IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Palo Alto Networks, Inc.,
Fortinet, Inc. and
WatchGuard Technologies, Inc.
Petitioners,

v.

Selective Signals, LLC
Patent Owner.

U.S. Patent No. 8,111,629
Issue Date: Feb. 7, 2012
Title: Media Session Identification Method for IP Networks

Inter Partes Review No.: Unassigned

DECLARATION OF SAMUEL H. RUSS IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 8,111,629

***Mail Stop* "*PATENT BOARD*"**
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

Page

I. Introduction

1. I am over the age of eighteen (18) and otherwise competent to make this declaration.

2. I have been retained by Palo Alto Networks (Petitioner) as an independent expert consultant in this proceeding before the United States Patent and Trademark Office. Although I am being compensated at my rate of $445 per hour for the time I spend on this matter, no part of my compensation depends on the outcome of this proceeding, and I have no other interest in this proceeding. To the best of my knowledge, I have no financial interest in Palo Alto Networks.

3. This declaration is in support of a Petition for *inter partes* review involving U.S. Patent No. 8,111,629 ("the '629 Patent") (PAN-1001). The '629 Patent is entitled "Media Session Identification Method for IP Networks" and lists Avi Oron as the inventor.

4. For the purposes of this *inter partes* review as discussed later, I have been instructed to assume that the effective filing date of the Claims of the '629 Patent challenged by the Petitioner in this *inter partes* review is no earlier than Nov. 6, 2006, the filing date of U.S. Provisional Patent Application No. Nov. 6, 2006.

5. I understand that according to USPTO records, the ‘629 Patent is currently assigned to Selective Signals, LLC (“Patent Owner”).

6. The ‘629 Patent is directed to identification of packets and packet sessions. I am familiar with the technology described in the ‘629 Patent as of the earliest possible priority date of Nov. 6, 2006.

7. In preparing this Declaration, I have reviewed the ‘629 Patent (PAN-1001), the file history of the ‘629 Patent (PAN-1002), and each of the documents cited herein, and I have considered these documents in light of the general knowledge in the art as of Nov. 6, 2006. In formulating my opinions, I have relied upon my experience in the relevant art. I have also considered the viewpoint of a person of ordinary skill in the art (“POSA”) in the field, as of Nov. 6, 2006.

8. I have been asked to provide my technical expertise, analysis, insights and opinions regarding the ‘629 Patent and relevant references that form the basis of the grounds of rejection set forth in the accompanying Petition for *inter partes* review of the ‘629 Patent. As described in detail below, I offer the following opinion in this Declaration:

- A POSA would have found Claims 15, 17, 18, 19, 20, 21 and 22 of the ‘629 Patent to be obvious over U.S. Pat No. 8,085,775 to Pappu issued on Dec. 27, 2011 from an application filed on Jul 31, 2006

("Pappu") alone, or in view of U.S. Pat Pub. No. 2006/0262789 published on November 23, 2006 from an application filed May 19, 2006 ("Peleg"). Pappu alone, or in view of Peleg teaches each element of Claims 15, 17, 18, 19, 20, 21 and 22 to a POSA and a POSA would have been motivated to combine the teachings of these references.

II. Materials Reviewed and Relied Upon

9. In formulating my opinions, I have considered and relied on statements in the documents identified below. These documents include patents, patent Applications, learned treatises, periodicals, pamphlets and other publications. I consider each of the references below as a reliable authority for the statements on which I rely.

Exhibit #	***Description***
1001	U.S. Patent No. 8,111,629 ("the '629 Patent")
1002	Select portions of prosecution history of the '629 Patent ("File History")
1004	U.S. Pat. No. 8,085,775 ("Pappu")
1005	U.S. Pat. No. 7,660,248 ("Duffield")
1006	U.S. Pat. Publication No. 20060262789 ("Peleg")
1007	U.S. Pat. Publication No. 20070076606 ("Olesinski")

Exhibit #	*Description*
1008	"Toward the Accurate Identification of Network Applications" ("Moore")
1009	"Automated Traffic Classification and Application Identification Using Machine Learning" ("Zander")
1010	"Evaluating Machine Learning Methods for Online Game Traffic Identification" ("Williams")

10. An additional basis for my opinions is my own experience as an engineering manager and as a staff expert while at Scientific-Atlanta. As noted above, I have managed the design of four electronic products that entered mass production (three set-top boxes, the Explorer 1840, 4200, and 8300, and one point-of-sale retail terminal, the e^{N}-Touch 1000). Two of those products (the 4200 and 8300) were produced in million-unit quantities. I also became a staff expert in home networking and DVR reliability and performance, as well as all aspects of design of computer-based systems, embedded systems, electronics manufacturing, and design for low product cost.

11. I have also relied on years of education, teaching, research, and experience concerning software, circuit design, signal integrity, computer architecture, digital logic design and synthesis, embedded systems, distributed computing, and computer design.

III. My Background and Qualifications

12. My qualifications for forming the opinions set forth in this declaration are summarized here and explained in more detail in my *curriculum vitae*, which is attached as part of Appendix A. Appendix A also includes a list of my publications and the cases in which I have testified at deposition, hearing, or trial during the past four years.

13. I received a Bachelor's degree in Electrical Engineering from the Georgia Institute of Technology ("Georgia Tech") in 1986 and a Ph.D. in Electrical Engineering from Georgia Tech in 1991.

14. From 2007 to the present, I have been a member of the faculty of the University of South Alabama as an Assistant and Associate Professor in the Department of Electrical and Computer Engineering. During that time, I have won awards for excellent teaching and have been actively publishing research in home networking and digital video recording (DVR) technologies. I am active in the Institute of Electrical and Electronic Engineers (IEEE) and am a Distinguished Lecturer for the IEEE Consumer Electronics Society. As a consultant, I have conducted briefings for members of the financial community on technology trends in the cable, satellite, and IPTV sectors.

