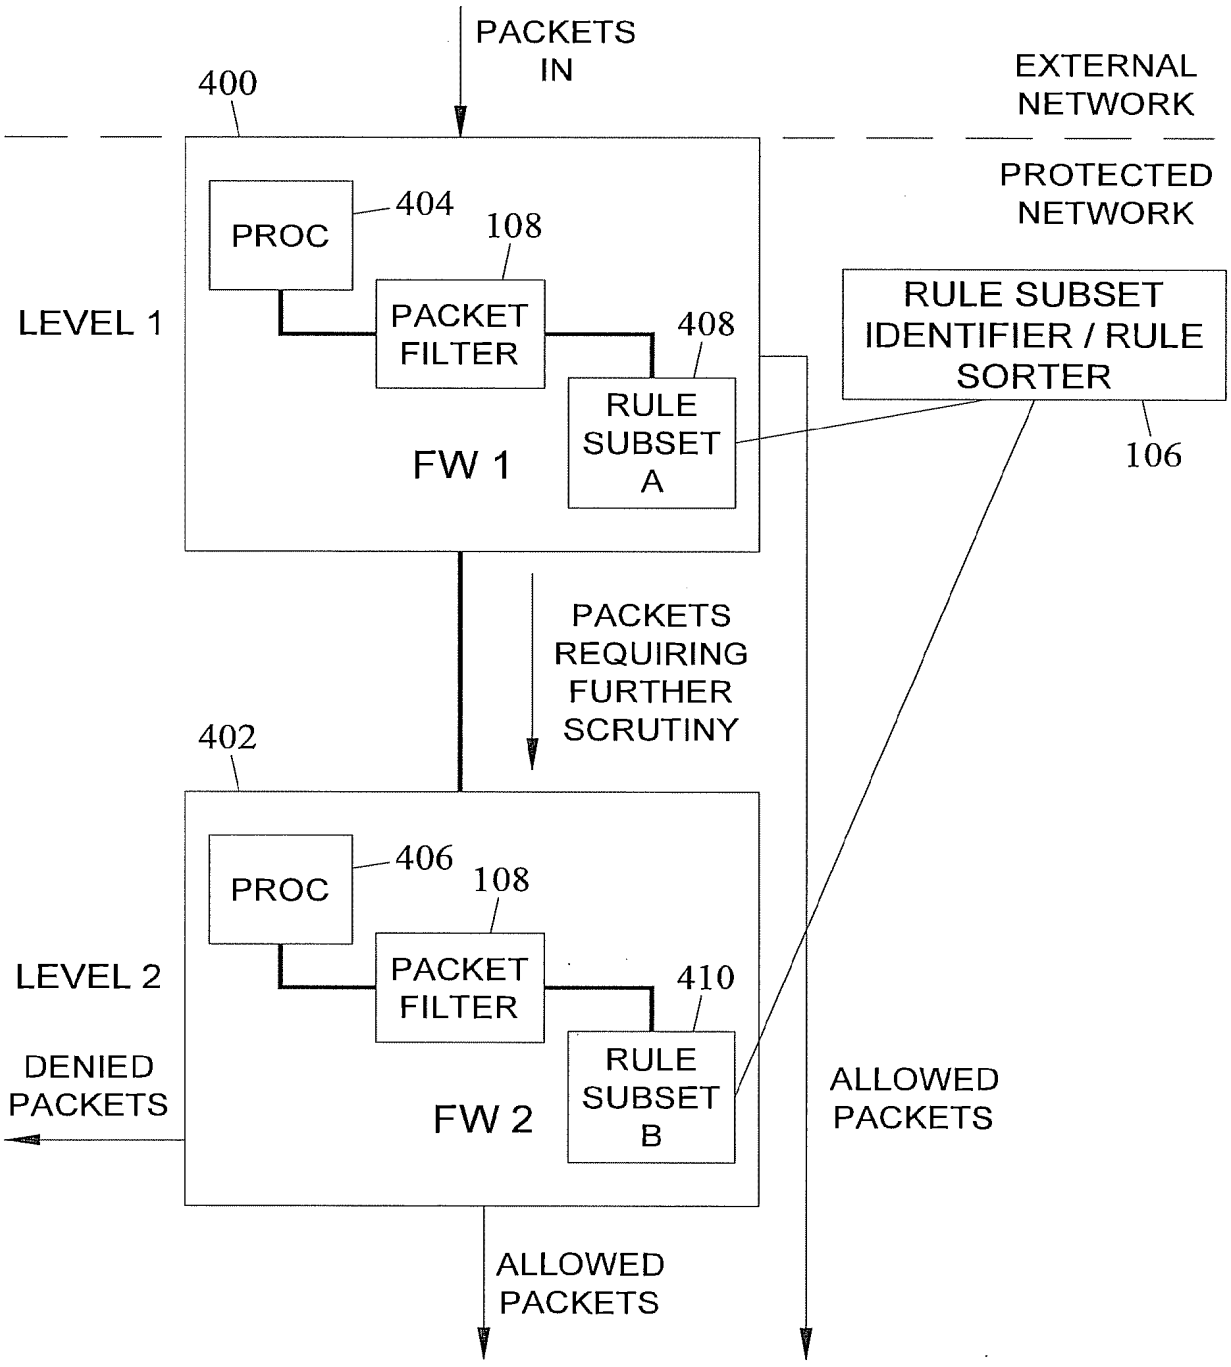
PACKETS IN
400
EXTERNAL NETWORK
PROTECTED NETWORK
PROC
404
108
PACKET FILTER
408
LEVEL 1
FW 1
RULE SUBSET A
RULE SUBSET IDENTIFIER / RULE SORTER
106
PACKETS REQUIRING FURTHER SCRUTINY
402
PROC
406
108
PACKET FILTER
410
LEVEL 2
DENIED PACKETS
FW 2
RULE SUBSET B
ALLOWED PACKETS
ALLOWED PACKETS

FIG. 4

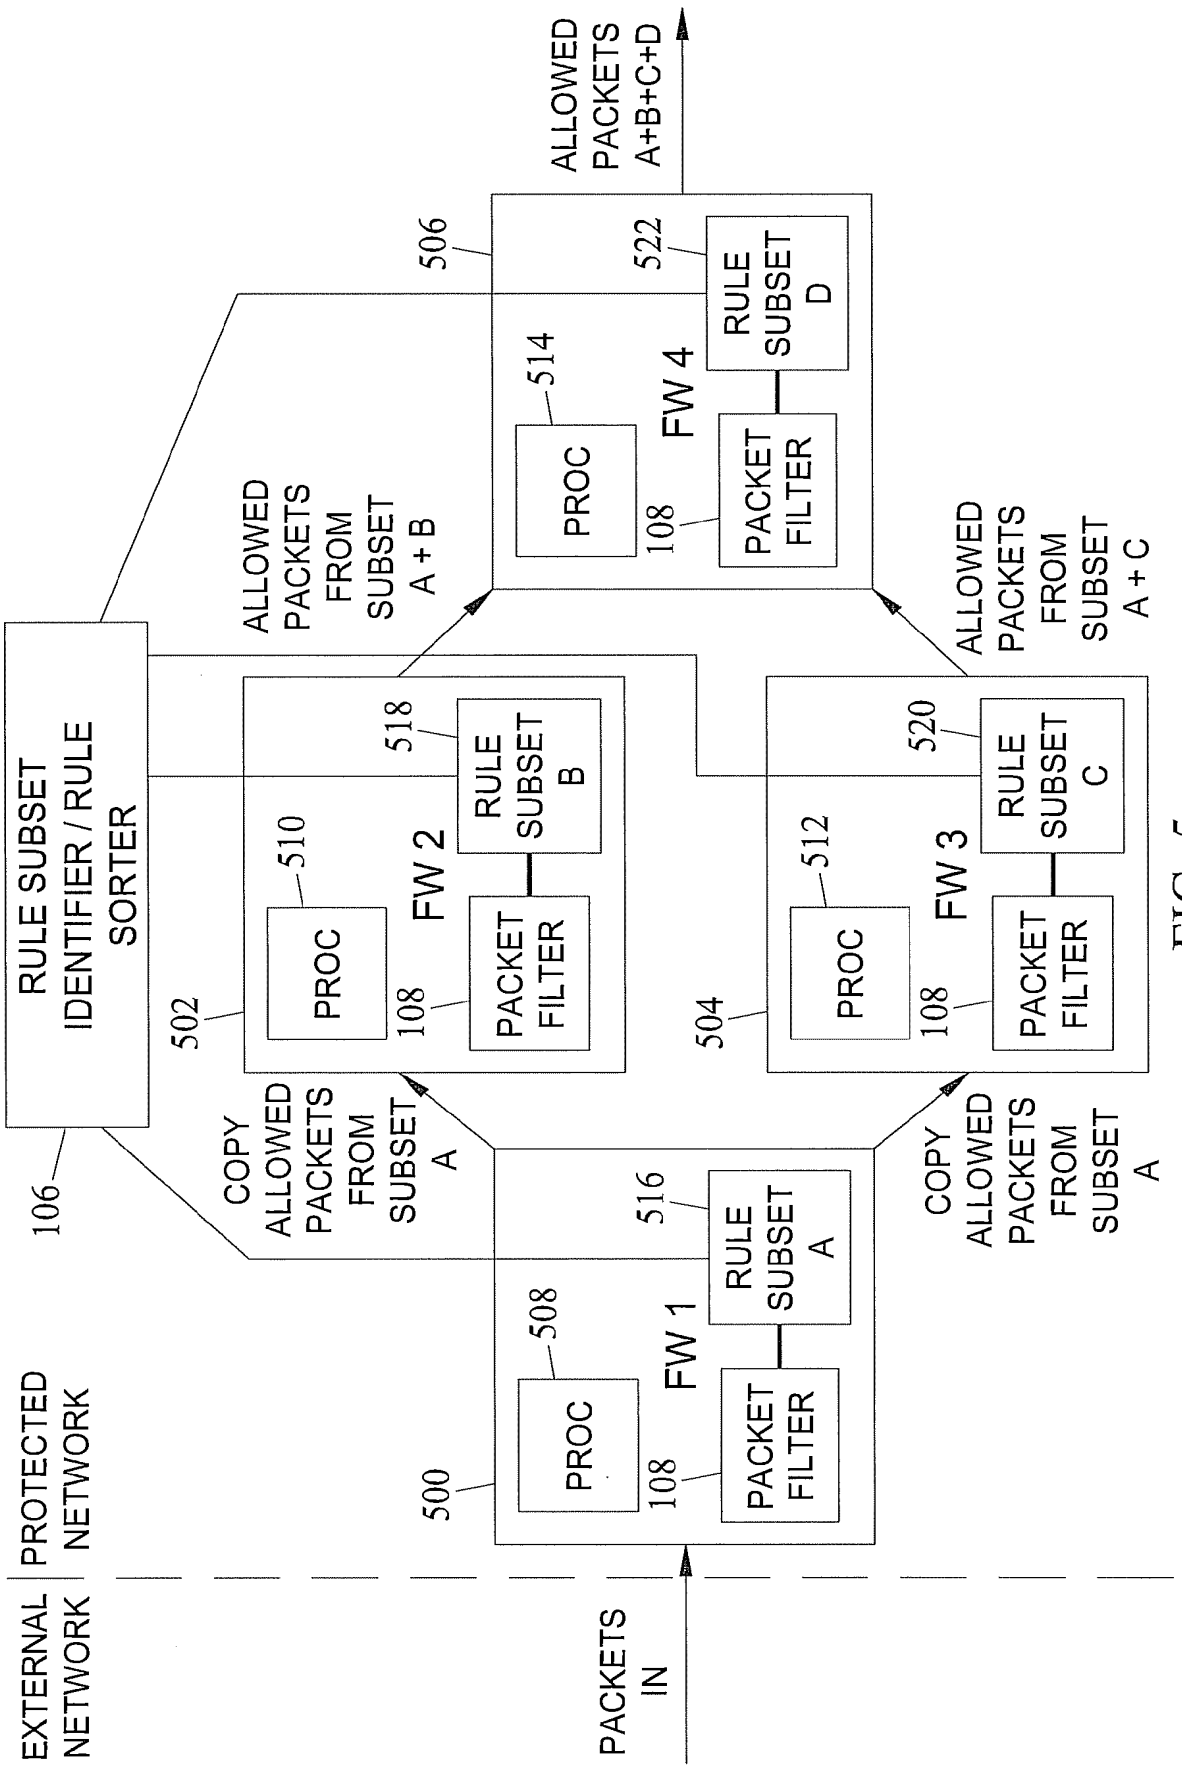
EXTERNAL NETWORK
PROTECTED NETWORK
PACKETS IN
106
RULE SUBSET IDENTIFIER / RULE SORTER
COPY ALLOWED PACKETS FROM SUBSET A
COPY ALLOWED PACKETS FROM SUBSET A
ALLOWED PACKETS FROM SUBSET A + B
ALLOWED PACKETS FROM SUBSET A + C
ALLOWED PACKETS A+B+C+D
500
PROC
508
PACKET FILTER
108
FW 1
RULE SUBSET A
516
502
PROC
510
PACKET FILTER
108
FW 2
RULE SUBSET B
518
504
PROC
512
PACKET FILTER
108
FW 3
RULE SUBSET C
520
506
PROC
514
PACKET FILTER
108
FW 4
RULE SUBSET D
522

FIG. 5

Espacenet

Bibliographic data: EP1484884 (A2) — 2004-12-08

Multi-layered firewall architecture

Inventor(s): SWANDER BRIAN D [US]; MAYFIELD PAUL G [US] ± (SWANDER, BRIAN D, ; MAYFIELD, PAUL G)

Applicant(s): MICROSOFT CORP [US] ± (MICROSOFT CORPORATION)

Classification: **- international:** ***G06F1/00; G06F11/00; G06F12/14; G06F13/00; G06F15/00; G06F17/00; H04L12/22; H04L12/56; H04L12/66; H04L29/06; H04L9/00; G06F;*** (IPC1-7): H04L29/06

- cooperative: **H04L63/0218**; **H04L63/0227**; **H04L63/0236**; **H04L63/0254**; H04L63/164

Application number: EP20040009147 20040416

Priority number (s): US20030456766 20030606

Also published as: EP1484884 (A3) US2005022010 (A1) US7509673 (B2) JP2004362581 (A) JP4572089 (B2) more

Abstract of EP1484884 (A2)

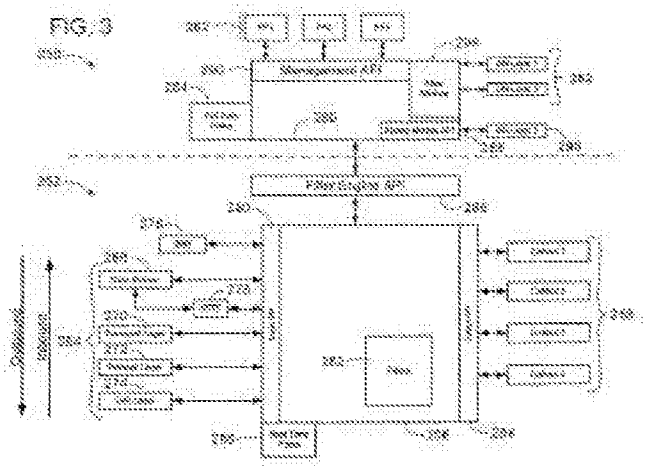

A method and system are provided for implementing a firewall architecture in a network device. The firewall architecture includes a plurality of network layers, a first firewall engine, and one or more callout modules. The layers send packets and packet information to the first firewall engine, maintain and pass packet context to subsequent layers, and process the packets. The first firewall engine compares the packet information to one or more installed filters and returns an action to the layers indicating how to treat the packet. The callouts provide additional functionality such as intrusion detection, logging, and parental control features.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 484 884 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
08.12.2004 Bulletin 2004/50

(51) Int Cl.[7]: **H04L 29/06**

(21) Application number: **04009147.2**

(22) Date of filing: **16.04.2004**

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IT LI LU MC NL PL PT RO SE SI SK TR
Designated Extension States:
AL HR LT LV MK

(30) Priority: **06.06.2003 US 456766**

(71) Applicant: **MICROSOFT CORPORATION**
Redmond, Washington 98052-6399 (US)

(72) Inventors:
- **Swander, Brian D.**
  Bellevue Washington 98004 (US)
- **Mayfield, Paul G.**
  Sammamish 98075 (US)

(74) Representative: **Grünecker, Kinkeldey,**
Stockmair & Schwanhäusser Anwaltssozietät
Maximilianstrasse 58
80538 München (DE)

(54) Multi-layered firewall architecture

(57) A method and system are provided for implementing a firewall architecture in a network device. The firewall architecture includes a plurality of network layers, a first firewall engine, and one or more callout modules. The layers send packets and packet information to the first firewall engine, maintain and pass packet context to subsequent layers, and process the packets. The first firewall engine compares the packet information to one or more installed filters and returns an action to the layers indicating how to treat the packet. The callouts provide additional functionality such as intrusion detection, logging, and parental control features.

FIG. 3

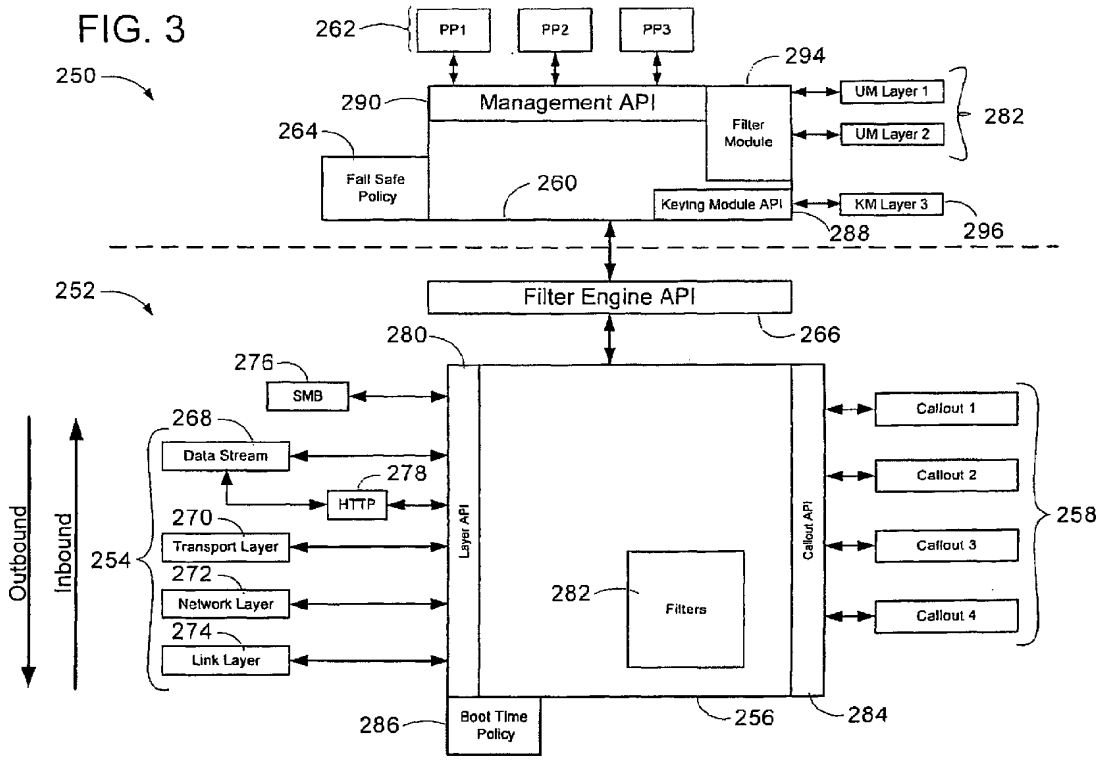

EP 1 484 884 A2

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 924

Description

RELATED APPLICATIONS

[0001] The present application contains subject matter related to that of patent application "Multi-Layer Based Method for Implementing Network Firewalls," attorney docket number 221039; "A Method and Framework for Integrating a Plurality of Network Policies," attorney docket number 221041; and "Method for Managing Network Filter Based Policy," attorney docket number 221037; filed on the same day as the present application, the disclosures of which are hereby expressly incorporated by reference.

TECHNICAL FIELD

[0002] This invention generally relates to computer systems and to network security. More particularly, this invention relates to a firewall architecture implemented in one or more network devices.

BACKGROUND OF THE INVENTION

[0003] Network protocols are designed to facilitate communication between network devices through an open exchange of data. While the open exchange of data greatly enhances the use of network devices to accomplish tasks, it also creates problems because network protocols are not designed for, and generally do not provide, network security. Computers coupled to both public and private networks, such as Local Area Networks (LANs), Wide Area Networks (WANs), intranets, and the Internet are susceptible to malicious attacks perpetrated by other network devices coupled either directly or indirectly to the network. Such malicious attacks include theft of data, Denial of Service (DOS) attacks, the proliferation of computer viruses, and the like. Other related issues arise when coupling computers to networks such as controlling access to undesirable or inappropriate web sites by children.

[0004] A firewall is a tool used to protect individual users, network devices, and networks in general, from malicious attacks, while also adding the ability to control the exchange of data over the network through implementation of a policy. The firewall implements the policy by examining network packets and determining, based on the examination, whether the packets should be permitted, or conversely blocked, from further traversing the network.

[0005] The policy implemented via the firewall is defined by one or more filters. Each filter includes filter parameters and an associated action. The filter parameters are used to identify network packets that are subject to the firewall policy and include information such as hardware addresses, e.g. Media Access Control (MAC) addresses, network addresses, e.g. Internet Protocol (IP) addresses, protocol type, e.g. Transport Control Protocol (TCP), port numbers, and the like. The action defines how packets with parameters that match the filter parameters should be treated. As a specific example, the filter includes as its parameters a Uniform Resource Locator (URL) address, e.g. "http://www.foo.com." The filter further associates the action of block, i.e. drop the packet, with that URL address. Whenever the firewall examines a packet and through that examination identifies the URL address "http://www.foo.com" as embedded in the packet, the firewall drops the packet thereby preventing it from traversing the network.

[0006] Network devices exchange data by sending and receiving packets through a network stack comprising a layered network architecture. While different network architecture models exist, most include at least an application layer, a transport layer, a network layer, and a link layer. Network packets traverse each layer sequentially and, as each layer is traversed, the packet is subject to processing. For outbound packets, the application layer processes data according to application protocols such as Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), and Simple Mail Transfer Protocol (SMTP) to name a few. Other layers, such as the network layer and the transport layer packetize the data by embedding it in TCP and IP headers. The layers perform reciprocal processing for inbound packets by, for example, parsing headers, unpacketizing data etc. The layered "stack" architecture and processing function performed by the layers results in a dynamic packet structure whereby packet content including the packet parameters change as the packet traverses the network protocol stack.

[0007] Firewalls examine packets at an inspection point located within the layered network stack. At one extreme, the inspection point is at the application layer. For example, the firewall is deployed as a Layered Service Provider (LSP). Packets at the application layer include the underlying data that will be transmitted to another network device or that has been received from another network device. Examining the packet at the application layer permits the firewall to identify application layer parameters, such as a URL address, and compare the application layer parameters to the filter parameters. However, other packet parameters such as IP addresses, port numbers, MAC addresses, and the like are not available because they have either not been added to outbound packets or have been parsed away from inbound packets.

[0008] At the other extreme, the firewall inspection point is implemented at lower levels of the network stack as an intermediate driver interposed between the link layer and the network layer. Packets at the lower levels of the network

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 925

stack include a maximum number of parameters, e.g. interface numbers, MAC addresses, IP addresses, protocol type, ports, and payload data. Although the packets include such parameters, it does not follow that the parameters are readily identifiable. After the firewall receives the packet, the firewall needs to parse and interpret the relevant packet parameter for comparison with the filter parameters. Thus, both the layers in the network stack and the firewall perform redundant packet parsing and interpretation functions.

SUMMARY OF THE INVENTION

[0009] The present invention is directed towards a firewall architecture that permits application of firewall filters to network packets at all layers within a network stack. The architecture includes a set layer processors capable of processing layer parameters from a network packet. The layer parameters are parameters associated with the layer processor and include parameters that the layer processor parses from, adds to, or otherwise derives from the network packet. The layer processors are further capable of issuing classification requests that include the layer parameters.

[0010] The architecture also includes a first firewall engine comprising a layer interface, a set of installed filters, and a lookup component. The layer interface receives the layer parameters sent as part of the classification request. The first firewall engine lookup component uses the layer parameters to identify one or more matching filters from the set of installed filters. Each matching filter includes an action that instructs the layer processor whether to permit the packet to further traverse the network or whether to block the packet. The first firewall engine returns at least one action to the layer processor via the layer interface.

[0011] In an embodiment of the invention, the layer processors maintain packet context. The layer processor receives the packet context from a previous layer processor and sends the packet context to the first firewall engine. The lookup component in the first firewall engine uses the packet context in combination with the layer parameters to identify the one or more matching filters. The layer processor also modifies the packet context by adding the layer parameters and then sending the modified packet context to a next layer processor.

[0012] In another embodiment of the invention, one or more callouts are included as part of the firewall architecture. Each of the one or more callout modules provide added functionality (beyond permit and block) such as packet logging functions, Internet Protocol Security (IPSec) verification, and execution of parental controls. The callout module is executed on the packet when it is identified as the action in one of the matching filters.

[0013] In another embodiment of the invention, a second firewall engine is provided. The second firewall engine adds new filters into the set of installed filters via a management API. The second firewall engine also includes an instance of the first firewall engine, referred to as a filter module that replicates the services of the first firewall engine within the second firewall engine.

[0014] Additional features and advantages of the invention will be made apparent from the following detailed description of illustrative embodiments, which proceeds with reference to the accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] While the appended claims set forth the features of the present invention with particularity, the invention, together with its objects and advantages, may be best understood from the following detailed description taken in conjunction with the accompanying drawings of which:

[0016] Figure 1 is a block diagram generally illustrating an exemplary computer system on which the present invention resides;

[0017] Figure 2 is a block diagram generally illustrating an exemplary network environment whereby the present invention is employed;

[0018] Figure 3 is a block diagram generally illustrating the firewall architecture of the present invention;

[0019] Figure 4 is a block diagram illustrating an exemplary filter used with the present invention;

[0020] Figure 5 is a block diagram illustrating an exemplary data structure for packet context used with the present invention;

[0021] Figure 6 is a block diagram illustrating an exemplary set of application programming interfaces used with the present invention;

[0022] Figure 7 is a block diagram illustrating an exemplary application programming interface used with the present invention;

[0023] Figure 8 is a block diagram illustrating functions performed by network layers according to the present invention; and

[0024] Figure 9 is a block diagram illustrating an exemplary set of callouts used with the present invention.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 926

DETAILED DESCRIPTION OF THE INVENTION

[0025] A firewall architecture that permits filtering at multiple layer processes, referred to herein as layers, is disclosed. The architecture includes a user mode process and a kernel mode process executing in an operating system. Alternatively, the firewall architecture is executed outside of the operating system in one or more program modules or within a single operating system mode.

[0026] The kernel mode process includes a plurality of layers comprising a protocol stack, a kernel firewall engine, and one or more callouts. The protocol stack includes an application layer, a transport layer, a network layer, and link layer. Additional layers are added or deleted from the architecture as desired. The layers each form a requesting layer capable of receiving a network packet and corresponding packet context data from a previous layer or process. The requesting layer then issues a classification request to the kernel firewall engine via a layer API. The classification request includes the packet received by the requesting layer, the packet context, and a set of layer parameters associated with the requesting layer. The kernel firewall engine processes the request and returns an action. The action, by way of example, instructs the requesting layer how to treat the packet (e.g. permit or block). If the action is permit, the requesting layer processes the packet according to a layer protocol, modifies the packet context to include the layer parameters, and passes the packet and packet context to a next layer. If the action is block, the requesting layer drops the packet and does not pass the packet to the next layer. The requesting layer may perform additional functions as a result of the block action, such as tearing down a TCP connection where a TCP packet is dropped.

[0027] The kernel firewall engine includes the layer API, a set of installed filters, and a callout API. Each of the set of installed filters includes a set of filter conditions and an associated action. The kernel firewall engine processes the classification request sent from the requesting layer by identifying one or more matching filters. The matching filters have filter conditions that match the layer parameters and packet context. Once the matching filters are identified, they are applied in order of filter priority. If the action of the filter being applied is permit or block, the action is returned to the requesting layer. If the action is callout, the classification request issued by the requesting layer, along with matching filter identification, is passed to one of the callout modules. The callout module performs its programmed function, and returns an action to the kernel firewall engine.

[0028] The kernel firewall engine applies matching filters in order of priority to the packet until at least a terminating action is designated. Once the terminating action is designated, it is returned to the requesting layer and no additional matching filters are applied. If no matching filters are identified for a packet, the requesting layer is notified that no matching filters were found and the requesting layer then decides how to treat the packet. Typically, where no matching filter is identified, the requesting layer treats the packet as if a permit action was returned.

[0029] Exemplary user mode processes include a user mode firewall engine and one or more policy providers. The policy providers obtain policy from any suitable source such as volatile or non-volatile memory or is user entered through a graphical user interface or similar tool. The policy is a source of information for rendering a new filter, including the set of filter conditions and associated actions.

[0030] The user mode also includes an instance of the kernel firewall engine within the user firewall engine thereby permitting user mode layers to be created. The user mode layers then use the user mode instance of the kernel firewall engine to install filters, and identify filters that match a set of parameters allowing the application of filtering within the user mode.

[0031] In an embodiment of the invention, a callout interface from the firewall engine to a set of callout modules enables virtually unlimited extension of the firewall capabilities. By way of example, an HTTP context callout provides a parental control feature by identifying acceptable and unacceptable URL addresses. An IPSec callout verifies that packets that are supposed to use IPSec have been properly subject to IPSec processing. A logging callout logs packets meeting established criteria thereby facilitating later examination of packets. An intrusion detection callout identifies suspicious packets based on known algorithms.

[0032] Turning to the drawings, wherein like reference numerals refer to like elements, the invention is illustrated as being implemented in a suitable computing environment. Although not required, the invention will be described in the general context of computer-executable instructions, such as program modules, being executed by a personal computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0033] Figure 1 illustrates an example of a suitable computing system environment 100 on which the invention may be implemented. The computing system environment 100 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the computing environment 100 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment 100.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 927

[0034] The invention is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

[0035] The invention may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

[0036] With reference to Figure 1, an exemplary system for implementing the invention includes a general purpose computing device in the form of a computer 110. Components of computer 110 may include, but are not limited to, a processing unit 120, a system memory 130, and a system bus 121 that couples various system components including the system memory to the processing unit 120. The system bus 121 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Associate (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus.

[0037] Computer 110 typically includes a variety of computer readable media. Computer readable media can be any available media that can be accessed by computer 110 and includes both volatile and nonvolatile media, removable and non-removable media. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computer 110. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of the any of the above should also be included within the scope of computer readable media.

[0038] The system memory 130 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 131 and random access memory (RAM) 132. A basic input/output system 133 (BIOS), containing the basic routines that help to transfer information between elements within computer 110, such as during start-up, is typically stored in ROM 131. RAM 132 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 120. By way of example, and not limitation, Figure 1 illustrates operating system 134, application programs 135, other program modules 136, and program data 137.

[0039] The computer 110 may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, Figure 1 illustrates a hard disk drive 141 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 151 that reads from or writes to a removable, nonvolatile magnetic disk 152, and an optical disk drive 155 that reads from or writes to a removable, nonvolatile optical disk 156 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 141 is typically connected to the system bus 121 through a non-removable memory interface such as interface 140, and magnetic disk drive 151 and optical disk drive 155 are typically connected to the system bus 121 by a removable memory interface, such as interface 150.

[0040] The drives and their associated computer storage media discussed above and illustrated in Figure 1, provide storage of computer readable instructions, data structures, program modules and other data for the computer 110. In Figure 1, for example, hard disk drive 141 is illustrated as storing operating system 144, application programs 145, other program modules 146, and program data 147. Note that these components can either be the same as or different from operating system 134, application programs 135, other program modules 136, and program data 137. Operating

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 928

system 144, application programs 145, other program modules 146, and program data 147 are given different numbers hereto illustrate that, at a minimum, they are different copies. A user may enter commands and information into the computer 110 through input devices such as a keyboard 162 and pointing device 161, commonly referred to as a mouse, trackball or touch pad. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 120 through a user input interface 160 that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 191 or other type of display device is also connected to the system bus 121 via an interface, such as a video interface 190. In addition to the monitor, computers may also include other peripheral output devices such as speakers 197 and printer 196, which may be connected through a output peripheral interface 195.

[0041] The computer 110 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 180. The remote computer 180 may be another personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the personal computer 110, although only a memory storage device 181 has been illustrated in Figure 1. The logical connections depicted in Figure 1 include a local area network (LAN) 171 and a wide area network (WAN) 173, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

[0042] When used in a LAN networking environment, the personal computer 110 is connected to the LAN 171 through a network interface or adapter 170. When used in a WAN networking environment, the computer 110 typically includes a modem 172 or other means for establishing communications over the WAN 173, such as the Internet. The modem 172, which may be internal or external, may be connected to the system bus 121 via the user input interface 160, or other appropriate mechanism. In a networked environment, program modules depicted relative to the personal computer 110, or portions thereof, may be stored in the remote memory storage device. By way of example, and not limitation, Figure 1 illustrates remote application programs 185 as residing on memory device 181. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

[0043] In the description that follows, the invention will be described with reference to acts and symbolic representations of operations that are performed by one or more computer, unless indicated otherwise, As such, it will be understood that such acts and operations, which are at times referred to as being computer-executed, include the manipulation by the processing unit of the computer of electrical signals representing data in a structured form. This manipulation transforms the data or maintains it at locations in the memory system of the computer, which reconfigures or otherwise alters the operation of the computer in a manner well understood by those skilled in the art. The data structures where data is maintained are physical locations of the memory that have particular properties defined by the format of the data. However, while the invention is being described in the foregoing context, it is not meant to be limiting as those of skill in the art will appreciate that various of the acts and operation described hereinafter may also be implemented in hardware.

[0044] A network environment wherein the firewall architecture of the present invention is used will now be described with reference to FIG. 2. The network is exemplary in nature as the firewall architecture of the present invention is implemented in any network device that is coupled to any network configuration. The network environment includes a private network 200 and a public network 202. The private network 200 and public network 202 are any suitable type such as Local Area Networks (LANs), Wide Area Networks (WANs), intranets, the Internet, or any combination thereof.

[0045] The network environment includes a plurality of network devices 204, 206, 208, 210 and 212. Network devices 204, 206 are coupled to the private network 200. Network devices 210, 212 are coupled to the public network 202. The network device 208 is coupled to, and provides an interface between, both the private network 200 and the public network 202. The network devices are coupled to the public and private networks using any suitable technology such as Ethernet, 1394, or 802 11 (b). The network devices are further implemented as any suitable computing devices such as personal computers, servers, handheld devices, printers, switches, routers, bridges, repeaters or the like.

[0046] The network device 208 includes a firewall 214 and one or more filters 216. The firewall 214 is a program module or a set of program modules, implemented according to the architecture of the present invention, that examine network packets exchanged between the network devices 204, 206, 208 coupled to the private network 200 and the network devices 210, 212 coupled to the public network 202. In an embodiment of the invention, the firewall 214 also examines locally destined network packets sent from, and destined to, network devices within the private network 200.

[0047] The firewall 214 is implemented in network device 208 to protect and control network traffic exchanged between the private network 200 and the public network 202, referred to as an edge firewall. Alternatively, the firewall 214 is implemented in, and protects, a single network device such as illustrated in the network device 210, referred to as a host firewall. The firewall is also capable of being implemented as a centrally administered set of hosts and/or edge firewalls in a synchronized manner, referred to as a distributed firewall. Placement of the network device or devices implementing the firewall 214 is preferably selected such that the firewall 214 examines all network traffic

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 929

destined for the network devices it is supposed to protect.

[0048] The filters 216 are implemented as part of the firewall 214. Alternatively, the filters 216 are implemented as part of separate data structure accessible by the firewall 214. The firewall 214 and filters 216 execute a firewall policy designed to protect network devices 204, 206, 208 from malicious attacks originating from network devices 210, 212 coupled to the public network. The firewall 214 also provides added functionality such as facilitating parental controls, intrusion detection, logging of network packets, and other added filter based functionality.

[0049] Each filter 216 includes a set of filter conditions and one or more associated actions. The filter conditions include parameters and information that can be parsed or otherwise obtained from network packets such as interface numbers, hardware addresses, network addresses, protocol type, port numbers, and payload data. The one or more associated actions define how the network device implementing the firewall should treat packets that match the filter conditions. Typical actions include permit, i.e. allow the packet to continue network traversal, and block, i.e. preclude further network traversal by dropping packet.

[0050] The firewall 214 examines network packets traversing the network when received at network device 208 and, by comparing the packet parameters to the filter conditions, identifies one or more matching filters. A matching filter results when the filter conditions match the packet parameters. The packet parameters, like filter conditions, include information that is parsed or otherwise obtained from packets. When the firewall identifies a matching filter, the one or more actions associated with the filter conditions are executed.

[0051] FIG. 3 shows an overall view an embodiment of the firewall architecture of the present invention. The architecture provides the ability to subject packets to filters at all layers of a network stack. The firewall provides central management capabilities that allow filters to be added and deleted and for filter conflicts to be identified and resolved. The firewall architecture is extensible in that filter layers are added and deleted as needed and is expanded to include specialized functionality beyond permit and block actions. Although the invention is described with specific reference to firewalls and firewall filters, it is also used to facilitate and manage other filters and policies. As specific examples, the present invention is suitable for facilitating and managing filters used with Quality of Service (QOS), Internet Protocol Security (IPSec) suite, as well as other encrypting, authenticating, and key management protocols.

[0052] The firewall architecture includes a user mode process 250 and a kernel mode process 252. The user mode process 250 and kernel mode process 252 execute as part of an operating system in a network device. Those skilled in the art will appreciate that the user mode process 250 and the kernel mode process 252 of the operating system include additional components that for simplicity are not shown. Alternatively, the firewall architecture is executed in whole or in part outside of the operating system as one or more program modules or application programs or within a single operating system process.

[0053] The kernel mode process 252 includes a network stack 254, a first firewall engine referred to herein as a kernel firewall engine 256, and optional callouts 258. Collectively, the kernel mode process 252 implements an established firewall policy by identifying matching filters for network packets, processing packets according to known protocols, and executing other actions on the packet as designated by matching filters.

[0054] The network stack 254 comprises a plurality of layer processes, referred to herein as layers, including a data stream layer 268, a transport layer 270, a network layer 272, and a link layer 274. The firewall architecture of the present invention is extensible and additional layers are dynamically added and deleted as needed. An example of an added layer includes file access layer 276, implemented according to a Server Master Block (SMB) protocol. The layers may work in cooperation with other program modules such as a Hypertext Transfer Protocol (HTTP) parser module 278.

[0055] The layers in the network stack 254 process inbound and outbound network packets. Outbound network packets are packets being transmitted from the network device implementing the firewall architecture onto the network. Inbound packets are packets received at the network device implementing the firewall architecture. As indicated by the respective arrows shown in Fig. 3, inbound packets traverse the network stack 254 from bottom to top and outbound packets traverse the network stack 254 from top to bottom.

[0056] Network packets traverse, and are processed by, the network layers sequentially. In accordance with known techniques, each layer in the network stack 254 is capable of receiving packets from a previous layer or process, processing the packet according to a specification or protocol, and sending the processed packet to a next layer or process. In accordance with the invention, each layer in the network stack 254 also maintains packet context, passes the packet context to the next layer, issues a classification request to the kernel firewall engine 256, and takes action on the packet according to the firewall policy..

[0057] The packet context is a data structure that follows the packet from layer to layer. Each layer maintains the context by adding into the context data structure a set of parameters that the layer is designed to process, e.g. the information that the layer is designed to parse from inbound packets, add to outbound packets or otherwise derive from contents of packets. An exemplary data structure used for the packet context is described with reference to Fig. 5.

[0058] One of the operations carried out by the layers of the network stack 254 is to invoke the kernel firewall engine by issuing the classification request. The classification request is a call by a layer in the network stack 254 requesting that any filters that match the packet be identified and any associated action, e.g. permit or block, be returned to the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 930

layer. The layer issuing the classification request is referred to herein as a requesting layer. Each layer also takes the action on the packet that is returned by the kernel firewall engine 256.

[0059] The kernel firewall engine 256 includes a layer API 280, a set of installed filters 282, and a callout API 284. The kernel firewall engine 256 performs various functions including (1) maintaining the set of installed filters 282 defining the firewall policy, (2) receiving classification requests from the layers in the network stack 254, (3) identifying one or more matching filters based on the classification request, and (4) instructing the requesting layer of any action to take on the packet based on the matching filters.

[0060] Each filter of the set of installed filters includes a set of filter conditions and one or more associated actions. As described with reference to FIG. 2, the filter conditions identify the network packets that are subject to the associated filter action. The actions specified in the set of installed filters 282 include permit and block. Additional functionality is added via optional callouts 258. An exemplary form of the filters is described with reference to Fig. 4.

[0061] The layer API 280 provides an interface between the layers in the network stack 254 and the kernel firewall engine 256. Through the layer API 280, the requesting layer issues the classification request to the kernel firewall engine 256. The classification request includes the packet as received by requesting layer, the packet context as received by the requesting layer, and layer parameters. The layer parameters are packet parameters processed by, e.g. added or parsed, by the requesting layer. As a specific example, source and destination Internet Protocol (IP) addresses are layer parameters sent by the network layer 272 when implementing the IP protocol. Layer parameters may also include information beyond the packet parameters that are added to or parsed from the packet. As a specific example, the layer parameters include a local address type. The local address type is determined by the IP layer and sent as part of the classification request. Local address types include unicast, broadcast, multicast, anycast and the like. A specific implementation of the layer API 280 is described with reference to FIG. 6.

[0062] Optionally, the callouts 258 are used to implement added functionality beyond the permit and block filter actions. A callout is executed when the kernel firewall engine 256 identifies a matching filter for the packet that includes, as the associated action, a callout to one of the callout modules. The kernel firewall engine sends the classification request as issued by the requesting layer, i.e. full packet, layer parameters, and packet context, to the callout module, along with an identification of the matching filter, via the callout API 284. In an embodiment of the invention, the firewall architecture includes a base set of callouts 258. Additional callouts, like layers, are added as needed to the extensible firewall architecture. A specific implementation of the callout API 284 is described with reference to Fig. 6.

[0063] The user mode process 250 includes a second firewall engine referred to as a user firewall engine 260 and one or more policy providers 262 identified as "PP1," "PP2," and "PP3." The policy providers 262 are processes that add firewall policy, i.e. installed filters 282, into the firewall architecture. Any process is used to accomplish this task. An example is a legacy IPSec policy service (LIPS). The legacy IPSec policy service adds filters defining network traffic that is supposed to use IPSec protocols such as Encapsulating Security Protocol (ESP) and Authentication Header Protocol (AH). As a specific example, the legacy IPSec policy adds a firewall policy indicating that all unsolicited inbound packets must be encrypted according to the ESP protocol. The policy further provides that any unsolicited inbound packet in clear text, i.e. an unencrypted packet, should be blocked. The policy providers 262 obtain the policy from any suitable source such as data in volatile or nonvolatile memory, or a Graphical User Interface (GUI) that permits an administrator or system users to directly enter policy. The user firewall engine 260 converts the policy into a new filter, i.e. defines the policy in terms of filter conditions and associated actions, and adds the new filter into the set of installed filters 282.

[0064] The user firewall engine 260 also performs filter arbitration and conflict resolution functions. When the policy provider 262 provides new policy to the user mode firewall engine 260, the user firewall engine determines whether the new filter resulting from the new policy conflicts with any of the installed filters 282. An example of a method of identifying and resolving conflicts that is suitable for use in the framework of the present invention is described in United States Patent Application entitled "Method for Managing Network Filter Based Policy," attorney docket number 221037.

[0065] The architecture further includes a filter engine API 266 that forms an interface between the user mode firewall engine 260 and the kernel firewall engine 256. The filter engine API 266 provides a mechanism for the user firewall engine 260 to add new filters to or remove existing filters from the set of installed filters 282, and to examine installed filters 282 so that filter conflicts can be detected and resolved. A management API 290, accessible by policy providers 262, also provides a mechanism to add and remove filters from the architecture.

[0066] The user mode firewall engine 260 also includes a filter module 294. The filter module 294 is an instance of the kernel firewall engine 256 in the user mode 250. The instance of the filter module 294 in the user mode firewall engine 260 permits the user firewall engine 260 to replicate the services of the kernel firewall engine 256 for one or more user mode layer processes referred to herein as user mode layers 282. The user mode layers 282 are added in the same way that kernel mode layers are created. Because the filter module 294 is the user mode instance of the kernel firewall engine 256, it will be understood that any functionality described herein for the kernel firewall engine 256 also applies to the filter module 294.

[0067] A keying module API 288 provides an interface between the user policy engine 260 and keying module 296.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 931

The keying module API 288 is used to signal the keying module of the need to establish an SA.

[0068] The firewall architecture includes a boot time policy 286 and a failsafe policy 264. The boot time policy 286 and the failsafe policy 264 are implemented to provide a basic set of installed filters 282 while the system is in a transitional state such as when the network device is initialized, e.g. turned on, reset, or restarted. During the transitional state, there is a latent period during which the user mode 250 is not established and it is not possible to install filters via policy providers 262 and the user firewall engine 260. As a result, the network device is vulnerable to malicious attacks from, for example, unsolicited inbound packets. The boot time policy 286 resides in the kernel mode 252 and is loaded into the kernel firewall engine 256 prior to the time the user mode 250 is established. The boot time policy 296 provides the basic installed filter 282 configurable to provide any suitable filter conditions and associated actions for protecting the system during the transitional state. An example of the boot time policy 286 is one that blocks all inbound and outbound network packets.

[0069] Once the user mode 250 is established, the transitional state continues while the policy providers 262 are identifying policy and attempting to add installed filters 282 into the kernel firewall engine 256. When the user mode 250 is established, the user firewall engine 260 first loads the fail-safe policy 264 into the kernel firewall engine 256 before any policy is accepted by the policy providers 262. The boot time policy 286 is then disabled. Like the boot time policy 286, the fail-safe policy 264 is a basic filter designed to protect the system from attacks during the transitional state when the policy providers and user firewall engine are still in the process of installing filters. Once the policy providers 262 successfully install their respective filters into the network device, the fail-safe policy 264 is disabled.

[0070] Referring to FIG. 4, the set of installed filters 282 will now be described. Each filter 310 of the set of installed filters 282 has a plurality of fields including a filter Id 312, a weight 314, one or more actions 316, a policy context 328, and a set of filter conditions 318. The filter Id 312 provides a unique identification to the filter. The filter Id 312 is used, for example, as a means for the kernel firewall engine 256 to return matching filter information to the user firewall engine 260 and the callouts 258. In an embodiment of the invention, the filter is assigned to one of the layers in the network stack 254. The filter Id 312 is used by the kernel firewall engine 256 to track which filter is assigned to which layer.

[0071] The weight field 314 includes a value that identifies priority of the filter 310. The higher the value in the weight field 314, the higher the priority of the filter. The filter priority determines the order whereby matching filters are applied to the packet by the kernel firewall engine 256.

[0072] In an embodiment of the invention, the filter with the highest priority, i.e. highest weight value, is applied first, then the next highest priority filter is applied and so on, until a matching filter with a terminating action is encountered. Terminating actions are described in more detail below. Once the matching filter with the terminating action is applied, the kernel firewall engine 256 stops applying matching filters. Thus, the action 316 specified by lower priority matching filters are not taken on the packet after the terminating action is applied. Alternatively, the firewall engine 256 identifies a single matching filter and returns a set of actions from the single matching filter.

[0073] The set of filter conditions 318 determine whether a packet matches the filter 310. Each filter condition 318 includes a type 320, data 322, and a layer Id:field Id 324. The type 320 defines the length and number of variables included in the corresponding data 322. The architecture provides for predefined known variable types such as Byte, Short, Long, 8 Bytes, String, Internet Protocol version 4 (IPv4) Address, Internet Protocol version 6 (IPv6) Address, IPv4 Address plus Mask, IPv6 Address plus Mask, and Address Range.

[0074] The data field 322 includes data matching the type. For example, if the type is IPv4 Address, an acceptable value for the data field 322 is a 32-bit number in range of 00.00.00.00 to 255.255.255.255 as expressed in dotted decimal notation. In some instances, the type 320 provides for multiple values in the data field 322. The Address Range, IPv4 Address plus Mask, and IPv6 Address plus mask types allow two IP address values, defining a beginning and ending range of IP addresses. For maximum flexibility, the architecture also permits user-defined types. Alternatively, additional types are manually added to the system architecture.

[0075] The layer Id:field Id 324 is used to identify an originating layer and a parameter from the originating layer, respectively. The originating layer and the parameter from the originating layer define packet parameters, i.e. layer parameters and packet context that the data 322 is compared against when identifying matching filters. The originating layer identifies a layer in the network stack. The parameter from the originating layer identifies a specific parameter associated with the originating layer. A specific example is illustrated by filter condition 326. The type is IPv4 thereby indicating the data 322 is a 32-bit IP address. The layer Id is "IP" representing that the 32 bit number is an IP, i.e. network, layer parameter. The field Id is "Src IP Addr" which in the example represents a source IP address. The IP address provided in the data field 322 is "123.3.2.1" indicating that any packet with that source IP address meets the filter condition thereby matching the filter. The filter 310 may include multiple filter conditions 318, in which case a packet matches the filter 310 only if all filter conditions are met.

[0076] The action 316 in the filter 310 is permit, block, or callout. If the action 324 in the filter 310 is permit or block and the packet matches the filter 310, the permit or block action is returned to the requesting layer by the kernel firewall engine 256. If the action 316 is callout, the kernel firewall engine 256 issues its own classification request that includes

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 932

the full packet, layer parameters, context and identification of the matching filter to the designated callout module 258. The callout module 258 performs its programmed function on the packet, e.g. intrusion detection. The callout may return an action (either permit, block) to the kernel firewall engine, which in turn relays the action to the requesting layer. The callout is also capable of maintaining packet context that is likewise returned to the requesting layer via the kernel firewall engine 256. The action may also be a value, such as a null value, indicating that no firewall policy exists.

[0077] Actions are designated terminating or non-terminating. If the action is terminating, once that action is identified for a packet, the action is returned to the requesting layer and no additional filters are applied to the packet. If the action is non-terminating, additional filter lookups continue until a terminating action for the packet is identified. As a default, permit and block are designated terminating actions.

[0078] The policy context 328 is used to store policy other than firewall policy, such as security policy or QOS policy. The policy context is any suitable data structure. For example, the policy context is a 64 bit number that is interpreted by a process that added the policy context. The policy context and/or action may be a null value.

[0079] FIG. 5 illustrates an example of a data structure 330 used for the packet context that is maintained by, and passed to, the layers in the network stack 254 and the callout modules 258. The packet context 330 follows the inbound or outbound network packet as it traverses the layers and includes one or more entries, labeled 336-340. Each entry includes a layer Id:field Id 332 and a corresponding value 334.

[0080] The layer Id:field Id 332 has the same meaning as the layer Id:field Id 324 provided as part of the filter conditions 318 in the filters 310 (FIG. 4). Namely, the layer Id:field Id 322 identifies the originating layer and the layer parameter from the originating layer for the data in the value field 334. The value field 334 includes a specific layer parameter.

[0081] As a specific example, entry 336 includes the layer Id:field Id 332 "NDIS: Src. MAC Addr." "NDIS" represents a Network Driver Interface Specification implementation of the Link layer 274 (Fig. 1). "Src MAC addr." represents a source MAC address. Thus, the layer:field Id 332 indicates that the data in the value field 334 is a source MAC address that was processed by the NDIS (Link) layer. The value field 334 includes the actual source MAC address which in the example is "00.08.74.4F.22.E5" as expressed in hexadecimal notation.

[0082] As a second example, entry 338 has a layer Id:field Id 332 of "NDIS:IF No." This again identifies the layer as NDIS, but in this case identifies the parameter as "IF No" representing an interface number as the specific NDIS parameter. The value field 334 includes the actual interface number, which in this case is 2.

[0083] As a third example, entry 340 has a layer Id:field Id 332 of "IP:Dst IP Addr." The "IP" represents the network layer using the IP protocol and the "Dst IP Addr" represents a destination IP address as the IP layer parameter. The value field 334 includes the actual destination IP address of "123.3.2.1."

[0084] Having described the underlying firewall architecture of the present invention, attention is drawn to the functional interfaces of the system and exemplary methods that are performed using the underlying firewall architecture described herein. The functional interfaces are implemented as a plurality of application programming interfaces (APIs) The APIs include the layer API 280, the callout API 284, the filter engine API 266, and the keying module API 288 as shown in FIG. 6 and FIG. 7.

[0085] The layer API 280 facilitates data exchange between each of the layers in the network stack 254 and the kernel firewall engine 256.. As shown, the layer API 280 includes a Classify method 402, an Add Layer method 404, and a Delete Layer method 406. The layer API 280 may also execute within the filter module 294 to facilitate communication between the user mode layers and the filter module 294.

[0086] The Classify method 402 is used by the requesting layer to send layer parameters, the packet as received by the requesting, and the packet context to the kernel firewall engine 256. The kernel firewall engine 256 compares the (1) layer parameters from the requesting layer and (2) packet context entries to the filter conditions 318 in each filter 310 assigned to the requesting layer to identify matching filters. The following is an exemplary implementation of the Classify method. It will be understood that the following methods are described as receiving or returning data values. According to known programming techniques, the methods may use pointers to data values instead of actual data values.

```
NTSTATUS
WFPClassify
(
```


CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 933

```
    IN ULONG                          LayerId,
    IN WFP_INCOMING_VALUES*           pInFixedValues,
    IN WFP_INCOMING_CONTEXT_VALUE*    pInContext,
    PVOID                             pPacket,
    OUT WFP_ACTION_TYPE*              pActionType,
    OUT UINT64*                       pOutContext
);
```

[0087] Where the following characterizes the recited parameters.

[0088] **LayerId** identifies the network layer issuing the classification request, i.e. the requesting layer. Referring to FIG. 3, the Layer Id identifies the layer as the data stream layer 268, the transport layer 270, the network layer 272 or the link layer 274. Other layers are valid if added to the system. For example, if the SMB layer 276 is added, it has its own unique identification. The firewall architecture of the present invention further permits multiple protocol implementations at a layer in the network stack 254. For example the stack has two transport layers 270, a first transport layer uses the TCP protocol and a second transport layer uses the UDP protocol. User mode layers 282 may also be added thereby forming a valid layer.

[0089] **pInFixedValues** includes a subset of the layer parameters processed by the requesting layer. The pInFixedValues, along with the packet context entries, are compared to the filter conditions to determine if the packet matches the filter. An example of default layer parameters included in the pInFixedValues for each layer are identified in Table A below. It will be understood that the following values are exemplary only as an advantage of the architecture is that it allows the layers to use any values accessible to that layer.

Table A

Layer	Default Layer Parameters
Link Layer	Source and Destination MAC Addresses; Interface Number
Network Layer	Source and Destination IP Addresses; Protocol Type; Local Address Type
Transport Layer	Source and Destination Port Numbers;
Application	Deciphered application layer protocol payload

[0090] **pInContext** includes the context data structure 330 (FIG. 5) as received by the requesting layer. The packet context is used by the kernel firewall engine 256 in conjunction with the layer parameters to identify matching packets.

[0091] **pPacket** includes the entire packet as received by the requesting layer. The pPacket is not used by the kernel firewall engine 256 to identify matching filters. As previously described, the kernel firewall engine 256 uses the pInFixedValues and pInContext to identify matching filters. The pPacket is included in the Classify method so that the kernel firewall engine 256 can send it to one or more callout modules 258 identified as the action 316 in matching filters.

[0092] **pActionType** includes the action 316 that is returned to the requesting layer. The action 316 returned is permit or block as identified in the matching filter, or callout modules executed by the matching filter.

[0093] **pOutContext** includes the policy context data. As previously described, the policy context is used to facilitate network policies associated with IPSec, QOS and any other non-firewall filter based policy.

[0094] The Add Layer 404 and Delete Layer 406 methods are used to add and remove a layer from the firewall architecture, respectively. The following is an exemplary form of the Add Layer 404 method.

```
NTSTATUS
AddExtensionLayer(OUT PULONG pLayerId);
```

[0095] Where the following characterizes the recited parameters.

[0096] **pLayerId** is a unique layer identification value returned to the layer being added, i.e. the layer executing the Add Layer method.

[0097] The following is an exemplary form of the Delete Layer 406 method.

```
NTSTATUS
RemoveExtensionLayer(ULONG LayerId);
```

[0098] Where the following characterizes the recited parameters.

[0099] **LayerId** identifies the layer being removed; i.e. the layer executing the Delete Layer method.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 934

[0100] The Callout API 284 facilitates data exchange between the kernel firewall engine 256 and the callouts 258. The Callout API 284 includes a Classify method 410, a Notify method 412, a Register Callout method 414, and a Deregister Callout method 416. The Classify method 410 of the Callout API 284 is similar to the Classify method 402 of the Layer API 280 except that it also includes matching filter data. The following is an exemplary form of the Classify method 410 used to execute a callout.

```
typedef NTSTATUS (*WFP_CALLOUT_CLASSIFY_FN)
(
        IN const WFP_INCOMING_VALUES*          fixedValues,
        IN WFP_INCOMING_CONTEXT_VALUE*         wfpContext,
        IN VOID*                               packet,
        IN WFP_FILTER*                         matchedFilter,
        OUT WFP_ACTION_TYPE*                   action,
        OUT UINT64*                            outContext
);
```

[0101] Where the following characterizes the recited parameters.

[0102] **fixedValues** includes the layer parameters sent from the requesting layer. The fixedValues is the same data provided by the requesting layer in pInFixedValues data sent as part of the Classify method 402 in the layer API 280.

[0103] wfpContext includes the context data structure 330 (FIG. 5). This data is the same as sent by the requesting layer in the pInContext sent as part of the Classify method 402 in the layer API 280.

[0104] **packet** includes the entire packet as received by the requesting layer. This data is the same as sent by the requesting layer in the pPacket sent as part of the Classify method 402 in the layer API 280.

[0105] **matchedFilter** identifies the filter requesting the callout. Typically, the matching filter is identified by the filter Id 312 of the matching filter 310 initiating the Classify method 410 of the callout API 284.

[0106] **pActionType** includes the action that is returned to the kernel firewall engine 256 from the callout 258. If the pActionType is permit or block, it is returned to the requesting layer as the pActionType returned by the Layer API 280. The callout can also return a continue action that instructs the kernel firewall engine 256 to continue applying matching filters to the packet.

[0107] **pOutContext** includes the policy context data, such as security or QOS policy data.

[0108] The Notify method 412 is used to notify a callout when a filter 310 is added or deleted to the set of installed filters 282 that identifies as one of its actions 316, the callout module 258. The notify provides the callout with an opportunity to take any required action such as allocating or de-allocating buffers that will be used by the callout 258 when it is executed by the kernel firewall engine 256. The following is an exemplary form of the Notify method.

```
typedef NTSTATUS (*WFP_CALLOUT_NOTIFY_FN)
(
        IN WFP_NOTIFY_ENUM    notify,
        IN WFP_FILTER*        filter
);
```

[0109] Where the following characterizes the recited parameters.

[0110] **notify** includes a numerical value that indicates whether the filter is being added or deleted. For example, a value of 1 indicates the filter is being added and a value of 2 indicates that the filter is being deleted.

[0111] **filter** identifies the filter being added or deleted by unique value. This may be accomplished by providing the filter Id 312 included as part of the filter 310.

[0112] The callout Registration method 414 and callout Deregistration 416 method are used to add and remove callout modules, respectively. An exemplary form the registration method 414 is as follows:

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 935

```
NTSTATUS WfpRegisterCallout
(
        IN const GUID*                  calloutId,
        IN const WFP_CALLOUT*           callout,
        IN const SECURITY_DESCRIPTOR*   sd
);
```

[0113] Where the following characterizes the recited parameters.

[0114] **callout Id** provides a unique identification for the registering callout module.

[0115] **callout** provides any callout specific information such as a driver service name, device name, and pointers to the callout classify and notify functions.

[0116] **sd** provides a security descriptor for the callout. The security descriptor provides an access control list (ACL) that specifies which processes can modify and delete the callout.

[0117] An exemplary form of the callout deregistration method 416 is as follows:

```
NTSTATUS WfpDeregisterCallout
(
        IN const GUID*      calloutId
);
```

[0118] Where the following characterizes the recited parameters.

[0119] **callout Id** is the unique Id of the callout to be removed.

[0120] The filter engine API 266 facilitates data exchange between the user mode firewall engine 260 and the kernel mode firewall 256 engine. As shown, the management API 266 includes an Add Filter method 418, a Delete Filter method 420, and an Enum Layer method 422. The methods of the filter engine API 266 are also provided as part of the management API 290.

[0121] The Add Filter 418 and Delete Filter 418 methods are used to add a new filter to the set of installed filters 282 and to delete an existing filter from the set of installed filters 282, respectively. The following is an exemplary form of the Add Filter method 418.

```
NTSTATUS
AddFilterToLayer
(
        ULONG               LayerId,
        WFP_FILTER*         pFilter
);
```

[0122] Where the following characterizes the recited parameters.

[0123] **LayerId** identifies the layer that the filter is assigned to.

[0124] **pFilter** is the filter 310 being added to the set of installed filters 282.

[0125] The following is an exemplary form of the Delete Filter method 420.

```
NTSTATUS
DeleteFilterFromLayer
(
        ULONG       LayerId,
        ULONG       FilterId
);
```


CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 936

[0126] Where the following characterizes the recited parameters.

[0127] **LayerId** identifies the layer that the filter is assigned to.

[0128] **pFilter** is the filter being deleted from the set of installed filters.

[0129] The Enum Layer method 422 provides a mechanism for the user firewall engine 260 to identify all filters matching a set of criteria. This allows the filter engine API to identify conflicting filters for filter arbitration and conflict resolution. The following is an exemplary form of the Enum Layer method 422.

```
LayerStartEnum
(
        PWFP_ENUM_TEMPLATE      pEnumTemplate,
        OUT PULONG              pMatchCount,
        OUT PWFP_ENUM_HANDLE    pEnumHandle
)
```

[0130] Wherein the following characterizes the recited parameters.

[0131] **pEnumTemplate** includes a data structure defining the filters that should be returned. Fore example, it include parameters that the filter conditions must match for the filter to be returned.

[0132] **pMatchCount** includes the number of filter matches based on the specified pEnumTemplate.

[0133] **pEnumHanIde** includes a reference to the matched filter entries.

[0134] The keying module API 288 provides an interface between keying modules and the user firewall engine 260. The keying module API 288 includes an IPSec SA Acquire method 430, an Expire Notify method 432, and an IPSec SA Acquire Complete method 434, a Keying Module Register method 436, a Keying Module Deregister method 438, an IPSec Inbound Get SPI method 440, an Add Inbound SA method 444, an Add Outbound SA method 446, an Inbound SA Expire 448 method, and a Keying Module Initiate method 350. The keying module API 288 may also include the methods previously described for the layer API, such as a Classify method 402, an Add Layer method 404 and a Delete Layer method 406.

[0135] The keying module API is used to facilitate use of known security protocols, such as defined by IPSec, used by an initiating computer and a responding computer. IPSec includes protocols such as AH and ESP. The ESP protocol is an authenticating and encrypting protocol that uses cryptographic mechanisms to provide integrity, source authentication, and confidentiality of data. The AH protocol is an authentication protocol that uses a hash signature in the packet header to validate the integrity of the packet data and authenticity of the sender.

[0136] An IKE protocol provides a method for an initiating computer and a responding computer to negotiate security settings used with the AH and ESP protocols. The negotiated security settings form a data structure called a security association (SA). The SA defines parameters such as an authentication algorithm, encryption algorithm, keys, and the lifetime of keys, used by ESP or AH to protect the contents of an IP packet. Because ESP and AH require an established SA, an IKE negotiation is executed before the ESP or AH protocols are used by the initiating and responding computer. A given SA is identified by a value known as a Security Parameter Index (SPI).

[0137] Each of the initiating and responding computers include an IPSec driver that determines, based on an IPSec policy, whether data sent between the initiating and responding computers requires encryption or authentication. The IPSec policy is a set of filters that defines how the network device uses IPSec and includes filter lists, authentication methods, and other information In an embodiment of the invention, the IPSec policy is defined by filters included in the set of installed filters in the kernel firewall engine or the filter module 294. The policy that applies to a packet is identified by invoking the keying module layer which uses the Classify method to identify matching filters.

[0138] The IPSec SA Acquire method 430 is called by the user firewall engine 260 (via a client proxy) to pass a driver acquire or an external initiate request a user mode layer 282 implemented as a keying module layer. The keying module layer returns this call and does the negotiation asynchronously. Once the keying module layer has completed the negotiation, the keying module layer calls the IPSec SA Acquire Complete method 343 to notify the user policy engine 260 that the negotiation is complete. The following is an exemplary form of the IPSec SA Acquire method.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 937

```
typedef WIN32_ERR (*PROCESS_IPSEC_SA_ACQUIRE0)
(
        IN FWP_IPSEC_ACQUIRE_CONTEXT0   ipsecContext,
        IN const FWP_IPSEC_SA_ACQUIRE0* acquire,
        IN FWP_IPSEC_SPI                inboundSAspi
);
```

[0139] Wherein the following characterizes the recited parameters.

[0140] **ipsecContext** is a handle to link the acquire with the SA being added.

[0141] **acquire** includes the necessary information for negotiating the SA according to known protocols such as IKE.

[0142] **inboundSAspi** includes a SPI that is used for an inbound SA.

[0143] The Expire Notify method 432 is called to pass an expire-notify to the keying module layer that added the inbound SA. The following is an exemplary form the Expire Notify method.

```
typedef VOID (*PROCESS_IPSEC_SA_EXPIRE0)
(
        IN const FWP_IPSEC_SA_EXPIRE_NOTIFY0* expireNotify
);
```

[0144] Wherein the following characterizes the recited parameters.

[0145] **expireNotify** contains information identifying the expiring SA. For example, in the case of an outbound SA, the SPI is provided.

[0146] The IPSec SA Acquire Complete 434 method is called by a keying module layer to close the user firewall engine's context after it has finished negotiation and added all the SAs, or after it has come across an error. After this method is executed, the keying module layer does not reuse the ipsecContext for any other API method. The following is an exemplary form the IPSec SA Acquire Complete method.

```
WIN32_ERR
FwpIPSecSAAcquireComplete0
(
        IN FWPM_ENGINE_HANDLE                   engineHandle,
        IN FWP_IPSEC_ACQUIRE_CONTEXT0           ipsecContext,
        IN const FWP_IPSEC_NEGOTIATION_STATUS0* status
);
```

[0147] Wherein the following characterizes the recited parameters.

[0148] **engincHandle** provides a handle to the user firewall engine 260.

[0149] **ipsecContext** is the context passed by user firewall engine with the IPSec Acquire method.

[0150] **status** provides status and other details of the SA negotiation. The status is returned by the user firewall engine 260 if the acquire was externally initiated via FwpKeyingModuleInitiate0.

[0151] The Keying Module Register method 436 is called by a keying module layer to register with the user firewall engine 260 and pass its function pointers. The following is an exemplary form of the Keying Module Register method.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 938

```
WIN32_ERR
FwpKeyingModuleRegister0
(
        IN FWPM_ENGINE_HANDLE                   engineHandle,
        IN const GUID*                          keyingModuleID,
        IN const FWP_KEYING_MODULE_INFO0*       keymodInfo
);
```

[0152] Wherein the following characterizes the recited parameters.

[0153] **engineHandle** provides the handle to the user firewall engine 260.

[0154] **keyingModuleID** is a unique ID for the keying module layer.

[0155] **keymodInfo** includes registration information about the keying module layer such as pointers to process IPSec SA Acquire and process IPSec SA Expire functions.

[0156] The Keying Module Deregister method 438 is called by the keying module layer to deregister the keying module from the user firewall engine 260. The following is an exemplary form of the Keying Module Deregister method.

```
WIN32_ERR
FwpKeyingModuleDeregister0
(
        IN FWPM_ENGINE_HANDLE       engineHandle,
        IN const GUID*              keyingModuleID
);
```

[0157] Wherein the following characterizes the recited parameters.

[0158] **engineHandle** is the handle to the user firewall engine 260.

[0159] **keyingModuleID** is the unique ID of the keying module layer.

[0160] The IPSec Inbound Get SPI method 440 is called by the keying module layer to obtain the SPI for a new inbound SA. The IPSec Inbound Get SPI method 440 is usually used when the keying module layer executes in a responding network device. The following is an exemplary form the IPSec Inbound Get SPI method.

```
WIN32_ERR
FwpIPSecSAInboundGetSpi0
(
        IN FWPM_ENGINE_HANDLE               engineHandle,
        IN const FWP_IPSEC_TRAFFIC0*        ipsecTrafficDescription,
        IN const FWP_IPSEC_UDP_ENCAP0*      udpEncapInfo,
        OUT FWP_IPSEC_SPI*                  inboundSpi
);
```

[0161] Wherein the following characterizes the recited parameters.

[0162] **engineHandle** is the handle to the user firewall engine 260.

[0163] **ipsecTrafficDescription** is a 5-tuple description for creating an inbound larval SA. The 5-tuple includes source and destination IP addresses, source and destination ports, and transport layer protocol type.

[0164] **udpEncapInfo** is UDP encapsulation data for creating the larval SA. UDP encapsulation is a known method of embedding a packet formatted according to a security protocol into an unencrypted UDP packet.

[0165] **inboundSpi** is the SPI for the inbound SA.

[0166] The Add Inbound SA method 444 is called by the keying module layer to add an inbound SA, i.e. update the larval SA. The user firewall engine 260 uses the SPI in the SA to map this call to its internal state, and ioctl the SA down to the IPSec driver. The following is an exemplary form of the Add Inbound SA method.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 939

```
WIN32_ERR
FwpIPSecSAInboundAdd0
(
        IN FWPM_ENGINE_HANDLE           engineHandle,
        IN const FWP_IPSEC_SA_STRUCT0*  inboundSA
);
```

[0167] Wherein the following characterizes the recited parameters.

[0168] **engineHandle** in the handle to the user firewall engine 260.

[0169] **inboundSA** includes the inbound SA.

[0170] The Add Outbound SA 446 method is called by a keying module layer to add an outbound SA. The user firewall engine uses an inbound SPI parameter to map this call to its internal state, and ioctl the SA down to the IPSec driver.. The following is an exemplary form the Add Outbound SA method.

```
WIN32_ERR
FwpIPSecSAOutboundAdd0
(
        IN FWPM_ENGINE_HANDLE           engineHandle,

        IN FWP_IPSEC_SPI                inboundSpi,
        IN const FWP_IPSEC_SA_STRUCT0*  outboundSA
);
```

[0171] Wherein the following characterizes the recited parameters.

[0172] **engineHandle** is the handle to the user firewall engine 260.

[0173] **inboundSpi** is the SPI for the inbound SA, with which an outbound SA is paired.

[0174] **outboundSA** includes the outbound SA.

[0175] The Inbound SA Expire method 448 is called by the keying module layer to expire the inbound SA that was previously added. The following is an exemplary form the Inbound SA Expire method 448.

```
WIN32_ERR
FwpIPSecSAInboundExpire0
(
        IN FWPM_ENGINE_HANDLE           engineHandle,
        IN const FWP_IPSEC_SA_EXPIRE0*  expire
);
```

[0176] Wherein the following characterizes the recited parameters.

[0177] **engineHandle** is the handle to the user firewall engine 260.

[0178] **expire** includes data for the SA to be expired.

[0179] The Keying Module Initiate method 450 is called by known external application like RAS, Winsock API, and the like to initiate the keying module layer and setup SAs before an application starts sending its network traffic. The user firewall engine 260 asynchronously pends the RPC call, gets the SPI from an IPSec driver, and passes the acquire to the appropriate keying module. Once the keying module layer calls FwpIPSecSAAcquireComplete0, the user firewall engine completes the asynchronous RPC with the negotiation status. The following is an exemplary form of the Keying Module Initiate method.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 940

```
WIN32_ERR
FwpKeyingModuleInitiate0
(
        IN FWPM_ENGINE_HANDLE                   engineHandle,
        IN const FWP_IPSEC_SA_ACQUIRE0*         acquire,
        IN HANDLE                               waitEvent,
        OUT FWP_IPSEC_NEGOTIATION_STATUS0*      negotiationStatus
);
```

[0180] Wherein the following characterizes the recited parameters.

[0181] **engineHandle** is the handle to the user firewall engine 260.

[0182] **acquire** includes data necessary for negotiating an SA.

[0183] **waitEvent** is a handle to an event that is triggered when negotiation status is available. If a client, i.e. calling external application, is not interested in waiting for the negotiation to complete, it can set this parameter to NULL. Internally the client proxy optionally passes this event to RPC and request it to set the event once the asynchronous RPC call completes.

[0184] **negotiationStatus** includes an outcome of the negotiation. The negotiationStatus is NULL if waitEvent is NULL. Otherwise negotiationStatus remains valid until the waitEvent is triggered.

[0185] FIG. 8 illustrates the methods used by the layers in the network stack 254 according to the present invention. The method illustrated in FIG. 8 may also be used by one or more user mode layers 282 in concert with the filter module 294 and the user firewall engine 260.

[0186] Each layer is capable of a plurality of functions including processing network packets, issuing classification requests to the kernel firewall engine 256, and managing packet context. In an embodiment of the invention, the functions are executed at each of the layers in a shim 460, 462, 464, 466 installed in the network stack 254. Alternatively, the functionality is directly built into the individual layers without the need for the shims.

[0187] The network stack 254 includes the data stream layer 268, the transport layer 270, the network layer 272, and the link layer 274. For the purposes of illustrating the present invention, the link layer 274 is implemented as a NDIS driver, the network layer 272 is implemented as an IP layer, the transport layer 270 is implemented as a TCP layer, and the data stream layer 268 is implemented as a HTTP layer. It will be understood that layers may be implemented according to any protocol. For example the transport layer also accommodates a User Datagram Protocol (UDP). The application layer supports File Transfer Protocol (FTP), Remote Procedure Call (RPC), Simple Mail Transfer Protocol (SMTP), Server Master Block (SMB) etc. As previously described, additional layers can be added to the architecture and layers may be deleted. For example, layers are added and deleted using the Add Layer and Delete Layer methods described with reference to FIG 6.

[0188] A network packet, labeled 468(a)-(d), illustrates the network packet as it traverses and is processed by the layers in the network stack 254. If the packet 468(a)-(d) is an inbound packet, it traverses the network stack from bottom to top. If the packet 468(a)-(d) is an outbound packet, it traverses the network stack from top to bottom. While such processing is well known, it is briefly described for the purpose of illustrating the present invention.

[0189] Assuming an application executing in a network device, such as a web browser, initiates a request for contents of a web page located on another network device, the application issues the request to the data stream layer 268. In the example, the data stream 264 formats the request according to the HTTP protocol and sends the request to the transport layer in packet 468(a). The transport layer 270 receives packet 468(a). The transport layer 270, which implements the TCP protocol, places the data in one or more packets, and each packet is provided with a TCP header. The TCP header includes information such as the source and destination ports, protocol type, i.e. TCP, sequence numbers, flags, and checksums. The transport layer then sends the packet, labeled 468(b) to the network layer.

[0190] The network layer implements the IP protocol and encapsulates the data within an IP header that includes the source and destination IP addresses, flags, checksums and other known information. The IP header also indicates whether the packet is fragmented. A packet is fragmented when the size of the IP packet exceeds a Maximum Transmission Unit (MTU) size for the network technology used to transmit the packet. For example, Ethernet technology specifies that the MTU is 1500 bytes. If the IP packet length exceeds the MTU, it is fragmented into two or more IP packets, each with its own IP header, all of equal or smaller length as compared to the MTU.

[0191] In an embodiment of the invention, the network layer is divided into first and second layers. The first layer, referred to as a fragment layer, processes IP packet fragments. The second layer, referred to as a fully assembled layer, processes full IP packets, e.g. before outbound IP packets are fragmented and after inbound IP packets are reassembled into a single IP packet. Following network layer processing and possible fragmentation, the packet 468

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 941

(c) is sent to the link layer 274. The link layer 274 further packetizes the data by providing a MAC header, with the source and destination MAC addresses, as well as other information. The packet is then sent to a network interface card (NIC) where it is physically transmitted onto the network.

[0192] Inbound packets are processed in a reciprocal fashion. Packet 468(d) is received by the NIC and sent to the link layer 274. The MAC header is removed and the packet 468(c) is sent to the network layer where the IP packet fragments are reassembled, if necessary, and the IP header is parsed. The network layer then sends the packet 468 (b) to the transport layer where the TCP header is removed and if the data stream was sent in a plurality of TCP packets, the data stream is reassembled. Finally, the data stream 468(a) is sent to the data stream layer 268 where the data is deciphered by the application protocol, in this case the HTTP protocol.

[0193] For each outbound packet, the layers of the network stack maintain packet context 470(a)-(c). For each inbound packet, the layers of the network stack maintain packet context 472(a)-(c). The packet context follows each packet as it traverses the network layers. Packet context is also passed to, and may be modified by, the callouts 258 (FIG. 1).

[0194] The packet context is updated as the packets are processed at each layer. Each layer adds its layer parameters to the packet context thereby providing this information to subsequent layers or processes. As shown, the link layer 274 adds source and destination MAC addresses and an interface number for an inbound packet as illustrated by context 472(a) That context is received by the network layer 272, which adds source and destination IP addresses as illustrated by context 472(b). The transport layer 266 receives the context and adds the port numbers as illustrated by context 472(a).

[0195] A similar process occurs for the context 470(a)-(c) associated with outbound packets. The data stream layer 268 adds information such as a URL address from the packet payload as illustrated by context 470(a), the transport layer 270 further adds source and destination port numbers as illustrated by context 470(b) and the network layer adds source and destination IP addresses as illustrated by context 470(c).

[0196] It will be understand that each layer can add any context information available to that layer. Typically, this includes any information that the layer is designed to process, i.e. add or parse or otherwise derive from packets.

[0197] When the packet and its corresponding context are received at each layer, the layer functions as a requesting layer by identifying the layer parameters and sending a classification request, labeled 474. The classification request 474 includes, the layer parameters 476, the packet context 478 received from the prior layer, and the full packet 480. An exemplary method used to issue the classification request is the Classify method 402 described with reference to the layer API 280 (Fig.7).

[0198] In response to each classify request, the kernel mode firewall engine 256 compares the layer parameters 476 and packet context 478 to the filter conditions 318 (FIG. 4) of the filters assigned to the requesting layer. The kernel firewall engine 256 sends a response, labeled 482, with the action 484 from the matching filters 310 with the highest weight 314 to the requesting layer. The kernel firewall engine 256 also returns the policy context 482. If no matching filters are identified by the kernel firewall engine 256, the kernel firewall engine notifies the requesting layer that no matching filters exist. The kernel firewall engine 256 continues to identify matching filters until a matching filter designates a terminating action, i.e. permit or block, or until all filters assigned to the requesting layer are checked, whichever comes first. Alternatively, the kernel firewall engine 256 identifies all matches and returns the actions to the requesting layer in a single response.

[0199] As generally illustrated, identifying layer parameters is done as part of the normal layer processing performed by the layers of the network stack 254. Additional packet parsing is not required thereby minimizing impact on system performance. Moreover, because the layers cooperate in maintaining packet context, it is possible to for the firewall engine 256 to compare filter conditions with packet parameters at layers that normally do not have access to those packet parameters. For example, the network layer 268 receives context for inbound packets from the link layer 274 that includes source and destination MAC addresses. Because the network layer 272 issues classification requests with network layer parameters, e.g. source and destination IP addresses, as well as the packet context, the kernel firewall engine 256 can filter on both IP and MAC addresses at the network layer 272 even though the MAC addresses are usually not available.

[0200] Referring to Fig. 9 and exemplary set of callout modules 258 included with the firewall architecture of the present invention will now be described. The callout modules 258 include HTTP context callout 500, intrusion detection callout 502, IPSec callout 504, and logging callout 506. The callout API 284 forms the interface between the kernel firewall engine 256 and the callouts 258.

[0201] The HTTP context callout 500 maintains a cache 508 of acceptable, or alternatively, unacceptable URL addresses. The HTTP context callout 508 periodically accesses a server 510 interfaced to the public network that maintains URL addresses and classifies them as acceptable or unacceptable. When the kernel firewall engine 256 executes the HTTP context callout, the callout examines the packet, deciphers the URL address if necessary, and determines if it is acceptable based on information in the cache 508. The HTTP callout then returns permit as the action 316 if the URL address is acceptable and block if the URL address is unacceptable to the kernel mode firewall engine 256, which

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 942

in turn returns the action 316 to the requesting layer via layer API 280. The HTTP context callout is useful in implementing a parental control function.

[0202] The intrusion detection callout 502 uses available algorithms and techniques to examine the packet to identify indicia of viruses or otherwise suspicious packets. If a suspicious packet is detected, an action 316 of block is returned. An example of a suspicious packet is a packet where all flags in IP and TCP headers are set to a value of one. The packet is suspicious because it is never valid and could indicate an attack signature. Alternatively, the intrusion detection callout 502 modifies packet context to flag the suspicious nature of the packet thereby deferring a decision on whether to block the packet to subsequent layers in the network stack.

[0203] The IPSec callout 504 is designed to determine to whether a proper security protocol was applied to the packet. The IPSec callout 504 communicates with an IPSec process and determines based on IPSec policy whether the packet was supposed to be subject to IPSec processing. If so, the IPSec callout 504 verifies, based on packet context, whether the packet was in fact subject to the IPSec processing. If the packet was supposed to be processed according to IPSec, but was not, e.g. the packet was in clear text, an action of block is returned. If the packet was subject to IPSec processing, the IPSec callout verifies that a proper Security Association (SA) was applied. If the proper SA was not applied, the callout returns an action 316 of block, otherwise an action 316 of permit is returned. Alternatively, the IPSec callout modifies the packet context deferring a decision on whether to block or permit the permit to a subsequent layer.

[0204] The logging callout 506 is used to save information pertaining to the packet, such as the full packet for later use. Such later use might be, for example, diagnostics for certain network traffic that is unexpectedly not traversing the network because it is being blocked or for tracking malicious attacks on the system.

[0205] All of the references cited herein, including are hereby incorporated in their entireties by reference.

[0206] In view of the many possible embodiments to which the principles of this invention may be applied, it should be recognized that the embodiment described herein with respect to the drawing figures is meant to be illustrative only and should not be taken as limiting the scope of invention. For example, those of skill in the art will recognize that the elements of the illustrated embodiment shown in software may be implemented in hardware and vice versa or that the illustrated embodiment can be modified in arrangement and detail without departing from the spirit of the invention. Therefore, the invention as described herein contemplates all such embodiments as may come within the scope of the following claims and equivalents thereof.

Claims

1. A firewall framework implemented within a computer system for providing multi-layering filtering of packet, comprising:

   a set of layer processors, each layer processor being capable of processing layer parameters for the packet associated with the layer processor and each layer processor being further capable of issuing a classification request that includes the layer parameters; and
   a first firewall engine including:

   a layer interface for receiving first layer parameters from a requesting layer processor and for returning an action to the requesting layer, the requesting layer processor being one of the set of layer processors;
   a set of installed filters; and
   a lookup component for identifying at least one matching filter from the set of installed filters and identifying from the matching filter the action to be returned by the layer interface.

2. The firewall framework of claim 1, wherein the requesting layer receives packet context including second layer parameters from a previous layer processor and the classification request includes the packet context, wherein the lookup component further uses the packet context for identifying the at least one matching filter.

3. The firewall framework of claim 2, wherein the firewall engine further comprises:

   a callout interface for sending the first layer parameters, the packet context, and the packet to a callout, wherein the callout analyzes the packet and returns the action to the firewall engine.

4. The firewall framework of claim 2, wherein the requesting layer passes the packet context to a next layer.

5. The firewall framework of claim 1, further comprising:

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 943

a policy provider for establishing a new filter; and
a second firewall engine for adding the new filter to the set of installed filers.

6. The firewall framework of claim 1, wherein each of the installed filters comprises:

a filter condition including a type-data pair, the type defining the size of the data and the data including a filter parameter, and
the action.

7. The firewall framework of claim 6, wherein the filter parameter includes a range of values.

8. The firewall framework of claim 6, wherein the filter comprises a weight factor defining a priority of the filter.

9. The firewall framework of claim 1, wherein the set of layer processors further comprises:

a link layer having layer parameters including an interface number and source and destination MAC addresses;
an network layer having layer parameters including a source and destination IP address;
a transport layer having layer parameters including source and destination ports; and
an application layer having layer parameters including a stream of data.

10. The firewall framework of claim 9, wherein the network layer is further divided into a fragment layer and a fully assembled layer, the fragment layer processing IP packet fragments and the fully assembled layer processing complete IP packets, each of the fragment layer and the fully assembled layer issuing the classify request.

11. The firewall framework of claim 1, wherein the first firewall engine executes in an operating system kernel mode.

12. The firewall framework of claim 1, wherein the first firewall engine executes in an operating system user mode.

13. A method of communicating between a first layer process and a firewall process in an operating system, comprising the steps of:

issuing, by the first layer process, a classify call having a plurality of parameters comprising, a protocol packet, at least one layer parameter. and a packet context from a second layer process;
receiving, by the firewall process, the classify call; and
issuing, by the firewall process, an action identified from a filter matching the at least one layer parameter.

14. The method of claim 13, wherein the first layer process is a new process, further comprising:

issuing by the first layer process, an add layer call.

15. The method of claim 13, wherein the first layer process is an existing process, further comprising:

issuing by the first layer process, a delete layer call.

16. A method of communicating between a firewall process and a callout process in an operating system, comprising the steps of:

issuing, by the firewall process, a classify call having a plurality of parameters comprising, a protocol packet, at least one layer parameter, a packet context, and a matching filter identification;
receiving, by the callout process, the classify call; and
issuing, by the callout process, an action identified from the plurality of parameters in the classify call.

17. The method of claim 16, wherein the callout process is a parental control module maintaining a cache of acceptable resource location, further comprising:

examining the packet to identify a resource location and comparing the resource location to the acceptable locations.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 944

18. The method of claim 17, wherein the action issued by the callout process is block and the protocol packet is prevented from network traversal.

19. The method of claim 17, wherein the action issued by the callout process is permit and the protocol packet is permitted to further traverse the network.

20. The method of claim 16, wherein the callout process is a logging module, further comprising storing in a memory the protocol packet.

21. The method of claim 16, wherein the callout process is a security module, further comprising:

determining, by the security module, that the protocol packet is required to conform to a security protocol;
verifying that the protocol packet conforms to the security protocol.

22. The method of claim 21, wherein the security protocol is an IPSec security protocol.

23. A computer-readable medium for executing computer-readable instructions for facilitating a firewall framework implemented within a computer system for providing multi-layering filtering of a packet, comprising:

a set of layer processors, each layer processor being capable of processing layer parameters for the packet associated with the layer processor and each layer processor being further capable of issuing a classification request that includes the layer parameters; and
a first firewall engine including:

a layer interface for receiving first layer parameters from a requesting layer processor and for returning an action to the requesting layer, the requesting layer processor being one of the set of layer processors;
a set of installed filters; and
a lookup component for identifying at least one matching filter from the set of installed filters and identifying from the matching filter the action to be returned by the layer interface.

24. The computer-readable medium of claim 23, wherein the requesting layer receives packet context including second layer parameters from a previous layer processor and the classification request includes the packet context, wherein the lookup component further uses the packet context for identifying the at least one matching filter.

25. The computer-readable medium of claim 24, wherein the firewall engine further comprises:

a callout interface for sending the first layer parameters, the packet context, and the packet to a callout, wherein the callout analyzes the packet and returns the action to the firewall engine.

26. The computer-readable medium of claim 24, wherein the requesting layer passes the packet context to a next layer.

27. The computer-readable medium of claim 23, further comprising:

a policy provider for establishing a new filter; and
a second firewall engine for adding the new filter to the set of installed filers.

28. A computer-readable medium for executing computer-readable instructions for communicating between a first layer process and a firewall process in an operating system, comprising the steps of:

issuing, by the first layer process, a classify call having a plurality of parameters comprising, a protocol packet, at least one layer parameter, and a packet context from a second layer process;
receiving, by the firewall process, the classify call; and
issuing, by the firewall process, an action identified from a filter matching the at least one layer parameter.

29. The computer-readable medium of claim 28, wherein the first layer process is a new process, further comprising:

issuing by the first layer process, an add layer call.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 945

30. The computer-readable medium of claim 28, wherein the first layer process is an existing process, further comprising:

issuing by the first layer process, a delete layer call.

31. A computer-readable medium for executing computer-executable instructions for communicating between a firewall process and a callout process in an operating system, comprising the steps of:

issuing, by the firewall process, a classify call having a plurality of parameters comprising, a protocol packet, at least one layer parameter, a packet context, and a matching filter identification;
receiving, by the callout process, the classify call; and
issuing, by the callout process, an action identified from the plurality of parameters in the classify call.

32. The computer-readable medium of claim 31, wherein the callout process is a logging module, further comprising storing in a memory the protocol packet.

33. The computer-readable medium of claim 31, wherein the callout process is an security module, further comprising:

determining, by the security module, that the protocol packet is required to conform to a security protocol;
verifying that the protocol packet conforms to the security protocol.

34. A functional interface for allowing a requesting layer to obtain policy for a packet, the requesting layer being one of a plurality layers, comprising:

a classify method comprising:

the packet received by the requesting layer;
a set of parameters associated with the packet, the set of parameters including data processed by the requesting layer,
a packet context received by the requesting layer from another layer of the plurality of layers; and
an action to be returned to the requesting layer identifying a first policy to be applied to the packet.

35. The functional interface of claim 34, further comprising:

an add layer method for allowing a new layer to be added to the plurality of layers.

36. The functional interface of claim 34, further comprising:

a delete layer method for removing an existing layer from the plurality of layers.

37. The functional interface of claim 34, further comprising:

a policy context defining a second policy to be applied to the packet.