15. From 2000 to 2007, I worked for Scientific-Atlanta (now Cisco's Service Provider Video Tech. Group), where I managed a cable set-top box (STB) design group that designed four STB models, including the Explorer 4200 (non-DVR) and 8300 (DVR) models. Both models sold several million units. As design-group manager, I was responsible for managing the design and prototyping activities of the group and for interfacing with other groups (especially integrated-circuit design, procurement, software developers, the factory where prototypes

were built, and product managers) and for maintaining the hardware and mechanical development schedule. Since the products were produced in extremely high volumes, the projects had very high visibility in the company, and therefore carried a great deal of responsibility.

16. Also while at Scientific-Atlanta, I became a staff expert in home networking, conducting demonstrations of wireless video technology and managing a group that developed a new coaxial home networking system. The coaxial system won a Technology and Engineering Emmy® Award in 2013. I became a staff expert in DVR reliability, and led a team that improved the software, hardware, repair, and manufacturing processes. I am a named inventor on fifty-one (51) patent applications that were filed while I was at Scientific-Atlanta, twenty eight (28) of which have issued as U.S. patents as of the writing of this report.

17. From 1999 to 2000, I was a Staff Electrical Engineer and then Matrix Manager at IVI Checkmate (now Ingenico), where I managed the hardware design team that completed the design of the e^N-Touch 1000 payment terminal. This terminal was in widespread use, for example, at the self-checkout at Home Depot.

18. I also served on the faculty of Mississippi State University from 1994 to 1999 as an Assistant Professor in the Department of Electrical & Computer

Engineering where I taught circuit board design and two-way interactive video classes, among other things.

19. I have also authored 32 journal articles and conference papers. A recent conference paper on digital video recording won second place in a "best paper" competition at the 2011 International Conference on Consumer Electronics in Las Vegas, NV.

20. My curriculum vitae, which includes a more detailed summary of my background and experience, as well as a complete list of my publication, is attached as Appendix A.

IV. Person of Ordinary Skill in the Art (POSA)

21. I am familiar with the knowledge and capabilities of one of ordinary skill in the art. Unless otherwise stated, my testimony below refers to the knowledge of one of ordinary skill in the art as of Nov. 6, 2006, the earliest possible effective filing date of the '629 Patent.

22. I have been informed and understand that a Person of Ordinary Skill in the Art ("POSA") is a hypothetical person who is presumed to be aware of all pertinent prior art, thinks along conventional wisdom in the art, and is a person of ordinary creativity.

23. With respect to the '629 Patent, a POSA in November 2006 would have been familiar with techniques for allocating network resources to packets. A POSA would have a working knowledge of identifying, classifying, and controlling media sessions, including those based on the IP protocol, inspection techniques, and provisioning. A POSA would have gained knowledge of these concepts through a mixture of training and work experience, such as by having a Bachelor's degree in computer science, computer communications, electrical engineering, or equivalent degree, coupled with two years of experience; or by obtaining a Master's degree in computer science, computer communications, or electrical engineering, but having no experience; or by having no formal education but experience in computer science or computer communications of at least four years.

24. A POSA may work as part of a multi-disciplinary team and draw upon not only his or her own skills, but also take advantage of certain specialized skills of others on the team to solve a given problem.

V. Relevant Legal Standards

25. I am not a lawyer and will not provide any legal opinions. Although I am not a lawyer, I have been informed and understand that certain legal standards are to be applied by technical experts in forming opinions regarding the meaning

and validity of patent claims. I have been asked to provide my opinions regarding whether the claims of the '629 Patent are anticipated or would have been obvious to a person having ordinary skill in the art at the time of the alleged invention, in light of the prior art.

26. I have been informed and understand that, to anticipate a claim under 35 U.S.C. § 102, a reference must teach every element of the claim either expressly or inherently to a person having ordinary skill in the relevant art.

27. Further, I have been informed and understand that a patent claim is not patentable under 35 U.S.C. § 103 if the differences between the patent claim and the prior art are such that the claimed subject matter as a whole would have been obvious at the time the claimed invention was made to a person having ordinary skill in the relevant art. Obviousness, as I have been informed and understand, is based on the scope and content of the prior art, the differences between the prior art and the claim, the level of ordinary skill in the art, and, to the extent that they exist, certain objective indicia of non-obviousness.

28. I understand that objective indicia can be important evidence regarding whether a patent is obvious or nonobvious, if it has an appropriate nexus to the claimed invention, i.e., is a result of the merits of a claimed invention (rather than the result of design needs or market-pressure advertising or similar activities).

Such indicia include: commercial success of products covered by the patent claims; a long-felt need for the invention; failed attempts by others to make the invention; copying of the invention by others in the field; unexpected results achieved by the invention as compared to the closest prior art; praise of the invention by the infringer or others in the field; the taking of licenses under the patent by others; expressions of surprise by experts and those skilled in the art at the making of the invention; and the patentee proceeded contrary to the accepted wisdom of the prior art.

29. I have been informed that whether there are any relevant differences between the prior art and the claimed invention is to be analyzed from the view of a person of ordinary skill in the relevant art at the time of the invention. As such, my opinions below as to a person of ordinary skill in the art are as of the time of the invention, even if not expressly stated as such; for example, even if stated in the present tense.

30. In analyzing the relevance of the differences between the claimed invention and the prior art, I have been informed that I must consider the impact, if any, of such differences on the obviousness or non-obviousness of the invention as a whole, not merely some portion of it. The person of ordinary skill faced with a

problem is able to apply his or her experience and ability to solve the problem and also look to any available prior art to help solve the problem.

31. I have been informed that a precise teaching in the prior art directed to the subject matter of the claimed invention is not needed. I have been informed that one may take into account the inferences and creative steps that a person of ordinary skill in the art would have employed in reviewing the prior art at the time of the invention. For example, if the claimed invention combined elements known in the prior art and the combination yielded results that were predictable to a person of ordinary skill in the art at the time of the invention, then this evidence would make it more likely that the claim was obvious. On the other hand, if the combination of known elements yielded unexpected or unpredictable results, or if the prior art teaches away from combining the known elements, then this evidence would make it more likely that the claim that successfully combined those elements was not obvious.

32. I have been informed and understand that there are recognized, exemplary, rationales for combining or modifying references to show obviousness of claimed subject matter. Some of the rationales include the following: combining prior art elements according to known methods to yield predictable results; simple substitution of one known element for another to yield predictable

results; use of a known technique to improve a similar device (method or product) in the same way; applying a known technique to a known device (method or product) ready for improvement to yield predictable results; choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations are predictable to one of ordinary skill in the art; and some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art teachings to arrive at the claimed invention.

VI. The '629 Patent

A. Summary

33. The '629 Patent is generally directed to a media session identification method that identifies media sessions from the flow of packets and assigns specific packets to the media sessions. *PAN-1001* 1:16-22. Specifically, the packets are analyzed to determine the packet traffic characteristics of each packet, the packets having similar traffic characteristics are grouped together, and the grouped packets are analyzed for session characteristics to thereby identify a session type, which is used to allocate network resources to the session. *PAN-1001* Abstract.

1. Problems addressed by the ’629 Patent

34. The ’629 Patent teaches that the continuous increase in IP based media traffic over the Internet and other networks has resulted in the need for IP traffic control tools to cope with increasingly congested networks. *PAN-1001* 1:23-28. A media session, such as a voice or video session, is expected to have a continuous stream of sampled data so it can be reconstructed at its destination to produce smooth sound or images or both. *PAN-1001* 1:36-39. As media at either end of a network is sampled and disassembled or reassembled by symmetrical session CODECs (coder & decoder), data streaming must comply with certain bandwidth demands and packet-rates so that the human ear or eye at the destination does not experience any information gaps or distortions. *PAN-1001* 1:40-45.

35. The ’629 Patent explains that IP traffic control tools apply specific provisioning, quality of service, and control actions based on the type of media session. For example, applications that are real-time sensitive such as Voice-over-IP and/or Video-over-IP tend also to be bandwidth availability sensitive. Therefore, maintaining the quality of service for such application sessions requires their identification, and possibly their classification. *PAN-1001* 1:50-55. By way of another example, the ’629 Patent explains that it may also be desirable to identify a session belonging to a particular application in order to implement and meet

interception requirements by law enforcement authorities, as well as to generate call detail records (CDRs) for billing and analysis purposes, and the like. *PAN-1001* 1:55-59.

2. Known prior art tools directed to the problem

36. The ’629 Patent describes known IP traffic control tools that address the provisioning, quality of service, and control actions, including various methods for analyzing IP traffic. The ’629 Patent explains that prior art technologies for analyzing and identifying IP traffic (including IP media sessions) include packet header analysis. *PAN-1001* 2:4-8. In packet header analysis, as described in the ‘629 patent, packets whose headers indicated the same destination were assumed to belong to the same session. *Id.*

37. The ’629 Patent also identifies the prior art techniques known as “Deep Packet Inspection” or “Packet Content Analysis.” These known techniques required digging into and analyzing the content, or payload, of the packets, with the aim of extrapolating signatures or fingerprints of the payloads. The signatures were then used to identify a specific protocol or application. Typically, the identification process included analysis of additional factors from the packet header (e.g., addresses and port numbers). *PAN-1001* 2:12-20. The ’629 Patent identifies several problems associated with the payload inspection techniques at the

time. First, packet payloads could change based on changes in the underlying application, such that a signature or fingerprint that had been correlated to a specific protocol or application could be rendered invalid upon the alteration of the payload due to the release of a new version of the application. *PAN-1001* 2:23-27. Second, the encryption of packets made inspection of the payload difficult and prevented obtaining a meaningful signature. *PAN-1001* 2:28-31.

38. The '629 Patent also identifies a widely used prior art technique called Complete Packet Inspection ("CPI") that overcame the two problems identified above. *PAN-1001* 2:4-8. CPI combined header analysis and deep packet inspection, otherwise known as "payload pattern search." The header analysis provided additional information that complemented the payload pattern search. *PAN-1001* 2:37-46. However, any techniques that required payload inspection required heavy processing power as well as excessive memory resources to implement. *PAN-1001* 2:47-68.

39. The '629 Patent also explained that it was *known* to use the observed behavior of packets to identify suspicious behavior:

> Some existing data security systems detect and prevent intrusion using simple predefined behavioristic rules and thresholds to alert for abnormal traffic behavior which may indicate possible attacks on the network

> (customarily implemented in Firewall technologies). Such suspicious behavior may include indications of denial of service (DoS) attacks.

PAN-1001 3:3-8.

3. Solution proposed by the ’629 Patent

40. The ’629 Patent purports to avoid the problems of the prior art by using the “regularity and similarity characteristics of packet traffic” to identify the packets belonging to a given session, including the use of plural header characteristics such as source address and destination address. *PAN-1001* 3:23-24, 6:8-17. The subsequent identification of the type of session permits the session to be provisioned and/or controlled (e.g., allocating appropriate resources or not allocating any resources at all, as required). *PAN-1001* 3:25-29. The ’629 Patent describes Fig. 6 as a simplified flow chart illustrating preferred embodiments, which includes categorizing passing packets using external packet characteristics; grouping those packets having similar values of characteristics, analyzing those grouped packets based on session characteristics, and assigning resources. *PAN-1001* 10:48-64, Fig. 6.