38. The functional interface of claim 37, wherein the first policy is a firewall policy and the second policy is a non-firewall policy.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 946

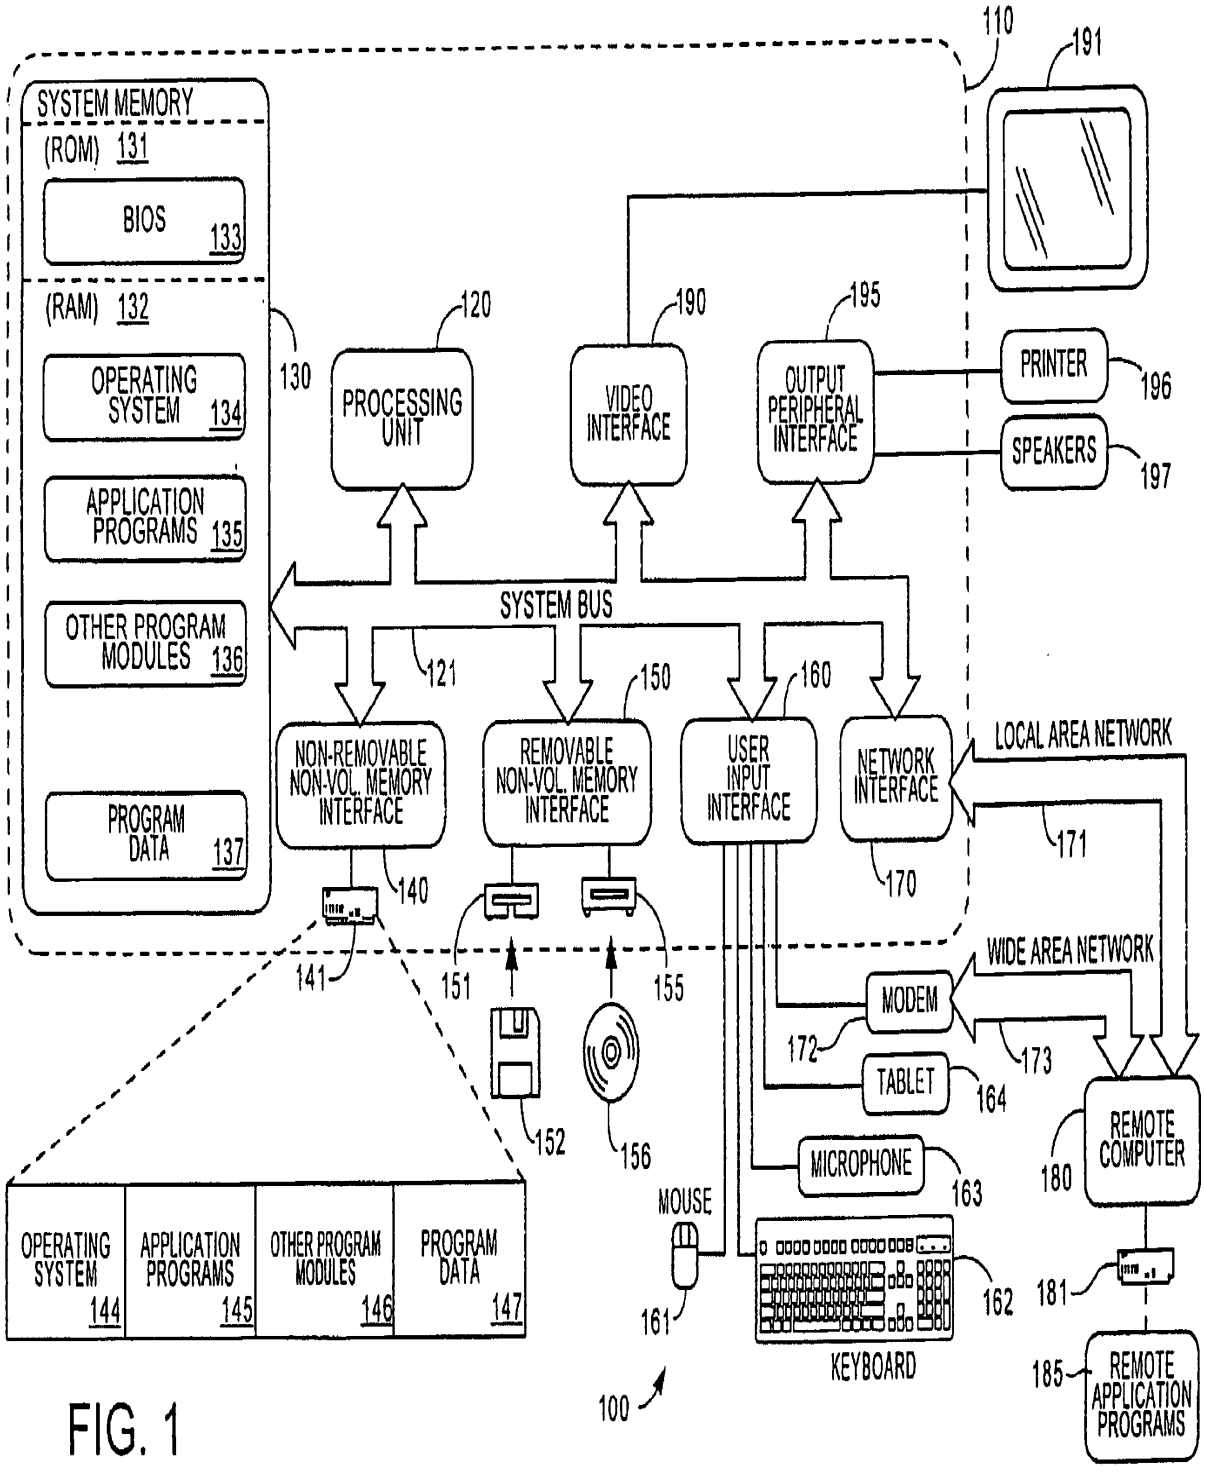
110
191
SYSTEM MEMORY
(ROM) 131
BIOS 133
(RAM) 132
OPERATING SYSTEM 134
APPLICATION PROGRAMS 135
OTHER PROGRAM MODULES 136
PROGRAM DATA 137
130
120
PROCESSING UNIT
190
VIDEO INTERFACE
195
OUTPUT PERIPHERAL INTERFACE
PRINTER
196
SPEAKERS
197
SYSTEM BUS
121
150
160
NON-REMOVABLE NON-VOL. MEMORY INTERFACE
REMOVABLE NON-VOL. MEMORY INTERFACE
USER INPUT INTERFACE
NETWORK INTERFACE
LOCAL AREA NETWORK
171
170
140
141
151
155
WIDE AREA NETWORK
MODEM
172
173
TABLET
164
REMOTE COMPUTER
180
MICROPHONE
163
152
156
MOUSE
161
OPERATING SYSTEM 144
APPLICATION PROGRAMS 145
OTHER PROGRAM MODULES 146
PROGRAM DATA 147
162
KEYBOARD
181
185
REMOTE APPLICATION PROGRAMS
100

FIG. 1

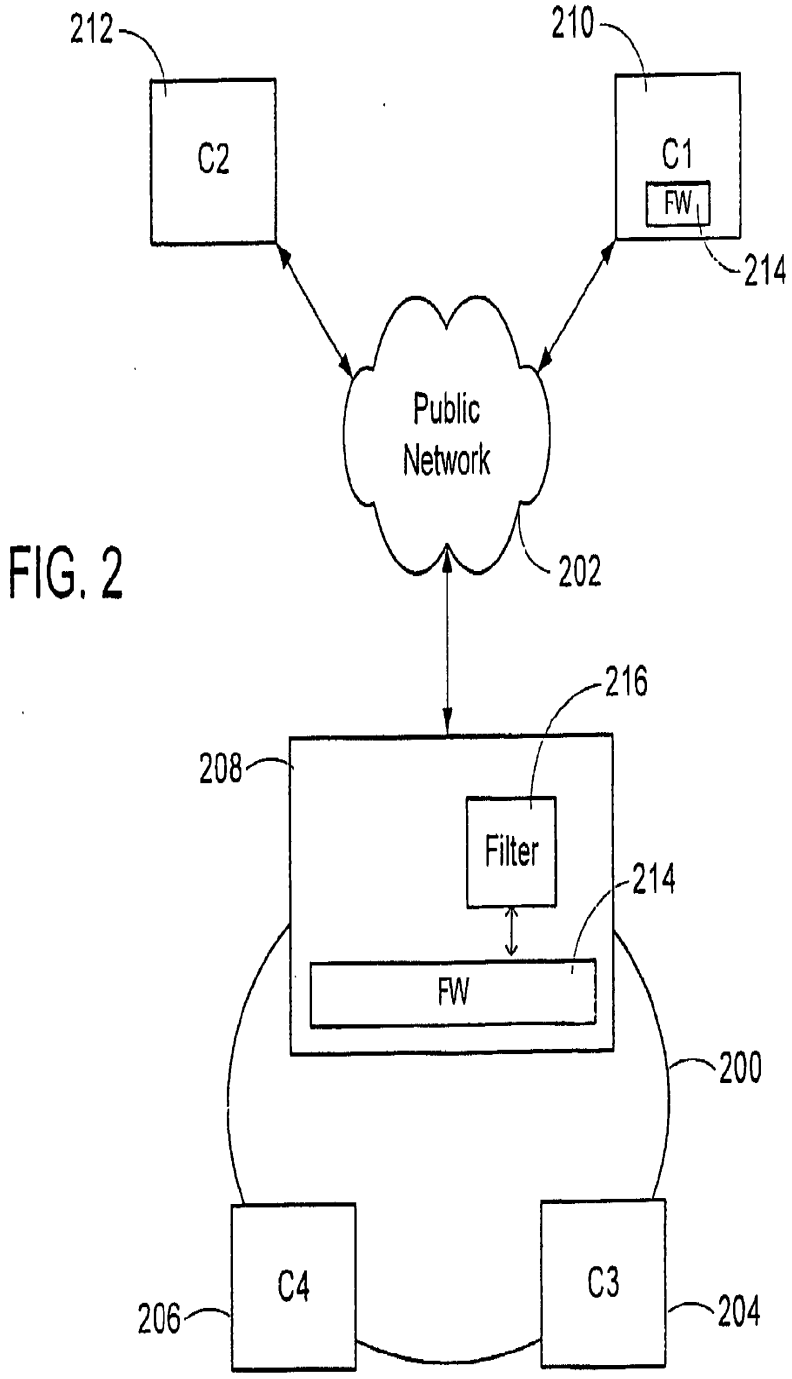

FIG. 2

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 948

FIG. 3

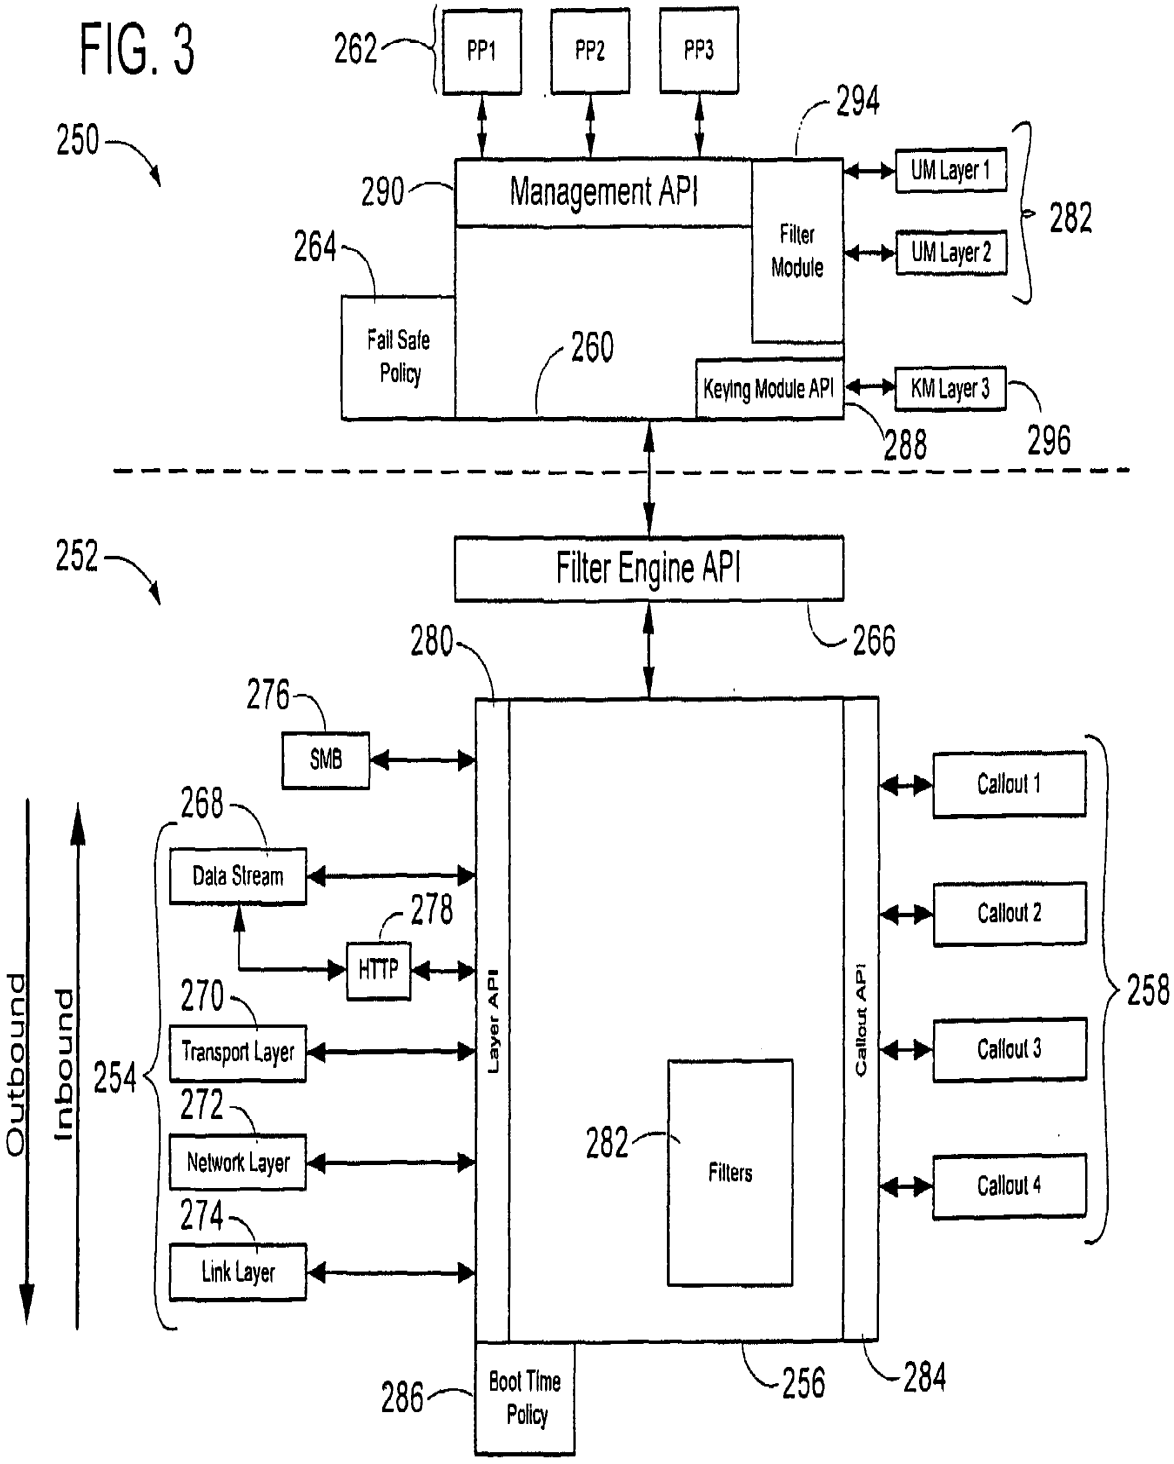

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 949

FIG. 4

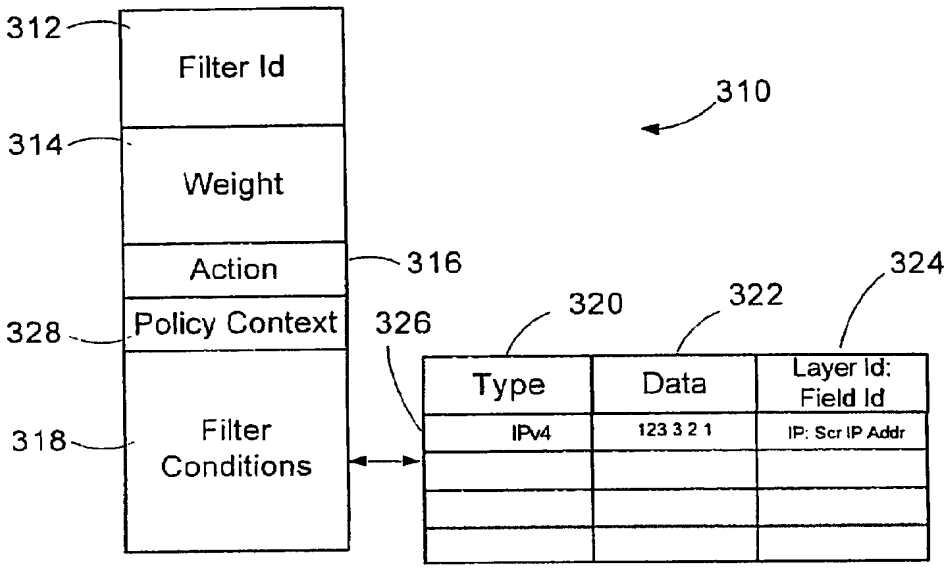

FIG. 5

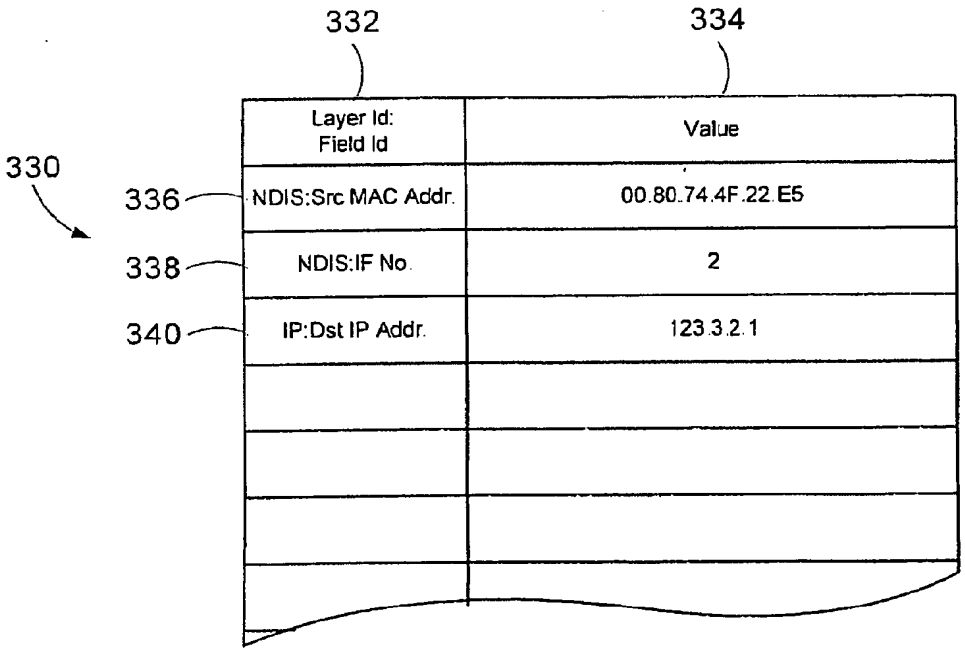

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 950

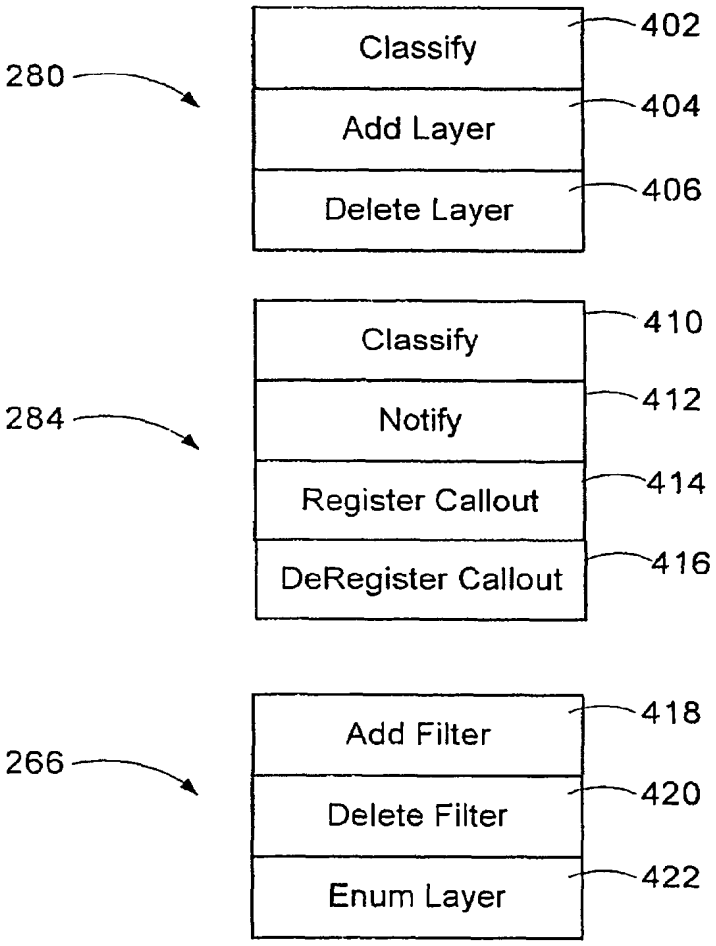

FIG. 6

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 951

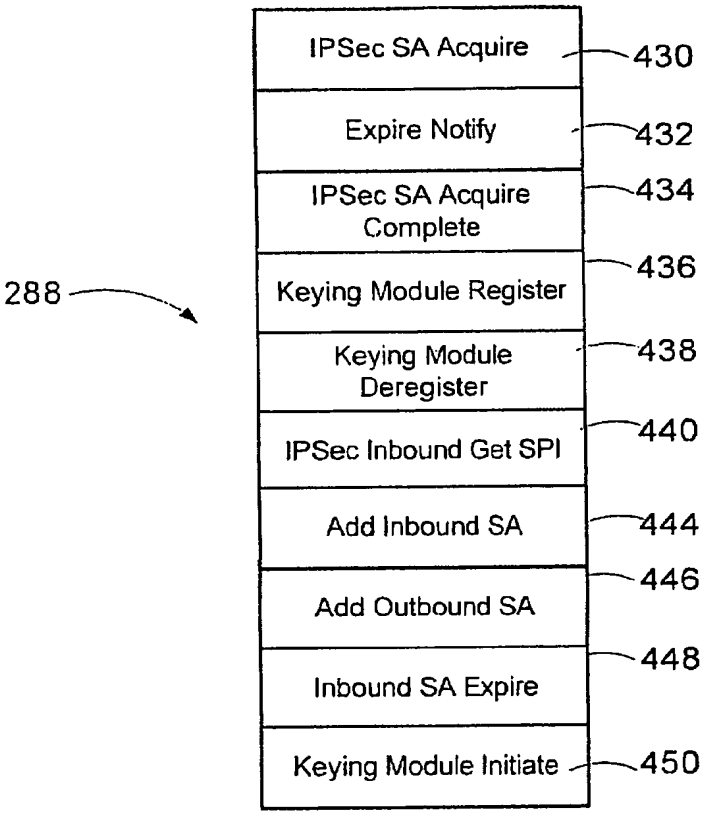

FIG. 7

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 952

FIG. 8

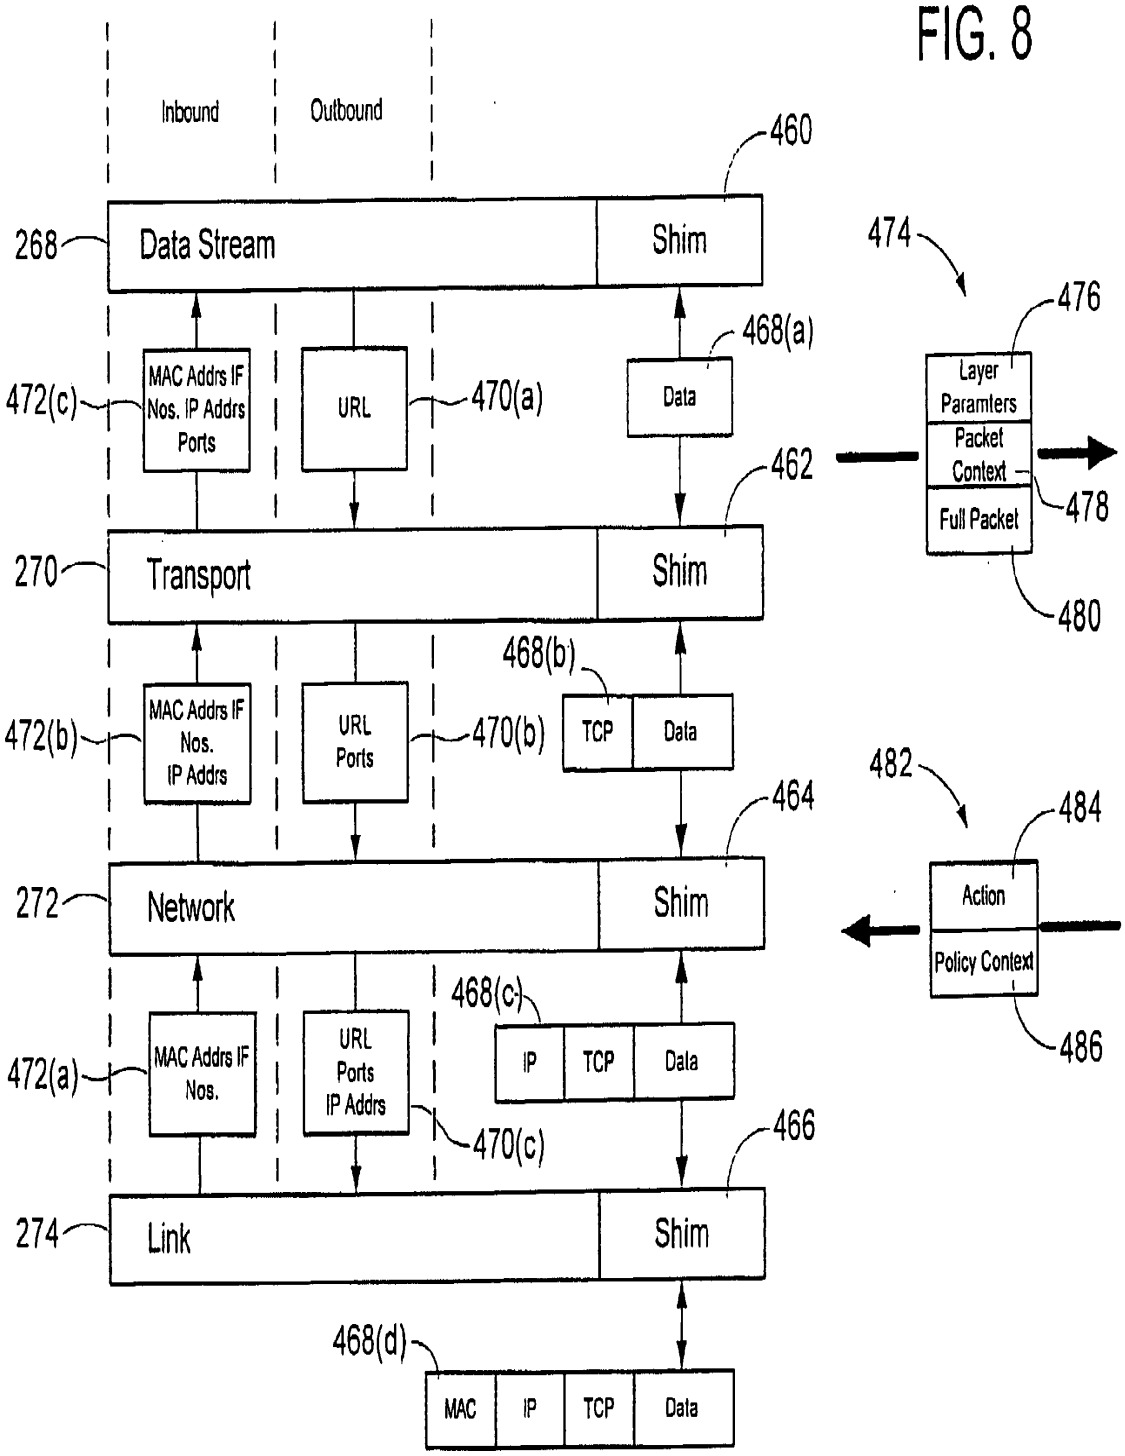
Inbound
Outbound
460
268
Data Stream
Shim
474
476
468(a)
472(c)
MAC Addrs IF Nos. IP Addrs Ports
URL
470(a)
Data
Layer Paramters
Packet Context
Full Packet
478
462
480
270
Transport
Shim
468(b)
472(b)
MAC Addrs IF Nos. IP Addrs
URL Ports
470(b)
TCP
Data
482
484
464
272
Network
Shim
Action
Policy Context
468(c)
486
472(a)
MAC Addrs IF Nos.
URL Ports IP Addrs
470(c)
IP
TCP
Data
466
274
Link
Shim
468(d)
MAC
IP
TCP
Data

FIG. 9

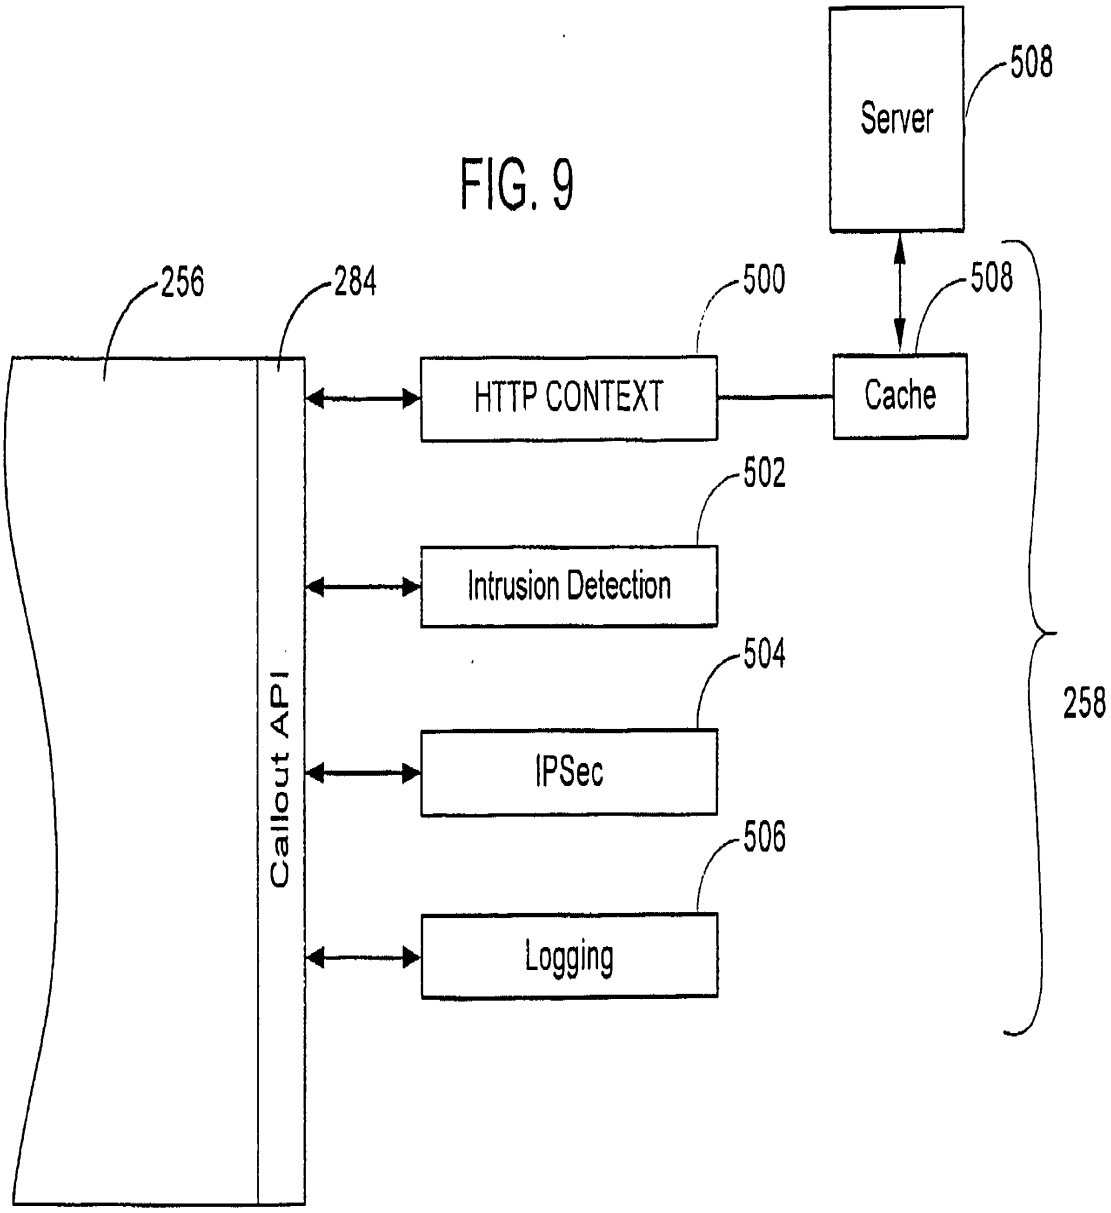

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 954

Australian Government

IP Australia

LETTERS PATENT

STANDARD PATENT

2006230171

I, Robyn Foster, the Commissioner of Patents, grant a Standard Patent with the following particulars:

Name and Address of Patentee(s):

Wake Forest University
Office of Tachnology Asset Management, Wake Forest University Health Sciences, 200 East First Street, Suite 101, Winston Salem, North Carolina, 27101, United States of America

Name of Actual Inventor(s):

Fulp, Errin W. and Tarsa, Stephen J..

Title of Invention:

Methods, systems, and computer program products for network firewall policy optimization

Term of Letters Patent:

Twenty years from 28 March 2006

Priority Details :

Number	*Date*	*Filed with*
60/665,664	28 March 2005	US

COMMONWEALTH OF AUSTRALIA
AUSTRALIA
PATENT OFFICE

Dated this 4th day of October 2012

R E Foster

Robyn Foster
Commissioner of Patents

PATENTS ACT 1990

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2006230171 B2**

(54) Title
Methods, systems, and computer program products for network firewall policy optimization

(51) International Patent Classification(s)
H04L 9/00 (2006.01)
G06F 9/00 (2006.01)
G06F 15/16 (2006.01)
G06F 15/18 (2006.01)
G06F 17/00 (2006.01)

(21) Application No: **2006230171**

(22) Date of Filing: **2006.03.28**

(87) WIPO No: **WO06/105093**

(30) Priority Data

(31) Number	(32) Date	(33) Country
60/665,664	**2005.03.28**	**US**

(43) Publication Date: **2006.10.05**
(44) Accepted Journal Date: **2012.06.21**

(71) Applicant(s)
Wake Forest University

(72) Inventor(s)
Tarsa, Stephen J.;Fulp, Errin W.

(74) Agent / Attorney
Griffith Hack, GPO Box 3125, Brisbane, QLD, 4001

(56) Related Art
WO 2006/093557 A2
US 2005/0125697 A1
EP 1006701 A2

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
5 October 2006 (05.10.2006)

PCT

(10) International Publication Number
WO 2006/105093 A2

(51) International Patent Classification:
G06F 15/16 (2006.01)

(21) International Application Number:
PCT/US2006/011291

(22) International Filing Date: 28 March 2006 (28.03.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/665,664 28 March 2005 (28.03.2005) US

(71) Applicant *(for all designated States except US)*: **WAKE FOREST UNIVERSITY** [US/US]; OFFICE OF TECHNOLOGY ASSET MANAGEMENT, WAKE FOREST UNIVERSITY HEALTH SCIECNES, 200 East First Street, Suite 101, Winston Salem, North Carolina 27101 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **FULP, Errin, W.** [US/US]; 308 N. Hawthorne Road, Winston-Salem, North Carolina 27104 (US). **TARSA, Stephen, J.** [US/US]; 135 King Phillips Path, Duxbury, Massachusetts 02332-3503 (US).

(74) Agent: **HUNT, Gregory, A.**; JENKINS, WILSON, TAYLOR & HUNT, P.A., Suite 1200 University Tower, 3100 Tower Boulevard, Durham, North Carolina 27707 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),

[Continued on next page]

(54) Title: METHODS, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR NETWORK FIREWALL POLICY OPTIMIZATION

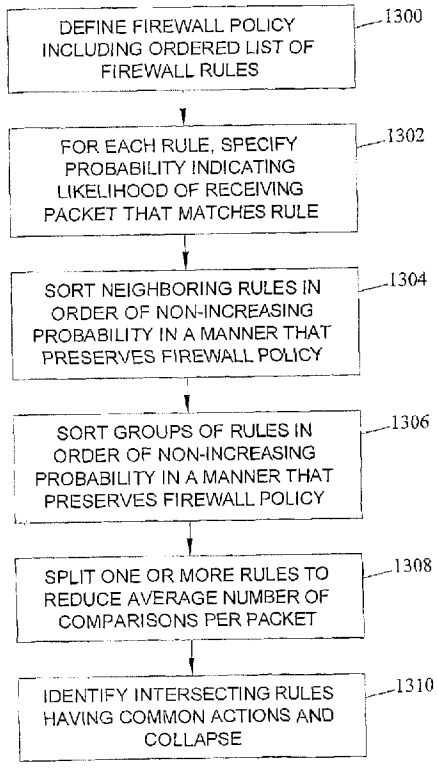

(57) Abstract: Methods, systems, and computer program products for firewall policy optimization are disclosed. According to one method, a firewall policy including an ordered list of firewall rules is defined. For each rule, a probability indicating a likelihood of receiving a packet matching the rule is determined. The rules are sorted in order of non-increasing probability in a manner that preserves the firewall policy.

WO 2006/105093 A2

European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

2006230171 06 Jun 2012

TITLE

METHODS, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR NETWORK FIREWALL POLICY OPTIMIZATION

RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Patent Application Serial No. 60/665,664, filed March 28, 2005; the disclosure of which is incorporated herein by reference in its entirety.

GOVERNMENT INTEREST

This invention was made with Government support under Grant No. DE-FG02-03ER25581 awarded by U.S. Department of Energy, Mathematical and Computational Information Sciences Division. The U.S. Government has certain rights in the invention.

TECHNICAL FIELD

The subject matter described herein relates to optimizing polices for network firewalls. More particularly, the subject matter described herein relates to methods, systems, and computer program products for optimizing performance of network firewalls.

BACKGROUND ART

Network firewalls remain the forefront defense for most computer systems. Guided by a security policy, these devices provide access control, auditing, and traffic control [3, 30, 31]. A security policy is a list of ordered rules, as seen in Table 1, that defines the action to perform on matching packets. For example, an accept action passes the packet into or from the secure network, while deny causes the packet to be discarded. In many implementations, the rule set is stored internally as a linked list. A packet is sequentially compared to the rules, starting with the first, until a match is found; otherwise, a default action is performed [30, 31]. This is referred to as a first-match policy and is used in many firewall systems including the Linux firewall implementation iptables [25].

CISCO EXHIBIT 1002
COMS ID No: ARCS-372276 Received by IP Australia: Time (H:m) 11:49 Date (Y-M-D) 2012-06-06 Filed: U.S. Patent No. 9565213
Page 959

No.	Proto.	Source IP	Source Port	Destination IP	Destination Port	Action	Prob.
1	UDP	1.1.*	*	*	80	deny	0.01
2	TCP	2.*	*	1.*	90	accept	0.02
3	UDP	*	*	1.*	*	accept	0.10
4	TCP	2.*	*	1.*	20	accept	0.17
5	UDP	1.*	*	*	*	accept	0.20
6	*	*	*	*	*	deny	0.50

Table 1: Example Security Policy Consisting of Multiple Ordered Rules

Traditional firewall implementations consist of a single, dedicated machine, similar to a router, that sequentially applies rules to each arriving packet. However, packet filtering represents a significantly higher processing load than routing decisions [24, 29, 31]. For example, a firewall that interconnects two 100 Mbps networks would have to process over 300,000 packets per second [30]. Successfully handling these traffic loads becomes more difficult as rule sets become more complex [4, 22, 31]. Furthermore, firewalls must be capable of processing even more packets as interface speeds increase. In a high-speed environment (e.g. Gigabit Ethernet), a single firewall can easily become a bottleneck and is susceptible to DoS attacks [4, 9, 13, 14]. An attacker could simply inundate the firewall with traffic, delaying or preventing legitimate packets from being processed.

One approach to increase firewall performance focuses on improving hardware design. Current research is investigating different distributed firewall designs to reduce processing delay [4, 9, 22], and possibly provide service differentiation [11]. Another approach focuses on improving performance via better firewall software [6, 7, 12, 16, 17, 24]. Similar to approaches that address the longest matching prefix problem for packet classification [8, 10, 25, 28], solutions typically represent the firewall rule set in different fashions (e.g. tree structures) to improve performance. While both approaches, together or separate, show great promise, each requires radical changes to the firewall system, and therefore are not amenable to current or legacy systems.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 960

Accordingly, there exists a need for improved methods, systems, and computer program products for network firewall policy optimization.

SUMMARY

The subject matter described herein includes methods, systems, and computer program products for network firewall policy optimization. According to one method, a firewall policy including an ordered list of firewall rules is defined. For each rule, a probability indicating a likelihood of receiving a packet matching the rule is defined. Neighboring rules are sorted in order of non-increasing probability in a manner that preserves the firewall policy.

As used herein, the term "firewall" refers to a logical entity that is adapted to filter packets at the ingress and/or egress portions of a network based on a policy. The term "firewall" is intended to include systems that block packets from entering or leaving a network and systems that perform intrusion detection and intrusion protection functions for packets entering or leaving a network. Thus, the methods and systems described herein for firewall policy optimization can be used to optimize policies for intrusion detection and intrusion protection firewalls.

The subject matter described herein includes a method to improve firewall performance and lower packet delay that can be applied to both legacy and current systems. In one implementation, a firewall rule set may be re-ordered to minimize the average number of rule comparisons to determine the action, while maintaining the integrity of the original policy. Integrity is preserved if the reordered and original rules always arrive at the same result. To maintain integrity, a firewall rule set may be modeled as a Directed Acyclical Graph (DAG), where vertices are firewall rules and edges indicate precedence relationships. Given this representation, any linear arrangement of the DAG is proven to maintain the original policy integrity. Unfortunately, determining the optimal rule order from all the possible linear arrangements is proven to be *NP*-complete, since it is equivalent to sequencing jobs with precedence constraints for a single machine [15]. Although determining the optimal order is *NP*-complete, a heuristic will be described below to order firewall rules that reduces the average number of comparisons while maintaining integrity. Simulation

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 961

results show the proposed reordering method yields rule orders that are comparable to optimal; thus, provides a simple means to significantly improve firewall performance and lower packet delay.

The subject matter described herein for network firewall policy optimization may be implemented using any combination of hardware, software, or firmware. For example, the subject matter described herein may be implemented using a computer program product comprising computer executable instructions embodied in a computer readable medium. Exemplary computer readable media suitable for implementing the subject matter described herein include disk memory devices, chip memory devices, programmable logic devices, application specific integrated circuits, and downloadable electrical signals. In addition, a computer program product that implements the subject matter described herein may be implemented using a single device or computing platform or may be distributed across multiple devices or computing platforms.

BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the subject matter described herein will now be explained with reference to the accompanying drawings of which:

Figure 1A is a block diagram of a rule list Directed Acyclical Graph (DAG) for the firewall rules in Table 1;

Figure 1B is a linear arrangement of firewall rules corresponding to the original rule order;

Figure 2A is a policy DAG representing a reordered version of the firewall rule set in Figure 1 after sorting;

Figure 2B is a policy DAG for the rule set illustrated in Table 1 illustrating an optimal rule order;

Figure 3A is a graph illustrating average numbers of packet comparisons versus intersection percentage for different orderings of a firewall policy;

Figure 3B is a graph illustrating the percent difference in average numbers of packet comparisons between sorted and optimal sortings of a firewall policy;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 962

Figure 4A is a graph of percent difference comparisons between original and sorted versions of a firewall policy;

Figure 4B is a graph illustrating the performance impact of considering rule intersection and actions in performing firewall rule sorting;

Figure 5A is a table illustrating a firewall policy before sorting;

Figure 5B is a table illustrating a sorted version of the firewall policy of Figure 5A;

Figure 6A is a policy trie representation of the firewall policy of Figure 5B;

Figure 6B is a table representing the firewall policy of Figure 5A after TCP and UDP sub-trie rotation;

Figure 7A is a policy trie illustrating the splitting of rule r_5 into rules r_{5a} and r_{5b};

Figure 7B is a table illustrating the firewall policy of Figure 5A after splitting rule r_5 into rules r_{5a} and r_{5b};

Figure 8A is a graph illustrating firewall performance with and without rule splitting;

Figure 8B is a graph illustrating locations and probabilities where splitting is better;

Figure 9A is a policy DAG of a firewall rule set;

Figure 9B is a policy DAG illustrating compression of rules r_1 and r_5;

Figure 10A is a policy DAG of a firewall rule list;

Figure 10B is an original policy trie before sorting;

Figure 10C is a policy trie illustrating the firewall rules after rotating sub-tries T_1 and T_2;

Figure 11A is a policy DAG illustrating a firewall policy;

Figure 11B is a policy trie illustrating firewall rules of the policy before sorting;

Figure 11C is a policy trie illustrating the firewall policy after exchanging nodes r_1 and r_2 and nodes r_7 and r_8;

Figure 11D is a policy trie illustrating the firewall policy after exchanging sub-tries T_3 and T_4;

Figure 11E is a policy trie illustrating the firewall policy after exchanging sub-tries T_1 and T_2;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 963

Figure 12A is a policy DAG illustrating an original order for a firewall policy;

Figure 12B is a policy DAG illustrating a sorted version of the firewall policy of Figure 12A after exchanging neighbors;

Figure 12C is a policy DAG illustrating a sorted version of the firewall policy after exchanging lists L_1 including rules r_1 and r_3 and lists L_2 including rules r_2 and r_4;

Figure 13 is a flow chart illustrating exemplary steps for network firewall policy optimization according to an embodiment of the subject matter described herein;

Figure 14 is a block diagram illustrating exemplary components of a system for firewall policy optimization according to an embodiment of the subject matter described herein; and

Figure 15 is a block diagram illustrating a plurality of firewall nodes for implementing a firewall policy according to an embodiment of the subject matter described herein.

DETAILED DESCRIPTION OF THE INVENTION

Modeling Firewall Security Policies

As described above, a firewall rule set, also known as a firewall policy, is traditionally an ordered list of firewall rules. Firewall policy models have been the subject of recent research [1, 2, 16]; however, the primary purpose is anomaly detection and policy verification. In contrast, the policy model described in this herein is designed for firewall performance optimization and integrity. Firewall performance refers to reducing the average number of comparisons required to determine an action, while integrity refers to maintaining the original policy intent. Although improving the worst-case performance is important, it is not possible without changing the list-based representation [16, 24].

Firewall Rule and Policy Models

In the examples described herein, a rule r is modeled as an ordered tuple of sets, $r = (r[1], r[2], ..., r[k])$. Order is necessary among the tuples since

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 964

comparing rules and packets requires the comparison of corresponding tuples. Each tuple $r[l]$ is a set that can be fully specified, specify a range, or contain wildcards '*' in standard prefix format. For the Internet, security rules are commonly represented as a 5-tuple consisting of: protocol type, IP source address, source port number, IP destination address, and destination port number [30, 31]. Given this model, the ordered tuples can be supersets and subsets of each other, which forms the basis of precedence relationships. In addition to the prefixes, each filter rule has an action, which is to accept or deny. However, the action will not be considered when comparing packets and rules. Similar to a rule, a packet (IP datagram) d can be viewed as an ordered k-tuple $d = (d[1], d[2], ..., d[k])$; however, ranges and wildcards are not possible for any packet tuple.

Using the previous rule definition, a standard security policy can be modeled as an ordered set (list) of n rules, denoted as $R = \{r_1, r_2, ..., r_n\}$. A packet d is sequentially compared against each rule r_i starting with the first, until a match is found ($d \Rightarrow r_i$) then the associated action is performed. A match is found between a packet and rule when every tuple of the packet is a subset of the corresponding tuple in the rule.

Definition Packet d matches r_i if

$$d \Rightarrow r_i \quad \text{iff} \quad d[l] \subseteq r_i[l], \quad l = 1, ..., k$$

The rule list R is comprehensive if for every possible legal packet d a match is found using R. Furthermore, two rule lists R and R' are equivalent if for every possible legal packet d the same action is performed by the two rule lists. If R and R' are different (e.g. a reorder) and the lists are equivalent, then the policy integrity is maintained.

As previously mentioned, a rule list has an implied precedence relationship where certain rules must appear before others if the integrity of the policy is to be maintained. For example consider the rule list in Table 1. Rule r_1 must appear before rules r_3 and r_5, likewise rule r_6 must be the last rule in the policy. If for example, rule r_3 was moved to the beginning of the policy, then it will shadow [2] the original rule r_1. However, there is no precedence relationship between rules r_2 and r_4 given in Table 1. Therefore, the relative ordering of these two rules will not impact the policy integrity and can be

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 965

changed to improve performance. The present example assumes that the original policy is free from any anomalies. Likewise, when a policy is reordered to improve performance it should not introduce any anomalies, which will occur if precedence relationships are not maintained. As a result, a model is needed to effectively represent precedence relationships.

Modeling Rule List Precedence Relationships

The precedence relationship between rules in a policy will be modeled as a Directed Acyclical Graph (DAG) [23, 18]. Such graphs have been successfully used to represent the relative order of individual tasks that must take place to complete a job (referred to as a task graph model). Since certain rules must appear before others to maintain policy integrity, this structure is well suited for modeling the precedence of firewall rules. Let $G = (R, E)$ be a rule list DAG for a rule list R, where vertices are rules and edges E are the precedence relationships (constraint). A precedence relationship, or edge, exists between rules r_i and r_j, if $i < j$, the actions for each rule are different, and the rules intersect.

Definition The intersection of rule r_i and r_j, denoted as $r_i \cap r_j$ is

$$r_i \cap r_j = (r_i[l] \cap r_j[l]), \qquad l = 1, \ldots, k$$

Therefore, the intersection of two rules results in an ordered set of tuples that collectively describes the packets that match both rules. The rules r_i and r_j intersect if every tuple of the resulting operation is non-empty. In contrast, the rules r_i and r_j do not intersect, denoted as $r_i \not\cap r_j$, if at least one tuple is the empty set. Note the intersection operation is symmetric; therefore, if r_i intersects r_j, then r_j will intersect r_i. The same is true for rules that do not intersect.

For example consider the rules given in Table 1, the intersection of r_1 and r_3 yields (UDP, 1.1.*, *, 1.*, 80). Again, the rule actions are not considered in the intersection or match operation. Since these two rules intersect, a packet can match both rules for example d = (UDP, 1.1.1.1, 80, 1.1.1.1, 80). Furthermore, the actions of the two rules are different. Therefore, the relative order must be maintained between these two rules and an edge drawn from r_1 to r_3 must be present in the rule list DAG, as seen in Figures 1A and 1B. More

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 966

particularly, Figure 1A illustrates a rule list DAG for the rules in Table 1. Figure 1B illustrates a linear arrangement of the rules in Figure 1A. In Figures 1A and 1B, the vertices represent rules, the circles represents an accept rules, and the squares represents deny rules. Edges that connect the vertices represent precedence requirements. As can be seen from Figure 1B, because of the edge between rules r_1 and r_3, precedence between rules r_1 and r_3 must be maintained. In contrast consider the intersection of rules r_3 and r_5. These two rules intersect, indicating packets belonging to the set (UDP, 1.*, *, 1.*, *) would match both rules. However, it does not matter which of the two rules a packet matches first, since the action is the same for both rules. Therefore, an edge does not exist between rules r_3 and r_5 in the diagram. Similarly, rules r_2 and r_4 do not intersect due to the fifth tuple (destination port). A packet cannot match both rules indicating the relative order can change; therefore, an edge will not exist between them.

The match operation can be used to identify precedence relationships, but it cannot do so in every case. Consider a partial-match example [1], where r_a = (UDP, *, 80, 10.*, 90, accept) and r_b = (UDP, 10.*, 80, *, 90, deny). The intersection of r_a and r_b is (UDP, 10.*, 80, 10.*, 90); therefore a packet, such as d = (UDP, 10.10.10.10, 80, 10.10.10.10, 90), can match both rules. If r_a appears before r_b then the packet d is accepted, but if r_b occurs before r_a then d is rejected. As a result, the order of r_a and r_b in the original policy must be maintained. However, the match operation is unable to identify the precedence in this example. A partial match exists in between rules r_3 and r_5 in Table 1, but as previously discussed an edge does not exist between the rules since the actions are the same.

Using the rule list DAG representation a linear arrangement is sought that improves the firewall performance. As depicted in Figure 1B, a linear arrangement (permutation or topological sort) is a list of DAG vertices where all the successors of a vertex appear in sequence after that vertex [23]. Therefore it follows that a linear arrangement of a rule list DAG represents a rule order, if the vertices are read from left to right. Furthermore, it is proven in the following theorem that any linear arrangement of a rule list DAG maintains integrity.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 967

Theorem Any linear arrangement of a rule list DAG maintains integrity.

Proof Assume a rule list DAG *G* is constructed from the security policy *R* that is free of anomalies. Consider any two rules r_i and r_j in the policy, where $i < j$. If an edge between r_i and r_j in *G* does not exist, then a linear arrangement of *G* can interchange the order of the two rules. An edge will not exist if the rules do not intersect; however, a reorder will not affect integrity since a packet cannot match both rules. Shadowing is not introduced due to the reorder since the intersection operation is symmetric. An edge will not exist if the two rules intersect but have the same action; however, a reorder will not affect integrity since the same action will occur regardless of which rule is matched first. If an edge does exist between the rules, then their relative order will be maintained in every linear arrangement of *G*; thus maintaining precedence and integrity.

Rule List Optimization

As described above, it is important to inspect packets as quickly as possible given increasing network speeds and QoS requirements. Using the rule list DAG to maintain policy integrity, a linear arrangement is sought that minimizes the average number of comparisons required. However, this will require information not present in the firewall rule list. Certain firewall rules have a higher probability of matching a packet than others. As a result, it is possible to develop a policy profile over time that indicates frequency of rule matches (similar to cache hit ratio). Let $P = \{p_1, p_2, ..., p_n\}$ be the policy profile, where p_i is the probability that a packet will match rule *i* (first match in the policy). Furthermore, assume a packet will always find a match, $\sum_{i=1}^{n} p_i = 1$; therefore *R* is comprehensive. Using this information, the average number of rule comparisons required is

$$E[n] = \sum_{i=1}^{n} i \cdot p_i \qquad (1)$$

For example, the average number of comparisons required for the rule set in Table 1 is 5.03.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 968

Given a rule list DAG $G = (R, E)$ and policy profile $P = \{p_1, p_2, ..., p_n\}$ a linear arrangement π of G is sought that minimizes Equation 1. In the absence of precedence relationships, the average number of comparisons is minimized if the rules are sorted in non-increasing order according to the probabilities [26], which is also referred to as Smith's algorithm [27]. Precedence constraints cause the problem to be more realistic; however, such constraints also make determining the optimal permutation more problematic.

Determining the optimal rule list permutation can be viewed as job scheduling for a single machine with precedence constraints [15, 21]. The notation for such scheduling problems is $\alpha|\beta|\gamma|\delta$, where α is the number of machines, β is the precedence (or absence of) which can be represented as a DAG, γ is a restriction on processing time, and δ is the optimality criterion [15]. Determining the optimal rule order is similar to the $1|\beta|1|\sum \omega C_i$ scheduling problem, or optimality criterion, where ω is a weight associated with a job (for example, importance) and C_i is the completion time. As previously noted, the $1|||\sum \omega C_i$ problem can be solved in linear time the using Smith's algorithm [27], which orders jobs according to non-decreasing $\frac{t_i}{\omega}$ ratio, where t_i is the processing time of job i. In this case set $t_i = 1$ and $\omega = p_i \ \forall i$. However, Lawler [19] and Lenstra et al. [21] proved $1|\beta|1|\sum \omega C_i$ to be *NP*-complete via the linear arrangement problem, which implies determining the optimal firewall rule order is also *NP*-complete. Note, determining the number of possible permutations has been proven to be *NP*-hard [5].

Theorem $1|\beta|1|\sum \omega C_i \ \alpha$ *Determining the optimal order of a firewall rule list*

Proof Consider the $1|\beta|1|\sum \omega C_i$ problem. Each of n jobs J_i, $i \in I$, has to be processed without preemption on a single machine that can handle at most one job at a time. For each $i \in I$, let ω be the associated weight. Furthermore, let $G = (V, E)$ be a DAG that represents the precedence order of the jobs J_i. Assume the processing time of each job equals 1 time unit, the weights to be $0 \leq \omega \leq 1$ such that $\sum \omega = 1$, and β, which is G, to be a rule list DAG. In this case, the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 969

optimization criterion $\sum \omega \cdot C_i$ is the same as $\sum p_i \cdot i$, which is given in equation 1. Clearly, the optimal firewall rule ordering problem has a solution if and only if $1|\beta|1|\sum \omega C_i$ has a solution. Therefore, determining the optimal permutation of firewall rules is *NP*-complete.

Exemplary Rule Sorting Algorithm

Although determining the optimal rule permutation was proven to be *NP*-complete, reducing the average number of comparisons required to process a packet remains an important objective. As previously discussed, a sorting algorithm must maintain the precedence relationships among rules. Of course an exhaustive search is possible if the number of rules is small (generate and test all possible topological orders); however as proven in the previous section, this is not feasible with a realistic rule list.

One exemplary algorithm starts with the original the rule set, then sorts neighboring rules based on nonincreasing probabilities. However, an exchange of neighbors should never occur if the rules intersect and have different actions. This test preserves any precedence relationships in the policy. For example, the following sorting algorithm uses such a comparison to determine if neighboring rules should be exchanged. Note a_i denotes the action associated with the i^{th} rule.

```
done = false
while(!done)
    done = true
    for(i = 1; i <n; i++)
        if(pi < pi+1 AND (ri ∩ ri+1 OR ai == ai+1))then
            exchange rules, actions, and probabilities
            done = false
        endif
    endfor
endwhile
```

Other sorting algorithms are possible if the completion time of the sort is an issue; however, the method presented is easy to implement and only requires a simple neighbor comparison. Assume the match probabilities for the rule list given in Table 1. Applying the sorting algorithm to this rule list results in the ordering depicted in Figure 2A, which has 10% fewer comparisons on

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 970

average. However when using the algorithm, it is possible that one rule can prevent another rule from being reordered. For example, rule r_1 prevents rule r_5 from being placed closer to the beginning of the rule set. However, if both rules r_1 and r_5 are placed closer to the beginning of the policy while maintaining their relative order, the average number of comparison will be further reduced, 15% fewer. This is the optimal order for the 6 rule set, which is depicted in Figure 2B. Although this simple sorting algorithm is unable to move groups of rules, it can still improve the performance of the firewall system. Its effectiveness is measured experimentally in the next section.

Experimental Results

In this section, the average number of rule comparisons required after sorting the rules is compared with average number of comparisons required using the original order and the optimal order. Note the optimal rule ordering was determined via an exhaustive search. As a result, it was not feasible to determine the optimal ordering once the number of rules equaled 15 given the number of permutations to consider. Two different variations of the rule sorting method described in the preceding section were implemented. The difference between the methods was the comparison used to determine if neighboring rules should be exchanged. The first sorting algorithm exchanged neighboring rules if they were out of order and did not intersect. The if-condition was as follows.

if ($p_i < p_{i+1}$ **AND** $r_i \cap r_{i+1}$) **then**

Therefore, rule actions were not considered and the method will be referred to as *non-action sort*. The second sorting method did consider the rule action (as described in the preceding section) and will be referred to as *action sort*. The comparison between the two sorting algorithms will indicate the importance of considering rule actions when ordering rules.

In the first experiment, lists of 10 firewall rules were generated with random precedence relationships. The match probability of each rule was given by a Zipf distribution [20], which assigns probabilities according to the rank of an item. For this simulation, the last rule had the highest probability

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 971

which is consistent with most policies (last rules are more general). As a result, the original order yields the worst average number of comparisons. The intersection percentage measured the percentage of rules that intersect in the policy. This metric gives a value for the dependency level in the policy; however note, rule action is not considered when calculating this metric. Rules were equally likely to have an accept or deny action and the results (average number of rule comparisons) for a particular intersection percentage were averaged over 1000 simulations.

Results of the first experiment are given in Figures 3A and 3B. More particularly, Figure 3A is a graph illustrating the average number of packet comparisons versus intersection percentage for the original rule sorting, the non-action sort, the action sort, and the optimal sort. Figure 3B is a graph illustrating the percent difference in the average number of comparisons required between the sorted and optimal firewall policy configurations. The average number of comparisons required was lower for the sorted and the optimal lists when the intersection percentage was low, as seen in Figure 3A. This was expected since there is a large number of possible rule permutations (few edges in the DAG). When the intersection percentage approached 100%, the values converged to the number required for the original list. This is due to the limited number of possible rule orders, only one order in the extreme case. The percent difference between the sorted and optimal order is shown in Figure 3B. At zero intersection percentage the sorted and optimal orders are equal, since any ordering is possible. Similarly at an intersection percentage of 100, the sorted and optimal orders are equal since only one order is possible. Between these two extremes, the action sorting algorithm remains close to the optimal value. In contrast, the non-action sort had a maximum difference of 33%.

The benefit of rule sorting on larger policies is also of interest; however as previously described, it was not possible to determine the optimal ordering once the number of rules approaches 15. Therefore, the second experiment used larger rule sets, but only compared the original order and the sorted orders. The number of rules ranged from 10 to 1000, while the matching probabilities and intersection percentages were the same as the previous

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 972

experiment. The results of this experiment are depicted in Figure 4A. The sorted rule sets always performed equal to or better than the original order, while the action sort consistently performed better than the non-action sort. As noted in the previous experiment, the percent difference is very large given few intersections (e.g. 80% decrease for 1000 rules with 0% dependency), but approaches zero as all the rules intersect. As the number of rules increases, sorting is increasingly beneficial, although only at low intersection percentages. The benefit of sorting drops more significantly as the intersection percentage increases with larger rule sets. This is primarily due to the low matching probabilities of each rule, which requires a complete reordering to have a significant impact on the average number of comparisons. As a result, large rule sets can benefit from sorting if the intersection percentage is low.

The impact of considering rule actions when sorting is illustrated in Figure 4B. In this experiment lists of 1000 rules were generated, where the intersection percentages ranged from 0% to 100% and the percentage of rules with the same action varied from 50% to 100%. The performance of the action sort to the original list was then compared. As the percentage of rules with the same action increased, the percent difference (reduction) in the average number of comparisons also increased. This is due to the increased number of permutations possible when the rule actions are increasingly the same (fewer edges in the rule list DAG to consider). This is depicted in Figure 4B where a policy with a 100% intersection percentage can significantly reduce the number of comparisons (80% reduction) if all the rules have the same action. This performance increase is not possible with the non-action sort. Therefore, considering the rule action increases the number of possible rule orders, thereby providing more possibilities to improve firewall performance.

Conclusions

Network firewalls enforce a security policy by comparing arriving packets to a list of rules. An action, such as accept or deny, is then performed on the packet based on the matching rule. Unfortunately packet filtering can impose significant delays on traffic due to the complexity and size of rule sets.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 973

Therefore, improving firewall performance is important, given network Quality of Service (QoS) requirements and increasing network traffic loads.

The sections above describe rule ordering methods to improve the performance of network firewalls. Assuming each rule has a probability of a packet matching, firewall rules should be sorted such that the matching probabilities are non-increasing. This reduces the average number of comparisons and the delay across the firewall. However, a simple sort is not possible given precedence relations across rules. It is common in a security policy that two rules may match the same packet yet have different actions. It is this precedence relationship between rules that must be maintained to preserve integrity. The method described above uses Directed Acyclical Graphs (DAG) to represent the precedence order rules must maintain. Given this representation, a topological sort can be used to determine the optimal order (minimum average number of comparisons); however, the examples above prove this problem to be *NP*-complete (similar to job scheduling for a single nonpreemptive machine with precedence constraints). As an alternative, a simple sorting method was introduced that maintained the precedence order of the rules. Simulation results indicate this method can significantly reduce the average number of comparisons required and is comparable to optimal ordering.

Several areas exist for future research in optimizing firewall rule lists. The sorting method proposed above is based on a simple algorithm. Although it can offer an improvement over the original rule order, algorithms that can move groups of rules could provide a larger reduction in the average number of comparisons. The effect of stateful firewalls should also be addressed in future research. Security can be enhanced with connection state and packet audit information. For example, a table can be used to record the state of each connection, which is useful for preventing certain types of attacks (e.g., TCP SYN flood) [30, 31]. The impact of such rules on the firewall needs to be investigated and whether sorting can be done on-line to reflect temporal changes. In addition, more research is needed to determine more accurate probability distributions for packet matching and dependency percentages.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 974

Given this information, better algorithms can be designed and more realistic simulations can be performed.

Trie-Based Network Firewall Optimization Methods and Rule Splitting

Overview

As described above, while assuring firewall policy integrity is critical, performance is equally important given increasing network speeds and traffic volumes. Firewall processing delay can be reduced via hardware as well as policy optimization, where rules are manipulated to reduce the number of comparisons. For a given firewall installation, it can be determined that certain security rules have a higher probability of matching a packet than others. As a result, it is possible to develop a *policy profile* over time that indicates frequency of rule matches (similar to cache hit ratio). Given this information, trie-based methods are described herein to minimize the average number of comparisons required per packet. Although most firewall systems still utilize an ordered list representation, the proposed enhancements still are applicable to current firewall systems since a policy trie can be converted into an ordered list using an inorder traversal.

Rule Sorting

As described above, rule sorting of neighboring rules based on ascending probabilities in a manner that considers rules intersection and rule actions can be used to improve firewall performance by reducing the number of packet comparisons. The methods described above use Directed Acyclic Goal Graphs (DAGs) to represent firewalls. The methods described in this section will use the rule sorting algorithms in combination with policy trie representations to perform sorting of groups of rules in a manner that improves firewall performance.

In this example, it is assumed that a security policy contains a list of n ordered rules, $\{r_1, r_2, ..., r_n\}$. In addition, let p_i represent the probability that a packet will match rule i (first match in the policy). Therefore, the policy is comprehensive, which means every packet will have a match, if $\sum_{i=1}^{n} p_i = 1$.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 975

Using this information, the average number of rule comparisons required can be calculated as

$$E[n] = \sum_{i=1}^{n} i \cdot p_i$$

The average number of comparisons is minimized if the rules are sorted in non-ascending order according to the probabilities. However, this can only be achieved if the sorting algorithm considers the probabilities and whether one rule is a superset of another. For example, sorting algorithms must use the following comparison to determine if neighboring rules should be interchanged Note, a_i denotes the action associated with the i^{th} rule.

if ($p_i < p_{i+1}$ **AND** $r_i \not\cap r_{i+1}$ **OR** $a_i == a_{i+1}$)) **then**

Sorting rules in this fashion can have positive impact on the average number of comparisons required. Figure 5A is a table illustrating a firewall policy prior to rule sorting. The expected number of comparisons for the policy of Figure 5A is 4.26 comparisons per packet. Figure 5B is a table illustrating the firewall policy of Figure 5A after sorting using the above-described algorithm. In Figure 5B, the expected number of comparisons per packet has been reduced to 3.94. Thus, using the method described above, the number of comparisons per packet can be reduced by 8%.

However, it can be seen from Figures 5A and 5B that one rule can prevent another rule from being reordered. For example in Figure 5B, rule r_1 prevents rule r_4 from being placed closer to the beginning of the rule set. However, if both rules r_1 and r_4 are placed closer to the beginning of the policy while maintaining their relative order, the average number of comparison will be reduced.

To solve this problem, rule sets can be sorted a using policy tries. First, the rule list is converted into an equivalent policy trie. Each sub-trie will have an associated probability $\overline{p}$ that is average probability (hit-ratio) of the rules comprising the sub-trie. Sub-tries can be rotated around their parent node to increase performance, using the method described earlier. Since the policy trie (or equivalent rule set) is tested from left to right, rotation should occur if the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 976

probability of the right sub-trie is greater than the left sub-trie. As a result, the average number of comparisons required will be reduced and the policy integrity is maintained. Figure 6A is a trie representation of the rules in Figure 5B. Figure 6B is a table representing resulting rule set after rotating the TCP and UDP sub-tries. The TCP sub-trie in Figure 6A has an expected packet matching probability of 0.095 while the UDP sub-trie has an expected packet matching probability of 0.155. Thus, the UDP sub-trie and the TCP sub-trie should be rotated so that the UDP sub-trie becomes the left sub-trie and the TCP sub-trie becomes the right sub-trie. In the table illustrated in Figure 6B, which lists the ordering of the rules after rotating the sub-tries, the average number of comparisons per packets has been reduced to 3.70.

Rule Splitting

Rule splitting takes a general rule and creates more specific rules that collectively perform the same action over the same set of packets. Here, rule splitting is used to reduce the average number of comparisons. For example in Figure 7A, rule r_5 is split into two separate rules, r_{5a} for UDP and r_{5b} for TCP. Figure 7B is a table corresponding to the policy trie of Figure 7A. Once the rules are positioned based on their probabilities and their relation to other rules, the average number of rule comparisons is reduced to 2.98 (after sort, trie rotation, and splitting) which is a 30% less.

It many not be advantageous to split a general rule since it adds another rule to the policy. For example, assume a policy contains 20 rules where the first 19 rules have the same probability. Assume the last rule can be split and the new specific rule has a probability that is $\frac{3}{4}$ of the last rule. The impact of the probability of the last rule and the location of the new split rule, m, is depicted in Figures 8A and 8B. More particularly, Figure 8A is a graph illustrating the average number of comparisons for firewalls with and without rule splitting. Figure 8B is a graph illustrating average number of comparisons versus different locations of the new rule and probabilities. The average number of comparisons is reduced as the split rule is located closer to the first rule (surface decreases as m approaches one). Furthermore, splitting yields

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 977

better results when the general rule has a high probability. However as illustrated in Figure 8B, the closer the specific rule is to the location of the original rule the average number of comparisons increases, which is the penalty of adding one more rule to the policy.

The effect of splitting a single rule can be described mathematically as follows. Consider n rules where rule r_n can be split into rules $r_{n,1}$ and $r_{n,r}$ (whose union is the original rule). In addition, assume the split rule $r_{n,1}$ will be located at the m^{th} position ($1 \leq m < n$), while rule $r_{n,r}$ will remain at the n^{th} location. The goal is to determine the best location m, which yields a lower average number of comparisons as compared to the original rule set. This can be defined mathematically in the following formula.

$$\sum_{i=1}^{n} i \cdot p_i > \sum_{i=1}^{m-1} i \cdot p_i + m \cdot p_{n,l} + \sum_{i=m}^{n-1} (i+1) \cdot p_i + (n+1) \cdot p_{n,r}$$

The left side of the inequality is the average number of comparisons for the original rule set. The right hand side of the inequality is the average number of comparisons with the specific rule at location m. If it is assumed that the rules located between m and n have an equal probability (denoted as p) the previous equation can be solved for m.

$$m < \frac{n \cdot p - n \cdot p_n + (n+1) \cdot p_{n,r}}{p - p_{n,l}}$$

The new rule must be located between the first and m^{th}; however, its final location will depend on the relationship with the other rules (cannot be placed before any rule for which it is a superset). This result can be applied iteratively to multiple rules and repeatedly to the same rule.

Rule Optimization for QoS

In the preceding sections, optimizing the entire security policy has been discussed. However, it may be desirable to optimize the rules for a certain type of traffic, in order to reduce the average number of comparisons this traffic encounters. This can be done by organizing the policy trie via traffic

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 978

classification, then optimizing the sub-tries. The result is an ordering where the rules for the most important traffic are tested first.

Network Firewall Optimization Using Rule Compression

Another method for network firewall optimization is referred to as rule compression. Given a security policy R, one objective in optimizing firewall performance is to reduce number of rules in the policy while maintaining the original integrity. A method for reducing the number of rules called *rule compression*, which combines and replaces two rules in the policy will now be described. Although compression does reduce the number of rules, it can only be done if the integrity of the original policy is maintained.

Consider two rules r_i and r_j in policy R, where r_i appears before r_j. The two rules can be compressed if and only if the following conditions are true

1. The two rules intersect, $r_i \cap r_j \neq \emptyset$ (an edge between r_i and r_j exists in the policy DAG) (Tuples that are proper subsets and supersets are considered. Compression may still occur if the rules do not intersect or are adjacent in the set space, the resulting rule may contain ranges.)
2. The action of r_i and r_j is the same.
3. Rule r_i is not dependent on any rule r_k that is directly or indirectly dependent on r_j (only one path from r_i to r_j exists in the policy DAG).

The result of compression is a new rule, $r_{i,j}$ whose tuples are the union of the corresponding tuples of rules r_i and r_j. The new rule $r_{i,j}$ replaces both r_i and r_j in R; however, the location of $r_{i,j}$ in R may require the relocation of other rules. Let D_j be the ordered set of rules that appear after r_i that directly or indirectly depend upon r_j. The new rule is placed at the original location of r_i in R and the rules D_j are placed before $r_{i,j}$.

For example, consider the policy given in Table 2. It is possible to compress rules r_1 and r_5, creating the new rule $r_{1,5}$ =[TCP, *, *, 2.*, *, accept]. The ordered set D_j in this example consists of the rules r_3 and r_4. As a result, the new rule $r_{1,5}$ is placed at the original location of r_1 while the rules in D_j are placed before the new rule. Figures 9A and 9B respectively illustrate the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 979

original policy DAG for Table 2 and the policy DAG after rules r_1 and r_5 are compressed.

No.	Proto.	Source IP	Source Port	Destination IP	Destination Port	Action
1	TCP	1.1.*	*	2.2.*	*	accept
2	TCP	1.*	*	*	80	deny
3	TCP	3.3.*	*	*	80	accept
4	TCP	3.*	*	*	*	deny
5	TCP	*	*	2.*	90	accept

Table 2: Example Policy Consisting of Multiple Ordered Rules

Theorem *Compressing rules r_i and r_j of policy R will maintain integrity if the three conditions are met and method described is used.*

Proof Before r_i and r_j are compressed, relocate the rules D_j immediately before r_i and D_i. This relocation will maintain integrity due to the third requirement for compression (no rule in D_i can be dependent on any rule in D_j). Now place r_j directly after r_i. This relocation will not affect integrity since no rule in D_i can have an edge to r_j (again, due to the third condition). Compressing the neighboring rules r_i and r_j creates $r_{i,j}$. This does not affect integrity since the resulting rule only matches the packets that match r_i or r_j in the original policy. Therefore, the result is the compression of rule r_i and r_j , and the integrity of the policy is maintained.

Determining if Compression Should Occur

Although compression may be possible, it may not improve the performance of the policy due to the relocation of rules. Compression should only occur if the average number of comparisons required for the new policy R_c is less than the original policy R; therefore, $E[R_c] < E[R]$.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 980

Firewall Rule Sub-Trie and Sub-List Ordering

Policy Sub-Trie Ordering

Additional methods for firewall optimization are referred to as firewall sub-trie and sub-list ordering. Given a policy trie T, it may be desirable to order all sub-tries such that the average number of tuple comparisons is reduced and policy integrity is maintained. Consider a policy trie T that contains sub-tries T_i and T_j having the same parent node, as seen in Figures 10A-10C. Let P_i denote the sum of the probabilities of the rules contained in sub-trie i, while C_i denotes the number of comparisons required to completely traverse sub-trie i (which is equal to the number of branches). In order to reduce the average number of tuple comparisons, sub-tries that share the same parent node should be ordered such that the following two conditions are observed.

1. Sub-tries that share a parent node are ordered such that the P_i values are non-ascending (higher match probabilities occur first, from left to right).
2. If sub-tries that share a parent node have the same probability (P_i equals P_j), then order the sub-tries such that the C_i values are non-descending (sub-tries consisting of fewer comparisons occur first, from left to right).

These conditions are recursively applied throughout the policy trie as seen in Figures 11A-11D. However, these conditions do not necessarily maintain policy integrity.

Consider a group of n sub-tries that share the same parent node that are numbered sequentially from left to right. Consider any two non-intersecting sub-tries T_i and T_{i+k} in this group that are out of order. If two sub-tries do not intersect, denoted as $T_i \pitchfork T_j$, then no rule in one sub-trie will intersect with any rule in the other sub-trie, $r_k \pitchfork r_l$, $\forall r_k \in T_i$, $\forall r_l \in T_j$. The following *integrity condition* must be observed when reordering the sub-tries. The two sub-tries can be exchanged (rotated about the parent) if and only the sub-tries T_i through T_{i+k-1} do not intersect with T_{i+k} and the sub-tries T_{i+1} through T_{i+k} do not intersect with T_i. Similar to finding the linear sequence of a policy DAGs, these conditions maintain the policy trie integrity. This is observed in Figure 10A,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 981

where no edge exists between any rule in T_1 and T_2. More particularly, Figure 10A is a policy DAG of a six-rule firewall list. Figure 10B is a policy trie representation of the firewall rule list represented by Figure 10A. In Figure 10B, the average number of tuple compares required per packet is 5.5. Figure 10C illustrates the policy trie after rotating sub-tries T_1 and T_2. In Figure 10C, the average number of tuple compares required is 4.5. In Figures 10A-10C, the following probabilities are assumed:

$P_1 = 0.5$

$C_1 = 4$

$P_2 = 0.5$; and

$C_2 = 2$.

The average number of tuple comparisons is an estimate based on the longest trie path.

Using the conditions for comparing and ordering, the following sorting algorithm sorts neighboring sub-tries that share the same parent node.

```
m = a parent node in T
done = false
while(!done)
    done = true
    for each sub-trie Ti having parent m AND a right neighboring sub-trie
        if((Pi < Pi+1 AND (Pi == Pi+1 AND Ci > Ci+1)) AND Ti ⋔ Ti+1) then
        interchange Ti and Ti+1
        done = false
        endif
    endfor
endwhile
```

Figures 11A-11D illustrate policy trie rotation for an eight-rule firewall policy. More particularly, Figure 11A illustrates the policy DAG representation of an eight-rule firewall policy. Figure 11B is a policy trie representation of the original firewall policy. In this example, ordering occurs in three stages, starting at the bottom level and moving towards the top level of the trie. In Figure 11C, the leaves are ordered at the lowest level. More particularly, nodes r_1 and r_2 are exchanged with nodes r_7 and r_8. The average number of tuple compares for the ordering represented by Figure 11C is 10.45. In Figure 11D, the sub-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 982

tries on the second lowest level (T_3, T_4, T_5, and T_6) are ordered. More particularly, Figure 11D illustrates the policy trie after exchanging sub-tries T_3 and T_4. The average number of tuple compares for the representation in Figure 11D is 10.275. In Figure 11E, the sub-tries on the second highest level, (T_1 and T_2) are ordered. Figure 11E illustrates the final ordering after exchanging T_1 and T_2. In the policy trie of Figure 11E, the average number of tuple compares required is 6.775. The average number of tuple comparisons is an estimate based on the longest trie path.

Policy Sub-List Ordering

The preceding section describes conditions for sorting policy sub-tries, which allow the exchange of groups of rules (sub-tries). The same conditions can be applied to list based policies, where sub-lists are ordered to improve the average number of rule comparisons.

In the sections above relating to rule sorting, a method is described to exchange neighboring rules in a list-based security policy, as seen in Figure 12B. However, it was observed that when ordering rules in this fashion it is possible that one rule can block the exchange of another. If groups of rules were allowed to be exchanged the average number of rule comparisons could be further reduced.

Given a policy list L, it may be desirable to order all sub-lists such that the average number of rule comparisons is reduced and policy integrity is maintained. Consider a policy list L that contains sub-tries L_i and L_j. Let P_i denote the sum of the probabilities of the rules contained in sub-list i, while C_i denotes the number of rules in sub-list i. In order to reduce the average number of rule comparisons, sub-lists should be ordered such that the following two conditions are observed.

1. Sub-lists are ordered such that the P_i values are non-ascending (higher match probabilities occur first, from left to right).
2. If sub-lists have the same probability (P_i equals P_j), then order the sub-lists such that the C_i values are non-descending (sub-lists consisting of fewer comparisons occur first, from left to right).

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 983

These conditions are applied throughout the policy list; however as with policy tries, these conditions do not necessarily maintain policy integrity.

Consider a group of n sub-lists that are numbered sequentially from left to right. Consider any two non-intersecting sub-lists L_i and L_{i+k} in this group that are out of order. If two sub-lists do not intersect, denoted as $L_i \pitchfork L_j$, then no rule in one sub-list will intersect with any rule in the other sub-list, $r_k \pitchfork r_l$, $\forall r_k \in L_i$, $\forall r_l \in L_j$. The following *integrity condition* must be observed when reordering the sub-lists. The two sub-lists can be exchanged if and only if the sub-lists L_i through L_{i+k-1} do not intersect with any rule in L_{i+k} and sub-lists L_{i+1} through L_{i+k} do not intersect with L_i. Similar to finding the linear sequence of a policy DAGs, these conditions maintain the policy trie integrity. For example, consider the sub-lists $L_1 = \{r_1, r_3\}$ and $L_2 = \{r_2, r_4\}$ given in Figure 12B. No edge exists between any rule in L_1 and L_2. The two sub-lists are out of order with respect to P and should be exchanged. The resulting list is given in Figure 12C. Figure 12A illustrates the original policy order and results in an average number of rule comparisons of 3.4. In the sorted version in Figure 12B, the average number of rule comparisons is 3.3. In Figure 12C, the average number of rule comparisons has been reduced to 2.7. Using the conditions for comparing and ordering, the following sorting algorithm sorts neighboring sub-lists.

```
done = false
while(!done)
    done = true
    for each sub-list Li having a right neighboring sub-list
        if ((Pi < Pi+1 OR (Pi == Pi+1 AND Ci > Ci+1)) AND Li ⋔ Li+1)then
            interchange rules and probabilities
            done = false
        endif
    endfor
endwhile
```


CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 984

Summary of Firewall Policy Optimization Techniques and Methods

As stated above, the methods and systems described herein for a network firewall policy optimization can be implemented using any combination of hardware, software, and/or firmware. Alternatively, the steps for firewall policy optimization can be performed manually with the number of rules being optimized is small enough. Figure 13 is a flow chart illustrating an exemplary process for firewall policy optimization according to an embodiment of the subject matter described herein. Referring to Figure 13, in block **1300**, a firewall policy including an ordered list of firewall rules is defined. In block **1302**, for each rule, a probability indicating a likelihood of receiving a packet that matches each rule is specified. In block **1304**, neighboring rules are sorted in order of non-increasing probability in a manner that preserves firewall policy. In block **1306**, groups of rules are sorted in order of non-increasing probability in a manner that preserves firewall policy. For example, any of the policy-trie-based methods described above may be used. In block **1308**, one or more rules are split to reduce the number of average packet comparisons. In block **1310**, intersecting rules having common actions are identified and collapsed into single rules.

As described above, the subject matter described herein may be implemented in hardware, software, and/or firmware. In one implementation, the subject matter described herein for firewall policy optimization may be implemented on a general purpose computing platform. Figure 14 is a block diagram of a general purpose computing platform including software for firewall policy optimization according to an embodiment of the subject matter described herein. Referring to Figure 14, computing platform **1400** may be a general purpose computer, such as a personal computer, that includes a microprocessor **1402**, memory **1404**, and I/O interfaces **1406**. Microprocessor **1400** may be any suitable microprocessor, such as any of the Intel or AMD families of microprocessors. Memory **1404** may include volatile memory for running programs and persistent memory, such as one or more disk storage devices. I/O interface **1406** may include interfaces with I/O devices, such as user input devices and output devices.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 985

In the illustrated example, software that may be resident in memory **1404** includes a firewall policy editor **1408** and a firewall policy optimizer **1410**. Firewall policy editor **1408** may allow a user to define a firewall policy. For example, firewall policy editor **1408** may allow a user to define a firewall policy by specifying different rules. The rules may be specified in a graphical manner, for example using policy DAGs as described above. Alternatively, firewall policy editor **1408** may allow a user to input rules in a tabular manner, as illustrated in any of the tables described herein. Firewall policy optimizer **1410** may implement any or all of the firewall policy optimization techniques to order rules entered via firewall policy editor **1408** in a manner that preserves policy integrity and that enhances firewall performance.

A firewall rule set that is optimized using the subject matter described herein may be implemented on any firewall system that includes one or more firewalls. For example, an optimized firewall rule set according to embodiments of the subject matter described herein may be implemented using any of the hierarchical, multi-node firewall systems described in commonly assigned, co-pending U.S. patent application no. 11/316,331, filed December 22, 2005, the disclosure of which is incorporated herein by reference in its entirety. Figure 15 illustrates an example of a multi-node firewall system suitable for implementing the firewall rules according to an embodiment of the subject matter described herein. Referring to Figure 15, a plurality of firewall nodes **1500**, **1502**, **1504**,and **1506** may collectively implement any of the optimized firewall policies described herein. For example, the firewall nodes may collectively implement different portions of a firewall policy data structure including an ordered list of firewall rules. A firewall policy engine **1508** resident on each firewall node may filter packets using its local portion **1510** of the firewall policy data structure. A control node **1512** may include a firewall policy optimizer **1514** that measures the average number of comparisons per packets for the current rule configuration and may dynamically reorder the rules to reduce the average number of comparisons per packet. For example, firewall policy optimizer **1514** may utilize any of the methods described herein to rearrange rules and improve firewall performance. Rules may be tested for rearrangement at user-specified

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 986

intervals or when average number of packet comparisons increases by a user-specified amount.

References

The disclosure of each of the following references is hereby incorporated herein by reference in its entirety.

[1] E. Al-Shaer and H. Hamed. Firewall Policy Management Advisor for Anomaly Detection and Rule Editing. In *Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management*, 2003.

[2] E. Al-Shaer and H. Hamed. Modeling and Management of Firewall Policies. *IEEE Transactions on Network and Service Management*, 1(1), 2004.

[3] S. M. Bellovin and W. Cheswick. Network Firewalls. *IEEE Communications Magazine*, pages 50–57, Sept. 1994.

[4] C. Benecke. A Parallel Packet Screen for High Speed Networks. In *Proceedings of the 15th Annual Computer Security Applications Conference*, 1999.

[5] G. Brightwell and P. Winkler. Counting Linear Extensions is #P-Complete. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*, 1991.

[6] M. Christiansen and E. Fleury. Using Interval Decision Diagrams for Packet Filtering. Technical report, BRICS, 2002.

[7] D. Decasper, Z. Dittia, G. Parulkar, and B. Plattner. Router Plugins: A Software Architecture for Next-Generation Routers. *IEEE/ACM Transactions on Networking*, 8(1), February 2000.

[8] M. Degermark, A. Brodnik, S. Carlsson, and S. Pink. Small Forwarding Tables for Fast Routing Lookups. In *Proceedings of ACM SIGCOMM*, pages 4 – 13, 1997.

[9] U. Ellermann and C. Benecke. Firewalls for ATM Networks. In Proceedings of *INFOSEC'COM*, 1998.

[10] A. Feldmann and S. Muthukrishnan. Tradeoffs for Packet Classification. In Proceedings of the *IEEE INFOCOM*, pages 397 – 413, 2000.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 987

[11] E. W. Fulp. Firewall Architectures for High Speed Networks. Technical Report 20026, Wake Forest University Computer Science Department, 2002.

[12] E. W. Fulp. Firewall Policy Models Using Ordered-Sets. Technical report, Wake Forest University Computer Science Department, 2004.

[13] R. Funke, A. Grote, and H.-U. Heiss. Performance Evaluation of Firewalls in Gigabit-Networks. In *Proceedings of the Symposium on Performance Evaluation of Computer and Telecommunication Systems*, 1999.

[14] S. Goddard, R. Kieckhafer, and Y. Zhang. An Unavailability Analysis of Firewall Sandwich Configurations. In *Proceedings of the 6th IEEE Symposium on High Assurance Systems Engineering*, 2001.

[15] R. L. Graham, E. L. Lawler, J. K. Lenstra, and A. H. G. R. Kan. Optimizing and Approximation in Determinstic Sequencing and Scheduling: A Survey. *Annals of Discrete Mathematics*, 5:287 – 326, 1979.

[16] A. Hari, S. Suri, and G. Parulkar. Detecting and Resolving Packet Filter Conflicts. In *Proceedings of IEEE INFOCOM*, pages 1203–1212, 2000.

[17] HiPAC. High Performance Packet Classification. http://www.hipac.org.

[18] E. Horowitz, S. Sahni, and D. Mehta. *Fundamentals of Data Structures in C++*. Compuer Science Press, 1995.

[19] E. L. Lawler. Sequencing Jobs to Minimize Total Weighted Completion Time Subject to Precedence Constraints. *Annals of Discrete Mathematics*, 2:75 – 90, 1978.

[20] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the Self-Similar Nature of Ethernet Traffic. *IEEE Transactions on Networking*, 2:1 – 15, 1994.

[21] J. K. Lenstra and A. H. G. R. Kan. Complexity of Scheduling under Precedence Constraints. *Operations Research*, 26(1):22 – 35, 1978.

[22] O. Paul and M. Laurent. A Full Bandwidth ATM Firewall. In Proceedings of the 6th European Symposium on *Research in Computer Security ESORICS'2000*, 2000.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 988

2006230171 08 Sep 2011

[23] B. R. Preiss. *Data Structures and Algorithms with Object-Oriented Design Patterns in C++*. John Wiley & Sons, 1999.

[24] L. Qui, G. Varghese, and S. Suri. Fast Firewall Implementations for Software and Hardware-Based Routers. In *Proceedings of ACM SIGMETRICS*, June 2001.

[25] V. P. Ranganath and D. Andresen. A Set-Based Approach to Packet Classification. In *Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems*, pages 889–894, 2003.

[26] R. Rivest. On Self-Organizing Sequential Search Heuristics. *Communications of the ACM*, 19(2), 1976.

[27] W. E. Smith. Various Optimizers for Single-Stage Production. *Naval Research Logistics Quarterly*, 3:59 – 66, 1956.

[28] V. Srinivasan, G. Varghese, S. Suri, and M. Waldvogel. Fast and Scalable Layer Four Switching. In *Proceedings of ACM SIGCOMM*, pages 191–202, 1998.

[29] S. Suri and G. Varghese. Packet Filtering in High Speed Networks. In *Proceedings of the Symposium on Discrete Algorithms*, pages 969 – 970, 1999.

[30] R. L. Ziegler. *Linux Firewalls*. New Riders, second edition, 2002.

[31] E. D. Zwicky, S. Cooper, and D. B. Chapman. *Building Internet Firewalls*. O'Reilly, 2000.

It is to be understood that, if any prior art publication is referred to herein, such reference does not constitute an admission that the publication forms a part of the common general knowledge in the art, in Australia or any other country.

In the claims which follow and in the preceding description of the invention, except where the context requires otherwise due to express language or necessary implication, the word "comprise" or variations such as "comprises" or "comprising" is used in an inclusive sense, i.e. to specify the presence of the stated features but not to preclude the presence or addition of further features in various embodiments of the invention.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 989

It will be understood that various details of the invention may be changed without departing from the scope of the invention. Furthermore, the foregoing description is for the purpose of illustration only, and not for the purpose of limitation.

5

2006230171 06 Jun 2012

CLAIMS

The claims defining the invention are as follows:

1. A method for producing a performance optimized firewall policy, the method comprising:
   (a) defining a firewall policy including an ordered list of firewall rules;
   (b) for each rule, defining a first match probability indicating the likelihood that the rule will be the first match for a given packet;
   (c) sorting neighboring rules in order of non-increasing probability for the first match probabilities and in a manner that preserves the firewall policy; and
   (d) sorting groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy, wherein at least one of the groups of two or more rules includes a first rule that prevents a second rule in the group from being moved, within the group, closer to a beginning of the ordered list that implements the firewall policy and wherein sorting the groups of rules includes moving the first and second rules closer to the beginning of the ordered list while maintaining their relative position.

2. The method of claim 1 wherein sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy includes rearranging the rules if a first rule has a lower probability than a second rule and the first rule is not a subset of the second rule with a different action.

3. The method of claim 1 comprising sorting groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy.

COMS ID No: ARCS-372276 Received by IP Australia: Time (H:m) 11:49 Date (Y-M-D) 2012-06-06

CISCO EXHIBIT 1002
Filed - Mistry - 9565213 Patent
Page 991

2006230171 06 Jun 2012

4. The method of claim 3 wherein sorting groups of two or more rules in a manner that preserves firewall policy includes assigning each group of rules to a sub-trie in a policy trie, computing a sum of packet matching probabilities for rules in each sub-trie, and, in response to determining that a first sub-trie has a lower sum than a second sub-trie, rotating the first and second sub-tries about a common parent so that rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy.

5. The method of claim 4 comprising, in response to determining that the first and second sub-tries have equal sums, determining whether the second sub-trie has an lower number of branches than the first sub-trie, and, in response to determining that the second sub-trie has a lower number of branches than the first sub-trie, rotating the first and second sub-tries about a common parent so that the rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy.

6. The method of claim 3 wherein sorting groups of two or more rules in a manner that preserves the firewall policy includes:

(a) identifying first and second sub-lists of rules in the ordered list of firewall rules, the first sub-list preceding the second sub-list in the ordered list of firewall rules;

(b) determining a sum of packet matching probabilities for each of the first and second sub-lists;

(c) determining whether the first and second sub-lists intersect; and

(d) in response to determining that the sum of the packet matching probabilities for the second sub-list is greater than the sum of the packet matching probabilities for the first sub-list and that the first and second sub-lists do not intersect, switching the order of the first and second sub-lists in the ordered list of firewall rules.

7. The method of claim 6 comprising, in response to determining that the sums of the packet matching probabilities for the first and sub-lists are

CISCO EXHIBIT 1002
COMS ID No: ARCS-372276 Received by IP Australia: Time (H:m) 11:49 Date (Y-M-D) 2012-06-06 Filed: 9565213 Patent
Page 992

2006230171 06 Jun 2012

equal, determining whether the second sub-list has a lower number of rules than the first sub-list, and, in response to determining that the second sub-list has a lower number of rules than the first sub-list, switching the order of the first and second sub-lists in the ordered list of firewall rules.

8. The method of claim 1 comprising splitting at least one of the firewall rules to reduce an average number of comparisons required per received packet.

9. The method of claim 1 comprising identifying intersecting rules having a common action and collapsing the rules into a single rule representing a union of the intersecting rules.

10. The method of claim 1 wherein sorting neighboring rules includes sorting rules such that comparisons for a particular class of packets are reduced.

11. The method of claim 1 comprising assigning the rules to vertices in a directed acyclical graph (DAG), determining precedence relationships between the rules, assigning the precedence relationships to edges that connect the vertices in the DAG, and wherein sorting the neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy includes sorting the neighboring rules in a manner that preserves the precedence relationships represented by the DAG.

12. The method of claim 1 wherein the rules comprise intrusion detection rules.

13. The method of claim 1 wherein the rules comprise intrusion protection rules.

CISCO EXHIBIT 1002
COMS ID No: ARCS-372276 Received by IP Australia: Time (H:m) 11:49 Date (Y-M-D) 2012-06-06 File No: 2006230171 Filing Date: 2006-09-06 9565213 Patent
Page 993

2006230171 06 Jun 2012

14. A system for firewall policy optimization, for implementing the steps of any one of claims 1-13.

15. A network firewall comprising:

(a) a firewall policy data structure for storing an ordered list of firewall rules;

(b) a firewall policy engine for implementing a firewall policy by comparing packets to the rules in the order specified by the firewall policy data structure; and

(c) a firewall policy optimizer for optimizing performance of the network firewall by reordering the rules in a manner that reduces an average number of rule comparisons per packet and that preserves the firewall policy, wherein the firewall policy optimizer is further configured to sort groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy, wherein at least one of the groups of two or more rules includes a first rule that prevents a second rule in the group from being moved, within the group, closer to a beginning of the ordered list that implements the firewall policy and wherein sorting the groups of rules includes moving the first and second rules closer to the beginning of the ordered list while maintaining their relative position.

16. The network firewall of claim 15 wherein the firewall policy optimizer is adapted to measure average number of comparisons per packet for a rule ordering specified by the firewall policy data structure and to dynamically re-order to rules to reduce the average number of comparisons per packet.

17. A computer program product comprising computer executable instructions embodied in a computer readable medium which, when executed on one or more computing devices or platforms, performs steps in accordance with any one of claims 1 - 13.

CISCO EXHIBIT 1002
COMS ID No: ARCS-372276 Received by IP Australia: Time (H:m) 11:49 Date (Y-M-D) 2012-06-06 File History of 9565213 Patent
Page 994

2006230171 06 Jun 2012

18. A method for producing a performance optimized firewall policy, the method substantially as herein described with reference to the accompanying drawings.

19. A network firewall substantially as herein described with reference to the accompanying drawings.

CISCO EXHIBIT 1002
COMS ID No: ARCS-372276 Received by IP Australia: Time (H:m) 11:49 Date (Y-M-D) 2012-06-06 Filed: Mid 2013-0009565213 Patent
Page 995

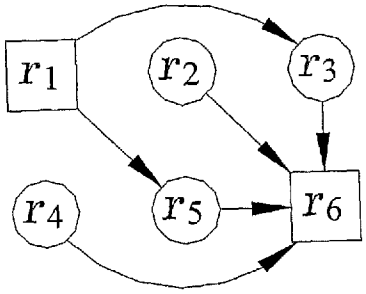

FIG. 1A

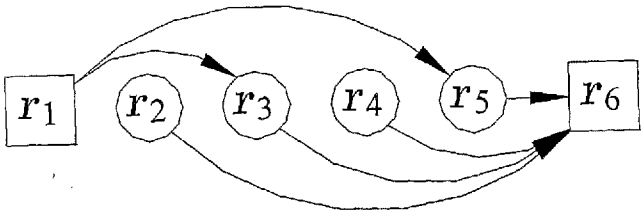

FIG. 1B

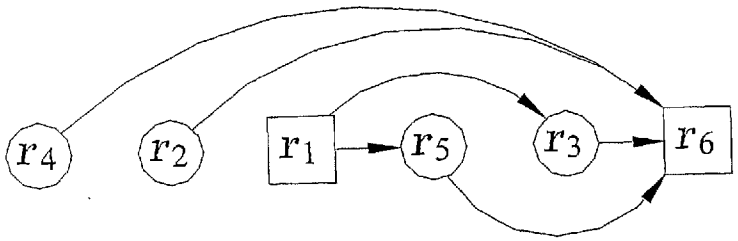

FIG. 2A

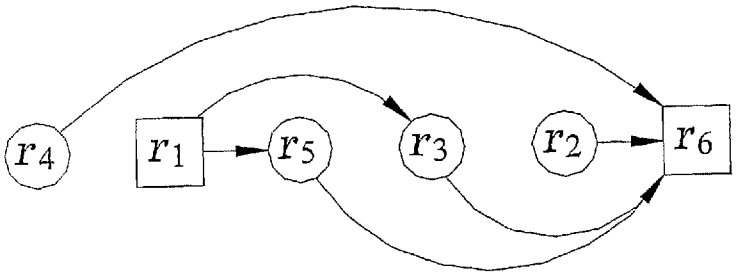

FIG. 2B

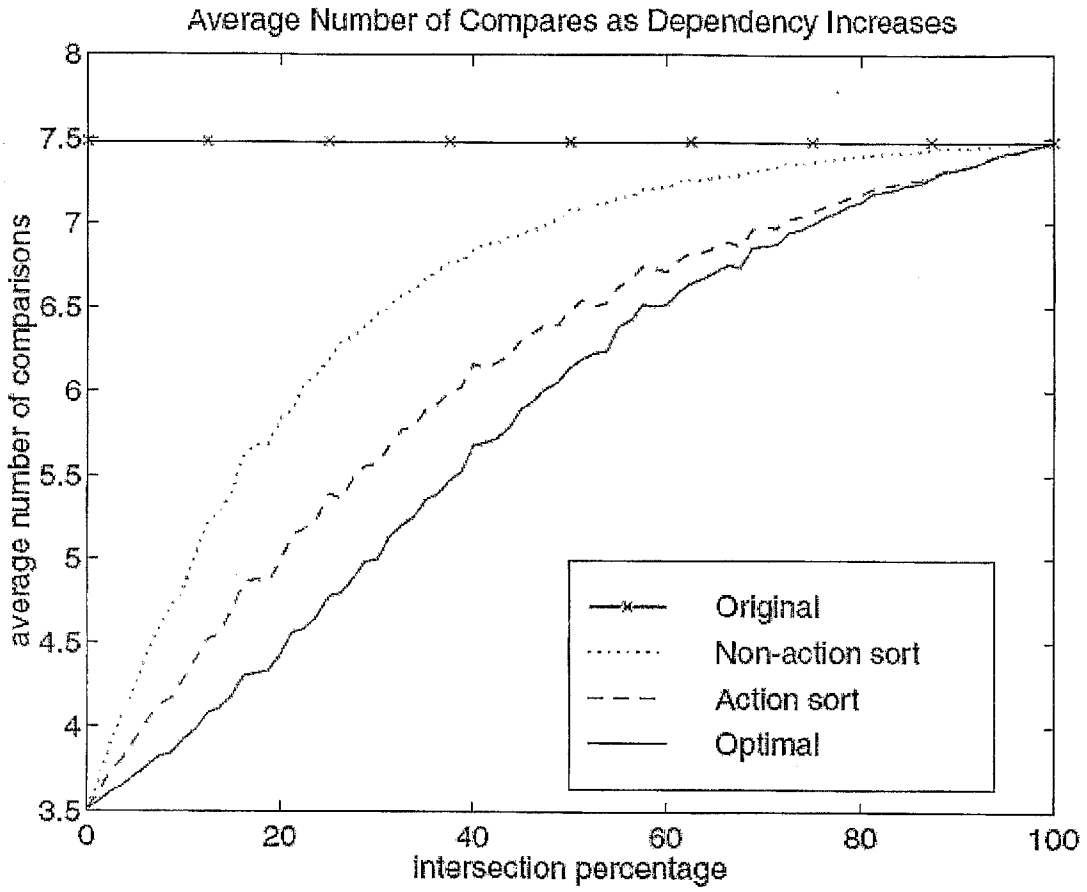
Average Number of Compares as Dependency Increases
average number of comparisons
intersection percentage
8
7.5
7
6.5
6
5.5
5
4.5
4
3.5
0
20
40
60
80
100
Original
Non-action sort
Action sort
Optimal

FIG. 3A

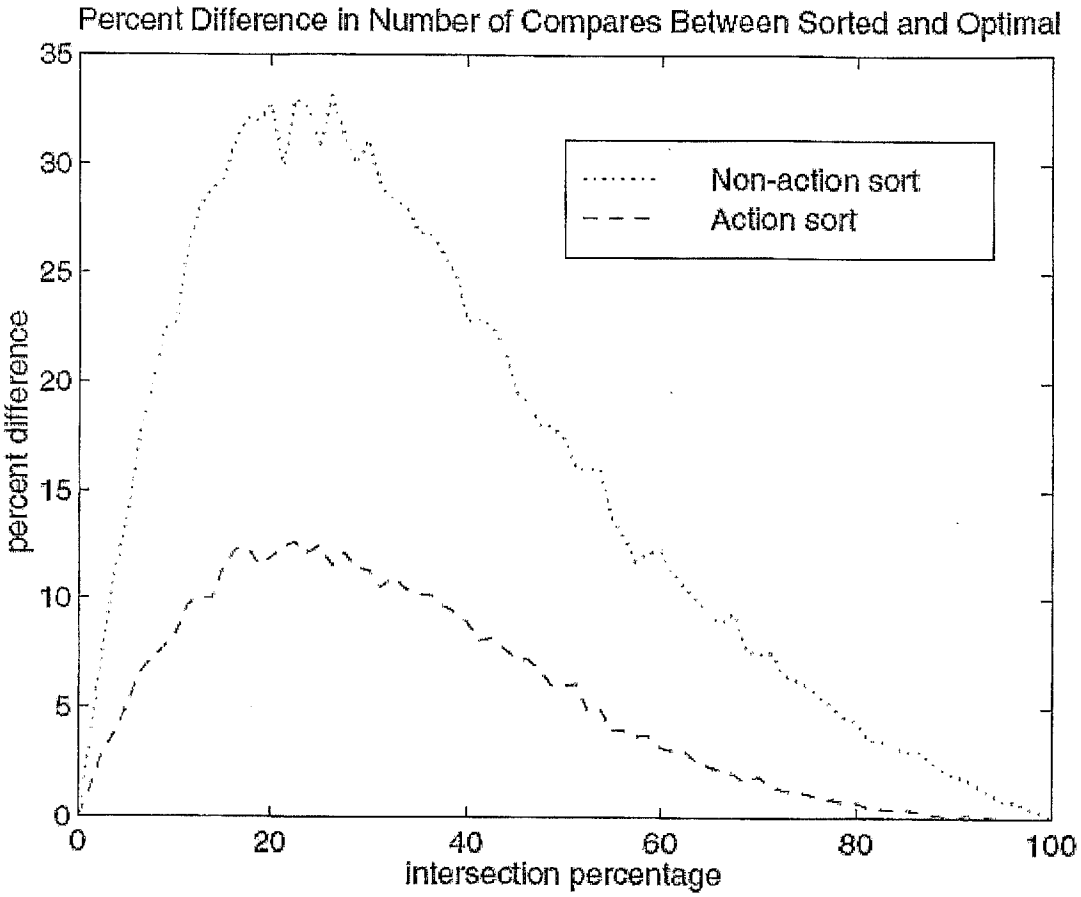
Percent Difference in Number of Compares Between Sorted and Optimal
percent difference
intersection percentage
Non-action sort
Action sort
0
5
10
15
20
25
30
35
0
20
40
60
80
100

FIG. 3B

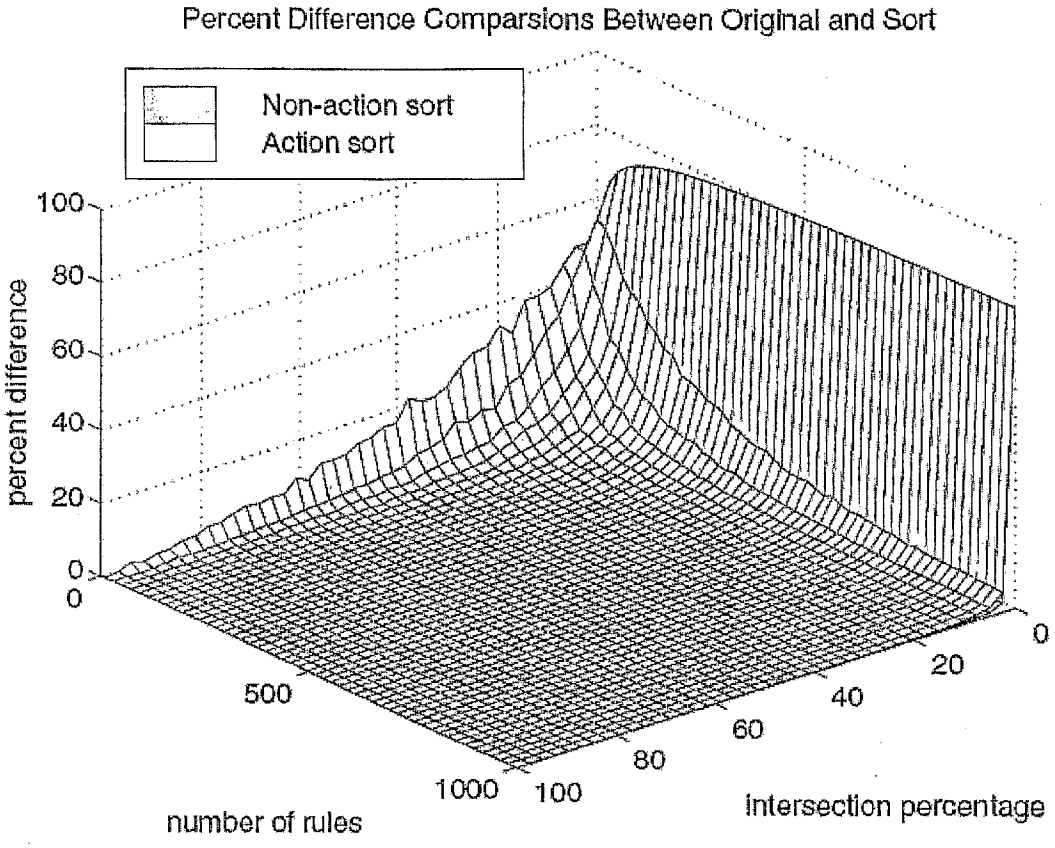
Percent Difference Comparsions Between Original and Sort
Non-action sort
Action sort
percent difference
100
80
60
40
20
0
0
500
1000
number of rules
0
20
40
60
80
100
intersection percentage

FIG. 4A

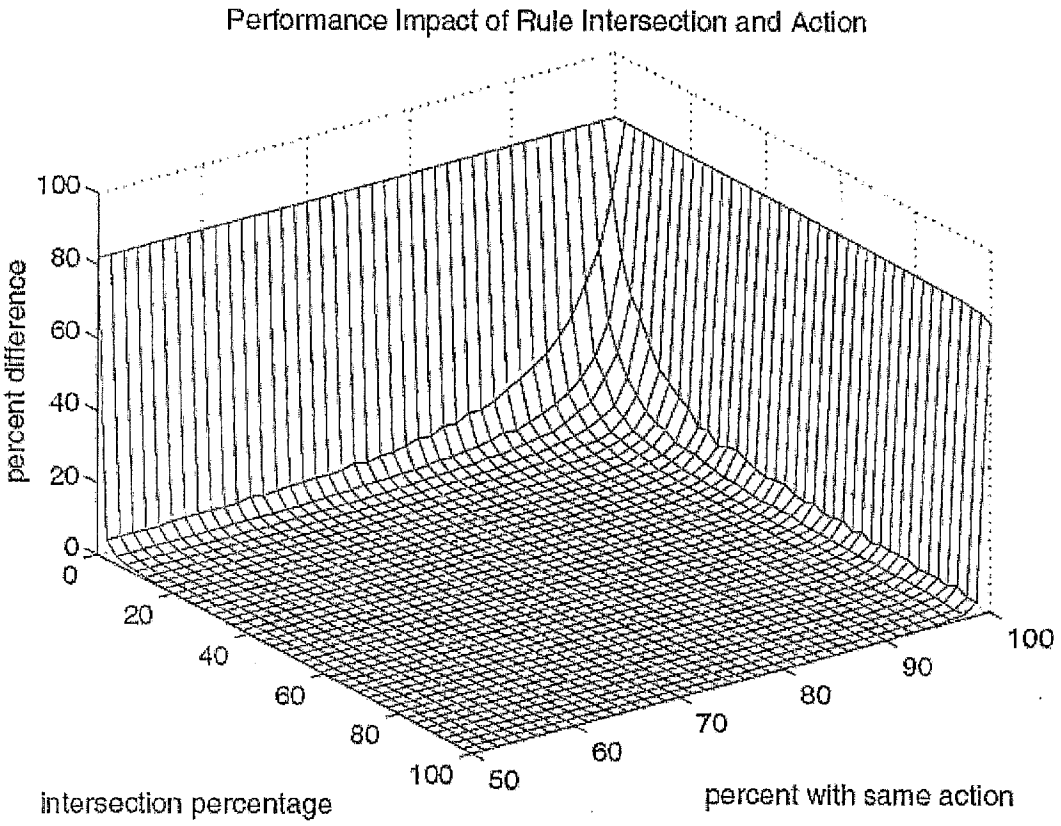
Performance Impact of Rule Intersection and Action
percent difference
100
80
60
40
20
0
0
20
40
60
80
100
intersection percentage
50
60
70
80
90
100
percent with same action

FIG. 4B

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 5A

r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 5B

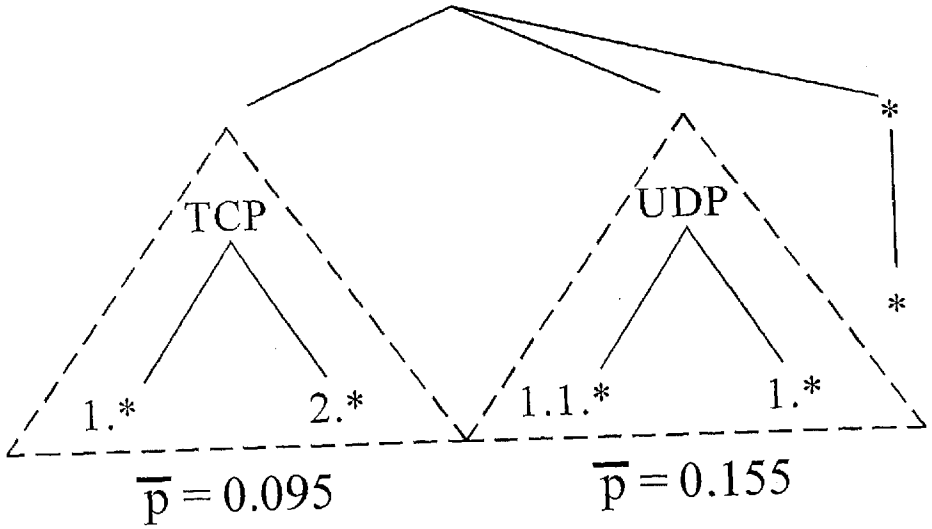

FIG. 6A

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 6B

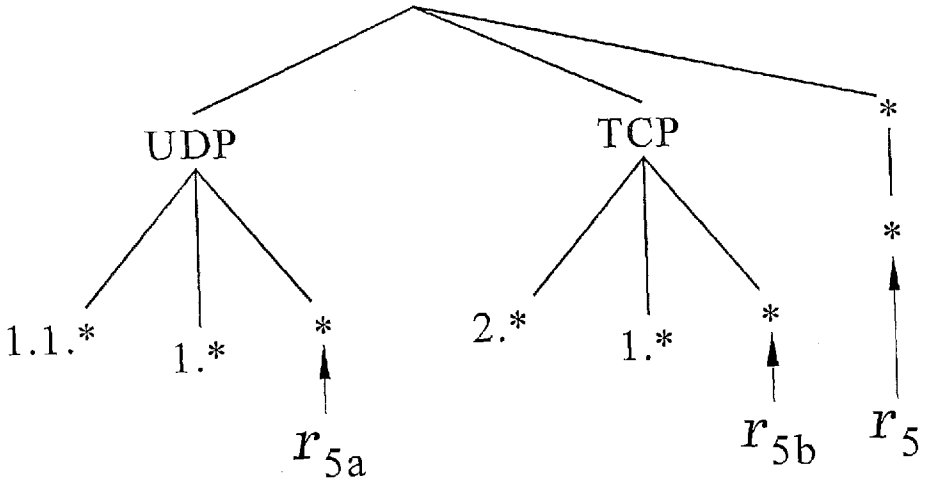

FIG. 7A

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	ACCEPT)	$p_4 = 0.30$
r_{5a}	(UDP,	*,	DENY)	$p_{5a} = 0.47$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_{5b}	(TCP,	*,	DENY)	$p_{5b} = 0.03$

FIG. 7B

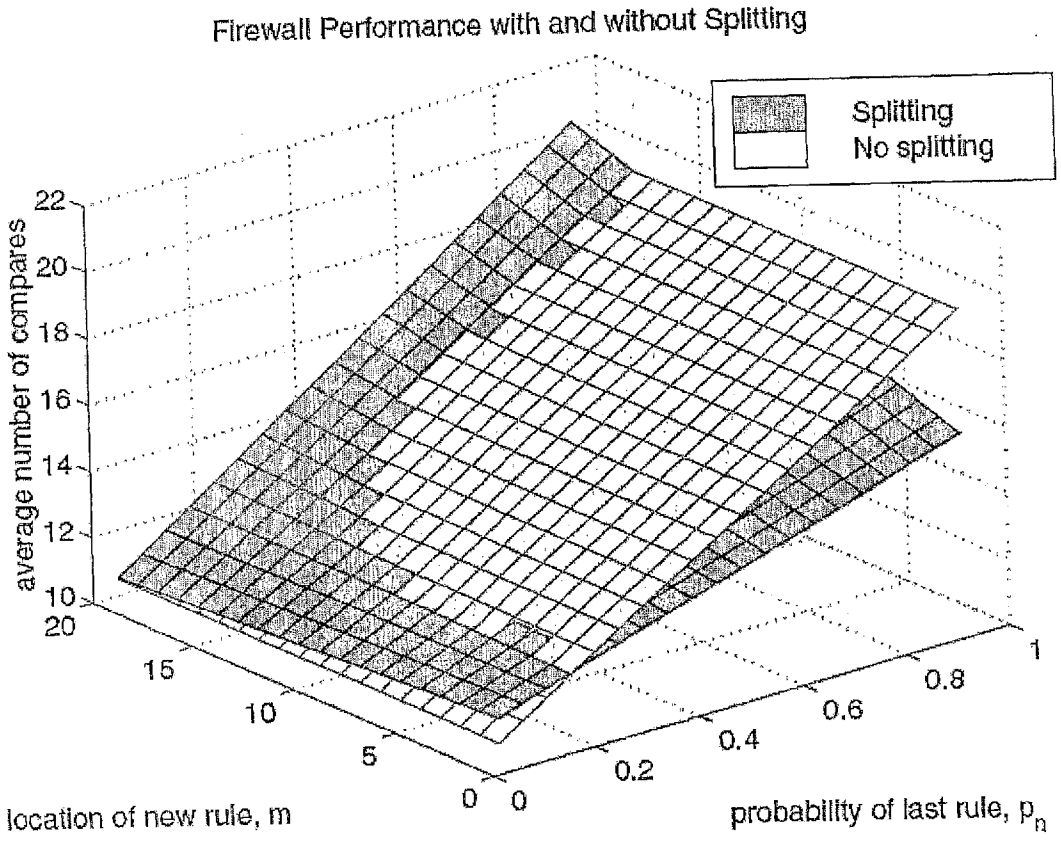
Firewall Performance with and without Splitting
Splitting
No splitting
average number of compares
22
20
18
16
14
12
10
20
15
10
5
0
location of new rule, m
0
0.2
0.4
0.6
0.8
1
probability of last rule, p_n

FIG. 8A

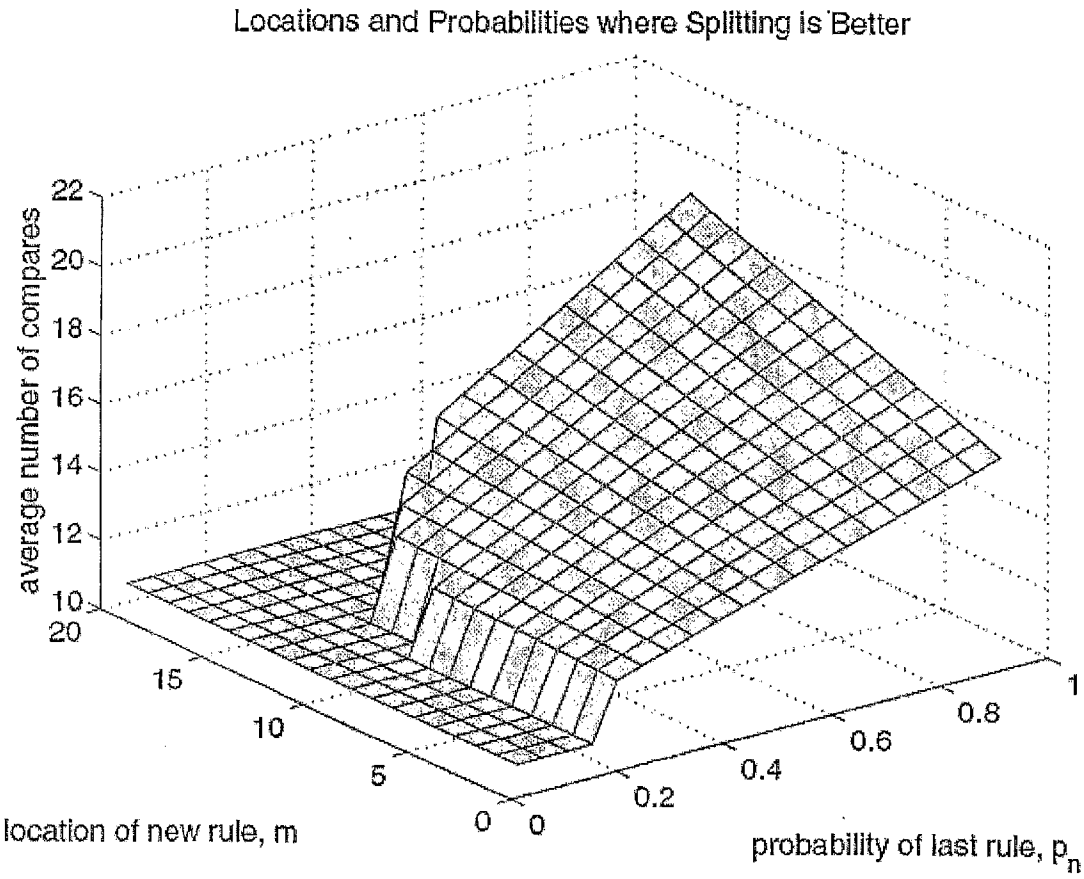
Locations and Probabilities where Splitting is Better
average number of compares
22
20
18
16
14
12
10
20
15
10
5
0
location of new rule, m
0
0.2
0.4
0.6
0.8
1
probability of last rule, p_n

FIG. 8B

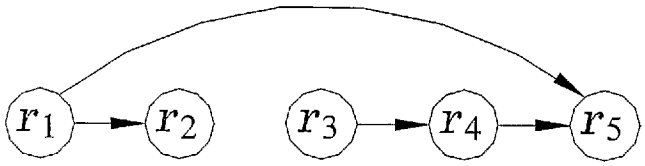

FIG. 9A

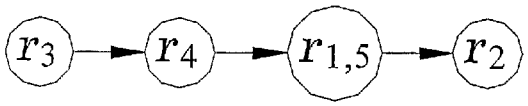

FIG. 9B

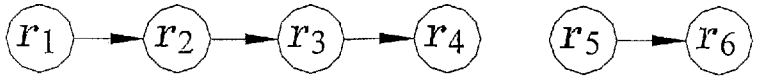

FIG. 10A

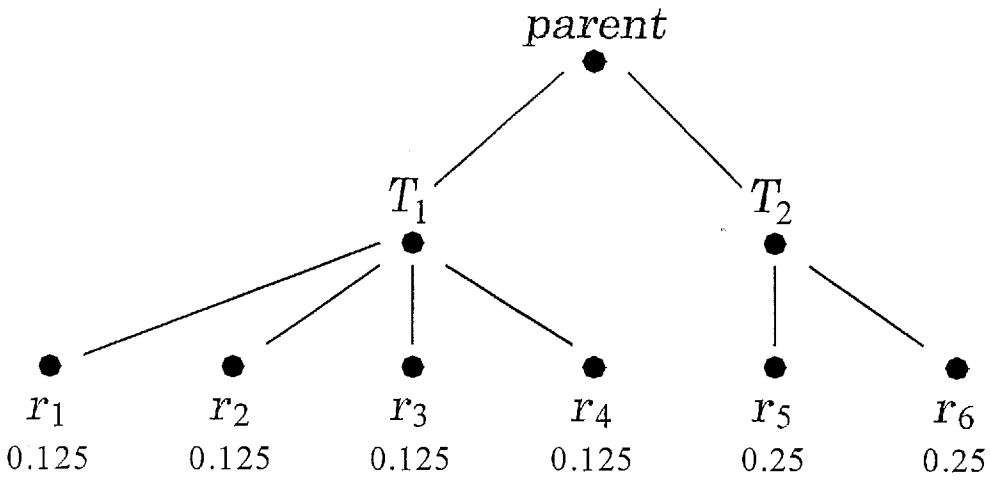

FIG. 10B

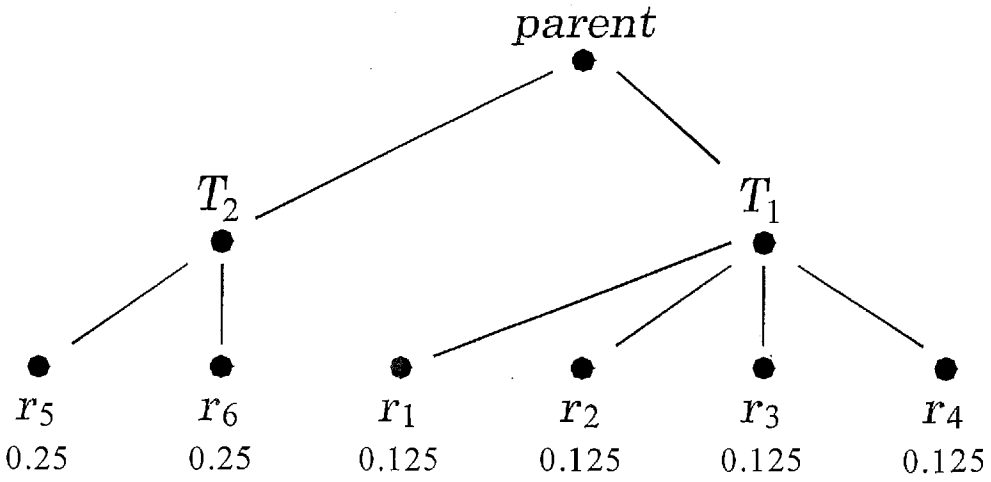

FIG. 10C

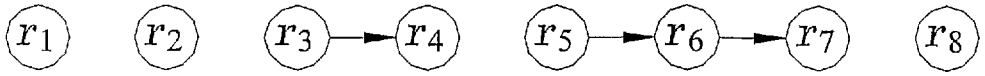

FIG. 11A

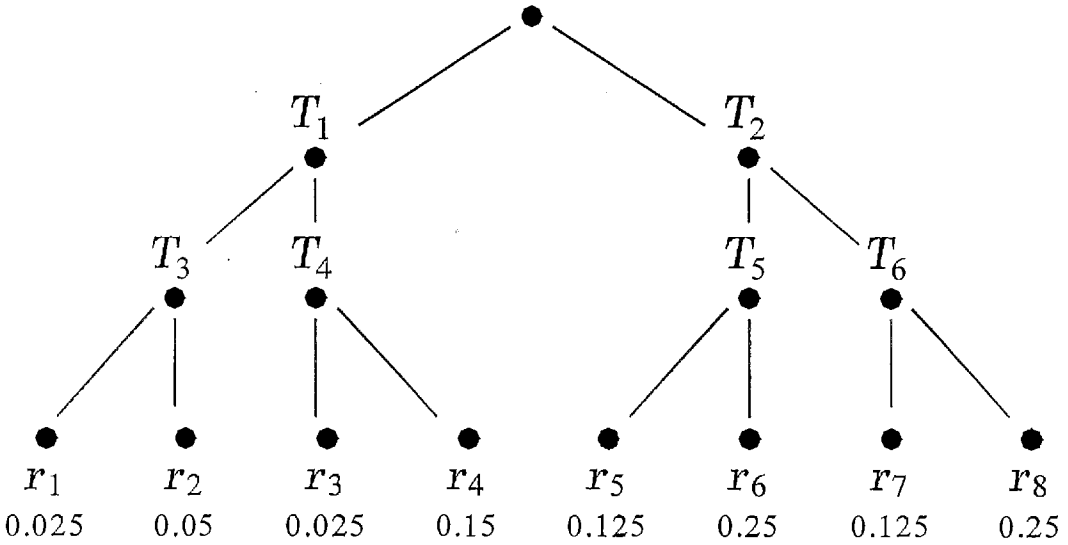

FIG. 11B

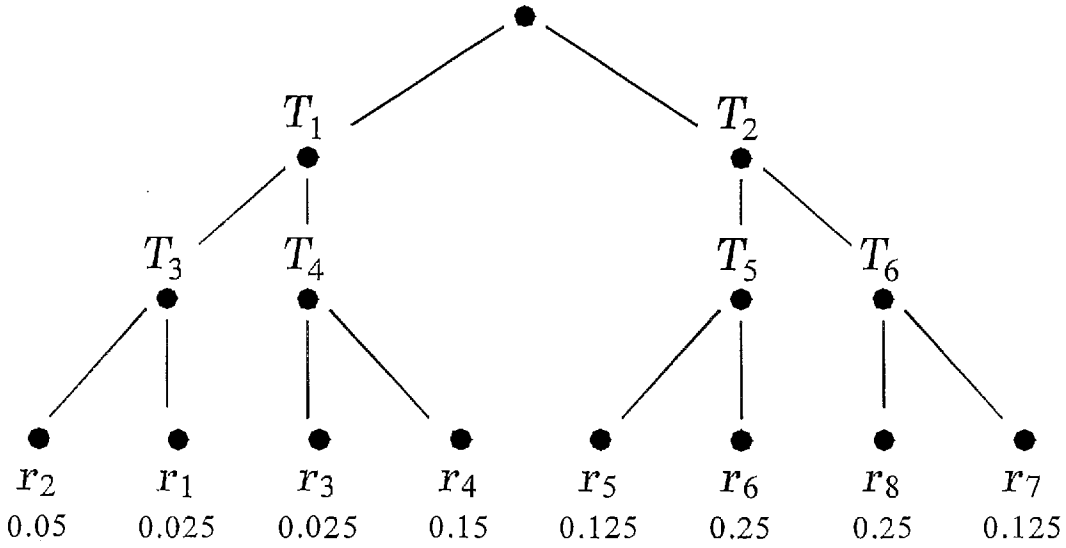

FIG. 11C

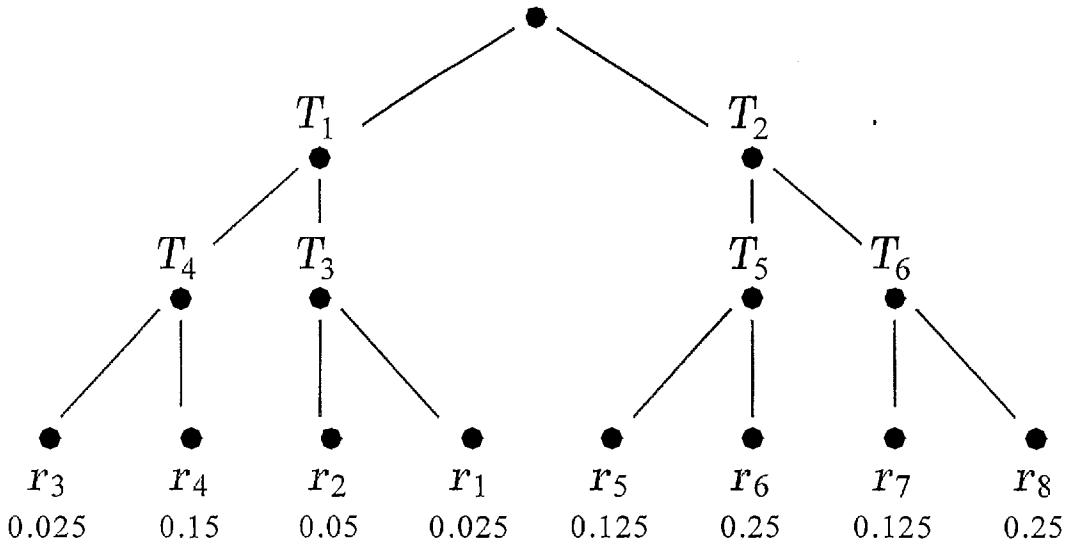

FIG. 11D

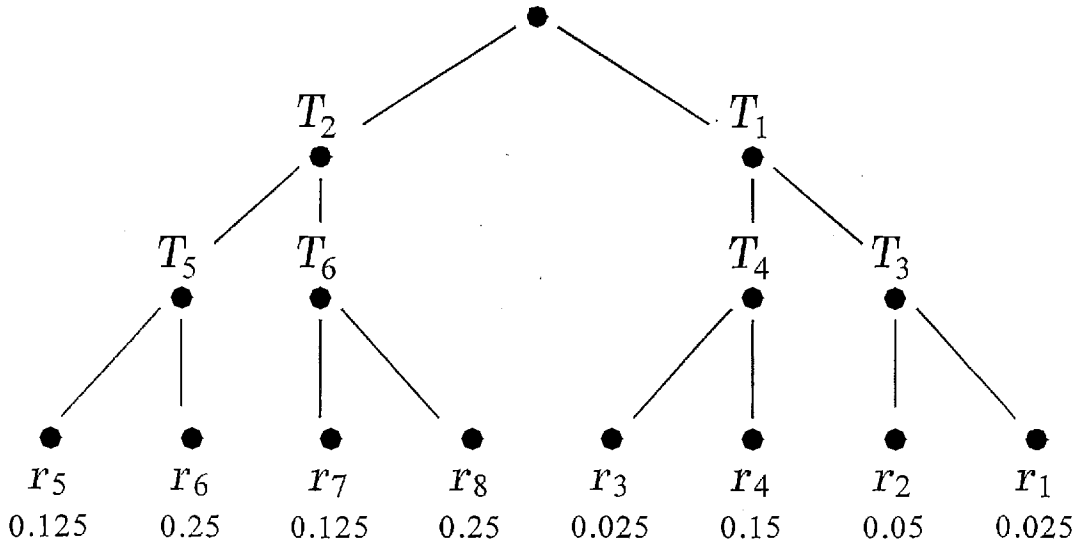

FIG. 11E

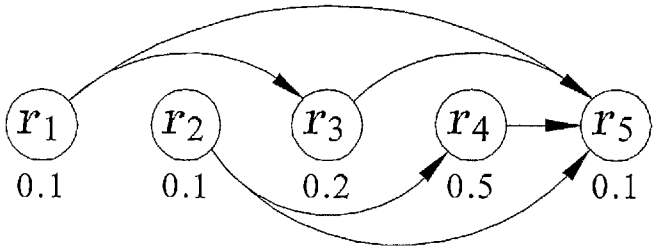

FIG. 12A

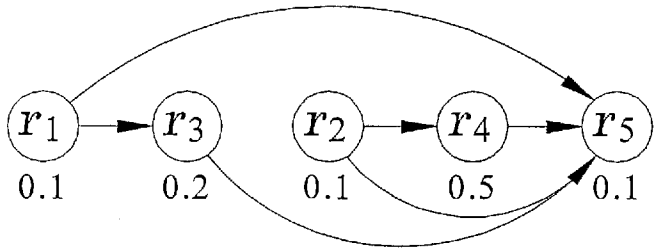

FIG. 12B

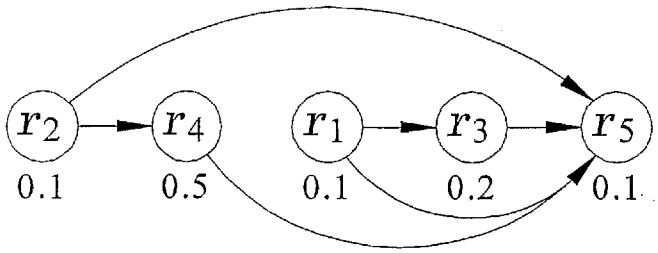

FIG. 12C

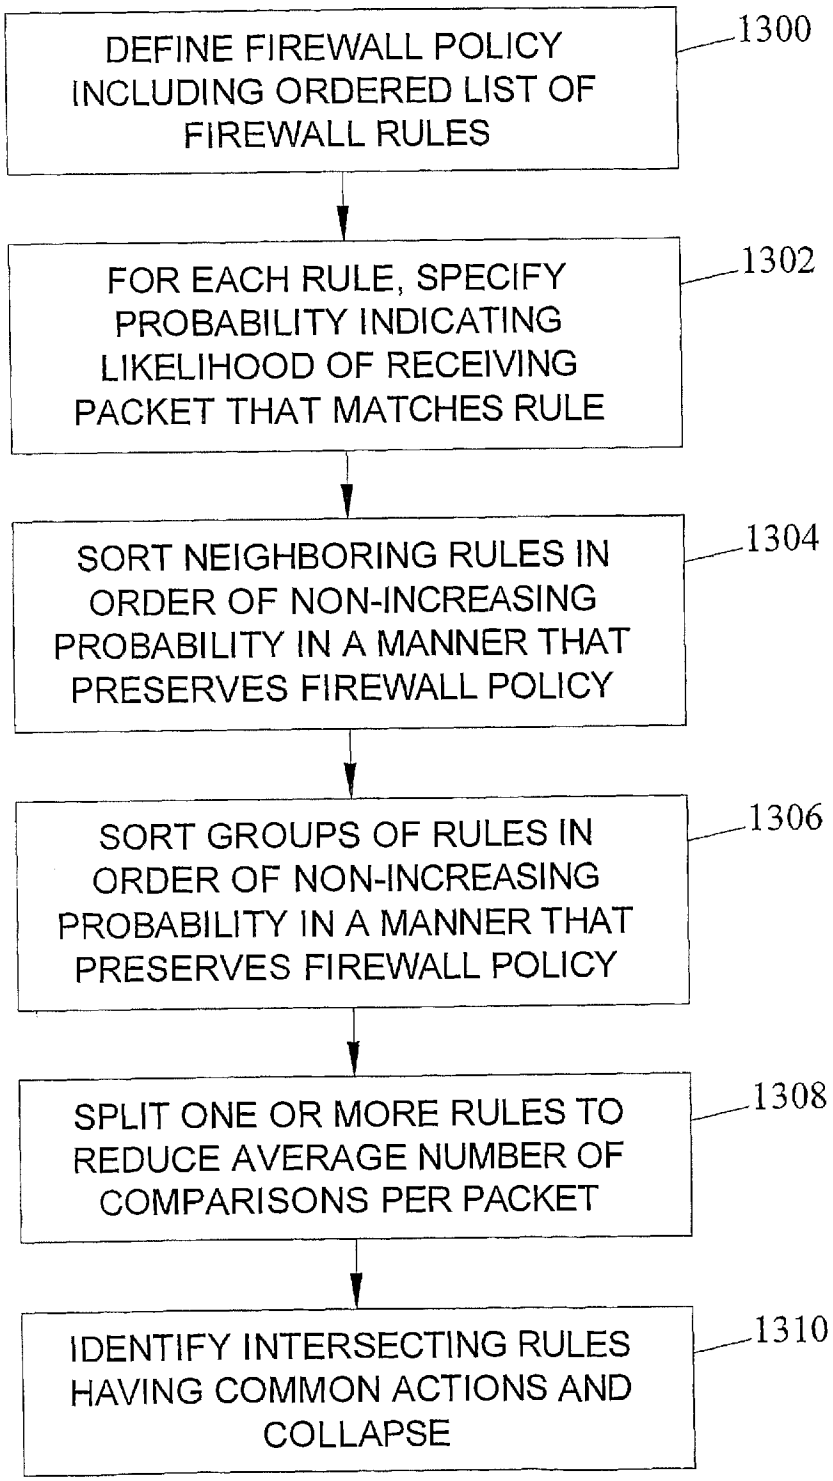
DEFINE FIREWALL POLICY INCLUDING ORDERED LIST OF FIREWALL RULES
1300
FOR EACH RULE, SPECIFY PROBABILITY INDICATING LIKELIHOOD OF RECEIVING PACKET THAT MATCHES RULE
1302
SORT NEIGHBORING RULES IN ORDER OF NON-INCREASING PROBABILITY IN A MANNER THAT PRESERVES FIREWALL POLICY
1304
SORT GROUPS OF RULES IN ORDER OF NON-INCREASING PROBABILITY IN A MANNER THAT PRESERVES FIREWALL POLICY
1306
SPLIT ONE OR MORE RULES TO REDUCE AVERAGE NUMBER OF COMPARISONS PER PACKET
1308
IDENTIFY INTERSECTING RULES HAVING COMMON ACTIONS AND COLLAPSE
1310

FIG. 13

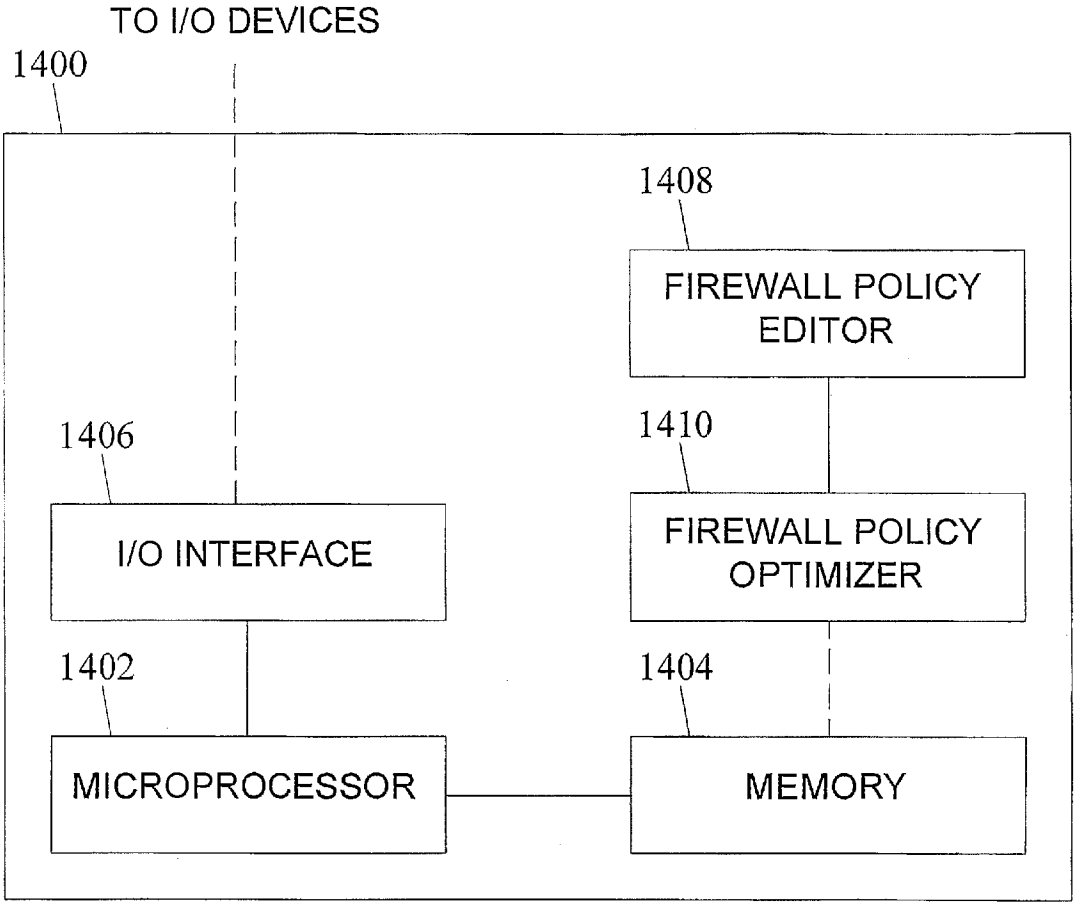
TO I/O DEVICES
1400
1408
FIREWALL POLICY EDITOR
1406
1410
I/O INTERFACE
FIREWALL POLICY OPTIMIZER
1402
1404
MICROPROCESSOR
MEMORY

FIG. 14

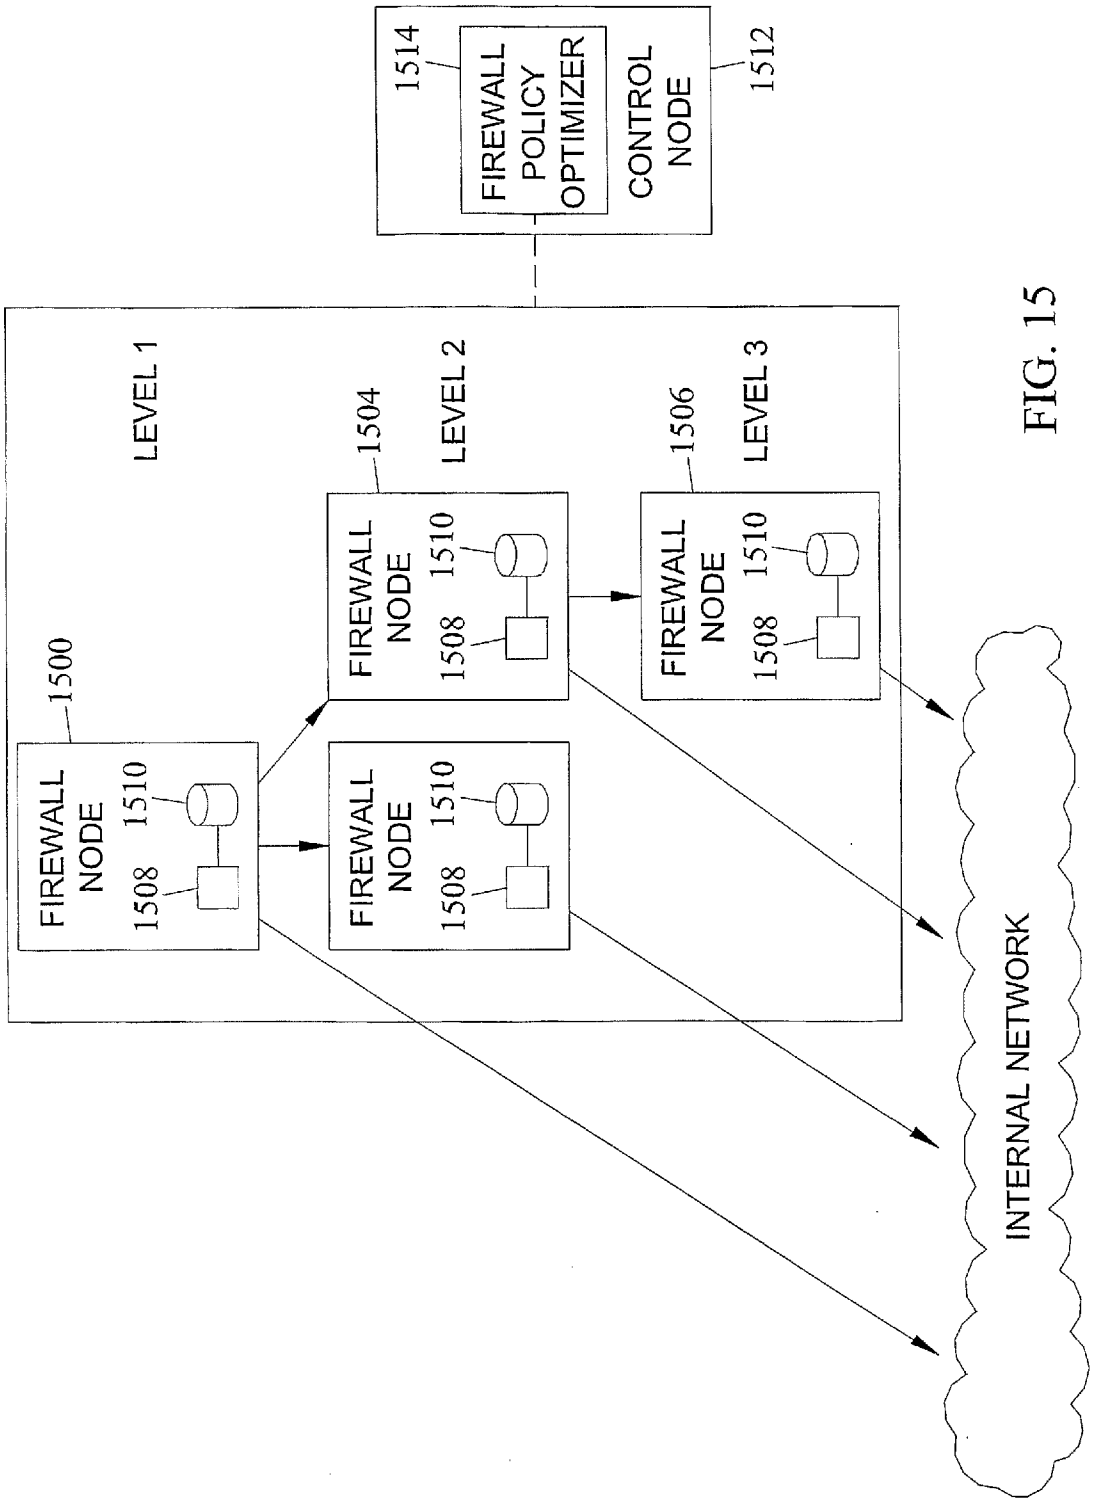
1514
FIREWALL POLICY OPTIMIZER
CONTROL NODE
1512
LEVEL 1
LEVEL 2
LEVEL 3
1500
1504
1506
FIREWALL NODE
1510
1508
INTERNAL NETWORK

FIG. 15

Espacenet

Bibliographic data: EP1864226 (B1) — 2013-05-15

METHODS, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR NETWORK FIREWALL POLICY OPTIMIZATION

Inventor(s): FULP ERRIN W [US]; TARSA STEPHEN J [US] ± (FULP, ERRIN, W, ; TARSA, STEPHEN, J)

Applicant(s): UNIV WAKE FOREST [US] ± (WAKE FOREST UNIVERSITY)

Classification: **- international:** ***H04L29/06***
- cooperative: **H04L63/0263**; **H04L63/20**

Application number: EP20060758213 20060328

Priority number (s): WO2006US11291 20060328 ; US20050665664P 20050328

Also published as: EP1864226 (A2) EP1864226 (A4) WO2006105093 (A2) WO2006105093 (A3) US2006248580 (A1) more

Abstract not available for EP1864226 (B1)
Abstract of corresponding document: WO2006105093 (A2)

Methods, systems, and computer program products for firewall policy optimization are disclosed. According to one method, a firewall policy including an ordered list of firewall rules is defined. For each rule, a probability indicating a likelihood of receiving a packet matching the rule is determined. The rules are sorted in order of non-increasing probability in a manner that preserves the firewall policy.

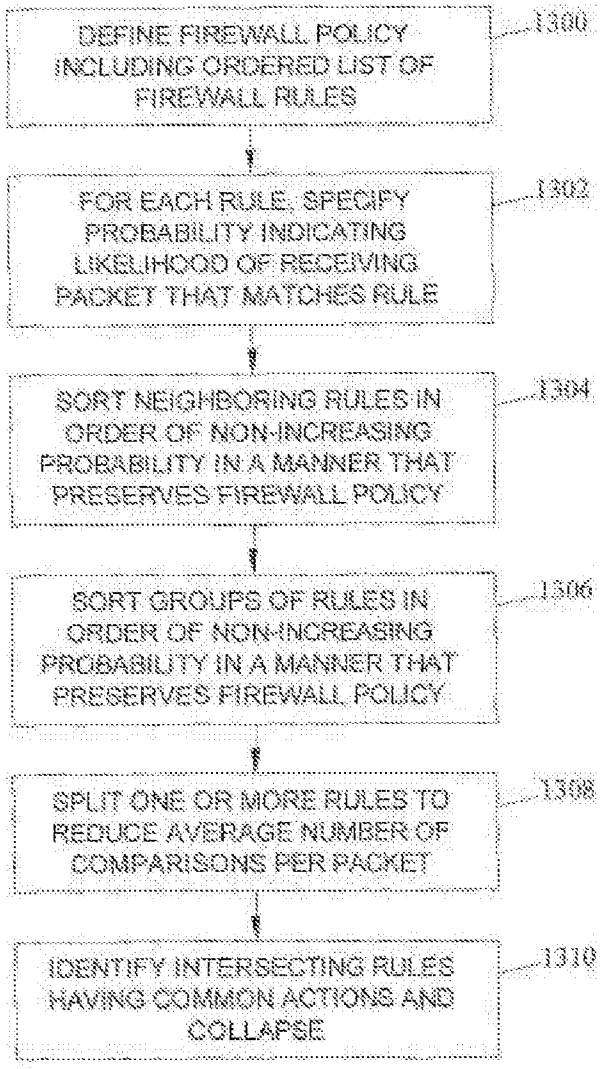
DEFINE FIREWALL POLICY INCLUDING ORDERED LIST OF FIREWALL RULES
1300
FOR EACH RULE, SPECIFY PROBABILITY INDICATING LIKELIHOOD OF RECEIVING PACKET THAT MATCHES RULE
1302
SORT NEIGHBORING RULES IN ORDER OF NON-INCREASING PROBABILITY IN A MANNER THAT PRESERVES FIREWALL POLICY
1304
SORT GROUPS OF RULES IN ORDER OF NON-INCREASING PROBABILITY IN A MANNER THAT PRESERVES FIREWALL POLICY
1306
SPLIT ONE OR MORE RULES TO REDUCE AVERAGE NUMBER OF COMPARISONS PER PACKET
1308
IDENTIFY INTERSECTING RULES HAVING COMMON ACTIONS AND COLLAPSE
1310

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 864 226 B1

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:
15.05.2013 Bulletin 2013/20

(21) Application number: **06758213.0**

(22) Date of filing: **28.03.2006**

(51) Int Cl.:
H04L 29/06 *(2006.01)*

(86) International application number:
PCT/US2006/011291

(87) International publication number:
WO 2006/105093 (05.10.2006 Gazette 2006/40)

(54) **METHODS, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR NETWORK FIREWALL POLICY OPTIMIZATION**

VERFAHREN, SYSTEME UND COMPUTERPROGRAMM ZUR OPTIMIERUNG DER FIREWALL-REGELUNG EINES NETZWERKES

PROCEDES, SYSTEMES ET PRODUITS DE PROGRAMMES INFORMATIQUES POUR L'OPTIMISATION DE LA POLITIQUE DE PARE-FEU D'UN RESEAU

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IS IT LI LT LU LV MC NL PL PT RO SE SI SK TR

(30) Priority: **28.03.2005 US 665664 P**

(43) Date of publication of application:
12.12.2007 Bulletin 2007/50

(73) Proprietor: **Wake Forest University**
Winston Salem, NC 27101 (US)

(72) Inventors:
- **FULP, Errin, W.**
  Winston-Salem, North Carolina 27104 (US)
- **TARSA, Stephen, J.**
  Duxbury, Massachusetts 02332-3503 (US)

(74) Representative: **Isarpatent**
Patent- und Rechtsanwälte
Postfach 44 01 51
80750 München (DE)

(56) References cited:
EP-A2- 1 006 701
US-A1- 2002 198 981
US-B1- 6 662 235
US-B2- 7 237 267
WO-A2-2006/093557
US-B1- 6 484 261
US-B1- 7 089 581

- **Errin W. Fulp: "Optimization of Network Firewall Policies Using Ordered Sets and Directed Acyclical Graphs", Wake Forest University , 2004, pages 1-12, XP002618343, Retrieved from the Internet: URL:www.cs.wfu.edu/~fulp/Papers/ewflist.ps [retrieved on 2011-01-24]**
- **Errin W. Fulp and Stephen J. Tarsa: "Network Firewall Policy Tries", Wake Forest University , 2004, pages 1-17, XP002618344, Retrieved from the Internet: URL:http://citeseerx.ist.psu.edu/viewdoc/s ummary?doi=10.1.1.143.5727 [retrieved on 2011-01-24]**
- **Lili Qiu, George Varghese, Subhash Suri: "Fast Firewall Implementations for Software and Hardware-based Routers", , 2001, pages 241-250, XP002618345, Retrieved from the Internet: URL: http://www.ieee-icnp.org/2001/papers/2001-26.pdf [retrieved on 2011-01-25]**
- **Errin W. Fulp: "Errin Fulp", , 2010, XP002618346, Retrieved from the Internet: URL: www.cs.wfu.edu/~fulp/ewfPub.html [retrieved on 2011-01-24]**

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 1 864 226 B1

Printed by Jouve, 75001 PARIS (FR)

Description

TECHNICAL FIELD

[0001] The subject matter described herein relates to optimizing polices for network firewalls. More particularly, the subject matter described herein relates to methods, systems, and computer program products for optimizing performance of network firewalls.

BACKGROUND ART

[0002] Network firewalls remain the forefront defense for most computer systems. Guided by a security policy, these devices provide access control, auditing, and traffic control [3, 30, 31]. A security policy is a list of ordered rules, as seen in Table 1, that defines the action to perform on matching packets. For example, an accept action passes the packet into or from the secure network, while deny causes the packet to be discarded. In many implementations, the rule set is stored internally as a linked list. A packet is sequentially compared to the rules, starting with the first, until a match is found; otherwise, a default action is performed [30, 31]. This is referred to as a first- match policy and is used in many firewall systems including the Linux firewall implementation iptables [25].

Table 1: Example Security Policy Consisting of Multiple Ordered Rules

		Source		Destination			
No.	Proto.	IP	Port	IP	Port	Action	Prob.
1	UDP	1.1.*	*	*	80	deny	0.01
2	TCP	2.*	*	1.*	90	accept	0.02
3	UDP	*	*	1.*	*	accept	0.10
4	TCP	2.*	*	1.*	20	accept	0.17
5	UDP	1.*	*	*	*	accept	0.20
6	*	*	*	*	*	deny	0.50

[0003] Traditional firewall implementations consist of a single, dedicated machine, similar to a router, that sequentially applies rules to each arriving packet. However, packet filtering represents a significantly higher processing load than routing decisions [24, 29, 31]. For example, a firewall that interconnects two 100 Mbps networks would have to process over 300,000 packets per second [30]. Successfully handling these traffic loads becomes more difficult as rule sets become more complex [4, 22, 31]. Furthermore, firewalls must be capable of processing even more packets as interface speeds increase. In a high-speed environment (e.g. Gigabit Ethernet), a single firewall can easily become a bottleneck and is susceptible to DoS attacks [4, 9, 13, 14]. An attacker could simply inundate the firewall with traffic, delaying or preventing legitimate packets from being processed.

[0004] One approach to increase firewall performance focuses on improving hardware design. Current research is investigating different distributed firewall designs to reduce processing delay [4, 9, 22], and possibly provide service differentiation [11]. Another approach focuses on improving performance via better firewall software [6, 7, 12, 16, 17, 24]. Similar to approaches that address the longest matching prefix problem for packet classification [8,10,25, 28], solutions typically represent the firewall rule set in different fashions (e.g. tree structures) to improve performance. While both approaches, together or separate, show great promise, each requires radical changes to the firewall system, and therefore are not amenable to current or legacy systems.

[0005] EP Patent Publication No. 1006701 discloses a packet data filter which stores ordered rules and sequentially applies the rules to received data packets to determine the disposition of the data packet. The packet filter maintains a match count in memory which indicates the number of times each rule matched an incoming data packet. Periodically, at the initiation of a user, or based on operating parameters of the filter, the rules are automatically re-ordered based on the match count. As a result of the re-ordering, rules with higher match counts are moved earlier in the sequential evaluation order and rules with lower match counts are moved later in the sequential evaluation order. As such, rules which are more likely to match incoming data packets are evaluated earlier, thus avoiding the evaluation of later rules.; In order to prevent a re-ordering which would change the overall security policy of the packet filter, pairs of rules are compared to determine if they conflict (i.e., the swapping of the two rules would result in a change in the overall security policy). During re-ordering, the swapping of conflicting rules is prevented.

[0006] US Patent Publication No. 2002/198981 discloses a method and system for determining whether to enforce a

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1018

plurality of filter rules for a packet including a key in a computer network is disclosed. Each of the plurality of filter rules has a priority. The method and system include accumulating statistics for each of the plurality of filter rules. The statistics indicate a frequency of enforcement for each of the plurality of filter rules. The method and system also include placing the plurality of filter rules in an order for testing against the key. The order is based on the frequency of each filter rule of the portion of the plurality of filter rules. Consequently, more frequently enforced filter rules may be tested first.

[0007] WO Publication No. 2006093557 discloses methods, systems, and computer program products for providing function-paralell firewalls are disclosed. According to one aspect, a function-parallel firewall includes a first firewall node for filtering received packets using a first portion of a rule set including a plurality of rules. The first portion includes less than all of the rules in the rule set. At least one second firewall node filters packets using a second portion of the rule set. The second portion includes at least one rule in the rule set that is not present in the first portion. The first and second portions together include all of the rules in the rule set.

[0008] Accordingly, there exists a need for improved methods, systems, and computer program products for network firewall policy optimization.

SUMMARY

[0009] The subject matter described herein includes methods, systems, and computer program products for network firewall policy optimization. According to one method, a firewall policy including an ordered list of firewall rules is defined. For each rule, a probability indicating a likelihood of receiving a packet matching the rule is defined. Neighbouring rules are sorted in order of non- increasing probability in a manner that preserves the firewall policy.

[0010] As used herein, the term "firewall" refers to a logical entity that is adapted to filter packets at the ingress and/or egress portions of a network based on a policy. The term "firewall" is intended to include systems that block packets from entering or leaving a network and systems that perform intrusion detection and intrusion protection functions for packets entering or leaving a network. Thus, the methods and systems described herein for firewall policy optimization can be used to optimize policies for intrusion detection and intrusion protection firewalls.

[0011] The subject matter described herein includes a method to improve firewall performance and lower packet delay that can be applied to both legacy and current systems. In one implementation, a firewall rule set may be reordered to minimize the average number of rule comparisons to determine the action, while maintaining the integrity of the original policy. Integrity is preserved if the reordered and original rules always arrive at the same result. To maintain integrity, a firewall rule set may be modeled as a Directed Acyclical Graph (DAG), where vertices are firewall rules and edges indicate precedence relationships. Given this representation, any linear arrangement of the DAG is proven to maintain the original policy integrity. Unfortunately, determining the optimal rule order from all the possible linear arrangements is proven to be NP- complete, since it is equivalent to sequencing jobs with precedence constraints for a single machine [15]. Although determining the optimal order is NP- complete, a heuristic will be described below to order firewall rules that reduces the average number of comparisons while maintaining integrity. Simulation results show the proposed reordering method yields rule orders that are comparable to optimal; thus, provides a simple means to significantly improve firewall performance and lower packet delay.

[0012] The subject matter described herein for network firewall policy optimization may be implemented using any combination of hardware, software, or firmware. For example, the subject matter described herein may be implemented using a computer program product comprising computer executable instructions embodied in a computer readable medium. Exemplary computer readable media suitable for implementing the subject matter described herein include disk memory devices, chip memory devices, programmable logic devices, application specific integrated circuits, and downloadable electrical signals. In addition, a computer program product that implements the subject matter described herein may be implemented using a single device or computing platform or may be distributed across multiple devices or computing platforms.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] Preferred embodiments of the subject matter described herein will now be explained with reference to the accompanying drawings of which:

Figure 1A is a block diagram of a rule list Directed Acyclical Graph (DAG) for the firewall rules in Table 1;
Figure 1 B is a linear arrangement of firewall rules corresponding to the original rule order;
Figure 2A is a policy DAG representing a reordered version of the firewall rule set in Figure 1 after sorting;
Figure 2B is a policy DAG for the rule set illustrated in Table 1 illustrating an optimal rule order;
Figure 3A is a graph illustrating average numbers of packet comparisons versus intersection percentage for different orderings of a firewall policy;
Figure 3B is a graph illustrating the percent difference in average numbers of packet comparisons between sorted

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1019

and optimal sortings of a firewall policy;
Figure 4A is a graph of percent difference comparisons between original and sorted versions of a firewall policy;
Figure 4B is a graph illustrating the performance impact of considering rule intersection and actions in performing firewall rule sorting;
Figure 5A is a table illustrating a firewall policy before sorting;
Figure 5B is a table illustrating a sorted version of the firewall policy of Figure 5A;
Figure 6A is a policy trie representation of the firewall policy of Figure 5B;
Figure 6B is a table representing the firewall policy of Figure 5A after TCP and UDP sub-trie rotation;
Figure 7A is a policy trie illustrating the splitting of rule r_5 into rules r_{5a} and r_{5b};
Figure 7B is a table illustrating the firewall policy of Figure 5A after splitting rule r_5 into rules r_{5a} and r_{5b};
Figure 8A is a graph illustrating firewall performance with and without rule splitting;
Figure 8B is a graph illustrating locations and probabilities where splitting is better;
Figure 9A is a policy DAG of a firewall rule set;
Figure 9B is a policy DAG illustrating compression of rules r_1 and r_5;
Figure 10A is a policy DAG of a firewall rule list;
Figure 10B is an original policy trie before sorting;
Figure 10C is a policy trie illustrating the firewall rules after rotating sub-tries T_1 and T_2;
Figure 11A is a policy DAG illustrating a firewall policy;
Figure 11B is a policy trie illustrating firewall rules of the policy before sorting;
Figure 11C is a policy trie illustrating the firewall policy after exchanging nodes r_1 and r_2 and nodes r_7 and r_8;
Figure 11 D is a policy trie illustrating the firewall policy after exchanging sub-tries T_3 and T_4;
Figure 11 E is a policy trie illustrating the firewall policy after exchanging sub-tries T_1 and T_2;
Figure 12A is a policy DAG illustrating an original order for a firewall policy;
Figure 12B is a policy DAG illustrating a sorted version of the firewall policy of Figure 12A after exchanging neighbors;
Figure 12C is a policy DAG illustrating a sorted version of the firewall policy after exchanging lists L_1 including rules r_1 and r_3 and lists L_2 including rules r_2 and r_4;
Figure 13 is a flow chart illustrating exemplary steps for network firewall policy optimization according to an embodiment of the subject matter described herein;
Figure 14 is a block diagram illustrating exemplary components of a system for firewall policy optimization according to an embodiment of the subject matter described herein; and
Figure 15 is a block diagram illustrating a plurality of firewall nodes for implementing a firewall policy according to an embodiment of the subject matter described herein.

DETAILED DESCRIPTION OF THE INVENTION

Modeling Firewall Security Policies

[0014] As described above, a firewall rule set, also known as a firewall policy, is traditionally an ordered list of firewall rules. Firewall policy models have been the subject of recent research [1, 2, 16]; however, the primary purpose is anomaly detection and policy verification. In contrast, the policy model described in this herein is designed for firewall performance optimization and integrity. Firewall performance refers to reducing the average number of comparisons required to determine an action, while integrity refers to maintaining the original policy intent. Although improving the worst-case performance is important, it is not possible without changing the list-based representation [16, 24].

Firewall Rule and Policy Models

[0015] In the examples described herein, a rule r is modeled as an ordered tuple of sets, $r = (r[1], r[2], ..., r[k])$. Order is necessary among the tuples since comparing rules and packets requires the comparison of corresponding tuples. Each tuple $r[i]$ is a set that can be fully specified, specify a range, or contain wildcards '*' in standard prefix format. For the Internet, security rules are commonly represented as a 5-tuple consisting of: protocol type, IP source address, source port number, IP destination address, and destination port number [30, 31]. Given this model, the ordered tuples can be supersets and subsets of each other, which forms the basis of precedence relationships. In addition to the prefixes, each filter rule has an action, which is to accept or deny. However, the action will not be considered when comparing packets and rules. Similar to a rule, a packet (IP datagram) d can be viewed as an ordered k-tuple $d = (d[1], d[2], ..., d[k])$; however, ranges and wildcards are not possible for any packet tuple.

[0016] Using the previous rule definition, a standard security policy can be modeled as an ordered set (list) of n rules, denoted as $R = \{r_1, r_2, ..., r_n\}$. A packet d is sequentially compared against each rule r_i starting with the first, until a match is found ($d \Rightarrow r_i$) then the associated action is performed. A match is found between a packet and rule when every tuple

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1020

of the packet is a subset of the corresponding tuple in the rule.
Definition Packet d matches r_i if

$$d \Rightarrow r_i \quad \text{iff} \quad d[l] \subseteq r_i[l], \quad l = 1, \ldots, k$$

[0017] The rule list R is comprehensive if for every possible legal packet d a match is found using R. Furthermore, two rule lists R and R' are equivalent if for every possible legal packet d the same action is performed by the two rule lists. If R and R' are different (e.g. a reorder) and the lists are equivalent, then the policy integrity is maintained.

[0018] As previously mentioned, a rule list has an implied precedence relationship where certain rules must appear before others if the integrity of the policy is to be maintained. For example consider the rule list in Table 1. Rule r_1 must appear before rules r_3 and r_5, likewise rule r_6 must be the last rule in the policy. If for example, rule r_3 was moved to the beginning of the policy, then it will shadow [2] the original rule r_1. However, there is no precedence relationship between rules r_2 and r_4 given in Table 1. Therefore, the relative ordering of these two rules will not impact the policy integrity and can be changed to improve performance. The present example assumes that the original policy is free from any anomalies. Likewise, when a policy is reordered to improve performance it should not introduce any anomalies, which will occur if precedence relationships are not maintained. As a result, a model is needed to effectively represent precedence relationships.

Modeling Rule List Precedence Relationships

[0019] The precedence relationship between rules in a policy will be modeled as a Directed Acyclical Graph (DAG) [23, 18]. Such graphs have been successfully used to represent the relative order of individual tasks that must take place to complete a job (referred to as a task graph model). Since certain rules must appear before others to maintain policy integrity, this structure is well suited for modeling the precedence of firewall rules. Let $G = (R, E)$ be a rule list DAG for a rule list R, where vertices are rules and edges E are the precedence relationships (constraint). A precedence relationship, or edge, exists between rules r_i and r_j, if $i < j$, the actions for each rule are different, and the rules intersect.
Definition The intersection of rule r_i and r_j, denoted as $r_i \cap r_j$ is

$$r_i \cap r_j = (r_i[l] \cap r_j[l]), \quad l = 1, \ldots, k$$

Therefore, the intersection of two rules results in an ordered set of tuples that collectively describes the packets that match both rules. The rules r_i and r_j intersect if every tuple of the resulting operation is non-empty. In contrast, the rules r_i and r_j do not intersect, denoted as $r_i \not\cap r_j$, if at least one tuple is the empty set. Note the intersection operation is symmetric; therefore, if r_i intersects r_j then r_j will intersect r_i. The same is true for rules that do not intersect.

[0020] For example consider the rules given in Table 1, the intersection of r_1 and r_3 yields (UDP, 1.1.*, *, 1.*, 80). Again, the rule actions are not considered in the intersection or match operation. Since these two rules intersect, a packet can match both rules for example d = (UDP, 1.1.1.1, 80, 1.1.1.1, 80). Furthermore, the actions of the two rules are different. Therefore, the relative order must be maintained between these two rules and an edge drawn from r_1 to r_3 must be present in the rule list DAG, as seen in Figures 1A and 1B. More particularly, Figure 1 A illustrates a rule list DAG for the rules in Table 1. Figure 1B illustrates a linear arrangement of the rules in Figure 1A. In Figures 1A and 1 B, the vertices represent rules, the circles represents an accept rules, and the squares represents deny rules. Edges that connect the vertices represent precedence requirements. As can be seen from Figure 1 B, because of the edge between rules r_1 and r_3, precedence between rules r_1 and r_3 must be maintained. In contrast consider the intersection of rules r_3 and r_5. These two rules intersect, indicating packets belonging to the set (UDP, 1.*, *, 1.*, *) would match both rules. However, it does not matter which of the two rules a packet matches first, since the action is the same for both rules. Therefore, an edge does not exist between rules r_3 and r_5 in the diagram. Similarly, rules r_2 and r_4 do not intersect due to the fifth tuple (destination port). A packet cannot match both rules indicating the relative order can change; therefore, an edge will not exist between them.

[0021] The match operation can be used to identify precedence relationships, but it cannot do so in every case. Consider a partial-match example [1], where r_a = (UDP, *, 80, 10.*, 90, accept) and r_b = (UDP, 10.*, 80, *, 90, deny). The intersection of r_a and r_b is (UDP, 10.*, 80, 10.*, 90); therefore a packet, such as d = (UDP, 10.10.10.10, 80, 10.10.10.10, 90), can match both rules. If r_a appears before r_b then the packet d is accepted, but if r_b occurs before r_a then d is rejected. As a result, the order of r_a and r_b in the original policy must be maintained. However, the match operation is unable to identify the precedence in this example. A partial match exists in between rules r_3 and r_5 in Table

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1021

1, but as previously discussed an edge does not exist between the rules since the actions are the same.

[0022] Using the rule list DAG representation a linear arrangement is sought that improves the firewall performance. As depicted in Figure 1 B, a linear arrangement (permutation or topological sort) is a list of DAG vertices where all the successors of a vertex appear in sequence after that vertex [23]. Therefore it follows that a linear arrangement of a rule list DAG represents a rule order, if the vertices are read from left to right. Furthermore, it is proven in the following theorem that any linear arrangement of a rule list DAG maintains integrity: **Theorem** Any linear arrangement of a rule list DAG maintains integrity.

[0023] **Proof** Assume a rule list DAG G is constructed from the security policy R that is free of anomalies. Consider any two rules r_i and r_j in the policy, where $i < j$. If an edge between r_i and r_j in G does not exist, then a linear arrangement of G can interchange the order of the two rules. An edge will not exist if the rules do not intersect; however, a reorder will not affect integrity since a packet cannot match both rules. Shadowing is not introduced due to the reorder since the intersection operation is symmetric. An edge will not exist if the two rules intersect but have the same action; however, a reorder will not affect integrity since the same action will occur regardless of which rule is matched first. If an edge does exist between the rules, then their relative order will be maintained in every linear arrangement of *G*; thus maintaining precedence and integrity.

Rule List Optimization

[0024] As described above, it is important to inspect packets as quickly as possible given increasing network speeds and QoS requirements. Using the rule list DAG to maintain policy integrity, a linear arrangement is sought that minimizes the average number of comparisons required. However, this will require information not present in the firewall rule list. Certain firewall rules have a higher probability of matching a packet than others. As a result, it is possible to develop a policy profile over time that indicates frequency of rule matches (similar to cache hit ratio). Let $P = \{p_1, p_2, ..., p_n\}$ be the policy profile, where p_i is the probability that a packet will match rule i (first match in the policy). Furthermore, assume a packet will always find a match, $\sum_{i=1}^{n} p_i = 1$; therefore R is comprehensive. Using this information, the average number of rule comparisons required is

$$E[n] = \sum_{i=1}^{n} i \cdot p_i \qquad (1)$$

For example, the average number of comparisons required for the rule set in Table 1 is 5.03.

[0025] Given a rule list DAG $G = (R, E)$ and policy profile $P = \{p_1, p_2, ..., p_n\}$ a linear arrangement π of G is sought that minimizes Equation 1. In the absence of precedence relationships, the average number of comparisons is minimized if the rules are sorted in non-increasing order according to the probabilities [26], which is also referred to as Smith's algorithm [27]. Precedence constraints cause the problem to be more realistic; however, such constraints also make determining the optimal permutation more problematic.

[0026] Determining the optimal rule list permutation can be viewed as job scheduling for a single machine with precedence constraints [15, 21]. The notation for such scheduling problems is $\alpha|\beta|\gamma|\delta$, where α is the number of machines, β is the precedence (or absence of) which can be represented as a DAG, γ is a restriction on processing time, and δ is the optimality criterion [15]. Determining the optimal rule order is similar to the $1|\beta|1|\sum \omega C_i$ scheduling problem, or optimality criterion, where ω is a weight associated with a job (for example, importance) and C_i is the completion time. As previously noted, the $1||| \sum \omega C_i$ problem can be solved in linear time the using Smith's algorithm [27], which orders jobs according to non-decreasing $\frac{t_i}{\omega}$ ratio, where t_i is the processing time of job i. In this case set $t_i = 1$ and $\omega = p_i \forall_i$. However, Lawler [19] and Lenstra et al. [21] proved $1|\beta|1| \sum \omega C_i$ to be *NP*-complete via the linear arrangement problem, which implies determining the optimal firewall rule order is also *NP*-complete. Note, determining the number of possible permutations has been proven to be *NP*-hard [5].

Theorem $1|\beta|1| \sum \omega C_i$ α *Determining the optimal order of a firewall rule list*

[0027] **Proof** Consider the $1|\beta|1| \sum \omega C_i$ problem. Each of n jobs J_i, $i \in I$, has to be processed without preemption on a single machine that can handle at most one job at a time. For each $i \in I$, let ω be the associated weight. Furthermore, let $G = (V, E)$ be a DAG that represents the precedence order of the jobs J_i. Assume the processing time of each job equals 1 time unit, the weights to be $0 \leq \omega \leq 1$ such that $\sum \omega = 1$, and β, which is G, to be a rule list DAG. In this case, the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1022

optimization criterion $\sum \omega \cdot C_i$ is the same as $\sum p_i \cdot i$, which is given in equation 1. Clearly, the optimal firewall rule ordering problem has a solution if and only if $1|\beta|1|\sum \omega C_i$ has a solution. Therefore, determining the optimal permutation of firewall rules is *NP*-complete.

Exemplary Rule Sorting Algorithm

[0028] Although determining the optimal rule permutation was proven to be *NP*-complete, reducing the average number of comparisons required to process a packet remains an important objective. As previously discussed, a sorting algorithm must maintain the precedence relationships among rules. Of course an exhaustive search is possible if the number of rules is small (generate and test all possible topological orders); however as proven in the previous section, this is not feasible with a realistic rule list.

[0029] One exemplary algorithm starts with the original the rule set, then sorts neighboring rules based on nonincreasing probabilities. However, an exchange of neighbors should never occur if the rules intersect and have different actions. This test preserves any precedence relationships in the policy. For example, the following sorting algorithm uses such a comparison to determine if neighboring rules should be exchanged. Note a_i denotes the action associated with the i^{th} rule.

```
done = false
while(!done)
    done = true
    for(i = 1; i <n; i++)
        if(pi < pi+1 AND (ri ⋔ ri+1 OR ai == ai+1))then
            exchange rules, actions, and probabilities
            done = false
        endif
    endfor
endwhile
```

[0030] Other sorting algorithms are possible if the completion time of the sort is an issue; however, the method presented is easy to implement and only requires a simple neighbor comparison. Assume the match probabilities for the rule list given in Table 1. Applying the sorting algorithm to this rule list results in the ordering depicted in Figure 2A, which has 10% fewer comparisons on average. However when using the algorithm, it is possible that one rule can prevent another rule from being reordered. For example, rule r_1 prevents rule r_5 from being placed closer to the beginning of the rule set. However, if both rules r_1 and r_5 are placed closer to the beginning of the policy while maintaining their relative order, the average number of comparison will be further reduced, 15% fewer. This is the optimal order for the 6 rule set, which is depicted in Figure 2B. Although this simple sorting algorithm is unable to move groups of rules, it can still improve the performance of the firewall system. Its effectiveness is measured experimentally in the next section.

Experimental Results

[0031] In this section, the average number of rule comparisons required after sorting the rules is compared with average number of comparisons required using the original order and the optimal order. Note the optimal rule ordering was determined via an exhaustive search. As a result, it was not feasible to determine the optimal ordering once the number of rules equaled 15 given the number of permutations to consider. Two different variations of the rule sorting method described in the preceding section were implemented. The difference between the methods was the comparison used to determine if neighboring rules should be exchanged. The first sorting algorithm exchanged neighboring rules if they were out of order and did not intersect. The if-condition was as follows.

if ($p_i < p_{i+1}$ **AND** $r_i \pitchfork r_{i+1}$) **then**

Therefore, rule actions were not considered and the method will be referred to as *non-action sort.* The second sorting method did consider the rule action (as described in the preceding section) and will be referred to as *action sort.* The comparison between the two sorting algorithms will indicate the importance of considering rule actions when ordering rules.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1023

[0032] In the first experiment, lists of 10 firewall rules were generated with random precedence relationships. The match probability of each rule was given by a Zipf distribution [20], which assigns probabilities according to the rank of an item. For this simulation, the last rule had the highest probability which is consistent with most policies (last rules are more general). As a result, the original order yields the worst average number of comparisons. The intersection percentage measured the percentage of rules that intersect in the policy. This metric gives a value for the dependency level in the policy; however note, rule action is not considered when calculating this metric. Rules were equally likely to have an accept or deny action and the results (average number of rule comparisons) for a particular intersection percentage were averaged over 1000 simulations.

[0033] Results of the first experiment are given in Figures 3A and 3B. More particularly, Figure 3A is a graph illustrating the average number of packet comparisons versus intersection percentage for the original rule sorting, the non-action sort, the action sort, and the optimal sort. Figure 3B is a graph illustrating the percent difference in the average number of comparisons required between the sorted and optimal firewall policy configurations. The average number of comparisons required was lower for the sorted and the optimal lists when the intersection percentage was low, as seen in Figure 3A. This was expected since there is a large number of possible rule permutations (few edges in the DAG). When the intersection percentage approached 100%, the values converged to the number required for the original list. This is due to the limited number of possible rule orders, only one order in the extreme case. The percent difference between the sorted and optimal order is shown in Figure 3B. At zero intersection percentage the sorted and optimal orders are equal, since any ordering is possible. Similarly at an intersection percentage of 100, the sorted and optimal orders are equal since only one order is possible. Between these two extremes, the action sorting algorithm remains close to the optimal value. In contrast, the non-action sort had a maximum difference of 33%.

[0034] The benefit of rule sorting on larger policies is also of interest; however as previously described, it was not possible to determine the optimal ordering once the number of rules approaches 15. Therefore, the second experiment used larger rule sets, but only compared the original order and the sorted orders. The number of rules ranged from 10 to 1000, while the matching probabilities and intersection percentages were the same as the previous experiment. The results of this experiment are depicted in Figure 4A. The sorted rule sets always performed equal to or better than the original order, while the action sort consistently performed better than the non-action sort. As noted in the previous experiment, the percent difference is very large given few intersections (e.g. 80% decrease for 1000 rules with 0% dependency), but approaches zero as all the rules intersect. As the number of rules increases, sorting is increasingly beneficial, although only at low intersection percentages. The benefit of sorting drops more significantly as the intersection percentage increases with larger rule sets. This is primarily due to the low matching probabilities of each rule, which requires a complete reordering to have a significant impact on the average number of comparisons. As a result, large rule sets can benefit from sorting if the intersection percentage is low.

[0035] The impact of considering rule actions when sorting is illustrated in Figure 4B. In this experiment lists of 1000 rules were generated, where the intersection percentages ranged from 0% to 100% and the percentage of rules with the same action varied from 50% to 100%. The performance of the action sort to the original list was then compared. As the percentage of rules with the same action increased, the percent difference (reduction) in the average number of comparisons also increased. This is due to the increased number of permutations possible when the rule actions are increasingly the same (fewer edges in the rule list DAG to consider). This is depicted in Figure 4B where a policy with a 100% intersection percentage can significantly reduce the number of comparisons (80% reduction) if all the rules have the same action. This performance increase is not possible with the non-action sort. Therefore, considering the rule action increases the number of possible rule orders, thereby providing more possibilities to improve firewall performance.

Conclusions

[0036] Network firewalls enforce a security policy by comparing arriving packets to a list of rules. An action, such as accept or deny, is then performed on the packet based on the matching rule. Unfortunately packet filtering can impose significant delays on traffic due to the complexity and size of rule sets. Therefore, improving firewall performance is important, given network Quality of Service (QoS) requirements and increasing network traffic loads.

[0037] The sections above describe rule ordering methods to improve the performance of network firewalls. Assuming each rule has a probability of a packet matching, firewall rules should be sorted such that the matching probabilities are non-increasing. This reduces the average number of comparisons and the delay across the firewall. However, a simple sort is not possible given precedence relations across rules. It is common in a security policy that two rules may match the same packet yet have different actions. It is this precedence relationship between rules that must be maintained to preserve integrity. The method described above uses Directed Acyclical Graphs (DAG) to represent the precedence order rules must maintain. Given this representation, a topological sort can be used to determine the optimal order (minimum average number of comparisons); however, the examples above prove this problem to be *NP*-complete (similar to job scheduling for a single nonpreemptive machine with precedence constraints). As an alternative, a simple sorting method was introduced that maintained the precedence order of the rules. Simulation results indicate this method can

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1024

significantly reduce the average number of comparisons required and is comparable to optimal ordering.

[0038] Several areas exist for future research in optimizing firewall rule lists. The sorting method proposed above is based on a simple algorithm. Although it can offer an improvement over the original rule order, algorithms that can move groups of rules could provide a larger reduction in the average number of comparisons. The effect of stateful firewalls should also be addressed in future research. Security can be enhanced with connection state and packet audit information. For example, a table can be used to record the state of each connection, which is useful for preventing certain types of attacks (e.g., TCP SYN flood) [30, 31]. The impact of such rules on the firewall needs to be investigated and whether sorting can be done on-line to reflect temporal changes. In addition, more research is needed to determine more accurate probability distributions for packet matching and dependency percentages. Given this information, better algorithms can be designed and more realistic simulations can be performed.

Trie-Based Network Firewall Optimization Methods and Rule Splitting Overview

[0039] As described above, while assuring firewall policy integrity is critical, performance is equally important given increasing network speeds and traffic volumes. Firewall processing delay can be reduced via hardware as well as policy optimization, where rules are manipulated to reduce the number of comparisons. For a given firewall installation, it can be determined that certain security rules have a higher probability of matching a packet than others. As a result, it is possible to develop a *policy profile* over time that indicates frequency of rule matches (similar to cache hit ratio). Given this information, trie-based methods are described herein to minimize the average number of comparisons required per packet. Although most firewall systems still utilize an ordered list representation, the proposed enhancements still are applicable to current firewall systems since a policy trie can be converted into an ordered list using an inorder traversal.

Rule Sorting

[0040] As described above, rule sorting of neighboring rules based on ascending probabilities in a manner that considers rules intersection and rule actions can be used to improve firewall performance by reducing the number of packet comparisons. The methods described above use Directed Acyclic Goal Graphs (DAGs) to represent firewalls. The methods described in this section will use the rule sorting algorithms in combination with policy trie representations to perform sorting of groups of rules in a manner that improves firewall performance.

[0041] In this example, it is assumed that a security policy contains a list of n ordered rules, $\{r_1, r_2, ..., r_n\}$. In addition, let p_i represent the probability that a packet will match rule i (first match in the policy). Therefore, the policy is comprehensive, which means every packet will have a match, if $\sum_{i=1}^{n} p_i = 1$. Using this information, the average number of rule comparisons required can be calculated as

$$E[n] = \sum_{i=1}^{n} i \cdot p_i$$

The average number of comparisons is minimized if the rules are sorted in non-ascending order according to the probabilities. However, this can only be achieved if the sorting algorithm considers the probabilities and whether one rule is a superset of another. For example, sorting algorithms must use the following comparison to determine if neighboring rules should be interchanged Note, a_i denotes the action associated with the i^{th} rule.

if ($p_i < p_{i+1}$ **AND** $r_i \not\cap r_{i+1}$ **OR** $a_i == a_{i+1}$)) **then**

[0042] Sorting rules in this fashion can have positive impact on the average number of comparisons required. Figure 5A is a table illustrating a firewall policy prior to rule sorting. The expected number of comparisons for the policy of Figure 5A is 4.26 comparisons per packet. Figure 5B is a table illustrating the firewall policy of Figure 5A after sorting using the above-described algorithm. In Figure 5B, the expected number of comparisons per packet has been reduced to 3.94. Thus, using the method described above, the number of comparisons per packet can be reduced by 8%.

[0043] However, it can be seen from Figures 5A and 5B that one rule can prevent another rule from being reordered. For example in Figure 5B, rule r_1 prevents rule r_4 from being placed closer to the beginning of the rule set. However, if both rules r_1 and r_4 are placed closer to the beginning of the policy while maintaining their relative order, the average number of comparison will be reduced.

[0044] To solve this problem, rule sets can be sorted using policy tries. First, the rule list is converted into an equivalent

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1025

policy trie. Each sub-trie will have an associated probability $\bar{p}$ that is average probability (hit-ratio) of the rules comprising the sub-trie. Sub-tries can be rotated around their parent node to increase performance, using the method described earlier. Since the policy trie (or equivalent rule set) is tested from left to right, rotation should occur if the probability of the right sub-trie is greater than the left sub-trie. As a result, the average number of comparisons required will be reduced and the policy integrity is maintained. Figure 6A is a trie representation of the rules in Figure 5B. Figure 6B is a table representing resulting rule set after rotating the TCP and UDP sub-tries. The TCP sub-trie in Figure 6A has an expected packet matching probability of 0.095 while the UDP sub-trie has an expected packet matching probability of 0.155. Thus, the UDP sub-trie and the TCP sub-trie should be rotated so that the UDP sub-trie becomes the left sub-trie and the TCP sub-trie becomes the right sub-trie. In the table illustrated in Figure 6B, which lists the ordering of the rules after rotating the sub-tries, the average number of comparisons per packets has been reduced to 3.70.

Rule Splitting

[0045] Rule splitting takes a general rule and creates more specific rules that collectively perform the same action over the same set of packets. Here, rule splitting is used to reduce the average number of comparisons. For example in Figure 7A, rule r_5 is split into two separate rules, r_{5a} for UDP and r_{5b} for TCP. Figure 7B is a table corresponding to the policy trie of Figure 7A. Once the rules are positioned based on their probabilities and their relation to other rules, the average number of rule comparisons is reduced to 2.98 (after sort, trie rotation, and splitting) which is a 30% less.

[0046] It many not be advantageous to split a general rule since it adds another rule to the policy. For example, assume a policy contains 20 rules where the first 19 rules have the same probability. Assume the last rule can be split and the new specific rule has a probability that is $\frac{3}{4}$ of the last rule. The impact of the probability of the last rule and the location of the new split rule, m, is depicted in Figures 8A and 8B. More particularly, Figure 8A is a graph illustrating the average number of comparisons for firewalls with and without rule splitting. Figure 8B is a graph illustrating average number of comparisons versus different locations of the new rule and probabilities. The average number of comparisons is reduced as the split rule is located closer to the first rule (surface decreases as m approaches one). Furthermore, splitting yields better results when the general rule has a high probability. However as illustrated in Figure 8B, the closer the specific rule is to the location of the original rule the average number of comparisons increases, which is the penalty of adding one more rule to the policy.

[0047] The effect of splitting a single rule can be described mathematically as follows. Consider n rules where rule r_n can be split into rules $r_{n,l}$ and $r_{n,r}$ (whose union is the original rule). In addition, assume the split rule $r_{n,l}$ will be located at the m^{th} position ($1 \leq m < n$), while rule $r_{n,r}$ will remain at the n^{th} location. The goal is to determine the best location m, which yields a lower average number of comparisons as compared to the original rule set. This can be defined mathematically in the following formula.

$$\sum_{i=1}^{n} i \cdot p_i > \sum_{i=1}^{m-1} i \cdot p_i + m \cdot p_{n,l} + \sum_{i=m}^{n-1} (i+1) \cdot p_i + (n+1) \cdot p_{n,r}$$

The left side of the inequality is the average number of comparisons for the original rule set. The right hand side of the inequality is the average number of comparisons with the specific rule at location m. If it is assumed that the rules located between m and n have an equal probability (denoted as p) the previous equation can be solved for m.

$$m < \frac{n \cdot p - n \cdot p_n + (n+1) \cdot p_{n,r}}{p - p_{n,l}}$$

The new rule must be located between the first and m^{th}; however, its final location will depend on the relationship with the other rules (cannot be placed before any rule for which it is a superset). This result can be applied iteratively to multiple rules and repeatedly to the same rule.

Rule Optimization for QoS

[0048] In the preceding sections, optimizing the entire security policy has been discussed. However, it may be desirable to optimize the rules for a certain type of traffic, in order to reduce the average number of comparisons this traffic

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1026

encounters. This can be done by organizing the policy trie via traffic classification, then optimizing the sub-tries. The result is an ordering where the rules for the most important traffic are tested first.

Network Firewall Optimization Using Rule Compression

[0049] Another method for network firewall optimization is referred to as rule compression. Given a security policy R, one objective in optimizing firewall performance is to reduce number of rules in the policy while maintaining the original integrity. A method for reducing the number of rules called *rule compression*, which combines and replaces two rules in the policy will now be described. Although compression does reduce the number of rules, it can only be done if the integrity of the original policy is maintained.

[0050] Consider two rules r_i and r_j in policy R, where r_i appears before r_j. The two rules can be compressed if and only if the following conditions are true

1. The two rules intersect, $r_i \cap r_j \neq 0$ (an edge between r_i and r_j exists in the policy DAG) (Tuples that are proper subsets and supersets are considered. Compression may still occur if the rules do not intersect or are adjacent in the set space, the resulting rule may contain ranges.)
2. The action of r_i and r_j is the same.
3. Rule r_i is not dependent on any rule r_k that is directly or indirectly dependent on r_j (only one path from r_i to r_j exists in the policy DAG).

[0051] The result of compression is a new rule, $r_{i,j}$ whose tuples are the union of the corresponding tuples of rules r_i and r_j. The new rule $r_{i,j}$ replaces both r_i and r_j in R; however, the location of $r_{i,j}$ in R may require the relocation of other rules. Let D_j be the ordered set of rules that appear after r_i that directly or indirectly depend upon r_j. The new rule is placed at the original location of r_i in R and the rules D_j are placed before $r_{i,j}$.

[0052] For example, consider the policy given in Table 2. It is possible to compress rules r_1 and r_5, creating the new rule $r_{1,5}$ =[TCP, *, *, 2.*, *, accept]. The ordered set D_j in this example consists of the rules r_3 and r_4. As a result, the new rule $r_{1,5}$ is placed at the original location of r_1 while the rules in D_j are placed before the new rule. Figures 9A and 9B respectively illustrate the original policy DAG for Table 2 and the policy DAG after rules r_1 and r_5 are compressed.

Table 2: Example Policy Consisting of Multiple Ordered Rules

		Source		Destination		
No.	Proto.	IP	Port	IP	Port	Action
1	TCP	1.1.*	*	2.2.*	*	accept
2	TCP	1.*	*	*	80	deny
3	TCP	3.3.*	*	*	80	accept
4	TCP	3.*	*	*	*	deny
5	TCP	*	*	2.*	90	accept

[0053] **Theorem** *Compressing rules r_i and r_j of policy R will maintain integrity if the three conditions are met and method described is used.*

[0054] **Proof** Before r_i and r_j are compressed, relocate the rules D_j immediately before r_i and D_i. This relocation will maintain integrity due to the third requirement for compression (no rule in D_i can be dependent on any rule in D_j). Now place r_j directly after r_i. This relocation will not affect integrity since no rule in D_i can have an edge to r_j (again, due to the third condition). Compressing the neighboring rules r_i and r_j creates $r_{i,j}$. This does not affect integrity since the resulting rule only matches the packets that match r_i or r_j in the original policy. Therefore, the result is the compression of rule r_i and r_j, and the integrity of the policy is maintained.