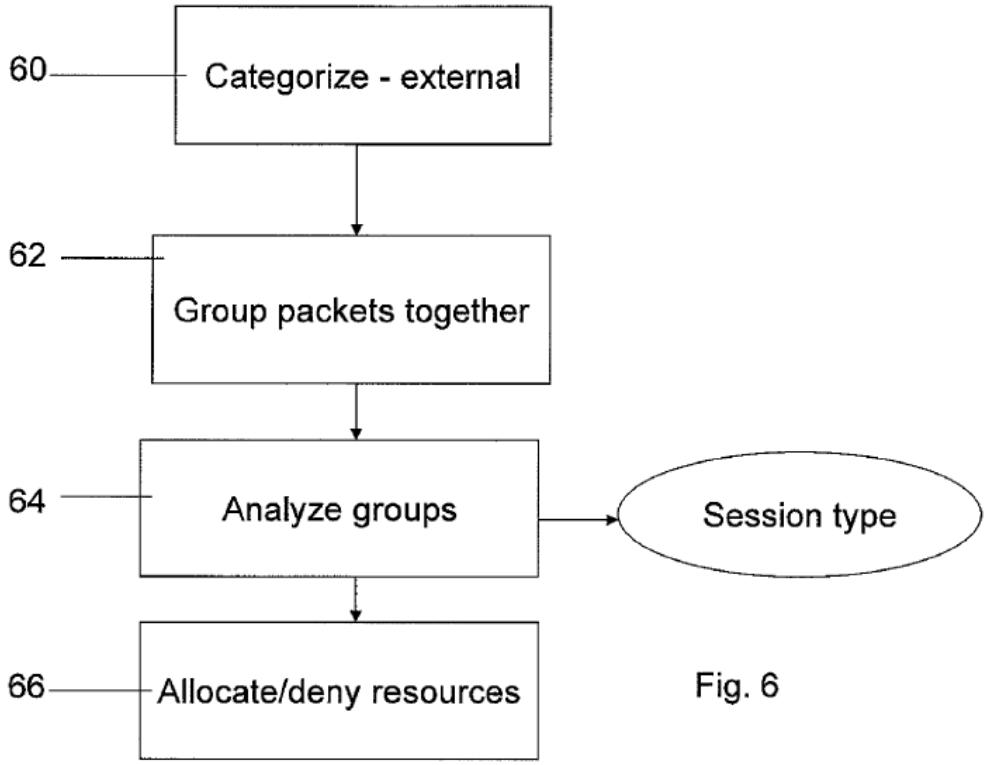

Fig. 6

PAN-1001 Fig. 6.

a. <u>The use of observed traffic characteristics</u>

41. The ’629 Patent explains that traffic characteristics “relate to those parameters inherent to packets traffic flow as it travels from source to destination” and that they include “packet length, inter-arrival period, bandwidth and statistical and other derivations thereof as well as certain header characteristics such as source address and destination address.” *PAN-1001* 6:11-17. The ’629 Patent further explains that traffic characteristics “exclude packet payload content whose inspection is common in deep packet inspection (packet content inspection) techniques.” *PAN-1001* 6:17-19. The traffic characteristics are used to identify

the session type so that appropriate resources can be made available (or not) for the session, or so that other provisioning and control actions may be applied to the session. *PAN-1001* 6:19-23. Thus, the '629 Patent describes grouping like packets together, without performing content inspection of the payload and with the possible use of "header characteristics such as source address and destination address." *PAN-1001* 6:16-17.

b. Grouping packets using traffic characteristics

42. The '629 Patent describes that in one embodiment, "a packet analyzer unit 12 and a session analyzer unit 14" work with a "resource assignment unit 16" to allocate network resources "to the identified media session so that individual packets are given a priority level according to the needs of the media session to which they are found to belong." *PAN-1001* 6:36-42.

43. The packet analyzer 12 analyzes incoming data packets and determines which packets belong to a common session. *PAN-1001* 6:43-44. The packet analyzer 12 includes a packet characteristic extractor 18, thresholder 20, session identifier 22 and marker 24. *PAN-1001* Fig. 1.

44. The '629 Patent explains that packet extractor 18 "extracts readily available packet traffic characteristics such as packet length or inter-arrival period, or source or destination addresses or other traffic characteristics of the packets."

PAN-1001 6:46-49. Thus, in one aspect, the packet extractor relies on the contents of packet headers, which are readily obtained, rather than the contents of the payload. Thresholder 20 sets up thresholds based on the characteristics found by the extractor. Generally, there is a certain amount of uniformity in packets belonging to the same session, thus they tend to be approximately the same length and to turn up at approximately regular intervals. *PAN-1001* 6:56-64. Packet Session identifier 22 uses these thresholds to identify packets belonging to a common session. *PAN-1001* 6:64-66. Packet marker 24 may insert an indication of the common session with which a given packet has been identified into its header. Thus packets believed to belong to the same session can be grouped together so that the session itself can be analyzed. *PAN-1001* 7:3-7. A POSA would understand that a network header is of a fixed length, and so would understand that the process of inserting an indication must involve altering a field in the header, and not the insertion of extra or additional data.

c. <u>Identifying session type using observed characteristics of the group</u>

45. The ’629 Patent explains that session analyzer 14 analyzes the grouped packets to obtain characteristics that could provide important information about the session type and its requirements. For example, the session could be

identified as a video session and therefore have a real-time quality of service requirement. *PAN-1001* 7:8-14.

46. The '629 Patent explains that the session analyzer may use the packet's traffic characteristics, such as packet length and/or inter-arrival period to indicate the session type. *PAN-1001* 7:23-28. Specifically, as explained in connection with FIG. 2, which illustrates "a more detailed embodiment of the apparatus of FIG. 1, "[p]acket length, inter-arrival period ('dt') and their derivatives may be accumulated" and used to derive "statistical parameters such as [a]verage [inter-arrival period], [a]verage length" and their variances and covariances. *PAN-1001* 8:28-30, 8:58-63. These derivatives can then be used to determine and categorize sessions. *PAN-1001* 8:66-9:2.

47. The session characteristics allow a session type to be identified. The session analyzer uses at least one member of the following information to identify a session type:

packet length;

inter-arrival period;

variance of inter-arrival period;

variance of length;

covariance of inter-arrival period;

covariance of length;

a statistical derivative obtained from multi records of packet length;

a statistical derivative obtained from multi records of inter-arrival period.

PAN-1001 4:22-34.

48. Thus, the session characteristics include statistics derived from the traffic characteristics of the packets in a session, e.g., average of packet length, average inter-arrival period, variance of packet length, etc.

d. Allocation of resources

49. The ’629 Patent teaches that once a session type is identified, network resources can be allocated appropriately:

> It is desirable to identify packets as belonging to a session to group the succession of IP packets and also to identify a particular application or type that the session belongs to in order to apply provisioning, quality of service and control actions to the session or stream; applications that are real time sensitive such as Voice-over-IP and/or Video-over-IP tend also to be bandwidth availability sensitive, therefore maintaining the quality of service for such application sessions require the identification, and possibly the classification, of their traffic streams. It may also be desirable to identify a session or a stream belonging to a particular application in order to implement and meet interception requirements by law enforcement authorities, as well as in the need to

> generate CDRs for billing and analysis purposes, and the like. In other cases, some applications may be perceived as prohibited or posing a security risk, that is creating `security holes`, and some networks may be interested in blocking such applications. Again, the best way to do so is to identify the session and the packets belonging to the session and then block packets belonging to the session. For example, many organizations treat Instant-Messaging (IM) applications as high security risks which might create a hole through which sensitive data can flow out of an organization. For those organizations, it is highly essential to identify data flow (e.g. file-transfers) within a media session.

PAN-1001 1:46-2:3.

50. The ’629 Patent explains that the described embodiments may be used with a wide variety of applications, such as “Session Border Controllers, firewalls & billing mediation servers.” *PAN-1001* 11:6-8.

B. Priority Date of Claims

51. The ’629 Patent issued as a National Phase Application of PCT/IL2007/001336 having an international filing date of November 1, 2007, which claims the benefit of U.S. Provisional Patent Application No. 60/856,795,

filed on November 6, 2006. For the purpose of this *inter partes* review, I have been instructed to assume that the effective filing date of each of the Challenged Claims is no earlier than November 6, 2006, the earliest claimed priority date.

VII. State of the Art

52. The following section describes the state of the art for allocating network resources to IP packets as of the November 2006, the presumed effective filing date of the '629 Patent. The prior art references, and the discussions of what was known to a POSA, provide the factual support for the general description of the state of the art at the time of the invention, and provide additional motivation to modify or combine the reference. Accordingly, these references should be considered by the Board.

1. Classifying of IP traffic

53. By the mid-2000s, it was known that there was a continuing increase in the number and variety of applications running over the Internet and over enterprise IP networks. The spectrum included interactive (e.g., telnet, instant messaging, games, etc.), bulk data transfer (e.g., ftp, P2P file downloads), corporate; (e.g., Lotus Notes, database transactions), and real-time applications (voice, video streaming, etc.) *PAN-1005* 1:8-15.

54. In some circumstances, network operators, particularly in enterprise networks, desired the ability to support different levels of Quality of Service (QoS) for different types of applications. This desire was driven by (i) the inherently different QoS requirements of different types of applications, e.g., low end-to-end delay for interactive applications, high throughput for file transfer applications, etc.; (ii) the different relative importance of different applications to the enterprise – e.g., Oracle database transactions are considered critical and therefore high priority, while traffic associated with browsing external web sites is generally less important; and (iii) the desire to optimize the usage of their existing network infrastructures under finite capacity and cost constraints, while ensuring good performance for important applications. *PAN-1005* 1:15-29; *PAN-1007* [0002].

55. Thus, in one approach, knowledge of traffic characteristics and/or traffic type could help optimize the usage of the communications network infrastructure employed, and help ensure a desired level of performance for applications/services important to the customers. *PAN-1007* [0003]; *PAN-1008* pp. 1-2; *PAN-1009* p. 1.

2. Classifying IP traffic using port numbers or payload inspection

56. Traditional methods of traffic detection and classification relied on monitoring logical port specifications typically carried in packet headers, or

content or payload inspection. *PAN-1007* [0003]; *PAN-1010* p. 1-2*; PAN-1009* p. 1; *PAN-1008* pp. 1-2.

57. By 2005, it was known that methods that relied solely on port analysis or content analysis were not reliable, due to the dynamic nature of ports and/or use of encryption. *PAN-1005* 3:54-5:5; *PAN-1008* p. 1.

58. For example, a large percentage of the traffic conveyed by communications networks consists of peer-to-peer (P2P) traffic. Because peer-to-peer traffic is conveyed between pairs of customer network nodes, it is not necessary that a well- known logical port be allocated, reserved, and assigned to traffic produced by applications generating peer-to-peer traffic and/or applications retrieving peer-to-peer content. Therefore, known approaches to traffic classification were no longer valid since logical ports are undefined for peer-to-peer applications and/or logical ports may be dynamically allocated as needed, such as in the case of the standard File Transfer Protocol ("FTP") and others. *PAN-1007* [0004].

59. Deep packet inspection techniques, as the name suggests, assumed the availability of unlimited resources to inspect entire packets to perform packet characterization. Therefore, deep packet inspection incurred high processing overheads and was subject to high costs. Deep packet inspection also suffered from

the high complexity associated with the requirement of inspecting packet payloads at high line rates. Thus, deep packet inspection was not suited at all for typical high throughput communications network nodes deployed in communications networks at the time. Deep packet inspection also suffered from a high maintenance overhead as the detection techniques rely on signatures, and peer-to-peer applications, especially, are known for concealing their identities. Indeed, a deep packet inspection detection signature that provides conclusive detection may not work in the future, and another conclusive signature would have to be found and coded into the system. *PAN-1007* [0010].

3. Classifying IP traffic using traffic statistics.

60. To avoid the problems associated with classifying IP traffic using port numbers or payload inspection, prior art techniques included using traffic statistics. One such technique included gathering statistical information from each packet using "flow level" analysis and inferring the type of application from the traffic statistics. *PAN-1005* 1:7-14, 6:28-47. Flow level statistics include packet length, 5-tuple information, and inter-arrival times between packets. *PAN-1005* 6:48-54; 7:16-20. This technique required data collected at a packet level, but then grouped packets into flows. The relative variance of these inter-arrival times was used as a measure of the burstiness (i.e., the number of high-bandwidth transmission over a

short period of time) of a traffic stream. Intraflow/connection features include loss rates, latencies, etc. *PAN-1005* 7:20-24.

61. Another type of traffic classification technique used a table created off-line that included traffic characteristics of flows. The packet flow classification system included a group of classification rules trained off-line on statistical trace traffic flow information for various traffic flows. A packet classifier sampled traffic flows on-line in real-time and compared the sampled traffic flow with the group of off-line trained classification rules to thereby classify the sampled flow. *PAN-1007* [0013]-[0014].