Determining if Compression Should Occur

[0055] Although compression may be possible, it may not improve the performance of the policy due to the relocation of rules. Compression should only occur if the average number of comparisons required for the new policy R_c is less than the original policy R; therefore, $E[R_c] < E[R]$.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1027

Firewall Rule Sub-Trie and Sub-List Ordering

Policy Sub-Trie Ordering

[0056] Additional methods for firewall optimization are referred to as firewall sub-trie and sub-list ordering. Given a policy trie T, it may be desirable to order all sub-tries such that the average number of tuple comparisons is reduced and policy integrity is maintained. Consider a policy trie T that contains sub-tries T_i and T_j having the same parent node, as seen in Figures 10A-10C. Let P_i denote the sum of the probabilities of the rules contained in sub-trie i, while C_i denotes the number of comparisons required to completely traverse sub-trie i (which is equal to the number of branches). In order to reduce the average number of tuple comparisons, sub-tries that share the same parent node should be ordered such that the following two conditions are observed.

1. Sub-tries that share a parent node are ordered such that the P_i values are non-ascending (higher match probabilities occur first, from left to right).
2. If sub-tries that share a parent node have the same probability (P_i equals P_j), then order the sub-tries such that the C_i values are non-descending (sub-tries consisting of fewer comparisons occur first, from left to right).

These conditions are recursively applied throughout the policy trie as seen in Figures 11A-11D. However, these conditions do not necessarily maintain policy integrity.

[0057] Consider a group of n sub-tries that share the same parent node that are numbered sequentially from left to right. Consider any two non-intersecting sub-tries T_i and T_{i+k} in this group that are out of order. If two sub-tries do not intersect, denoted as $T_i \pitchfork T_j$, then no rule in one sub-trie will intersect with any rule in the other sub-trie, $r_k \pitchfork r_l$, $\forall r_k \in T_i$, $\forall r_l \in T_j$. The following *integrity condition* must be observed when reordering the sub-tries. The two sub-tries can be exchanged (rotated about the parent) if and only the sub-tries T_i through T_{i+k-1} do not intersect with T_{i+k} and the sub-tries T_{i+1} through T_{i+k} do not intersect with T_i. Similar to finding the linear sequence of a policy DAGs, these conditions maintain the policy trie integrity. This is observed in Figure 10A, where no edge exists between any rule in T_1 and T_2. More particularly, Figure 10A is a policy DAG of a six-rule firewall list. Figure 10B is a policy trie representation of the firewall rule list represented by Figure 10A. In Figure 10B, the average number of tuple compares required per packet is 5.5. Figure 10C illustrates the policy trie after rotating sub-tries T_1 and T_2. In Figure 10C, the average number of tuple compares required is 4.5. In Figures 10A-10C, the following probabilities are assumed:

$P_1 = 0.5$
$C_1 = 4$
$P_2 = 0.5$; and
$C_2 = 2$.

The average number of tuple comparisons is an estimate based on the longest trie path.

[0058] Using the conditions for comparing and ordering, the following sorting algorithm sorts neighboring sub-tries that share the same parent node.

```
m = a parent node in T
done = false
while(!done)
    done = true
    for each sub-trie Ti having parent m AND a right neighboring sub-trie
        if((Pi < Pi+1 AND (Pi == Pi+1 AND Ci > Ci+1)) AND Ti ⋔ Ti+1) then
        interchange Ti and Ti+1
        done = false
        endif
    endfor
endwhile
```


CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1028

[0059] Figures 11A-11D illustrate policy trie rotation for an eight-rule firewall policy. More particularly, Figure 11A illustrates the policy DAG representation of an eight-rule firewall policy. Figure 11B is a policy trie representation of the original firewall policy. In this example, ordering occurs in three stages, starting at the bottom level and moving towards the top level of the trie. In Figure 11C, the leaves are ordered at the lowest level. More particularly, nodes r_1 and r_2 are exchanged with nodes r_7 and r_8. The average number of tuple compares for the ordering represented by Figure 11C is 10.45. In Figure 11D, the sub-tries on the second lowest level (T_3, T_4, T_5, and T_6) are ordered. More particularly, Figure 11 D illustrates the policy trie after exchanging sub-tries T_3 and T_4. The average number of tuple compares for the representation in Figure 11 D is 10.275. In Figure 11 E, the sub-tries on the second highest level, (T_1 and T_2) are ordered. Figure 11E illustrates the final ordering after exchanging T_1 and T_2. In the policy trie of Figure 11E, the average number of tuple compares required is 6.775. The average number of tuple comparisons is an estimate based on the longest trie path.

Policy Sub-List Ordering

[0060] The preceding section describes conditions for sorting policy sub-tries, which allow the exchange of groups of rules (sub-tries). The same conditions can be applied to list based policies, where sub-lists are ordered to improve the average number of rule comparisons.

[0061] In the sections above relating to rule sorting, a method is described to exchange neighboring rules in a list-based security policy, as seen in Figure 12B. However, it was observed that when ordering rules in this fashion it is possible that one rule can block the exchange of another. If groups of rules were allowed to be exchanged the average number of rule comparisons could be further reduced.

[0062] Given a policy list L, it may be desirable to order all sub-lists such that the average number of rule comparisons is reduced and policy integrity is maintained. Consider a policy list L that contains sub-tries L_i and L_j. Let P_i denote the sum of the probabilities of the rules contained in sub-list i, while C_i denotes the number of rules in sub-list i. In order to reduce the average number of rule comparisons, sub-lists should be ordered such that the following two conditions are observed.

1. Sub-lists are ordered such that the P_i values are non-ascending (higher match probabilities occur first, from left to right).
2. If sub-lists have the same probability (P_i equals P_j), then order the sub-lists such that the C_i values are non-descending (sub-lists consisting of fewer comparisons occur first, from left to right).

These conditions are applied throughout the policy list; however as with policy tries, these conditions do not necessarily maintain policy integrity.

[0063] Consider a group of n sub-lists that are numbered sequentially from left to right. Consider any two non-intersecting sub-lists L_i and L_{i+k} in this group that are out of order. If two sub-lists do not intersect, denoted as $L_i \pitchfork L_j$, then no rule in one sub-list will intersect with any rule in the other sub-list, $r_k \pitchfork r_l, \forall r_k \in L_i, \forall r_l \in L_j$. The following *integrity condition* must be observed when reordering the sub-lists. The two sub-lists can be exchanged if and only if the sub-lists L_i through L_{i+k-1} do not intersect with any rule in L_{i+k} and sub-lists L_{i+1} through L_{i+k} do not intersect with L_i. Similar to finding the linear sequence of a policy DAGs, these conditions maintain the policy trie integrity. For example, consider the sub-lists $L_1 = \{r_1, r_3\}$ and $L_2 = \{r_2, r_4\}$ given in Figure 12B. No edge exists between any rule in L_1 and L_2. The two sub-lists are out of order with respect to P and should be exchanged. The resulting list is given in Figure 12C. Figure 12A illustrates the original policy order and results in an average number of rule comparisons of 3.4. In the sorted version in Figure 12B, the average number of rule comparisons is 3.3. In Figure 12C, the average number of rule comparisons has been reduced to 2.7. Using the conditions for comparing and ordering, the following sorting algorithm sorts neighboring sub-lists.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1029

```
done = false
while(!done)
    done = true
    for each sub-list Li having a right neighboring sub-list
        if ((Pi < Pi+1 OR (Pi == Pi+1 AND Ci > Ci+1)) AND Li ⌒ Li+1)then
            interchange rules and probabilities
            done = false
        endif
    endfor
endwhile
```

Summary of Firewall Policy Optimization Techniques and Methods

[0064] As stated above, the methods and systems described herein for a network firewall policy optimization can be implemented using any combination of hardware, software, and/or firmware. Alternatively, the steps for firewall policy optimization can be performed manually with the number of rules being optimized is small enough. Figure 13 is a flow chart illustrating an exemplary process for firewall policy optimization according to an embodiment of the subject matter described herein. Referring to Figure 13, in block **1300**, a firewall policy including an ordered list of firewall rules is defined. In block **1302**, for each rule, a probability indicating a likelihood of receiving a packet that matches each rule is specified. In block **1304**, neighboring rules are sorted in order of non-increasing probability in a manner that preserves firewall policy. In block **1306**, groups of rules are sorted in order of non-increasing probability in a manner that preserves firewall policy. For example, any of the policy-trie-based methods described above may be used. In block **1308**, one or more rules are split to reduce the number of average packet comparisons. In block **1310**, intersecting rules having common actions are identified and collapsed into single rules.

[0065] As described above, the subject matter described herein may be implemented in hardware, software, and/or firmware. In one implementation, the subject matter described herein for firewall policy optimization may be implemented on a general purpose computing platform. Figure 14 is a block diagram of a general purpose computing platform including software for firewall policy optimization according to an embodiment of the subject matter described herein. Referring to Figure 14, computing platform **1400** may be a general purpose computer, such as a personal computer, that includes a microprocessor **1402**, memory **1404**, and I/O interfaces **1406**. Microprocessor **1400** may be any suitable microprocessor, such as any of the Intel or AMD families of microprocessors. Memory **1404** may include volatile memory for running programs and persistent memory, such as one or more disk storage devices. I/O interface **1406** may include interfaces with I/O devices, such as user input devices and output devices.

[0066] In the illustrated example, software that may be resident in memory 1404 includes a firewall policy editor 1408 and a firewall policy optimizer 1410. Firewall policy editor 1408 may allow a user to define a firewall policy. For example, firewall policy editor 1408 may allow a user to define a firewall policy by specifying different rules. The rules may be specified in a graphical manner, for example using policy DAGs as described above. Alternatively, firewall policy editor 1408 may allow a user to input rules in a tabular manner, as illustrated in any of the tables described herein. Firewall policy optimizer 1410 may implement any or all of the firewall policy optimization techniques to order rules entered via firewall policy editor 1408 in a manner that preserves policy integrity and that enhances firewall performance.

[0067] A firewall rule set that is optimized using the subject matter described herein may be implemented on any firewall system that includes one or more firewalls. For example, an optimized firewall rule set according to embodiments of the subject matter described herein may be implemented using any of the hierarchical, multi-node firewall systems described in commonly assigned, co- pending U.S. patent application no. 11/316,331, filed December 22, 2005. Figure 15 illustrates an example of a multi-node firewall system suitable for implementing the firewall rules according to an embodiment of the subject matter described herein. Referring to Figure 15, a plurality of firewall nodes 1500, 1502, 1504, and 1506 may collectively implement any of the optimized firewall policies described herein. For example, the firewall nodes may collectively implement different portions of a firewall policy data structure including an ordered list of firewall rules. A firewall policy engine 1508 resident on each firewall node may filter packets using its local portion 1510

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1030

of the firewall policy data structure. A control node 1512 may include a firewall policy optimizer 1514 that measures the average number of comparisons per packets for the current rule configuration and may dynamically reorder the rules to reduce the average number of comparisons per packet. For example, firewall policy optimizer 1514 may utilize any of the methods described herein to rearrange rules and improve firewall performance. Rules may be tested for rearrangement at user-specified intervals or when average number of packet comparisons increases by a user-specified amount.

References

[0068] The disclosure of each of the following references is hereby incorporated herein by reference in its entirety.

[1] E. Al-Shaer and H. Hamed. Firewall Policy Management Advisor for Anomaly Detection and Rule Editing. In Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management, 2003.
[2] E. Al-Shaer and H. Hamed. Modeling and Management of Firewall Policies. IEEE Transactions on Network and Service Management, (1), 2004.
[3] S. M. Bellovin and W. Cheswick. Network Firewalls. IEEE Communications Magazine, pages 50-57, Sept. 1994.
[4] C. Benecke. A Parallel Packet Screen for High Speed Networks. In Proceedings of the 15th Annual Computer Security Applications Conference, 1999.
[5] G. Brightwell and P. Winkler. Counting Linear Extensions is #P-Complete. In Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing, 1991.
[6] M. Christiansen and E. Fleury. Using Interval Decision Diagrams for Packet Filtering. Technical report, BRICS, 2002.
[7] D. Decasper, Z. Dittia, G. Parulkar, and B. Plattner. Router Plugins: A Software Architecture for Next-Generation. Routers. IEEE/ACM Transactions on Networking, 8(1), February 2000.
[8] M. Degermark, A. Brodnik, S. Carlsson, and S. Pink. Small Forwarding Tables for Fast Routing Lookups. In Proceedings of ACM SIGCOMM, pages 4-13, 1997.
[9] U. Ellermann and C. Benecke. Firewalls for ATM Networks. In Proceedings of INFOSEC'COM, 1998.
[10] A. Feldmann and S. Muthukrishnan. Tradeoffs for Packet Classification. In Proceedings of the IEEE INFOCOM, pages 397 - 413, 2000.
[11] E. W. Fulp. Firewall Architectures for High Speed Networks. Technical Report 20026, Wake Forest University Computer Science Department, 2002.
[12] E. W. Fulp. Firewall Policy Models Using Ordered-Sets. Technical report, Wake Forest University Computer Science Department, 2004.
[13] R. Funke, A. Grote, and H.-U. Heiss. Performance Evaluation of Firewalls in Gigabit-Networks. In Proceedings of the Symposium on Performance Evaluation of Computer and Telecommunication Systems, 1999.
[14] S. Goddard, R. Kieckhafer, and Y. Zhang. An Unavailability Analysis of Firewall Sandwich Configurations. In Proceedings of the 6th IEEE Symposium on High Assurance Systems Engineering, 2001.
[15] R. L. Graham, E. L. Lawler, J. K. Lenstra, and A. H. G. R. Kan. Optimizing and Approximation in Determinstic Sequencing and Scheduling: A Survey. Annals of Discrete Mathematics, 5:287 - 326, 1979.
[16] A. Hari, S. Suri, and G. Parulkar. Detecting and Resolving Packet Filter Conflicts. In Proceedings of IEEE INFOCOM, pages 1203-1212, 2000.
[17] HiPAC. High Performance Packet Classification. http://www.hipac.org.
[18] E. Horowitz, S. Sahni, and D. Mehta. Fundamentals of Data Structures in C++. Compuer Science Press, 1995.
[19] E. L. Lawler. Sequencing Jobs to Minimize Total Weighted Completion Time Subject to Precedence Constraints. Annals of Discrete Mathematics, 2:75 - 90, 1978.
[20] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the Self-Similar Nature of Ethernet Traffic. IEEE Transactions on Networking, 2:1-15, 1994.
[21] J. K. Lenstra and A. H. G. R. Kan. Complexity of Scheduling under Precedence Constraints. Operations Research, 26(1):22 - 35, 1978.
[22] O. Paul and M. Laurent. A Full Bandwidth ATM Firewall. In Proceedings of the 6th European Symposium on Research in Computer Security ESORICS'2000, 2000.
[23] B. R. Preiss. Data Structures and Algorithms with Object-Oriented Design Patterns in C++. John Wiley & Sons, 1999.
[24] L. Qui, G. Varghese, and S. Suri. Fast Firewall Implementations for Software and Hardware-Based Routers. In Proceedings of ACM SIGMETRICS, June 2001.
[25] V. P. Ranganath and D. Andresen. A Set-Based Approach to Packet Classification. In Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems, pages 889-894, 2003.
[26] R. Rivest. On Self-Organizing Sequential Search Heuristics. Communications of the ACM, 19(2), 1976.
[27] W. E. Smith. Various Optimizers for Single-Stage Production. Naval Research Logistics Quarterly, 3:59 - 66,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1031

1956.
[28] V. Srinivasan, G. Varghese, S. Suri, and M. Waldvogel. Fast and Scalable Layer Four Switching. In Proceedings of ACM SIGCOMM, pages 191-202, 1998.
[29] S. Suri and G. Varghese. Packet Filtering in High Speed Networks. In Proceedings of the Symposium on Discrete Algorithms, pages 969 - 970, 1999.
[30] R. L. Ziegler. Linux Firewalls. New Riders, second edition, 2002.
[31] E. D. Zwicky, S. Cooper, and D. B. Chapman. Building Internet Firewalls. O'Reilly, 2000.

[0069] It will be understood that various details of the invention may be changed without departing from the scope of the invention. Furthermore, the foregoing description is for the purpose of illustration only, and not for the purpose of limitation.

Claims

1. A method for producing a performance optimized firewall policy, the method comprising:

   (a) defining a firewall policy including an ordered list of firewall rules;
   (b) for each rule, defining a probability indicating the likelihood of receiving a packet matching each rule; and
   (c) sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy;

   the method **characterized in that**

   (d) sorting groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy, wherein at least one of the groups of two or more rules includes a first rule that prevents a second rule in the group from being moved, within the group, closer to a beginning of the ordered list that implements the firewall policy and wherein sorting the groups of rules includes moving the first and second rules closer to the beginning of the ordered list while maintaining their relative order.

2. The method of claim 1 wherein sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy includes, for each pair of neighboring rules, rearranging the rules if a first rule of the pair has a lower probability than a second rule of the pair and the first rule of the pair is not a subset of the second rule of the pair with a different action.

3. The method of claim 1, wherein sorting groups of two or more rules in a manner that preserves firewall policy includes assigning each group of rules to a sub-trie in a policy trie, computing a sum of packet matching probabilities for rules in each sub-trie, and, in response to determining that a first sub-trie has a lower sum than a second sub-trie, rotating the first and second sub-tries about a common parent so that rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy.

4. The method of claim 3, comprising, in response to determining that the first and second sub-tries have equal sums, determining whether the second sub-trie has an lower number of branches than the first sub-trie, and, in response to determining that the second sub-trie has a lower number of branches than the first sub-trie, rotating the first and second sub-tries about a common parent so that the rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy.

5. The method of claim 1 wherein sorting groups of two or more rules in a manner that preserves the firewall policy includes:

   (e) identifying first and second sub-lists of rules in the ordered list of firewall rules, the first sub-list preceding the second sub-list in the ordered list of firewall rules;
   (f) determining a sum of packet matching probabilities for each of the first and second sub-lists;
   (g) determining whether the first and second sub-lists intersect; and
   (h) in response to determining that the sum of the packet matching probabilities for the second sub-list is greater than the sum of the packet matching probabilities for the first sub-list and that the first and second sub-lists do not intersect, switching the order of the first and second sub-lists in the ordered list of firewall rules.

6. The method of claim 5, comprising, in response to determining that the sums of the packet matching probabilities

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1032

for the first and sub-lists are equal, determining whether the second sub-list has a lower number of rules than the first sub-list, and, in response to determining that the second sub-list has a lower number of rules than the first sub-list, switching the order of the first and second sub-lists in the ordered list of firewall rules.

7. The method of claim 1, comprising splitting at least one of the firewall rules to reduce an average number of comparisons required per received packet.

8. The method of claim 1, comprising identifying intersecting rules having a common action and collapsing the rules into a single rule representing a union of the intersecting rules.

9. The method of claim 1, wherein sorting neighboring rules includes sorting rules such that comparisons for a particular class of packets are reduced.

10. The method of claim 1, comprising assigning the rules to vertices in a directed acyclical graph, DAG, determining precedence relationships between the rules, assigning the precedence relationships to edges that connect the vertices in the DAG, and wherein sorting the neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy includes sorting the neighboring rules in a manner that preserves the precedence relationships represented by the DAG.

11. The method of claim 1, wherein the rules comprise intrusion detection rules.

12. The method of claim 1, wherein the rules comprise intrusion protection rules.

13. A network firewall, comprising:

(a) a firewall policy data structure (1510) for storing an ordered list of firewall rules;
(b) a firewall policy engine (1508) for implementing a firewall policy by comparing packets to the rules in the order specified by the firewall policy data structure (1510); and
(c) a firewall policy optimizer (1514) for optimizing performance of the network firewall by reordering the rules in a manner that reduces an average number of rule comparisons per packet and hat preserves the firewall policy;

the network firewall **characterized in that** the firewall policy optimizer (1514) is adapted to (d) sort groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy, wherein at least one of the groups of two or more rules includes a first rule that prevents a second rule in the group from being moved, within the group, closer to a beginning of the ordered list that implements the firewall policy and wherein sorting the groups of rules includes moving the first and second rules closer to the beginning of the ordered list while maintaining their relative order.

14. The network firewall of claim 13, wherein the firewall policy optimizer (1514) is adapted to measure average number of comparisons per packet for a rule ordering specified by the firewall policy data structure (1510) and to dynamically re-order to rules to reduce the average number of comparisons per packet.

Patentansprüche

1. Verfahren zum Generieren einer leistungsoptimierten Firewall-Vorschrift, wobei das Verfahren Folgendes umfasst:

(a) Definieren einer Firewall-Vorschrift, die eine geordnete Liste mit Firewall-Regeln umfasst;
(b) für jede Regel, Definieren einer Wahrscheinlichkeit, welche die Wahrscheinlichkeit des Empfangens eines mit jeder Regel übereinstimmenden Paketes anzeigt; und
(c) Sortieren benachbarter Regeln in der Reihenfolge einer nicht-zunehmenden Wahrscheinlichkeit in einer Weise, welche die Firewall-Vorschrift wahrt;

wobei das Verfahren **gekennzeichnet ist durch**:

(d) Sortieren von Gruppen von zwei oder mehr Regeln in der Reihenfolge einer nicht-zunehmenden Wahrscheinlichkeit in einer Weise, welche die Firewall-Vorschrift wahrt, wobei mindestens eine der Gruppen von zwei oder mehr Regeln eine erste Regel umfasst, die verhindert, dass eine zweite Regel in der Gruppe innerhalb

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1033

der Gruppe näher zu einem Anfang der geordneten Liste, welche die Firewall-Vorschrift implementiert, verschoben wird, und wobei das Sortieren der Gruppen von Regeln umfasst, die erste und die zweite Regel näher an den Anfang der geordneten Liste zu bewegen, während ihre relative Reihenfolge beibehalten wird.

2. Verfahren nach Anspruch 1, wobei das Sortieren benachbarter Regeln in der Reihenfolge einer nicht-zunehmenden Wahrscheinlichkeit in einer Weise, welche die Firewall-Vorschrift wahrt, für jedes Paar benachbarter Regeln umfasst, die Regeln umzuordnen, wenn eine erste Regel das Paares einer geringere Wahrscheinlichkeit als eine zweite Regel das Paares hat und die erste Regel das Paares keine Teilmenge der zweiten Regel das Paares mit einer anderen Aktion ist.

3. Verfahren nach Anspruch 1, wobei das Sortieren von Gruppen von zwei oder mehr Regeln in einer Weise, welche die Firewall-Vorschrift wahrt, Folgendes umfasst: Zuordnen jeder Gruppe von Regeln zu einem Sub-Trie in einem Vorschrifts-Trie; Berechnen einer Summe von Wahrscheinlichkeiten der Paketübereinstimmung für Regeln in jedem Sub-Trie; und in Reaktion auf die Feststellung, dass ein erster Sub-Trie eine kleinere Summe als ein zweiter Sub-Trie hat, Drehen des ersten und des zweiten Sub-Tries um einen gemeinsamen Ursprung, so dass Regeln in dem zweiten Sub-Trie vor den Regeln in dem ersten Sub-Trie in der Firewall-Vorschrift angewendet werden.

4. Verfahren nach Anspruch 3, das in Reaktion auf die Feststellung, dass der erste und der zweite Sub-Trie gleiche Summen haben, das Feststellen umfasst, ob der zweite Sub-Trie eine kleinere Anzahl von Abzweigungen als der erste Sub-Trie hat, und in Reaktion auf die Feststellung, dass der zweite Sub-Trie eine kleinere Anzahl von Abzweigungen als der erste Sub-Trie hat, das Drehen des ersten und des zweiten Sub-Tries um einen gemeinsamen Ursprung umfasst, so dass die Regeln in dem zweiten Sub-Trie vor den Regeln in dem ersten Sub-Trie in der Firewall-Vorschrift angewendet werden.

5. Verfahren nach Anspruch 1, wobei das Sortieren von Gruppen von zwei oder mehr Regeln in einer Weise, welche die Firewall-Vorschrift wahrt, Folgendes umfasst:

(e) Identifizieren einer ersten und einer zweiten Teilliste von Regeln in der geordneten Liste mit Firewall-Regeln, wobei die erste Teilliste der zweiten Teilliste in der geordneten Liste mit Firewall-Regeln vorausgeht;
(f) Bestimmen einer Summe von Wahrscheinlichkeiten der Paketübereinstimmung für die erste und für die zweite Teilliste;
(g) Feststellen, ob sich die erste und die zweite Teilliste überschneiden; und
(h) in Reaktion auf die Feststellung, dass die Summe der Wahrscheinlichkeiten der Paketübereinstimmung für die zweite Teilliste größer als die Summe der Wahrscheinlichkeiten der Paketübereinstimmung für die erste Teilliste ist und dass sich die erste und die zweite Teilliste nicht überschneiden, Wechseln der Reihenfolge der ersten und der zweiten Teilliste in der geordneten Liste mit Firewall-Regeln.

6. Verfahren nach Anspruch 5, das, in Reaktion auf die Feststellung, dass die Summen der Wahrscheinlichkeiten der Paketübereinstimmung für die erste und Teilliste gleich sind, das Feststellen umfasst, ob die zweite Teilliste eine kleinere Anzahl von Regeln als die erste Teilliste hat, und in Reaktion auf die Feststellung, dass die zweite Teilliste eine kleinere Anzahl von Regeln als die erste Teilliste hat, das Wechseln der Reihenfolge der ersten und der zweiten Teilliste in der geordneten Liste mit Firewall-Regeln umfasst.

7. Verfahren nach Anspruch 1, das umfasst, mindestens eine der Firewall-Regeln zu teilen, um eine durchschnittliche Anzahl von Vergleichen, die je empfangenem Paket erforderlich sind, zu verringern.

8. Verfahren nach Anspruch 1, das umfasst, sich überschneidende Regeln zu identifizieren, die eine gemeinsame Aktion haben, und die Regeln zu einer einzigen Regel zusammenzulegen, die eine Vereinigung der sich überschneidenden Regeln darstellt.

9. Verfahren nach Anspruch 1, wobei das Sortieren benachbarter Regeln umfasst, die Regeln so zu sortieren, dass Vergleiche für eine bestimmte Klasse von Paketen verringert werden.

10. Verfahren nach Anspruch 1, das Folgendes umfasst: Zuweisen der Regeln zu Scheitelpunkten in einem gerichteten azyklischen Graph (DAG); Bestimmen von Präzedenzbeziehungen zwischen den Regeln; Zuweisen der Präzedenzbeziehungen zu Flanken, welche die Scheitelpunkte in dem DAG verbinden; und wobei das Sortieren der benachbarten Regeln in der Reihenfolge einer nicht-zunehmenden Wahrscheinlichkeit in einer Weise, welche die Firewall-Vorschrift wahrt, umfasst, die benachbarten Regeln in einer Weise zu sortieren, welche die durch den DAG

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1034

dargestellten Präzedenzbeziehungen wahrt.

11. Verfahren nach Anspruch 1, wobei die Regeln Eindringdetektionsregeln umfassen.

12. Verfahren nach Anspruch 1, wobei die Regeln Eindringschutzregeln umfassen.

13. Netzwerk-Firewall, die Folgendes umfasst:

(a) eine Firewall-Vorschrifts-Datenstruktur (1510) zum Speichern einer geordneten Liste mit Firewall-Regeln;
(b) eine Firewall-Vorschrifts-Engine (1508) zum Implementieren einer Firewall-Vorschrift durch Vergleichen von Paketen mit den Regeln in der durch die Firewall-Vorschrifts-Datenstruktur (1510) spezifizierten Reihenfolge; und
(c) einen Firewall-Vorschrifts-Optimierer (1514) zum Optimieren der Leistung der Netzwerk-Firewall durch Umordnen der Regeln in einer Weise, die eine durchschnittliche Anzahl von Regelvergleichen je Paket verringert und die Firewall-Vorschrift wahrt;

wobei die Netzwerk-Firewall **dadurch gekennzeichnet ist, dass** der Firewall-Vorschrifts-Optimierer (1514) dafür geeignet ist:

(d) Gruppen von zwei oder mehr Regeln in der Reihenfolge einer nicht-zunehmenden Wahrscheinlichkeit in einer Weise zu sortieren, welche die Firewall-Vorschrift wahrt, wobei mindestens eine der Gruppen von zwei oder mehr Regeln eine erste Regel umfasst, die verhindert, dass eine zweite Regel in der Gruppe innerhalb der Gruppe näher zu einem Anfang der geordneten Liste, welche die Firewall-Vorschrift implementiert, verschoben wird, und wobei das Sortieren der Gruppen von Regeln umfasst, die erste und die zweite Regel näher an den Anfang der geordneten Liste zu bewegen, während ihre relative Reihenfolge beibehalten wird.

14. Netzwerk-Firewall nach Anspruch 13, wobei der Firewall-Vorschrifts-Optimierer (1514) dafür geeignet ist, die durchschnittliche Anzahl von Vergleichen je Paket für eine durch die Firewall-Vorschrifts-Datenstruktur (1510) spezifizierte Regelordnung zu messen und die Regeln dynamisch umzuordnen, um die durchschnittliche Anzahl von Vergleichen je Paket zu verringern.

Revendications

1. Procédé de génération d'une politique de pare-feu à performances optimisées, le procédé consistant à :

(a) définir une politique de pare-feu comprenant une liste ordonnée de règles de pare-feu ;
(b) pour chaque règle, définir une probabilité indiquant la probabilité de recevoir un paquet correspondant à chaque règle ; et
(c) trier des règles voisines dans l'ordre de probabilité non croissante d'une manière qui préserve la politique de pare-feu ;

le procédé étant **caractérisé par** le fait de

(d) trier des groupes de deux règles ou plus dans l'ordre de probabilité non croissante d'une manière qui préserve la politique de pare-feu, dans lequel au moins l'un des groupes de deux règles ou plus comprend une première règle qui empêche une seconde règle du groupe d'être déplacée, au sein du groupe, plus près d'un début de la liste ordonnée qui met en oeuvre la politique de pare-feu, et dans lequel le tri des groupes de règles consiste à déplacer les première et seconde règles plus près du début de la liste ordonnée, tout en maintenant leur ordre relatif.

2. Procédé selon la revendication 1, dans lequel le tri de règles voisines dans l'ordre de probabilité non croissante d'une manière qui préserve la politique de pare-feu consiste, pour chaque paire de règles voisines, à réarranger les règles si une première règle de la paire a une probabilité plus faible qu'une seconde règle de la paire et si la première règle de la paire n'est pas un sous-ensemble de la seconde règle de la paire avec une action différente.

3. Procédé selon la revendication 1, dans lequel le tri de groupes de deux règles ou plus d'une manière qui préserve la politique de pare-feu consiste à affecter chaque groupe de règles à un sous-arbre dans un arbre de politique,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1035

calculer une somme de probabilités de correspondance de paquet pour les règles de chaque sous-arbre, et, en réponse à la détermination que le premier sous-arbre a une somme inférieure à un second sous-arbre, faire tourner les premier et second sous-arbres autour d'un parent commun, de sorte que les règles du second sous-arbre soient appliquées avant les règles du premier sous-arbre dans la politique de pare-feu.

4. Procédé selon la revendication 3, consistant, en réponse à la détermination que les premier et second sous-arbres ont des sommes égales, à déterminer si le second sous-arbre a un plus petit nombre de branches que le premier sous-arbre, et, en réponse à la détermination que le second sous-arbre a un plus petit nombre de branches que le premier sous-arbre, faire tourner les premier et second sous-arbres autour d'un parent commun, de sorte que les règles du second sous-arbre soient appliquées avant les règles du premier sous-arbre dans la politique de pare-feu.

5. Procédé selon la revendication 1, dans lequel le tri de groupes de deux règles ou plus d'une manière qui préserve la politique de pare-feu consiste à :

(e) identifier de première et seconde sous-listes de règles dans la liste ordonnée de règles de pare-feu, la première sous-liste précédant la seconde sous-liste dans la liste ordonnée de règles de pare-feu ;
(f) déterminer une somme de probabilités de correspondance de paquet pour chacune des première et seconde sous-listes ;
(g) déterminer si les première et seconde sous-listes se croisent ; et
(h) en réponse à la détermination que la somme des probabilités de correspondance de paquet pour la seconde sous-liste est supérieure à la somme des probabilités de correspondance de paquet pour la première sous-liste et que les première et seconde sous-listes ne se croisent pas, inverser l'ordre des première et seconde sous-listes dans la liste ordonnée de règles de pare-feu.

6. Procédé selon la revendication 5, consistant, en réponse à la détermination que les sommes des probabilités de correspondance de paquet pour les première et seconde sous-listes sont égales, déterminer si la seconde sous-liste a un plus petit nombre de règles que la première sous-liste, et, en réponse à la détermination que la seconde sous-liste a un plus petit nombre de règles que la première sous-liste, inverser l'ordre des première et seconde sous-listes dans la liste ordonnée de règles de pare-feu.

7. Procédé selon la revendication 1, consistant à diviser au moins l'une des règles de pare-feu afin de réduire un nombre moyen de comparaisons nécessaire par paquet reçu.

8. Procédé selon la revendication 1, consistant à identifier des règles qui se croisent ayant une action commune et à regrouper les règles en une règle unique qui représente une union des règles qui se croisent.

9. Procédé selon la revendication 1, dans lequel le tri de règles voisines consiste à trier des règles de telle sorte que les comparaisons relatives à une catégorie particulière de paquets sont réduites.

10. Procédé selon la revendication 1, consistant à affecter les règles à des sommets dans un graphe acyclique dirigé, DAG, déterminer des relations de priorité entre les règles, affecter les relations de priorité à des bords qui relient les sommets dans le DAG, et dans lequel le tri des règles voisines dans l'ordre de probabilité non croissante d'une manière qui préserve la politique de pare-feu consiste à trier les règles voisines d'une manière qui préserve les relations de priorité représentées par le DAG.

11. Procédé selon la revendication 1, dans lequel les règles comprennent des règles de détection des intrusions.

12. Procédé selon la revendication 1, dans lequel les règles comprennent des règles de protection contre les intrusions.

13. Pare-feu de réseau, comprenant :

(a) une structure de données de politique de pare-feu (1510) pour stocker une liste ordonnée de règles de pare-feu ;
(b) un moteur de politique de pare-feu (1508) pour mettre en oeuvre une politique de pare-feu en comparant des paquets aux règles dans l'ordre spécifié par la structure de données de politique de pare-feu (1510) ; et
(c) un optimiseur de politique de pare-feu (1514) pour optimiser la performance du pare-feu de réseau en réordonnant les règles d'une manière qui réduit un nombre moyen de comparaisons de règles par paquet et qui préserve la politique de pare-feu ;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1036

le pare-feu de réseau étant **caractérisé en ce que** l'optimiseur de politique de pare-feu (1514) est conçu pour :

(d) trier des groupes de deux règles ou plus dans l'ordre de probabilité non croissante d'une manière qui préserve la politique de pare-feu, dans lequel au moins l'un des groupes de deux règles ou plus comprend une première règle qui empêche une seconde règle du groupe d'être déplacée, au sein du groupe, plus près d'un début de la liste ordonnée qui met en oeuvre la politique de pare-feu, et dans lequel le tri des groupes de règles consiste à déplacer les première et seconde règles plus près du début de la liste ordonnée, tout en maintenant leur ordre relatif.

14. Pare-feu de réseau selon la revendication 13, dans lequel l'optimiseur de politique de pare-feu (1514) est conçu pour mesurer un nombre moyen de comparaisons par paquet pour un ordonnancement de paquets spécifié par la structure de données de politique de pare-feu (1510) et pour réordonner de manière dynamique les règles afin de réduire le nombre moyen de comparaisons par paquet.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1037

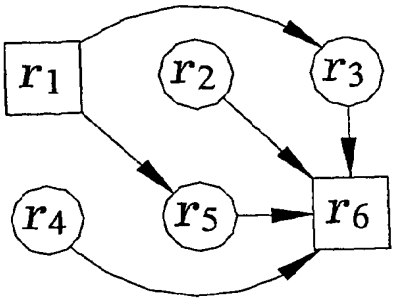

FIG. 1A

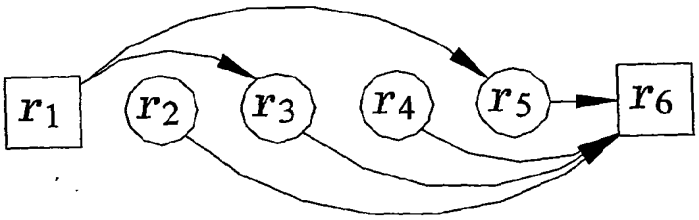

FIG. 1B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1038

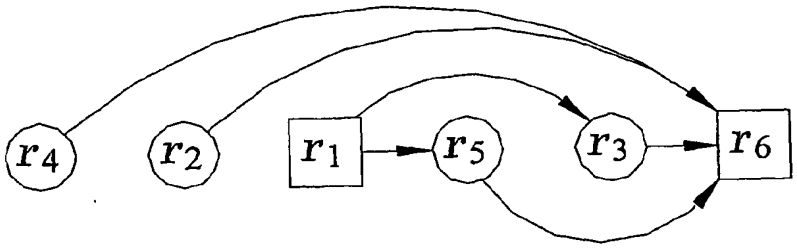

FIG. 2A

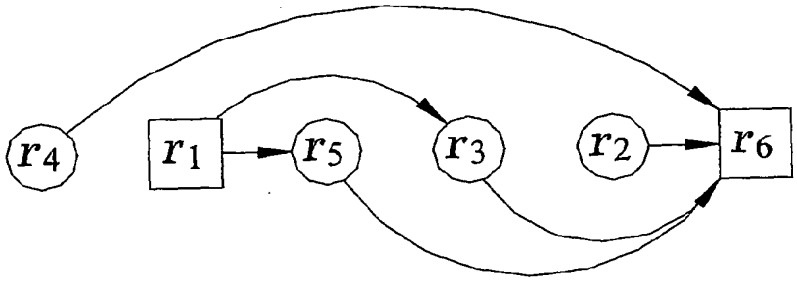

FIG. 2B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1039

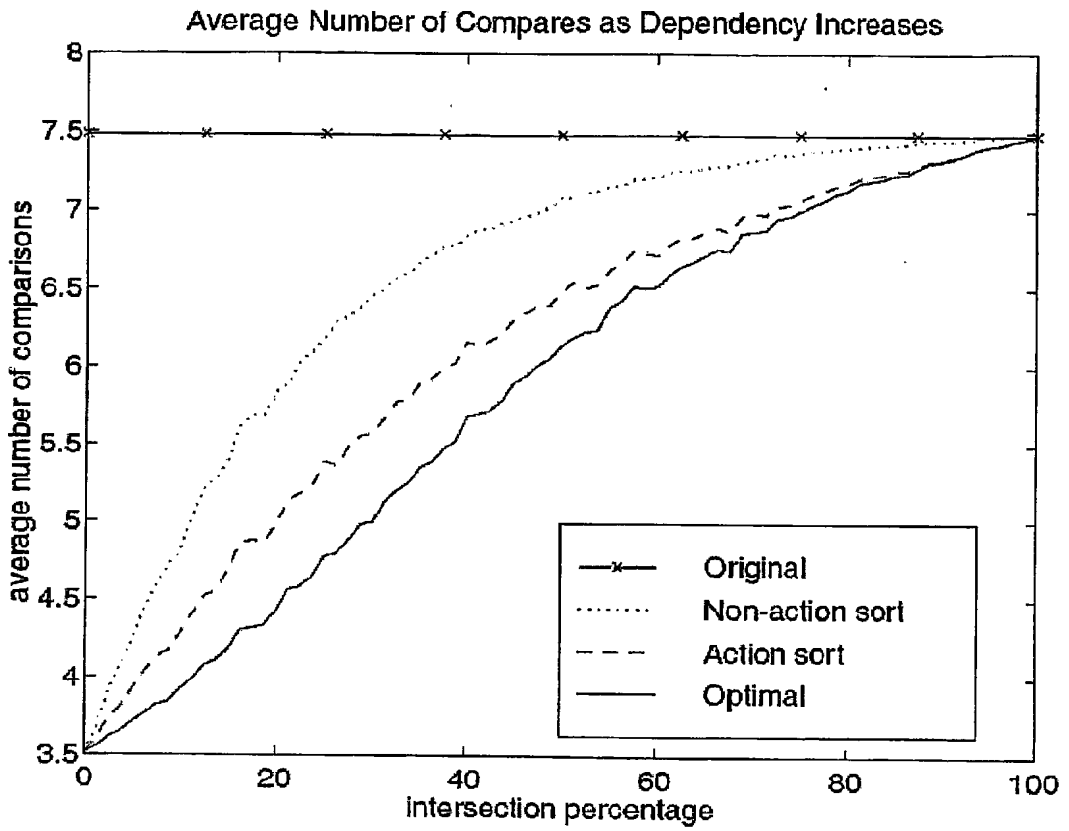

FIG. 3A

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1040

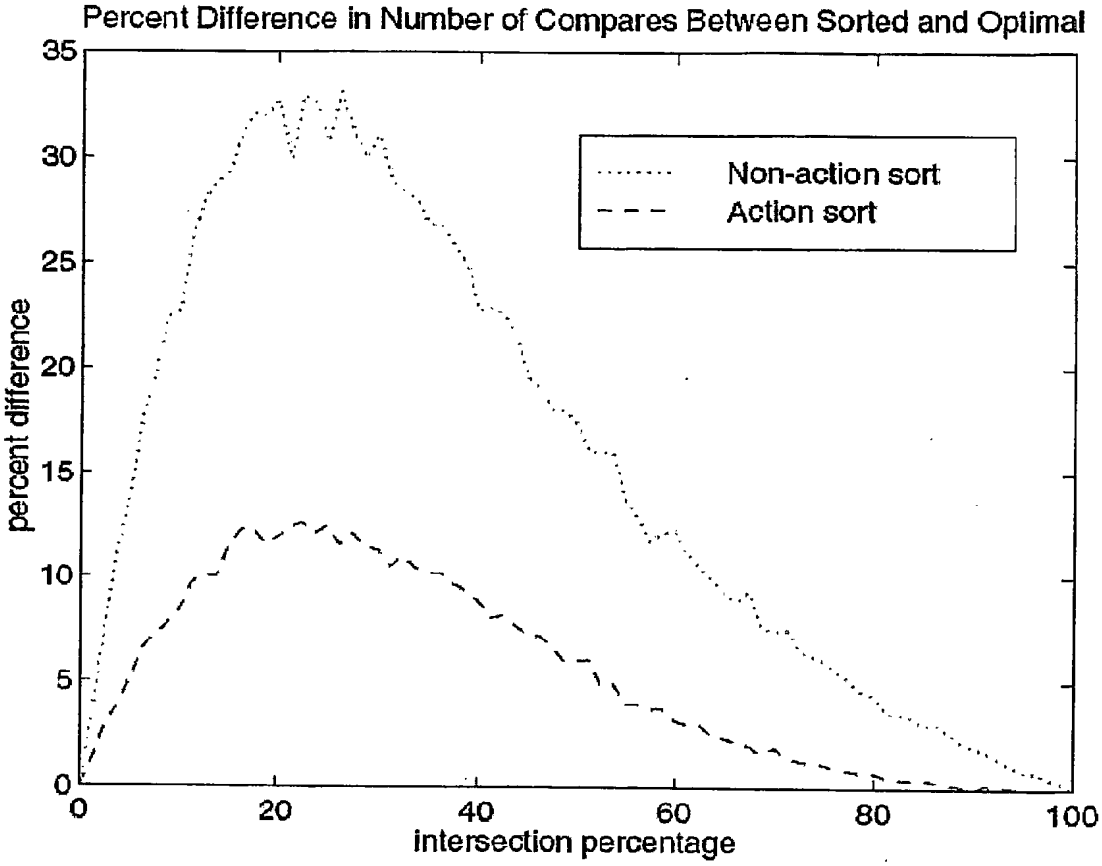

FIG. 3B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1041

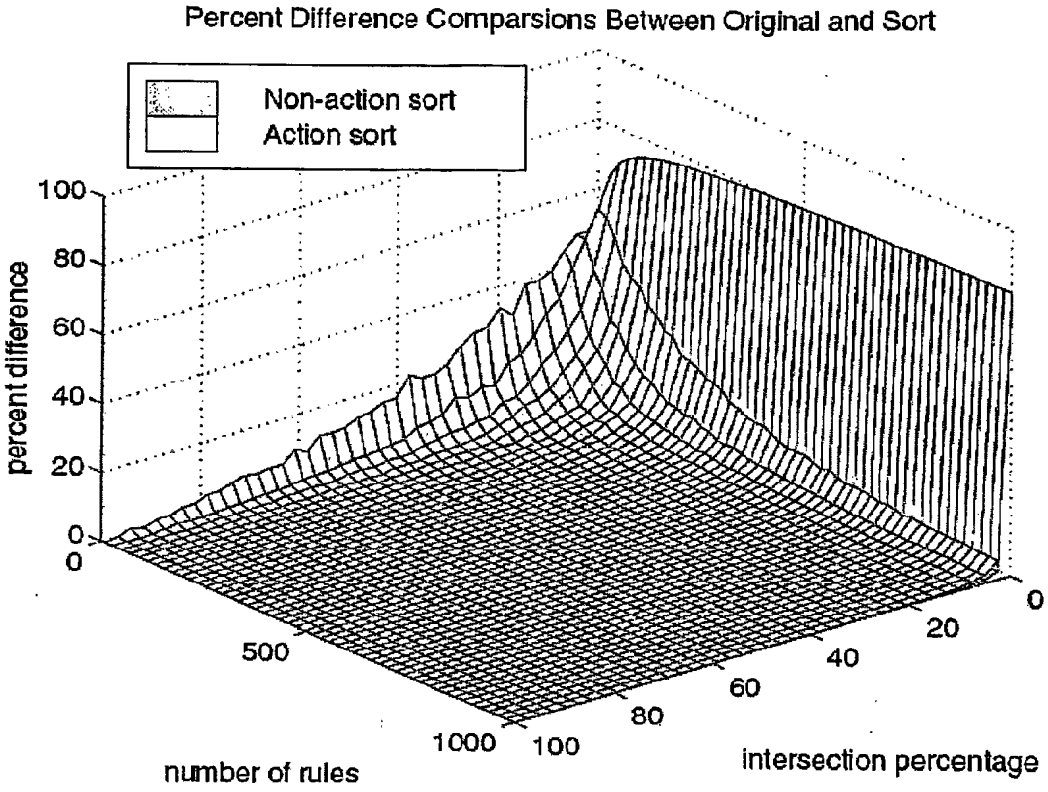

FIG. 4A

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1042

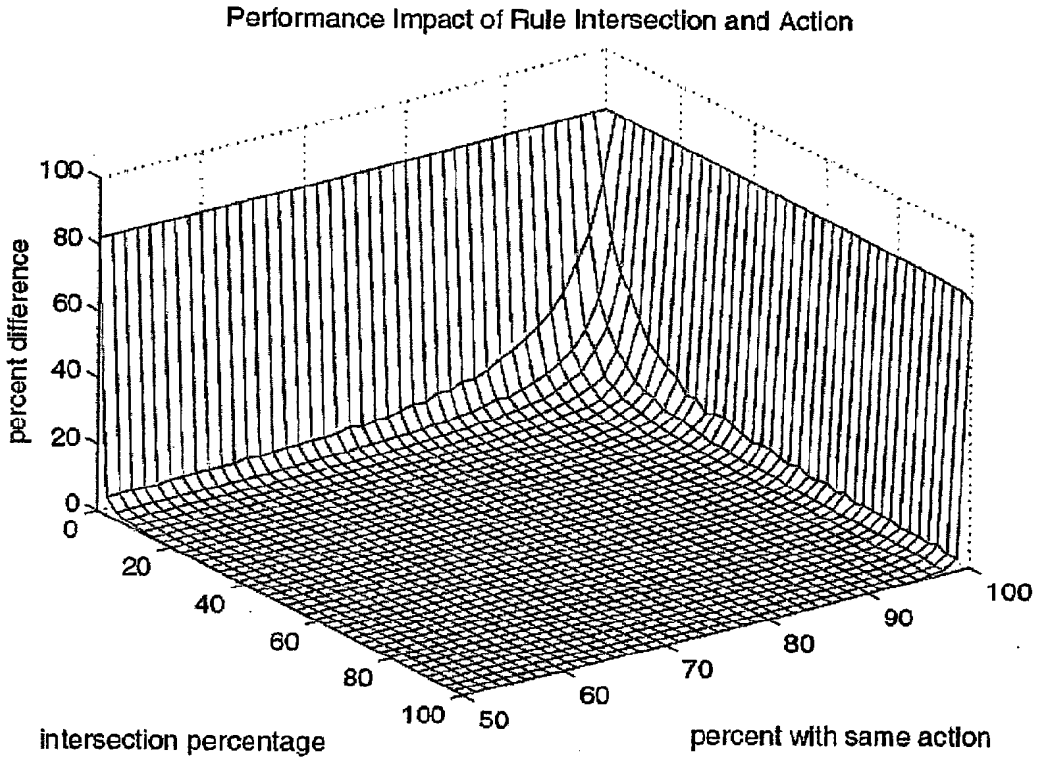

FIG. 4B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1043

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 5A

r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 5B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1044

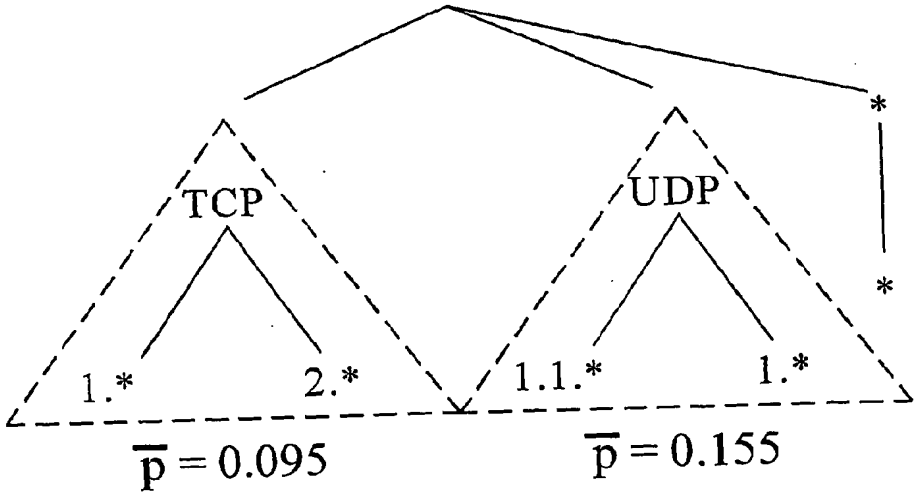

FIG. 6A

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 6B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1045

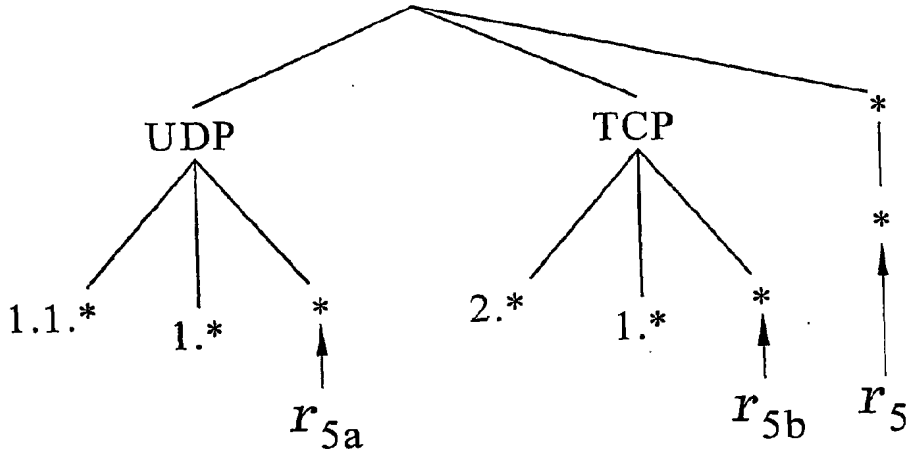

FIG. 7A

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	ACCEPT)	$p_4 = 0.30$
r_{5a}	(UDP,	*,	DENY)	$p_{5a} = 0.47$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_{5b}	(TCP,	*,	DENY)	$p_{5b} = 0.03$

FIG. 7B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1046

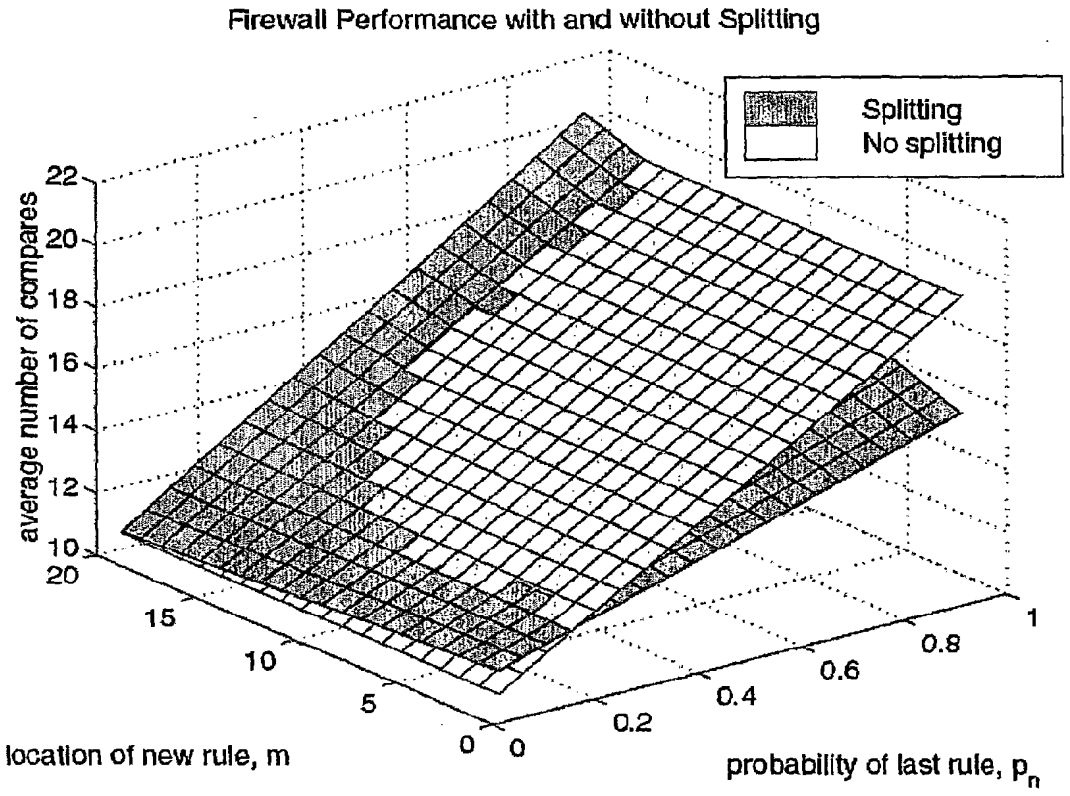

FIG. 8A

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1047

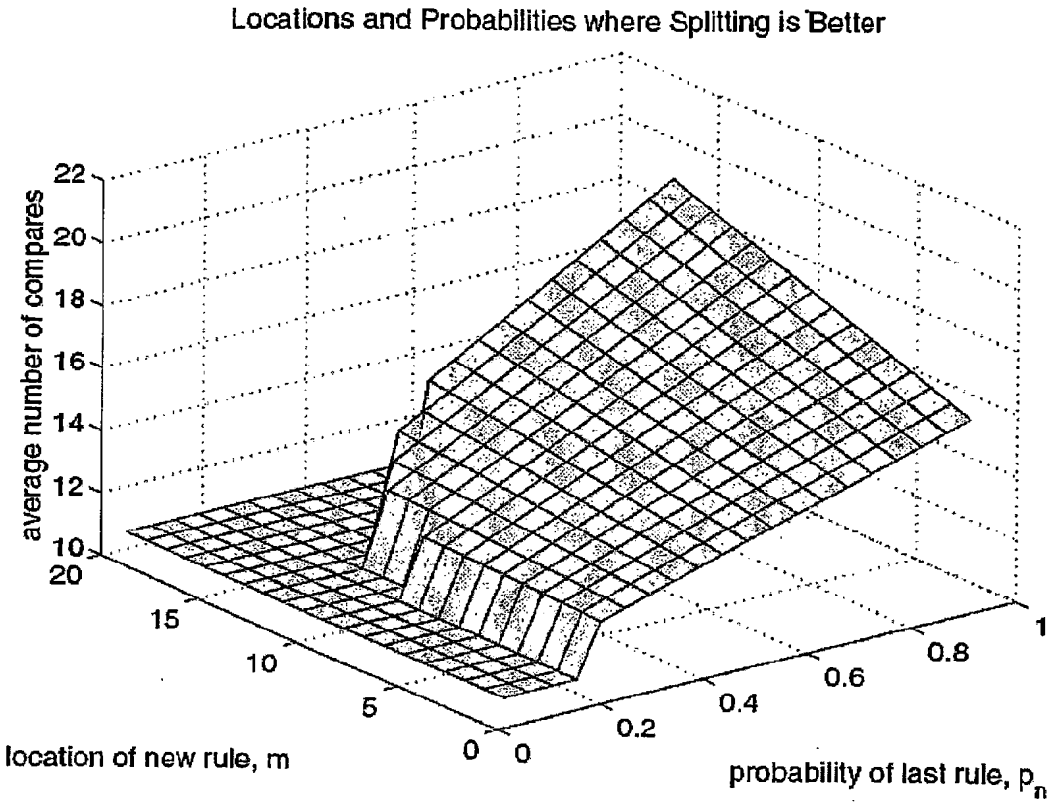

FIG. 8B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1048

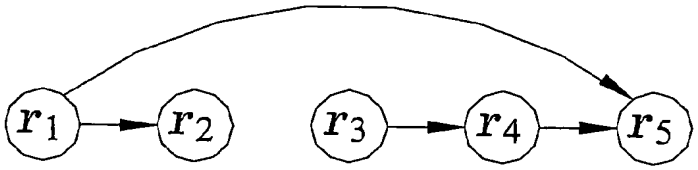

FIG. 9A

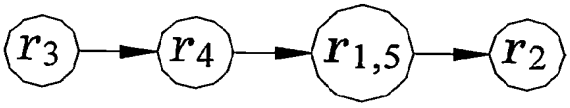

FIG. 9B

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1049

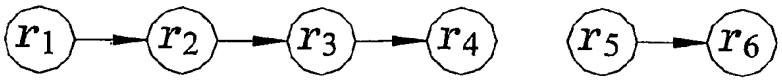

FIG. 10A

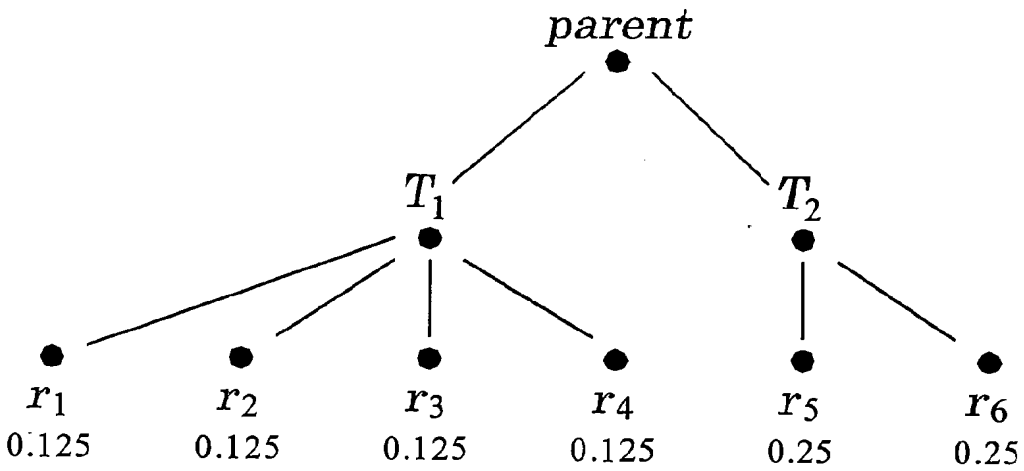

FIG. 10B

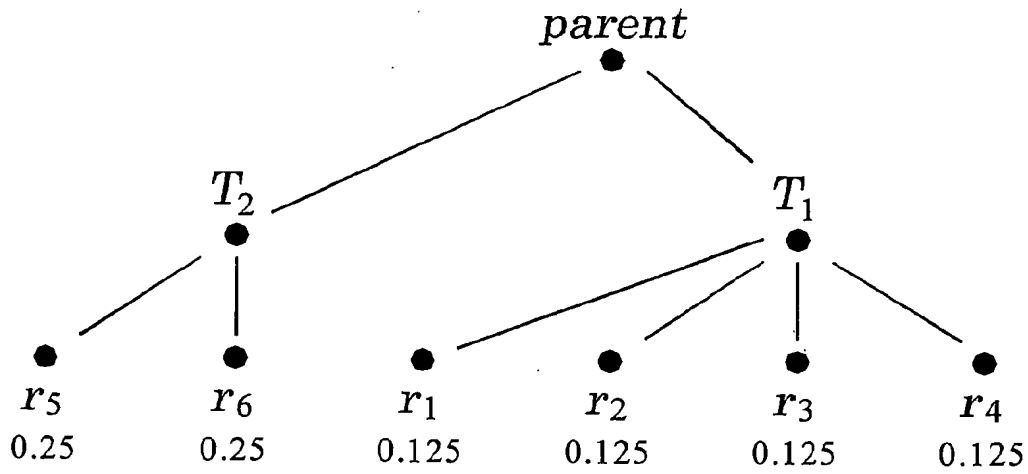

FIG. 10C

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1050

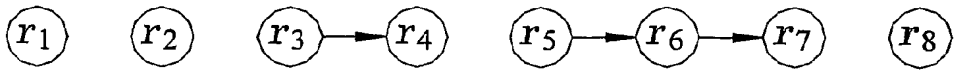

FIG. 11A

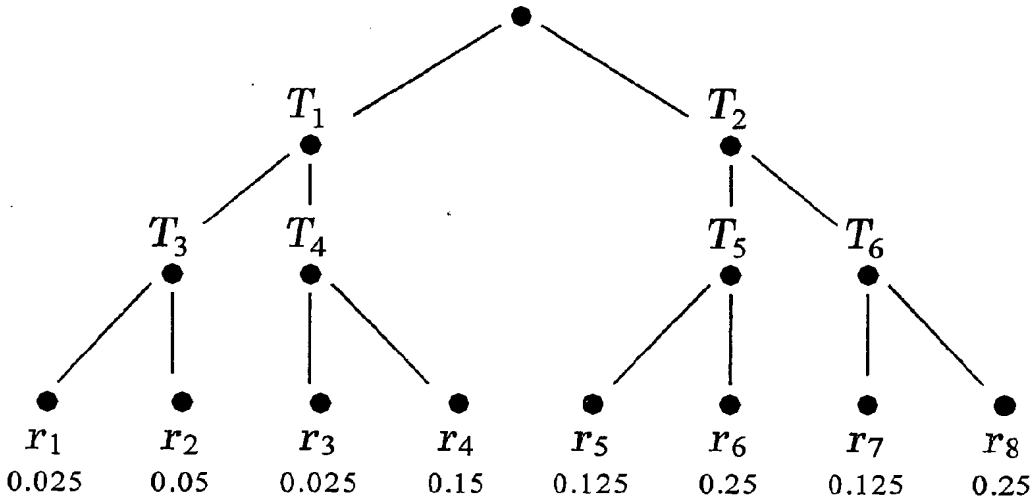

FIG. 11B

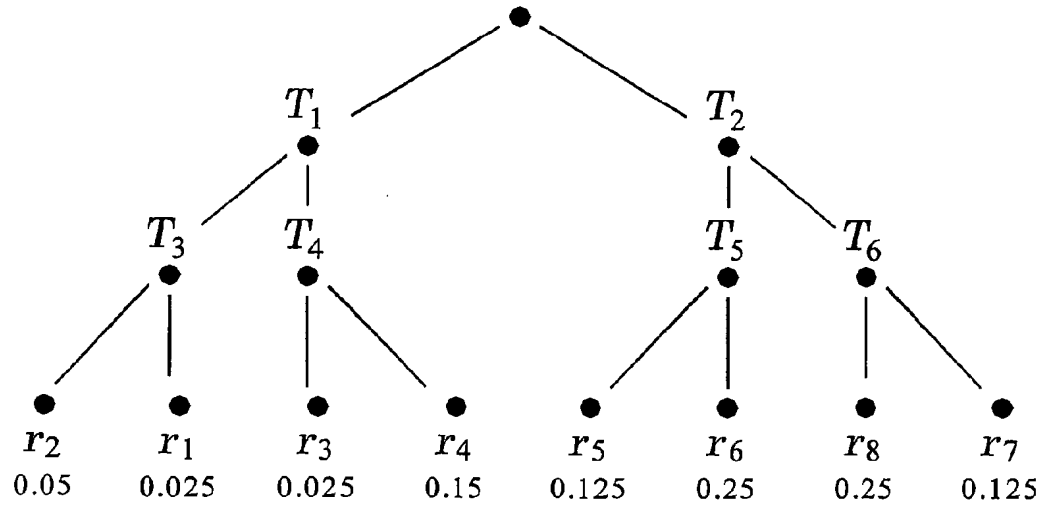

FIG. 11C

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1051

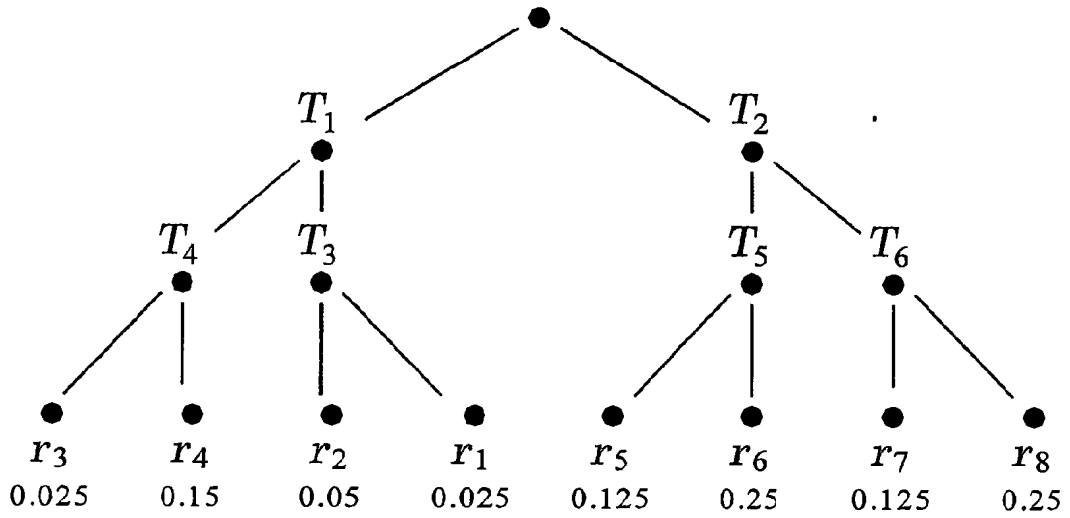

FIG. 11D

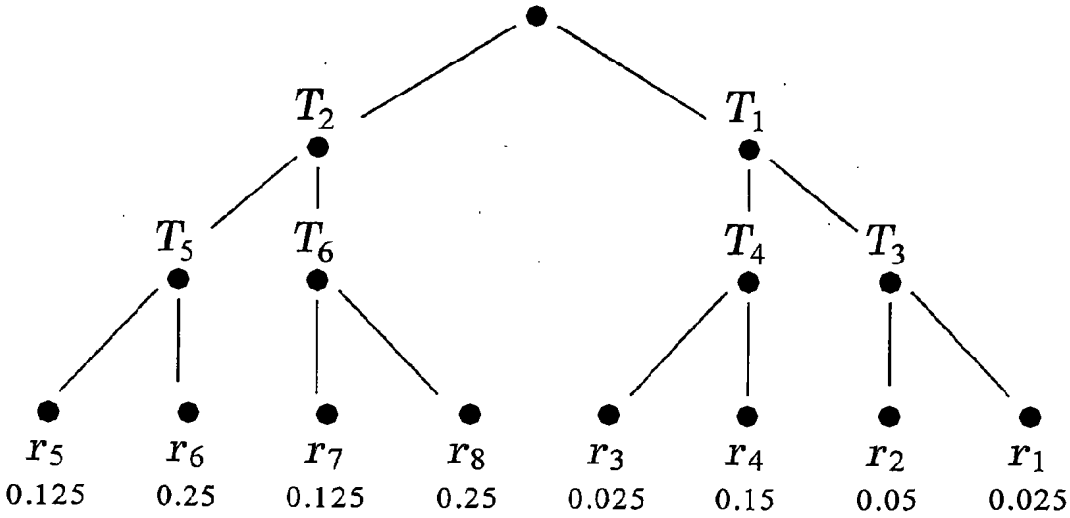

FIG. 11E

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1052

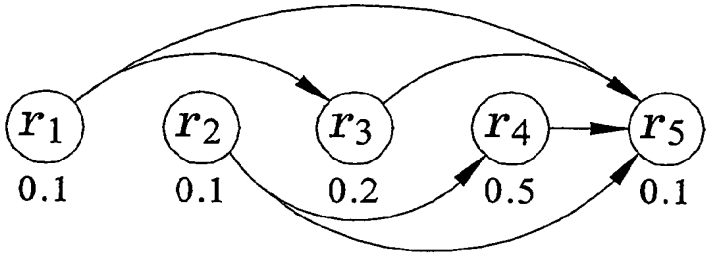

FIG. 12A

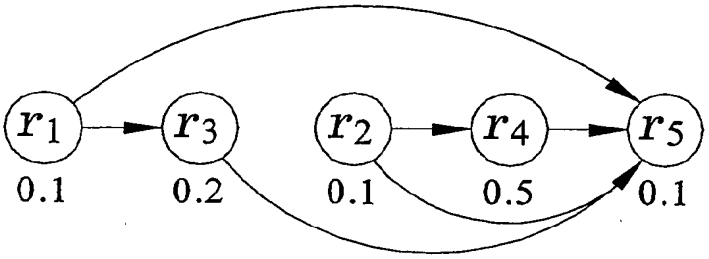

FIG. 12B

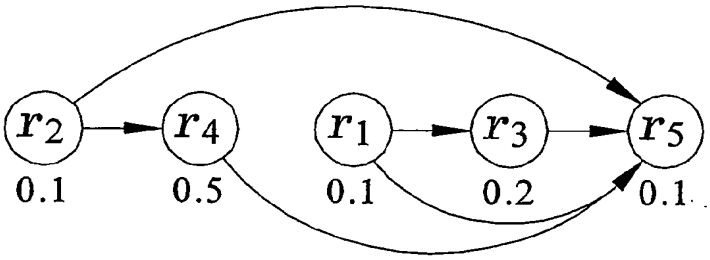

FIG. 12C

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1053

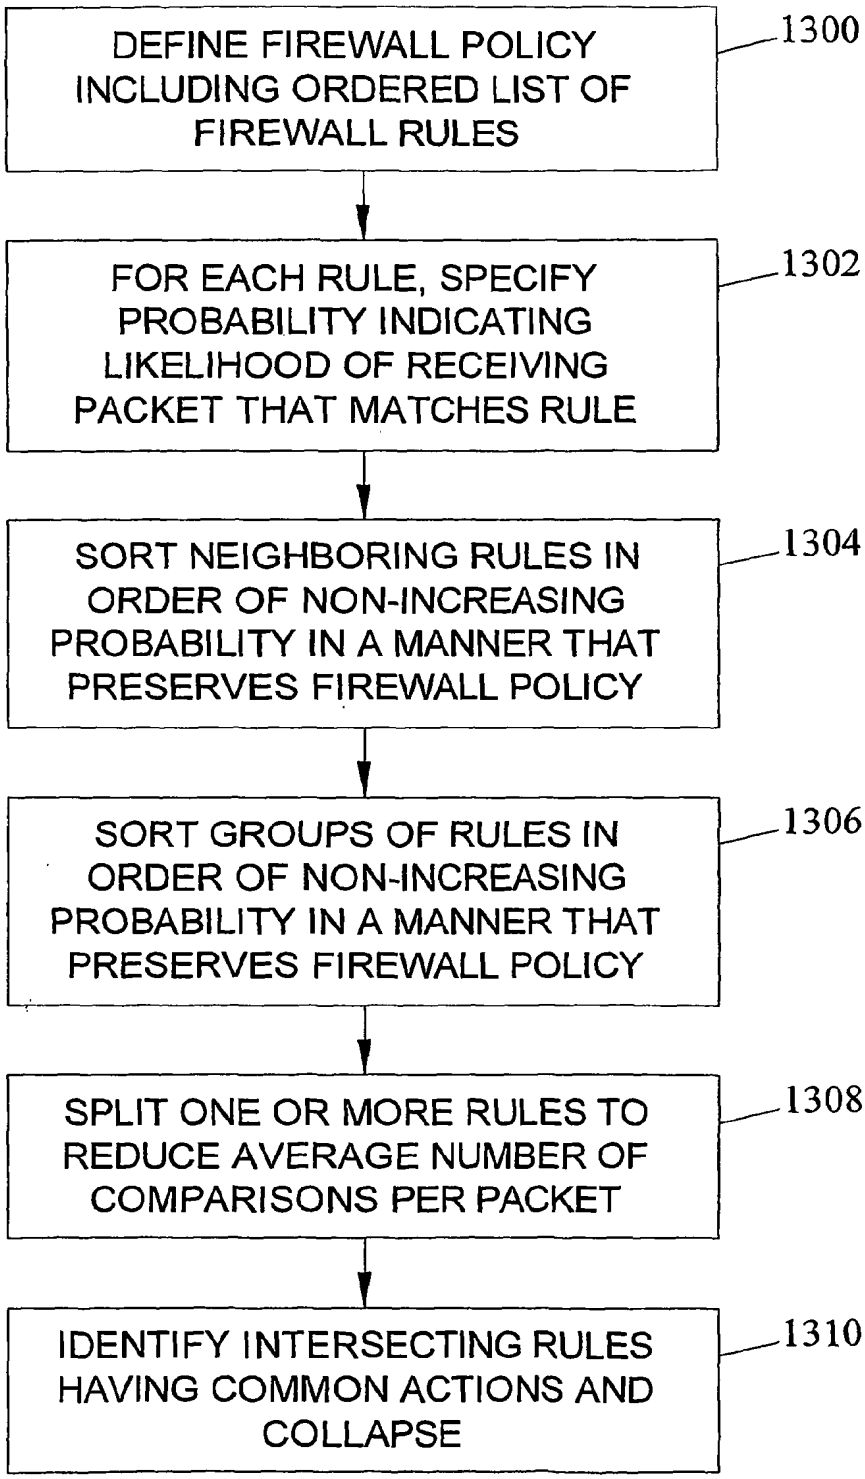

FIG. 13

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1054

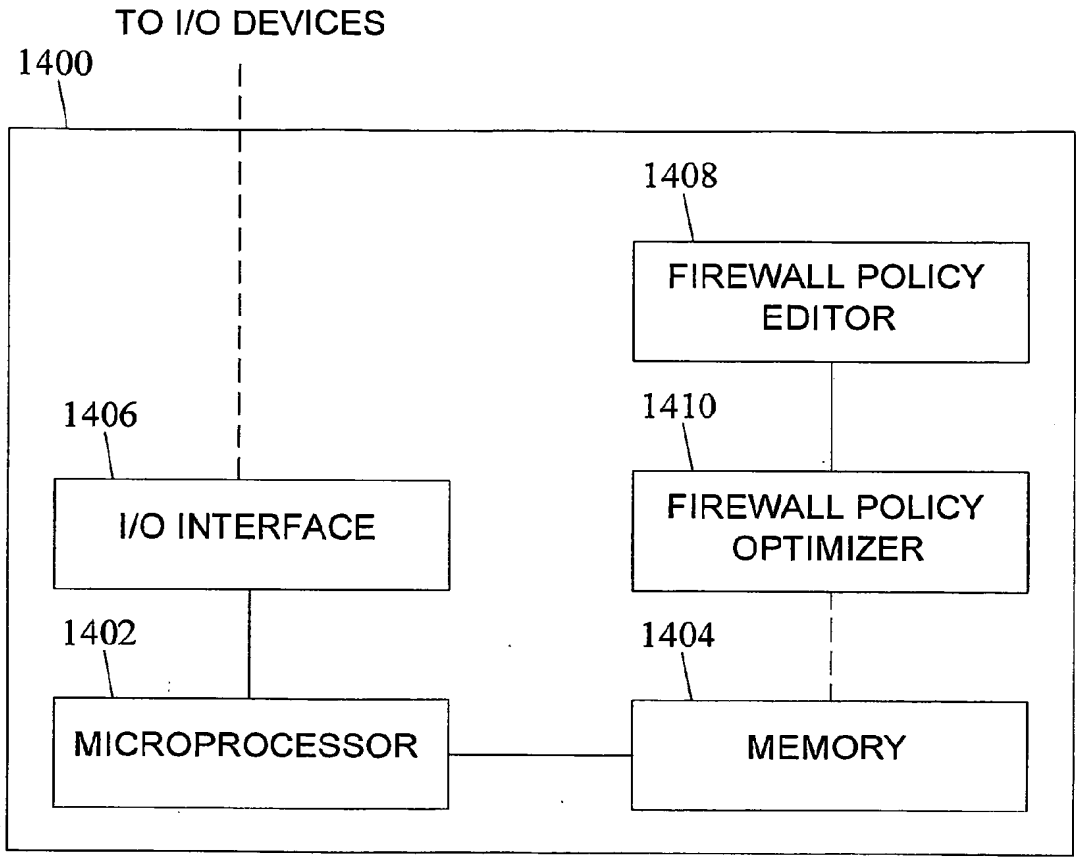

FIG. 14

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1055

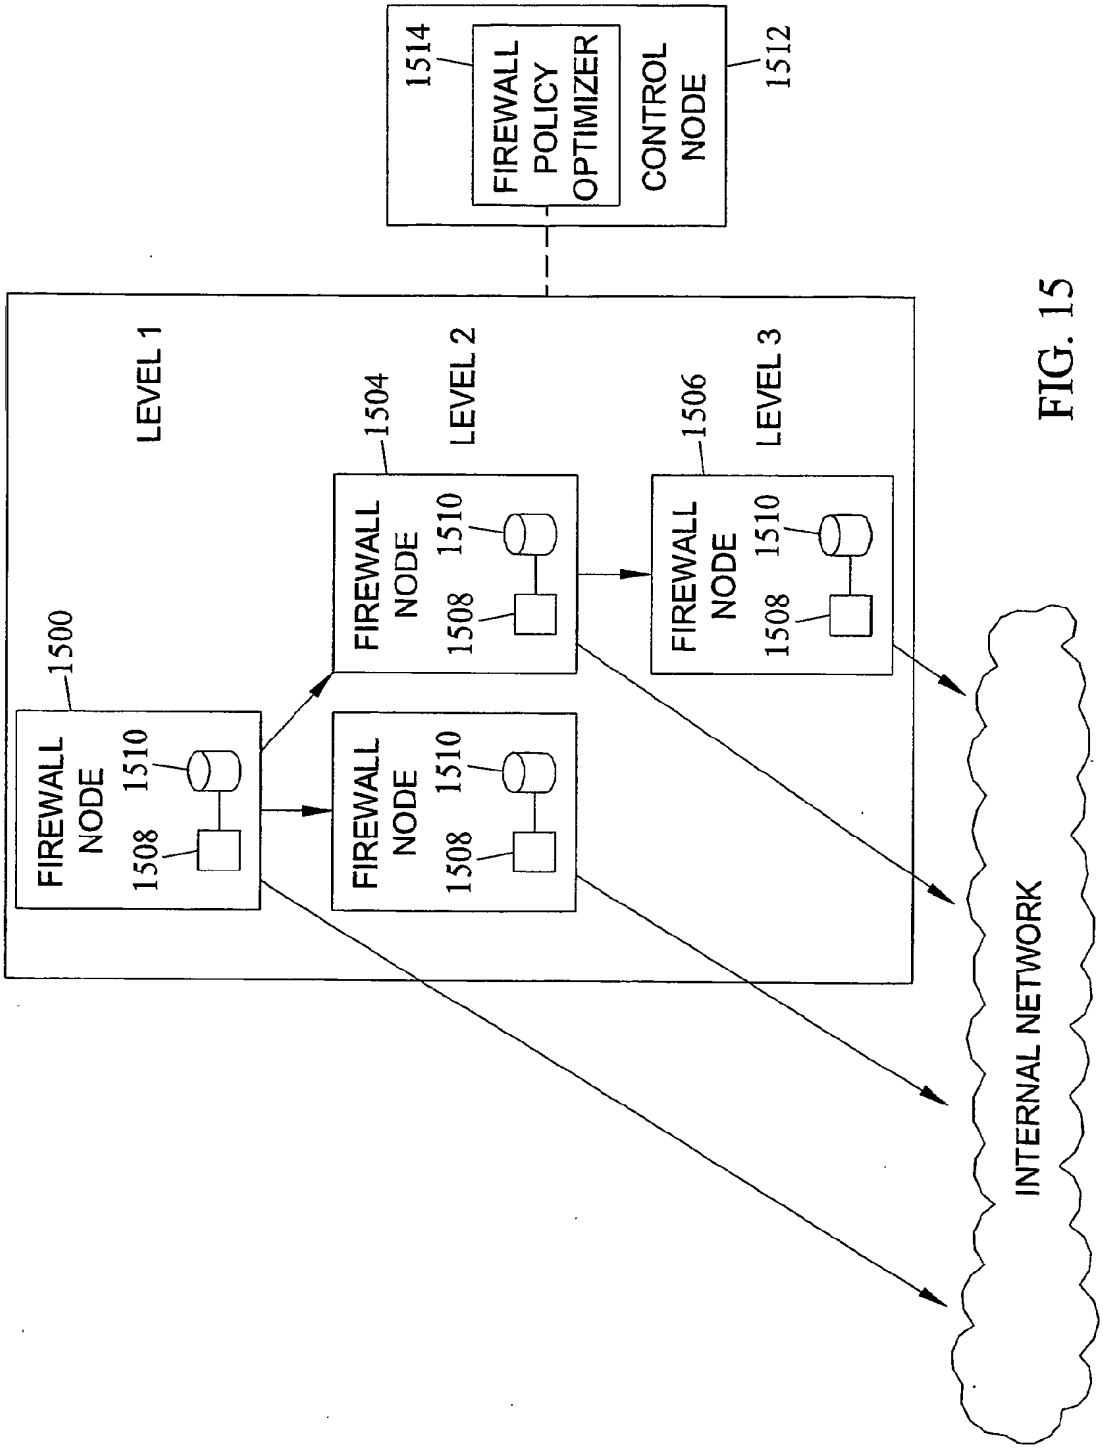

FIG. 15

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1056

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- EP 1006701 A **[0005]**
- US 2002198981 A **[0006]**
- WO 2006093557 A **[0007]**
- US 31633105 A **[0067]**

Non-patent literature cited in the description

- **E. AL-SHAER ; H. HAMED.** Firewall Policy Management Advisor for Anomaly Detection and Rule Editing. *Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management,* 2003 **[0068]**
- **E. AL-SHAER ; H. HAMED.** Modeling and Management of Firewall Policies. *IEEE Transactions on Network and Service Management,* 2004 **[0068]**
- **S. M. BELLOVIN ; W. CHESWICK.** Network Firewalls. *IEEE Communications Magazine,* September 1994, 50-57 **[0068]**
- **C. BENECKE.** A Parallel Packet Screen for High Speed Networks. *Proceedings of the 15th Annual Computer Security Applications Conference,* 1999 **[0068]**
- **G. BRIGHTWELL ; P. WINKLER.** Counting Linear Extensions is #P-Complete. *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing* **[0068]**
- **M. CHRISTIANSEN ; E. FLEURY.** Using Interval Decision Diagrams for Packet Filtering. *Technical report, BRICS,* 2002 **[0068]**
- **D. DECASPER ; Z. DITTIA ; G. PARULKAR ; B. PLATTNER.** Router Plugins: A Software Architecture for Next-Generation. Routers. *IEEE/ACM Transactions on Networking,* February 2000, vol. 8 (1 **[0068]**
- **M. DEGERMARK ; A. BRODNIK ; S. CARLSSON ; S. PINK.** Small Forwarding Tables for Fast Routing Lookups. *Proceedings of ACM SIGCOMM,* 1997, 4-13 **[0068]**
- **U. ELLERMANN ; C. BENECKE.** Firewalls for ATM Networks. *Proceedings of INFOSEC'COM,* 1998 **[0068]**
- **A. FELDMANN ; S. MUTHUKRISHNAN.** Tradeoffs for Packet Classification. *Proceedings of the IEEE INFOCOM,* 2000, 397-413 **[0068]**
- **E. W. FULP.** Firewall Architectures for High Speed Networks. *Technical Report,* 2002 **[0068]**
- **E. W. FULP.** Firewall Policy Models Using Ordered-Sets. *Technical report,* 2004 **[0068]**
- **R. FUNKE ; A. GROTE ; H.-U. HEISS.** Performance Evaluation of Firewalls in Gigabit-Networks. *Proceedings of the Symposium on Performance Evaluation of Computer and Telecommunication Systems,* 1999 **[0068]**
- **S. GODDARD ; R. KIECKHAFER ; Y. ZHANG.** An Unavailability Analysis of Firewall Sandwich Configurations. *Proceedings of the 6th IEEE Symposium on High Assurance Systems Engineering,* 2001 **[0068]**
- Optimizing and Approximation in Determinstic Sequencing and Scheduling: A Survey. **R. L. GRAHAM ; E. L. LAWLER ; J. K. LENSTRA ; A. H. G. R. KAN.** Annals of Discrete Mathematics. 1979, vol. 5, 287-326 **[0068]**
- **A. HARI ; S. SURI ; G. PARULKAR.** Detecting and Resolving Packet Filter Conflicts. *Proceedings of IEEE INFOCOM,* 2000, 1203-1212 **[0068]**
- *High Performance Packet Classification, http://www.hipac.org* **[0068]**
- **E. HOROWITZ ; S. SAHNI ; D. MEHTA.** Fundamentals of Data Structures in C++. Compuer Science Press, 1995 **[0068]**
- Sequencing Jobs to Minimize Total Weighted Completion Time Subject to Precedence Constraints. **E. L. LAWLER.** Annals of Discrete Mathematics. 1978, vol. 2, 75-90 **[0068]**
- **W. E. LELAND ; M. S. TAQQU ; W. WILLINGER ; D. V. WILSON.** On the Self-Similar Nature of Ethernet Traffic. *IEEE Transactions on Networking,* 1994, vol. 2, 1-15 **[0068]**
- **J. K. LENSTRA ; A. H. G. R. KAN.** Complexity of Scheduling under Precedence Constraints. *Operations Research,* 1978, vol. 26 (1), 22-35 **[0068]**
- **O. PAUL ; M. LAURENT.** A Full Bandwidth ATM Firewall. *Proceedings of the 6th European Symposium on Research in Computer Security ESORICS'2000,* 2000 **[0068]**
- **B. R. PREISS.** Data Structures and Algorithms with Object-Oriented Design Patterns in C++. John Wiley & Sons, 1999 **[0068]**

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1057

- **L. QUI ; G. VARGHESE ; S. SURI.** Fast Firewall Implementations for Software and Hardware-Based Routers. *Proceedings of ACM SIGMETRICS,* June 2001 **[0068]**
- **V. P. RANGANATH ; D. ANDRESEN.** A Set-Based Approach to Packet Classification. *Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems,* 2003, 889-894 **[0068]**
- **R. RIVEST.** On Self-Organizing Sequential Search Heuristics. *Communications of the ACM,* 1976, vol. 19 (2 **[0068]**
- **W. E. SMITH.** Various Optimizers for Single-Stage Production. *Naval Research Logistics Quarterly,* 1956, vol. 3, 59-66 **[0068]**
- **V. SRINIVASAN ; G. VARGHESE ; S. SURI ; M. WALDVOGEL.** Fast and Scalable Layer Four Switching. *Proceedings of ACM SIGCOMM,* 1998, 191-202 **[0068]**
- **S. SURI ; G. VARGHESE.** Packet Filtering in High Speed Networks. *Proceedings of the Symposium on Discrete Algorithms,* 1999, 969-970 **[0068]**
- **R. L. ZIEGLER.** Linux Firewalls. New Riders. 2002 **[0068]**
- **E. D. ZWICKY ; S. COOPER ; D. B. CHAPMAN.** *Building Internet Firewalls,* 2000 **[0068]**

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1058

Office de la propriété intellectuelle du Canada
Un organisme d'Industrie Canada

Canadian Intellectual Property Office
An Agency of Industry Canada

Brevet canadien / Canadian Patent

Le commissaire aux brevets a reçu une demande de délivrance de brevet visant une invention. Ladite requête satisfait aux exigences de la *Loi sur les brevets*. Le titre et la description de l'invention figurent dans le mémoire descriptif, dont une copie fait partie intégrante du présent document.

Le présent brevet confère à son titulaire et à ses représentants légaux, pour une période expirant vingt ans à compter de la date du dépôt de la demande au Canada, le droit, la faculté et le privilège exclusif de fabriquer, construire, exploiter et vendre à d'autres, pour qu'ils l'exploitent, l'objet de l'invention, sauf jugement en l'espèce rendu par un tribunal compétent, et sous réserve du paiement des taxes périodiques.

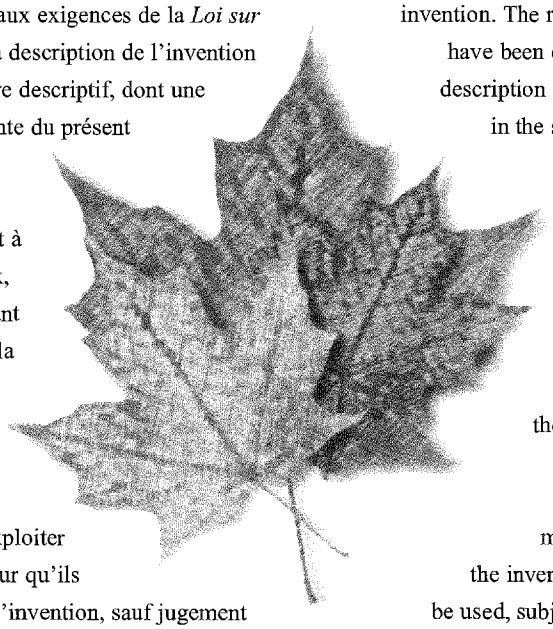

The Commissioner of Patents has received a petition for the grant of a patent for an invention. The requirements of the *Patent Act* have been complied with. The title and a description of the invention are contained in the specification, a copy of which forms an integral part of this document.

The present patent grants to its owner and to the legal representatives of its owner, for a term which expires twenty years from the filing date of the application in Canada, the exclusive right, privilege and liberty of making, constructing and using the invention and selling it to others to be used, subject to adjudication before any court of competent jurisdiction, and subject to the payment of maintenance fees.

BREVET CANADIEN **2,600,236** CANADIAN PATENT

Date à laquelle le brevet a été accordé et délivré	**2014/08/12**	Date on which the patent was granted and issued
Date du dépôt de la demande	**2006/03/28**	Filing date of the application
Date à laquelle la demande est devenue accessible au public pour consultation	**2006/10/05**	Date on which the application was made available for public inspection

Commissaire aux brevets / Commissioner of Patents

Canada

3256 (CIPO 91) 04/14

OPIC CIPO

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1059

Office de la Propriété Intellectuelle du Canada
Un organisme d'Industrie Canada

Canadian Intellectual Property Office
An agency of Industry Canada

CA 2600236 C 2014/08/12

(11)(21) 2 600 236

(12) BREVET CANADIEN
CANADIAN PATENT

(13) C

(86) Date de dépôt PCT/PCT Filing Date: 2006/03/28
(87) Date publication PCT/PCT Publication Date: 2006/10/05
(45) Date de délivrance/Issue Date: 2014/08/12
(85) Entrée phase nationale/National Entry: 2007/09/04
(86) N° demande PCT/PCT Application No.: US 2006/011291
(87) N° publication PCT/PCT Publication No.: 2006/105093
(30) Priorité/Priority: 2005/03/28 (US60/665,664)

(51) Cl.Int./Int.Cl. *G06F 15/16* (2006.01)
(72) Inventeurs/Inventors:
FULP, ERRIN W., US;
TARSA, STEPHEN J., US

(73) Propriétaire/Owner:
WAKE FOREST UNIVERSITY, US

(74) Agent: NORTON ROSE FULBRIGHT CANADA LLP/S.E.N.C.R.L., S.R.L.

(54) Titre : PROCEDES, SYSTEMES ET PRODUITS DE PROGRAMMES INFORMATIQUES POUR L'OPTIMISATION DE LA POLITIQUE DE PARE-FEU D'UN RESEAU

(54) Title: METHODS, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR NETWORK FIREWALL POLICY OPTIMIZATION

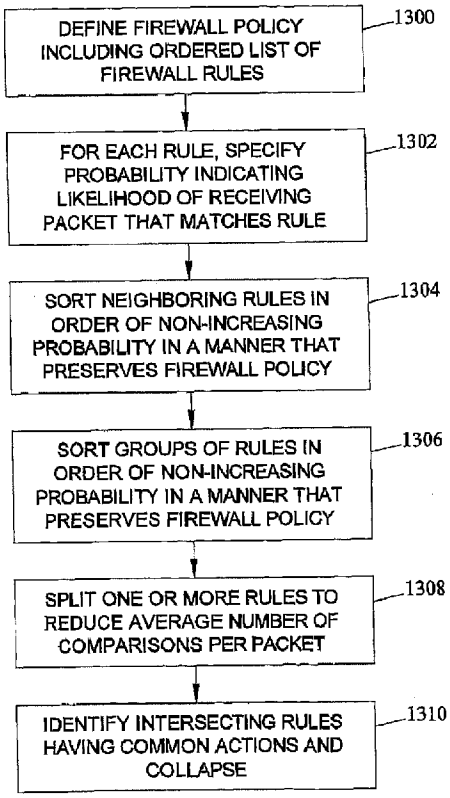

(57) Abrégé/Abstract:
Methods, systems, and computer program products for firewall policy optimization are disclosed. According to one method, a firewall policy including an ordered list of firewall rules is defined. For each rule, a probability indicating a likelihood of receiving a packet matching the rule is determined. The rules are sorted in order of non-increasing probability in a manner that preserves the firewall policy.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1060

METHODS, SYSTEMS, AND COMPUTER PROGRAM PRODUCTS FOR NETWORK FIREWALL POLICY OPTIMIZATION

GOVERNMENT INTEREST

This invention was made with Government support under Grant No. DE-FG02-03ER25581 awarded by U.S. Department of Energy, Mathematical and Computational Information Sciences Division. The Government has certain rights in the invention.

TECHNICAL FIELD

The subject matter described herein relates to optimizing polices for network firewalls. More particularly, the subject matter described herein relates to methods, systems, and computer program products for optimizing performance of network firewalls.