62. Accordingly, acquired knowledge derived from the off-line statistical analysis was used via the trained classification rules to perform on-line, real-time, classification of traffic in a managed communications network. *PAN-1007* [0057].

4. Allocation of resources

63. Based on the classification of the traffic flow, specific actions can be taken, including marking packets of the sampled flows, assignment of packet processing priorities, assignment of sampled flow packets to specific classes-of-service, special processing of different traffic types, precise monitoring of traffic, etc.; wherein special packet processing could include discarding packets of the sampled flow. *PAN-1007* [0052].

VIII. Claim Construction

64. I have been informed and understand that in order to properly evaluate the '629 Patent, the terms of the claims must first be interpreted. I have also been informed and understand that during an *inter partes* review proceeding in which the target patent has not expired, the terms of the claims of the '629 Patent are to be given their broadest reasonable interpretation (BRI) that is consistent with their plain meaning in light of the specification. The plain meaning is the ordinary and customary meaning to those skilled in the art.

IX. Overview

65. The Challenged Claims are unpatentable as anticipated and/or obvious over the prior art. The Challenged Claims are directed to prior art systems and combinations of conventional components that perform conventional functions regarding identifying and classifying network traffic.

Independent Claim 15 recites:

A method of identifying session type, comprising

obtaining passing packets of respectively unknown sessions and unknown session types;

obtaining traffic packet characteristics of said passing packets of respectively unknown session types;

comparing said obtained packets with each other using respectively obtained traffic packet characteristics;

grouping together those packets having similar values of said traffic packet characteristics into a presumed session; and

analyzing said grouped packets of said presumed session for session characteristic;

using said session characteristics to identify a session type of said presumed session.

66. During prosecution, in distinguishing from the prior art, the Applicant characterized Claim 15 and others as describing:

(1) "that the packets are of unknown session;"

(2) "that the traffic characteristics[1] are extracted from the packets;"

[1] Traffic characteristics are described by the '629 Patent as including header information: "Traffic characteristics relate to those parameters inherent to packets traffic flow as it travels from source to destination and includes features such as packet length, inter-arrival period, bandwidth and statistical and other derivations

(3) “that the packets are grouped into presumed sessions based on similarity of the traffic characteristics,” and,

(4) “[f]ollowing the grouping, the characteristics of the presumed session are then studied to find out a type of this presumed group.”

PAN-1002 response to Office Action of June 15, 2011 page 11.

67. However, at the time of the filing of the ’629 Patent, it was well known to use traffic characteristics to group together packets having similar characteristics, and to analyze characteristics of the grouped packets to determine the type of session in order to, e.g., allocate appropriate resources such as bandwidth. U.S. Pat. No. 8,085,775 (“Pappu”) and U.S. Pat. Publication No. 20060262789 (“Peleg”) are such prior references that disclose these features.

68. Pappu, as part of its process for “identifying, classifying and controlling flows in a network,” similarly describes:

(1) receiving and processing packets, including packets that do “not match any previously established flow” (i.e., packets of an unknown session). *PAN-1004* 7:4-6.

---

thereof **as well as certain header characteristics such as source address and destination address**. *PAN-1001* 6:11-22 (emphasis added).

(2) that a "packet typically comprises a header portion" which "contains information that is used [i.e., extracted] by line cards 202," including, e.g., a "source address and destination address" (i.e., traffic characteristics). *PAN-1004* 5:52-54.

(3) "the packets are grouped into a flow [i.e., presumed session] if those packets share a sufficient amount of header information." *PAN-1004* 6:17-26.

(4) "determining what behaviors those flows are exhibiting" and that "[t]he type of traffic that a particular flow represents can be estimated from that flow's behavioral statistics." *PAN-1004* 6:46-47, 8:26-27.

69. Pappu and Peleg disclose the claims of '629 Patent and are directed to the same problem (a method of allocating resources in a network based on media session type) and describe the same solution (using observed characteristics to group packets into common media sessions, to determine the type of media session, and to use the type of media session to allocate resources.)

X. Challenge 1: Claims 15, 17, 18, 19, 20, 21 and 22 of the '629 Patent are invalid as being obvious over Pappu alone, or in combination with Peleg

A. Pappu alone, or in view of Peleg, Teach(es) Every Element of Claims 15, 17, 18, 19, 20, 21 and 22.

1. The Pappu reference

70. Pappu is titled "Identifying Flows Based on Behavior Characteristics And Applying User-Defined Actions" and is generally directed to identifying, classifying, and controlling information packet flows in a network.

a. Problems identified in the prior art

71. Pappu describes the importance of identifying and classifying packet flows in a network. For example, Pappu describes reasons for identifying VOIP traffic and peer-to-peer (P2P) traffic. *PAN-1004* 1:14-65. Pappu also describes the importance of using behavioral characteristics of a packet flow, rather than inspecting the payload of the packet, or relying solely on header information in determining the type of traffic. For example, Pappu describes that complicated inspection of a payload is expensive. *PAN-1004* 1:40-43. In addition, Pappu explains that prior art techniques relying on the identification of a port number in the header are not effective against protocols that mask ports or dynamically change ports (e.g., TCP). *PAN-1004* 1:48-49, 2:3-7. Further, Pappu explains that it was common in the prior art to use distinct signatures in the payload of a packet to identify the packet's flow. However, this technique is effective for only a

relatively short period of time because signatures change on a fairly constant basis for some protocols. In addition, this technique cannot be used in encrypted systems since the use of encryption obfuscates the signature. *PAN-1004* 2:10-26.

b. Pappu's solution uses observed characteristics

72. Pappu proposes a solution that avoids these known problems in the prior art, describing a mechanism for using observed behavior of a flow to effectively identify, classify, and control information packet flow in a network:

> This mechanism may be applied to any type of network traffic that cannot otherwise be statically classified, including, but certainly not limited to, P2P traffic, online gaming traffic, and VOIP traffic. In one embodiment of the invention, flows are identified and classified based upon their observed behavior. Because flows are identified and classified based upon their observed behavior, and because their behavior cannot be hidden, the traffic type that a flow represents can be estimated with relatively high accuracy even if the contents of the packets that represent the traffic type are dissimilar. Thus, regardless of how applications attempt to conceal (e.g., by using encryption) the type of traffic represented by packets that originate from those applications, the traffic types that flows represent can be estimated with high accuracy, and packets that belong to those flows can

be handled in a desired manner appropriate to those traffic types.

PAN-1004 2:37-53:

c. Grouping packets into packet flows based on traffic characteristics

73. Pappu discloses that a flow is a set of packets that are related in some manner, and in one embodiment, packets are grouped into a flow if those packets share a sufficient amount of header information. *PAN-1004* 6:17-26.

74. More specifically, packets are deemed to belong to the same flow if they have the "five tuple" in common. *PAN-1004* 6:15-22. Pappu explains that a "five tuple" consists of the following five pieces of information:

(1) a source address of the entity sending the packet;

(2) a source port;

(3) a destination address of the entity that is to receive the packet);

(4) a destination port number; and

(5) an indication of the transport layer protocol that is to be used.

PAN-1004 5:56-67, 6:22-26.

75. Pappu explains that if two or more packets have the same source address, the same source port number, the same destination address, the same destination port number, and the same transport layer protocol, then those packets

are deemed to belong to the same flow. *PAN-1004* 6:23-26. Pappu explains that by "grouping packets into flows, it is possible to aggregate individual packets in a meaningful way to allow the traffic flowing through router 102 to be understood at a higher level." *PAN-1004* 6:29-32.

d. Identifying flow type using observed characteristics of the group

76. Pappu discloses that although a packet's five tuple may determine the flow to which that packet belongs, the packet's five tuple alone does not necessarily determine which actions router 102 should apply to that packet. The five tuple lacks the information that categorizes the traffic as being P2P, VOIP, online gaming traffic, or another kind of traffic, primarily because of the random nature of transport layer port numbers that are present in the five tuples of packets that belong to such flows. *PAN-1004* 8:16-25.

77. Pappu teaches that the flows that pass through router 102 may represent many different types of traffic. For example, the flows may represent web browsing traffic, P2P traffic, Voice Over IP (VOIP) traffic, File Transfer Protocol (FTP) traffic, Hypertext Transfer Protocol (HTTP) traffic, gaming traffic, video traffic, etc. To make the best use of available resources, and to best control the traffic that passes through router 102, router 102 would benefit from an ability

to identify different types of traffic, and to take specified appropriate actions relative to those types of traffic. *PAN-1004* 6:33-42.

78. In operation, Pappu creates a flow block for each previously identified flow. A flow block is a data structure containing behavioral statistics of the flow. When a packet arrives at a router it is checked to see if it matches an existing packet flow. If the packet does not match an existing flow block, then a "flow manager" creates a new flow block for a new flow to which the packet is deemed to belong. *PAN-1004* 7:9-16. As the packets of a flow are processed, a set of behavioral statistics are maintained for the flow, as shown in Block 302 in Fig. 3 below:

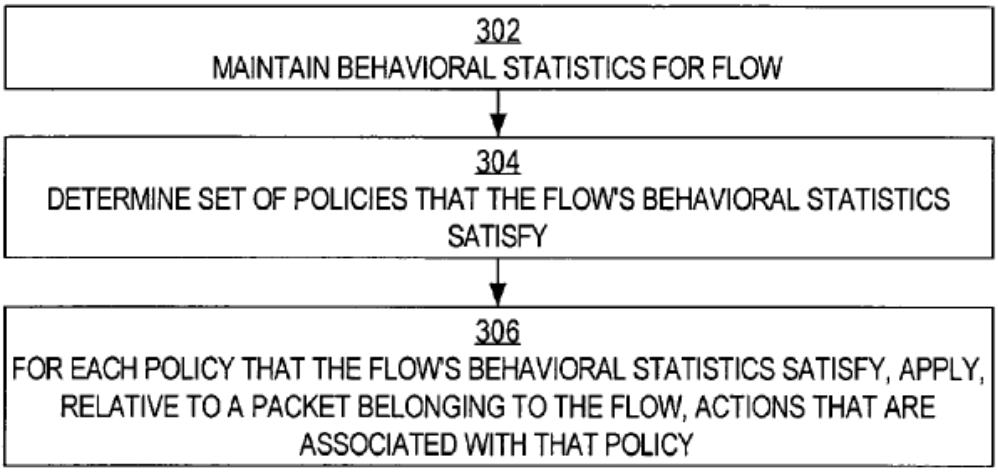

FIG. 3

PAN-1004 Fig. 3.

79. These behavioral statistics reflect the empirical behavior of the flow and are stored in the flow block for that flow. *PAN-1004* 7:21-24. FIG. 4 illustrates a flow block identifying behavioral statistics.