BACKGROUND ART

Network firewalls remain the forefront defense for most computer systems. Guided by a security policy, these devices provide access control, auditing, and traffic control [3, 30, 31]. A security policy is a list of ordered rules, as seen in Table 1, that defines the action to perform on matching, packets. For example, an accept action passes the packet into or from the secure network, while deny causes the packet to be discarded. In many implementations, the rule set is stored internally as a linked list. A packet is sequentially compared to the rules, starting with the first, until a match is found; otherwise, a default action is performed [30, 31]. This is referred to as a first- match policy and is used in many firewall systems including the Linux firewall implementation iptables [25].

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1061

No.	Proto.	Source IP	Source Port	Destination IP	Destination Port	Action	Prob.
1	UDP	1.1.*	*	*	80	deny	0.01
2	TCP	2.*	*	1.*	90	accept	0.02
3	UDP	*	*	1.*	*	accept	0.10
4	TCP	2.*	*	1.*	20	accept	0.17
5	UDP	1.*	*	*	*	accept	0.20
6	*	*	*	*	*	deny	0.50

Table 1: Example Security Policy Consisting of Multiple Ordered Rules

Traditional firewall implementations consist of a single, dedicated machine, similar to a router, that sequentially applies rules to each arriving packet. However, packet filtering represents a significantly higher processing load than routing decisions [24, 29, 31]. For example, a firewall that interconnects two 100 Mbps networks would have to process over 300,000 packets per second [30]. Successfully handling these traffic loads becomes more difficult as rule sets become more complex [4, 22, 31]. Furthermore, firewalls must be capable of processing even more packets as interface speeds increase. In a high-speed environment (e.g. Gigabit Ethernet), a single firewall can easily become a bottleneck and is susceptible to DoS attacks [4, 9, 13, 14]. An attacker could simply inundate the firewall with traffic, delaying or preventing legitimate packets from being processed.

One approach to increase firewall performance focuses on improving hardware design. Current research is investigating different distributed firewall designs to reduce processing delay [4, 9, 22], and possibly provide service differentiation [11]. Another approach focuses on improving performance via better firewall software [6, 7, 12, 16, 17, 24]. Similar to approaches that address the longest matching prefix problem for packet classification [8, 10, 25, 28], solutions typically represent the firewall rule set in different fashions (e.g. tree structures) to improve performance. While both approaches, together or separate, show great promise, each requires radical changes to the firewall system, and therefore are not amenable to current or legacy systems.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1062

Accordingly, there exists a need for improved methods, systems, and computer program products for network firewall policy optimization.

SUMMARY

The subject matter described herein includes methods, systems, and computer program products for network firewall policy optimization. According to one method, a firewall policy including an ordered list of firewall rules is defined. For each rule, a probability indicating a likelihood of receiving a packet matching the rule is defined. Neighboring rules are sorted in order of non-increasing probability in a manner that preserves the firewall policy.

As used herein, the term "firewall" refers to a logical entity that is adapted to filter packets at the ingress and/or egress portions of a network based on a policy. The term "firewall" is intended to include systems that block packets from entering or leaving a network and systems that perform intrusion detection and intrusion protection functions for packets entering or leaving a network. Thus, the methods and systems described herein for firewall policy optimization can be used to optimize policies for intrusion detection and intrusion protection firewalls.

The subject matter described herein includes a method to improve firewall performance and lower packet delay that can be applied to both legacy and current systems. In one implementation, a firewall rule set may be re-ordered to minimize the average number of rule comparisons to determine the action, while maintaining the integrity of the original policy. Integrity is preserved if the reordered and original rules always arrive at the same result. To maintain integrity, a firewall rule set may be modeled as a Directed Acyclical Graph (DAG), where vertices are firewall rules and edges indicate precedence relationships. Given this representation, any linear arrangement of the DAG is proven to maintain the original policy integrity. Unfortunately, determining the optimal rule order from all the possible linear arrangements is proven to be *NP*-complete, since it is equivalent to sequencing jobs with precedence constraints for a single machine [15]. Although determining the optimal order is *NP*-complete, a heuristic will be described below to order firewall rules that reduces the average number of comparisons while maintaining integrity. Simulation

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1063

results show the proposed reordering method yields rule orders that are comparable to optimal; thus, provides a simple means to significantly improve firewall performance and lower packet delay.

The subject matter described herein for network firewall policy optimization may be implemented using any combination of hardware, software, or firmware. For example, the subject matter described herein may be implemented using a computer program product comprising computer executable instructions embodied in a computer readable medium. Exemplary computer readable media suitable for implementing the subject matter described herein include disk memory devices, chip memory devices, programmable logic devices, application specific integrated circuits, and downloadable electrical signals. In addition, a computer program product that implements the subject matter described herein may be implemented using a single device or computing platform or may be distributed across multiple devices or computing platforms.

BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the subject matter described herein will now be explained with reference to the accompanying drawings of which:

Figure 1A is a block diagram of a rule list Directed Acyclical Graph (DAG) for the firewall rules in Table 1;

Figure 1B is a linear arrangement of firewall rules corresponding to the original rule order;

Figure 2A is a policy DAG representing a reordered version of the firewall rule set in Figure 1 after sorting;

Figure 2B is a policy DAG for the rule set illustrated in Table 1 illustrating an optimal rule order;

Figure 3A is a graph illustrating average numbers of packet comparisons versus intersection percentage for different orderings of a firewall policy;

Figure 3B is a graph illustrating the percent difference in average numbers of packet comparisons between sorted and optimal sortings of a firewall policy;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1064

Figure 4A is a graph of percent difference comparisons between original and sorted versions of a firewall policy;

Figure 4B is a graph illustrating the performance impact of considering rule intersection and actions in performing firewall rule sorting;

Figure 5A is a table illustrating a firewall policy before sorting;

Figure 5B is a table illustrating a sorted version of the firewall policy of Figure 5A;

Figure 6A is a policy trie representation of the firewall policy of Figure 5B;

Figure 6B is a table representing the firewall policy of Figure 5A after TCP and UDP sub-trie rotation;

Figure 7A is a policy trie illustrating the splitting of rule r_5 into rules r_{5a} and r_{5b};

Figure 7B is a table illustrating the firewall policy of Figure 5A after splitting rule r_5 into rules r_{5a} and r_{5b};

Figure 8A is a graph illustrating firewall performance with and without rule splitting;

Figure 8B is a graph illustrating locations and probabilities where splitting is better;

Figure 9A is a policy DAG of a firewall rule set;

Figure 9B is a policy DAG illustrating compression of rules r_1 and r_5;

Figure 10A is a policy DAG of a firewall rule list;

Figure 10B is an original policy trie before sorting;

Figure 10C is a policy trie illustrating the firewall rules after rotating sub-tries T_1 and T_2;

Figure 11A is a policy DAG illustrating a firewall policy;

Figure 11B is a policy trie illustrating firewall rules of the policy before sorting;

Figure 11C is a policy trie illustrating the firewall policy after exchanging nodes r_1 and r_2 and nodes r_7 and r_8;

Figure 11D is a policy trie illustrating the firewall policy after exchanging sub-tries T_3 and T_4;

Figure 11E is a policy trie illustrating the firewall policy after exchanging sub-tries T_1 and T_2;

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1065

Figure 12A is a policy DAG illustrating an original order for a firewall policy;

Figure 12B is a policy DAG illustrating a sorted version of the firewall policy of Figure 12A after exchanging neighbors;

Figure 12C is a policy DAG illustrating a sorted version of the firewall policy after exchanging lists L_1 including rules r_1 and r_3 and lists L_2 including rules r_2 and r_4;

Figure 13 is a flow chart illustrating exemplary steps for network firewall policy optimization according to an embodiment of the subject matter described herein;

Figure 14 is a block diagram illustrating exemplary components of a system for firewall policy optimization according to an embodiment of the subject matter described herein; and

Figure 15 is a block diagram illustrating a plurality of firewall nodes for implementing a firewall policy according to an embodiment of the subject matter described herein.

DETAILED DESCRIPTION OF THE INVENTION

Modeling Firewall Security Policies

As described above, a firewall rule set, also known as a firewall policy, is traditionally an ordered list of firewall rules. Firewall policy models have been the subject of recent research [1, 2, 16]; however, the primary purpose is anomaly detection and policy verification. In contrast, the policy model described in this herein is designed for firewall performance optimization and integrity. Firewall performance refers to reducing the average number of comparisons required to determine an action, while integrity refers to maintaining the original policy intent. Although improving the worst-case performance is important, it is not possible without changing the list-based representation [16, 24].

Firewall Rule and Policy Models

In the examples described herein, a rule r is modeled as an ordered tuple of sets, $r = (r[1], r[2], ..., r[k])$. Order is necessary among the tuples since

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1066

comparing rules and packets requires the comparison of corresponding tuples. Each tuple $r[i]$ is a set that can be fully specified, specify a range, or contain wildcards '*' in standard prefix format. For the Internet, security rules are commonly represented as a 5-tuple consisting of: protocol type, IP source address, source port number, IP destination address, and destination port number [30, 31]. Given this model, the ordered tuples can be supersets and subsets of each other, which forms the basis of precedence relationships. In addition to the prefixes, each filter rule has an action, which is to accept or deny. However, the action will not be considered when comparing packets and rules. Similar to a rule, a packet (IP datagram) d can be viewed as an ordered k-tuple $d = (d[1], d[2], ..., d[k])$; however, ranges and wildcards are not possible for any packet tuple.

Using the previous rule definition, a standard security policy can be modeled as an ordered set (list) of n rules, denoted as $R = \{r_1, r_2, ..., r_n\}$. A packet d is sequentially compared against each rule r_i starting with the first, until a match is found ($d \Rightarrow r_i$) then the associated action is performed. A match is found between a packet and rule when every tuple of the packet is a subset of the corresponding tuple in the rule.

Definition Packet d matches r_i if

$$d \Rightarrow r_i \quad \text{iff} \quad d[i] \subseteq r_i[i], \quad i = 1, ..., k$$

The rule list R is comprehensive if for every possible legal packet d a match is found using R. Furthermore, two rule lists R and R' are equivalent if for every possible legal packet d the same action is performed by the two rule lists. If R and R' are different (e.g. a reorder) and the lists are equivalent, then the policy integrity is maintained.

As previously mentioned, a rule list has an implied precedence relationship where certain rules must appear before others if the integrity of the policy is to be maintained. For example consider the rule list in Table 1. Rule r_1 must appear before rules r_3 and r_5, likewise rule r_6 must be the last rule in the policy. If for example, rule r_3 was moved to the beginning of the policy, then it will shadow [2] the original rule r_1. However, there is no precedence relationship between rules r_2 and r_4 given in Table 1. Therefore, the relative ordering of these two rules will not impact the policy integrity and can be

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1067

changed to improve performance. The present example assumes that the original policy is free from any anomalies. Likewise, when a policy is reordered to improve performance it should not introduce any anomalies, which will occur if precedence relationships are not maintained. As a result, a model is needed to effectively represent precedence relationships.

Modeling Rule List Precedence Relationships

The precedence relationship between rules in a policy will be modeled as a Directed Acyclical Graph (DAG) [23, 18]. Such graphs have been successfully used to represent the relative order of individual tasks that must take place to complete a job (referred to as a task graph model). Since certain rules must appear before others to maintain policy integrity, this structure is well suited for modeling the precedence of firewall rules. Let $G = (R, E)$ be a rule list DAG for a rule list R, where vertices are rules and edges E are the precedence relationships (constraint). A precedence relationship, or edge, exists between rules r_i and r_j, if $i < j$, the actions for each rule are different, and the rules intersect.

Definition The intersection of rule r_i and r_j, denoted as $r_i \cap r_j$ is

$$r_i \cap r_j = (r_i[l] \cap r_j[l]), \qquad l = 1, \ldots, k$$

Therefore, the intersection of two rules results in an ordered set of tuples that collectively describes the packets that match both rules. The rules r_i and r_j intersect if every tuple of the resulting operation is non-empty. In contrast, the rules r_i and r_j do not intersect, denoted as $r_i \not\cap r_j$, if at least one tuple is the empty set. Note the intersection operation is symmetric; therefore, if r_i intersects r_j, then r_j will intersect r_i. The same is true for rules that do not intersect.

For example consider the rules given in Table 1, the intersection of r_1 and r_3 yields (UDP, 1.1.*, *, 1.*, 80). Again, the rule actions are not considered in the intersection or match operation. Since these two rules intersect, a packet can match both rules for example d = (UDP, 1.1.1.1, 80, 1.1.1.1, 80). Furthermore, the actions of the two rules are different. Therefore, the relative order must be maintained between these two rules and an edge drawn from r_1 to r_3 must be present in the rule list DAG, as seen in Figures 1A and 1B. More

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1068

particularly, Figure 1A illustrates a rule list DAG for the rules in Table 1. Figure 1B illustrates a linear arrangement of the rules in Figure 1A. In Figures 1A and 1B, the vertices represent rules, the circles represents an accept rules, and the squares represents deny rules. Edges that connect the vertices represent precedence requirements. As can be seen from Figure 1B, because of the edge between rules r_1 and r_3, precedence between rules r_1 and r_3 must be maintained. In contrast consider the intersection of rules r_3 and r_5. These two rules intersect, indicating packets belonging to the set (UDP, 1.*, *, 1.*, *) would match both rules. However, it does not matter which of the two rules a packet matches first, since the action is the same for both rules. Therefore, an edge does not exist between rules r_3 and r_5 in the diagram. Similarly, rules r_2 and r_4 do not intersect due to the fifth tuple (destination port). A packet cannot match both rules indicating the relative order can change; therefore, an edge will not exist between them.

The match operation can be used to identify precedence relationships, but it cannot do so in every case. Consider a partial-match example [1], where r_a = (UDP, *, 80, 10.*, 90, accept) and r_b = (UDP, 10.*, 80, *, 90, deny). The intersection of r_a and r_b is (UDP, 10.*, 80, 10.*, 90); therefore a packet, such as d = (UDP, 10.10.10.10, 80, 10.10.10.10, 90), can match both rules. If r_a appears before r_b then the packet d is accepted, but if r_b occurs before r_a then d is rejected. As a result, the order of r_a and r_b in the original policy must be maintained. However, the match operation is unable to identify the precedence in this example. A partial match exists in between rules r_3 and r_5 in Table 1, but as previously discussed an edge does not exist between the rules since the actions are the same.

Using the rule list DAG representation a linear arrangement is sought that improves the firewall performance. As depicted in Figure 1B, a linear arrangement (permutation or topological sort) is a list of DAG vertices where all the successors of a vertex appear in sequence after that vertex [23]. Therefore it follows that a linear arrangement of a rule list DAG represents a rule order, if the vertices are read from left to right. Furthermore, it is proven in the following theorem that any linear arrangement of a rule list DAG maintains integrity.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1069

Theorem Any linear arrangement of a rule list DAG maintains integrity.

Proof Assume a rule list DAG *G* is constructed from the security policy *R* that is free of anomalies. Consider any two rules r_i and r_j in the policy, where $i < j$. If an edge between r_i and r_j in *G* does not exist, then a linear arrangement of *G* can interchange the order of the two rules. An edge will not exist if the rules do not intersect; however, a reorder will not affect integrity since a packet cannot match both rules. Shadowing is not introduced due to the reorder since the intersection operation is symmetric. An edge will not exist if the two rules intersect but have the same action; however, a reorder will not affect integrity since the same action will occur regardless of which rule is matched first. If an edge does exist between the rules, then their relative order will be maintained in every linear arrangement of *G*; thus maintaining precedence and integrity.

Rule List Optimization

As described above, it is important to inspect packets as quickly as possible given increasing network speeds and QoS requirements. Using the rule list DAG to maintain policy integrity, a linear arrangement is sought that minimizes the average number of comparisons required. However, this will require information not present in the firewall rule list. Certain firewall rules have a higher probability of matching a packet than others. As a result, it is possible to develop a policy profile over time that indicates frequency of rule matches (similar to cache hit ratio). Let $P = \{p_1, p_2, ..., p_n\}$ be the policy profile, where p_i is the probability that a packet will match rule *i* (first match in the policy). Furthermore, assume a packet will always find a match, $\sum_{i=1}^{n} p_i = 1$; therefore *R* is comprehensive. Using this information, the average number of rule comparisons required is

$$E[n] = \sum_{i=1}^{n} i \cdot p_i \qquad (1)$$

For example, the average number of comparisons required for the rule set in Table 1 is 5.03.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1070

Given a rule list DAG $G = (R, E)$ and policy profile $P = \{p_1, p_2, ..., p_n\}$ a linear arrangement π of G is sought that minimizes Equation 1. In the absence of precedence relationships, the average number of comparisons is minimized if the rules are sorted in non-increasing order according to the probabilities [26], which is also referred to as Smith's algorithm [27]. Precedence constraints cause the problem to be more realistic; however, such constraints also make determining the optimal permutation more problematic.

Determining the optimal rule list permutation can be viewed as job scheduling for a single machine with precedence constraints [15, 21]. The notation for such scheduling problems is $\alpha|\beta|\gamma|\delta$, where α is the number of machines, β is the precedence (or absence of) which can be represented as a DAG, γ is a restriction on processing time, and δ is the optimality criterion [15]. Determining the optimal rule order is similar to the $1|\beta|1|\sum \omega C_i$ scheduling problem, or optimality criterion, where ω is a weight associated with a job (for example, importance) and C_i is the completion time. As previously noted, the $1||\sum \omega C_i$ problem can be solved in linear time the using Smith's algorithm [27], which orders jobs according to non-decreasing $\frac{t_i}{\omega}$ ratio, where t_i is the processing time of job i. In this case set $t_i = 1$ and $\omega = p_i \ \forall i$. However, Lawler [19] and Lenstra et al. [21] proved $1|\beta|1|\sum \omega C_i$ to be *NP*-complete via the linear arrangement problem, which implies determining the optimal firewall rule order is also *NP*-complete. Note, determining the number of possible permutations has been proven to be *NP*-hard [5].

Theorem $1|\beta|1|\sum \omega C_i$ α *Determining the optimal order of a firewall rule list*

Proof Consider the $1|\beta|1|\sum \omega C_i$ problem. Each of n jobs J_i, $i \in I$, has to be processed without preemption on a single machine that can handle at most one job at a time. For each $i \in I$, let ω be the associated weight. Furthermore, let $G = (V, E)$ be a DAG that represents the precedence order of the jobs J_i. Assume the processing time of each job equals 1 time unit, the weights to be $0 \leq \omega \leq 1$ such that $\sum \omega = 1$, and β, which is G, to be a rule list DAG. In this case, the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1071

optimization criterion $\sum \omega \cdot C_i$ is the same as $\sum p_i \cdot i$, which is given in equation 1. Clearly, the optimal firewall rule ordering problem has a solution if and only if $1|\beta|1|\sum \omega C_i$ has a solution. Therefore, determining the optimal permutation of firewall rules is *NP*-complete.

Exemplary Rule Sorting Algorithm

Although determining the optimal rule permutation was proven to be *NP*-complete, reducing the average number of comparisons required to process a packet remains an important objective. As previously discussed, a sorting algorithm must maintain the precedence relationships among rules. Of course an exhaustive search is possible if the number of rules is small (generate and test all possible topological orders); however as proven in the previous section, this is not feasible with a realistic rule list.

One exemplary algorithm starts with the original the rule set, then sorts neighboring rules based on nonincreasing probabilities. However, an exchange of neighbors should never occur if the rules intersect and have different actions. This test preserves any precedence relationships in the policy. For example, the following sorting algorithm uses such a comparison to determine if neighboring rules should be exchanged. Note a_i denotes the action associated with the i^{th} rule.

```
done = false
while(!done)
    done = true
    for(i = 1; i <n; i++)
        if(p_i < p_i+1 AND (r_i ∩ r_i+1 OR a_i == a_i+1))then
            exchange rules, actions, and probabilities
            done = false
        endif
    endfor
endwhile
```

Other sorting algorithms are possible if the completion time of the sort is an issue; however, the method presented is easy to implement and only requires a simple neighbor comparison. Assume the match probabilities for the rule list given in Table 1. Applying the sorting algorithm to this rule list results in the ordering depicted in Figure 2A, which has 10% fewer comparisons on

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1072

average. However when using the algorithm, it is possible that one rule can prevent another rule from being reordered. For example, rule r_1 prevents rule r_5 from being placed closer to the beginning of the rule set. However, if both rules r_1 and r_5 are placed closer to the beginning of the policy while maintaining their relative order, the average number of comparison will be further reduced, 15% fewer. This is the optimal order for the 6 rule set, which is depicted in Figure 2B. Although this simple sorting algorithm is unable to move groups of rules, it can still improve the performance of the firewall system. Its effectiveness is measured experimentally in the next section.

Experimental Results

In this section, the average number of rule comparisons required after sorting the rules is compared with average number of comparisons required using the original order and the optimal order. Note the optimal rule ordering was determined via an exhaustive search. As a result, it was not feasible to determine the optimal ordering once the number of rules equaled 15 given the number of permutations to consider. Two different variations of the rule sorting method described in the preceding section were implemented. The difference between the methods was the comparison used to determine if neighboring rules should be exchanged. The first sorting algorithm exchanged neighboring rules if they were out of order and did not intersect. The if-condition was as follows.

if ($p_i < p_{i+1}$ **AND** $r_i \cap r_{i+1}$) **then**

Therefore, rule actions were not considered and the method will be referred to as *non-action sort*. The second sorting method did consider the rule action (as described in the preceding section) and will be referred to as *action sort*. The comparison between the two sorting algorithms will indicate the importance of considering rule actions when ordering rules.

In the first experiment, lists of 10 firewall rules were generated with random precedence relationships. The match probability of each rule was given by a Zipf distribution [20], which assigns probabilities according to the rank of an item. For this simulation, the last rule had the highest probability

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1073

which is consistent with most policies (last rules are more general). As a result, the original order yields the worst average number of comparisons. The intersection percentage measured the percentage of rules that intersect in the policy. This metric gives a value for the dependency level in the policy; however note, rule action is not considered when calculating this metric. Rules were equally likely to have an accept or deny action and the results (average number of rule comparisons) for a particular intersection percentage were averaged over 1000 simulations.

Results of the first experiment are given in Figures 3A and 3B. More particularly, Figure 3A is a graph illustrating the average number of packet comparisons versus intersection percentage for the original rule sorting, the non-action sort, the action sort, and the optimal sort. Figure 3B is a graph illustrating the percent difference in the average number of comparisons required between the sorted and optimal firewall policy configurations. The average number of comparisons required was lower for the sorted and the optimal lists when the intersection percentage was low, as seen in Figure 3A. This was expected since there is a large number of possible rule permutations (few edges in the DAG). When the intersection percentage approached 100%, the values converged to the number required for the original list. This is due to the limited number of possible rule orders, only one order in the extreme case. The percent difference between the sorted and optimal order is shown in Figure 3B. At zero intersection percentage the sorted and optimal orders are equal, since any ordering is possible. Similarly at an intersection percentage of 100, the sorted and optimal orders are equal since only one order is possible. Between these two extremes, the action sorting algorithm remains close to the optimal value. In contrast, the non-action sort had a maximum difference of 33%.

The benefit of rule sorting on larger policies is also of interest; however as previously described, it was not possible to determine the optimal ordering once the number of rules approaches 15. Therefore, the second experiment used larger rule sets, but only compared the original order and the sorted orders. The number of rules ranged from 10 to 1000, while the matching probabilities and intersection percentages were the same as the previous

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1074

experiment. The results of this experiment are depicted in Figure 4A. The sorted rule sets always performed equal to or better than the original order, while the action sort consistently performed better than the non-action sort. As noted in the previous experiment, the percent difference is very large given few intersections (e.g. 80% decrease for 1000 rules with 0% dependency), but approaches zero as all the rules intersect. As the number of rules increases, sorting is increasingly beneficial, although only at low intersection percentages. The benefit of sorting drops more significantly as the intersection percentage increases with larger rule sets. This is primarily due to the low matching probabilities of each rule, which requires a complete reordering to have a significant impact on the average number of comparisons. As a result, large rule sets can benefit from sorting if the intersection percentage is low.

The impact of considering rule actions when sorting is illustrated in Figure 4B. In this experiment lists of 1000 rules were generated, where the intersection percentages ranged from 0% to 100% and the percentage of rules with the same action varied from 50% to 100%. The performance of the action sort to the original list was then compared. As the percentage of rules with the same action increased, the percent difference (reduction) in the average number of comparisons also increased. This is due to the increased number of permutations possible when the rule actions are increasingly the same (fewer edges in the rule list DAG to consider). This is depicted in Figure 4B where a policy with a 100% intersection percentage can significantly reduce the number of comparisons (80% reduction) if all the rules have the same action. This performance increase is not possible with the non-action sort. Therefore, considering the rule action increases the number of possible rule orders, thereby providing more possibilities to improve firewall performance.

Conclusions

Network firewalls enforce a security policy by comparing arriving packets to a list of rules. An action, such as accept or deny, is then performed on the packet based on the matching rule. Unfortunately packet filtering can impose significant delays on traffic due to the complexity and size of rule sets.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1075

Therefore, improving firewall performance is important, given network Quality of Service (QoS) requirements and increasing network traffic loads.

The sections above describe rule ordering methods to improve the performance of network firewalls. Assuming each rule has a probability of a packet matching, firewall rules should be sorted such that the matching probabilities are non-increasing. This reduces the average number of comparisons and the delay across the firewall. However, a simple sort is not possible given precedence relations across rules. It is common in a security policy that two rules may match the same packet yet have different actions. It is this precedence relationship between rules that must be maintained to preserve integrity. The method described above uses Directed Acyclical Graphs (DAG) to represent the precedence order rules must maintain. Given this representation, a topological sort can be used to determine the optimal order (minimum average number of comparisons); however, the examples above prove this problem to be *NP*-complete (similar to job scheduling for a single nonpreemptive machine with precedence constraints). As an alternative, a simple sorting method was introduced that maintained the precedence order of the rules. Simulation results indicate this method can significantly reduce the average number of comparisons required and is comparable to optimal ordering.

Several areas exist for future research in optimizing firewall rule lists. The sorting method proposed above is based on a simple algorithm. Although it can offer an improvement over the original rule order, algorithms that can move groups of rules could provide a larger reduction in the average number of comparisons. The effect of stateful firewalls should also be addressed in future research. Security can be enhanced with connection state and packet audit information. For example, a table can be used to record the state of each connection, which is useful for preventing certain types of attacks (e.g., TCP SYN flood) [30, 31]. The impact of such rules on the firewall needs to be investigated and whether sorting can be done on-line to reflect temporal changes. In addition, more research is needed to determine more accurate probability distributions for packet matching and dependency percentages.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1076

Given this information, better algorithms can be designed and more realistic simulations can be performed.

Trie-Based Network Firewall Optimization Methods and Rule Splitting

Overview

As described above, while assuring firewall policy integrity is critical, performance is equally important given increasing network speeds and traffic volumes. Firewall processing delay can be reduced via hardware as well as policy optimization, where rules are manipulated to reduce the number of comparisons. For a given firewall installation, it can be determined that certain security rules have a higher probability of matching a packet than others. As a result, it is possible to develop a *policy profile* over time that indicates frequency of rule matches (similar to cache hit ratio). Given this information, trie-based methods are described herein to minimize the average number of comparisons required per packet. Although most firewall systems still utilize an ordered list representation, the proposed enhancements still are applicable to current firewall systems since a policy trie can be converted into an ordered list using an inorder traversal.

Rule Sorting

As described above, rule sorting of neighboring rules based on ascending probabilities in a manner that considers rules intersection and rule actions can be used to improve firewall performance by reducing the number of packet comparisons. The methods described above use Directed Acyclic Goal Graphs (DAGs) to represent firewalls. The methods described in this section will use the rule sorting algorithms in combination with policy trie representations to perform sorting of groups of rules in a manner that improves firewall performance.

In this example, it is assumed that a security policy contains a list of n ordered rules, $\{r_1, r_2, ..., r_n\}$. In addition, let p_i represent the probability that a packet will match rule i (first match in the policy). Therefore, the policy is comprehensive, which means every packet will have a match, if $\sum_{i=1}^{n} p_i = 1$.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1077

Using this information, the average number of rule comparisons required can be calculated as

$$E[n] = \sum_{i=1}^{n} i \cdot p_i$$

The average number of comparisons is minimized if the rules are sorted in non-ascending order according to the probabilities. However, this can only be achieved if the sorting algorithm considers the probabilities and whether one rule is a superset of another. For example, sorting algorithms must use the following comparison to determine if neighboring rules should be interchanged Note, a_i denotes the action associated with the i^{th} rule.

if ($p_i < p_{i+1}$ **AND** $r_i \not\subset r_{i+1}$ **OR** $a_i == a_{i+1}$)) **then**

Sorting rules in this fashion can have positive impact on the average number of comparisons required. Figure 5A is a table illustrating a firewall policy prior to rule sorting. The expected number of comparisons for the policy of Figure 5A is 4.26 comparisons per packet. Figure 5B is a table illustrating the firewall policy of Figure 5A after sorting using the above-described algorithm. In Figure 5B, the expected number of comparisons per packet has been reduced to 3.94. Thus, using the method described above, the number of comparisons per packet can be reduced by 8%.

However, it can be seen from Figures 5A and 5B that one rule can prevent another rule from being reordered. For example in Figure 5B, rule r_1 prevents rule r_4 from being placed closer to the beginning of the rule set. However, if both rules r_1 and r_4 are placed closer to the beginning of the policy while maintaining their relative order, the average number of comparison will be reduced.

To solve this problem, rule sets can be sorted a using policy tries. First, the rule list is converted into an equivalent policy trie. Each sub-trie will have an associated probability $\overline{p}$ that is average probability (hit-ratio) of the rules comprising the sub-trie. Sub-tries can be rotated around their parent node to increase performance, using the method described earlier. Since the policy trie (or equivalent rule set) is tested from left to right, rotation should occur if the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1078

probability of the right sub-trie is greater than the left sub-trie. As a result, the average number of comparisons required will be reduced and the policy integrity is maintained. Figure 6A is a trie representation of the rules in Figure 5B. Figure 6B is a table representing resulting rule set after rotating the TCP and UDP sub-tries. The TCP sub-trie in Figure 6A has an expected packet matching probability of 0.095 while the UDP sub-trie has an expected packet matching probability of 0.155. Thus, the UDP sub-trie and the TCP sub-trie should be rotated so that the UDP sub-trie becomes the left sub-trie and the TCP sub-trie becomes the right sub-trie. In the table illustrated in Figure 6B, which lists the ordering of the rules after rotating the sub-tries, the average number of comparisons per packets has been reduced to 3.70.

Rule Splitting

Rule splitting takes a general rule and creates more specific rules that collectively perform the same action over the same set of packets. Here, rule splitting is used to reduce the average number of comparisons. For example in Figure 7A, rule r_5 is split into two separate rules, r_{5a} for UDP and r_{5b} for TCP. Figure 7B is a table corresponding to the policy trie of Figure 7A. Once the rules are positioned based on their probabilities and their relation to other rules, the average number of rule comparisons is reduced to 2.98 (after sort, trie rotation, and splitting) which is a 30% less.

It many not be advantageous to split a general rule since it adds another rule to the policy. For example, assume a policy contains 20 rules where the first 19 rules have the same probability. Assume the last rule can be split and the new specific rule has a probability that is $\frac{3}{4}$ of the last rule. The impact of the probability of the last rule and the location of the new split rule, *m*, is depicted in Figures 8A and 8B. More particularly, Figure 8A is a graph illustrating the average number of comparisons for firewalls with and without rule splitting. Figure 8B is a graph illustrating average number of comparisons versus different locations of the new rule and probabilities. The average number of comparisons is reduced as the split rule is located closer to the first rule (surface decreases as *m* approaches one). Furthermore, splitting yields

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1079

better results when the general rule has a high probability. However as illustrated in Figure 8B, the closer the specific rule is to the location of the original rule the average number of comparisons increases, which is the penalty of adding one more rule to the policy.

The effect of splitting a single rule can be described mathematically as follows. Consider n rules where rule r_n can be split into rules $r_{n,1}$ and $r_{n,r}$ (whose union is the original rule). In addition, assume the split rule $r_{n,1}$ will be located at the m^{th} position ($1 \leq m < n$), while rule $r_{n,r}$ will remain at the n^{th} location. The goal is to determine the best location m, which yields a lower average number of comparisons as compared to the original rule set. This can be defined mathematically in the following formula.

$$\sum_{i=1}^{n} i \cdot p_i > \sum_{i=1}^{m-1} i \cdot p_i + m \cdot p_{n,l} + \sum_{i=m}^{n-1} (i+1) \cdot p_i + (n+1) \cdot p_{n,r}$$

The left side of the inequality is the average number of comparisons for the original rule set. The right hand side of the inequality is the average number of comparisons with the specific rule at location m. If it is assumed that the rules located between m and n have an equal probability (denoted as p) the previous equation can be solved for m.

$$m < \frac{n \cdot p - n \cdot p_n + (n+1) \cdot p_{n,r}}{p - p_{n,l}}$$

The new rule must be located between the first and m^{th}; however, its final location will depend on the relationship with the other rules (cannot be placed before any rule for which it is a superset). This result can be applied iteratively to multiple rules and repeatedly to the same rule.

Rule Optimization for QoS

In the preceding sections, optimizing the entire security policy has been discussed. However, it may be desirable to optimize the rules for a certain type of traffic, in order to reduce the average number of comparisons this traffic encounters. This can be done by organizing the policy trie via traffic

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1080

classification, then optimizing the sub-tries. The result is an ordering where the rules for the most important traffic are tested first.

Network Firewall Optimization Using Rule Compression

Another method for network firewall optimization is referred to as rule compression. Given a security policy R, one objective in optimizing firewall performance is to reduce number of rules in the policy while maintaining the original integrity. A method for reducing the number of rules called *rule compression*, which combines and replaces two rules in the policy will now be described. Although compression does reduce the number of rules, it can only be done if the integrity of the original policy is maintained.

Consider two rules r_i and r_j in policy R, where r_i appears before r_j. The two rules can be compressed if and only if the following conditions are true

1. The two rules intersect, $r_i \cap r_j \neq \emptyset$ (an edge between r_i and r_j exists in the policy DAG) (Tuples that are proper subsets and supersets are considered. Compression may still occur if the rules do not intersect or are adjacent in the set space, the resulting rule may contain ranges.)
2. The action of r_i and r_j is the same.
3. Rule r_i is not dependent on any rule r_k that is directly or indirectly dependent on r_j (only one path from r_i to r_j exists in the policy DAG).

The result of compression is a new rule, $r_{i,j}$ whose tuples are the union of the corresponding tuples of rules r_i and r_j. The new rule $r_{i,j}$ replaces both r_i and r_j in R; however, the location of $r_{i,j}$ in R may require the relocation of other rules. Let D_j be the ordered set of rules that appear after r_i that directly or indirectly depend upon r_j. The new rule is placed at the original location of r_i in R and the rules D_j are placed before $r_{i,j}$.

For example, consider the policy given in Table 2. It is possible to compress rules r_1 and r_5, creating the new rule $r_{1,5}$ =[TCP, *, *, 2.*, *, accept]. The ordered set D_j in this example consists of the rules r_3 and r_4. As a result, the new rule $r_{1,5}$ is placed at the original location of r_1 while the rules in D_j are placed before the new rule. Figures 9A and 9B respectively illustrate the

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1081

original policy DAG for Table 2 and the policy DAG after rules r_1 and r_5 are compressed.

No.	Proto.	Source IP	Source Port	Destination IP	Destination Port	Action
1	TCP	1.1.*	*	2.2.*	*	accept
2	TCP	1.*	*	*	80	deny
3	TCP	3.3.*	*	*	80	accept
4	TCP	3.*	*	*	*	deny
5	TCP	*	*	2.*	90	accept

Table 2: Example Policy Consisting of Multiple Ordered Rules

Theorem *Compressing rules r_i and r_j of policy R will maintain integrity if the three conditions are met and method described is used.*

Proof Before r_i and r_j are compressed, relocate the rules D_j immediately before r_i and D_i. This relocation will maintain integrity due to the third requirement for compression (no rule in D_i can be dependent on any rule in D_j). Now place r_j directly after r_i. This relocation will not affect integrity since no rule in D_i can have an edge to r_j (again, due to the third condition). Compressing the neighboring rules r_i and r_j creates $r_{i,j}$. This does not affect integrity since the resulting rule only matches the packets that match r_i or r_j in the original policy. Therefore, the result is the compression of rule r_i and r_j, and the integrity of the policy is maintained.

Determining if Compression Should Occur

Although compression may be possible, it may not improve the performance of the policy due to the relocation of rules. Compression should only occur if the average number of comparisons required for the new policy R_c is less than the original policy R; therefore, $E[R_c] < E[R]$.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1082

Firewall Rule Sub-Trie and Sub-List Ordering

Policy Sub-Trie Ordering

Additional methods for firewall optimization are referred to as firewall sub-trie and sub-list ordering. Given a policy trie T, it may be desirable to order all sub-tries such that the average number of tuple comparisons is reduced and policy integrity is maintained. Consider a policy trie T that contains sub-tries T_i and T_j having the same parent node, as seen in Figures 10A-10C. Let P_i denote the sum of the probabilities of the rules contained in sub-trie i, while C_i denotes the number of comparisons required to completely traverse sub-trie i (which is equal to the number of branches). In order to reduce the average number of tuple comparisons, sub-tries that share the same parent node should be ordered such that the following two conditions are observed.

1. Sub-tries that share a parent node are ordered such that the P_i values are non-ascending (higher match probabilities occur first, from left to right).
2. If sub-tries that share a parent node have the same probability (P_i equals P_j), then order the sub-tries such that the C_i values are non-descending (sub-tries consisting of fewer comparisons occur first, from left to right).

These conditions are recursively applied throughout the policy trie as seen in Figures 11A-11D. However, these conditions do not necessarily maintain policy integrity.

Consider a group of n sub-tries that share the same parent node that are numbered sequentially from left to right. Consider any two non-intersecting sub-tries T_i and T_{i+k} in this group that are out of order. If two sub-tries do not intersect, denoted as $T_i \pitchfork T_j$, then no rule in one sub-trie will intersect with any rule in the other sub-trie, $r_k \pitchfork r_l$, $\forall r_k \in T_i$, $\forall r_l \in T_j$. The following *integrity condition* must be observed when reordering the sub-tries. The two sub-tries can be exchanged (rotated about the parent) if and only the sub-tries T_i through T_{i+k-1} do not intersect with T_{i+k} and the sub-tries T_{i+1} through T_{i+k} do not intersect with T_i. Similar to finding the linear sequence of a policy DAGs, these conditions maintain the policy trie integrity. This is observed in Figure 10A,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1083

where no edge exists between any rule in T_1 and T_2. More particularly, Figure 10A is a policy DAG of a six-rule firewall list. Figure 10B is a policy trie representation of the firewall rule list represented by Figure 10A. In Figure 10B, the average number of tuple compares required per packet is 5.5. Figure 10C illustrates the policy trie after rotating sub-tries T_1 and T_2. In Figure 10C, the average number of tuple compares required is 4.5. In Figures 10A-10C, the following probabilities are assumed:

$P_1 = 0.5$

$C_1 = 4$

$P_2 = 0.5$; and

$C_2 = 2$.

The average number of tuple comparisons is an estimate based on the longest trie path.

Using the conditions for comparing and ordering, the following sorting algorithm sorts neighboring sub-tries that share the same parent node.

```
m = a parent node in T
done = false
while(!done)
    done = true
    for each sub-trie Ti having parent m AND a right neighboring sub-trie
        if((Pi < Pi+1 AND (Pi == Pi+1 AND Ci > Ci+1)) AND Ti ⋔ Ti+1) then
        interchange Ti and Ti+1
        done = false
        endif
    endfor
endwhile
```

Figures 11A-11D illustrate policy trie rotation for an eight-rule firewall policy. More particularly, Figure 11A illustrates the policy DAG representation of an eight-rule firewall policy. Figure 11B is a policy trie representation of the original firewall policy. In this example, ordering occurs in three stages, starting at the bottom level and moving towards the top level of the trie. In Figure 11C, the leaves are ordered at the lowest level. More particularly, nodes r_1 and r_2 are exchanged with nodes r_7 and r_8. The average number of tuple compares for the ordering represented by Figure 11C is 10.45. In Figure 11D, the sub-

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1084

tries on the second lowest level (T_3, T_4, T_5, and T_6) are ordered. More particularly, Figure 11D illustrates the policy trie after exchanging sub-tries T_3 and T_4. The average number of tuple compares for the representation in Figure 11D is 10.275. In Figure 11E, the sub-tries on the second highest level, (T_1 and T_2) are ordered. Figure 11E illustrates the final ordering after exchanging T_1 and T_2. In the policy trie of Figure 11E, the average number of tuple compares required is 6.775. The average number of tuple comparisons is an estimate based on the longest trie path.

Policy Sub-List Ordering

The preceding section describes conditions for sorting policy sub-tries, which allow the exchange of groups of rules (sub-tries). The same conditions can be applied to list based policies, where sub-lists are ordered to improve the average number of rule comparisons.

In the sections above relating to rule sorting, a method is described to exchange neighboring rules in a list-based security policy, as seen in Figure 12B. However, it was observed that when ordering rules in this fashion it is possible that one rule can block the exchange of another. If groups of rules were allowed to be exchanged the average number of rule comparisons could be further reduced.

Given a policy list L, it may be desirable to order all sub-lists such that the average number of rule comparisons is reduced and policy integrity is maintained. Consider a policy list L that contains sub-tries L_i and L_j. Let P_i denote the sum of the probabilities of the rules contained in sub-list i, while C_i denotes the number of rules in sub-list i. In order to reduce the average number of rule comparisons, sub-lists should be ordered such that the following two conditions are observed.

1. Sub-lists are ordered such that the P_i values are non-ascending (higher match probabilities occur first, from left to right).
2. If sub-lists have the same probability (P_i equals P_j), then order the sub-lists such that the C_i values are non-descending (sub-lists consisting of fewer comparisons occur first, from left to right).

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1085

These conditions are applied throughout the policy list; however as with policy tries, these conditions do not necessarily maintain policy integrity.

Consider a group of n sub-lists that are numbered sequentially from left to right. Consider any two non-intersecting sub-lists L_i and L_{i+k} in this group that are out of order. If two sub-lists do not intersect, denoted as $L_i \cap L_j$, then no rule in one sub-list will intersect with any rule in the other sub-list, $r_k \cap r_l$, $\forall r_k \in L_i$, $\forall r_l \in L_j$. The following *integrity condition* must be observed when reordering the sub-lists. The two sub-lists can be exchanged if and only if the sub-lists L_i through L_{i+k-1} do not intersect with any rule in L_{i+k} and sub-lists L_{i+1} through L_{i+k} do not intersect with L_i. Similar to finding the linear sequence of a policy DAGs, these conditions maintain the policy trie integrity. For example, consider the sub-lists $L_1 = \{r_1, r_3\}$ and $L_2 = \{r_2, r_4\}$ given in Figure 12B. No edge exists between any rule in L_1 and L_2. The two sub-lists are out of order with respect to P and should be exchanged. The resulting list is given in Figure 12C. Figure 12A illustrates the original policy order and results in an average number of rule comparisons of 3.4. In the sorted version in Figure 12B, the average number of rule comparisons is 3.3. In Figure 12C, the average number of rule comparisons has been reduced to 2.7. Using the conditions for comparing and ordering, the following sorting algorithm sorts neighboring sub-lists.

```
done = false
while(!done)
    done = true
    for each sub-list Li having a right neighboring sub-list
        if ((Pi < Pi+1 OR (Pi == Pi+1 AND Ci > Ci+1)) AND Li ∩ Li+1)then
            interchange rules and probabilities
            done = false
        endif
    endfor
endwhile
```


CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1086

Summary of Firewall Policy Optimization Techniques and Methods

As stated above, the methods and systems described herein for a network firewall policy optimization can be implemented using any combination of hardware, software, and/or firmware. Alternatively, the steps for firewall policy optimization can be performed manually with the number of rules being optimized is small enough. Figure 13 is a flow chart illustrating an exemplary process for firewall policy optimization according to an embodiment of the subject matter described herein. Referring to Figure 13, in block **1300**, a firewall policy including an ordered list of firewall rules is defined. In block **1302**, for each rule, a probability indicating a likelihood of receiving a packet that matches each rule is specified. In block **1304**, neighboring rules are sorted in order of non-increasing probability in a manner that preserves firewall policy. In block **1306**, groups of rules are sorted in order of non-increasing probability in a manner that preserves firewall policy. For example, any of the policy-trie-based methods described above may be used. In block **1308**, one or more rules are split to reduce the number of average packet comparisons. In block **1310**, intersecting rules having common actions are identified and collapsed into single rules.

As described above, the subject matter described herein may be implemented in hardware, software, and/or firmware. In one implementation, the subject matter described herein for firewall policy optimization may be implemented on a general purpose computing platform. Figure 14 is a block diagram of a general purpose computing platform including software for firewall policy optimization according to an embodiment of the subject matter described herein. Referring to Figure 14, computing platform **1400** may be a general purpose computer, such as a personal computer, that includes a microprocessor **1402**, memory **1404**, and I/O interfaces **1406**. Microprocessor **1400** may be any suitable microprocessor, such as any of the Intel or AMD families of microprocessors. Memory **1404** may include volatile memory for running programs and persistent memory, such as one or more disk storage devices. I/O interface **1406** may include interfaces with I/O devices, such as user input devices and output devices.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1087

In the illustrated example, software that may be resident in memory **1404** includes a firewall policy editor **1408** and a firewall policy optimizer **1410**. Firewall policy editor **1408** may allow a user to define a firewall policy. For example, firewall policy editor **1408** may allow a user to define a firewall policy by specifying different rules. The rules may be specified in a graphical manner, for example using policy DAGs as described above. Alternatively, firewall policy editor **1408** may allow a user to input rules in a tabular manner, as illustrated in any of the tables described herein. Firewall policy optimizer **1410** may implement any or all of the firewall policy optimization techniques to order rules entered via firewall policy editor **1408** in a manner that preserves policy integrity and that enhances firewall performance.

A firewall rule set that is optimized using the subject matter described herein may be implemented on any firewall system that includes one or more firewalls. For example, an optimized firewall rule set according to embodiments of the subject matter described herein may be implemented using any of the hierarchical, multi-node firewall systems described in commonly assigned, co-pending U.S. patent application no. 11/316,331, filed December 22, 2005, the disclosure of which is incorporated herein by reference in its entirety. Figure 15 illustrates an example of a multi-node firewall system suitable for implementing the firewall rules according to an embodiment of the subject matter described herein. Referring to Figure 15, a plurality of firewall nodes **1500**, **1502**, **1504**,and **1506** may collectively implement any of the optimized firewall policies described herein. For example, the firewall nodes may collectively implement different portions of a firewall policy data structure including an ordered list of firewall rules. A firewall policy engine **1508** resident on each firewall node may filter packets using its local portion **1510** of the firewall policy data structure. A control node **1512** may include a firewall policy optimizer **1514** that measures the average number of comparisons per packets for the current rule configuration and may dynamically reorder the rules to reduce the average number of comparisons per packet. For example, firewall policy optimizer **1514** may utilize any of the methods described herein to rearrange rules and improve firewall performance. Rules may be tested for rearrangement at user-specified

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1088

intervals or when average number of packet comparisons increases by a user-specified amount.

References

The disclosure of each of the following references is hereby incorporated herein by reference in its entirety.

[1] E. Al-Shaer and H. Hamed. Firewall Policy Management Advisor for Anomaly Detection and Rule Editing. In *Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management*, 2003.

[2] E. Al-Shaer and H. Hamed. Modeling and Management of Firewall Policies. *IEEE Transactions on Network and Service Management*, 1(1), 2004.

[3] S. M. Bellovin and W. Cheswick. Network Firewalls. *IEEE Communications Magazine*, pages 50–57, Sept. 1994.

[4] C. Benecke. A Parallel Packet Screen for High Speed Networks. In *Proceedings of the 15th Annual Computer Security Applications Conference*, 1999.

[5] G. Brightwell and P. Winkler. Counting Linear Extensions is #P-Complete. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*, 1991.

[6] M. Christiansen and E. Fleury. Using Interval Decision Diagrams for Packet Filtering. Technical report, BRICS, 2002.

[7] D. Decasper, Z. Dittia, G. Parulkar, and B. Plattner. Router Plugins: A Software Architecture for Next-Generation Routers. *IEEE/ACM Transactions on Networking*, 8(1), February 2000.

[8] M. Degermark, A. Brodnik, S. Carlsson, and S. Pink. Small Forwarding Tables for Fast Routing Lookups. In *Proceedings of ACM SIGCOMM*, pages 4 – 13, 1997.

[9] U. Ellermann and C. Benecke. Firewalls for ATM Networks. In Proceedings of *INFOSEC'COM*, 1998.

[10] A. Feldmann and S. Muthukrishnan. Tradeoffs for Packet Classification. In Proceedings of the *IEEE INFOCOM*, pages 397 – 413, 2000.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1089

[11] E. W. Fulp. Firewall Architectures for High Speed Networks. Technical Report 20026, Wake Forest University Computer Science Department, 2002.

[12] E. W. Fulp. Firewall Policy Models Using Ordered-Sets. Technical report, Wake Forest University Computer Science Department, 2004.

[13] R. Funke, A. Grote, and H.-U. Heiss. Performance Evaluation of Firewalls in Gigabit-Networks. In *Proceedings of the Symposium on Performance Evaluation of Computer and Telecommunication Systems*, 1999.

[14] S. Goddard, R. Kieckhafer, and Y. Zhang. An Unavailability Analysis of Firewall Sandwich Configurations. In *Proceedings of the 6th IEEE Symposium on High Assurance Systems Engineering*, 2001.

[15] R. L. Graham, E. L. Lawler, J. K. Lenstra, and A. H. G. R. Kan. Optimizing and Approximation in Determinstic Sequencing and Scheduling: A Survey. *Annals of Discrete Mathematics*, 5:287 – 326, 1979.

[16] A. Hari, S. Suri, and G. Parulkar. Detecting and Resolving Packet Filter Conflicts. In *Proceedings of IEEE INFOCOM*, pages 1203–1212, 2000.

[17] HiPAC. High Performance Packet Classification. http://www.hipac.org.

[18] E. Horowitz, S. Sahni, and D. Mehta. *Fundamentals of Data Structures in C++*. Compuer Science Press, 1995.

[19] E. L. Lawler. Sequencing Jobs to Minimize Total Weighted Completion Time Subject to Precedence Constraints. *Annals of Discrete Mathematics*, 2:75 – 90, 1978.

[20] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the Self-Similar Nature of Ethernet Traffic. *IEEE Transactions on Networking*, 2:1 – 15, 1994.

[21] J. K. Lenstra and A. H. G. R. Kan. Complexity of Scheduling under Precedence Constraints. *Operations Research*, 26(1):22 – 35, 1978.

[22] O. Paul and M. Laurent. A Full Bandwidth ATM Firewall. In Proceedings of the 6th European Symposium on *Research in Computer Security ESORICS'2000*, 2000.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1090

[23] B. R. Preiss. *Data Structures and Algorithms with Object-Oriented Design Patterns in C++*. John Wiley & Sons, 1999.

[24] L. Qui, G. Varghese, and S. Suri. Fast Firewall Implementations for Software and Hardware-Based Routers. In *Proceedings of ACM SIGMETRICS*, June 2001.

[25] V. P. Ranganath and D. Andresen. A Set-Based Approach to Packet Classification. In *Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems*, pages 889–894, 2003.

[26] R. Rivest. On Self-Organizing Sequential Search Heuristics. *Communications of the ACM*, 19(2), 1976.

[27] W. E. Smith. Various Optimizers for Single-Stage Production. *Naval Research Logistics Quarterly*, 3:59 – 66, 1956.

[28] V. Srinivasan, G. Varghese, S. Suri, and M. Waldvogel. Fast and Scalable Layer Four Switching. In *Proceedings of ACM SIGCOMM*, pages 191–202, 1998.

[29] S. Suri and G. Varghese. Packet Filtering in High Speed Networks. In *Proceedings of the Symposium on Discrete Algorithms*, pages 969 – 970, 1999.

[30] R. L. Ziegler. *Linux Firewalls*. New Riders, second edition, 2002.

[31] E. D. Zwicky, S. Cooper, and D. B. Chapman. *Building Internet Firewalls*. O'Reilly, 2000.

It will be understood that various details of the invention may be changed without departing from the scope of the invention. Furthermore, the foregoing description is for the purpose of illustration only, and not for the purpose of limitation.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1091

CLAIMS

What is claimed is:

1. A computer-implemented method for producing a performance optimized firewall policy, the method comprising:

(a) defining a firewall policy including an ordered list of firewall rules;

(b) for each rule, defining a probability indicating the likelihood of receiving a packet matching each rule;

(c) sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy; and

(d) sorting groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy.

2. The method of claim 1, wherein sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy includes, for each pair of neighboring rules, rearranging the rules if a first rule of the pair has a lower probability than a second rule of the pair and the first rule of the pair is not a subset of the second rule of the pair with a different action.

3. The method of claim 1, wherein sorting groups of two or more rules in a manner that preserves firewall policy includes assigning each group of rules to a sub-trie in a policy trie, computing a sum of packet matching probabilities for rules in each sub-trie, and, in response to determining that a first sub-trie has a lower sum than a second sub-trie, rotating the first and second sub-tries about a common parent so that rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy.

4. The method of claim 3, comprising, in response to determining that the first and second sub-tries have equal sums, determining whether the second sub-trie has a lower number of branches than the first sub-trie, and, in response to determining that the second sub-trie has a lower number of branches than the first sub-trie, rotating

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1092

the first and second sub-tries about a common parent so that the rules in the second sub-trie will be applied before the rules in the first sub-trie in the firewall policy.

5. The method of claim 1, wherein sorting groups of two or more rules in a manner that preserves the firewall policy includes:

(a) identifying first and second sub-lists of rules in the ordered list of firewall rules, the first sub-list preceding the second sub-list in the ordered list of firewall rules;

(b) determining a sum of packet matching probabilities for each of the first and second sub-lists;

(c) determining whether the first and second sub-lists intersect; and

(d) in response to determining that the sum of the packet matching probabilities for the second sub-list is greater than the sum of the packet matching probabilities for the first sub-list and that the first and second sub-lists do not intersect, switching the order of the first and second sub-lists in the ordered list of firewall rules.

6. The method of claim 5, comprising, in response to determining that the sums of the packet matching probabilities for the first and sub-lists are equal, determining whether the second sub-list has a lower number of rules than the first sub-list, and, in response to determining that the second sub-list has a lower number of rules than the first sub-list, switching the order of the first and second sub-lists in the ordered list of firewall rules.

7. The method of claim 1, comprising splitting at least one of the firewall rules to reduce an average number of comparisons required per received packet.

8. The method of claim 1, comprising identifying intersecting rules having a common action and collapsing the rules into a single rule representing a union of the intersecting rules.

9. The method of claim 1, wherein sorting neighboring rules includes sorting rules such that comparisons for a particular class of packets are reduced.

10. The method of claim 1, comprising assigning the rules to vertices in a directed acyclical graph (DAG). determining precedence relationships between the rules,

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1093

assigning the precedence relationships to edges that connect the vertices in the DAG, and wherein sorting the neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy includes sorting the neighboring rules in a manner that preserves the precedence relationships represented by the DAG.

11. The method of claim 1, wherein the rules comprise intrusion detection rules.

12. The method of claim 1, wherein the rules comprise intrusion protection rules.

13. A system for firewall policy optimization, for implementing the steps of any one of claims 1-12.

14. A network firewall comprising:

(a) a firewall policy data structure for storing an ordered list of firewall rules;

(b) a firewall policy engine for implementing a firewall policy by comparing packets to the rules in the order specified by the firewall policy data structure; and

(c) a firewall policy optimizer for optimizing performance of the network firewall by reordering the rules in a manner that reduces an average number of rule comparisons per packet and that preserves the firewall policy, the firewall policy optimizer sorting groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy.

15. The network firewall of claim 14, wherein the firewall policy optimizer is adapted to measure average number of comparisons per packet for a rule ordering specified by the firewall policy data structure and to dynamically re-order to rules to reduce the average number of comparisons per packet

16. A computer program product comprising computer executable instructions embodied in a computer readable medium for performing steps comprising:

(a) defining a firewall policy including an ordered list of firewall rules;

(b) for each rule, defining a probability indicating the likelihood of receiving a packet matching each rule;

(c) sorting neighboring rules in order of non-increasing probability in a manner that preserves the firewall policy; and

(d) sorting groups of two or more rules in order of non-increasing probability in a manner that preserves the firewall policy.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1095

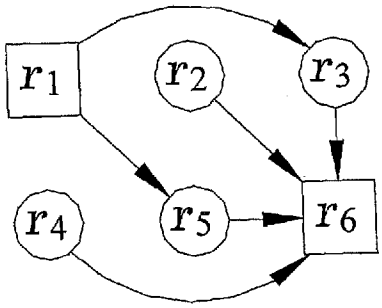

FIG. 1A

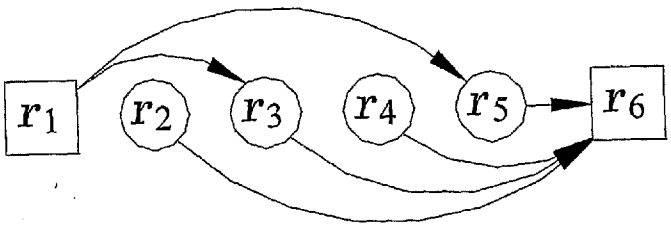

FIG. 1B

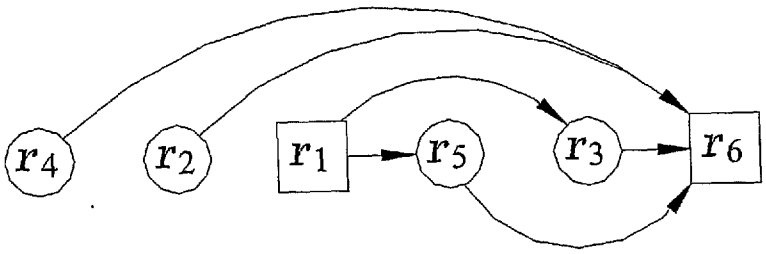

FIG. 2A

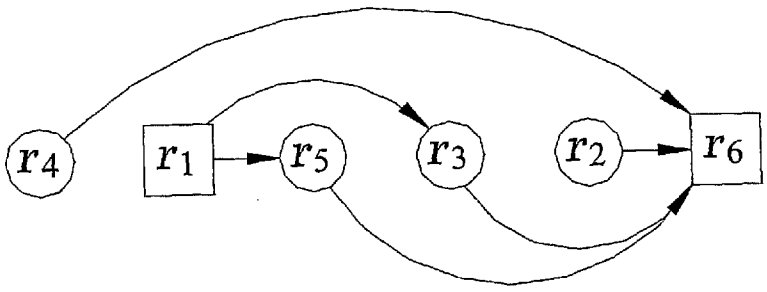

FIG. 2B

3/19

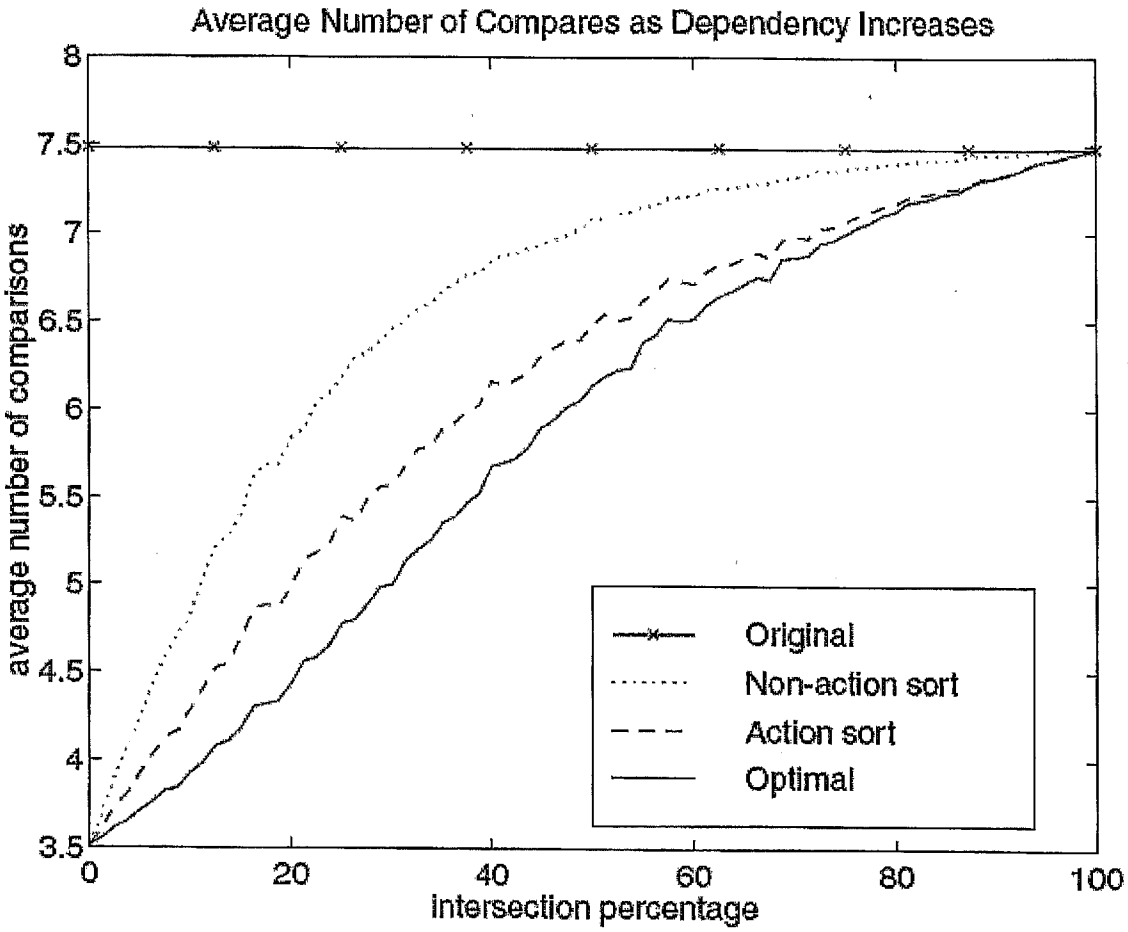
Average Number of Compares as Dependency Increases
average number of comparisons
intersection percentage
8
7.5
7
6.5
6
5.5
5
4.5
4
3.5
0
20
40
60
80
100
Original
Non-action sort
Action sort
Optimal

FIG. 3A

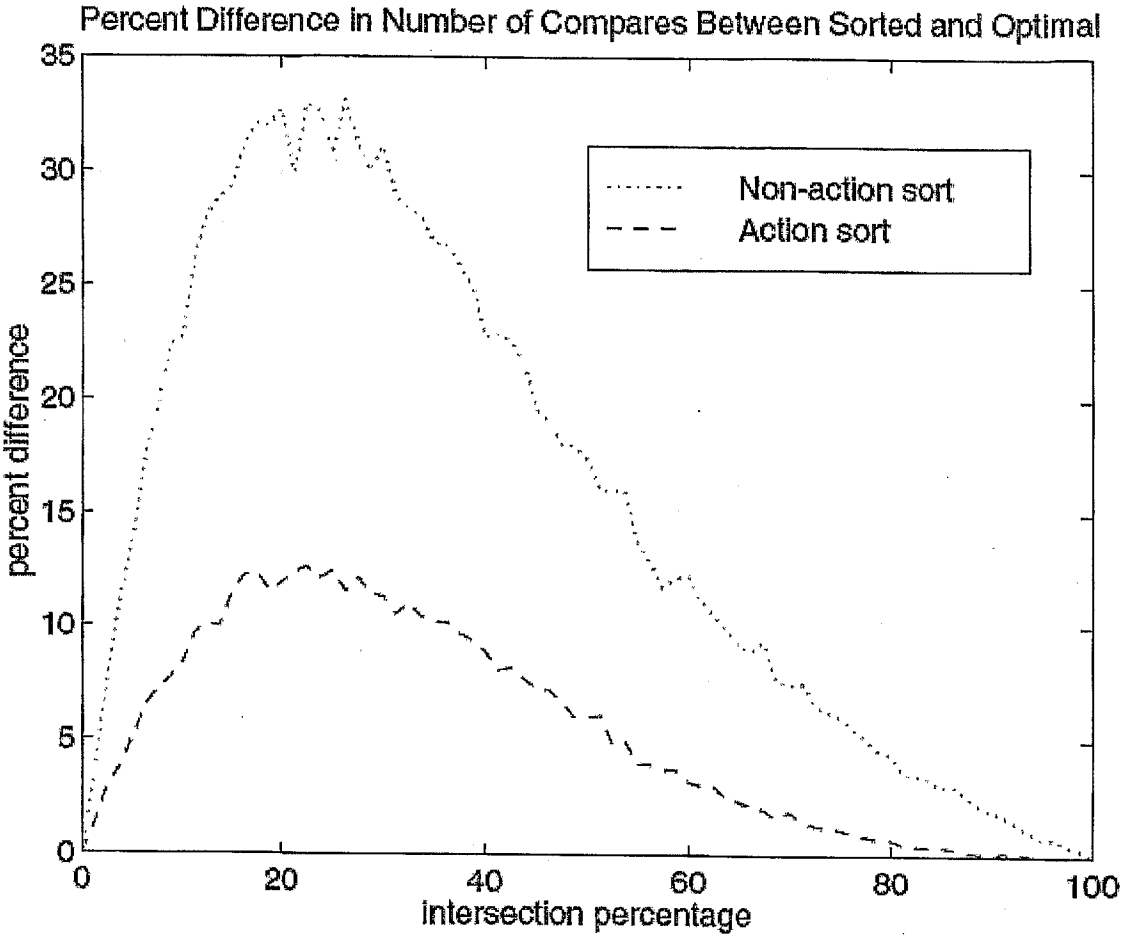
Percent Difference in Number of Compares Between Sorted and Optimal
percent difference
intersection percentage
Non-action sort
Action sort
0
5
10
15
20
25
30
35
0
20
40
60
80
100

FIG. 3B

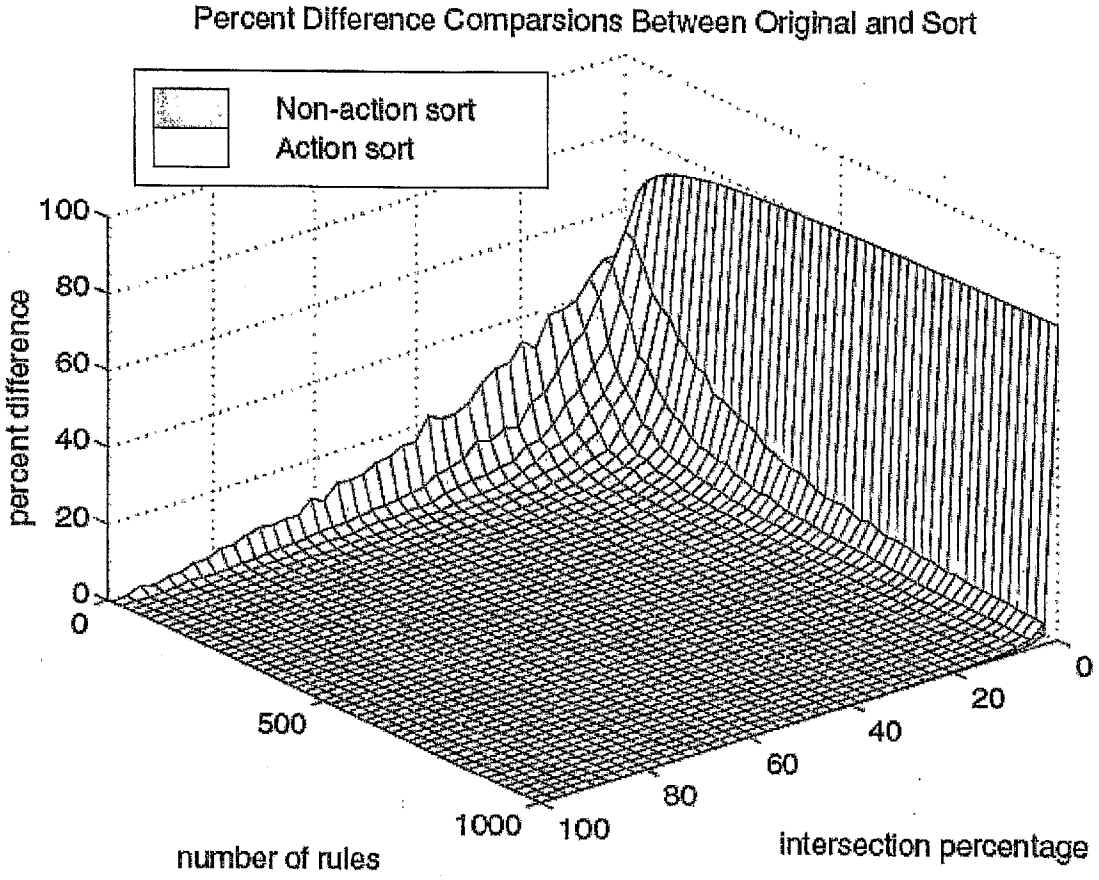
Percent Difference Comparsions Between Original and Sort
Non-action sort
Action sort
percent difference
100
80
60
40
20
0
0
500
1000
number of rules
0
20
40
60
80
100
intersection percentage

FIG. 4A

6/19

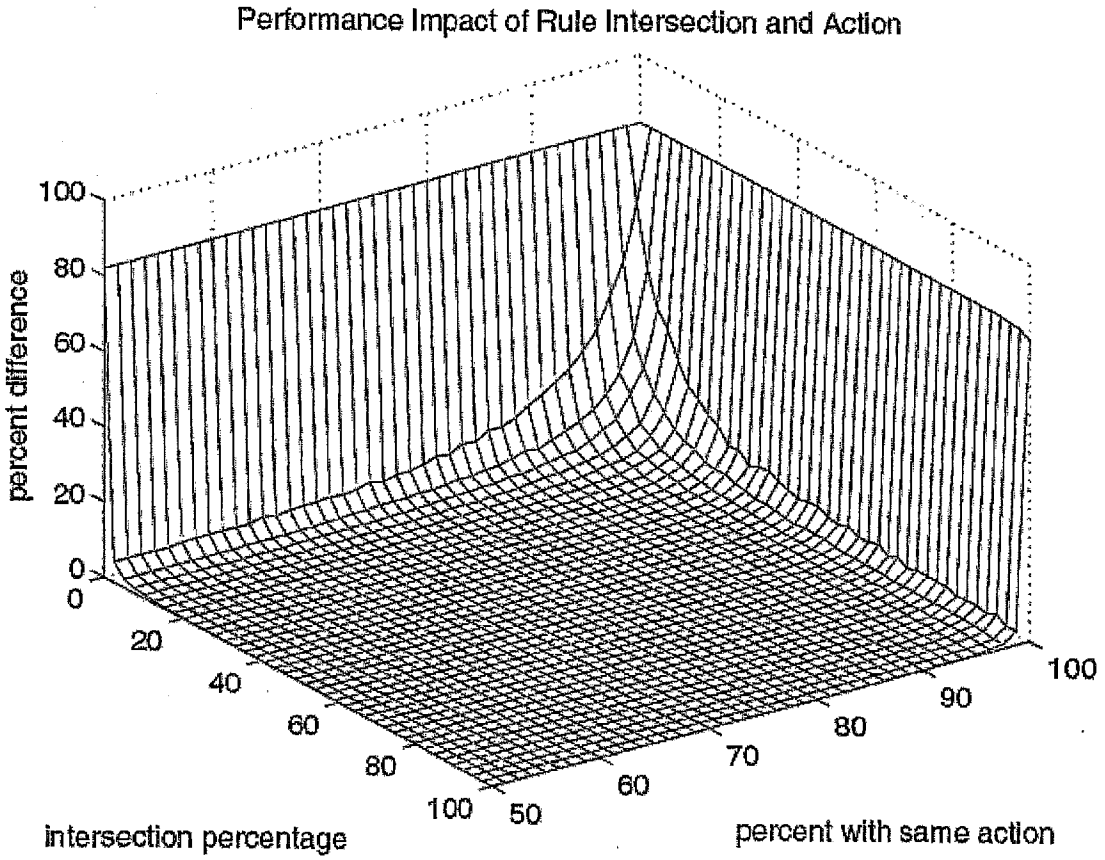
Performance Impact of Rule Intersection and Action
percent difference
100
80
60
40
20
0
0
20
40
60
80
100
intersection percentage
50
60
70
80
90
100
percent with same action

FIG. 4B

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 5A

r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 5B

8/19

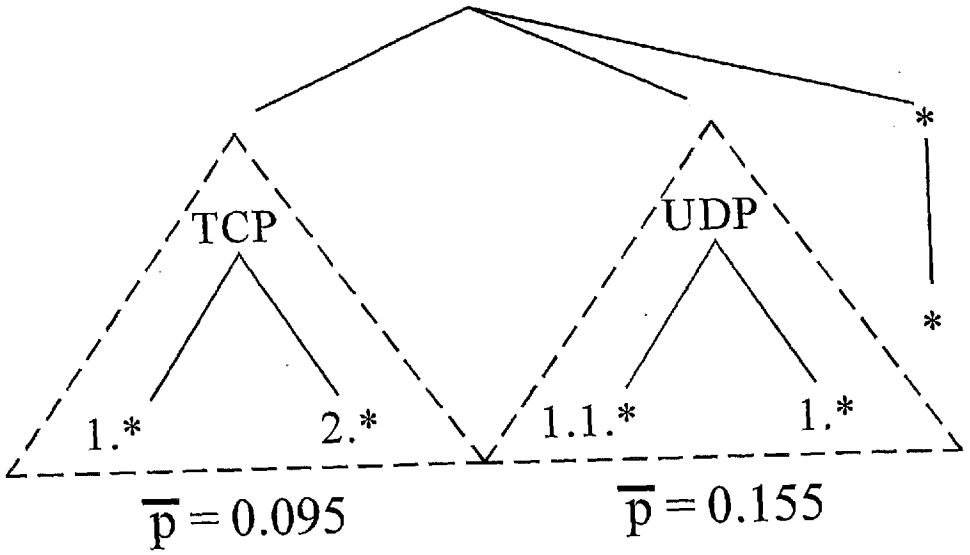

FIG. 6A

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	DENY)	$p_4 = 0.30$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_5	(*,	*,	DENY)	$p_5 = 0.50$

FIG. 6B

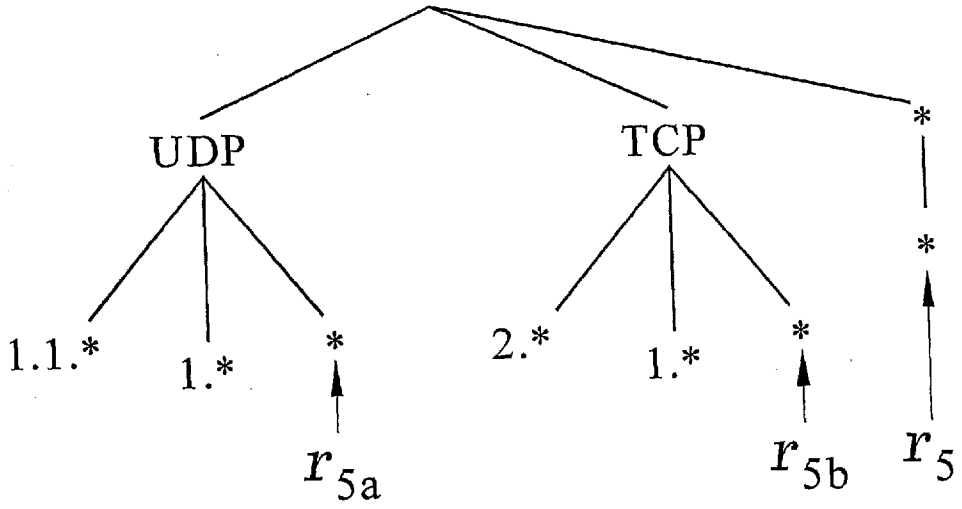

FIG. 7A

r_1	(UDP,	1.1.*,	ACCEPT)	$p_1 = 0.01$
r_4	(UDP,	1.*,	ACCEPT)	$p_4 = 0.30$
r_{5a}	(UDP,	*,	DENY)	$p_{5a} = 0.47$
r_3	(TCP,	2.*,	ACCEPT)	$p_3 = 0.17$
r_2	(TCP,	1.*,	ACCEPT)	$p_2 = 0.02$
r_{5b}	(TCP,	*,	DENY)	$p_{5b} = 0.03$

FIG. 7B

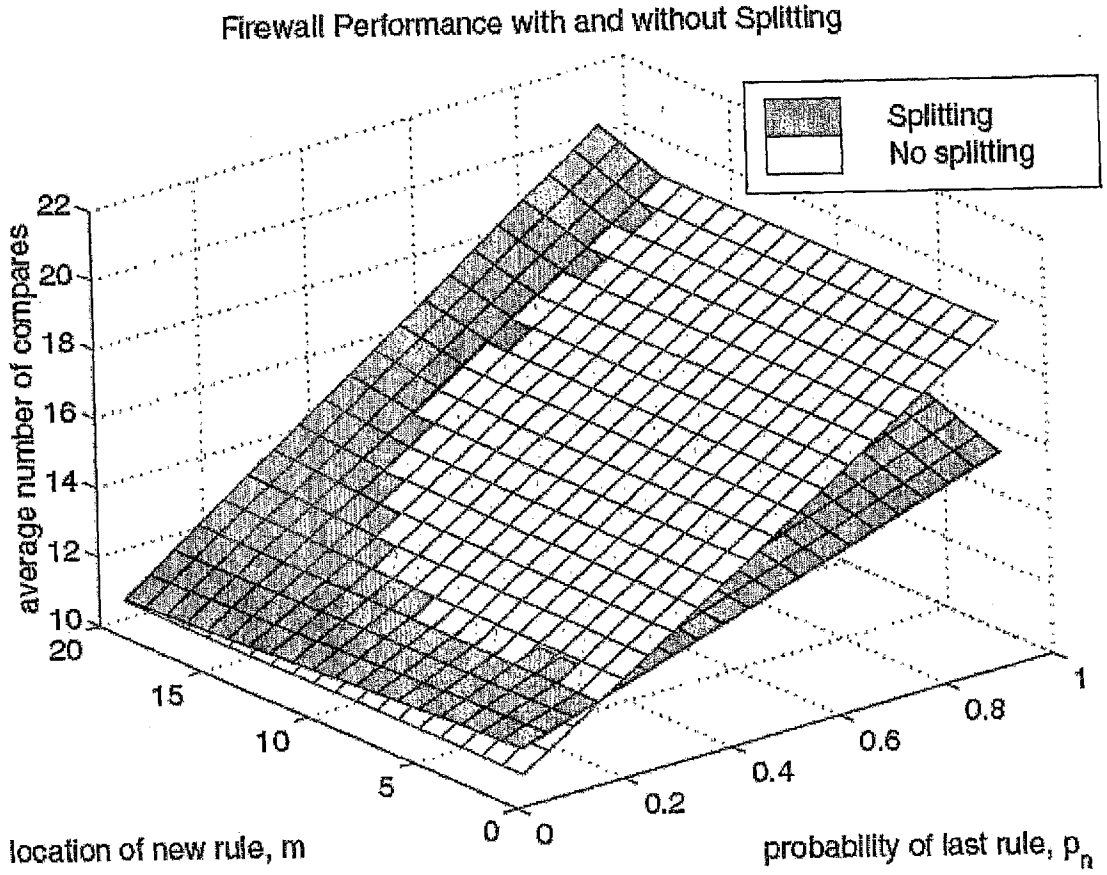
Firewall Performance with and without Splitting
Splitting
No splitting
average number of compares
22
20
18
16
14
12
10
20
15
10
5
0
location of new rule, m
0
0.2
0.4
0.6
0.8
1
probability of last rule, p_n

FIG. 8A

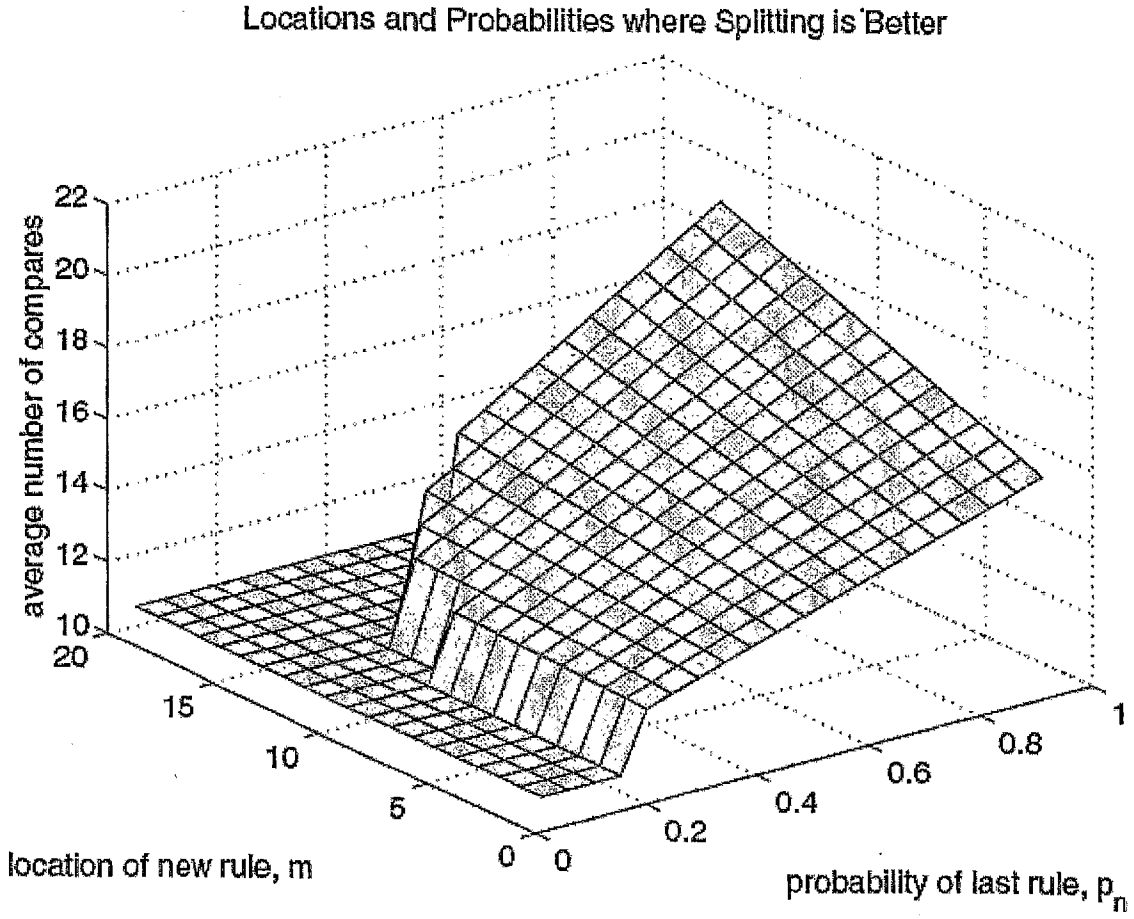
Locations and Probabilities where Splitting is Better
average number of compares
22
20
18
16
14
12
10
20
15
10
5
0
location of new rule, m
0
0.2
0.4
0.6
0.8
1
probability of last rule, p_n

FIG. 8B

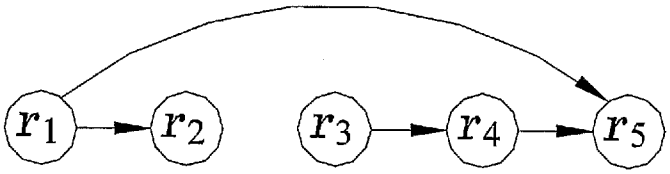

FIG. 9A

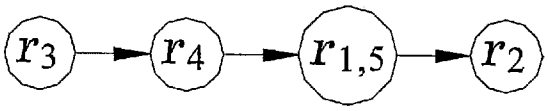

FIG. 9B

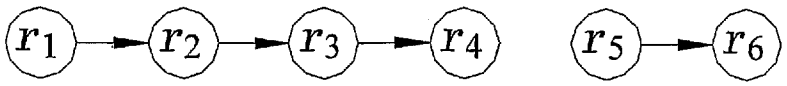

FIG. 10A

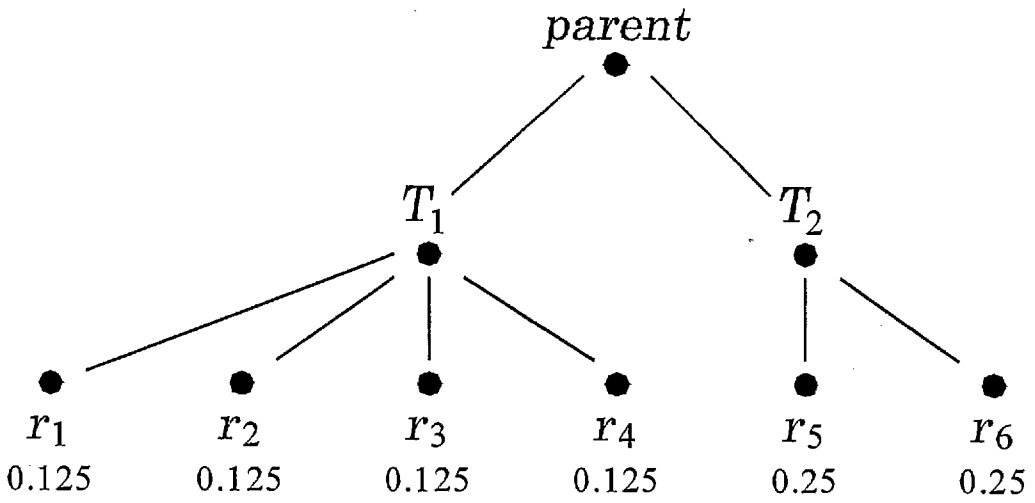

FIG. 10B

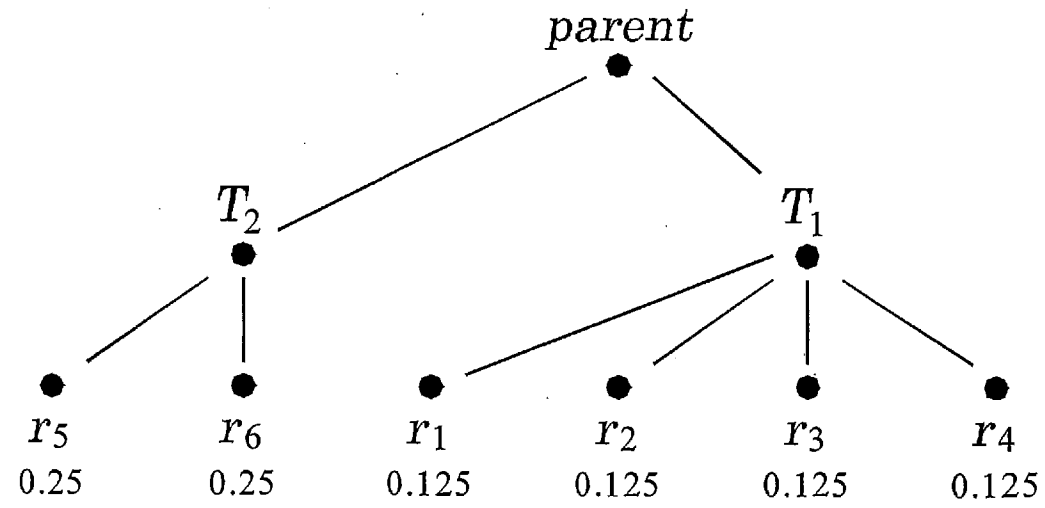

FIG. 10C

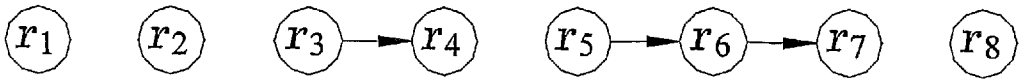

FIG. 11A

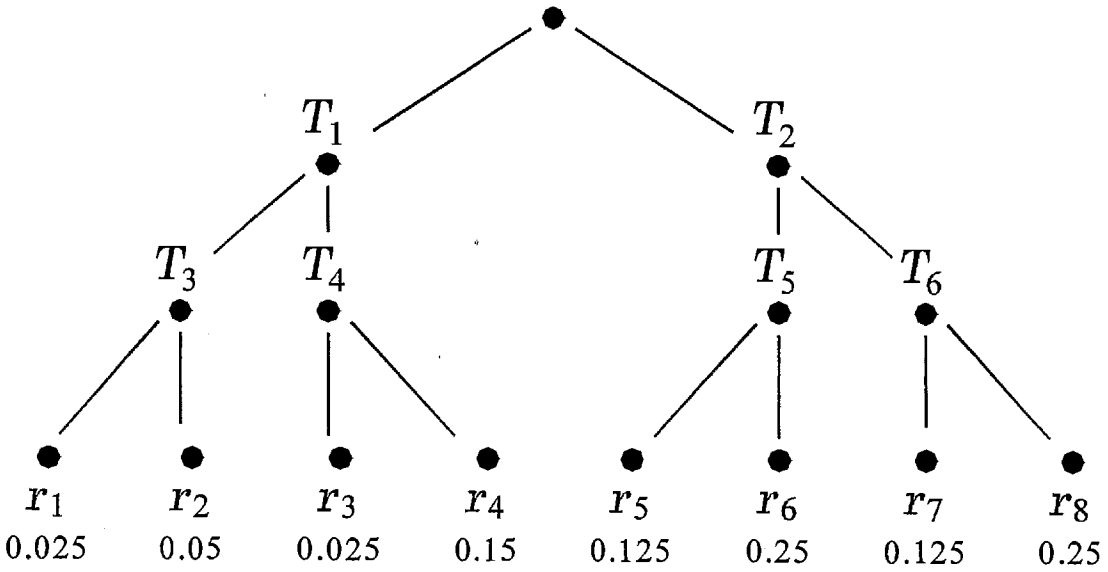

FIG. 11B

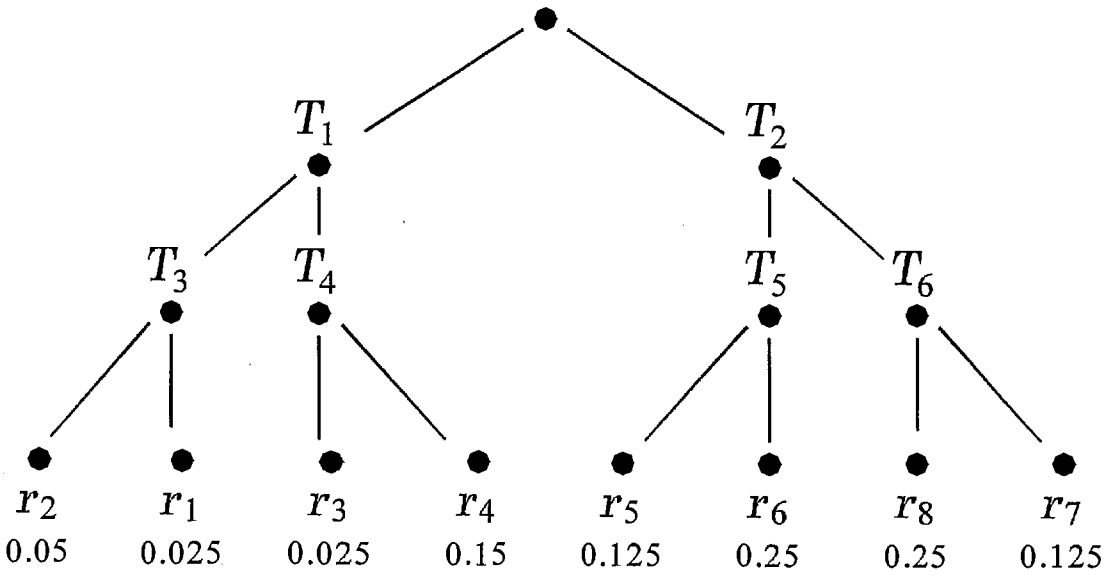

FIG. 11C

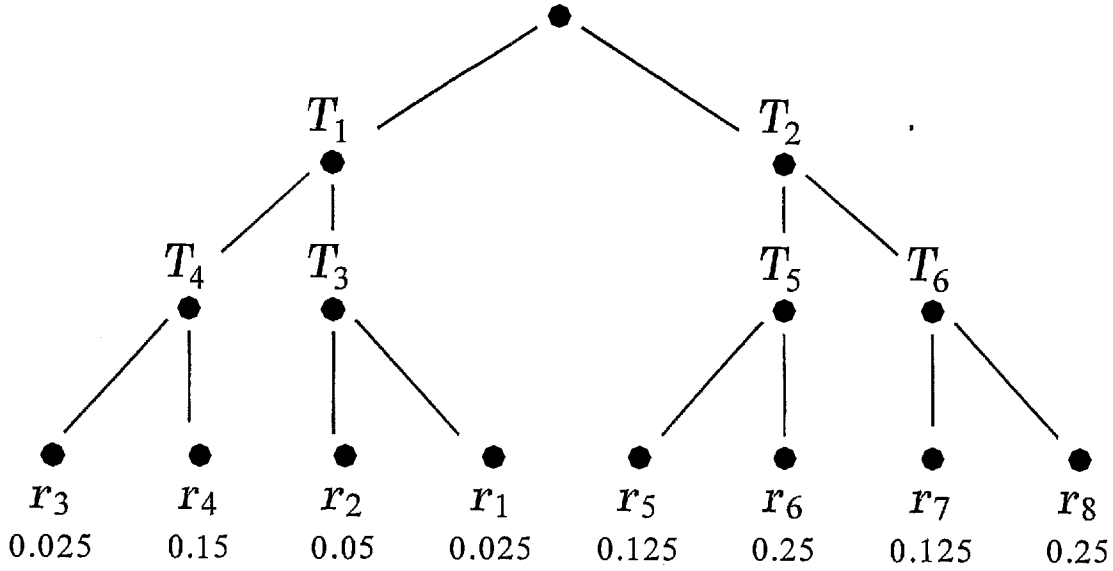

FIG. 11D

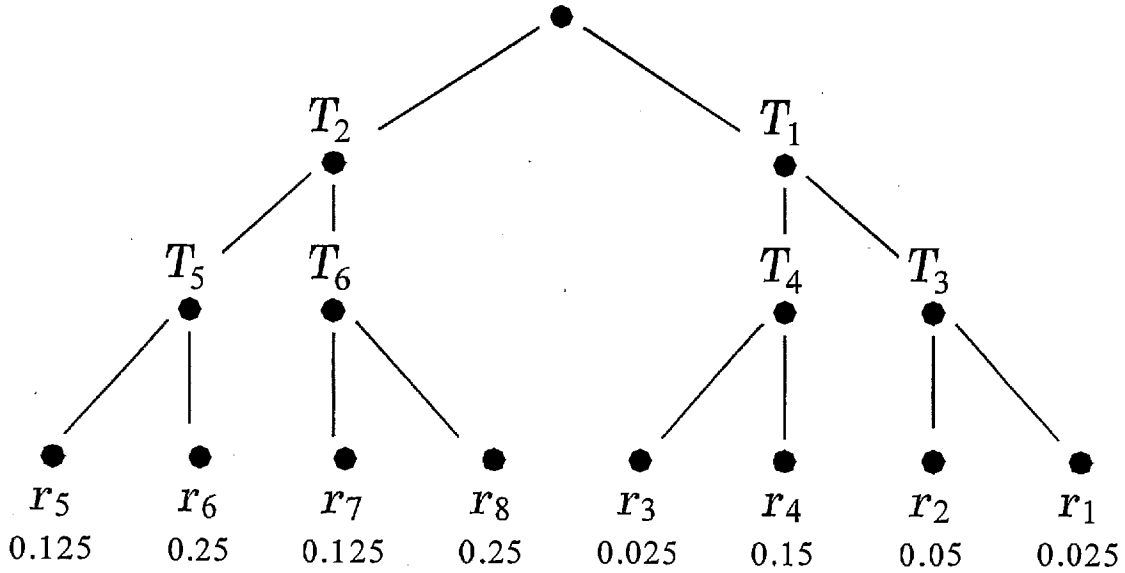

FIG. 11E

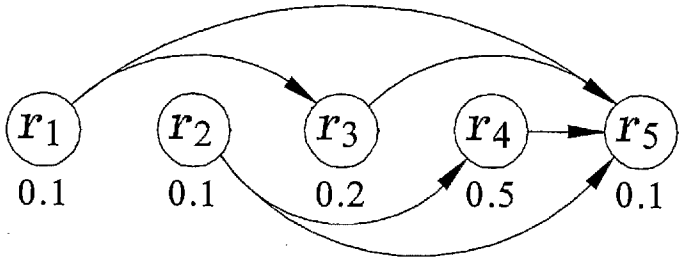

FIG. 12A

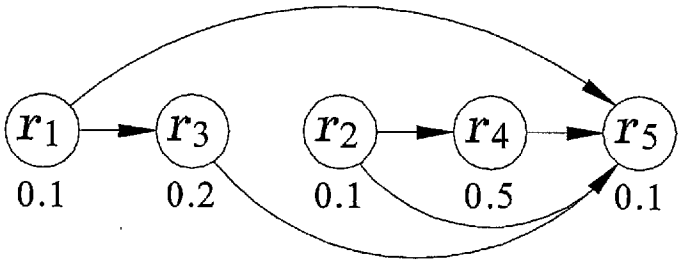

FIG. 12B

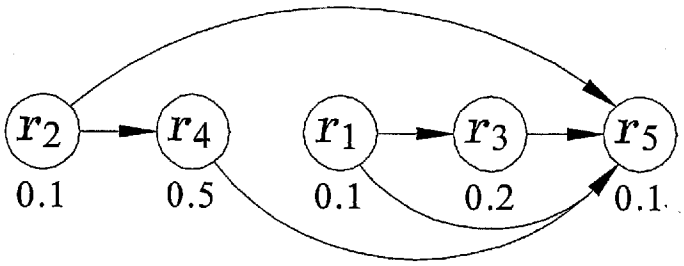

FIG. 12C

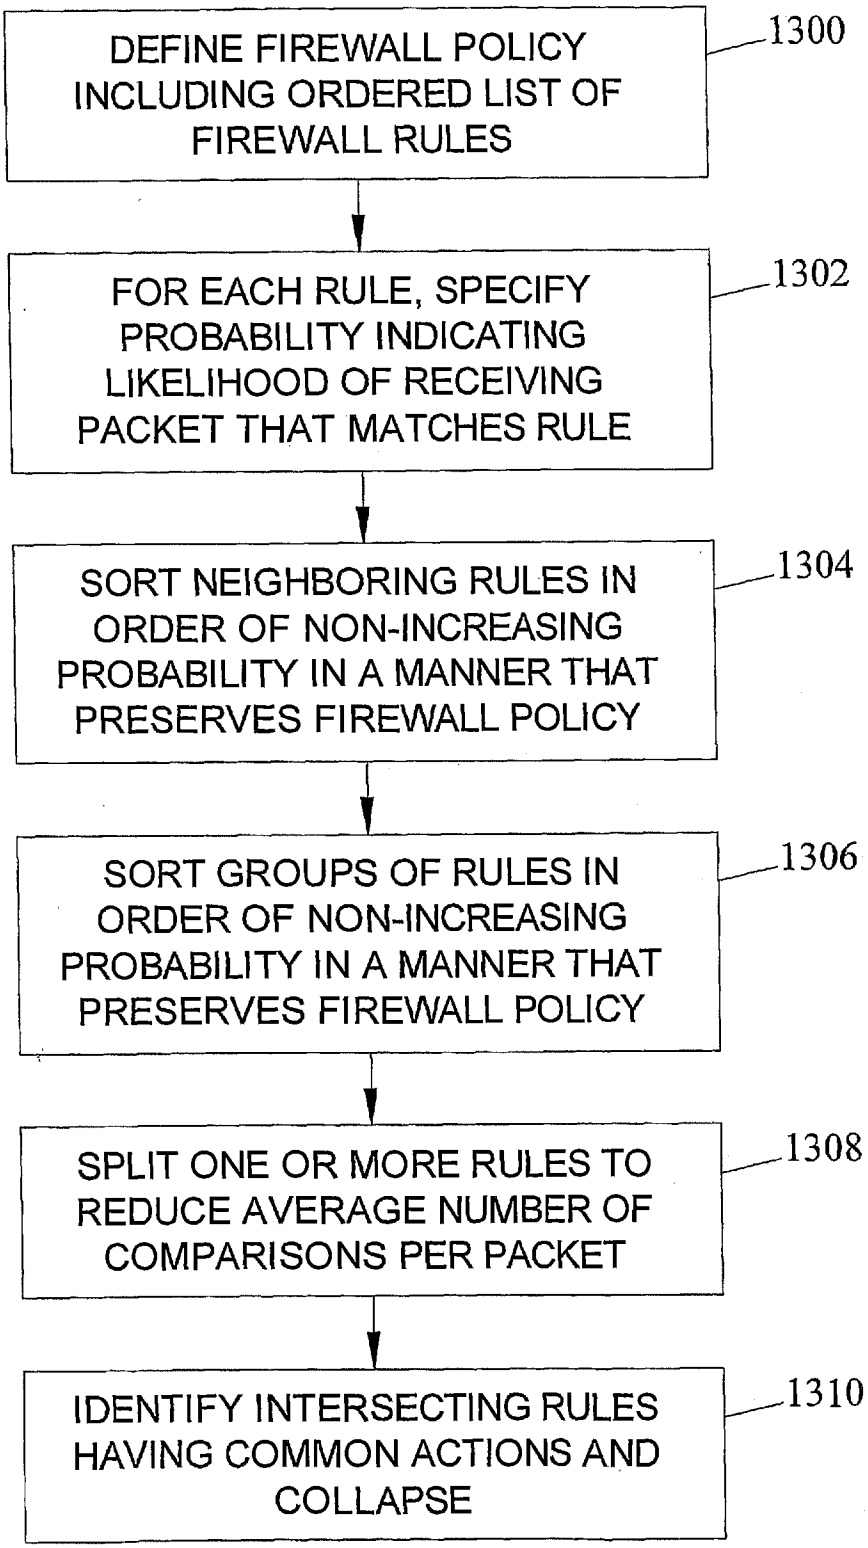
DEFINE FIREWALL POLICY INCLUDING ORDERED LIST OF FIREWALL RULES
1300
FOR EACH RULE, SPECIFY PROBABILITY INDICATING LIKELIHOOD OF RECEIVING PACKET THAT MATCHES RULE
1302
SORT NEIGHBORING RULES IN ORDER OF NON-INCREASING PROBABILITY IN A MANNER THAT PRESERVES FIREWALL POLICY
1304
SORT GROUPS OF RULES IN ORDER OF NON-INCREASING PROBABILITY IN A MANNER THAT PRESERVES FIREWALL POLICY
1306
SPLIT ONE OR MORE RULES TO REDUCE AVERAGE NUMBER OF COMPARISONS PER PACKET
1308
IDENTIFY INTERSECTING RULES HAVING COMMON ACTIONS AND COLLAPSE
1310

FIG. 13

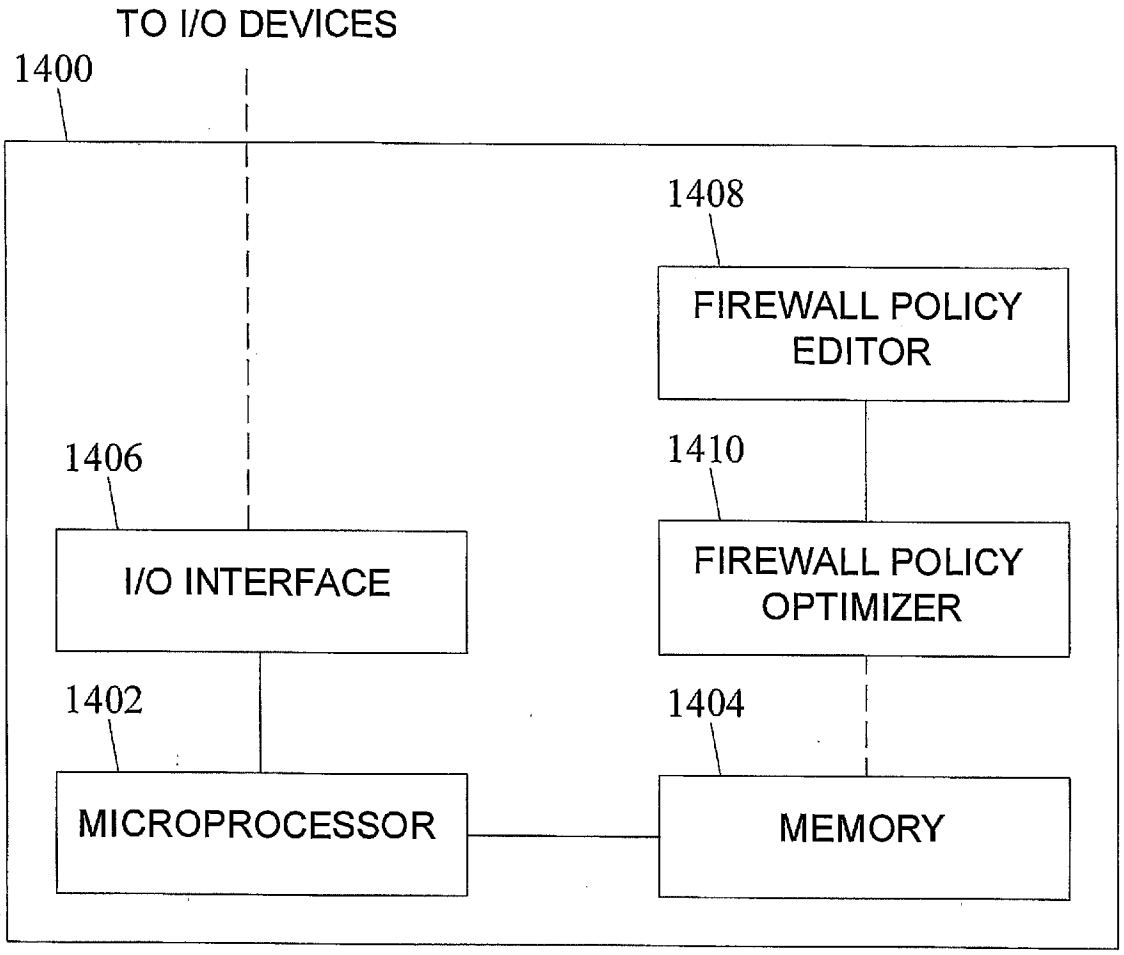
TO I/O DEVICES
1400
1408
FIREWALL POLICY EDITOR
1406
1410
I/O INTERFACE
FIREWALL POLICY OPTIMIZER
1402
1404
MICROPROCESSOR
MEMORY

FIG. 14

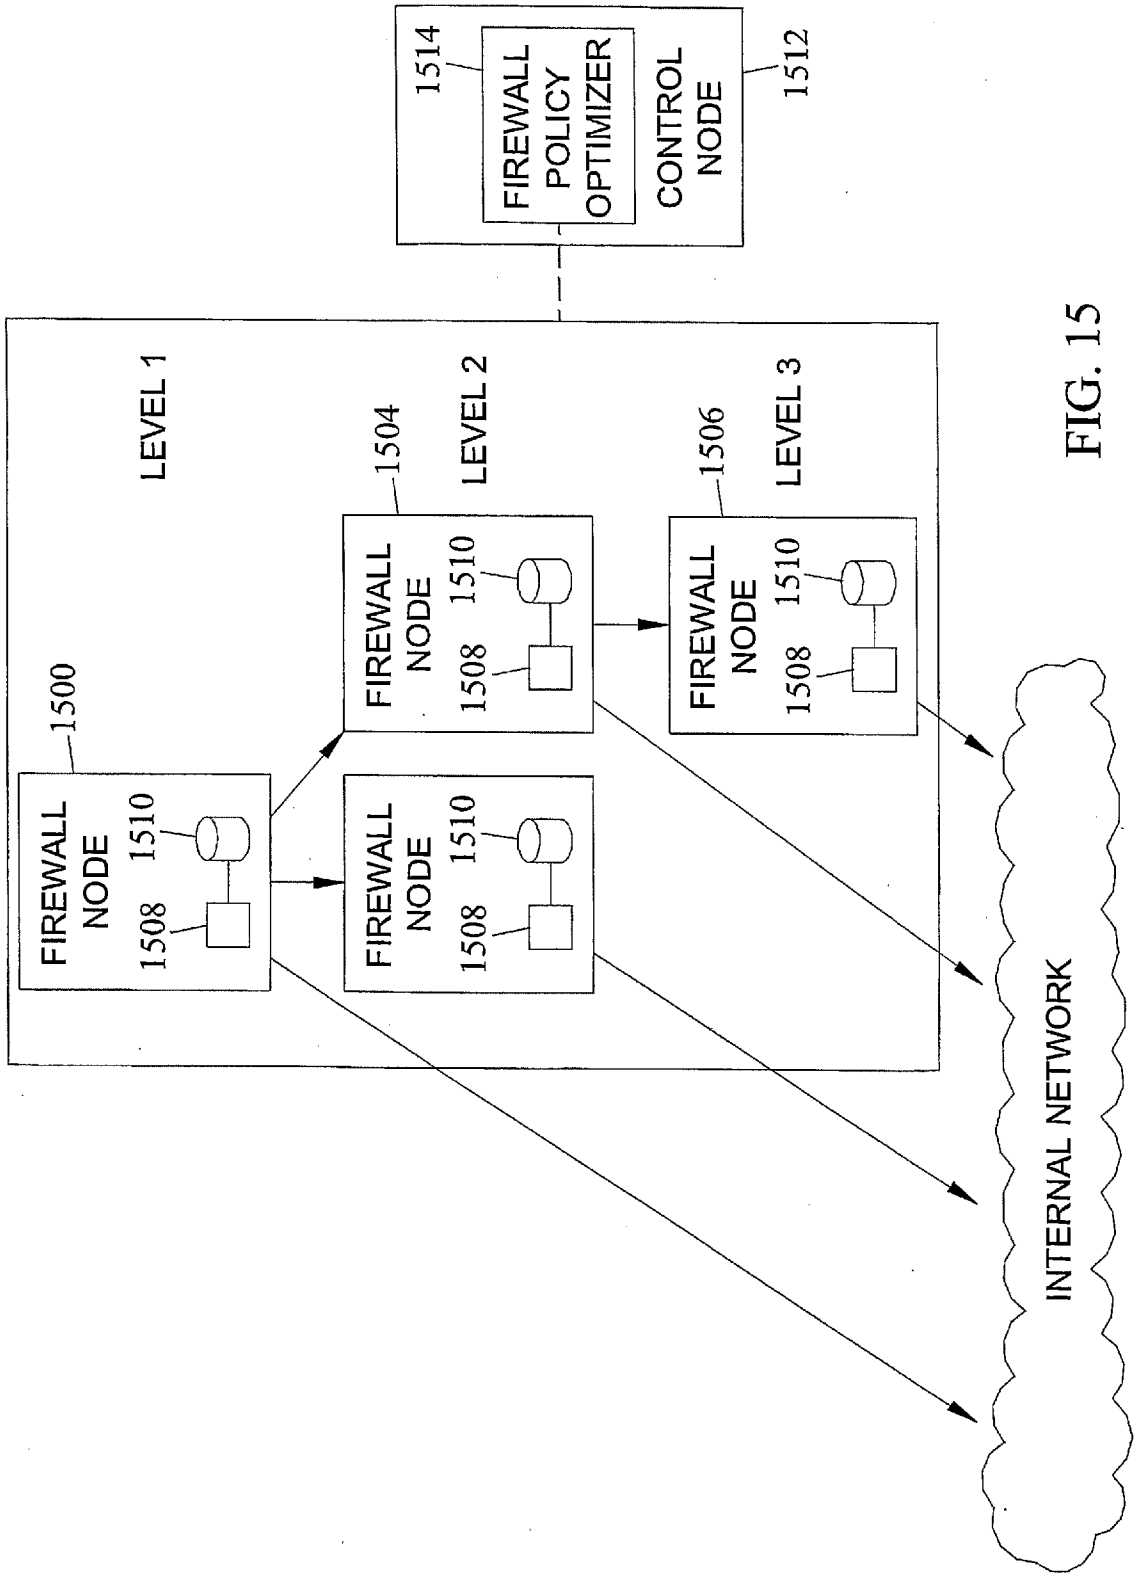
1514
FIREWALL POLICY OPTIMIZER
CONTROL NODE
1512
LEVEL 1
LEVEL 2
LEVEL 3
1500
1504
1506
FIREWALL NODE
1510
1508
INTERNAL NETWORK

FIG. 15

Espacenet

Bibliographic data: WO2012146265 (A1) — 2012-11-01

CORRELATION OF MEDIA PLANE AND SIGNALING PLANE OF MEDIA SERVICES IN A PACKET-SWITCHED NETWORK

Inventor(s): SCHOLZ HENDRIK [DE]; WALLBAUM MICHAEL [DE]; KRUEGER MICHAEL [DE]; BASTIAN JAN [DE] ± (SCHOLZ, HENDRIK, ; WALLBAUM, MICHAEL, ; KRUEGER, MICHAEL, ; BASTIAN, JAN)

Applicant(s): VOIPFUTURE GMBH [DE]; SCHOLZ HENDRIK [DE]; WALLBAUM MICHAEL [DE]; KRUEGER MICHAEL [DE]; BASTIAN JAN [DE] ± (VOIPFUTURE GMBH, ; SCHOLZ, HENDRIK, ; WALLBAUM, MICHAEL, ; KRUEGER, MICHAEL, ; BASTIAN, JAN)

Classification: **- international:** ***H04L12/24; H04L29/06;*** *H04L12/26*
- cooperative: **H04L41/5032**; **H04L43/062**; **H04L43/12**; **H04L65/1006**; **H04L65/1009**; **H04L65/1069**; **H04L65/607**; **H04L65/80**

Application number: WO2011EP02145 20110428

Priority number(s): WO2011EP02145 20110428

Also published as: US2014089504 (A1) EP2702740 (A1) CA2834168 (A1)

Abstract of WO2012146265 (A1)

This invention relates to methods for correlating media streams and signaling sessions of services, for example, in a passive monitoring system of a packet-switched network. Furthermore, the invention also relates to an implementation of these methods in hardware and software, and provides a signaling plane probe, a media plane probe and a correlation unit. Moreover, a passive monitoring system comprising one or more of these hardware devices is provided. To correlate media streams and

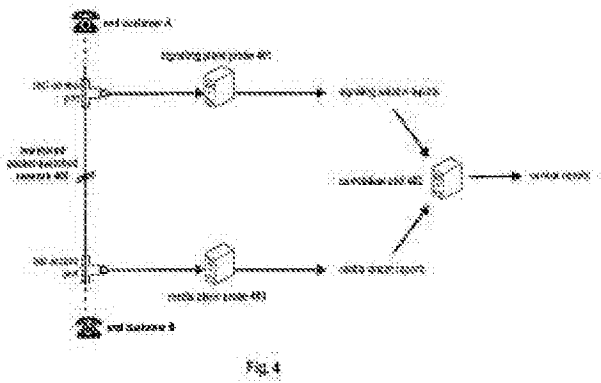

signaling sessions of services, the invention proposes to independently generate correlation keys in a media plane probe for monitored media streams and correlation keys for signaling sessions that are monitored by a signaling plane probe in a fashion that matching correlation keys are generated for a respective service. By identifying reports on the media streams and reports on the signaling session that contain matching correlation keys respectively, the media streams and the signaling streams are correlated in a correlation unit.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
1 November 2012 (01.11.2012)

WIPO | PCT

(10) International Publication Number
WO 2012/146265 A1

(51) International Patent Classification:
H04L 29/06 (2006.01) *H04L 12/26* (2006.01)
H04L 12/24 (2006.01)

(21) International Application Number:
PCT/EP2011/002145

(22) International Filing Date:
28 April 2011 (28.04.2011)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant *(for all designated States except US)*: **VOIPFUTURE GMBH** [DE/DE]; Wendenstrasse 4, 20097 Hamburg (DE).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **SCHOLZ, Hendrik** [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE). **WALLBAUM, Michael** [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE). **KRUEGER, Michael** [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE). **BASTIAN, Jan** [DE/DE]; C/o Voipfuture Gmbh, Wendenstrasse 4, 20097 Hamburg (DE).

(74) Agent: **BOEHMERT & BOEHMERT**; Felix Hermann, Hollerallee 32, 28209 Bremen (DE).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

(54) **Title**: CORRELATION OF MEDIA PLANE AND SIGNALING PLANE OF MEDIA SERVICES IN A PACKET-SWITCHED NETWORK

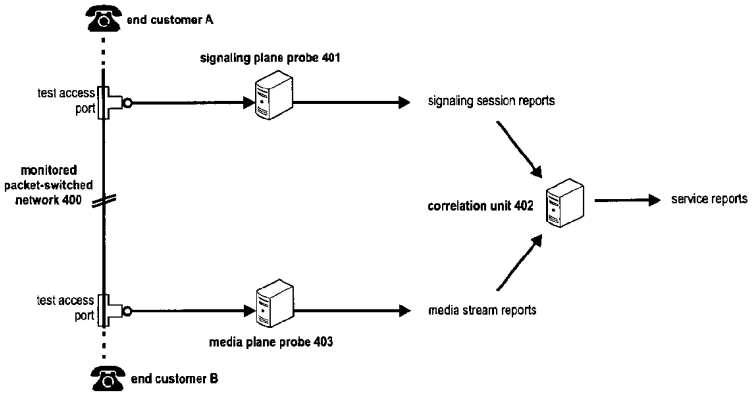

Fig. 4

(57) **Abstract**: This invention relates to methods for correlating media streams and signaling sessions of services, for example, in a passive monitoring system of a packet-switched network. Furthermore, the invention also relates to an implementation of these methods in hardware and software, and provides a signaling plane probe, a media plane probe and a correlation unit. Moreover, a passive monitoring system comprising one or more of these hardware devices is provided. To correlate media streams and signaling sessions of services, the invention proposes to independently generate correlation keys in a media plane probe for monitored media streams and correlation keys for signaling sessions that are monitored by a signaling plane probe in a fashion that matching correlation keys are generated for a respective service. By identifying reports on the media streams and reports on the signaling session that contain matching correlation keys respectively, the media streams and the signaling streams are correlated in a correlation unit.

WO 2012/146265 A1

Correlation of Media Plane and Signaling Plane of Media Services in a Packet-Switched Network

FIELD OF THE INVENTION

This invention relates to methods for correlating media streams and signaling sessions of services, for example, in a passive monitoring system of a packet-switched network. Furthermore, the invention also relates to an implementation of these methods in hardware and software, and provides a signaling plane probe, a media plane probe and a correlation unit. Moreover, a passive monitoring system comprising one or more of these hardware devices is provided.

TECHNICAL BACKGROUND

VoIP Services

VoIP ("Voice over IP" – IP denoting the Internet Protocol) services may be considered to consist of a signaling plane and a media plane. On the signaling plane various protocols describe the session (call) flow in terms of involved parties, intermediary VoIP entities (i.e. VoIP proxies, routers) and the characteristics of the VoIP service (call).The media plane typically carries the media information (i.e. audio and/or video data) between the involved parties. Neither the media plane nor the signaling plane alone is sufficient to carry a VoIP service. On the signaling plane protocols like SIP (see IETF RFC 3261, "SIP: Session Initiation Protocol", available at http://www.ietf.org) or ITU-T recommendation H.323 (see H.323, "Packet-based multimedia communications systems", Edition 7, 2009, available at http://www.itu.int) are commonly used, whereas protocols like RTP (Real-time Transport Protocol, see IETF RFC 3550, "RTP: A Transport Protocol for Real-Time Applications", available at http://www.ietf.org), MSRP (see IETF RFC 4975, "The Message Session Relay Protocol (MSRP)", available at http://www.ietf.org) or ITU recommendation T.38 (see T.38, "Procedures for real-time Group 3 facsimile communication over IP networks", Edition 5 (2007) or Edition 6 (2010), available at http://www.itu.int) may be present on the media plane.

In contrast to the traditional PSTN (Public Switched Telephone Network) network both planes may be on different infrastructure using different protocols and even take different routes through a network.

Correlation Information

The property of disjunct signaling stream(s) and media stream(s) of a VoIP service within packet-switched networks may necessitate a correlation mechanism. This correlation may only be based on the common parameters carried on both planes of the session. Monitoring equipment used in VoIP networks may be limited to either the signaling or media plane due to the network structure. A correlation of signaling stream(s) and media stream(s) is thus needed outside the VoIP network in between the monitoring components.

In RTP-based VoIP networks the data basis is rather limited as defined by the RTP protocol and the underlying transport and Internet layer protocols. An RTP stream which typically carries voice or video information from one party to another of a VoIP service is made up of plural RTP packets. The header of each RTP packet does not contain an identifier which can be used to distinctively associate single RTP packets or the entire RTP stream to a VoIP service (defined by the VoIP signaling). The same way the VoIP signaling protocol does not carry information to relate a signaling session (signaling plane of the VoIP service) to RTP protocol header information.

The parties originating or receiving a VoIP service actively dictate the transport and Internet layer address (IP address and port) on which they want to receive media streams. This configuration is done on the signaling plane. Using SIP as a signaling protocol, the address tuple(s) to be used on the media plane is/are defined in a so-called SDP body (see IETF RFC 3264, "An Offer/Answer Model with the Session Description Protocol (SDP)", available at http://www.ietf.org). The address tuple(s) - usually consisting of an IP address and port – is/are the sole information to correlate media plane of a VoIP service with the signaling plane of the VoIP service, and vice versa.

Challenges in Correlation of Signaling Plane and Media Plane

Beside the limited information available for correlation, the correlation of the signaling plane and media plane imposes also numerous further challenges, especially in passive mid-point monitoring scenarios. While an address tuple (IP address and port) may be enough information to identify a single RTP stream at a given point in time, multiple streams

identified by different address tuples may exist in a regular VoIP call. In fact at least two media streams are needed to allow two-way communication between the caller and callee.

It will be outlined in further detail in the following how multiple media streams may occur in a single VoIP service and how they are communicated between the parties involved (which may not always be the case).

Mid-point monitoring equipment (i.e. the signaling and media probes) is introduced into the network path which active components (e.g. VoIP phones and servers) use to communicate with each other. This passive equipment is not involved in the communication itself and has to derive all information from the network packets it observes on said network path. The correlation challenge increases as the passive equipment lacks information the active components internally have. A VoIP phone may internally assign a media port to a signaling session but never send out this information combined so that a mid-point monitoring probe could benefit from this information.

Time Constraints due to Port Reuse

The IP address and port combination transmitted inside the signaling plane of a VoIP service as well as the IP address and port observed on the media plane of the VoIP service may be reused. For exemplary purposes it can be assumed that each VoIP device has a single IP address. The port range from which the client can draw is limited to roughly 65000 ports by the underlying transport layer protocol UDP (see RFC 768, "User Datagram Protocol", available at http://www.ietf.org). In practice additional specifications and configuration parameters significantly further limit this port range. RFC 3550 which defines the Real-Time Transport Protocol mandates that only even port numbers may be used to transport media streams. The odd ports are reserved for Real-Time Transport Control Protocol (RTCP) packets. This standardization alone effectively halves the number of ports available for VoIP communication.

As a result a VoIP device has to re-use a certain IP address and port combination, thus eliminating a unique identification criterion. Essentially the IP address and port combination can only be assumed to be unique for a limited amount of time. The time-frame depends on the utilization of the VoIP device as well as the number of IP addresses required in an Internet Service Provider's network. If a correlation of the signaling and media plane takes place without a time criteria, it would therefore produce so-called false positives: Multiple media streams would match the requested IP address and port combination, since the signaling plane may have caused re-use of the very same IP address and/or port for another VoIP service.

Media Streams

Multiple streams may exist per direction of each VoIP service. The caller sends data to the callee and vice versa resulting in at least one media stream per direction. Desirable features, such as for example codec changes, may trigger additional streams. On the network multiple streams may use the same IP address and port tuple depending on the media protocol. For example, each RTP stream is additionally identified by a unique identifier - the Sender Synchronization Source (SSRC) – which is used to distinguish between multiple streams.

Fig. 1 shows exemplary situations where multiple media streams may be triggered. The focus is on the media plane traffic and its timing. The x-axis denotes the time basis, starting on the left side. The vertical lines in Fig. 1 denote signaling events and the bold horizontal lines are RTP streams present in the respective exemplified VoIP service example (please note that the different VoIP service examples are separated by broken horizontal lines). Only the RTP stream direction from the callee to the caller is shown per VoIP service example. The RTP stream from the caller to the callee can be assumed to be less complex (e.g. as shown in the "normal" VoIP service example).

The three vertical lines denote important events in the phase of a call. The leftmost line denotes the start timestamp. This is the time at which the caller submitted the call setup request to the network. The connected point is the timestamp at which a callee or his endpoint/handset picked up and moved the service into a fully established state. The service ends when one of the parties sends a termination request.

The most common VoIP services will be the trivial session examples where there is only one RTP stream per direction (see "normal" or "early media"). In a "normal" VoIP service flow both parties only start sending once the session is fully established on the signaling plane. More often the callee would start sending early media, i.e. when sending ring-tones or (pre-call) announcements. In "session forking" VoIP service example there may be multiple concurrent or sequential RTP streams at the beginning of a service e.g. due ring-tones being sent to multiple VoIP devices of the callee, while one is answered at the connect event

During a service media codecs may be changed (see "code change" VOIP service example) which is commonly resulting in a new media stream. Another service scenario may be the VoIP call being put on hold and picked up at a later time essentially generating multiple media streams (see "call on hold" VoIP service example). Silence suppression and other RTP features may cause multiple streams sporadically during a service (see "silence suppression" VoIP service example). Also retransmissions on the signaling plane may

desynchronize the signaling plane and media plane (see "SIP retransmissions" VoIP service examples).

Geographic Distribution or Media Plane and Signaling Plane

As indicated above, signaling plane data and media plane data of VoIP services may take different routes. In examples, where SIP is used for session setup, this is indeed the case as those sessions usually contain signaling-only components (i.e. SIP proxies, REGISTRARs). As a result passive mid-point monitoring solutions may be split geographically along the routes of the VoIP traffic.

An exemplary passive mid-point monitoring scenario for this case is shown in Fig. 2. Fig. 2 shows three geographically distributed POPs (Point of Presence) which are linked over the carrier's internal network. The signaling data flows on the signaling plane between POP A and B via the signaling-only POP C. POP C would typically host a centralized routing entity which reduces the complexity in the POPs A and B. To prevent back-hauling traffic from POP A to POP C before sending it to POP B the VoIP system is configured in a way that would allow the media to flow directly on the media plane between POPs A and B.

Typical mid-point monitoring locations would be inside each POP. Due to the layout of the network no media traffic would be visible in POP C. To allow for successful correlation of signaling and media streams by the correlation mechanism it is therefore desirable to support distributed networks. If the same session traverses multiple POPs it would be desirable for the correlation mechanism to identify this and support correlation of the session legs across multiple sites.

Limitations/Challenges in Correlation of Signaling Plane and Media Plane

The correlation may be solely based on the IP address and port used on the transport and Internet layer. On the media plane each RTP stream is destined to exactly one IP address and port on the transport layer. On the signaling plane one IP address and port is typically exchanged inside the Session Description Body of the SDP protocol during the VoIP service setup. However this IP address and port combination may not be the combination used throughout the VoIP service. In this case the above mentioned correlation mechanism fails: It fails upon searching a media stream for a specific signaling session of the VoIP service. The same way it fails to locate a signaling session of the VoIP service for a specific media stream on the media plane. The reasons for using another IP address port combination than advertised upon session setup are manifold.

Network Address Translation (NAT)

Network Address Translation (NAT, RFC 1631) was developed to solve Internet address space depletion and routing problems. In today's Internet one may distinguish between internal addresses which are often only unique locally (i.e. typically used inside company or home networks), and external addresses of global validity and uniqueness as used on the public Internet.

Fig. 3 shows an exemplary network using internal and external addresses. The internal network on the left side uses IP addresses from the so-called private space, as for example defined in IETF RFC 1918, "Address Allocation for Private Internets", available at http://www.ietf.org. The external network on the right side is the public Internet with globally reachable addresses. A router with embedded Network Address Translation (NAT) functionality is located in between the two networks. The router does not directly route packets in between the two connected networks, but "translates" (i.e. replaces) IP source and destination addresses within the routed IP packets on demand.

Essentially, the NAT functionality of the router hides the entire internal network (10.0.0.0/8 in this example) behind one public IP address (192.0.2.1). Every IP connection from within the internal network to the Internet will seem to originate from the sole public IP address (192.0.2.1) of the entire network.

For each node inside the internal network this imposes some limitations. Although Internet access works in general, a node in the internal network may not be reachable from the outside. This is due to the fact that IP traffic from an outside node (i.e. a node in the external network) would be destined to the public IP address (192.0.2.1), but the router with this IP address does not know at this point in time where to forward the traffic to on the internal network, given that there was no prior communication from an node inside of the internal network and the external node.

Additionally the internal node may not be aware that it is behind a router which performs the NAT. As required by the transport protocols, the internal node will use the internal IP address for communication purposes. On the VoIP signaling and media layer (i.e. session layer or application layer), the application will also use the internal IP address (i.e. 10.0.0.123) instead of the public IP address. VoIP signaling protocols (e.g. SIP or H.323) contain various IP addresses on the application layer. An intermediary NAT device (i.e. the router in Fig. 3) will not perform translation of the addresses on the session layer or application layer, but only translates IP addresses in the network layer. The internal IP addresses inside the VoIP

signaling (i.e. SIP or H.323 messages) impose a problem to the application as internal IP addresses are most likely not reachable from the public Internet.

VoIP devices involved in the processing of a certain call may detect this problem by comparing the IP addresses seen on the Internet and transport layers with the IP addresses set on the application layer. The following message is an exemplary SIP message as seen on the network and transport layer, which means that both the Internet and transport layer information (IP addresses and ports) are visible:

```
U 2010/09/09 13:57:31.925226 192.0.2.99:37682 -> 10.0.0.1:5060
```

The Internet and transport layers show that the UDP packet (U) was received from the external address 192.0.2.99 (port 37682).

The following message represents the same exemplary SIP message as seen on the application layer, which means that both the Internet and transport layer information (IP addresses and ports) as well as the application layer (SIP message contents) are visible:

```
INVITE sip:echo@example.com SIP/2.0
Via: SIP/2.0/UDP
192.168.1.216;branch=z9hG4bK3BB64874;rport=37682;received=192.0
.2.99
CSeq: 7726 INVITE
To: <sip:echo@example.com>
Content-Type: application/sdp
From: "John Doe" <sip:jdoe@example.com>;tag=693C7725
Call-ID: 134172356l@192.168.1.216
Subject: sip:jdoe@example.com
Content-Length: 230
User-Agent: kphone/4.2
Contact: "John Doe" <sip:jdoe@192.168.1.216;transport=udp>
v=0
o=username 0 0 IN IP4 192.168.1.216
s=The Funky Flow
c=IN IP4 192.168.1.216
t=0 0
m=audio 60942 RTP/AVP 0 97 8 3
a=rtpmap:0 PCMU/8000
a=rtpmap:3 GSM/8000
a=rtpmap:8 PCMA/8000
a=rtpmap:97 iLBC/8000
a=fmtp:97 mode=30
```

In contrast to the information available at the network and transport layers claiming that the message was received from the external address 192.0.2.99 (port 37682), the application claims that the source was the internal address 192.168.1.216 (port 5060 - the port 5060 is not viewed for the source address as it is the default port).

In general the translation of the network (IP) address also applies to the port numbers which are used on the transport layer. This feature is named Port Address Translation (PAT). NAT and PAT commonly happen at the same time in the same device.

PAT becomes necessary if multiple devices attempt to use the same port. Since different devices have different IP addresses their IP address and port tuple is unique. If packet streams from both devices traverse the same NAT device, the NAT device has to translate the port of one of the two packet streams. Otherwise there would be two distinct packet streams sharing the same external IP address of the NAT device and the same port. While this may not be a problem if the destination IP address and port tuple of the destination device are different for the two packet streams, it adds overhead which the NAT device wants to prevent. By rewriting the source port of one of the packet streams the address tuples of the external IP address and ports become unique again.

However, in networks where NAT (and PAT) is performed to for example hide private networks from the outside Internet, the correlation of signaling sessions of VoIP services and the associated media plane streams is becoming complex.

Multiple Communication Paths

Numerous solutions have been developed to allow traffic of VoIP services to operate in networks with Network Address Translation (NAT). These mechanisms require the involved caller and callee to exchange information, for instance on additional transport protocols or addresses which can be used to reach the remote party. These include – but are not limited to - IPv6 (Internet Protocol Version 6, see IETF RFC 2460, "Internet Protocol, Version 6 (IPv6), Specification"), TURN (Traversal Using Relays around NAT, see IETF RFC 5766, "Traversal Using Relays around NAT (TURN):Relay Extensions to Session Traversal Utilities for NAT (STUN)") addresses or ICE (Interactive Connection Establishment, see IETF RFC 5245, "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols") candidates (all RFCs available at http://www.ietf.org).

These additional tuples of IP addresses and ports communicated in these solutions establish different data paths in media plane and signaling plane of a VoIP service and thereby break

up the previous one-to-one relation of signaling session to media stream to a given VoIP device. Previously one media stream to the given VoIP device could be distinctively identified by exactly one IP address and port. In the new scenario on the signaling plane a list of IP address and port tuples is transported on which the VoIP device may receive media thereby increasing the complexity of correlation between media plane and signaling plane of a VoIP service.

SUMMARY OF THE INVENTION

One object of the invention is to suggest correlation mechanisms that allow overcoming one or more of the limitations and challenges on the correlation of signaling plane and media plane of services.

The object is solved by the subject matter of the independent claims. Advantageous embodiments of the invention are subject to the dependent claims.

One of the various aspects of the invention is to suggest a definition of correlation keys that allow the correlation of media streams of services and the associated signaling session of the respective services based on data records (reports) that provide information on the media plane and signaling plane of the respective services. These correlation keys may be for example provided in form of hash keys. A correlation key may be associated to a given direction of the media data paths of the media streams belonging to a given service (e.g. caller-to-callee or callee-to-caller for a simple bi-directional service, e.g. VoIP call). Hence, assuming that the service comprises at least two media streams of different directions, at least two correlation keys are provided on the signaling plane for later correlation of the media plane and signaling plane of the service. On the media plane one correlation key is provided per media stream.

The correlation keys for use in correlation are generated by a media plane probe and a signaling plane probe independently from each other, but the two probes generate matching correlation keys for a given media service (i.e. the signaling plane probe and the media plane probe independently each generate the same correlation key for a given direction/media data path of the media service, given that there are not any exceptional situations preventing the match of the correlation keys).

In response to the detection of the end of a signaling session, the signaling plane probe generates for this signaling session at least one correlation key for each "direction" of the service to which the signaling session belongs. Each correlation key (e.g. hash key) is generated by the signaling plane probe based on a timestamp indicative of the point in time

at which the end of the respective signaling session has been detected by the signaling plane probe and using information on respective direction of the media plane data flow (e.g. caller-to-callee or callee-to-caller for a simple bi-directional service, e.g. VoIP call). The correlation keys generated for the different directions of the service corresponding to the signaling session the end of which has been detected are then collated in a signaling session report, which may be for example stored in a database or a flat file.

For the media plane, the media plane probe generates a correlation key (e.g. hash key) for each media stream. The correlation keys for the media steams are generated upon the media plane probe detecting the end of the respective media stream. The correlation key for a respective terminated media stream is generated by the media plane probe based on the direction of the respective media stream, e.g. based on the media data path of a media stream, and on a timestamp indicative of the point in time at which the respective media stream has ended. The timestamp at which the respective media stream has ended is for example the time at which the last packet of the media stream has been received by the media plane probe. The media plane probe stores the correlation key (e.g. hash key) for the media stream in a media stream report, which may be for example stored in a database or exported as a flat file.

Another aspect of the invention is the use of the correlation keys comprised in the signaling session reports and the media stream reports to identify the media streams and signaling session belonging to a single media service. As indicated above, the correlation keys generated by the signaling plane probe and the media plane probe should be matching each other (i.e. the correlation keys of a given signaling session report should be found in the media stream reports, respectively, the correlation key of a media stream report should be found in one of the signaling session reports), provided that signaling session and related media streams are terminated nearly at the same point in time, i.e. within a given timeframe. Hence, a correlation unit can identify which media streams belong to which signaling session (or vice versa) and may therefore generate a service record that summarizes a service based on information contained in the media stream reports and information contained in the signaling session report of the respective service.

According to one exemplary embodiment of the invention, a method for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network is provided. It is assumed for exemplary purposes that one service has one signaling session.

In this method a signaling plane probe generates at least one correlation key for each direction of a service based on the information of the signaling session of the respective

service monitored at the signaling plane probe. As noted above the direction may be equivalent to a respective media data path of a media stream belonging to a given service. The correlation keys for a respective service may be generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and using information on the direction of the respective service. Furthermore, the signaling plane probe stores a signaling session report for each signaling session in a database. Each of the signaling session reports comprises at least two correlation keys generated for the respective directions of the service.

Moreover, in this method a media plane probe also generates a correlation key for each media stream monitored by the media plane probe. The media plane probe generates the correlation key using a timestamp indicative of the point in time of the end of the respective media stream as detected by the media plane probe and using information on the media data path of the respective media stream. The media plane probe stores a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

A correlation unit further correlates the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports.

In a more detailed exemplary embodiment of the invention, the generation of the respective correlation keys for the respective media streams by the media plane probe is based on the media data path information of the respective media stream and the timestamp indicative of the point in time at which the respective media stream ended.

According to another exemplary embodiment of the invention, information on the direction of the respective service – based on which the correlation keys are generated by signaling plane probe for a respective service – is information on one or more media data paths of media streams of the service learned from messages of the signaling session of the service. In other words, the direction of the service may be for example given by the media data path which is yielding a given direction. For example in a simple VoIP call between two peer devices, there are two directions: caller-to-callee and callee-to-caller as noted above.

In a further embodiment of the invention, the signaling plane probe monitors signaling messages of the respective signaling sessions of the services, and detects media streams belonging to each service associated to a respective monitored signaling session based on the signaling messages of the respective signaling sessions. The correlation keys for a respective one of the signaling sessions may be generated by the signaling plane probe

based on information on the media data path of the media streams belonging to the service associated to the respective one of the signaling sessions and the timestamp indicative of the point in time at which the end of the respective signaling session has been detected by the signaling plane probe. For example, in one possible implementation information on a media data path of a media stream is a destination IP address and port number of the media stream.

In another embodiment, the timestamps used for the generation of the correlation keys by the signaling plane probe and the media plane probe could be for example rounded (e.g. in a configurable or fixed interval) before the generation of the correlation keys, in order to ensure that timing differences in the events monitored in the media plane and signaling plane of the services do not lead to generation of non-matching correlation keys for a given service. Alternatively, in another exemplary implementation, the timestamps – by default – have a granularity that is higher than one second, e.g. a granularity in the range of one to ten seconds.

In a further embodiment, instead of storing a single correlation key per direction of a signaling session or per media stream, additional correlation keys may be generated and stored. The input for these additional correlation keys are the same IP address and port information, but a different timestamp. For example, the signaling plane probe could calculate an additional correlation key using a timestamp earlier than the timestamp indicative of the end of the respective signaling session.

In case of using a rounded timestamp this earlier timestamp may be the next earlier rounded timestamp before the rounded timestamp indicative of the end of the respective signaling session. In one exemplary implementation, the signaling plane probe could generate two timestamps from a (non-rounded) timestamp indicative of the end of the respective signaling session which are each used together with the same IP address and port information to generate two respective correlation keys for a given direction: One of these two generated timestamps used for calculation of the correlation keys may be the timestamp indicative of the end of the respective signaling session rounded to the nearest lower value according to the rounding interval, while the other of these two generated timestamps may be the timestamp indicative of the end of the respective signaling session rounded to the nearest higher value according to the rounding interval.

All correlation keys that are generated by the signaling plane probe for a given direction may be stored as part of the signaling session records.

In a similar fashion as described above for the signaling plane probe and the directions indicated in the signaling, also the media plane probe may optionally generate more than one correlation key for a given media data path/media stream. In case of using rounded timestamps also the media plane probe could generate two timestamps from a (non-rounded) timestamp indicative of the end of the respective media stream which are each used together with the same IP address and port information of the media stream to generate two respective correlation keys. One of the two timestamps generated by the media plane probe may be the timestamp indicative of the end of the respective media stream rounded to the nearest lower value according to the rounding interval, while the other of these two generated timestamps may be the timestamp indicative of the end of the respective media stream rounded to the nearest higher value according to the rounding interval. All correlation keys that are generated by the media plane probe for a given media stream may be stored as part of the media stream records.

In case of rounding timestamps in the media plane probe and the signaling plane probe, the rounding interval in both probes should be identical. The gain of having both probes generating additional correlation keys and correlating a multitude of correlation keys may not reflect in the increased processing needs.

In a further exemplary embodiment of the invention, the correlation of the media streams and the signaling session of the respective services is performed by a correlation unit periodically or in response to a query.

Another exemplary embodiment of the invention is related to a method for correlating media streams and signaling sessions of services. According to this method a correlation unit correlates the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services. For correlating media streams and the signaling sessions, the correlation unit identifies one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports. Further, the correlation unit generates a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained in the signaling session report of said at least one service.

In one exemplary implementation, these steps are performed by a correlation unit periodically for all signaling session reports and media stream reports generated within a configurable time interval, or in response to a request.

In case the correlation unit is unable to match one or more media streams to said signaling session report, the correlation unit may optionally generate a service record without quality information of the media streams of the service or may optionally generate a corresponding indication of the user/operator of the correlation unit.

A service report generated for a service by the correlation unit may comprise – according to one exemplary implementation – information on the service and the quality information of the media streams of the service.

In one exemplary implementation, the correlation keys provided in a respective signaling session report may have been generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the corresponding service has been detected and using information on the media streams of the service. Similarly, the correlation keys provided in a respective media stream report may have been generated using a timestamp indicative of the point in time at which the respective media stream ended and using information on the media data path of the media stream.

In one further exemplary implementation, the correlation unit first attempts a correlation of media stream reports and signaling session reports using the correlation keys that have been generated based on the timestamps indicative of the end of the signaling session, respectively, the end of the media stream(s). In case the correlation is unsuccessful, the correlation unit may further use optionally available additional correlation keys contained in the signaling session reports and/or media plane reports which are based on other timestamps, e.g. before and after the actual end of the respective signaling session or media stream(s), respectively. This layered approach allows to first conduct the correlation using the correlation keys which are more likely to yield a result and as a fallback conduct the correlation using a wider timeframe.

According to a more detailed exemplary embodiment of the invention, the correlation of the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions by the correlation unit comprises matching the respective correlation keys of the media stream reports of media streams are to the respective correlation keys of the signaling session reports, or vice versa. A match of the correlation keys of a one of the media stream reports and one of the signaling session reports (i.e. the

two correlation keys are identical) yields that the media stream of said one media stream report belongs to the same (media) service to which the signaling session of said one signaling session report belongs.

In another exemplary embodiment of the invention, for correlating the media streams and the signaling sessions of one or more services by the correlation unit, same may obtain from a signaling session report correlation keys identifying the media streams belonging to the same service as the signaling session. Further the correlation unit could then identify those media stream reports which include one of the correlation keys obtained from the signaling session report.

Alternatively, in another exemplary embodiment, the correlation of the media streams and the signaling sessions of one or more services could also be implemented by the correlation unit obtaining from a media stream report a correlation key identifying the media stream, and further the correlation unit identifying a signaling session report among the signaling session reports which includes the same correlation key obtained from the signaling session report.

According to one exemplary embodiment of the invention, a method for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network is provided. It is assumed for exemplary purposes that one service has one signaling session.

Another embodiment of the invention is related to a passive monitoring system for correlating media streams and signaling sessions of services in a packet-switched network. The passive monitoring system comprises a signaling plane probe that generates at least one correlation key for each direction of a service based on the information of the signaling session of the respective service monitored at the signaling plane probe. As noted above, there may be one or more correlation keys for each direction of the service. The correlation keys for a respective service are generated by the signaling plane probe using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and further using information on the direction of the respective service (e.g. information on the media data path of the given direction, such as a destination IP address and UDP port number). The signaling plane probe stores a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service.

Furthermore, the passive monitoring system also comprises a media plane probe for generating a correlation key for each media stream monitored by the media plane probe. The

correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream. The media plane probe stores a media stream report for each media stream in a database. Each media stream report comprises at least the correlation key generated for the respective media.

The passive monitoring system further comprises a correlation unit for correlating the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports. The correlation unit may for example use one of the different correlation methods described herein for this purpose.

Another embodiment of the invention is related to the design of the signaling plane probe. According to this embodiment of the invention, the signaling plane probe comprises a network interface card for receiving packets of signaling sessions of multiple media services and a monitoring unit for monitoring the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of the signaling sessions. The monitoring unit of the signaling plane probe is further capable of detecting the point in time at which a respective signaling session of the respective service terminates.

The signaling plane probe also comprises a processing unit for generating at least one correlation key for each direction of a service based on a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the monitoring unit and based on information on the direction of the respective service. Moreover, the monitoring unit stores a signaling session report for each signaling session in a database, wherein each signaling session report comprises at least two correlation keys generated for the respective directions of the service.

Another embodiment of the invention is related to a media plane probe which comprises a network interface card for receiving packets of media streams of multiple media services, and a monitoring unit for monitoring the media streams based on the packets of the media streams received by the media plane probe. Thereby the media plane probe obtains information on a media data path a respective one of the media stream. Further, the monitoring unit detects the point in time at which a respective media stream of the respective service terminates, and a processing unit of the media plane probe generates a correlation key for each media stream monitored by the monitoring unit. This correlation key may be for example generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream.

The media plane probe's monitoring unit further stores a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

The invention according to another embodiment is related to a correlation unit capable of correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network. The correlation unit comprises a database access unit for accessing a database comprising signaling session reports for signaling sessions of services and a database comprising media stream reports of media streams of the services. The signaling session reports each comprise at least one correlation key for each direction of a respective service, and the media stream reports each comprise one correlation key indicating the direction of the media stream. A processing unit of the correlation unit is used to correlate the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services. The processing unit identifies one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports. The correlation unit may further comprise a service record generation unit (e.g. implemented by a processing unit) for generating a service record for a respective service based on information comprised in the media stream reports and information comprised the signaling session report of said at least one service.

In a further embodiment of the invention, the correlation unit is further adapted, e.g. by comprised respective configured means, to perform the steps of the correlation methods according to one of the different embodiments described herein. Furthermore, also the signaling plane probe and the media plane probe may be adapted, e.g. by comprised respective configured means, to implement the different functions thereof as described herein.

Further embodiments of the invention relate to the implementation of the different functions of media plane probe, signaling plane probe and correlation unit in software. In this context, another embodiment of the invention provides a computer-readable medium storing instructions that, when executed by a processor of a signaling plane probe, cause the signaling plane probe to receive packets of signaling sessions of multiple media services, and to monitor the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of the signaling sessions. The execution of the instructions further cause the signaling plane probe to detect the point in time at which a respective signaling session of the respective

service terminates, and to generate at least one correlation key for each direction of a service based a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected and based on information on the direction of the respective service. Moreover, the executed instructions may further cause the signaling plane probe to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service.

Another exemplary embodiment of the invention is providing a computer-readable medium storing instructions that, when executed by a processor of a media plane probe, cause the media plane probe to receive packets of media streams of multiple media services, and to monitor the media streams based on the packets of the media streams received by the media plane probe to thereby obtain information on a media data path a respective one of the media streams. Moreover, the execution of the instructions cause the media plane probe to detect the point in time at which a respective media stream of the respective service terminates, and to generate a correlation key for each media stream monitored by the monitoring unit, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream. Furthermore, the execution of the instructions may further cause the media plane probe to store a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

Another embodiment of the invention is providing a computer-readable medium storing instructions that, when executed by a correlation unit, cause the correlation unit to correlate media streams and signaling sessions of services in a passive monitoring system of a packet-switched network. The execution of the instructions cause the correlation unit to access a database comprising signaling session reports for signaling sessions of services and a database comprising media stream reports of media streams of the services. The signaling session reports each comprise at least one correlation key for each direction of a respective service, and the media stream reports each comprise one correlation key indication the direction of the media stream. The execution of the instructions further cause the correlation unit to correlate the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services and to identify one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports.

The execution of the instructions also cause the correlation unit to generate a service record for a respective one of said at least one service based on information comprised in the media stream reports and information comprised the signaling session report of said at least one service.

Another embodiment is providing a computer readable medium storing instruction that, when executed by the processor of the correlation unit, cause the correlation unit to perform the steps of the method for correlating media streams and signaling sessions of services according to one of the different embodiments described herein.

BRIEF DESCRIPTION OF THE FIGURES

In the following the invention is described in more detail in reference to the attached figures and drawings. Similar or corresponding details in the figures are marked with the same reference numerals.

Fig. 1 shows exemplary situations where multiple media streams may be triggered in a media service,

Fig. 2 shows three geographically distributed POPs (Point of Presence) which are linked over the carrier's internal network,

Fig. 3 shows an exemplary network using internal and external addresses,

Fig. 4 shows an exemplary overview on a system according to an embodiment of the invention comprising signaling plane probe, a correlation unit and a media plane probe in a packet-switched network,

Fig. 5 shows a flow chart of the operation of the correlation unit according to an exemplary embodiment of the invention,

Fig. 6 shows the structure of a correlation unit according to an exemplary embodiment of the invention,

Figs. 7 & 8 shows two exemplary flow charts of the operation of a signaling plane probe according to different exemplary embodiments of the invention,

Fig. 9 shows the structure of a signaling plane probe according to an embodiment of the invention,

Fig 10 shows the structure of a media plane probe according to an embodiment of the invention,

Fig. 11 shows a flow chart of the operation of a media plane probe according to an embodiment of the invention,

Fig. 12 shows a network in which locations are indicated, where signaling and media plane probes of the invention may be advantageously located according to an embodiment of the invention,

Fig. 13 illustrates the use of rounded timestamps within a signaling plane probe and a media plane probe to highlight potential sources of error in correlation, and

Fig. 14 illustrates the use of two rounded timestamps for generating respective correlation keys according to an exemplary embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

The following paragraphs will describe various embodiments of the invention. For exemplary purposes only, most of the embodiments are outlined in relation mainly to VoIP services as an example for a media service, but the invention is not limited to VoIP services. For example, the invention may also be used for other types of media services provided in packet-switched networks. The invention may be advantageously used in media services which have signaling plane and media plane separated and/or taking different routes through the packet-switched network.

The explanations given in the Technical Background section above are intended to better understand the often VoIP call-related exemplary embodiments described herein and should not be understood as limiting the invention to the described specific implementations of processes and functions in a packet-switched communication network.

As indicated above, the invention generally relates to methods for correlating at least one media data path and a signaling session of a service in a passive monitoring system that is monitoring a packet-switched network. Furthermore, the invention also relates to an implementation of these methods in hardware and software, and provides a signaling plane probe, a media plane probe and a correlation unit, as well as passive monitoring system comprising same. Based on the assumption of disjunctive signaling and media components of a service, the correlation mechanism suggested herein supports correlation across

multiple hosts connected via a network (e.g. a Local Area Network (LAN), Wide Area Network (WAN) or the Internet).

One of the various aspects of the invention is to suggest a definition of correlation keys that allow the correlation of media streams of services and the associated signaling session of the respective services based on data records (reports) that provide information on the media plane and signaling plane of the respective services. In the following these correlation keys are implemented as in form of hash keys for exemplary purposes. The correlation keys may be used to query database(s) to correlate media stream reports for media streams monitored by a media plane probe and a signaling session report on the signaling session provided by a signaling plane probe that belongs to the same service – it should be noted that for the proper functioning of the invention, it is not decisive whether media stream reports and signaling session reports are stored in the same or different databases, or whether there are even multiple database partitions provided. The media stream reports and signaling session reports may even be stored in files.

The correlation keys are generated by one or more media plane probes and one or more signaling plane probes that monitor a packet-switched network independently from each other. This independent calculation of the correlation keys facilitates distributed correlation spanning across multiple geographically distributed components or even across multiple vendors of software components. Furthermore, as will become more apparent from the following, for a given media data path – respectively a corresponding given direction –, the parameters that are used for generating correlation keys at the signaling plane probe(s) and the media plane probe(s) ensure that the identical correlation keys are generated for the given media data path by the signaling plane probe(s) and the media plane probe(s) to facilitate correlation. A respective correlation key may be associated to a given direction of the media data paths of the media streams belonging to a given service. For example, assuming for exemplary purposes a simple bi-directional service, e.g. VoIP call, there are two directions: caller-to-callee and callee-to-caller.

The signaling plane probe may generate a correlation key for each direction of the service. For example, the signaling plane probe generates one correlation key for the caller-to-callee direction and another correlation key for the callee-to-caller direction. However, the signaling plane probe may also generate multiple correlation keys per direction, as will be outlined below in further detail. Consequently, in case it is assumed that a service comprises at least two media streams of different directions, then at least two correlation keys are calculated by the signaling plane probe for later correlation of the media plane and signaling plane of the service. The correlation key for a given direction or media data path may be generated based on information on the direction or media data path, such as for example a tuple comprising

the destination IP address and UDP port for a media stream in the service's media plane corresponding to the given direction or media data path. In addition, the generation of the correlation key is further based on a timestamp. With respect to the generation of the correlation key by the signaling plane probe, the timestamp may for example indicate the point in time at which the signaling plane probe has detected the end of the signaling session (and thus the service). In another exemplary implementation, the timestamp used in the generation of a correlation key for a media data path of a media stream in the media plane of the service could also be the point in time of the termination (end) of the media stream in the media plane detected by the signaling plane probe.

The signaling plane probe may generate the correlation keys for the different endpoints (the endpoint indicate a direction of the related media data flow) of a service upon detection of the end of a signaling session and additionally or alternatively, upon detection of the end of a media stream in the service's media plane. Typically but not limited thereto, there is one endpoint per direction, e.g. the caller and callee. The callee may be represented by multiple endpoints in case of more complex media service scenarios, e.g. when session forking is used (see Fig. 1).

The end of a signaling session may be for example detected in the signaling plane probe by detecting a signaling protocol specific hang-up message, a terminal failure to connect the involved endpoints or a session timeout. Each correlation key (e.g. hash key) is then generated by the signaling plane probe based on a timestamp indicative of the point in time at which the end of the respective signaling session has been detected by the signaling plane probe and using information on respective direction of the media plane data flow (e.g. caller-to-callee or callee-to-caller for a simple bi-directional service, e.g. VoIP call).

The direction may be for example given by a media data path of a media stream of the respective service. A media data path could be for example indicated by a tuple comprising the destination IP address and UDP port number for transmitting the media packets of a media stream of the service. The signaling plane probe may thereby generate only one single correlation key per direction or multiple correlation keys per direction. In one exemplary embodiment of the invention, the signaling plane probe generates one correlation key per direction, where the correlation key for a respective direction is based on a media data path of a media stream. This media stream of which the media data path information is used to generate the correlation key should be existent at the end of the signaling session (according to the session signaling monitored by the signaling plane probe), in order to be able to match the correlation key to a corresponding correlation key in a media stream report generated by a media plane probe. The media data path(s) of a service and optionally their

existence (e.g. start time and termination time) could be for example learned by the signaling plane probe from the messages of the signaling session of the service.

The correlation keys generated for the different directions of the service corresponding to the signaling session the end of which has been detected are then comprised in a signaling session report, which may be for example stored in a database or (flat) file.

Alternatively, the signaling plane probe may also generate a correlation key for each media stream in the media plane of the service. In this case it may be assumed that the signaling plane traffic, i.e. the packets of the signaling protocol of the signaling plane, allow the signaling plane probe to detect the termination (i.e. end) of respective media streams in the media plane. In this case, the detection of the end of a media stream could trigger the generation of a correlation key for the respective media stream based on information on its media data path and a timestamp indicating the point in time at which the signaling plane probe detected the termination of the media stream.

Further, in another alternative implementation, the detection of the end of a signaling session, respectively, service as well as the end of a media stream as detected by a signaling plane probe could be used to trigger the generation of correlation keys in the signaling plane probe. For example, during the ongoing service the detection of the termination of a media stream in the media plane as learned from the signaling plane traffic could trigger the signaling plane probe to generate and include respective correlation keys for the media streams terminated during the ongoing session into the signaling session report of the service, while the detection of the end of the signaling session, and thus the service, could be used as a trigger for generating correlation keys for the respective media streams being "alive" in the media plane of the service at the signaling plane probe and their inclusion into the signaling session report of the service.

In the media plane, the media plane probe generates a correlation key (e.g. hash key) for each media stream. The correlation keys for the media steams are generated upon the media plane probe detecting the end of the respective media stream. The correlation key for a respective terminated media stream is calculated by the media plane probe based on the direction of the respective media stream, e.g. based on the media data path of a media stream, and on a timestamp indicative of the point in time at which the respective media stream ended. This point in time is for example the point in time at which the last packet of the media stream has been received at the media plane probe. In this latter case, the media plane probe may thus maintain a timestamp indicating the time of reception of the last packet of the media stream at the media plane probe.

The timestamp used for the generation of the correlation key may indicate another point in time than the point in time at which the media plane probe detects the end of the media stream. For example, in a packet switched network, there may be interruption in the packet stream of the media stream so that there may be a timer that measures the time since the last reception of a packet of the media stream. In case the time since the last reception of a packet of the media stream is exceeding a threshold (e.g. the timer "expires") – for example after 15 or 30 seconds, or the like – then the media plane probe detects the end of the media stream. The timestamp used in the correlation key generation may nevertheless be the point in time at which the last packet of the media stream has been received. The media plane probe stores the correlation key (e.g. hash key) for the media stream in a media stream report, which may be for example stored in a database or exported as a flat file.

Furthermore the timestamp used for the generation of a correlation key may be rounded. This counters differences in the time basis (local time) of multiple signaling and media probes as slightly different timestamps still yield the same rounded timestamp. Additional correlation keys may be generated with different timestamps, e.g. indicative of the – optionally rounded – end of a respective signaling session and/or media stream or a timeframe before or thereafter. These additional correlation keys may be stored in the respective signaling session reports and/or media session reports.

A media stream may also be interrupted due to codec features such as Voice Activity Detection (VAD). If a sender of a media stream is mute the codec may not encode any data resulting in an interruption of the media stream. This is typically indicated in the packet stream transporting the media and should not be considered by the media plane probe as an indication of the end of the media stream itself. In case a timer mechanism is used to detect terminated media streams, the timer should be configured such that codec features that lead to a temporary interruption of the media stream are not triggering the termination of the media stream (according to the detection mechanism for detecting terminated media streams in the media plane probe).

Furthermore, in another exemplary embodiment, the correlation of media streams of services and the signaling sessions of the services may also be based on a compound keys used as a correlation keys. This exemplary alternative implementation may however require a more complex database structure of the media stream reports and signaling session reports, as will be outlined below.

Another further aspect of the invention is the use of the correlation keys comprised in the signaling session reports and the media stream reports to identify the media streams and signaling session belonging to the respective services. This correlation of media streams and

signaling session could be done for example periodically or on request. The correlation keys generated by the signaling plane probe and the media plane probe should be matching each other (i.e. the correlation keys of a given signaling session report should be found in the media stream reports, respectively, the correlation key of a media stream report should be found in one of the signaling session reports). Hence, a correlation unit can identify which media streams belong to which signaling session (or vice versa) based on the correlation keys contained in the media stream reports and the signaling session reports and may therefore generate a service record for a respective service based on information contained in the media stream reports and information contained the signaling session report of the service.

The correlation methods according to the different embodiments of the invention proposed herein are based on a probe monitoring the signaling plane traffic of services (i.e. the signaling plane probe) identifying new services and associated media data paths of the media plane by detecting their corresponding signaling sessions and determining for each service correlation keys allowing to identify at least the media streams of the media plane existent upon termination of the service. It should be noted that "at least the media streams" means that it is optionally possible to generate a correlation key for each media data path of the media plane learned from the signaling plane traffic of the service, in case it is possible to detect the termination of the media data paths of the media plane in the signaling plane traffic. In this case the correlation key for a respective media stream may be based on the information on the media data path of the respective media stream visible in the signaling plane of the service and the point in time at which the signaling plane probe detects the termination of the media stream in the signaling plane.

Regarding the use of the point in time of the termination of the signaling session (e.g. a signaling message in the signaling plane indicating "hang up") for the generation of correlation keys it should be noted that the termination of the signaling session also means that all existent media data paths, respectively, the corresponding media streams of the media plane are terminated. Hence, in case it is not possible to detect the termination of the media data paths of the media plane in the signaling plane traffic, it is assumed to be possible to detect at least the end of the signaling session and to thereby conclude on the termination of the respective media data paths of the media plane and to thereby facilitate correlation of the media streams that have been "alive" at the end of the service and the corresponding signaling session of the service.

As noted above, another possibility may be to use compound keys as correlation keys. A compound key is defining a set of information comprised e.g. in a given set of fields of a database entry of a signaling session report and media stream report that – together with

additional criteria on the point in time of the termination of a media stream or the service – allows to correlate media streams and signaling sessions of the services. In one exemplary embodiment, the one or more databases that are used to store the signaling session reports and the media stream reports are partitioned into partitions covering only a given amount of time (i.e. time interval) of signaling session reports and/or media stream reports, respectively. The "size", i.e. the time interval covered by the partitions may be for example determined based on the average time interval in which port reuse occurs in the target monitored packet-switched network (e.g. the partition could cover ½ of the average time interval in which port reuse occurs). In the given partitions it may thus be assumed that the media data path information, such as destination IP address and UDP port number are unique. In this case, the signaling session reports and media stream reports may not include separate correlation keys, but the correlation key can be generated on-demand as a compound key from the information on the media data paths contained in the signaling session reports and media stream reports.

Hence, it would be sufficient that the signaling session reports and media stream reports store information on the media data paths of the service (e.g. destination IP address and UDP port number) and a timestamp that indicates the point in time of termination of the respective media data path in the media plane, as detected by the signaling plane probe from the monitored signaling sessions and the media plane probe from the monitored media streams. As noted above, the signaling plane probe can also detect the point in time of termination of the media data path(s) in the media plane by detecting the point in time of termination of the signaling session, which is – in most practical cases – coinciding with the termination of all media streams of the service in the media plane as well.

In this case, the correlation could be realized by first selecting the database partition of the media plane reports according to a timestamp of the termination of a media data path/signaling session within a given signaling session report, and by then determining the media stream reports that comprise media stream information matching those indicated in the given signaling session report.

Fig. 4 shows a packet switched network 400 according to an exemplary embodiment of the invention that comprises a signaling plane probe 401, a correlation unit 402 and a media plane probe 403. Please note that more than one signaling plane probe 401 and media plane probe 403 may be provided. Optionally, there may be multiple correlation units 402 present, working independent of each other and not impacting the performance or correlation accuracy of any other component.

The signaling monitoring probe 401 and the media plane probe 402 are monitoring the traffic on a packet-switched network 400, which may be for example an IP based packet switched network. It is assumed for exemplary purposes that the signaling plane probe 401 and the media plane probe 403 monitor the traffic of the packet-switched network 400 at different locations of the packet-switched network 400, but observe the same network traffic.

In one exemplary implementation, the signaling plane probe 401 and the media plane probe 403 are probes of a minimal non-intrusive, passive monitoring system. The respective monitoring probes 401, 403 are connected to the packet-switched network 400 via respective test access ports (TAPs). A TAP is a passive network device that mirrors network traffic of the packet-switched network without interference of the original network traffic. It provides a copy of every packet sent or received on the network, by duplicating the data packets observed on the physical layer of the network. Instead of a TAP a Switch Port Analyzer (SPAN) port may be used to monitor the traffic of the packet-switched network.

In packet-switched environments, the signaling data of the signaling plane and media data of the media plane is packaged into packets on the transmitter side. Packets are transported over the packet-switched network independent of each other. The receiver of the packets – identified by the destination IP address and port – may have to regroup and sometimes reorder packets in order to compile streams from single packets.

On the signaling plane multiple packets belong to one signaling session. The first packet of a signaling session – usually containing a "call setup" message – starts a new signaling session (respectively a new service). At this point, the signaling plane probe 401 generates a session entry along with auxiliary data, i.e. a unique correlation key identifying the session internally (signaling session identifier). Packets following the initial signaling packet are matched to one of the currently ongoing signaling sessions (respectively services) and signaling sessions managed by the signaling plane probe. Packets which do not start a new signaling session and do not belong to existing sessions may be discarded.

Matching single signaling packets to existing signaling sessions is specified in the standards which define the signaling protocol itself. Transport layer information like IP addresses and protocol-specific information (e.g. SIP Call-IDs and tags) may be used.

A signaling session may be torn down upon reception of a special packet – a "call hang-up" message. This message again is specific to the signaling protocol used. Upon reception of this message the signaling plane probe 401 can conclude on the termination (end) of the signaling session and generates the correlation keys for use in correlation of media plane and signaling plane of the services. Furthermore, the signaling session may be removed

from its internal storage. Additionally the signaling monitoring component may remove sessions from its internal storage, if it identifies stale sessions.

The media plane 403 of the service may for example identify the media plane data in form of RTP streams. RTP streams are – analog to the signaling sessions – made up from individual RTP packets. The transmitter of an RTP stream generates RTP packets by splitting the media data into smaller pieces. In general these packets are relatively small to reduce the delay between sender and receiver. As a result many packets are created and form a single RTP stream – typically 50 packets per second. Single RTP packets can be grouped into an RTP stream by media plane probe 403. When using RTP, the media plane probe 403 has transport layer information as well as the RTP protocol header which contains a so called Synchronization source (SSRC). The SSRC is a 32-bit numeric identifier which is contained in every single packet and deemed mandatory by the RTP. When media plane probe 403 first receives a packet which does not match an existing RTP stream in its internal storage it creates an RTP stream. Consecutive RTP packets can then be matched to the existing streams by comparing the transport layer information and SSRC value. If no packets are received for a single RTP stream for a certain time the stream is considered ended and can be removed from the internal storage. A short interruption in the packet stream should not trigger the media plane probe 403 to mark an RTP stream as ended, thus the time before a media plane probe 403 can consider an RTP stream as ended is multiple seconds and has to take features like voice activity detection into account.

The signaling plane probe 401 and the media plane probe 403 generate information about the signaling session and media streams, respectively. This generated information is comprised in signaling session reports and media stream reports, respectively. In order to track running sessions and media streams the entities 401, 403 store information internally while the signaling sessions and media streams are in progress. When a signaling session or media stream is terminated the respective monitoring component 401, 403 may generate a signaling session report with signaling information, respectively media stream report with information on the media stream. As will be outlined below in further detail, the signaling session report provides a summary of the signaling session (service) and includes the correlation keys (e.g. hash keys) for correlation. Similarly, the media stream reports summarize the media stream and also include a respective correlation key (e.g. hash key) to facilitate correlation.

Each signaling session report for a signaling session may be for example written by the signaling plane probe 401 upon completion of a signaling session and contains data which describes the signaling session it monitored. For example, in one embodiment of the invention, the data stored in the signaling session report and describing the signaling session

include – but are not limited to – a set of caller and callee specific properties which will be available per endpoint (caller or callee). These endpoint-specific properties may for example include at least one of the following:

- endpoint address/name (e.g. SIP URI),
- phone number,
- User-Agent (endpoint sóftware and version),
- capabilities/features,
- supported media types,
- IP addresses and ports for signaling and media streams (from SDP body), and
- hash key of the media path (for example described by the destination IP address and UDP port of the expected media stream)
- signaling routing information.

Additionally, one or more of the following signaling session properties which are not endpoint-specific may also be included in the signaling plane report:

- timestamps: e.g. session start, Post-Dial-Delay (PDD), connect time, session end;
- information on the session flow (e.g. signaling message(s) with timestamps),
- session state (e.g. successful, unsuccessful, timed out, reason for failure), and
- events during session (media on hold, codec renegotiation, session forking, authentication)

In one exemplary embodiment of the invention, the signaling session report comprises all of the above noted endpoint-specific properties, and – optionally – all signaling session properties listed above, as far as they are available from the signaling protocol messages.

Depending on the specific implementation and capabilities of the signaling plane probe, the signaling session reports may contain information about multiple endpoints and/or multiple media streams per endpoint. In this case the session end timestamp may be available per involved endpoint, and not per session.

For example, the most trivial VoIP service is a phone call between two parties (including two media streams, one in each direction). In this case, the signaling session report could store each party as one endpoint with the above mentioned information and the media stream's destination IP address and port information. If the VoIP service would contain a video call (i.e. one audio and video stream per direction) the endpoint information may be extended to contain the destination IP address(es) and port numbers of the two media streams per direction (one IP address & port tuple for audio, one IP address & port tuple pair for video). Considering a VoIP service, the respective tuples of destination IP address and UDP port number of the media streams of the session may be for example learned by the signaling plane probe from the SDP body of SIP signaling.

Likewise, the signaling plane probe may extend the signaling session report to include additional endpoints, if multiple endpoints were part of the VoIP service. Multiple endpoints could be for example involved in a VoIP service, in case a caller tries to reach a callee which has call forwarding configured. The caller may then hear a ring-tone from the callee's landline, followed by a ring-tone from the callee's mobile phone. Finally a voicemail box picks up the session generating a third endpoint on the callee's side.

The media plane probe 403 monitors the media traffic, i.e. the individual media streams that convey the media information of the different services. The media plane probe 403 identifies a new media stream of a media service based on the network traffic it observes on the monitored packet switched network 400. The identification of a new media stream for a given service may happen shortly after this service is set up on the signaling plane. Throughout the duration of the media stream information on the media stream may be stored in an internal data record on the media steam. The media plane probe 403 may note for each stream the last RTP packet received for the respective media stream in an internal data record on the media stream in order to allow detecting the media stream's termination.

The media plane probe 403 may optionally perform a real-time analysis of the media streams monitored on the packet switched network and may generate a fine-grained media stream quality report with media stream related information. Fine-grained media stream quality reports could be for example written in predetermined time intervals, e.g. every five seconds, and could provide a summary of the media stream in the respective interval as well as quality information of the media stream in the respective interval. The information of the fine-grained media stream quality reports of a media stream could be used by the media plane probe 403 to generate a single media stream report on the media stream upon the end of the media stream is detected by the media plane probe 403. The media stream report and optionally also the fine-grained media stream quality reports of a media stream may be stored in individual databases.

The media stream report generated by the media plane probe 403 upon detecting the end of the media stream may contain data which describes the unidirectional media stream it monitored as well as correlation key that can be for example generated based on the destination IP address and UDP port number of the media stream and a timestamp that indicates the point in time at which the media stream ended/terminated. The media plane probe 403 could for example learn the destination IP address and UDP port number from the IP header encapsulating the RTP packets on the network layer. The characteristics of each media stream can be grouped into two categories. The transport and RTP header information characterize the media stream, e.g. for correlation purposes, whereas the quality information enable quality related processing of the media stream. In general the transport information is the same for all media packets processed as part of a single media stream. The stream quality information includes information covering all media packets and may be aggregated or stored on a per-packet basis.

In one exemplary embodiment of the invention, the media stream report generated by the media plane probe 403 for a given media stream may for example include one or more of the following transport media information:

- source and destination IP address,
- source and destination UDP port,
- RTP header Synchronization source (SSRC) value,
- RTP payload type, and
- correlation key(s) (e.g. hash value(s)) calculated based on destination IP address, destination UDP port and a timestamp indicating the point in time at which the of the media stream terminated (e.g. the timestamp of the last packet received for the media stream) or timestamp indicative of the timeframe before or after the termination timestamp.

The stream quality information that may be added to the media stream report by the media plane probe 403 may for example include/indicate one or more of the following:

- timestamp information (e.g. start and/or end time),
- RTP format violations,
- protocol conformance (e.g. sequence errors and/or DSCP errors),

- transport information (e.g. packet interval and jitter statistics, information on packet loss, information on duplicate packets, etc.), and

- MOS values (e.g. per time-slice (e.g. every 5 seconds), minimum/average/maximum value for entire stream duration, etc.),

- events (e.g. comfort noise, silence suppression, DTMF tones, codec changes, marker bits).

In one exemplary embodiment of the invention, the media stream report comprises all of the above noted transport media information and optionally further all stream quality information listed above, as far as they are available to the media plane probe from the packets of the media streams.

In addition to "measured values" (e.g. jitter between multiple media packets) the media stream report may also contain calculated information like MOS (Mean Opinion Score) values and information about impairments experienced by the receiver of the stream.

As further shown in Fig. 4, a correlation unit 402 correlates the media stream reports and signaling session reports generated by the signaling plane probe 401 and the media plane probe 403 in order to identify the media streams and the signaling session that belong to a respective service. It should be noted that each service is assumed to have one signaling session, but may have one or more media streams.

Generation of Correlation Key

The correlation key that is used for correlation may be built from:

- the timestamp that indicates the end of the signaling session, respectively the end of a given media stream, and

- information on a media stream, such as for example the destination IP address and UDP port of the media stream.

The correlation of signaling plane and media plane by the correlation unit 402 is based on this compound correlation key.

As noted earlier, the correlation key may be for example a hash key generated based on the before-mentioned parameters. The hash function to calculate the hash key (sometimes also called hash value) from the given data should be known to all parties involved (i.e. signaling plane probe 401, and media plane probe 403 and optionally also correlation unit 402) so that

the result for a given input is the same for all parties. Therefore, the hash keys are determined by the signaling plane probe 401 and media plane probe 403 using media data path information in addition to a timestamp that indicates point in time when the signaling session, respectively, the media stream ended/terminated. The correlation unit 402 may generate additional correlation keys based on the timestamps indicative of the timeframe before or after the timestamp that indicates the point in time when the respective session ended.

The hash function to determine the hash key may be a simple concatenation of one or more of the input parameters. For example, a concatenation of a destination IP address, destination UDP port number and a timestamp may (optionally including a delimiter character in between the parameters) be used for defining the hash keys. Of course more complex mechanisms may also be used to calculate a hash key representing the signaling session identifiers and/or the media path identifiers. For example the MD5 message digest function (see IETF RFC 1321, "The MD5 Message-Digest Algorithm", 1992, available at http://www.ietf.org) could be used to create a valid hash key. The MD5 message digest function may be fed with a concatenation of a destination IP address, destination UDP port number and the timestamp indicating the time of termination of the signaling session, respectively, the media stream and the resulting string can be used as a correlation key. Another hash function that could be used may be selected from the SHA cryptographic hash function family (see FIPS PUB-180, "Secure Hash Standard (SHS)", 2008 available at http://www.itl.nist.gov/fipspubs/). Furthermore, in another embodiment, a linear hash function may be used.

As outlined with respect to Fig. 1 earlier herein, ephemeral media streams may exist at any point during a service, such as a VoIP service. The ephemeral media streams may be generated by early media streams, service forking and codec changes. As a result the start and end timestamps of media streams may not necessarily concur with the start and end time of the VoIP service itself. Yet the timestamps of the start or end of a signaling session or media stream are an important piece of information for a successful correlation, as it significantly reduces the search window and amount of data to search. In typical scenarios, it can be expected that the start time of signaling and media streams deviate significantly and are thus not ideal candidates for a correlation mechanism.

At the end of a VoIP service both caller and callee have an established signaling session and send each other media on the media plane. Once a party hangs up the session a hang-up signal is send on the signaling plane and at the same time stops sending media packets. Once the remote party receives the hang-up signal it also stops transmitting media. The timestamps of the end of the signaling session of a service and the end of the corresponding

media streams in the media plane of the service are commonly very close to each other compared to the start of a media service in signaling plane and media plane. All these events will be subject to the same network- and application-imposed delays. Still the time-frame in which these happen is typically in the few milliseconds range or in the sub-second-range on long-distance or low bandwidth networks. This makes the timestamps of the end of the signaling session and the end of the media streams a good choice for calculating a correlation key.

Nevertheless, since signaling plane probe 401 and media plane probe 403 independently determine the correlation keys based on timestamps (and even minor deviations of the timestamps in the millisecond-range could lead non-matching correlation keys), the granularity of the timestamps used for the generation of the correlation keys may be chosen to account for those deviations of events in signaling plane and media plane events (i.e. deviations in the end of the signaling session in the signaling plane and the end of the corresponding media streams in the media plane of a service). Therefore the timestamps could have a granularity for example in the range of one or more seconds or at least 250 milliseconds or 500 milliseconds. In one exemplary embodiment, UNIX timestamps (which is indicating the seconds passed since January 1, 1970, 00:00:00) are used for the calculation of the correlation keys. Alternatively, also timestamps with a high resolution (e.g. resolution of one or ten or hundred milliseconds) could be used, but the timestamps may be rounded for the calculation of the correlation keys. It should be noted that the choice of the granularity or rounding interval may also impact the performance of a database or databases storing the media stream reports and signaling session reports for correlation. The larger the granularity, respectively the rounding interval and thus the search window the more data have to be processed to match signaling sessions and media streams.

Rounding timestamps could be for example implemented by rounding the timestamp to the nearest n second interval. Assuming for example that the interval is 5 seconds the following UNIX timestamps would be rounded as exemplified in Table 1.

Unix timestamp value	rounded value (5 sec)	description
1283936131	1283936130	rounded down
1283936133	1283936135	rounded up
1283936135	1283936135	unchanged
1283936136	1283936135	rounded down
1283936139	1283936140	rounded up

Table 1

The rounding interval can be any time interval in seconds or even milliseconds whereas the specific value is a trade-off between speed and accuracy. As noted above, the larger the interval the more data may have to be processed to match signaling sessions and media streams. The smaller the search window is the more likely it is to miss a match due to misaligned timestamps or derivations due to the delays outlined above.

The best rounding interval depends on the network monitored. Reasons for increased time differences are possible derivations from a shared time source (NTP) or differences seen in signaling and media events, i.e. the delay between a call being terminated and the endpoints stopping to send media.

For a smaller search window of three seconds the rounded values are different, as for example shown in Table 2.

Unix timestamp value	rounded value (3 sec)	description
1283936131	1283936130	rounded down
1283936133	1283936133	unchanged
1283936135	1283936136	rounded up
1283936136	1283936136	unchanged
1283936139	1283936139	unchanged

Table 2

Using normalized "session end" timestamp (i.e. timestamps of low granularity or timestamps that have been rounded) may be advantageous to counter the effects (e.g. false positive hash key matches) of the re-used port numbers. Since the end of a service (and/or end of media streams) becomes part of the correlation criteria only the port numbers that are used and valid at this point in time are considered in the correlation.

Another issue that could complicate the correlation processing is that deviations in detecting the end of a media stream, respectively a signaling session at the media plane probe 403 and the signaling plane probe 401 leads to the two entities rounding the timestamps to different values. For example, the signaling plane probe 401 could round the timestamp up, while the media plane probe 403 rounds the timestamp down to the next rounding interval value due to (minor) deviations in the timestamps at the signaling plane probe 401 and media plane probe 403. In another exemplary embodiment, additional correlation keys may be therefore generated by either one or both, the signaling plane probe 401 and media plane probe 403.

The signaling plane probe 401 and/or media plane probe 403 may use timestamps indicative of the rounding intervals before or after the timestamp indicative of the point in time where a

particular signaling session or media steam ended. For example a session ended at a point in time identified by the Unix timestamp value 1283936135. The correlation unit applies the rounding mechanism based on 3 second intervals. The rounded value thus is 1283936136. The correlation unit may generate additional keys with the timeframes before or after the mentioned rounded value. In this case the timestamps used for generation of the additional keys are 1283936133 (1283936136 – 3 seconds) and 1283936139 (1283936136 + 3 seconds). Of course, especially in cases where the rounding interval is short, it may also be desirable to more rounded timestamps rounded to the next and subsequent rounding interval(s) based on which correlation keys are calculated.

As noted above, the inclusion of multiple correlation keys calculated based on different timestamps for a given direction of a signaling session and/or media stream may for example be useful to counter correlation failures due to applying different rounding of timestamps in the signaling plane probe 401 and media plane probe 403 as will be outlined in further detail based on Fig. 13 and Fig. 14 in the following. The end of the signaling session $t_{signaling}$ of a service as observed by the signaling plane probe 401 and the end of the one or more media streams t_{stream} of the service as observed in the media plane probe 403 may not perfectly match each other as shown in Fig. 13. Even in case signaling plane probe 401 and media plane probe 403 calculate correlation keys based on the same IP address and UDP port tuple, the use of different timestamps $t_{signaling}$ and t_{stream} for the calculation of the correlation key will result in the correlation keys generated by signaling plane probe 401 and media plane probe 403 not matching each other. Additionally the use of rounded timestamps may not necessarily resolve this problem, as it cannot always be assured that the timestamps $t_{signaling}$ and t_{stream} are rounded to the same value. As also highlighted in Fig. 13 it could well happen that signaling plane probe 401 rounds the timestamp $t_{signaling}$ to the next lower rounded value T_1 in accordance with the rounding interval R, while the media plane probe 403 rounds the timestamp t_{stream} to the next higher rounded value T_2 in accordance with the rounding interval R. Again, correlation keys calculated on the different rounded timestamps T_1 and T_2 would not match each other.

As noted above, one possible solution to this problem is to include additional correlation keys at least to one of the signaling session reports and media stream reports. For example, in one exemplary embodiment, the media plane probe 403 could include two correlation keys per direction to the signaling session report. For a given media data path of a media stream (e.g. IP address and UDP port tuple) observed by the signaling plane probe, the signaling

plane probe could calculate one correlation key $C_{stream,1}$ on the media data path information and a rounded timestamp T_1, i.e. using the timestamp indicative of the end of the signaling session $t_{signaling}$ rounded to the next lower value according to the rounding interval, and the signaling plane probe could further calculate one correlation key $C_{stream,2}$ on the media data path information and a rounded timestamp T_2, i.e. using the timestamp indicative of the end of the signaling session $t_{signaling}$ rounded to the next higher value according to the rounding interval. Both correlation keys could then be included for the direction/media data path to the signaling session report.

A further implementation may include even further correlation keys, e.g. a third correlation key that is based on the rounded timestamp T_0, indicative of the next lower timestamp to T_1 according to the rounding interval.

The media plane probe 403 may calculate only one correlation key C_{stream} based on a rounded timestamp T_1 or T_2, i.e. based on the timestamp indicative of the end of the media stream (media data path) t_{stream} rounded to the next lower or higher value according to the rounding interval (and rounding rules). This correlation key for the media stream would then be included to the media stream report.

The correlation of the signaling session to the media stream(s) of the service would be successful in this embodiment, since either the correlation key $C_{stream,1}$ or the correlation key $C_{stream,2}$ within the signaling session report for the given direction/media stream would correspond to the correlation key C_{stream} included in the media stream's media stream report.

Optionally, in another embodiment of the invention, also the media plane probe 403 may calculate two or more correlation keys for each media stream, one correlation key based on the rounded timestamp T_1, i.e. based on the timestamp indicative of the end of the media stream (media data path) t_{stream} rounded to the next lower value according to the rounding interval, and the other correlation key based on the rounded timestamp T_2, i.e. based on the timestamp indicative of the end of the media stream (media data path) t_{stream} rounded to the next higher value according to the rounding interval. Both timestamps would then be included in the media stream report for the media stream by the media plane probe. In this exemplary implementation, the correlation may check whether the two correlation keys of the service's signaling session report for the direction/media stream match the two correlation keys within

the media stream report of the service's media stream to confirm the signaling session report and media stream report belonging to the same service.

Please note that instead of or in addition to the media plane probe 403 including two or more correlation keys to the media stream reports, also the signaling plane probe 401 could calculate two or more correlation keys per direction/media data path that is to be indicated in the signaling session report. Including multiple correlation keys per direction/media data path in the signaling session reports may be advantageous as it reduces the overhead information, since there are commonly less signaling session reports than media stream reports generated in the passive monitoring system.

Furthermore, in another exemplary implementation, the database(s) for signaling session reports and media stream reports may be further partitioned to reduce the search space for correlation. For example, the (respective) database could be partitioned in chunks of multiple hours, e.g. 4 or 6 hours or the like. The signaling session reports and media stream reports may further include as a separate element the timestamps of the end of the signaling session, respectively, the end of the media streams, based on which the database(s) can be partitioned and based on which the correct partition could be chosen for correlation later. Hence, in a more advanced implementation, besides the correlation keys a service end timestamp can be used as a search criteria to limit the amount of data that are to be searched in a database significantly.

In another exemplary implementation, the signaling plane probe may also generate a respective correlation key for all media streams of a service, in case the end of media streams is detectable within the session signaling, as noted previously. In this case the signaling session reports could further include a timestamp that indicates the end of each media stream, in order to allow choosing the correct partition of media stream reports in the database for correlation.

Correlation Mechanisms

For correlation of the media streams and the signaling sessions, different possibilities exist how the correlation by the correlation unit can be implemented. For example, according to exemplary embodiments of the invention, the correlation unit may correlate the media streams and the signaling session of the respective services periodically (or proactive) or in response to a query (i.e. on demand).

For example, a proactive correlation process could be utilized by a third party entity which has access to both signaling session reports and media stream reports. Their data may be

present in one or more databases to ease access to the raw data. As noted above, the correlation aims at correlating signaling and media information of the services and at merging them into combined service reports that may be then accessed by a third party, e.g. for evaluation of the services.

The proactive correlation process periodically takes all new signaling session reports which outline – but not fully describe – the media services, e.g. the monitored VoIP sessions. For each signaling session report the correlation process draws the calculated correlation keys from the signaling session report. There are at least two correlation keys in each signaling session report, at least one per direction. Additional correlation keys may optionally exist in the signaling session report, for instance if it describes a video call or a session with multiple endpoints. For each correlation key contained in a respective signaling session report the correlation process searches the media stream report database for the media streams with the same correlation key. If the correlation keys in the signaling session report are generated based on a timestamp indicating the end of the signaling session, every media stream of the service that existed at the end of the service in the media plane should have the same correlation key and will be in the result set of the search. If the correlation keys in the signaling session report also contain correlation keys for media streams that have been terminated during the ongoing service, the search may also further detect those media streams in the media stream report database based on the correlation key.

If a search returned two or more media stream reports, the correlation process encountered one of two scenarios. Either both media streams belong to the signaling session (i.e. service) or one of them was returned due to a correlation key collision. Both cases may be recovered by comparing the contents of the media stream report against the signaling session report. If a media stream in the result set starts before or ends after a signaling session, it was added to the result set due to a correlation key collision. Due to time information it cannot match the signaling session and is deleted from the result set by the correlation unit. If the media stream result set still contains multiple entries after the previous correlation key collision check, all remaining media streams (i.e. their media stream reports) belong to the media service (e.g. the VoIP session).

If a search returned no usable entry, i.e. no signaling session reports and media stream reports comprising matching correlation keys have been found, the correlation unit may optionally use additional correlation keys contained in either the signaling session reports and/or media stream reports to run another search, as explained with respect to Fig. 13 and Fig. 14 above.

The correlation process may optionally select only the most significant media stream for each direction for inclusion to the service report on the media service. The parameters to define the significance of the media streams may be for example configurable by the user/administrator. One parameter for deciding the significance of a media stream in the result set could be the duration of a media stream. For example, the service report generated for a given media service by the correlation unit could include the media stream with the longest duration, in case multiple media stream reports match a correlation key of a signaling session report. Another option would be to include in the service report the last media stream that has been alive in the media service, in case multiple media stream reports match a correlation key of a signaling session report. The correlation process may optionally add information from multiple media streams to the service report, e.g. when the service consisted of an audio and video part.

For a given signaling session (i.e. service), the correlation process may run at least twice, one time for each correlation key in a signaling session report. For example, in case the media service is a VoIP service, the correlation process may be for example performed by the correlation unit once for the caller's media stream and once for the callee's media stream, both identified by their respective correlation keys in the signaling session report.

The correlation unit combines the information contained in the media stream reports that are correlated to a respective signaling session report and the information contained in the signaling session report to provide a service report, providing an outline of the service. The content of this combined service report may be for example the sum or a subset of information of the signaling and media stream reports of a given service. Furthermore, the content of the service reports for the respective services may be configurable and/or depending on the application.

In another exemplary embodiment of the invention, the correlation unit may perform the correlation based on the media stream reports as starting points which are matched to their corresponding signaling session report. Ideally, each media stream of a media stream report should match to one signaling session report based on its correlation key. This implementation may optimize the search time in the correlation processing. The above correlation mechanism starting from the signaling session reports searches the rather large media stream report database at least twice for each signaling session. If the correlation process is based on the correlation keys within the media stream reports entries as starting point, the correlation would periodically take all new media stream reports, obtain the correlation keys thereof and search the signaling session reports for matching signaling sessions. At this point the same steps are taken as outlined above.

Another option may be to correlate media stream reports and signaling session reports on demand. An on-demand correlation employs the same mechanism as the proactive correlation mechanism described above, but instead of running periodically, e.g. when introducing new signaling session reports or media stream reports into the database(s), the on-demand correlation process runs when a user requests particular data.

In an exemplary implementation, the database may contain independent sets of signaling session reports and media stream reports. An operator could request service reports of services linked to some given search criteria, e.g. a specific user, a specific source and/destination IP address, a certain service type, services using a particular software application, codec, etc. The user could be for example specified by providing a username, phone number, the user's SIP URI, etc. or another identifier of the user that can be found in a signaling session report. The correlation unit would then search for the signaling session reports for all signaling sessions matching the particular search criteria. The search criteria may be for example provided to the correlation unit through a web application running on a browser.

For each signaling session report matching the search criteria, the correlation unit subsequently searches for the matching media stream reports as described above for the proactive correlation process. The input for this on-demand correlation is a set of signaling session reports matching the search criteria and not all signaling session reports as in the proactive correlation procedure described above. The on-demand correlation process generates for each signaling session report matching the search criteria a service report, for example by enriching the information of the signaling session report with the matching media stream report information and returns the service report for each matching signaling session report. The returned data may be for example viewed by the web application as individual service reports. Optionally the on-demand correlation process may also store the generated service reports in a database. This would save the on-demand correlation processing time in case an operator requests the same information multiple times.

Correlation Unit

Next, the correlation unit 402 as shown in Fig. 4 is discussed in further detail with respect to Fig. 5 and Fig. 6. Fig. 6 shows the structure of a correlation unit 402 according to an exemplary embodiment of the invention.

The correlation unit 402 may be reachable via the packet-switched network, e.g. through a web based application. For this purpose, the correlation unit 402 may comprise one or more network interface cards 601. The correlation unit 402 has a receiving unit 602 and a

transmitter unit 603, for example for receiving correlation related messages, such as for example on-demand correlation requests, and for responding thereto. Furthermore, the receiving unit 602 and a transmitter unit 603 may transmit and receive messages for accessing remote database(s) that contain(s) the signaling session reports and the media stream reports. The receiving unit 602 and transmitter unit 603 may receive and transmit the correlation related messages from its network interface card(s) 601.

The correlation unit 402 may also have a processing unit 604. The processing unit 604 may be used to determine correlation respective signaling sessions to their corresponding media streams. Furthermore, the processing unit 604 may be used to generate services reports. The service reports may for example comprise information on the respective services and the quality information of the media streams of the respective services. This information may be obtained from the signaling session report and media stream reports belonging to a given service. For correlation, it may be necessary to query one or more databases that comprise the signaling session reports and the media stream reports as provided by one or more signaling plane probes 401 and media plane probes 403. The correlation unit 402 may therefore comprise a database access unit 605 which provides functionality to access these one or more databases. In one exemplary implementation, the database(s) may be located remote from the correlation unit 403.

The correlation unit 402 may also have a storage unit 606, e.g. a RAM, for (temporarily) storing information and data related to the correlation processing. For example, the result set of media stream reports for correlation process, the data of the signaling session reports, and the service report generated for a respective service may be stored in the storage unit 606.

Fig. 5 shows a flow chart of the operation of the correlation unit 402 according to an exemplary embodiment of the invention. Exemplarily, the correlation processing for determining the media streams belonging to one given signaling session report is illustrated in Fig. 5. In case a correlation should be performed for more than one signaling session (report), the procedure of Fig. 5 is repeated for each signaling session report. For example, when using a proactive correlation mechanism, the correlation unit 402 could obtain all signaling session reports that have been generated since the previous correlation process from a signaling session database and process each of the signaling session reports according to Fig. 5. In another embodiment related to an on-demand correlation process, the correlation unit 402 could also receive a correlation request that is indicating one or more search criteria, and may first – not shown in Fig. 5 – query a signaling session database for all signaling session reports that match the search criteria. The signaling session reports matching the search one or more criteria may then be each processed according to Fig. 5 and the service report(s) are returned to the querying party.

The correlation unit 402 obtains 501 a signaling session report to be processed and generates 502 a service report for the service corresponding to the given signaling session. This service report may include information contained in the signaling session report and will be further supplemented with additional information as described below. Furthermore, the correlation unit 402 obtains 503 a correlation key contained in the signaling session report – as noted previously there should be at least two correlation keys contained in the signaling session report. Using the obtained correlation key, the correlation unit 402 queries 504 a media report database to receive all media stream reports that match the correlation key i.e. that contain the same correlation key.

As noted above, it may happen that the query results in the database returning more than one media stream report for the given correlation key, for example due to a hash collision of the correlation key. Therefore, the correlation unit 402 checks 505 whether the returned result set of the database query comprises more than one media stream report. If not (no), the correlation unit 402 can immediately enrich 508 the service report with information on the media stream as contained in the one media stream report returned from the media stream database.

In case the result set returned by the database comprises multiple media stream reports, the correlation unit 402 may try to detect 506 and resolve a hash collision. For example, the hash collision could be resolved by checking further information in the media stream reports of the result set for parameters that identify the media stream (report) that belongs to the correlation key that has been queried. For example, in case the media stream reports contain timestamps of the beginning and/or end of the media stream, same could be compared to respective timestamps contained in the signaling session report, which yield the beginning and/or end of the signaling session. The media stream report(s) of media stream(s) that start before or end after the signaling session may be obtained from the database due to a hash collision. Accordingly, the correlation unit 402 may select 507 only the most relevant media stream report, i.e. the media stream report that is most likely corresponding to the service of the signaling session report from the result set and may add 508 at least a part of its information to the service report. Further, the correlation unit checks 509 whether there is a further correlation key contained in the signaling session report and, if so (yes), repeats steps 503 to 508. These steps are repeated for each further correlation key(s) comprised in the signaling session report.

Once all correlation keys contained in the signaling session report have been processed, the service report for the signaling session is complete, and may be optionally stored 510 by the correlation unit 402 in a service report database.

Depending on the correlation type or application, the procedure of Fig. 5 could be supplemented with additional steps. For example, in case of a proactive/periodic correlation, the correlation unit 402 may check whether there is another signaling session report to be matched to its corresponding media stream reports, and if so the procedure of Fig. 5 would be executed for each signaling session report to be processed. In case of performing a correlation on demand, the correlation unit 402 may likewise check whether there is another signaling session report to be matched to its corresponding media stream reports, and if so the procedure of Fig. 5 would be executed for each signaling session report to be processed. Once all signaling session reports have been processed the correlation unit 402 may generate a correlation response message comprising the service reports generated by the correlation unit 402 in response to the search criteria provided in a correlation request, and returns the correlation response message to the querying entity, e.g. another device/computer.

Signaling Plane Probe

Fig. 9 shows an exemplary implementation of a signaling plane probe 401 and Figs. 7 and 8 show two exemplary flow charts of the operation of same according to different embodiments of the invention. In general, the signaling plane probe 401 observes one or more network links to obtain a copy of all packets passing these network links/interfaces. For this purpose, the signaling plane probe 401 comprises a monitoring unit 901 that receives 701 packets of a service on the packet-switched network 400. For example, this monitoring unit 901 may receive the packets from a one or more specialized packet capture cards 906 (network interface card - NIC) which is capable of receiving and processing packets on the physical interface and to provide them to the application layer, nearly without requiring CPU processing time and operation system functionality. An optional filter in each packet capture card of the signaling plane probe 401 may be used to prevent unwanted packets from being copied from the network interfaces to the processing processes. These filters may be for example implemented as stateless or stateful Internet Protocol (IP) or payload filters. For example, this filter may check 702, whether the currently processed packet is a signaling packet or not. If the packet does not belong to a signaling session, i.e. is not packet of a signaling related protocol, the next packet may be processed.

Each signaling packet received at the signaling plane probe 401 will be processed by a detection unit 902 which analyzes the contents of the packets provided by the monitoring unit 901 so as to identify and process packets as signaling packets. Specifically, the detection unit 902 detects signaling packets and matches these to signaling sessions of the media service, and is capable of identifying one or more media data paths of a media plane of the

service based on signaling information of the service comprised in the received packets. Optionally, the detection unit 902 may be provided in one or more specialized packet capture cards.

For example, the detection unit 902 identifies packets containing signaling information (signaling packets) in a parsing process based on their header information. If parsed successfully as a signaling packet, the signaling packet is forwarded to a state engine. The state engine is designed to identify 703 and form signaling sessions from individual signaling packets.

The state engine (which may be considered an internal database with correlation related information maintained in the storage unit 903) of the detection unit 902 has a processing facility and a state table which holds information about signaling sessions (signaling session records). Based on the protocol-specific identifiers in the signaling packets the detection unit 902 can identify the start and end of a signaling session, respectively media service, or "append" a respective processed signaling packet to an existing signaling session. Signaling packets may also contain information on the start and termination of media streams in the media plane of the media service, which are events that may be tracked by the detection unit 902. Depending on whether the identified signaling session is already registered in the signaling plane probe 401 (step 704), the detection unit 902 registers 705 the new signaling session in the state table, including the generation of a new signaling session record for the signaling session.

The signaling session record may be supplemented with information on the signaling session while same is running to account – for example – for endpoint address/name (e.g. SIP URI); the phone number of the user; the user-agent (endpoint software and version); capabilities/features; supported media types; IP addresses and ports for signaling and media streams (from SDP body); signaling routing information; timestamps: e.g. session start, Post-Dial-Delay (PDD), connect time, session end; information on the session flow (e.g. signaling message(s) with timestamps); session state (e.g. successful, unsuccessful, timed out, reason for failure), and/or events during session (media on hold, codec renegotiation, session forking, authentication), etc.

Furthermore, by analyzing 706 the contents of the signaling packets, e.g. a SDP body of a SIP messages comprised in signaling packets, the detection unit 902 may further learn one or more media data paths of the media plane of the service belonging to the media service that is associated to a given signaling session. The signaling plane probe 401 may further comprise a storage unit 903, e.g. a volatile memory, to store the information describing the state of the respective signaling sessions in the state table. In case the signaling packet

indicates the start of a new media stream (step 706: yes), the detection unit 902 may add 707 information on the media data path (e.g. IP address and port number of the destination, and optionally the source of the media stream) of this media stream to the signaling session record. Likewise, in case the signaling packet indicates the end/termination of a new media stream (step 708: yes), the detection unit 902 may remove the information on the respective media stream from the signaling session record maintained in storage unit 903. In case the signaling packet indicates the termination of the signaling session (step 710: yes), the detection unit 902 may either itself or trigger the processing unit 904 to generate 711 a respective correlation key for each media data path indicated in the signaling session record. The correlation key for a respective media data path of a media stream may be generated for example based on information on the directivity of the media stream (e.g. as given by the tuple of destination IP address and UDP port number) and a timestamp indicating the point in time at which the signaling session was terminated. For example, this timestamp may indicate the point in time of reception of the signaling packet indicating the event of the termination of the signaling session at the signaling plane probe 401. For example, the network interface card(s) 906 of the signaling plane probe could generate a timestamp indicating the time of reception of the respective signaling plane packet and could make same available to the other components of the signaling plane probe 401. Furthermore, the detection unit 902 or the processing unit 904 may further generate 912 a signaling session report based on the information on the signaling session gathered in the signaling session record. Further, the signaling session report also included the correlation keys for the media data paths of the service that exist at the end of the service according to the signaling plane information. The signaling session report may then be stored in a signaling session report database, for example using the database access unit 907 of the signaling plane probe 401. Moreover, the signaling session is deregistered 713 from the state table.

Once a given signaling packet has been processed, it may be discarded. It should be also noted that the order of the checks in steps 706, 708 and 710 and the respective actions to be taken is only exemplary. It is only important that the information on the media streams of the service that exist in the media plane are tracked by the signaling plane probe 401, and that the signaling plane probe 401 is able to detect the termination of the signaling session.

Optionally (not shown in Fig. 7), in another exemplary embodiment, the detection unit 902 of the signaling plane probe 401 periodically (e.g. every 5, 10 or 15 seconds) traverses all entries in the state table in order to find and clean up stale signaling sessions. For this purpose, the signaling plane probe 401 may add a timestamp to each signaling session record in the state table that is updated each time a new signaling packet/message of the respective signaling session is received. Based on this timestamp allows to identify stale

signaling session by simply checking based on the signaling sessions' timestamps whether the last signaling packet/message for a given signaling session has been received longer than a threshold time period ago. For example, stale sessions might be those which did not receive new signaling packets/messages for 86400 seconds or signaling sessions may be defined stale and are terminated after a certain time. This clean-up removes unnecessary information from the state.

Fig. 8 shows another flow chart of the operation of a signaling plane probe 401 according to another embodiment of the invention. In the example discussed with respect to Fig. 7, the signaling plane probe 401 keeps track of the media streams currently existing in the media plane of the service, eliminating those media streams from the signaling session record that are terminated during the ongoing signaling session/service. In the flow chart of Fig. 8, which is corresponding to Fig. 7 to a large extent, the termination of one or more media streams in the media plane, as detected in step 708, causes the detection unit 902 or the processing unit 904 (triggered by the detection unit 902) to generate 801 a correlation key for each terminated media stream. Similar to the other correlation keys, the correlation key for a terminated media stream may be generated based on its media data path information (e.g. destination IP address and UDP port number) and a timestamp that indicates the point in time the media stream was terminated. This timestamp may indicate the point in time of reception of the signaling packet indicating the event of the termination of the media stream at the signaling plane probe 401. The correlation key generated is then added 802 to the signaling session record of the signaling session by the detection unit 902 or the processing unit 904 and is later then also stored in the signaling session report in step 712 together with the other correlation keys either generated in step 801 or step 711.

Media Plane Probe

Next, the media plane probe 403 according to an exemplary embodiment of the invention is described with respect to Fig. 10 and Fig. 11. Fig 10 shows an exemplary embodiment of a media plane probe 403, while Fig. 11 shows a flow chart of its exemplary operation according to an embodiment of the invention.

Similar to the signaling plane probe 401, the media plane probe 403 observes one or more network links to obtain a copy of all packets passing these interfaces. The media plane probe 403 may observe one or more network links to obtain a copy of all packets passing these interfaces. The media plane probe 403 comprises a monitoring unit 1001 that receives 1101 packets of a service on the packet-switched network 401. For example, this monitoring unit 1001 received the packets from one or more specialized packet capture cards 1006 (network interface card - NIC) which are capable of receiving and processing packets on the physical

interface and to provide them to the application layer, nearly without requiring CPU processing time and operation system functionality. An optional filter in the packet capture card of the media plane probe 403 may be used to prevent unwanted packets from being copied from the network interfaces to the processing processes. This filter may be for example implemented as stateless or stateful Internet Protocol (IP) or payload filter. For example, this filter may check 1102, whether the currently processed packet is a media plane packet (media packet) or is some other packet. If the packet does not belong to the media plane, i.e. is not packet of a media related protocol, such as for example RTP, the next packet may be processed.

Each media packet received will be processed by the media plane probe 403. The detection unit 1002 of the media plane probe 403 is designed to detect and form media streams from single media packets. The media plane probe 403 may further comprise a processing unit 1002 that may hold a state table which stores information about the media streams monitored by the media plane probe 403 in a local storage unit 1003, e.g. a volatile memory. Based on the protocol-specific identifiers in the media streams the media plane probe 403 (e.g. based on header information in the RTP packets and/or lower layer headers) the detection unit 1002 identifies 1103 the media data path, i.e. the media stream, the media packet belongs to. If the media data path has not been yet registered (step 1104: no), i.e. the processed media packet starts a new media stream (with the given media data path) the detection unit 1002 generates a new media data record in the state table to account for the new media data path observed by the media plane probe 403. This media stream record may contain information on the media stream such as for example the destination (and optionally the source) IP address and port combination of the media stream, the SSRC in the packet header (in case RTP is used for media transport), etc. If the media data path is already known (step 1104: yes), the media stream record already exists. In any case, the detection unit 1002 updates a timestamp 1106 in the media stream record to indicate the point in time at which the last media packet of the media stream has been received. This timestamp may be used for calculating the correlation key for the media stream. Furthermore, the detection unit 1002 may (re)start a timer 1108 for the media stream record, which when timing out is indicating the termination of the media stream as will be outlined below. Alternatively, no timer is used, but the media plane probe 403 periodically checks for media streams the last media packet of which was received more than a threshold time period ago, in order to determine terminated media streams.

The media stream records of the media streams may for example indicate the respective media data paths thereof by an IP address and port tuple of the destination device, i.e. the receiving device of the media stream. The media plane probe's detection unit 1002 stores

information describing the state of the detected media streams in the media stream records maintained in the state table. Optionally, the individual media streams (i.e. their media data paths) may be identified by a media path identifier (e.g. stream hash key) as unique key. Further optionally, upon detecting a new media stream, respectively media path (step 1104: no), the processing unit 1004 could generate 1105 a media path identifier (e.g. a media path identifier (e.g. stream hash key)) for the media stream, respectively media path and stores 1106 the media stream record in the state table as outlined above.

Optionally, the processing unit 1004 could analyze every received packet for quality information 1107 and – further optionally – store them in the storage unit 1003 for a given time period. Periodically the processing unit 1004 may aggregate or compress the quality information to form fine-grained media stream quality reports with media stream related information. Fine-grained media stream quality reports could be for example written in predetermined time intervals, e.g. every five seconds, and could provide a summary of the media stream in the respective interval and the aggregated quality information of the media stream in the respective interval. Upon generation of a respective fine-grained media stream quality report, the packets used for the generation of this fine-grained media stream quality report can be removed from the storage unit 1003.

The processing unit 1004 of the media plane probe may for example periodically check 1109 (e.g. every 5, 10 or 15 seconds); whether a timer set of a media stream (record) in step 1108 has expired. Alternatively, as noted above, the media plane probe 403 may also check timestamps indicating the last update (i.e. last reception of a media packet) of the respective media stream reports in the state table. A media stream may be for example considered terminated, if no new media packets have been received by the monitoring unit 1001 for the existing media stream for a given amount of time (e.g. 60 seconds). Accordingly, if using a timer to monitor the termination, the timer value could be set to this given amount of time. Further alternatively, instead of periodically checking the media stream records or expiry of their timers, the expiry of a timer for a media stream (record) could cause an interrupt invoking the processing unit 1004 to perform steps 1110 to 1112 as outlined below.

Upon detecting the termination of a media stream, the processing unit 1004 processes the respective information in the media stream record of this media stream the state table. The processing unit 1004 generates 1110 a correlation key for the media stream based on its media data path information (e.g. the destination IP address and UDP port number) and the timestamp in the media stream record that is indicating the point in time of receiving the last media packet of the media stream (see step 1106). Furthermore, the processing unit 1004 generates 1111 a media stream report for the media data path, respectively media stream, which includes the generated correlation key and further information on the media stream.

This further information may be for example transport media information and optionally stream quality information. For example, in case RTP is used for media transport, the transport media information could include the source and destination IP address of the RTP stream, the source and destination UDP port of the RTP stream, the RTP header Synchronization source (SSRC) value and the RTP payload type. Furthermore, the processing unit 1004 may also generate information on the media stream quality that is included to the media stream report as stream quality information. This stream quality information could for example include timestamp information (e.g. start and/or end time of the media stream), an indication of RTP format violations (e.g. the number thereof), information on the protocol conformance (e.g. sequence errors and/or DSCP errors), and transport information (e.g. packet interval and jitter statistics, information on packet loss, information on duplicate packets, etc.), MOS values (e.g. per time-slice (e.g. every 5 seconds), minimum/average/maximum value for entire stream duration, etc.), and/or events (e.g. comfort noise, silence suppression, DTMF tones, codec changes, marker bits). The contents of the media stream reports may be configurable by an operator and may be adapted to customers' needs.

The media stream record may be further stored in a media stream report database using the database access unit 1007 of the media plane probe 403. Finally, the processing unit 1004 of the media plane probe 403 deregisters 1112 the media data path from the state table, e.g. by deleting the corresponding media stream record.

Steps 1110 to 1112 are performed for every media stream that is considered terminated by the media plane probe.

Optimizations may be used to speed up processing of entries in the state tables inside the signaling plane probe 401, the correlation unit 402 and the media plane probe 403. A simple list would be sufficient to implement the features as outlined above. Add, modify, delete and search operations on simple lists may however be too slow. These operations may be implemented utilizing more complex but faster data structures, e.g. hash tables or data structures heavily dependent on CPU and Operating System specific behavior and functionality. The Intel Thread Building Block (TBB) library for instance can be used to implement in-memory databases used for the state tables in the signaling plane probe 401, the correlation unit 402 and the media plane probe 403. The Intel Thread Building Block (TBB) library may also be used to aid implementation of the processing queues inside and in between the components outlined above.

Please note that in case the signaling plane probe 401 and the media plane probe 403 should perform passive non-intrusive monitoring of the signaling plane traffic, respectively

media plane traffic, a separate network interface card is employed for the transmission of correlation related messages (i.e. the network interface card(s) for monitoring the traffic are only employed for packet reception but not for packet transmission).

Please also note that in the procedures described above the detection unit 902, 1002 and the processing unit 904, 1004 may use the database access unit 907, 1007 to access the database/state table in the storage unit 903, 1003.

Application

The correlation mechanism for signaling plane and media plane of a service may be used in manifold applications. For example, the correlation mechanism is useful in small networks but is more important in larger scale networks where the various signaling and media components are decomposed and present on different, sometimes geographically split, network entities. This is the case in IP Multimedia Subsystem (IMS) or Next Generation Network (NGN) setups.

Fig. 12 shows a simplified network in which locations are indicated where correlation may be desired and can be implemented using the correlation mechanism according to one of the various embodiments described herein. Possible monitoring locations are for example (and not limited thereto):

- The Access Session Border Controller (SBC) network interfaces with end customers (shown of the left side of Fig. 12),
- a decomposed gateways system, e.g. to the PSTN, is shown as Media Gateway (MGW) and a Media Gateway Controller (MGC) – see In the top right section of Fig. 12, and
- an Interconnection SBC interfaces to one or more interconnection or transit partners (see the lower right side of Fig. 12.

It should be noted that the location of the monitoring probes for signaling plane and media plane may also be subject to requirements on the monitoring solution.

In some cases, signaling plane and media plane may originate and/or terminate at the same endpoint (device). In case of an end customer this endpoint most likely uses the same IP address on the signaling and media plane. These streams may be for example observed, if the signaling plane probes 401 and the media plane probes 403 are placed close to the end customer devices. The Access Session Border Controller may be considered to be close to the carrier internal network's "border" towards end customer devices and to therefore receive

traffic from the end customers on one side and transmits this traffic on the other side towards the carrier internal network. Most commonly these Session Border Controllers (SBC) are used to separate the devices outside the own trust domain from the inside network. Although the traffic is in one geographic location the SBC may use different IP addresses to load balance traffic. In this case a simple IP address based correlation mechanism not using the principles and ideas outlined herein would already fail whereas the correlation mechanism described herein would succeed. The signaling plane probes 401 and media plane probes 403 may be for example placed in a way that they observe the signaling streams and media streams respectively.

In the upper right corner of the diagram shown in Fig. 12, a decomposed gateway is shown which uses different physical entities as media and signaling components (MGC and MGW, respectively). In between those devices a third party protocol may be used to synchronize the devices. For the sake of a minimum size network diagram only one of each component types is shown but in real life scenarios multiple media components are controlled by one signaling component (as it is the case in IMS, NGN). The media components may be co-located in the same location with the signaling component or spread across the internal network of the carrier. In either case the correlation task is rather complex as it involves multiple IP addresses on different physical network links and thus may require multiple signaling plane probes 401 as well as media plane probes 403.

In a third exemplary scenario the correlation has to correlate signaling and media streams from different locations. A media plane probe 403 component on the internal side of the interconnection point (lower right corner of the diagram – Interconnection SBC) observes a media stream from the Access SBC (middle left). The corresponding signaling session is observed next to the aforementioned Access SBC by a separate signaling plane probe 401.

The location of the correlation unit 402 does not impact what signaling sessions or media streams can be correlated. The correlation unit 402 should have access to the signaling session reports and media stream reports of the signaling plane probes 401 and media plane probes 403 and can be implemented as distinct physical machine or co-located on an existing signaling plane probe 401 or media plane probe 403.

The terms probe and correlation unit used herein can be considered to refer to dedicated hardware, a software program or a combination thereof to implement the desired functionality of the respective probe, respectively, correlation unit. Accordingly, another embodiment of the invention relates to the implementation of the above described various embodiments using hardware and software. It is recognized that probes and correlation unit of the invention may be implemented by having computing devices (processors) execute a software

program that causes the desired functionality being executed. A computing device or processor may for example be general purpose processors, digital signal processors (DSP), application specific integrated circuits (ASIC), field programmable gate arrays (FPGA) or other programmable logic devices, etc. Some of the computing devices may store the software program internally.

The various embodiments of the invention may also be performed or embodied by a combination of computing devices and software programs providing the desired functionality stored on any kind of computer readable storage media, for example RAM, EPROM, EEPROM, flash memory, registers, hard disks, CD-ROM, DVD, etc.

It should be further noted that the individual features of the different embodiments of the invention may individually or in arbitrary combination be subject matter to another invention.

It would be appreciated by a person skilled in the art that numerous variations and/or modifications may be made to the present invention as shown in the specific embodiments without departing from the spirit or scope of the invention as broadly described. The present embodiments are, therefore, to be considered in all respects to be illustrative and not restrictive.

CLAIMS

1. A method for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network, the method comprising:

a) generating at least one correlation key for each direction of a service based on the information of the signaling session of the respective service monitored at the signaling plane probe, wherein the correlation keys for a respective service are generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and using information on the direction of the respective service,

b) storing, by the signaling plane probe, a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service,

c) generating a correlation key for each media stream monitored by the media plane probe, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,

d) storing, by the media plane probe, a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media, and

e) correlating the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports.

2. The method according to claim 1, wherein the generation of the respective correlation keys for the respective media streams by the media plane probe is based on the media data path information of the respective media stream and the timestamp indicative of the point in time at which the respective media stream ended.

3. The method according to claim 1 or 2, wherein information on the direction of the respective service, based on which the correlation keys are generated by

signaling plane probe for a respective service, is information on one or more media data paths of media streams of the service learned from messages of the signaling session of the service.

4. The method according to one of claims 1 to 3, further comprising the steps of:

f) monitoring, by the signaling plane probe, signaling messages of the respective signaling sessions of the services, and

g) detecting, by the signaling plane probe, media streams belonging to each service associated to a respective monitored signaling session based on the signaling messages of the respective signaling sessions,

wherein the correlation keys for a respective one of the signaling sessions are generated based on information on the media data path of the media streams belonging to the service associated to the respective one of the signaling sessions and the timestamp indicative of the point in time at which the end of the respective signaling session has been detected by the signaling plane probe.

5. The method according to one of claims 1 to 4, wherein information on a media data path of a media stream is a destination IP address and port number of the media stream.

6. The method according to one of claims 1 to 5, wherein the timestamps used for the generation of the correlation keys by the signaling plane probe and the media plane probe are rounded before the generation of the correlation keys or have a granularity in the range of one to ten seconds.

7. The method according to one of claims 1 to 6, wherein the signaling plane probe generates for each direction of the service at least one correlation key using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and being rounded to the next lower timestamp value according to a rounding interval, and at least another correlation key using the timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and being rounded to the next higher timestamp value according to the rounding interval.

8. The method according to one of claims 1 to 7, wherein the media plane probe generates for each media stream at least one correlation key using a timestamp indicative of the point in time at which the respective media ended and being rounded to the next lower timestamp value according to a rounding interval, and at least another correlation key using the timestamp indicative of the point in time at which the respective media stream ended and being rounded to the next higher timestamp value according to the rounding interval.

9. The method according to one of claims 1 to 8, wherein the correlation of the media streams and the signaling session of the respective services is performed by a correlation unit periodically or in response to a query.

10. A method for correlating media streams and signaling sessions of services, the method comprising the steps of:

a) correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services, wherein correlating the media streams and the signaling sessions comprises identifying one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys contained in the media stream reports and signaling session reports, and

b) generating a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained the signaling session report of said at least one service.

11. The method according to claim 10, wherein correlation keys provided in a respective signaling session report have been generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the corresponding service has been detected and using information on the media streams of the service.

12. The method according to claim 10 or 11, wherein correlation key provided in a respective media stream report has been generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the media stream.

13. The method according to one of claims 10 to 12, wherein correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions comprises:

matching the respective correlation keys of the media stream reports of media streams are to the respective correlation keys of the signaling session reports, or vice versa,

wherein a match of the correlation keys of a one of the media stream reports and one of the signaling session reports yields the media stream of said one media stream report belonging to the same service to which the signaling session of said one signaling session report belongs.

14. The method according to one of claims 10 to 13, wherein the service report comprises information on the service and the quality information of the media streams of the service.

15. The method according to one of claims 10 to 13, wherein correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions comprises:

obtaining from a signaling session report correlation keys identifying the media streams belonging to the same service as the signaling session, and

identifying those media stream reports which include one of the correlation keys obtained from the signaling session report.

16. The method according to one of claims 10 to 13, wherein correlating the media streams and the signaling sessions of one or more services based on the correlation keys comprised within media stream reports of the media streams and signaling session reports of the signaling sessions comprises:

obtaining from a media stream report a correlation key identifying the media stream, and

identifying a signaling session report among the signaling session reports which includes the same correlation key obtained from the signaling session report.

17. The method according to one of claims 10 to 16, wherein steps a) and b) are performed by a correlation unit periodically for all signaling session reports and media stream reports generated within a configurable time interval, or in response to a request.

18. The method according to one of claims 10 to 17, further comprising steps a) to d) of claim 1 or the steps of the method according to one of claims 2 to 9.

19. A passive monitoring system for correlating media streams and signaling sessions of services in a packet-switched network, the passive monitoring system comprising:

a signaling plane probe for generating at least one correlation key for each direction of a service based on the information of the signaling session of the respective service monitored at the signaling plane probe, wherein the correlation keys for a respective service are generated using a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the signaling plane probe and using information on the direction of the respective service,

wherein the signaling plane probe is further adapted to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service,

a media plane probe for generating a correlation key for each media stream monitored by the media plane probe, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,

wherein the media plane probe is adapted to store a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media, and

a correlation unit for correlating the media streams and the signaling sessions based on the correlation keys comprised within media stream reports and signaling session reports.

20. A signaling plane probe comprising:

a network interface card for receiving packets of signaling sessions of multiple media services,

monitoring unit for monitoring the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of the signaling sessions, wherein the monitoring unit is further adapted to detect the point in time at which a respective signaling session of the respective service terminates, and

a processing unit for generating at least one correlation key for each direction of a service based on a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected by the monitoring unit and based on information on the direction of the respective service,

wherein the monitoring unit is adapted to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises at least two correlation keys generated for the respective directions of the service.

21. A media plane probe comprising:

a network interface card for receiving packets of media streams of multiple media services,

monitoring unit for monitoring the media streams based on the packets of the media streams received by the media plane probe to thereby obtain information on a media data path a respective one of the media streams, wherein the monitoring unit is further adapted to detect the point in time at which a respective media stream of the respective service terminates, and

a processing unit for generating a correlation key for each media stream monitored by the monitoring unit, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,

wherein the monitoring unit is adapted to store a media stream report for each media stream in a database, wherein each media stream report comprises at least the correlation key generated for the respective media.

22. A correlation unit for correlating media streams and signaling sessions of services in a passive monitoring system of a packet-switched network, the correlation unit comprising:

a database access unit for accessing a database containing signaling session reports for signaling sessions of services and a database comprising media stream reports of media streams of the services, wherein the signaling session reports each contain at least one correlation key for each direction of a respective service, and the media stream reports each contain one correlation key indication the direction of the media stream,

a processing unit for correlating the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services, wherein the processing unit is adapted to identify one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys contained in the media stream reports and signaling session reports, and

a service record generation unit for generating a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained the signaling session report of said at least one service.

23. The correlation unit according to claim 22, further comprising means to perform the steps of the method according to one of claims 10 to 17.

24. A computer-readable medium storing instructions that, when executed by a processor of a signaling plane probe, cause the signaling plane probe:

to receive packets of signaling sessions of multiple media services,

to monitor the signaling sessions based on the packets of the signaling sessions received by the signaling plane probe to thereby obtain information on a respective one of the signaling sessions,

to detect the point in time at which a respective signaling session of the respective service terminates,

to generate at least one correlation key for each direction of a service based a timestamp indicative of the point in time at which the end of the respective signaling session of the respective service has been detected and based on information on the direction of the respective service, and

to store a signaling session report for each signaling session in a database, wherein each signaling session report comprises the at least two correlation keys generated for the respective directions of the service.

25. A computer-readable medium storing instructions that, when executed by a processor of a media plane probe, cause the media plane probe:

to receive packets of media streams of multiple media services,

to monitor the media streams based on the packets of the media streams received by the signaling plane probe to thereby obtain information on a media data path a respective one of the media streams,

to detect the point in time at which a respective media stream of the respective service terminates, and

to generate a correlation key for each media stream monitored by the monitoring unit, wherein the correlation key is generated using a timestamp indicative of the point in time at which the respective media stream has ended and using information on the media data path of the respective media stream,

to store a media stream report for each media stream in a database, wherein each media stream report contains at least the correlation key generated for the respective media.

26. A computer-readable medium storing instructions that, when executed by a correlation unit, cause the correlation unit to correlate media streams and signaling sessions of services in a passive monitoring system of a packet-switched network, by:

accessing a database containing signaling session reports for signaling sessions of services and a database containing media stream reports of media streams of the services, wherein the signaling session reports each contain at least one correlation key for each direction of a respective service, and the media stream reports each contain one correlation key indication the direction of the media stream,

correlating the media streams and the signaling sessions of the services based on the correlation keys comprised within media stream reports and signaling session reports to thereby identify the media stream reports and signaling session report belonging to a respective one of said one or more services,

identifying one or more media streams of a respective service and the corresponding signaling session of the respective service based on the correlation keys comprised in the media stream reports and signaling session reports, and

generating a service record for a respective one of said at least one service based on information contained in the media stream reports and information contained the signaling session report of said at least one service.

27. The computer-readable medium according to claim 26, further storing instruction that, when executed by the processor of the correlation unit, cause the correlation unit to perform the steps of the method according to one of claims 10 to 17.

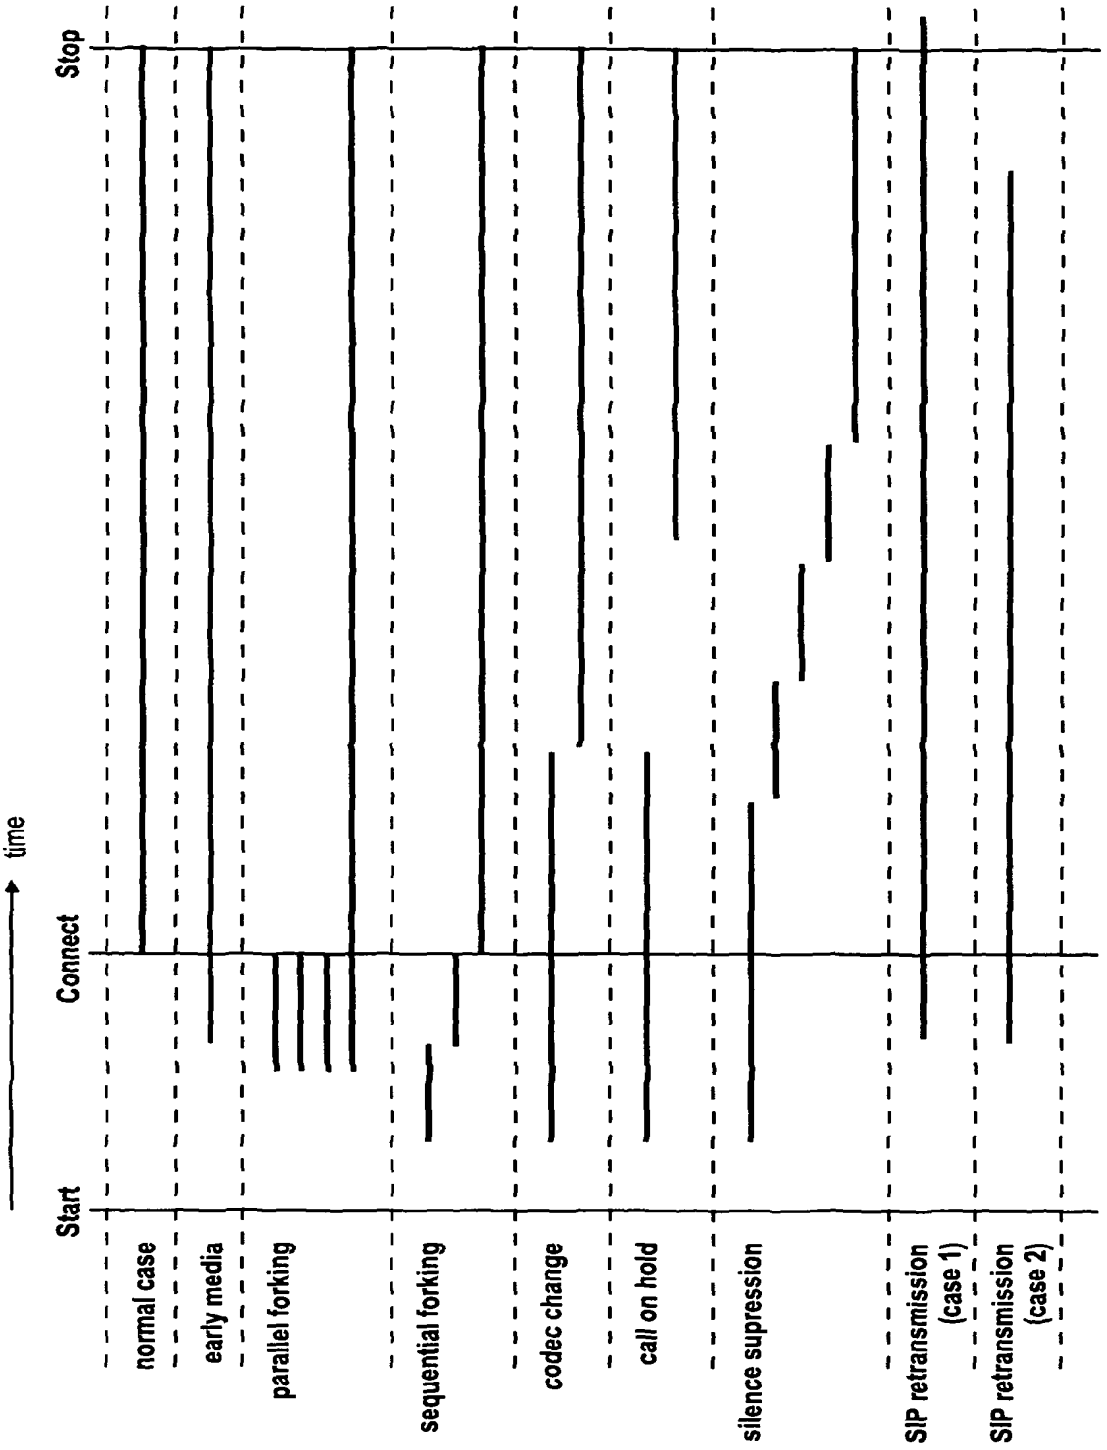
time
Start
Connect
Stop
normal case
early media
parallel forking
sequential forking
codec change
call on hold
silence supression
SIP retransmission (case 1)
SIP retransmission (case 2)

Fig. 1

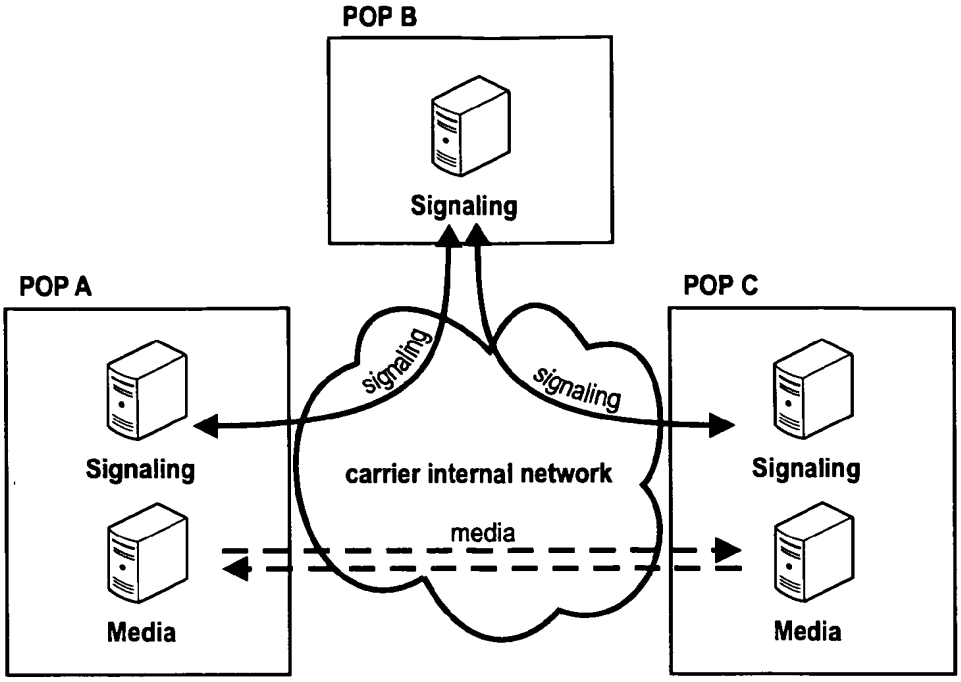
POP B
Signaling
POP A
POP C
signaling
signaling
Signaling
carrier internal network
Signaling
media
Media
Media

Fig. 2

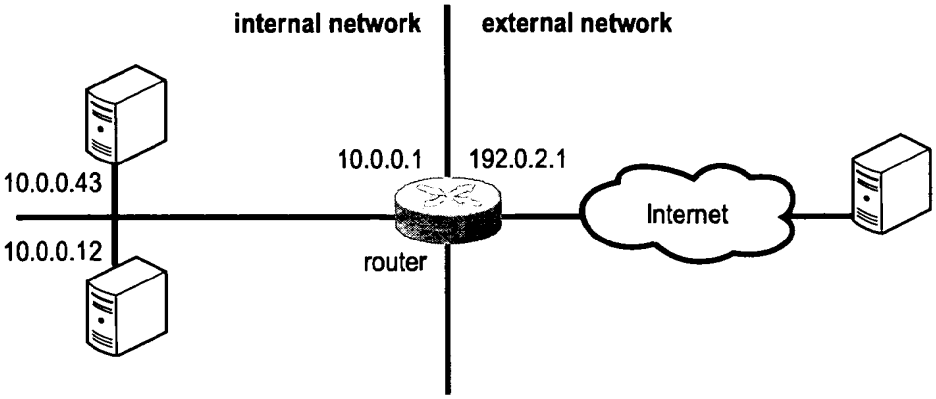
internal network
external network
10.0.0.1
192.0.2.1
10.0.0.43
10.0.0.12
router
Internet

Fig. 3

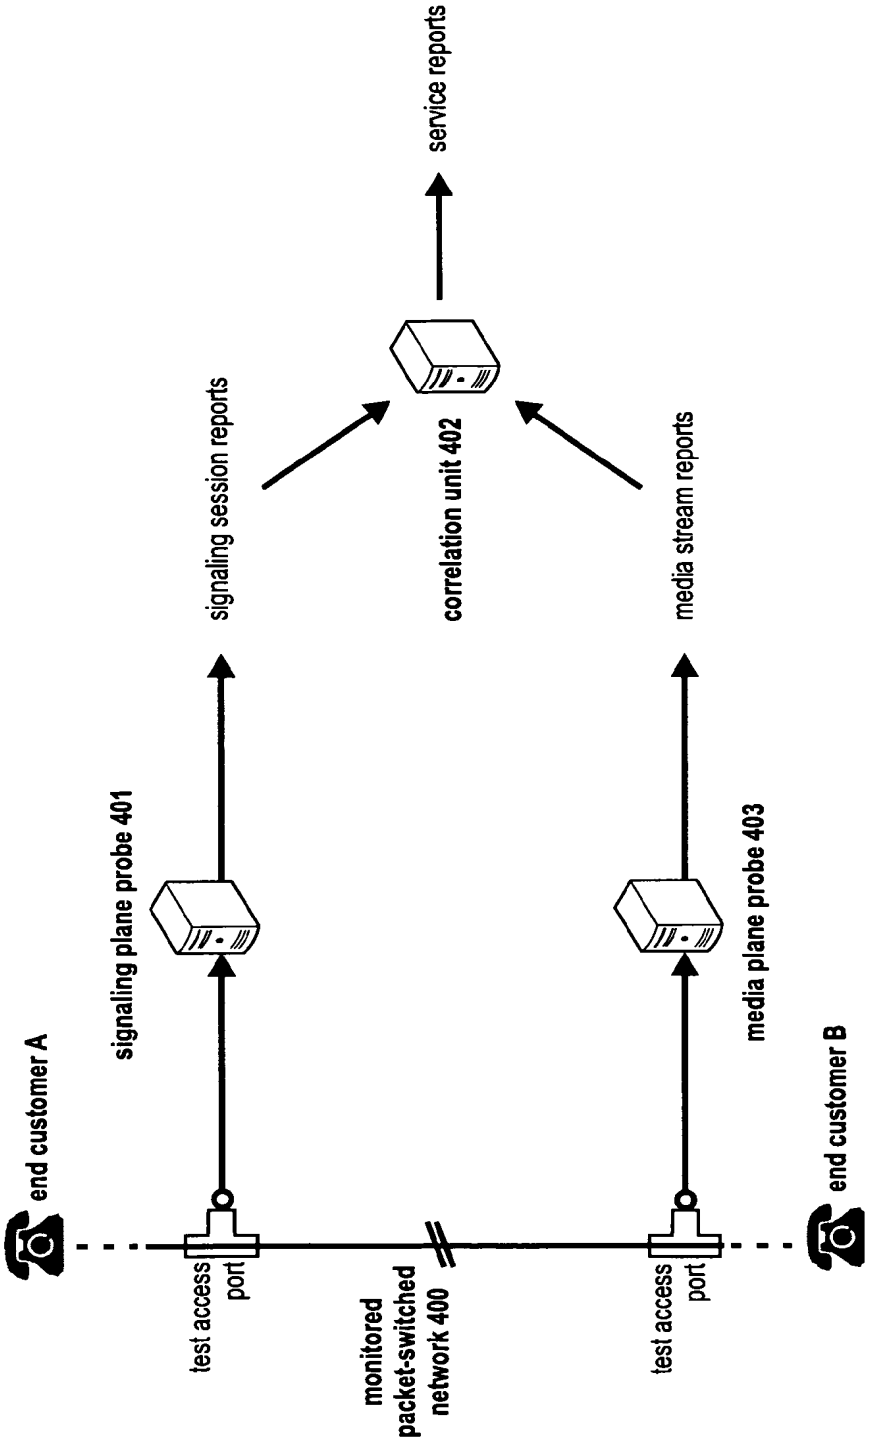
service reports
correlation unit 402
signaling session reports
media stream reports
signaling plane probe 401
media plane probe 403
end customer A
end customer B
test access port
test access port
monitored packet-switched network 400

Fig. 4

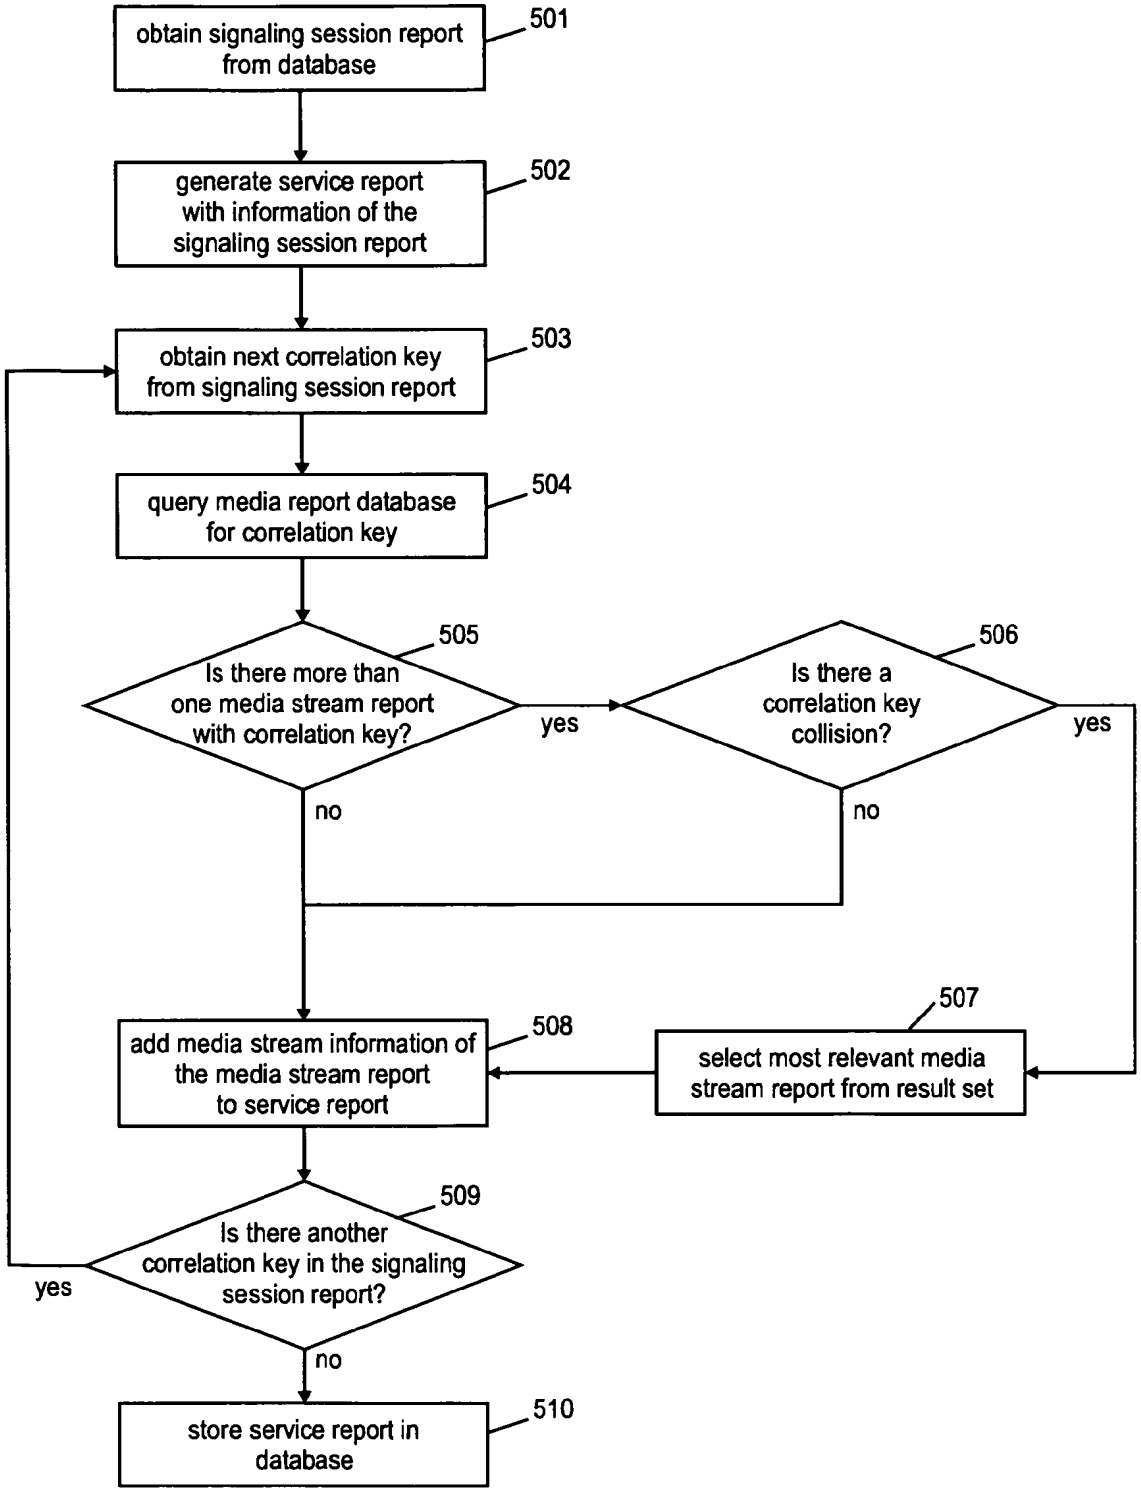
obtain signaling session report from database
501
generate service report with information of the signaling session report
502
obtain next correlation key from signaling session report
503
query media report database for correlation key
504
Is there more than one media stream report with correlation key?
505
yes
no
Is there a correlation key collision?
506
yes
no
add media stream information of the media stream report to service report
508
select most relevant media stream report from result set
507
Is there another correlation key in the signaling session report?
509
yes
no
store service report in database
510

Fig. 5

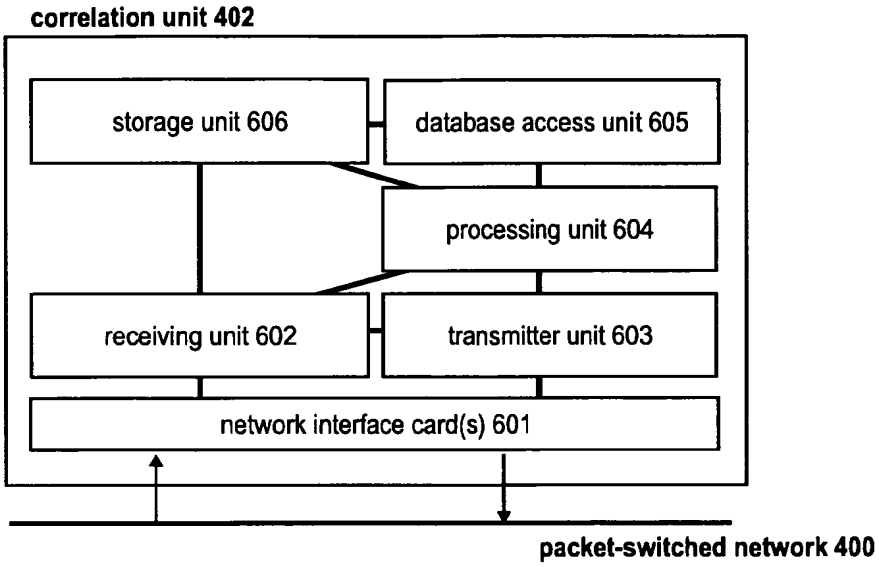
correlation unit 402
storage unit 606
database access unit 605
processing unit 604
receiving unit 602
transmitter unit 603
network interface card(s) 601
packet-switched network 400

Fig. 6

Fig. 7

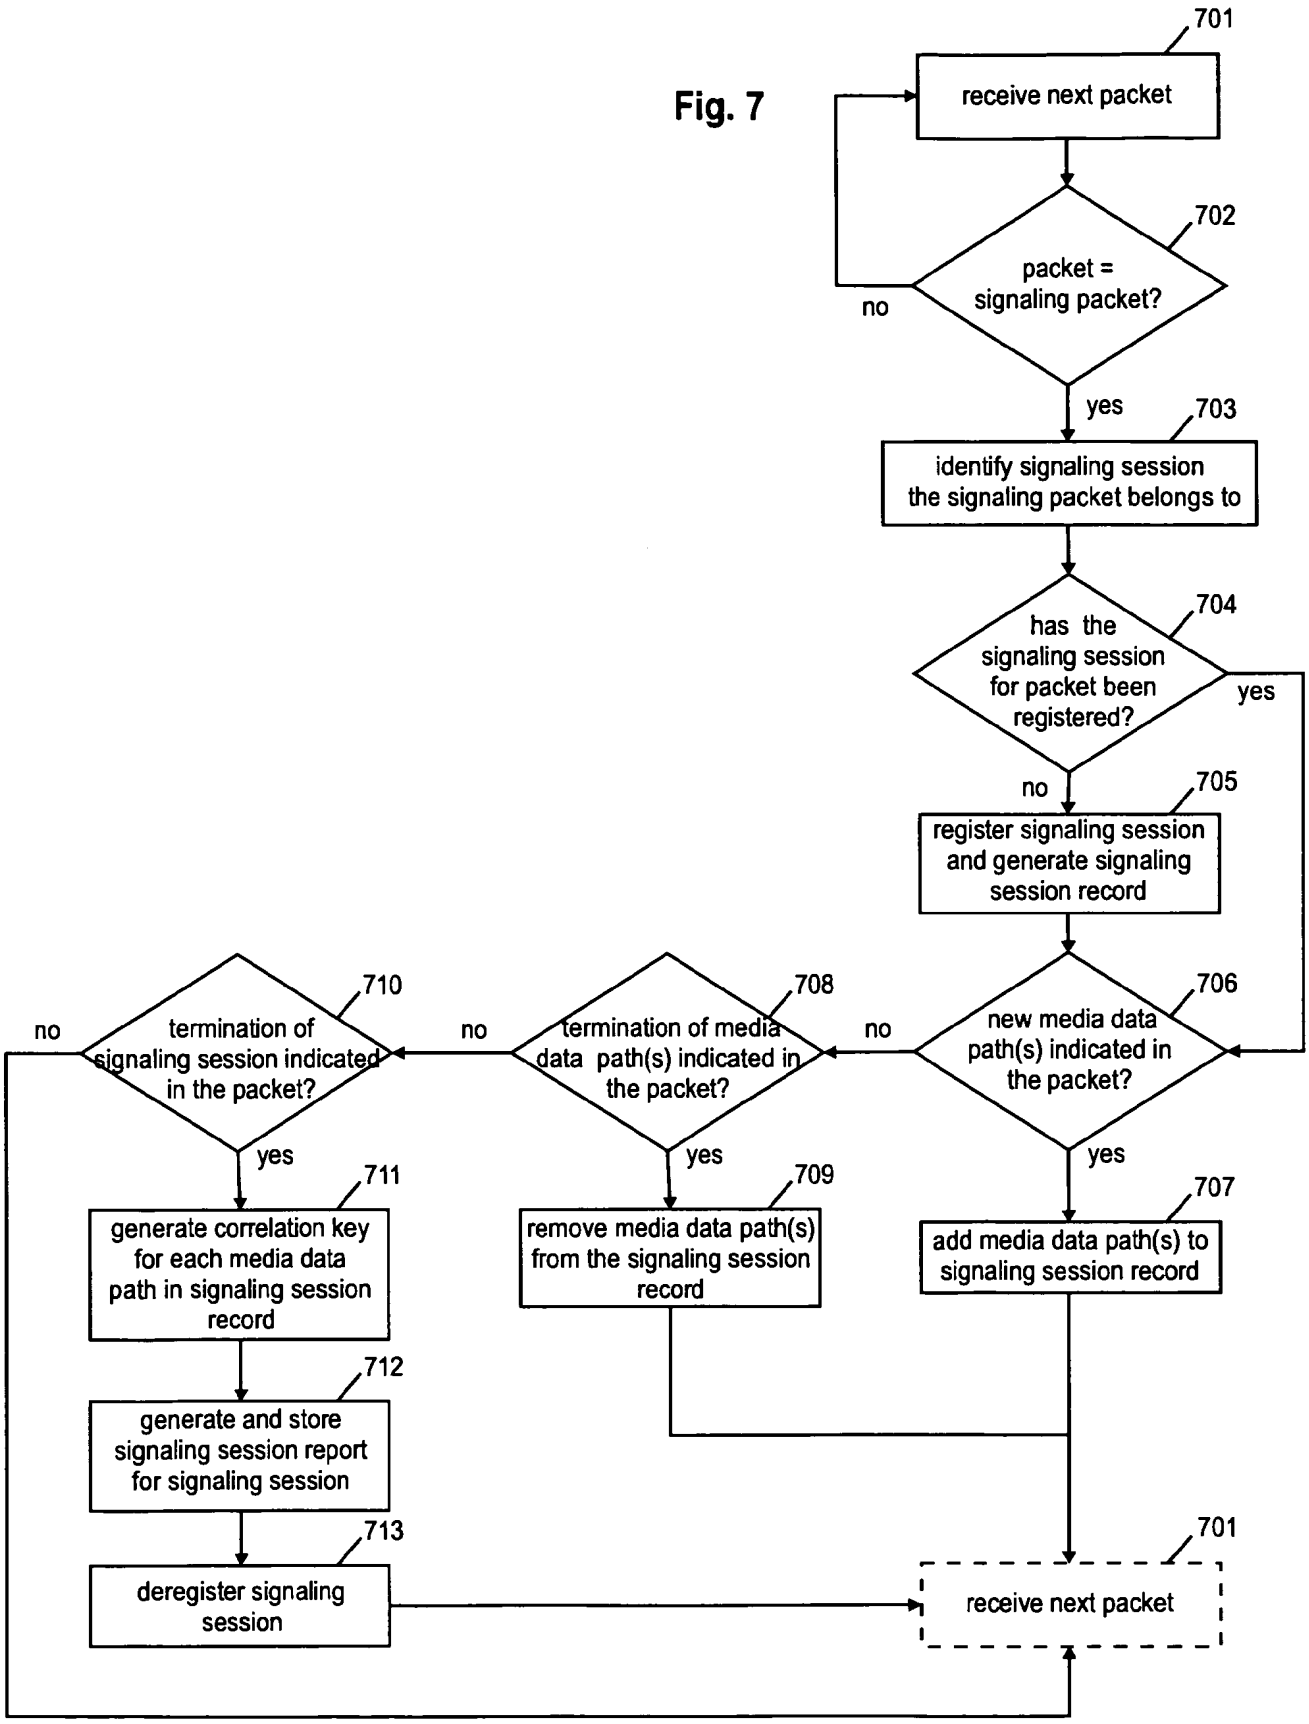
701
receive next packet
702
packet = signaling packet?
no
yes
703
identify signaling session the signaling packet belongs to
704
has the signaling session for packet been registered?
yes
no
705
register signaling session and generate signaling session record
706
new media data path(s) indicated in the packet?
no
yes
707
add media data path(s) to signaling session record
708
termination of media data path(s) indicated in the packet?
no
yes
709
remove media data path(s) from the signaling session record
710
termination of signaling session indicated in the packet?
no
yes
711
generate correlation key for each media data path in signaling session record
712
generate and store signaling session report for signaling session
713
deregister signaling session
701
receive next packet

Fig. 8

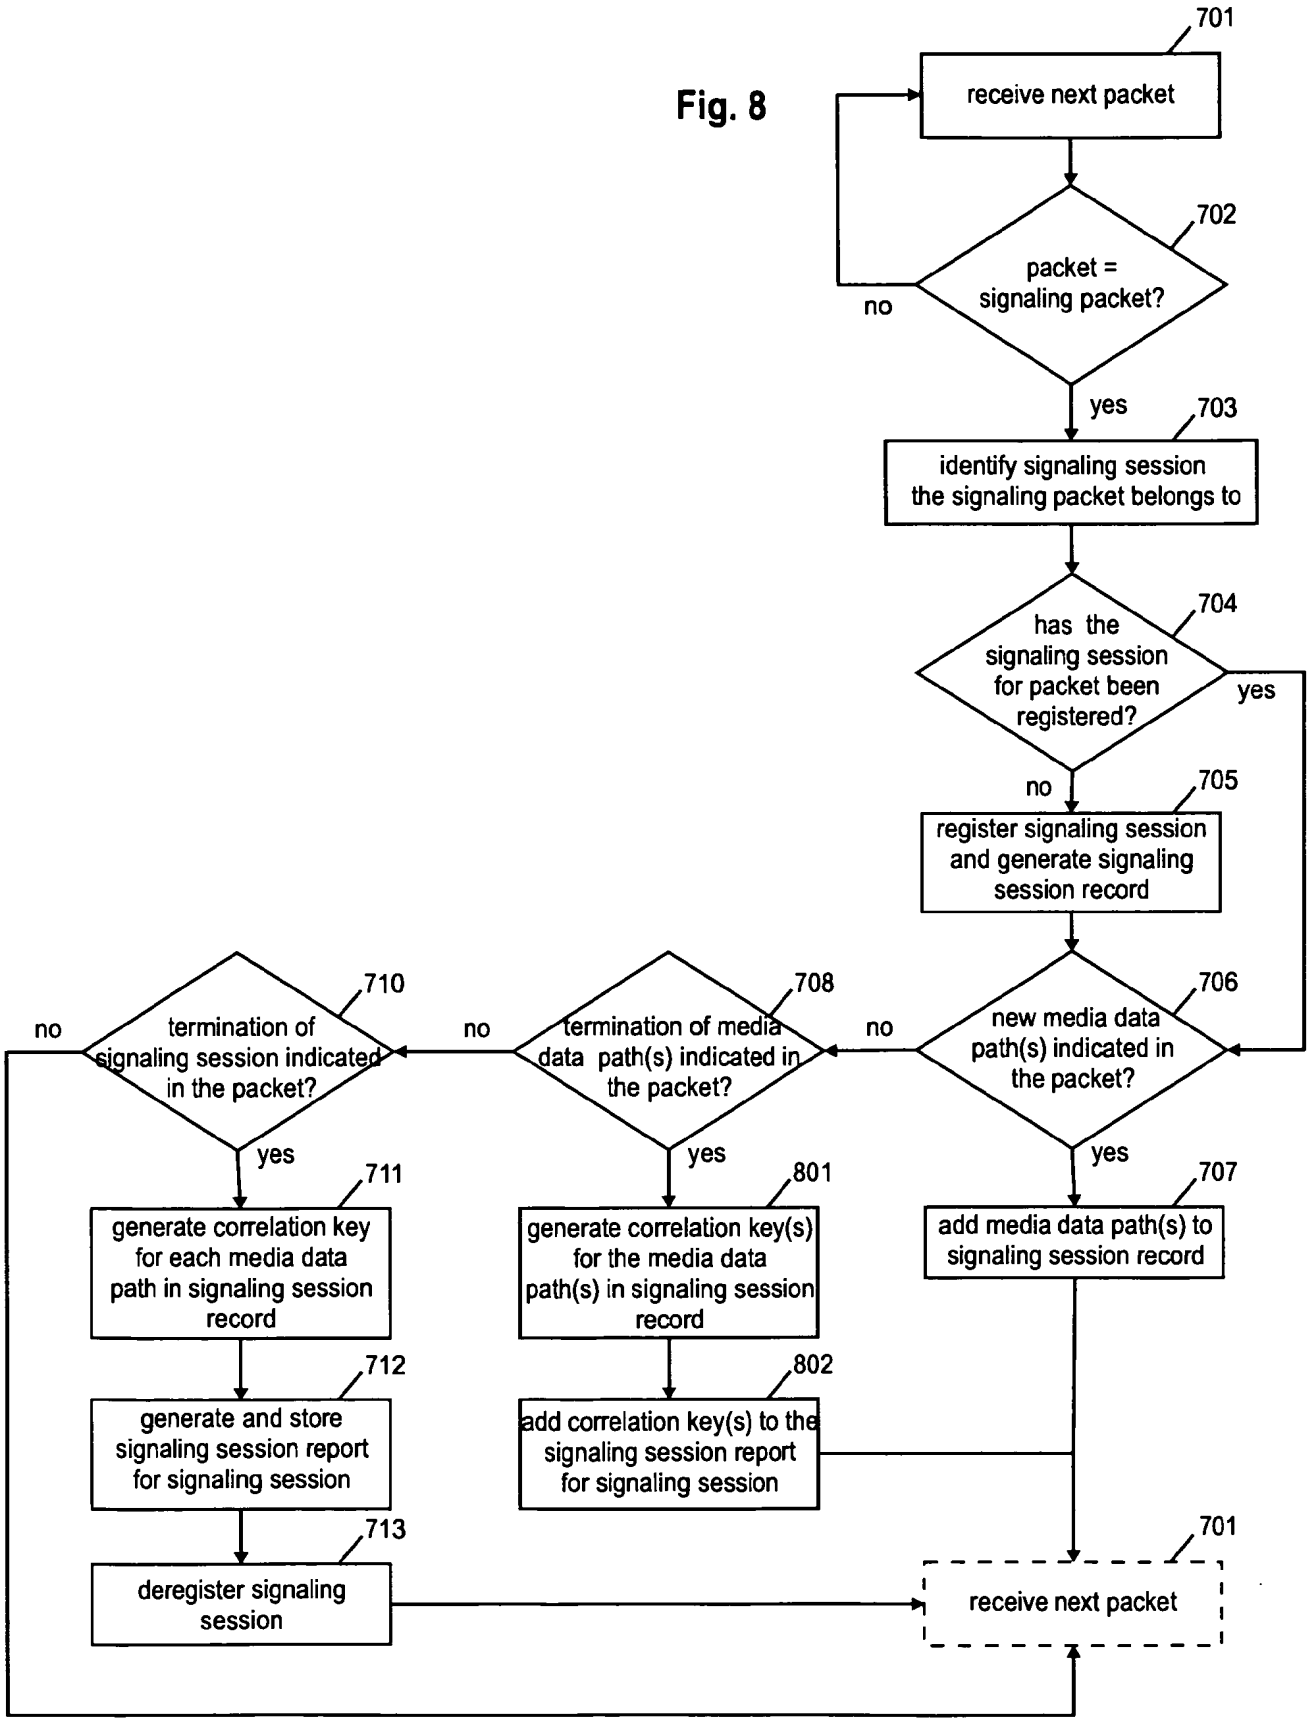
701
receive next packet
702
packet = signaling packet?
no
yes
703
identify signaling session the signaling packet belongs to
704
has the signaling session for packet been registered?
yes
no
705
register signaling session and generate signaling session record
706
new media data path(s) indicated in the packet?
no
yes
707
add media data path(s) to signaling session record
708
termination of media data path(s) indicated in the packet?
no
yes
801
generate correlation key(s) for the media data path(s) in signaling session record
802
add correlation key(s) to the signaling session report for signaling session
710
termination of signaling session indicated in the packet?
no
yes
711
generate correlation key for each media data path in signaling session record
712
generate and store signaling session report for signaling session
713
deregister signaling session
701
receive next packet

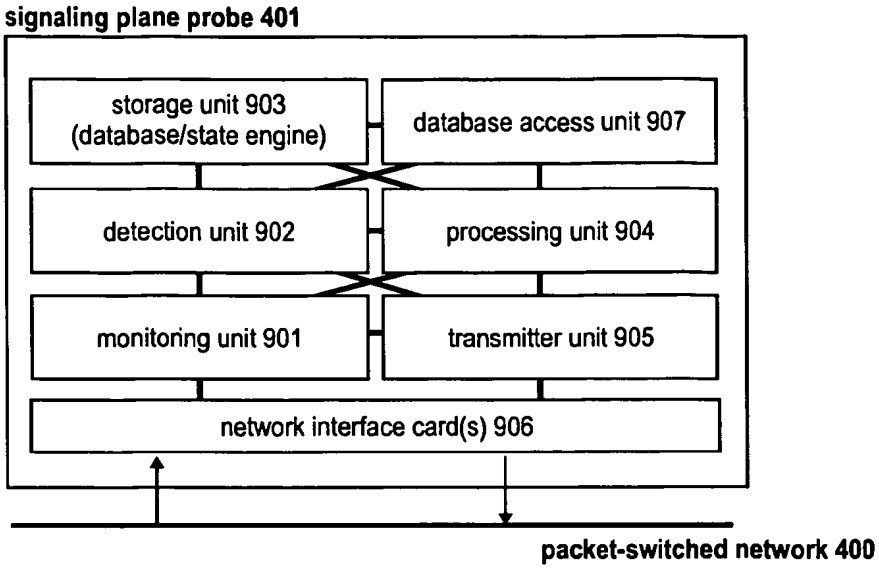
signaling plane probe 401
storage unit 903 (database/state engine)
database access unit 907
detection unit 902
processing unit 904
monitoring unit 901
transmitter unit 905
network interface card(s) 906
packet-switched network 400

Fig. 9

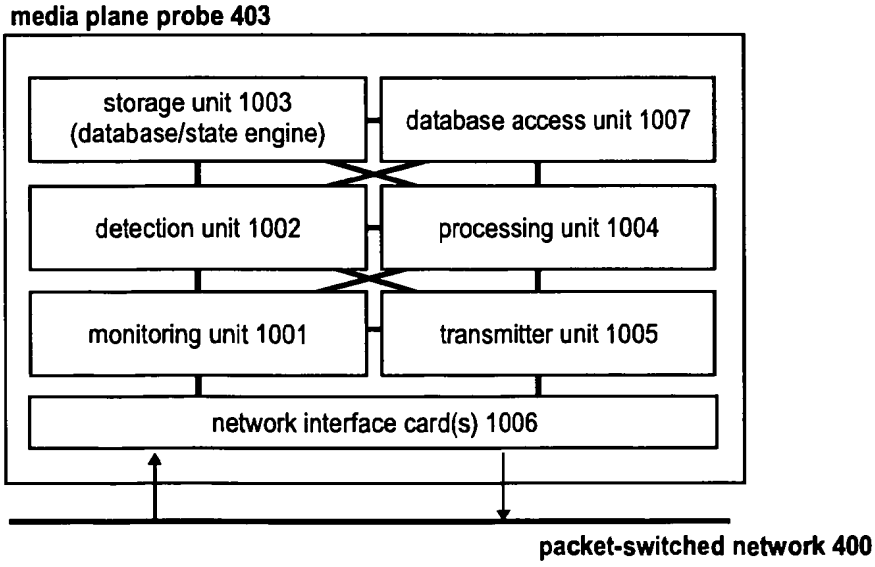
media plane probe 403
storage unit 1003 (database/state engine)
database access unit 1007
detection unit 1002
processing unit 1004
monitoring unit 1001
transmitter unit 1005
network interface card(s) 1006
packet-switched network 400

Fig. 10

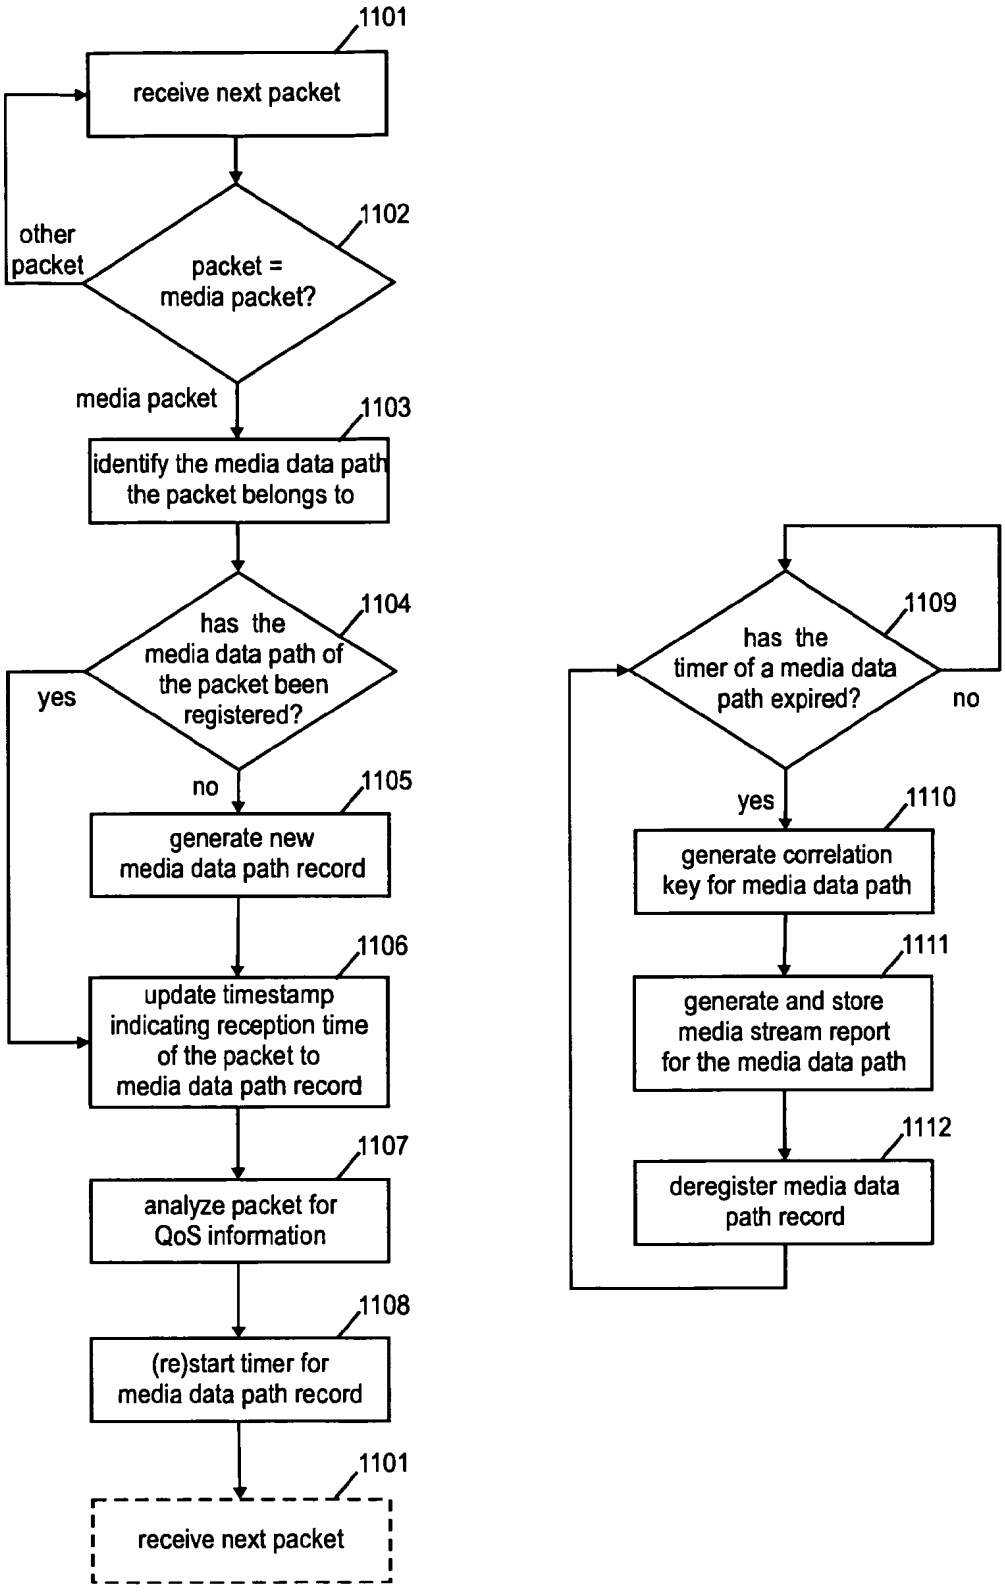
1101
receive next packet
1102
other packet
packet = media packet?
media packet
1103
identify the media data path the packet belongs to
1104
has the media data path of the packet been registered?
yes
no
1105
generate new media data path record
1106
update timestamp indicating reception time of the packet to media data path record
1107
analyze packet for QoS information
1108
(re)start timer for media data path record
1101
receive next packet
1109
has the timer of a media data path expired?
no
yes
1110
generate correlation key for media data path
1111
generate and store media stream report for the media data path
1112
deregister media data path record

Fig. 11

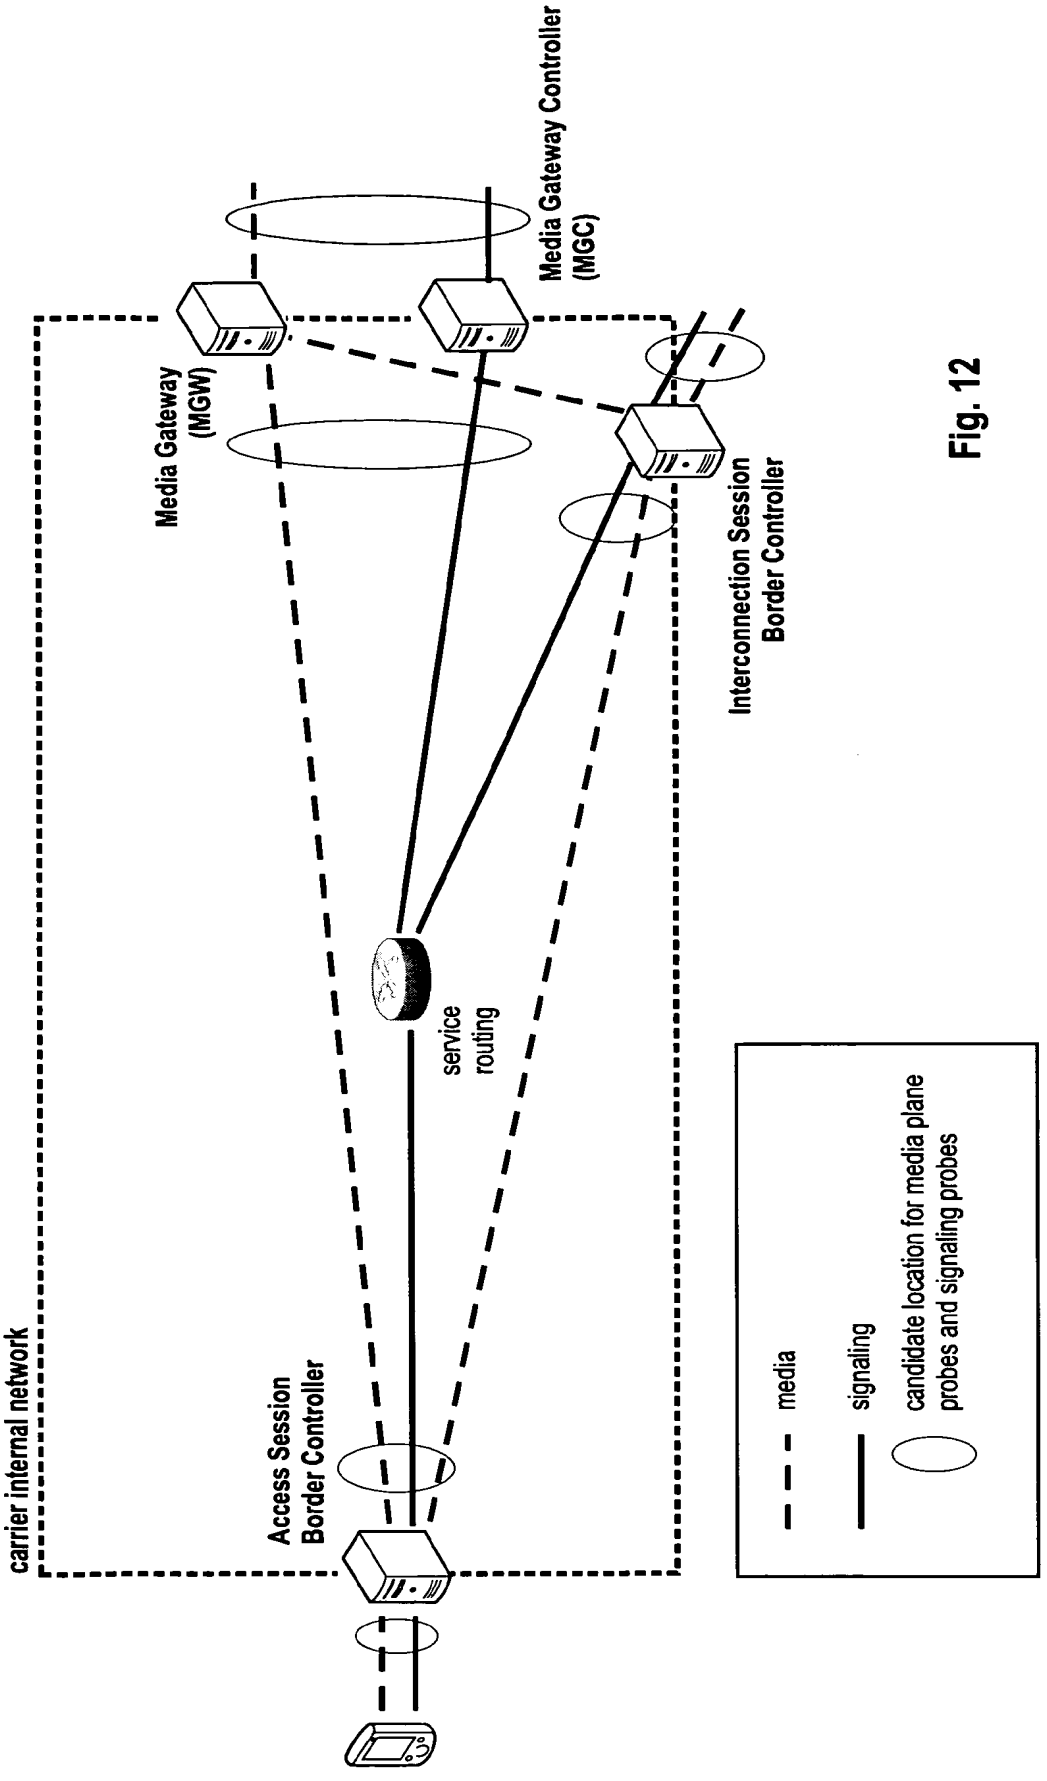
Media Gateway Controller
(MGC)
Media Gateway
(MGW)
Interconnection Session
Border Controller
service
routing
carrier internal network
Access Session
Border Controller
media
signaling
candidate location for media plane
probes and signaling probes

Fig. 12

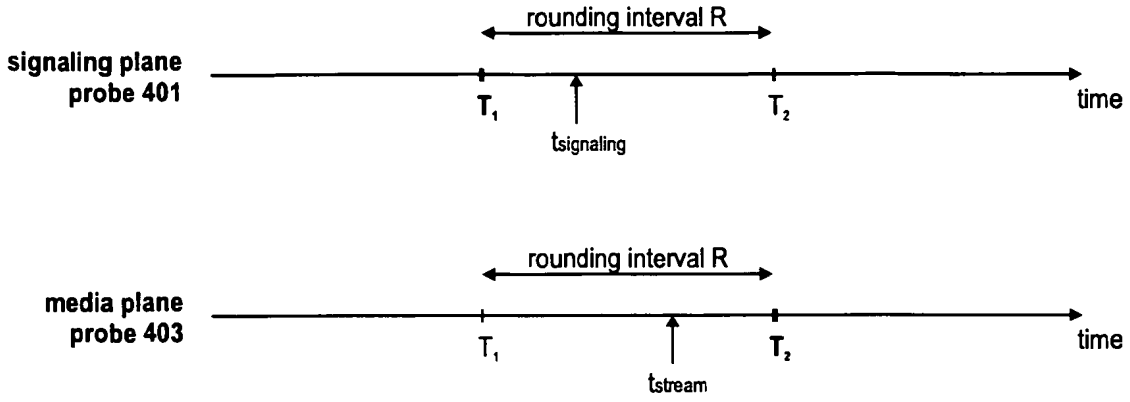
rounding interval R
signaling plane
probe 401
T1
tsignaling
T2
time
rounding interval R
media plane
probe 403
T1
tstream
T2
time

Fig. 13

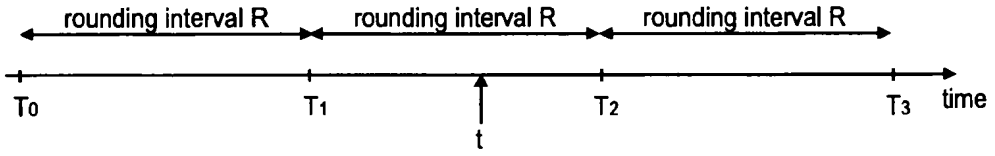
rounding interval R
rounding interval R
rounding interval R
T0
T1
t
T2
T3
time

Fig. 14

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2011/002145

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L29/06 H04L12/24
ADD. H04L12/26

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)
EPO-Internal, COMPENDEX, INSPEC, IBM-TDB, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	ATSUSHI KOBAYASHI ET AL: "VoIP Measurement Architecture Using Data Mediation", 29 October 2009 (2009-10-29), IP OPERATIONS AND MANAGEMENT, SPRINGER BERLIN HEIDELBERG, BERLIN, HEIDELBERG, PAGE(S) 15 - 26, XP019133039, ISBN: 978-3-642-04967-5 page 18, line 5 - page 20, line 35 -----	1-27
A	US 2008/151918 A1 (FOTI GEORGE [CA]) 26 June 2008 (2008-06-26) the whole document -----	1-27
A	WO 2010/106075 A1 (KONINKL KPN NV [NL]; TNO [NL]; STOKKING HANS MAARTEN [NL]; WALRAVEN FA) 23 September 2010 (2010-09-23) page 4, line 16 - page 8, line 26 -----	1-27

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
10 January 2012	17/01/2012

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Aura Marcos, F

3

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/EP2011/002145

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2008151918 A1	26-06-2008	US 2008151918 A1	26-06-2008
		WO 2008078209 A2	03-07-2008
WO 2010106075 A1	23-09-2010	EP 2409432 A1	25-01-2012
		KR 20110125674 A	21-11-2011
		WO 2010106075 A1	23-09-2010

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1192

Electronic Acknowledgement Receipt

EFS ID:	25169929
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	Bradley C. Wright/Holly Parker
Filer Authorized By:	Bradley C. Wright
Attorney Docket Number:	007742.00029
Receipt Date:	11-MAR-2016
Filing Date:	16-APR-2014
Time Stamp:	11:49:19
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Non Patent Literature	Statement_17761019.pdf	66900 0c2d5d1785db2e8352fd788df77b2eb9cd7979e4	no	1

Warnings:	
Information:	

2	Non Patent Literature	12871806_17748444.pdf	486111 d98688f170b4b3d00908b0480ae591948adf4064	no	8
Warnings:					
Information:					
3	Non Patent Literature	128718063_17748442.pdf	1820688 8cd6d88bcf14c95f95d05c5f0936c59af596164f	no	27
Warnings:					
Information:					
4	Non Patent Literature	128718063_17748729.pdf	125961 d0e799e835b53bc9542ab489078408f1ddfd2faa	no	4
Warnings:					
Information:					
5	Non Patent Literature	128718063_17748727.pdf	1457696 fdaf7e08ea22271791869cdfea47c8ae05e27489	no	16
Warnings:					
Information:					
6	Non Patent Literature	128718063_17748748.pdf	202311 d1be9bb8cdc413159ad1666cfa4319695486413e	no	6
Warnings:					
Information:					
7	Non Patent Literature	128718063_17748741.pdf	114514 a2c06b74ce562b85ec3c311b441b4dac68e23a60	no	3
Warnings:					
Information:					
8	Non Patent Literature	128718063_17748734.pdf	245205 61c594a1c82e88ac4bd5d5a245b0f5a2f4879713	no	3
Warnings:					
Information:					
9	Non Patent Literature	128718063_17748766.pdf	282709 389e10afbd02afbfcc755285dd4bda35484710a3	no	4
Warnings:					
Information:					
10	Non Patent Literature	128718063_17748764.pdf	108988 d96beeb8b5e3d9e49b59799b7caefc7e504f61d8	no	3
Warnings:					
Information:					

11	Non Patent Literature	128718063_17748760.pdf	2533792 45ea01d4b457774a01670e82e8a04ecd358d5db1	no	62
Warnings:					
Information:					
12	Non Patent Literature	128718063_17748773.pdf	48581 9eef5fc00a34cc7e354f17b4d82fc48b3cdb54f7	no	1
Warnings:					
Information:					
13	Non Patent Literature	128718063_17750290.pdf	436983 dcf36d568d3932bafc15b58fc4fe86344a11e88d	no	15
Warnings:					
Information:					
14	Non Patent Literature	128718063_17750237.pdf	1206684 0044a226d9c011010f6d8ecf62b589c51a6611e1	no	27
Warnings:					
Information:					
15	Non Patent Literature	128718063_17750240.pdf	637605 57627bd99c5d17db648705cbc74e424bf87acfc8	no	21
Warnings:					
Information:					
16	Non Patent Literature	128718063_17750245.pdf	607940 e3193f3218896730ca77045b69d11f333eb71489	no	20
Warnings:					
Information:					
17	Non Patent Literature	128718063_17750248.pdf	659756 aa9de2722748dd0111873f944e94c2f1ba9de19d	no	10
Warnings:					
Information:					
18	Non Patent Literature	128718063_17750229.pdf	238710 8c34046e74fcbd728967753cd8a303b7c952903d	no	6
Warnings:					
Information:					
19	Non Patent Literature	128718063_17750226.pdf	160790 2e0bf87d22fa9e1e6fc994dc8d5eb4440f2921a6	no	3
Warnings:					
Information:					

20	Non Patent Literature	128718063_17750235.pdf	121567 96ad7b5996c3eac66a9ec1489d78ab155c57da3f	no	2
Warnings:					
Information:					
21	Non Patent Literature	128718063_17750279.pdf	891534 83d6eb67909e68839fe14d24cdfd2300f76b215d	no	27
Warnings:					
Information:					
22	Non Patent Literature	128718063_17750277.pdf	918733 eca7ef41a2f4fd1439ea36b27914b8c2c7394e2d	no	28
Warnings:					
Information:					
23	Foreign Reference	EP1006701A2_17690010.pdf	779545 f2da2b4b1c1220d3c5ce4d5982f3fece2efff587	no	15
Warnings:					
Information:					
24	Foreign Reference	WO2006105093A2_17690109.pdf	2750332 d8f7a6023f473296c2e990d0aadf80b708b2c230	no	62
Warnings:					
Information:					
25	Foreign Reference	WO2006093557A2_17690117.pdf	3097033 8c3b461d2321cbe7ec5f8f229c40c85cbf0f9440	no	72
Warnings:					
Information:					
26	Foreign Reference	128718063_17748754.pdf	3103892 30b788b05754645d32577638f64f4565dd006372	no	74
Warnings:					
Information:					
27	Foreign Reference	WO2011038420A2_17690375.pdf	2007006 6c4c9c4f60be9e685c65752da01e8856c0f6cc73	no	34
Warnings:					
Information:					
28	Foreign Reference	EP1484884A2_17690384.pdf	1943998 a07deab9e2c710976dc4bd58d26f822f871d1c5d	no	32
Warnings:					
Information:					

29	Foreign Reference	128718063_17748768.pdf	2289665 5ceeb63cf91f9aad8cf39d604e915d3d8d04db49	no	60
Warnings:					
Information:					
30	Foreign Reference	EP1864226B1_17690406.pdf	2030938 b36cae0ea8bd094f45d8a5285b295af62b60849a	no	44
Warnings:					
Information:					
31	Foreign Reference	128718063_17750217.pdf	2385931 edf220ad5b5dc6ca0e270598a7fffb36d983a9ce	no	56
Warnings:					
Information:					
32	Foreign Reference	WO2012146265A1_17715972.pdf	4103078 4a100c960bd1ab5910a84302cda0371ff1bb080d	no	78
Warnings:					
Information:					
		Total Files Size (in bytes):	37865176		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of: **Steven Rogers et al.**			
Appl. No.:	14/253,992	Confirmation No.:	5163
Filed:	April 16, 2014	Art Unit:	2497
For:	Methods and Systems For Protecting A Secured Network	Examiner:	Huang, Cheng-Feng

AMENDMENT IN RESPONSE TO NON-FINAL OFFICE ACTION

UNDER 37 C.F.R. § 1.111

United States Patent and Trademark Office
Customer Service Window, Mail Stop Amendment
Randolph Building
401 Dulany Street
Alexandria, VA 22314

Dear Commissioner:

INTRODUCTORY COMMENTS

In response to the Office Action mailed February 26, 2016, please amend the instant application as follows:

Amendments to the Claims are reflected in the listing of claims which begins on page 2 of this paper.

Remarks/Arguments begin on page 8 of this paper.

Applicant requests any necessary extension of time for the submission of this paper. If any fees are required or if an overpayment is made, the Commissioner is authorized to debit or credit our Deposit Account No. 19-0733, accordingly.

AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Currently amended) A ~~method,~~ method comprising:

receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy that ~~includes~~ comprises at least one rule specifying application-layer packet-header information and a packet transformation function to be performed on packets comprising the application-layer packet-header information;

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway;

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets comprising the application-layer packet-header information; ~~and~~

performing, by the packet security gateway and on a packet-by-packet basis, the packet transformation function on each of the one or more packets comprising the application-layer packet-header ~~information.~~ information; and

routing, by the packet security gateway and on a packet-by-packet basis, to a monitoring device each of the one or more packets corresponding to the application-layer packet-header information in response to the performing the packet transformation function.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1199

2. (Currently amended) The method of claim 1, wherein the at least one rule indicates ~~that~~ performing an accept packet transformation function on the packets ~~having~~ comprising the application-layer packet-header ~~information should have an accept packet transformation function performed on them,~~ information, and wherein the performing the packet transformation function ~~on each of the one or more packets comprising the application-layer packet header information~~ comprises ~~the packet security gateway forwarding~~ forwarding, by the packet security gateway, the each of the one or more packets comprising the application-layer packet-header information toward its respective destination.

3. (Currently amended) The method of claim 1, wherein the at least one rule indicates ~~that~~ performing a deny packet transformation function on the packets ~~having~~ comprising the application-layer packet-header ~~information should have a deny packet transformation function performed on them,~~ information, and wherein the performing the packet transformation function on ~~each of the one or more packets comprising the application layer packet header information~~ comprises ~~the packet security gateway dropping~~ dropping, by the packet security gateway, the each of the one or more packets comprising the application-layer packet-header information.

4. (Original) The method of claim 1, wherein the application-layer packet-header information identifies one or more hypertext transfer protocol (HTTP) packets, and wherein identifying the one or more packets comprising the application-layer packet-header information comprises determining by the packet security gateway that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets.

5. (Original) The method of claim 4, wherein the one or more HTTP packets comprise an HTTP GET method call, and wherein determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets comprises determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP GET method call.

6. (Original) The method of claim 5, wherein the HTTP GET method call specifies a uniform resource identifier (URI), and wherein determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP GET method call comprises determining that the one or more packets comprising the application-layer packet-header information originated from or are destined for a network address corresponding to the URI.

7. (Original) The method of claim 4, wherein the one or more HTTP packets comprise an HTTP PUT method call, and wherein determining that the one or more packets comprising the application-layer packet-header information are amongst the one or more HTTP packets comprises determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP PUT method call.

8. (Original) The method of claim 7, wherein the HTTP PUT method call specifies a uniform resource identifier (URI), and wherein determining that the one or more packets comprising the application-layer packet-header information are associated with the HTTP PUT method call comprises determining that the one or more packets comprising the application-layer packet-header information originated from or are destined for a network address corresponding to the URI.

9. (Original) The method of claim 1, wherein the at least one rule comprises a five-tuple specifying one or more transport-layer protocols, a range of source addresses, a range of source ports, a range of destination addresses, and a range of destination ports, and wherein identifying the one or more packets comprising the application-layer packet-header information comprises determining by the packet security gateway that each of the one or more packets comprising the application-layer packet-header information corresponds to at least one of the one or more transport-layer protocols, has a source address within the range of source addresses, a source port within the range of source ports, a destination address within the range of destination addresses, and a destination port within the range of destination ports.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1201

10. (Currently amended) A ~~method,~~ method comprising:

receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy that ~~includes~~ comprises at least one rule specifying packet-identification criteria and a packet transformation function comprising a packet digest logging function to be performed on packets corresponding to the packet-identification criteria, wherein the packet-identification criteria comprises a Differentiated Service Code Point (DSCP) selector;

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway;

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets corresponding to the packet-identification criteria; and

performing, by the packet security gateway and on a packet-by-packet basis, the packet digest logging function on each of the one or more packets corresponding to the packet-identification criteria.

11. (Original) The method of claim 10, wherein performing the packet digest logging function on each of the one or more packets corresponding to the packet-identification criteria comprises:

identifying a subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria; and

generating, for each of the one or more packets corresponding to the packet-identification criteria, a record comprising the subset of information specified by the packet digest logging function.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1202

12. (Original) The method of claim 11, wherein the subset of information comprises at least one of a portion of data from the packet, a source address of the packet, a source port of the packet, a destination address of the packet, a destination port of the packet, a transport-protocol type of the packet, a uniform resource identifier (URI) from the packet, an arrival time of the packet, a size of the packet, a flow direction of the packet, an identifier of an interface of the packet security gateway that received the packet, or one or more media access control (MAC) addresses associated with the packet.

13. (Original) The method of claim 11, comprising:

temporarily storing data in a memory of the packet security gateway, the data comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria; and

utilizing, by the packet security gateway, the data to generate a message comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria.

14. (Currently amended) The method of ~~claim 13~~ claim 13, comprising communicating, by the packet security gateway and to the security policy management server, the message comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria.

15. (Original) The method of claim 11, comprising reformatting, for each of the one or more packets corresponding to the packet-identification criteria, the subset of information specified by the packet digest logging function in accordance with a logging system standard.

16. (Original) The method of claim 15, wherein reformatting the subset of information specified by the packet digest logging function in accordance with the logging system standard comprises reformatting the subset of information specified by the packet digest logging function in accordance with a syslog logging system standard.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1203

17. (Original) The method of claim 10, wherein the packet-identification criteria comprises a five-tuple specifying one or more transport-layer protocols, a range of source addresses, a range of source ports, a range of destination addresses, and a range of destination ports, and wherein identifying the one or more packets corresponding to the packet-identification criteria comprises determining by the packet security gateway that each of the one or more packets corresponding to the packet-identification criteria corresponds to at least one of the one or more transport-layer protocols, has a source address within the range of source addresses, a source port within the range of source ports, a destination address within the range of destination addresses, and a destination port within the range of destination ports.

18. (Original) The method of claim 10, wherein the packet-identification criteria comprises application-layer packet-header information, and wherein identifying the one or more packets corresponding to the packet-identification criteria comprises determining by the packet security gateway that each of the one or more packets corresponding to the packet-identification criteria comprises at least a portion of the application-layer packet-header information.

19. (Previously canceled).

20. (Previously canceled).

21. (New) The method of claim 1, wherein the application-layer packet-header information comprises a Differentiated Service Code Point (DSCP) selector.

22. (New) The method of claim 10, further comprising:

routing, by the packet security gateway and on a packet-by-packet basis, to a monitoring device each of the one or more packets corresponding to the packet-identification criteria in response to the performing the packet digest logging function.

REMARKS/ARGUMENTS

The Office Action of February 26, 2016 has been carefully reviewed and these remarks are responsive thereto. Claims 1-3, 10, and 14 have been amended, claims 21-22 have been added, and no claims have been presently canceled nor withdrawn. Claims 19 and 20 were previously canceled. No new subject matter has been added. Thus, claims 1-18 and 21-22 remain pending in this application, with claims 1 and 10 being independent. Reconsideration and allowance of the instant application are respectfully requested.

Non-Statutory Double Patenting Rejection

Claims 1-18 stand rejected under the non-statutory judicially created doctrine of double patenting over claims 1 and 10 of U.S. Pat. No. 9,137,205. Applicant acknowledges the non-statutory double patenting rejection and reserves the right to submit an appropriate terminal disclaimer, if the non-statutory double patenting rejection is maintained, upon indication that the application is otherwise in condition for allowance. Until such time, Applicant respectfully requests that the double patenting rejections of claims 1-18 be held in abeyance.

Rejections Under 35 U.S.C. § 103

Claims 1-6, 9, 10, 17, and 18 stand rejected under 35 U.S.C. § 103 as being unpatentable over U.S. Pub. No. 2009/0328219, hereinafter *Narayanaswamy*, in view of U.S. Pub. No. 2008/0229415, hereinafter *Kapoor*. Claims 7, 8, and 11-14 stand rejected under 35 U.S.C. § 103 as being unpatentable over *Narayanaswamy*, in view of *Kapoor*, further in view of U.S. Pub. No. 2014/0215574, hereinafter *Erb*. Claims 7, 8, and 11-14 stand rejected under 35 U.S.C. § 103 as being unpatentable over *Narayanaswamy*, in view of *Kapoor*, further in view of *Erb*, and further in view of U.S. Pub. No. 2011/0185055, hereinafter *Nappier*. Applicant respectfully traverses this rejection for at least the following reasons.

Amended independent claim 1, now recites, in part:

> **routing**, by the packet security gateway and on a packet-by-packet basis, **to a monitoring device** each of the one or more packets corresponding to the application-layer packet-header

> information **in response to the performing the packet transformation function**.

(emphasis added). Neither *Narayanaswamy* nor *Kapoor* teach or suggest the features above as recited in claim 1. Thus, the cited documents, alone or in any alleged combination, fail to teach all the features recited in claim 1. Accordingly, Applicant respectfully requests reconsideration and withdrawal of the rejection for at least these reasons.

Amended independent claim 10, now recites, in part:

> receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy that comprises at least one rule specifying packet-identification criteria and a packet transformation function comprising a packet digest logging function to be performed on packets corresponding to the packet-identification criteria, **wherein the packet-identification criteria comprises a Differentiated Service Code Point (DSCP) selector**.

(emphasis added). Neither *Narayanaswamy* nor *Kapoor* teach or suggest the features above as recited in claim 10. Thus, the cited documents, alone or in any alleged combination, fail to teach all the features recited in claim 10. Accordingly, Applicant respectfully requests reconsideration and withdrawal of the rejection for at least these reasons.

Each of dependent claims 2-9 and 11-18 depends on one of independent claims 1 and 10 and is therefore distinguishable over *Narayanaswamy* and *Kapoor* by virtue of its dependence, and in view of the additional features recited therein. Furthermore, *Erb* and *Nappier* do not teach or suggest the features discussed above in reference to claims 1 and 10. Therefore, the dependent claims are distinguishable over the cited documents, alone and in combination at least based on their respective dependencies from the allowable independent claims, and further in view of the additional features recited therein.

New Claims

New claims 21-22 are fully supported and distinguishable over the art of record.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1206

CONCLUSION

All issues having been addressed, Applicant respectfully submits that the instant application is in condition for allowance, and respectfully solicits prompt notification of the same. However, if for any reason the Examiner believes the application is not in condition for allowance or there are any questions, the Examiner is requested to contact the undersigned at (202) 824-3160.

Respectfully submitted,

Date: August 23, 2016 By: /Bradley C. Wright/
Bradley C. Wright
Registration No.: 38,061

BANNER & WITCOFF, LTD.
1100 13th St. NW
STE 1200
Washington, DC 20005
(202) 824-3000

BCW/JAL/hp

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/026339

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L29/06 H04L12/26
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, INSPEC, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2004/093513 A1 (CANTRELL CRAIG [US] ET AL) 13 May 2004 (2004-05-13)	1,31,32
Y	abstract paragraph [0012] - paragraph [0044]; figure 1 paragraph [0065] - paragraph [0069]	2-30, 33-43
X	US 2012/023576 A1 (SORENSEN AMANDA [US] ET AL) 26 January 2012 (2012-01-26)	1,31,32
Y	abstract paragraph [0006] - paragraph [0012]; figure 3	2-30, 33-43
X	US 2008/163333 A1 (KASRALIKAR RAHUL [US]) 3 July 2008 (2008-07-03)	1,31,32
A	abstract paragraph [0009] - paragraph [0010]; figure 2	2-30, 33-43

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
1 June 2016	09/06/2016

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	San Millán Maeso, J

1

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2016/026339

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004093513 A1	13-05-2004	AR 042020 A1	08-06-2005
		AU 2003290674 A1	03-06-2004
		CN 1720459 A	11-01-2006
		EP 1558937 A2	03-08-2005
		JP 2006506853 A	23-02-2006
		JP 2010268483 A	25-11-2010
		KR 20050086441 A	30-08-2005
		KR 20100132079 A	16-12-2010
		US 2004093513 A1	13-05-2004
		US 2005028013 A1	03-02-2005
		US 2005044422 A1	24-02-2005
		WO 2004045126 A2	27-05-2004
US 2012023576 A1	26-01-2012	AU 2011279907 A1	28-02-2013
		CA 2805823 A1	26-01-2012
		SG 187068 A1	28-02-2013
		US 2012023576 A1	26-01-2012
		WO 2012012280 A2	26-01-2012
US 2008163333 A1	03-07-2008	NONE	

Information on Search Strategy - Pilot phase (see OJ 2015, A86)
The type of information contained in this sheet may change during the pilot for improving the usefulness of this new service.

Application Number
PCT/US2016/026339

TITLE: RULE-BASED NETWORK-THREAT DETECTION

APPLICANT: CENTRIPETAL NETWORKS, INC.

IPC CLASSIFICATION: H04L29/06, H04L12/26

EXAMINER: San Millán Maeso, J

CONSULTED DATABASES:

CLASSIFICATION SYMBOLS DEFINING EXTENT OF THE SEARCH:

IPC:

CPC: TXTE, EPODOC, INSPEC, WPI, XPESP, XPI3E, XPIEE, H04L63/0227, H04L43/06, H04L12/2602, H04L43/00, H04L63/0263, H04L63/1416, H04L63/1425, H04L63/1441

FI/F-TERMS:

KEYWORDS OR OTHER ELEMENTS FEATURING THE INVENTION:

EPO FORM P04442

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7215637	B1	2007-05-08	Ferguson et al.	
	2	8176561	B1	2012-05-08	Hurst et al.	
	3	8117655	B2	2012-02-14	Spielman	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20010039579	A1	2001-11-08	TRCKA et al.	
	2	20030105976	A1	2003-06-05	Copeland	
	3	20030154399	A1	2003-08-14	Zuk et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	4	20040250124	A1	2004-12-09	Chesla et al.	
	5	20060053491	A1	2006-03-09	Khuti et al.	
	6	20060137009	A1	2006-06-22	Chesla	
	7	20070211644	A1	2007-09-13	Ottamalika et al.	
	8	20100095367	A1	2010-04-15	Narayanaswamy	
	9	20130117852	A1	2013-05-09	Stute	
	10	20140075510	A1	2014-03-13	Sonoda et al.	
	11	20150334125	A1	2015-11-19	BARTOS et al.	
	12	20120106354	A1	2012-05-03	Pleshek et al.	
	13	20110185055	A1	2011-07-28	NAPPIER et al.	
	14	20130047020	A1	2013-02-21	Hershko et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	15	20090240698	A1	2009-09-24	Shukla et al.	
	16	20040098511	A1	2004-05-20	Lin et al.	
	17	20100232445	A1	2010-09-16	Bellovin	
	18	20130305311	A1	2013-11-14	Puttaswamy Naga et al.	
	19	20020186683	A1	2002-12-12	Buck et al.	
	20	20090222877	A1	2009-09-03	Diehl et al.	
	21	20100107240	A1	2010-04-29	Thaler et al.	
	22	20040093513	A1	2004-05-13	Cantrell et al.	
	23	20120023576	A1	2012-01-26	Sorensen et al.	
	24	20080163333	A1	2008-07-03	Kasralikar	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	2016 Apr 15 - (US) Notice of Allowance - App 14/855,374	
	2	2015 Nov 02 - (AU) Office Action - App 2013372879	
	3	2016 Apr 26 - (US) Office Action - App 14/745,207	
	4	2016 May 06 - (US) Office Action - App 14/714,207	
	5	2016 May 13 - (US) Office Action - App 13/940,240	
	6	2016 Jun 14 - (US) Office Action - App 14/625,486	
	7	2016 Feb 25 - (AU) Office Action - App 2014249055	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	8	2016 Feb 24 - (AU) Office Action - App 2014228257
	9	2016 Jun 09 - (WO) International Search Report - PCT/US2016/026339
	10	2016 Jun 16 - (CA) Office Action - App 2,888,935
	11	2016 Jul 11 - (EP) Office Action - App 14720824.3
	12	2016 Jul 22 - (US) Office Action - App 14/921,718
	13	2016 Jul 20 - (AU) Office Action - App 2013335255

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☐ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

☒ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

☒ A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Bradley C. Wright/	Date (YYYY-MM-DD)	2016-08-23
Name/Print	Bradley C. Wright	Registration Number	38,061

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Patent Application Fee Transmittal

Application Number:	14253992
Filing Date:	16-Apr-2014
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Filer:	Bradley C. Wright/Holly Parker
Attorney Docket Number:	007742.00029

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension - 3 months with $0 paid	2253	1	700	700
Miscellaneous:				
Submission- Information Disclosure Stmt	2806	1	90	90
		Total in USD ($)		**790**

Electronic Acknowledgement Receipt

EFS ID:	26717220
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	Bradley C. Wright/Holly Parker
Filer Authorized By:	Bradley C. Wright
Attorney Docket Number:	007742.00029
Receipt Date:	23-AUG-2016
Filing Date:	16-APR-2014
Time Stamp:	10:13:14
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	DA
Payment was successfully received in RAM	$790
RAM confirmation Number	082316INTEFSW00008711190733
Deposit Account	190733
Authorized User	Bradley Wright

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

37 CFR 1.17 (Patent application and reexamination processing fees)

37 CFR 1.19 (Document supply fees)

37 CFR 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		Amendment_in_Response_to_Non-Final_Office_Action_Under_37_CFR__1_18013835.pdf	55032 bc8f7d5b8ff0b76f5522319601f667a41b33d2f0	yes	10

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Amendment/Req. Reconsideration-After Non-Final Reject	1	1
	Claims	2	7
	Applicant Arguments/Remarks Made in an Amendment	8	10

Warnings:					
Information:					
2	Non Patent Literature	Notice_of_Allowance_and_Fees_Due_PTOL-85__INTVSUMEX__892__1449_17820023.pdf	1978709 4d4e3f669e324489b5cc5e9526926f72d6499efd	no	50
Warnings:					
Information:					
3	Non Patent Literature	AU_2013_372879_1st_report_17844165.PDF	180230 0db14b50d6b54fb780bc134451640ce58c5b8f84	no	3
Warnings:					
Information:					
4	Non Patent Literature	Final_Rejection__892__1449_17836139.pdf	2221034 630c07a4f486c227bd2b154764c1d52922647ed1	no	59
Warnings:					
Information:					
5	Non Patent Literature	Non-Final_Rejection__1449_17863955.pdf	1980426 cc2fae7e71b3f218bf284f104b06ea34ced4be21	no	59

Warnings:					
Information:					
6	Non Patent Literature	Non-Final_Rejection__892__1449_17907654.pdf	3501971 596402d50d623fd856dc9e75f0b2e46796297687	no	86
Warnings:					
Information:					
7	Non Patent Literature	AU2014249055OA_17930089.PDF	131541 2c1fd926dd8411ad1cff21014df6bf3f90176215	no	3
Warnings:					
Information:					
8	Non Patent Literature	AU2014228257OA_17930107.PDF	114304 7ec5e241c75e7cf6e0e9496d6097111f08feb5bd	no	2
Warnings:					
Information:					
9	Non Patent Literature	17920268_1_17936575.pdf	154479 327c57ef9aaf24c59b40ae7d298d36d2ea40f261	no	3
Warnings:					
Information:					
10	Non Patent Literature	17949695_1_17954003.pdf	91536 4f342d662e17ef4b0f97d50d89ea43a15f098bd2	no	4
Warnings:					
Information:					
11	Non Patent Literature	Examination_report_1107_17971763.PDF	233390 ea2ab8ea9b22be899b5e4a32d079f436a8187228	no	7
Warnings:					
Information:					
12	Non Patent Literature	Non-Final_Rejection__892__14493_17967478.pdf	1157498 772a49050d9b3c93da1b4a96b8be50ca296697d5	no	36
Warnings:					
Information:					

13	Non Patent Literature	AUFirstEXam_17988909.pdf	265946 b205b9fb42ea35b902bd6fa3d3f1032d26b65832	no	3
Warnings:					
Information:					
14	Information Disclosure Statement (IDS) Form (SB08)	Information_Disclosure_Statement_Fillable_PDF_18013713.pdf	1036153 dfe05ed9da575abf36d70db157b196148a8760d3	no	7
Warnings:					
Information:					
15	Non Patent Literature	Non-Final_Rejection__892__14493_18013755.pdf	3410436 79f3af09ba3089cbed3dacc044d94d080725b457	no	88
Warnings:					
Information:					
16	Fee Worksheet (SB06)	fee-info.pdf	33032 796d5236b4b6c0b07d5a6d4a15c675ea6e527e83	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	16545717		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number	Filing Date	☐ To be Mailed
14/253,992	04/16/2014	

ENTITY: ☐ LARGE ☒ SMALL ☐ MICRO

APPLICATION AS FILED – PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 =	*	X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 =	*	X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	

APPLICATION AS AMENDED – PART II

AMENDMENT **08/23/2016**

	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
Total (37 CFR 1.16(i))	* 20	Minus	** 20	= 0	X $40 =	0
Independent (37 CFR 1.16(h))	* 3	Minus	***3	= 0	X $210 =	0
☐ Application Size Fee (37 CFR 1.16(s))						
☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
					TOTAL ADD'L FEE	0

AMENDMENT

	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =	
Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =	
☐ Application Size Fee (37 CFR 1.16(s))						
☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
					TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.

** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".

*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".

The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

LIE
William Phillips

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Doc Code: ECOMM.AUTH/ECOMM.WTDW
Doc Description: Internet Communications Authorized/Internet Communications Authorization Withdrawn

PTO/SB/439 (11-15)

AUTHORIZATION FOR INTERNET COMMUNICATIONS IN A PATENT APPLICATION OR REQUEST TO WITHDRAW AUTHORIZATION FOR INTERNET COMMUNICATIONS		
	Application No.	14/253,992
	Filing Date	April 16, 2014
	First Named Inventor	Steven Rogers et al.
	Art Unit	2497
	Examiner Name	Cheng-Feng Huang
	Practitioner Docket No.	007742.00029

To: Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

I. To authorize permission for Internet Communications.

☑ Recognizing that Internet communications are not secure, I hereby authorize the USPTO to communicate with the undersigned and practitioners in accordance with 37 CFR 1.33 and 37 CFR 1.34 concerning any subject matter of this application via video conferencing, instant messaging, or electronic mail. I understand that a copy of these communications will be made of record in the application file. (MPEP 502.03)

II. To withdraw authorization for Internet Communications.

☐ The authorization given on _____, to the USPTO to communicate with the undersigned and any practitioner in accordance with 37 CFR 1.33 and 37 CFR 1.34 concerning any subject matter of this application via Internet communications such as video conferencing, instant messaging, or electronic mail is hereby withdrawn. I understand that the withdrawal is effective when approved rather than when received.

I am the

☐ applicant.

☑ attorney or agent of record. Registration number 38,061.

☐ attorney or agent acting under 37 CFR 1.34. Registration number _____.

/Bradley C. Wright/ — Signature

November 8, 2016 — Date

Bradley C. Wright — Typed or printed name

202-824-3000 — Telephone Number

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. Juristic entities must be represented by a patent practitioner (see 37 CFR 1.31, which is applicable to any paper filed on or after September 16, 2012, that is presented on behalf of a juristic entity, regardless of application filing date). Submit multiple forms if more than one signature is required, see below*.

☐ * Total of _____ forms are submitted.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	27448998
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	Bradley C. Wright/Holly Parker
Filer Authorized By:	Bradley C. Wright
Attorney Docket Number:	007742.00029
Receipt Date:	08-NOV-2016
Filing Date:	16-APR-2014
Time Stamp:	15:01:24
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Internet Communications Authorized	sb0439_18124270.PDF	110315 88f9a1a2ec0e7fc473fcc7c6ccce5a3efedd6bd6	no	2

Warnings:

Information:	
Total Files Size (in bytes):	110315

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Doc Code: DIST.E.FILE **Document Description: Electronic Terminal Disclaimer - Filed**	PTO/SB/26 U.S. Patent and Trademark Office Department of Commerce

Electronic Petition Request	**TERMINAL DISCLAIMER TO OBVIATE A DOUBLE PATENTING REJECTION OVER A "PRIOR" PATENT**
Application Number	14253992
Filing Date	16-Apr-2014
First Named Inventor	Steven Rogers
Attorney Docket Number	007742.00029
Title of Invention	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

☒ Filing of terminal disclaimer does not obviate requirement for response under 37 CFR 1.111 to outstanding Office Action

☒ This electronic Terminal Disclaimer is not being used for a Joint Research Agreement.

Owner	Percent Interest
Centripetal Networks, Inc.	100%

The owner(s) with percent interest listed above in the instant application hereby disclaims, except as provided below, the terminal part of the statutory term of any patent granted on the instant application which would extend beyond the expiration date of the full statutory term of prior patent number(s)

9137205

as the term of said prior patent is presently shortened by any terminal disclaimer. The owner hereby agrees that any patent so granted on the instant application shall be enforceable only for and during such period that it and the prior patent are commonly owned. This agreement runs with any patent granted on the instant application and is binding upon the grantee, its successors or assigns.

In making the above disclaimer, the owner does not disclaim the terminal part of the term of any patent granted on the instant application that would extend to the expiration date of the full statutory term of the prior patent, "as the term of said prior patent is presently shortened by any terminal disclaimer," in the event that said prior patent later:
- expires for failure to pay a maintenance fee;
- is held unenforceable;
- is found invalid by a court of competent jurisdiction;
- is statutorily disclaimed in whole or terminally disclaimed under 37 CFR 1.321;
- has all claims canceled by a reexamination certificate;
- is reissued; or
- is in any manner terminated prior to the expiration of its full statutory term as presently shortened by any terminal disclaimer.

◉ Terminal disclaimer fee under 37 CFR 1.20(d) is included with Electronic Terminal Disclaimer request.

○	I certify, in accordance with 37 CFR 1.4(d)(4), that the terminal disclaimer fee under 37 CFR 1.20(d) required for this terminal disclaimer has already been paid in the above-identified application.

Applicant claims the following fee status:

◉ Small Entity

○ Micro Entity

○ Regular Undiscounted

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

THIS PORTION MUST BE COMPLETED BY THE SIGNATORY OR SIGNATORIES

I certify, in accordance with 37 CFR 1.4(d)(4) that I am:

◉ An attorney or agent registered to practice before the Patent and Trademark Office who is of record in this application

Registration Number 38061

○ A sole inventor

○ A joint inventor; I certify that I am authorized to sign this submission on behalf of all of the inventors as evidenced by the power of attorney in the application

○ A joint inventor; all of whom are signing this request

Signature	/Bradley C. Wright/
Name	Bradley C. Wright

*Statement under 37 CFR 3.73(b) is required if terminal disclaimer is signed by the assignee (owner).
Form PTO/SB/96 may be used for making this certification. See MPEP § 324.

Electronic Patent Application Fee Transmittal

Application Number:	14253992
Filing Date:	16-Apr-2014
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Filer:	Bradley C. Wright/Holly Parker
Attorney Docket Number:	007742.00029

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
STATUTORY OR TERMINAL DISCLAIMER	2814	1	160	160
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension-of-Time:				
Miscellaneous:				
			Total in USD ($)	**160**

Doc Code: DISQ.E.FILE
Document Description: Electronic Terminal Disclaimer – Approved

Application No.: 14253992

Filing Date: 16-Apr-2014

Applicant/Patent under Reexamination: Rogers

Electronic Terminal Disclaimer filed on November 8, 2016

☒ APPROVED

This patent is subject to a terminal disclaimer

☐ DISAPPROVED

Approved/Disapproved by: Electronic Terminal Disclaimer automatically approved by EFS-Web

U.S. Patent and Trademark Office

Electronic Acknowledgement Receipt

EFS ID:	27450294
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	Bradley C. Wright/Holly Parker
Filer Authorized By:	Bradley C. Wright
Attorney Docket Number:	007742.00029
Receipt Date:	08-NOV-2016
Filing Date:	16-APR-2014
Time Stamp:	15:49:30
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	DA
Payment was successfully received in RAM	$160
RAM confirmation Number	110916INTEFSW00002413190733
Deposit Account	190733
Authorized User	Bradley Wright

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

37 CFR 1.17 (Patent application and reexamination processing fees)

37 CFR 1.19 (Document supply fees)

37 CFR 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Electronic Terminal Disclaimer-Filed	eTerminal-Disclaimer.pdf	33528 c616a100f1944e86c26f21bdf7e3bd75ec8e04ea	no	2
Warnings:					
Information:					
2	Fee Worksheet (SB06)	fee-info.pdf	30740 ad9b10799ba10b37033b1cff1f21b60e8c85d958	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	64268		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

NOTICE OF ALLOWANCE AND FEE(S) DUE

22907 7590 11/16/2016

BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051

EXAMINER
HUANG, CHENG-FENG

ART UNIT	PAPER NUMBER
2497	

DATE MAILED: 11/16/2016

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
14/253,992	04/16/2014	Steven Rogers	007742.00029	5163

TITLE OF INVENTION: METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	SMALL	$480	$0	$0	$480	02/16/2017

THE APPLICATION IDENTIFIED ABOVE HAS BEEN EXAMINED AND IS ALLOWED FOR ISSUANCE AS A PATENT. PROSECUTION ON THE MERITS IS CLOSED. THIS NOTICE OF ALLOWANCE IS NOT A GRANT OF PATENT RIGHTS. THIS APPLICATION IS SUBJECT TO WITHDRAWAL FROM ISSUE AT THE INITIATIVE OF THE OFFICE OR UPON PETITION BY THE APPLICANT. SEE 37 CFR 1.313 AND MPEP 1308.

THE ISSUE FEE AND PUBLICATION FEE (IF REQUIRED) MUST BE PAID WITHIN THREE MONTHS FROM THE MAILING DATE OF THIS NOTICE OR THIS APPLICATION SHALL BE REGARDED AS ABANDONED. THIS STATUTORY PERIOD CANNOT BE EXTENDED. SEE 35 U.S.C. 151. THE ISSUE FEE DUE INDICATED ABOVE DOES NOT REFLECT A CREDIT FOR ANY PREVIOUSLY PAID ISSUE FEE IN THIS APPLICATION. IF AN ISSUE FEE HAS PREVIOUSLY BEEN PAID IN THIS APPLICATION (AS SHOWN ABOVE), THE RETURN OF PART B OF THIS FORM WILL BE CONSIDERED A REQUEST TO REAPPLY THE PREVIOUSLY PAID ISSUE FEE TOWARD THE ISSUE FEE NOW DUE.

HOW TO REPLY TO THIS NOTICE:

I. Review the ENTITY STATUS shown above. If the ENTITY STATUS is shown as SMALL or MICRO, verify whether entitlement to that entity status still applies.

If the ENTITY STATUS is the same as shown above, pay the TOTAL FEE(S) DUE shown above.

If the ENTITY STATUS is changed from that shown above, on PART B - FEE(S) TRANSMITTAL, complete section number 5 titled "Change in Entity Status (from status indicated above)".

For purposes of this notice, small entity fees are 1/2 the amount of undiscounted fees, and micro entity fees are 1/2 the amount of small entity fees.

II. PART B - FEE(S) TRANSMITTAL, or its equivalent, must be completed and returned to the United States Patent and Trademark Office (USPTO) with your ISSUE FEE and PUBLICATION FEE (if required). If you are charging the fee(s) to your deposit account, section "4b" of Part B - Fee(s) Transmittal should be completed and an extra copy of the form should be submitted. If an equivalent of Part B is filed, a request to reapply a previously paid issue fee must be clearly made, and delays in processing may occur due to the difficulty in recognizing the paper as an equivalent of Part B.

III. All communications regarding this application must give the application number. Please direct all communications prior to issuance to Mail Stop ISSUE FEE unless advised to the contrary.

IMPORTANT REMINDER: Utility patents issuing on applications filed on or after Dec. 12, 1980 may require payment of maintenance fees. It is patentee's responsibility to ensure timely payment of maintenance fees when due.

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), to: Mail **Mail Stop ISSUE FEE**
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450
or Fax (571)-273-2885

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

22907 7590 11/16/2016
BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission
I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being facsimile transmitted to the USPTO (571) 273-2885, on the date indicated below.

(Depositor's name)
(Signature)
(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
14/253,992	04/16/2014	Steven Rogers	007742.00029	5163

TITLE OF INVENTION: METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	SMALL	$480	$0	$0	$480	02/16/2017

EXAMINER	ART UNIT	CLASS-SUBCLASS
HUANG, CHENG-FENG	2497	726-001000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).
- ❑ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.
- ❑ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-02 or more recent) attached. **Use of a Customer Number is required.**

2. For printing on the patent front page, list
(1) The names of up to 3 registered patent attorneys or agents OR, alternatively,
(2) The name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.
1 _____
2 _____
3 _____

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document has been filed for recordation as set forth in 37 CFR 3.11. Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE (B) RESIDENCE: (CITY and STATE OR COUNTRY)

Please check the appropriate assignee category or categories (will not be printed on the patent) : ❑ Individual ❑ Corporation or other private group entity ❑ Government

4a. The following fee(s) are submitted:
- ❑ Issue Fee
- ❑ Publication Fee (No small entity discount permitted)
- ❑ Advance Order - # of Copies _____

4b. Payment of Fee(s): **(Please first reapply any previously paid issue fee shown above)**
- ❑ A check is enclosed.
- ❑ Payment by credit card. Form PTO-2038 is attached.
- ❑ The director is hereby authorized to charge the required fee(s), any deficiency, or credits any overpayment, to Deposit Account Number _____ (enclose an extra copy of this form).

5. **Change in Entity Status** (from status indicated above)
- ❑ Applicant certifying micro entity status. See 37 CFR 1.29 — NOTE: Absent a valid certification of Micro Entity Status (see forms PTO/SB/15A and 15B), issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment.
- ❑ Applicant asserting small entity status. See 37 CFR 1.27 — NOTE: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.
- ❑ Applicant changing to regular undiscounted fee status. — NOTE: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

NOTE: This form must be signed in accordance with 37 CFR 1.31 and 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Authorized Signature _____ Date _____

Typed or printed name _____ Registration No. _____

CISCO EXHIBIT 1002
File History of 9565213 Patent

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
14/253,992	04/16/2014	Steven Rogers	007742.00029	5163

22907 7590 11/16/2016

BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051

EXAMINER	
HUANG, CHENG-FENG	

ART UNIT	PAPER NUMBER
2497	

DATE MAILED: 11/16/2016

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(Applications filed on or after May 29, 2000)

The Office has discontinued providing a Patent Term Adjustment (PTA) calculation with the Notice of Allowance.

Section 1(h)(2) of the AIA Technical Corrections Act amended 35 U.S.C. 154(b)(3)(B)(i) to eliminate the requirement that the Office provide a patent term adjustment determination with the notice of allowance. See Revisions to Patent Term Adjustment, 78 Fed. Reg. 19416, 19417 (Apr. 1, 2013). Therefore, the Office is no longer providing an initial patent term adjustment determination with the notice of allowance. The Office will continue to provide a patent term adjustment determination with the Issue Notification Letter that is mailed to applicant approximately three weeks prior to the issue date of the patent, and will include the patent term adjustment on the patent. Any request for reconsideration of the patent term adjustment determination (or reinstatement of patent term adjustment) should follow the process outlined in 37 CFR 1.705.

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at 1-(888)-786-0101 or (571)-272-4200.

OMB Clearance and PRA Burden Statement for PTOL-85 Part B

The Paperwork Reduction Act (PRA) of 1995 requires Federal agencies to obtain Office of Management and Budget approval before requesting most types of information from the public. When OMB approves an agency request to collect information from the public, OMB (i) provides a valid OMB Control Number and expiration date for the agency to display on the instrument that will be used to collect the information and (ii) requires the agency to inform the public about the OMB Control Number's legal significance in accordance with 5 CFR 1320.5(b).

The information collected by PTOL-85 Part B is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450. Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Notice of Allowability	Application No. 14/253,992	Applicant(s) ROGERS ET AL.	
	Examiner CHENG-FENG HUANG	Art Unit 2497	AIA (First Inventor to File) Status Yes

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☒ This communication is responsive to *08/23/2016*.
   ☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on_____.
2. ☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.
3. ☒ The allowed claim(s) is/are *1-10,12-14 and 16-18*. As a result of the allowed claim(s), you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to PPHfeedback@uspto.gov.
4. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

   Certified copies:

   a) ☐ All b) ☐ Some *c) ☐ None of the:
   1. ☐ Certified copies of the priority documents have been received.
   2. ☐ Certified copies of the priority documents have been received in Application No. _____ .
   3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

   * Certified copies not received: _____.

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
   ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____.

   Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).
6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

1. ☐ Notice of References Cited (PTO-892)
2. ☒ Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date See Continuation Sheet
3. ☐ Examiner's Comment Regarding Requirement for Deposit of Biological Material
4. ☒ Interview Summary (PTO-413), Paper No./Mail Date *11/8/2016* .
5. ☒ Examiner's Amendment/Comment
6. ☒ Examiner's Statement of Reasons for Allowance
7. ☐ Other _____.

/HARUNUR RASHID/ Primary Examiner, Art Unit 2497	/CHENG-FENG HUANG/ Examiner, Art Unit 2497

Continuation of Attachment(s) 2. Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date: 3/11/2016, 3/11/2016, and 8/23/2016.

CISCO EXHIBIT 1002
File History of 9565213 Patent
Page 1241

DETAILED ACTION

1. The present application, filed on or after March 16, 2013, is being examined under the first inventor to file provisions of the AIA.

Response to Amendment

2. This is a reply to the application filed on 8/23/2016, in which, claim(s) 1-18 and 21-22 are pending. Claim(s) 1, and 10 are independent. Claim(s) 1-3, 10, and 14 are amended. Claim(s) 21-22 are newly added.

Information Disclosure Statement

3. The information disclosure statement (IDS) submitted on 3/11/2016, 3/11/2016, and 8/23/2016 has been reviewed. The submission is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Terminal Disclaimer

4. The terminal disclaimer filed on 11/8/2016 disclaiming the terminal portion of any patent granted on this application which would extend beyond the expiration date of US Patent No. 9,137,205 has been reviewed and is accepted. The terminal disclaimer has been recorded.

Response to Arguments

Claim Rejections - 35 U.S.C. § 103:

5. Applicants' arguments with respect to claims rejected under prior art have been fully considered and the rejection of 35 U.S.C. § 103 have been withdrawn in view of the examiner's amendment to claim.

EXAMINER'S AMENDMENT

6. An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

7. Authorization for this examiner's amendment was given in a telephone interview with Attorney Bradley Wright (Reg. 38,061) on 11/8/2016.

The application has been amended as follows:

- ***Please replace Claim 1 with:***

1. (Currently Amended) A method comprising:

receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy that comprises at least one rule specifying application-layer packet-header information and a packet

transformation function comprising a packet digest logging function to be performed on packets comprising the application-layer packet-header information;

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway;

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets comprising the application-layer packet-header information;

performing, by the packet security gateway and on a packet-by-packet basis, the packet transformation function on each of the one or more packets comprising the application-layer packet-header ~~information; and~~ information, wherein the performing the packet transformation function comprises

identifying a subset of information specified by the packet digest logging function for each of the one or more packets comprising the application-layer packet-header information;

generating, for each of the one or more packets comprising the application-layer packet-header information, a record comprising the subset of information specified by the packet digest logging function; and

reformatting, for each of the one or more packets comprising the application-layer packet-header information, the subset of information specified by the packet digest logging function in accordance with a logging system standard; and

routing, by the packet security gateway and on a packet-by-packet basis, to a monitoring device each of the one or more packets corresponding to the application-layer packet-header information in response to the performing the packet transformation function.

- ***Please replace Claim 10 with:***

10. (Currently amended) A method comprising:

receiving, by each of a plurality of packet security gateways associated with a security policy management server and from the security policy management server, a dynamic security policy that comprises at least one rule specifying packet-identification criteria and a packet transformation function comprising a packet digest logging function to be performed on packets corresponding to the packet-identification criteria, wherein the packet-identification criteria comprises a Differentiated Service Code Point (DSCP) selector;

receiving, by a packet security gateway of the plurality of packet security gateways, packets associated with a network protected by the packet security gateway;

identifying, by the packet security gateway, from amongst the packets associated with the network protected by the packet security gateway, and on a packet-by-packet basis, one or more packets corresponding to the packet-identification criteria;

performing, by the packet security gateway and on a packet-by-packet basis, the packet digest logging function on each of the one or more packets corresponding to the packet-identification ~~criteria.~~ criteria, wherein the performing the packet digest logging function comprises:

identifying a subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria;

generating, for each of the one or more packets corresponding to the packet-identification criteria, a record comprising the subset of information specified by the packet digest logging function; and

reformatting, for each of the one or more packets corresponding to the packet-identification criteria, the subset of information specified by the packet digest logging function in accordance with a logging system standard; and

routing, by the packet security gateway and on a packet-by-packet basis, to a monitoring device each of the one or more packets corresponding to the packet-identification criteria in response to the performing the packet digest logging function.

- ***Please cancel Claim 11;***

- ***Please replace Claim 12 with:***

12. (Currently amended) The method of ~~claim 11,~~ claim 10, wherein the subset of information comprises at least one of a portion of data from the packet, a source address of the packet, a source port of the packet, a destination address of the packet, a destination port of the packet, a transport-protocol type of the packet, a uniform resource identifier (URI) from the packet, an arrival time of the packet, a size of the packet, a flow direction of the packet, an identifier of an interface of the packet security gateway that received the packet, or one or more media access control (MAC) addresses associated with the packet.

- ***Please replace Claim 13 with:***

13. (Currently amended) The method of ~~claim 11,~~ claim 10, comprising:

temporarily storing data in a memory of the packet security gateway, the data comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria; and

utilizing, by the packet security gateway, the data to generate a message comprising the subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria.

- ***Please cancel Claim 15;***

- ***Please replace Claim 16 with:***

16. (Currently amended) The method of ~~claim 15,~~ claim 10, wherein reformatting the subset of information specified by the packet digest logging function in accordance with the logging system standard comprises reformatting the subset of information specified by the packet digest logging function in accordance with a syslog logging system standard.

- ***Please cancel Claims 21-22;***

Allowable Subject Matter

8. **Claims 1-10, 12-14, and 16-18 are allowed.**

9. The following is an examiner's statement of reasons for allowance:

Independent Claim(s) 1, and 10 and their respective dependent claims are allowable over prior arts since the prior arts taken individually or in combination fails to particular discloses, fairly suggest or render obvious the following italic limitations:

In claim **1**:

*"**wherein the performing the packet transformation function comprises***

identifying a subset of information specified by the packet digest logging function for each of the one or more packets comprising the application-layer packet-header information;

generating, for each of the one or more packets comprising the application-layer packet-header information, a record comprising the subset of information specified by the packet digest logging function; and

reformatting, for each of the one or more packets comprising the application-layer packet-header information, the subset of information specified by the packet digest logging function in accordance with a logging system standard;" in combination with other limitations recited as specified in the independent claim(s).

In claim **10**:

*"**wherein the performing the packet digest logging function comprises:***

identifying a subset of information specified by the packet digest logging function for each of the one or more packets corresponding to the packet-identification criteria;

generating, for each of the one or more packets corresponding to the packet-identification criteria, a record comprising the subset of information specified by the packet digest logging function; and

reformatting, for each of the one or more packets corresponding to the packet-identification criteria, the subset of information specified by the packet digest logging function in accordance with a logging system standard;" in combination with other limitations recited as specified in the independent claim(s).

10. **The closest prior art of record teaches**:

- Krishna Narayanaswamy (US 2009/0328219 A1, cited by the applicant in the 03/13/2015 IDS) teaches dynamic policy provisioning. A network security device may comprise a memory that stores a first policy that identifies a first set of patterns that correspond to a first set of network attacks and a second policy, and a control unit that applies the first policy to the network traffic to detect the first set of network attacks. The control unit, while applying the first policy, monitors parameters corresponding to one or more resources and dynamically determines whether to apply a second policy to the network traffic based on the parameters. The control

unit, based on the dynamic determination, applies the second policy to the network traffic to detect a second set of network attacks and forwards the network traffic based on the application of the second policy. In this manner, the network security device may implement the dynamic policy provisioning techniques.

- Kapoor et al. (US 2008/0229415 A1) teaches A flow processing for switching, security, and other network applications, including a facility that processes a data flow to address patterns relevant to a variety of conditions are directed at internal network security, virtualization, and web connection security. A flow processing facility for inspecting payloads of network traffic packets detects security threats and intrusions across accessible layers of the IP-stack by applying content matching and behavioral anomaly detection techniques based on regular expression matching and self-organizing maps. Exposing threats and intrusions within packet payload at or near real-time rates enhances network security from both external and internal sources while ensuring security policy is rigorously applied to data and system resources.
- Erb et al. (US 2014/0215574 A1) teaches A hosted storage service stores a virtual data object that corresponds to data. The virtual data object includes metadata that enables access to the data in a delegated storage service but does not include the data. A delegate storage service stores the data. The hosted storage service receives a request for access to the

virtual object and sends a response that includes metadata to access the data in a delegated storage service. The delegate storage service receives a request for access to the data based on the metadata. In response to receiving the request for access to the data object, the delegate storage service sends the data to the client application.

- Nappier et al. (US 2011/0185055 A1, cited by the applicant in the 07/21/2015 IDS) teaches a log correlation engine which analyzes various event logs that describe activity in a network to learn relationships between network identities and network addresses and generates alerts in response to discovering changes in the learned relationships. For example, the log correlation engine may identify authentication events described in the logs to map network identities to IP addresses, and may further analyze the logs to map the IP addresses to Ethernet addresses. Thus, the log correlation engine may discover new and changed relationships between the network identities, the IP addresses, and the Ethernet addresses.

11. Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled “Comments on Statement of Reasons for Allowance.”

Conclusion

12. Any inquiry concerning this communication or earlier communications from the examiner should be directed to CHENG-FENG HUANG whose telephone number is (571)272-6186. The examiner can normally be reached on 7:30 am ~ 5:00 pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Hadi Armouche can be reached on 571-270-3618. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/CHENG-FENG HUANG/
Examiner, Art Unit 2497

/HARUNUR RASHID/
Primary Examiner, Art Unit 2497

Examiner-Initiated Interview Summary	Application No.	Applicant(s)	
	14/253,992	ROGERS ET AL.	
	Examiner	Art Unit	
	CHENG-FENG HUANG	2497	

All participants (applicant, applicant's representative, PTO personnel):

(1) *CHENG-FENG HUANG*.

(2) *Bradley Wright (Reg. No. 38,061)*.

(3) *Rashid Harunur (Primary)*.

(4) *Javier Liceaga*.

Date of Interview: *08 November 2016*.

Type: ☒ Telephonic ☐ Video Conference
☐ Personal [copy given to: ☐ applicant ☐ applicant's representative]

Exhibit shown or demonstration conducted: ☐ Yes ☒ No.
If Yes, brief description: _____.

Issues Discussed ☐101 ☐112 ☐102 ☒103 ☒Others
(For each of the checked box(es) above, please describe below the issue and detailed description of the discussion)

Claim(s) discussed: *1,10,11,15 and 22*.

Identification of prior art discussed: *Narayanaswamy and Kapoor* .

Substance of Interview
(For each issue discussed, provide a detailed description and indicate if agreement was reached. Some topics may include: identification or clarification of a reference or a portion thereof, claim interpretation, proposed amendments, arguments of any applied references etc...)

Examiners first thanked Attorney for taking this interview. Examiners suggested to move up claims 11, 15 and 22 to claim 10, and make claim 1 with similar limitations as claims 11 and 15 to put the application in the condition of allowance. Examiners also suggested to file electronic Terminal Disclaimer ASAP. Agreement was reached..

Applicant recordation instructions: It is not necessary for applicant to provide a separate record of the substance of interview.

Examiner recordation instructions: Examiners must summarize the substance of any interview of record. A complete and proper recordation of the substance of an interview should include the items listed in MPEP 713.04 for complete and proper recordation including the identification of the general thrust of each argument or issue discussed, a general indication of any other pertinent matters discussed regarding patentability and the general results or outcome of the interview, to include an indication as to whether or not agreement was reached on the issues raised.

☐ Attachment

/HARUNUR RASHID/ Primary Examiner, Art Unit 2497	/CHENG-FENG HUANG/ Examiner, Art Unit 2497

Doc code: IDS
Doc description: Information Disclosure Statement (IDS) Filed

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /C.H./

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	1	R.L. Graham et al., "Optimization and Approximation in Deterministic Sequencing and Scheduling: A Survey", Annals of Discrete Mathematics, 5: 287-326, 1979	
	2	E.L. Lawler, "Sequencing Jobs to Minimize Total Weighted Completion oTime Subject to Precedence Constraints", Annals of Discrete Mathematics, 2: 75-90, 1978	
	3	J.K. Lenstra et al., "Complexity of Scheduling Under Precedence Constraints", Operations Research, 26(1): 22-35,1978	
	4	R. Rivest, "On Self-Organizing Sequential Search Heuristics", Communications of the ACM, 19(2): 1976	
	5	W.E. Smith, "Various Optimizers for Single-Stage Productions", Naval Research Logistics Quarterly, 3: 59-66, 1956	
	6	Bellion, "HIgh Performance Packet Classification", http://www.hipac.org (Publication Date Unknown)	
	7	2011 Oct 18 - (EP) Communication Pursuant to Article 94(3) - App 06 758 213.0	
	8	2011 Jun 09 - (US) Notice of Allowance - App 11/390,976	
	9	2012 Jun 26 - (EP) Extended Search Report - App 05857614.1	
	10	2012 Jun 09 - (AU) Notice of Acceptance - App 2006230171	
	11	2011 Nov 11 - (AU) Second Office Action - App 2006230171	

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /C.H./

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	12	2013 Jan 17 - (CA) Office Action - App 2,600,236	
	13	2013 Jan 16 - (CA) Office Action - App 2,594,020	
	14	2012 Nov 20 - (EP) Communication under rule 71(3) - App 06 758 213.0	
	15	2013 Apr 18 - (EP) Decision to Grant a European Patent - App 06758212.0	
	16	2011 Aug 25 - (US) Non Final Rejection - App 12/871,806	
	17	2012 Feb 06 - (US) Final Rejection - App 12/871,806	
	18	2012 Aug 07 - (US) Non Final Rejection - App 12/871,806	
	19	2012 Nov 26 - (US) Final Rejection - App 12/871,806	
	20	2013 Apr 04 - (US) Notice of Allowance - App 12/871,806	
	21	2015 Jan 14 - (EP) Extended Search Report - App 10819667.6	
	22	2014 May 26 - (CA) Office Action - App 2010297968	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	23	2015 May 25 - (AU) Notice of Acceptance - App 2010297968	
	24	2015 May 14 - (US) Non Final Rejection - App 13/940,240	
	25	2015 Nov 27 - (US) Final Rejection - App 13/940,240	
	26	2016 Jan 28 - (WO) International Search Report and Written Opinion - App PCT/US2015/062691	
	27	Statement RE: Related Application, dated July 24, 2015	
	28	2015 Dec 22 - (US) Final Office Action - App 14/714,207	

If you wish to add additional non-patent literature document citation information please click the Add button [Add]

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	11/08/2016

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S131	1427	((Steven) near2 (Rogers)).INV.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 13:44
S132	274	((Sean) near2 (Moore)).INV.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 13:44
S133	168	((David) near2 (Ahn)).INV.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 13:44
S134	7	((Peter) near2 (Geremia)).INV.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 13:45
S135	24	((Centripetal)).AS.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 13:45
S136	1842	S131 S132 S133 S134 S135	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 13:45
S137	4	S136 and ((security adj policy adj management adj server) with (security adj policy adj management adj server)).clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 13:58
S138	4	S136 and ((packet adj security adj	US-PGPUB;	OR	ON	2016/11/02

		gateway) with (security adj policy adj management adj server)).clm.	USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB			14:01
S139	4	S136 and ((dynamic adj security adj policy) with (packet adj transformation adj function)).clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 14:02
S140	4	S136 and ((dynamic adj security adj policy)).clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 14:02
S141	1	S136 and ((packet adj security adj gateway) with (rout$3)).clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 14:03
S142	365	("20120113987" | "20010039624" | "20030142681" | "20100082811" | "20080005795" | "20120314617" | "8806638" | "20110185055" | "20120331543" | "20070240208" | "6317837" | "20050010765" | "20140115654" | "7263099" | "20100095367" | "8176561" | "20060248580" | "7684400" | "20050138204" | "20060212572" | "20100296441" | "7539186" | "20060146879" | "8306994" | "20040177139" | "20040205360" | "20140283004" | "20150237012" | "7107613" | "7237267" | "7331061" | "8495725" | "20060053491" | "20120023576" | "20130305311" | "20110270956" | "20130059527" | "20030035370" | "20110088092" | "20090328219" | "8935785" | "20030154297" | "20020049899" | "9137205" | "20140215574" | "20050114704" | "20050183140" | "20060104202" | "20090138938" | "20140201123" | "6098172" | "7089581" | "20040093513" | "7215637" | "20060195896" | "20100011434" | "20120264443" | "7818794" | "20140366132" | "20080229415" | "20020038339" | "20040010712" | "20050251570" | "20060070122" | "20100199346" | "20110055923" | "6484261" | "7299353" | "20090240698" | "20110185055" | "20120106354" | "20140075510" | "6226372" |	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 14:04

		"US20060195896" | "20040151155" | "20050141537" | "20080072307" | "20100132027" | "20060114899" | "20130061294" | "8726379" | "20050286522" | "20140281030" | "20150304354" | "7227842" | "7954143" | "20020186683" | "20030154399" | "20040098511" | "20110055916" | "20030212900" | "20050117576" | "20030145225" | "20030097590" | "7478429" | "8004994" | "20100242098" | "20040193943" | "20050229246" | "20140283030" | "8037517" | "9124552" | "20080163333" | "20090222877" | "20100107240" | "20130117852" | "8117655" | "20060048142" | "20030188192" | "7721084" | "7710885" | "20040088542" | "20080235755" | "20020165949" | "20020198981" | "20050125697" | "20100303240" | "9094445" | "9160713" | "20010039579" | "20040250124" | "20130047020" | "US20060248580" | "20030123456" | "20070083924" | "20100211678" | "8856926" | "20090172800" | "20120240135" | "6611875" | "20130254766" | "20060136987" | "20030120622" | "20040073655" | "20080077705" | "20080301765" | "20100011433" | "20100268799" | "6662235" | "8042167" | "20030105976" | "20060137009" | "20070211644" | "20100232445" | "20150334125").PN.				
S143	3	S142 and ((dynamic adj security adj policy)).clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 14:05
S144	12	((11/316331) or (11/390976) or (14/714207) or (12/871806) or (13/940240) or (14/625486) or (14/855374) or (14/745207) or (14/921718)).APP.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2016/11/02 14:16
S145	0	S144 and ((dynamic adj security adj policy)).clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/02 14:16
S146	0	S144 and ((packet adj security adj gateway)).clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT;	OR	ON	2016/11/02 14:16

			IBM_TDB			
S147	1	(14/253992).APP.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2016/11/08 16:48
S148	20	(US-20140115654-$ or US-20140283004-$ or US-20030074473-$ or US-20050091437-$ or US-20110173692-$ or US-20090328219-$ or US-20130254766-$ or US-20110055916-$ or US-20030123456-$ or US-20050141537-$ or US-20070240208-$ or US-20060136987-$ or US-20060146879-$ or US-20120240185-$ or US-20080262990-$ or US-20080229415-$ or US-20140215574-$ or US-20130066833-$ or US-20110185055-$).did. or (US-9137205-$).did.	US-PGPUB; USPAT	OR	ON	2016/11/08 16:49
S149	2793	((DSCP adj field) (differentiated adj service adj code adj point))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 16:50
S150	157	(packet adj security adj gateway)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 17:42
S151	10	S149 and S150	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 17:44
S152	4	S151 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 17:44
S153	79920	(H04L63/20 or H04L67/02 or H04L63/0209).cpc.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 17:59
S154	2	S153 and ((packet adj transformation adj function) with (((application-$layer) (application near layer)) near ((packet-$header) (packet near	US-PGPUB; USPAT; USOCR; FPRS;	OR	ON	2016/11/08 17:59

		header))))	EPO; JPO; DERWENT; IBM_TDB			
S155	184	("20120113987" | "20010039624" | "20030142681" | "20100082811" | "20080005795" | "20120314617" | "8806638" | "20110185055" | "20120331543" | "20070240208" | "6317837" | "20060248580" | "7684400" | "20050138204" | "20060212572" | "20100296441" | "7539186" | "20060146879" | "8306994" | "20110270956" | "20130059527" | "20030035370" | "20110088092" | "20090328219" | "8935785" | "20030154297" | "20020049899" | "9137205" | "20060195896" | "20100011434" | "20120264443" | "7818794" | "20140366132" | "6226372" | "US20060195896" | "20040151155" | "20050141537" | "20080072307" | "20100132027" | "20060114899" | "20130061294" | "8726379" | "20110055916" | "20030212900" | "20050117576" | "20030145225" | "20030097590" | "7478429" | "8004994" | "20100242098" | "20060048142" | "20030188192" | "7721084" | "7710885" | "20040088542" | "20080235755" | "US20060248580" | "20030123456" | "20070083924" | "20100211678" | "8856926" | "20090172800" | "20120240135" | "6611875" | "20130254766" | "20060136987").PN.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 19:13
S156	4	S155 and ((packet adj header) with application)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 19:13
S157	4	S156 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 19:13
S158	1	(identifying with (subset near information) with specif$3 with (packet near digest near logging near function) with criteria)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 19:25
S159	1	(generating with record with (subset near information) with specif$3 with (packet near digest near logging near function))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO;	OR	ON	2016/11/08 19:28

			DERWENT; IBM_TDB			
S160	0	(reformatting with (subset near information) with specif$3 with (packet near digest near logging near function)) with (logging near standard)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 19:28
S161	2129120	((denial near service) Dos DDos)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 19:30
S162	10784	S161 and ((generating creat$3 produc$3) near (record log message digest))	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 19:30
S163	217	S162 and (HTTP near GET)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 19:30
S164	197	S163 and @ad<"20140416"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2016/11/08 19:30

EAST Search History (Interference)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S165	1	(identifying with (subset near information) with specif$3 with (packet near digest near logging near function) with criteria).clm.	US-PGPUB; USPAT	OR	ON	2016/11/08 19:25
S167	1	(generating with record with (subset near information) with specif$3 with (packet near digest near logging near function)).clm.	US-PGPUB; USPAT	OR	ON	2016/11/08 19:28
S168	0	(reformatting with (subset near information) with specif$3 with (packet near digest near logging near function)) with (logging near standard).clm.	US-PGPUB; USPAT	OR	ON	2016/11/08 19:30

11/ 8/ 2016 7:32:45 PM
C:\ Users\ chuang2\ Documents\ EAST\ Workspaces\ 14253992.wsp

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	7215637	B1	2007-05-08	Ferguson et al.	
	2	8176561	B1	2012-05-08	Hurst et al.	
	3	8117655	B2	2012-02-14	Spielman	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20010039579	A1	2001-11-08	TRCKA et al.	
	2	20030105976	A1	2003-06-05	Copeland	
	3	20030154399	A1	2003-08-14	Zuk et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	4	20040250124	A1	2004-12-09	Chesla et al.	
	5	20060053491	A1	2006-03-09	Khuti et al.	
	6	20060137009	A1	2006-06-22	Chesla	
	7	20070211644	A1	2007-09-13	Ottamalika et al.	
	8	20100095367	A1	2010-04-15	Narayanaswamy	
	9	20130117852	A1	2013-05-09	Stute	
	10	20140075510	A1	2014-03-13	Sonoda et al.	
	11	20150334125	A1	2015-11-19	BARTOS et al.	
	12	20120106354	A1	2012-05-03	Pleshek et al.	
	13	20110185055	A1	2011-07-28	NAPPIER et al.	
	14	20130047020	A1	2013-02-21	Hershko et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	15	20090240698	A1	2009-09-24	Shukla et al.	
	16	20040098511	A1	2004-05-20	Lin et al.	
	17	20100232445	A1	2010-09-16	Bellovin	
	18	20130305311	A1	2013-11-14	Puttaswamy Naga et al.	
	19	20020186683	A1	2002-12-12	Buck et al.	
	20	20090222877	A1	2009-09-03	Diehl et al.	
	21	20100107240	A1	2010-04-29	Thaler et al.	
	22	20040093513	A1	2004-05-13	Cantrell et al.	
	23	20120023576	A1	2012-01-26	Sorensen et al.	
	24	20080163333	A1	2008-07-03	Kasralikar	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /C.H./

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2]	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	2016 Apr 15 - (US) Notice of Allowance - App 14/855,374	
	2	2015 Nov 02 - (AU) Office Action - App 2013372879	
	3	2016 Apr 26 - (US) Office Action - App 14/745,207	
	4	2016 May 06 - (US) Office Action - App 14/714,207	
	5	2016 May 13 - (US) Office Action - App 13/940,240	
	6	2016 Jun 14 - (US) Office Action - App 14/625,486	
	7	2016 Feb 25 - (AU) Office Action - App 2014249055	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	8	2016 Feb 24 - (AU) Office Action - App 2014228257
	9	2016 Jun 09 - (WO) International Search Report - PCT/US2016/026339
	10	2016 Jun 16 - (CA) Office Action - App 2,888,935
	11	2016 Jul 11 - (EP) Office Action - App 14720824.3
	12	2016 Jul 22 - (US) Office Action - App 14/921,718
	13	2016 Jul 20 - (AU) Office Action - App 2013335255

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	11/08/2016

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	14253992	ROGERS ET AL.
	Examiner	Art Unit
	CHENG-FENG HUANG	2497

CPC- SEARCHED		
Symbol	Date	Examiner
H04L63/20 or H04L67/02 or H04L63/0209	10/08/2015	C. H.
H04L63/20 or H04L67/02 or H04L63/0209	11/8/2016	C. H.

CPC COMBINATION SETS - SEARCHED		
Symbol	Date	Examiner

US CLASSIFICATION SEARCHED			
Class	Subclass	Date	Examiner

SEARCH NOTES		
Search Notes	Date	Examiner
EAST Searches (including inventor and assignee search)	10/08/2015	C. H.
EAST classification search (keywords combined with CPC)	10/08/2015	C. H.
NPL Searches	10/08/2015	C. H.
PLUS Search Performed	10/08/2015	C. H.
Update EAST Searches	02/18/2016	C. H.
Update NPL Searches	02/18/2016	C. H.
Above Searches Updated	11/8/2016	C. H.

INTERFERENCE SEARCH			
US Class/ CPC Symbol	US Subclass / CPC Group	Date	Examiner
	General interference and searches of claim (PGPUB,USPAT)	11/8/2016	C. H.

/CHENG-FENG HUANG/ Examiner.Art Unit 2497	/HARUNUR RASHID/ Primary Examiner, Art Unit 2497

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	6098172		2000-08-01	Coss et al.	
	2	7331061	B1	2008-02-12	Ramsey et al.	
	3	7299353	B2	2007-11-20	Le Pennec et al.	
	4	7227842	B1	2007-06-05	Ji et al.	
	5	7107613	B1	2006-09-12	Chen et al.	
	6	7263099	B1	2007-08-28	Woo et al.	
	7	7237267	B2	2007-06-26	Rayes et al.	
	8	7089581	B1	2006-08-08	Nagai et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	9	6662235	B1	2003-12-09	Callis et al.	
	10	6484261	B1	2002-11-19	Wiegel	
	11	8042167	B2	2011-10-18	Fulp et al.	
	12	8037517	B2	2011-10-11	Fulp et al.	
	13	7954143	B2	2011-05-31	Aaron	
	14	8495725	B2	2013-07-23	Ahn	
	15	9124552	B2	2015-09-01	Moore	
	16	9160713	B2	2015-10-13	Moore	
	17	9094445	B2	2015-07-28	Moore et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	1	20060070122	A1	2006-03-30	Bellovin	
	2	20050010765	A1	2005-01-13	Swander et al.	
	3	20020198981	A1	2002-12-26	Corl et al.	
	4	20090138938	A1	2009-05-28	Harrison et al.	
	5	20060104202	A1	2006-05-18	Reiner	
	6	20050229246	A1	2005-10-13	Rajagopal et al.	
	7	20040177139	A1	2004-09-09	Schuba et al.	
	8	20040010712	A1	2004-01-15	Hui et al.	
	9	20030120622	A1	2003-06-26	Nurmela et al.	
	10	20080301765	A1	2008-12-04	Nicol et al.	
	11	20050251570	A1	2005-11-10	Heasman et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	12	20050183140	A1	2005-08-18	Goddard	
	13	20050125697	A1	2005-06-09	Tahara	
	14	20040193943	A1	2004-09-30	Angelino et al.	
	15	20020165949	A1	2002-11-07	Na et al.	
	16	20020038339	A1	2002-03-28	Xu	
	17	20050114704	A1	2005-05-26	Swander	
	18	20100268799	A1	2010-10-21	Maestas	
	19	20110055923	A1	2011-03-03	Thomas	
	20	20100303240	A1	2010-12-02	Beachem et al.	
	21	20100199346	A1	2010-08-05	Ling et al.	
	22	20100011433	A1	2010-01-14	HARRISON et al.	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	23	20080077705	A1	2008-03-27	Li et al.	
	24	20050286522	A1	2005-12-29	Paddon et al.	
	25	20040205360	A1	2004-10-14	Norton et al.	
	26	20140281030	A1	2014-09-18	Cui et al.	
	27	20040073655	A1	2004-04-15	Kan et al.	
	28	20150237012	A1	2015-08-20	Moore	
	29	20150304354	A1	2015-10-22	Rogers et al.	
	30	20140115654	A1	2014-04-24	Rogers et al.	
	31	20140201123	A1	2014-07-17	Ahn et al.	
	32	20140283004	A1	2014-09-18	Moore	
	33	20140283030	A1	2014-09-18	Moore et al.	

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /C.H./

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	20010079361	KR	A	2001-08-22	Kim, Sang Wook		
	2	1006701	EP	A2	2000-06-07	Lucent Technologies Inc		
	3	2006105093	WO	A2	2006-10-05	Univ Wake Forest		
	4	2006093557	WO	A2	2006-09-08	Univ Wake Forest		
	5	2005328336	AU	B2	2011-09-15	Univ Wake Forest		
	6	2011038420	WO	A2	2011-03-31	Great Wall Systems		
	7	1484884	EP	A2	2004-12-08	Microsoft Corp		
	8	2006230171	AU	B2	2012-06-21	Univ Wake Forest		
	9	1864226	EP	B1	2013-05-15	Univ Wake Forest		

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	10	2600236	CA	A1	2006-10-05	Univ Wake Forest		
	11	2012146265	WO	A1	2012-11-01	Voipfuture Gmbh		

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]
	1	2011 Apr 27 - (WO) International Search Report and Written Opinion - App PCT/US2010/054520	
	2	2011 Mar 04 - (US) Notice of Allowance - App 11/316,331	
	3	2011 Mar 03 - (EP) Communication Pursuant to Rules 70(2) and 70a(2) - App 06758213.0	
	4	2011 Feb 14 - (EP) Search Report - App 06758213.0	
	5	Fulp, Errin: "Errin Fulp," XP002618346, www.cs.wfu.edu/fulp/ewfPub.html, pgs. 1-5 (Copyright 2010)	
	6	2010 Sept 30 - (US) Office Action - App 11/390,976	
	7	2010 Sept 10 - (AU) Office Action - App 2006230171	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	8	2010 Aug 20 - (AU) Office Action - App 2005328336	
	9	2010 Jun 23 - (US) Final Rejection - App 11/316,331	
	10	2010 Apr 29 - (US) Interview Summary - App 11/390,976	
	11	2010 Mar 26 - (US) Final Rejection - App 11/390,976	
	12	2009 Sept 14 (US) Office Action - App 11/316,331	
	13	2009 Jun 24 - (US) Office Action - App 11/390,976	
	14	2008 Jul 03 - (WO) Written Opinion of the International Searching Authority - App PCT/US06/11291	
	15	2007 Aug 31 - (EP) Communication Pursuant to Rules 109 and 110 - App 05857614.1	
	16	Acharya et al, "OPTWALL: A Hierarchical Traffic-Aware Firewall," Department of Computer Science, Telecommunications Program, University of Pittsburgh, pgs. 1-11 (2007)	
	17	2006 Sept 11 - (WO) Written Opinion of the International Searching Authority - App PCT/US05/47008	
	18	Tarsa et al., "Balancing Trie-Based Policy representations for Network Firewalls," Department of Computer Science, Wake Forest University, pgs. 1-6 (2006)	

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /C.H./

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	19	Fulp, "Trie-Based Policy Representations for Network Firewalls," Proceedings of the IEEE International Symposium on Computer Communications (2005)	
	20	E. Fulp, "Optimization of Network Firewall Policies Using Ordered Sets and Directed Acyclical Graphs", Technical Report, Computer Scient Department, Wake Forest University, January 2004	
	21	E. Fulp et al., "Network Firewall Policy Tries", Technical Report, Computer Science Department, Wake Forest University, 2004	
	22	E. Al-Shaer et al., "Modeling and Management of Firewall Policies", IEEE Transactions on Network and Service Management, 1(1): 2004	
	23	E.W. Fulp, "Firewall Architectures for High Speed Networks", U.S. Department of Energy Grant Application, Funded September 2003	
	24	E. Al-Shaer et al., "Firewall Policy Advisor for Anomaly Discovery and Rule Editing", Proceedings of the IFIP/IEEE International Symposium on Integrated Network Management, 2003	
	25	V.P. Ranganath, "A Set-Based Approach to Packet Classification", Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems, 889-894, 2003	
	26	M. Christiansen et al., "Using IDDsfor Packet Filtering", Technical Report, BRICS, October 2002	
	27	Lee et al., "Development Framework for Firewall Processors," IEEE, pgs. 352-355 (2002)	
	28	L. Qui et al., "Fast Firewall Implementations for Software and Hardware-Based Routers", Proceedings of ACM Sigmetrics, June 2001	
	29	D. Eppstein et al., "Internet Packet Filter Management and Rectangle Geometry", Proceedings of the Symposium on Discrete Algorithms, 827-835, 2001	

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /C.H./

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	30	E. Fulp, "Preventing Denial of Service Attacks on Quality of Service", Proceedings of the 2001 DARPA Information Survivability Conference and Exposition II, 2001	
	31	S. Goddard et al., "An Unavailability Analysis of Firewall Sandwich Configurations", Proceedings of the 6th IEEE Symposium on High Assurance Systems Engineering, 2001	
	32	G.V. Rooij, "Real Stateful TCP Packet Filtering in IP Filter", Proceedings of the 10th USENIX Security Symposium, 2001	
	33	P. Warkhede et al., "Fast Packet Classification for Two-Dimensional Conflict-Free Filters", Proceedings of IEEE INFOCOM, 1434-1443, 2001	
	34	D. Decasper et al., "Router Plugins: A Software Architecture for Next-Generation Routers", IEEE/ACM Transactions on Networking, 8(1): February 2000	
	35	A. Feldmann et al., "Tradeoffs for Packet Classification", Proceedings of the IEEE INFOCOM, 397-413, 2000	
	36	X. Gan et al., "LSMAC vs. LSNAT: Scalable Cluster-based Web servers", Journal of Networks, Software Tools, and Applications, 3(3): 175-185, 2000	
	37	A. Hari et al., "Detecting and Resolving Packet Filter Conflicts", Proceedings of IEEE INFOCOM, 1203-1212, 2000	
	38	O. Paul et al., "A full Bandwidth ATM Firewall", Proceedings of the 6th European Symposium on Research in Computer Security ESORICS'2000, 2000	
	39	J. Xu et al., "Design and Evaluation of a High-Performance ATM Firewall Switch and Its Applications", IEEE Journal on Selected Areas in Communications, 17(6): 1190-1200, June 1999	
	40	C. Benecke, "A Parallel Packet Screen for High Speed Networks", Proceedings of the 15th Annual Computer Security Applications Conference, 1999	

INFORMATION DISCLOSURE STATEMENT BY APPLICANT
(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	41	R. Funke et al., "Performance Evaluation of Firewalls in Gigabit-Networks", Proceedings of the Symposium on Performance Evaluation of Computer and Telecommunication Systems, 1999	
	42	S. Suri et al., "Packet Filtering in High Speed Networks", Proceedings of the Symposium on Discrete Algorithms, 969-970, 1999	
	43	U. Ellermann et al., "Firewalls for ATM Networks", Proceedings of INFOSEC'COM, 1998	
	44	V. Srinivasan et al., "Fast and Scalable Layer Four Switching", Proceedings of ACM SIGCOMM, 191-202, 1998	
	45	M. Degermark et al., "Small Forwarding Tables for Fast Routing Lookups", Proceedings of ACM SIGCOMM, 4-13, 1997	
	46	S,M. Bellovin et al., "Network Firewalls", IEEE Communications Magazine, 50-57, 1994	
	47	W.E. Leland et al., "On the Self-Similar Nature of Ethernet Traffic", IEEE Transactions on Networking, 2(1); 15, 1994	
	48	G. Brightwell et al., "Counting Linear Extensions is #P-Complete", Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing, 1991	
	49	M. Al-Suwaiyel et al., "Algorithms for Trie Compaction", ACM Transactions on Database Systems, 9(2): 243-263, June 1984	
	50	D. Comer, "Analysis of a Heuristic for Full Trie Minimization", ACM Transactions on Database Systems, 6(3): 513-537, Sept. 1981	

If you wish to add additional non-patent literature document citation information please click the Add button Add

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /C.H./

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

EXAMINER SIGNATURE			
Examiner Signature	/CHENG-FENG HUANG/	Date Considered	11/08/2016

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	14253992	ROGERS ET AL.
	Examiner	**Art Unit**
	CHENG-FENG HUANG	2497

| CPC | | |

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	14253992	ROGERS ET AL.
	Examiner	**Art Unit**
	CHENG-FENG HUANG	2497

US ORIGINAL CLASSIFICATION		INTERNATIONAL CLASSIFICATION										
CLASS	SUBCLASS	CLAIMED					NON-CLAIMED					
726	1	H	0	4	L	29 / 06 (2006.01.01)						

CROSS REFERENCE(S)					
CLASS	SUBCLASS (ONE SUBCLASS PER BLOCK)				

/CHENG-FENG HUANG/ Examiner.Art Unit 2497	11/8/2016	**Total Claims Allowed:**	
(Assistant Examiner)	(Date)	16	
/HARUNUR RASHID/ Primary Examiner.Art Unit 2497	11/09/2016	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	12

Issue Classification	Application/Control No.	Applicant(s)/Patent Under Reexamination
	14253992	ROGERS ET AL.
	Examiner	**Art Unit**
	CHENG-FENG HUANG	2497

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☒ T.D. ☐ R.1.47

Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original
1	1	15	17												
2	2	16	18												
3	3		19												
4	4		20												
5	5		21												
6	6		22												
7	7														
8	8														
9	9														
10	10														
	11														
11	12														
12	13														
13	14														
	15														
14	16														

/CHENG-FENG HUANG/ Examiner.Art Unit 2497 (Assistant Examiner)	11/8/2016 (Date)	**Total Claims Allowed:** 16	
/HARUNUR RASHID/ Primary Examiner.Art Unit 2497 (Primary Examiner)	11/09/2016 (Date)	O.G. Print Claim(s) 1	O.G. Print Figure 12

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

BIB DATA SHEET

CONFIRMATION NO. 5163

SERIAL NUMBER	FILING or 371(c) DATE RULE	CLASS	GROUP ART UNIT	ATTORNEY DOCKET NO.
14/253,992	04/16/2014	726	2497	007742.00029

APPLICANTS

Centripetal Networks, Inc., Reston, VA;

INVENTORS

Steven Rogers, Leesburg, VA;
Sean Moore, Hollis, NH;
David K. Ahn, Winston-Salem, NC;
Peter P. Geremia, Portsmouth, NH;

**** CONTINUING DATA ***************************

**** FOREIGN APPLICATIONS ***************************

**** IF REQUIRED, FOREIGN FILING LICENSE GRANTED **** ** SMALL ENTITY **
05/02/2014

Foreign Priority claimed ☐ Yes ☑ No 35 USC 119(a-d) conditions met ☐ Yes ☑ No Verified and Acknowledged /Cheng-Feng Huang/ Examiner's Signature	☐ Met after Allowance Initials	STATE OR COUNTRY	SHEETS DRAWINGS	TOTAL CLAIMS	INDEPENDENT CLAIMS
		VA	13	~~20~~ 16	~~3~~ 2

ADDRESS

BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051
UNITED STATES

TITLE

METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK

FILING FEE RECEIVED	FEES: Authority has been given in Paper	☐ All Fees
730	No._____ to charge/credit DEPOSIT ACCOUNT	☐ 1.16 Fees (Filing)
	No._____ for following:	☐ 1.17 Fees (Processing Ext. of time)
		☐ 1.18 Fees (Issue)
		☐ Other _____
		☐ Credit

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	14253992	ROGERS ET AL.
	Examiner	**Art Unit**
	CHENG-FENG HUANG	2497

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☒ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	10/09/2015	02/18/2016	11/08/2016						
1	1	÷	✓	=						
2	2	÷	✓	=						
3	3	÷	✓	=						
4	4	÷	✓	=						
5	5	÷	✓	=						
6	6	÷	✓	=						
7	7	÷	✓	=						
8	8	÷	✓	=						
9	9	÷	✓	=						
10	10	÷	✓	=						
	11	÷	✓	-						
11	12	÷	✓	=						
12	13	÷	✓	=						
13	14	÷	✓	=						
	15	÷	✓	-						
14	16	÷	✓	=						
15	17	÷	✓	=						
16	18	÷	✓	=						
	19	÷	-	-						
	20	÷	-	-						
	21			-						
	22			-						

PART B - FEE(S) TRANSMITTAL

Complete and send this form, together with applicable fee(s), to: Mail **Mail Stop ISSUE FEE**
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450
or Fax (571)-273-2885

INSTRUCTIONS: This form should be used for transmitting the ISSUE FEE and PUBLICATION FEE (if required). Blocks 1 through 5 should be completed where appropriate. All further correspondence including the Patent, advance orders and notification of maintenance fees will be mailed to the current correspondence address as indicated unless corrected below or directed otherwise in Block 1, by (a) specifying a new correspondence address; and/or (b) indicating a separate "FEE ADDRESS" for maintenance fee notifications.

CURRENT CORRESPONDENCE ADDRESS (Note: Use Block 1 for any change of address)

BANNER & WITCOFF, LTD.
1100 13th St. NW
STE 1200
Washington, DC 20005

Note: A certificate of mailing can only be used for domestic mailings of the Fee(s) Transmittal. This certificate cannot be used for any other accompanying papers. Each additional paper, such as an assignment or formal drawing, must have its own certificate of mailing or transmission.

Certificate of Mailing or Transmission

I hereby certify that this Fee(s) Transmittal is being deposited with the United States Postal Service with sufficient postage for first class mail in an envelope addressed to the Mail Stop ISSUE FEE address above, or being facsimile transmitted to the USPTO (571) 273-2885, on the date indicated below.

(Depositor's name)
(Signature)
(Date)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
14/253,992	04/16/2014	Steven Rogers	007742.00029\US	5163

TITLE OF INVENTION: Methods and Systems for Protecting a Secured Network

APPLN. TYPE	ENTITY STATUS	ISSUE FEE DUE	PUBLICATION FEE DUE	PREV. PAID ISSUE FEE	TOTAL FEE(S) DUE	DATE DUE
nonprovisional	SMALL	$480.00			$480.00	02/16/2017

EXAMINER	ART UNIT	CLASS-SUBCLASS
C. Huang	2497	726-001000

1. Change of correspondence address or indication of "Fee Address" (37 CFR 1.363).
- ☐ Change of correspondence address (or Change of Correspondence Address form PTO/SB/122) attached.
- ☐ "Fee Address" indication (or "Fee Address" Indication form PTO/SB/47; Rev 03-02 or more recent) attached. **Use of a Customer Number is required.**

2. For printing on the patent front page, list
(1) The names of up to 3 registered patent attorneys or agents OR, alternatively,
(2) The name of a single firm (having as a member a registered attorney or agent) and the names of up to 2 registered patent attorneys or agents. If no name is listed, no name will be printed.

1 Banner & Witcoff, Ltd.
2
3

3. ASSIGNEE NAME AND RESIDENCE DATA TO BE PRINTED ON THE PATENT (print or type)

PLEASE NOTE: Unless an assignee is identified below, no assignee data will appear on the patent. If an assignee is identified below, the document has been filed for recordation as set forth in 37 CFR 3.11. Completion of this form is NOT a substitute for filing an assignment.

(A) NAME OF ASSIGNEE: Centripetal Networks, Inc.

(B) RESIDENCE: (CITY and STATE OR COUNTRY): Reston, Virginia

Please check the appropriate assignee category or categories (will not be printed on the patent) : ☐ Individual ☒ Corporation or other private group entity ☐ Government

4a. The following fee(s) are submitted:
- ☒ Issue Fee
- ☐ Publication Fee (No small entity discount permitted)
- ☐ Advance Order - # of Copies _____

4b. Payment of Fee(s): **(Please first reapply any previously paid issue fee shown above)**
- ☐ A check is enclosed.
- ☐ Payment by credit card. Form PTO-2038 is attached.
- ☒ The Director is hereby authorized to charge the required fee(s), any deficiency, or credit any overpayment, to Deposit Account Number 19-0733 (enclose an extra copy of this form).

5. **Change in Entity Status** (from status indicated above)
- ☐ Applicant certifying micro entity status. See 37 CFR 1.29 — NOTE: Absent a valid certification of Micro Entity Status (see forms PTO/SB/15A and 15B), issue fee payment in the micro entity amount will not be accepted at the risk of application abandonment.
- ☐ Applicant asserting small entity status. See 37 CFR 1.27 — NOTE: If the application was previously under micro entity status, checking this box will be taken to be a notification of loss of entitlement to micro entity status.
- ☐ Applicant changing to regular undiscounted fee status. — NOTE: Checking this box will be taken to be a notification of loss of entitlement to small or micro entity status, as applicable.

NOTE: This form must be signed in accordance with 37 CFR 1.31 and 1.33. See 37 CFR 1.4 for signature requirements and certifications.

Authorized Signature /Bradley C. Wright/ Date December 16, 2016

Typed or printed name Bradley C. Wright Registration No. 38,061

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20110141900	A1	2011-06-16	Jayawardena et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

	1	2016 Oct 05 - (US) Notice of Allowance - App 14/698,560	
	2	2016 Sept 26 - (CA) Office Action - App 2,902,158	
	3	2016 Oct 26 - (US) Office Action - App 13/940,240	
	4	2016 Nov 21 - (US) Office Action - App 14/745,207	
	5	2016 Dec 05 - (US) Notice of Allowance - App 14/714,207	

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature		Date Considered	

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT (Not for submission under 37 CFR 1.99)		
	Application Number	14253992
	Filing Date	2014-04-16
	First Named Inventor	Steven Rogers
	Art Unit	2497
	Examiner Name	C. Huang
	Attorney Docket Number	007742.00029\US

CERTIFICATION STATEMENT

Please see 37 CFR 1.97 and 1.98 to make the appropriate selection(s):

That each item of information contained in the information disclosure statement was first cited in any communication from a foreign patent office in a counterpart foreign application not more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(1).

OR

☒ That no item of information contained in the information disclosure statement was cited in a communication from a foreign patent office in a counterpart foreign application, and, to the knowledge of the person signing the certification after making reasonable inquiry, no item of information contained in the information disclosure statement was known to any individual designated in 37 CFR 1.56(c) more than three months prior to the filing of the information disclosure statement. See 37 CFR 1.97(e)(2).

See attached certification statement.

☒ The fee set forth in 37 CFR 1.17 (p) has been submitted herewith.

A certification statement is not submitted herewith.

SIGNATURE

A signature of the applicant or representative is required in accordance with CFR 1.33, 10.18. Please see CFR 1.4(d) for the form of the signature.

Signature	/Bradley C. Wright/	Date (YYYY-MM-DD)	2016-12-16
Name/Print	Bradley C. Wright	Registration Number	38,061

This collection of information is required by 37 CFR 1.97 and 1.98. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 1 hour to complete, including gathering, preparing and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these record s.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

(PATENT)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:			
Steven Rogers et al.		Atty. Docket No.:	007742.00029
Appl. No.:	14/253,992	Confirmation No.:	5163
Filed:	April 16, 2014	Art Unit:	2497
For:	Methods and Systems For Protecting A Secured Network	Examiner:	Huang, Cheng-Feng

APPLICANT'S COMMENTS ON
STATEMENT OF REASONS FOR ALLOWANCE
UNDER 37 C.F.R. § 1.104(e)

United States Patent and Trademark Office
Customer Service Window, Mail Stop Issue Fee
Randolph Building
401 Dulany Street
Alexandria, VA 22314

Dear Commissioner:

In response to the Notice of Allowance mailed November 16, 2016, Applicant submits the enclosed response to the Examiner's Statement of Reasons for Allowance, included in the Notice of Allowance.

Applicant requests any necessary extension of time for the submission of this paper. If any fees are required or if an overpayment is made, the Commissioner is authorized to debit or credit our Deposit Account No. 19-0733, accordingly.

REMARKS

Applicant thanks the Examiner for the allowance of this application. In brief response to the Statement of Reasons for Allowance from the Notice of Allowance, Applicant respectfully submits that specific excerpts quoted or paraphrased in the Statement are not the full claim language for all the claims, and are not necessarily the only reasons for distinguishing the claims from the various references, as there are many other features in the independent and dependent claims. Also, to the extent there is any deviation between the quoted excerpts or paraphrased language and the language of the claims, Applicant submits that the language of the claims will control for establishing the requirements of the various claims.

Respectfully submitted,

Date: December 16, 2016 By: /Bradley C. Wright/
Bradley C. Wright
Registration No.: 38,061

BANNER & WITCOFF, LTD.
1100 13th St. NW
STE 1200
Washington, DC 20005
(202) 824-3000

BCW/JAL/hp

Electronic Patent Application Fee Transmittal

Application Number:	14253992
Filing Date:	16-Apr-2014
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Filer:	Bradley C. Wright/Holly Parker
Attorney Docket Number:	007742.00029

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
UTILITY APPL ISSUE FEE	2501	1	480	480

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension-of-Time:				
Miscellaneous:				
SUBMISSION- INFORMATION DISCLOSURE STMT	2806	1	90	90
		Total in USD ($)		**570**

Electronic Acknowledgement Receipt

EFS ID:	27824034
Application Number:	14253992
International Application Number:	
Confirmation Number:	5163
Title of Invention:	METHODS AND SYSTEMS FOR PROTECTING A SECURED NETWORK
First Named Inventor/Applicant Name:	Steven Rogers
Customer Number:	22907
Filer:	Bradley C. Wright/Holly Parker
Filer Authorized By:	Bradley C. Wright
Attorney Docket Number:	007742.00029
Receipt Date:	19-DEC-2016
Filing Date:	16-APR-2014
Time Stamp:	09:28:17
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	DA
Payment was successfully received in RAM	$570
RAM confirmation Number	121916INTEFSW00007350190733
Deposit Account	190733
Authorized User	Bradley Wright

The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:

37 CFR 1.17 (Patent application and reexamination processing fees)

37 CFR 1.19 (Document supply fees)

37 CFR 1.20 (Post Issuance fees)

37 CFR 1.21 (Miscellaneous fees and charges)

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Issue Fee Payment (PTO-85B)	Issue_Fee_Transmittal_18179481.pdf	33499 58d2be8b3ba44f23602a60e63a6d44e34931e579	no	1
Warnings:					
Information:					
2	Non Patent Literature	Notice_of_Allowance_and_Fees_Due_PTOL-85__INTVSUMEX__892__1449_18076047.pdf	1537046 d8a8ad44e3593cdfa76b2c0fdf963036e70bdeb1	no	38
Warnings:					
Information:					
3	Non Patent Literature	OTT01_-7861591-v1-PAT_83575W-1_Corr_from_PTO_-_Office_Action_18101654.PDF	285929 c87d71516445919f7dfca6a0e1dd62aa55e1463a	no	5
Warnings:					
Information:					
4	Non Patent Literature	18105700_1_18116178.pdf	1041266 9b9dace0d9e37e051b27e2662c18d1d91f311243	no	30
Warnings:					
Information:					
5	Non Patent Literature	Non-Final_Rejection__892__1449_18140604.pdf	1584757 a988730b4239ed429a4344dc8b1ddd5c65372a63	no	42
Warnings:					
Information:					
6	Non Patent Literature	Notice_of_Allowance_and_Fees_Due_PTOL-85__1449_18158549.pdf	695760 461671807412c0444aa64839e61ec9579560a367	no	14
Warnings:					
Information:					

7	Information Disclosure Statement (IDS) Form (SB08)	Information_Disclosure_Statement_Fillable_PDF_18177940.pdf	1035067 e8d4c966866f25efddf510cc416ed179c963ea03	no	4
Warnings:					
Information:					
8	Miscellaneous Incoming Letter	DRAFT_Comments_on_Statement_of_Reasons_for_Allowance_18179487.pdf	20940 e26e89b6e1d765d5de91e2c0a5d0331a3fbc185d	no	2
Warnings:					
Information:					
9	Fee Worksheet (SB06)	fee-info.pdf	32274 0ec59e8a07c802460e9add1ef0f25ae0953aa0a6	no	2
Warnings:					
Information:					
			Total Files Size (in bytes):	6266538	

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111

If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371

If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office

If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/08a (03-15)
Approved for use through 07/31/2016. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1					

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	20110141900	A1	2011-06-16	Jayawardena et al.	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1							

If you wish to add additional Foreign Patent Document citation information please click the Add button Add

NON-PATENT LITERATURE DOCUMENTS Remove

Examiner Initials*	Cite No	Include name of the author (in CAPITAL LETTERS), title of the article (when appropriate), title of the item (book, magazine, journal, serial, symposium, catalog, etc), date, pages(s), volume-issue number(s), publisher, city and/or country where published.	T[5]

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	2497
Examiner Name	C. Huang
Attorney Docket Number	007742.00029\US

	1	2016 Oct 05 - (US) Notice of Allowance - App 14/698,560	
	2	2016 Sept 26 - (CA) Office Action - App 2,902,158	
	3	2016 Oct 26 - (US) Office Action - App 13/940,240	
	4	2016 Nov 21 - (US) Office Action - App 14/745,207	
	5	2016 Dec 05 - (US) Notice of Allowance - App 14/714,207	

If you wish to add additional non-patent literature document citation information please click the Add button Add

EXAMINER SIGNATURE

Examiner Signature	/CHENG-FENG HUANG/	Date Considered	12/28/2016

*EXAMINER: Initial if reference considered, whether or not citation is in conformance with MPEP 609. Draw line through a citation if not in conformance and not considered. Include copy of this form with next communication to applicant.

[1] See Kind Codes of USPTO Patent Documents at www.USPTO.GOV or MPEP 901.04. [2] Enter office that issued the document, by the two-letter code (WIPO Standard ST.3). [3] For Japanese patent documents, the indication of the year of the reign of the Emperor must precede the serial number of the patent document. [4] Kind of document by the appropriate symbols as indicated on the document under WIPO Standard ST.16 if possible. [5] Applicant is to place a check mark here if English language translation is attached.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
14/253,992	04/16/2014	Steven Rogers	007742.00029	5163

22907 7590 01/04/2017

BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051

EXAMINER
HUANG, CHENG-FENG

ART UNIT	PAPER NUMBER
2497	

NOTIFICATION DATE	DELIVERY MODE
01/04/2017	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

eofficeaction@bannerwitcoff.com
GPD@bannerwitcoff.com

Corrected Notice of Allowability	Application No. 14/253,992	Applicant(s) ROGERS ET AL.	
	Examiner CHENG-FENG HUANG	Art Unit 2497	AIA (First Inventor to File) Status Yes

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☒ This communication is responsive to *12/19/2016*.
   ☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on_____.
2. ☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.
3. ☒ The allowed claim(s) is/are *1-10, 12-14 and 16-18*. As a result of the allowed claim(s), you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to PPHfeedback@uspto.gov.
4. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

   Certified copies:

   a) ☐ All b) ☐ Some *c) ☐ None of the:
   1. ☐ Certified copies of the priority documents have been received.
   2. ☐ Certified copies of the priority documents have been received in Application No. _____ .
   3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

   * Certified copies not received: _____.

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
   ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____.

   Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).
6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

1. ☐ Notice of References Cited (PTO-892)
2. ☒ Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date 12/19/2016
3. ☐ Examiner's Comment Regarding Requirement for Deposit of Biological Material
4. ☐ Interview Summary (PTO-413), Paper No./Mail Date _____ .
5. ☐ Examiner's Amendment/Comment
6. ☐ Examiner's Statement of Reasons for Allowance
7. ☐ Other _____.

/HARUNUR RASHID/ Primary Examiner, Art Unit 2497	/CHENG-FENG HUANG/ Examiner, Art Unit 2497

PTO/SB/08a (01-10)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

INFORMATION DISCLOSURE STATEMENT BY APPLICANT

(Not for submission under 37 CFR 1.99)

Application Number	14253992
Filing Date	2014-04-16
First Named Inventor	Steven Rogers
Art Unit	TBA 2497
Examiner Name	TBA Cheng-Feng Huang
Attorney Docket Number	007742.00029

U.S.PATENTS Remove

Examiner Initial*	Cite No	Patent Number	Kind Code[1]	Issue Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
	1	6226372		2001-05-01	Beebe et al.	

If you wish to add additional U.S. Patent citation information please click the Add button. Add

U.S.PATENT APPLICATION PUBLICATIONS Remove

Examiner Initial*	Cite No	Publication Number	Kind Code[1]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear
Change(s) applied to document, /C.L./ 12/5/2016	1	20060195896 ~~US20060041058~~		2006-08-31	Fulp et al.	
	2	20060248580 ~~US20060092485~~		2006-11-02	Fulp et al.	
	3	20110055916		2011-03-03	Ahn	

If you wish to add additional U.S. Published Application citation information please click the Add button. Add

FOREIGN PATENT DOCUMENTS Remove

Examiner Initial*	Cite No	Foreign Document Number[3]	Country Code[2] i	Kind Code[4]	Publication Date	Name of Patentee or Applicant of cited Document	Pages,Columns,Lines where Relevant Passages or Relevant Figures Appear	T[5]
	1	1677484	EP	A2	2006-07-05	Microsoft Corp [us]		☐

ALL REFERENCES CONSIDERED EXCEPT WHERE LINED THROUGH. /CH/

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	ISSUE DATE	PATENT NO.	ATTORNEY DOCKET NO.	CONFIRMATION NO.
14/253,992	02/07/2017	9565213	007742.00029	5163

22907 7590 01/18/2017

BANNER & WITCOFF, LTD.
1100 13th STREET, N.W.
SUITE 1200
WASHINGTON, DC 20005-4051

ISSUE NOTIFICATION

The projected patent number and issue date are specified above.

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(application filed on or after May 29, 2000)

The Patent Term Adjustment is 0 day(s). Any patent to issue from the above-identified application will include an indication of the adjustment on the front page.

If a Continued Prosecution Application (CPA) was filed in the above-identified application, the filing date that determines Patent Term Adjustment is the filing date of the most recent CPA.

Applicant will be able to obtain more detailed information by accessing the Patent Application Information Retrieval (PAIR) WEB site (http://pair.uspto.gov).

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Application Assistance Unit (AAU) of the Office of Data Management (ODM) at (571)-272-4200.

APPLICANT(s) (Please see PAIR WEB site http://pair.uspto.gov for additional applicants):

Steven Rogers, Leesburg, VA;
Centripetal Networks, Inc., Reston, VA;
Sean Moore, Hollis, NH;
David K. Ahn, Winston-Salem, NC;
Peter P. Geremia, Portsmouth, NH;

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The USA offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to encourage and facilitate business investment. To learn more about why the USA is the best country in the world to develop technology, manufacture products, and grow your business, visit SelectUSA.gov.