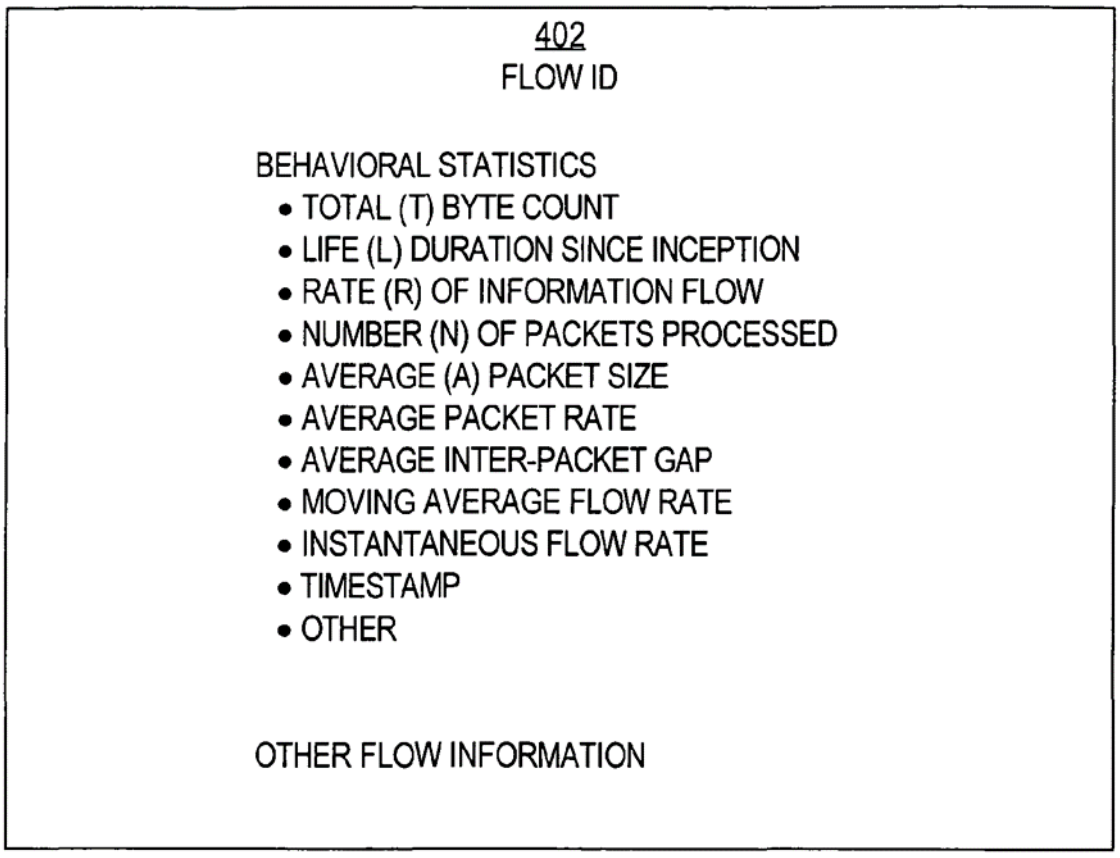

FIG. 4

80. Pappu explains that the processing of flows is described as occurring at routers, but that "the invention also may be implemented within other types of

network elements such as hubs, switches, load balancers, gateways, firewalls, etc." *PAN-1004* 4:9-13.

81. Pappu explains that the type of traffic that a particular flow represents can then be estimated from that flow's behavioral statistics:

> Therefore, even if two packets have different five tuples and therefore belong to different flows despite the fact that both of those packets represent the same type of traffic (e.g., P2P, VOIP, etc.), techniques described herein enable routers to handle both of those packets in a similar, consistent, and uniform manner. Routers which employ the techniques described herein may classify and handle packets in a manner that takes into account information beyond that which is contained in the packets themselves (some kinds of traffic cannot be identified only on the basis of the information contained in the packets that form that traffic).

PAN-1004 8:26-38.

82. After a flow block is created and behavioral statistics observed, the router determines a set of policies, reflective of the traffic/flow type, that the behavioral statistics identified in the flow block satisfy. The router accomplishes this by determining, for each specified policy, whether the behavioral statistics satisfy all of that policy's criteria. After determining the set of satisfied policies,

the router applies, relative to each packet, all of the user-specified actions that are associated with each of the satisfied policies as shown in Blocks 304 and 306 of Fig. 3 below. *PAN-1004* 15:16-22.

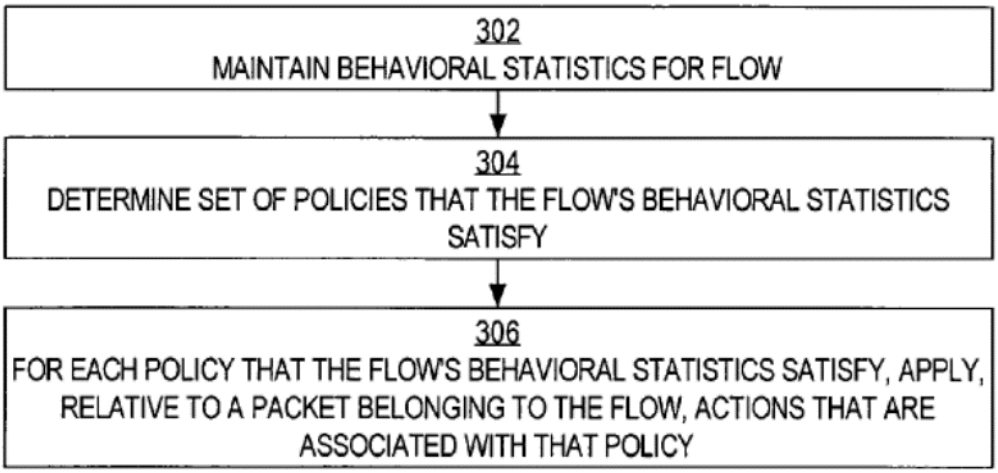

FIG. 3

83. Pappu describes that in one embodiment, each policy comprises a set of one or more user-specified criteria. The criteria may involve conditions which may or may not be satisfied by a given flow's behavioral statistics. The conditions may be specified in terms of whether a particular behavioral statistic is greater than, less than, or equal to a specified threshold, value, or parameter. The conditions may be formulated in such a way that the behavioral statistics of flows that represent a particular type of traffic tend to satisfy those conditions, while the behavioral statistics of flows that do not comprise that particular type of traffic

tend not to satisfy those conditions. *PAN-1004* 8:54-67. For example, a policy may be defined in a manner such that flows which represent VOIP traffic tend to satisfy (by their behavioral statistics) the policy's criteria, but flows which do not represent VOIP traffic tend to not satisfy the policy's criteria. *PAN-1004* 9:23-27.

e. <u>Allocation of resources</u>

84. Pappu describes one embodiment in which a user programs router 102 with one or more user-specified associations or mappings between policies and sets of actions. *PAN-1004* 10:59-62. For example, the set of actions may be directed to a quality of service associated with the flow. In one embodiment, a flow may be associated with a "high priority" flow type (e.g. VoIP). When a flow is associated with this flow type, router 102 buffers and forwards that flow's packets in a manner such that the flow's packets are forwarded before the packets of any other flow that is not also associated with that flow type. *PAN-1004* 12:6-13.

2. The Peleg Reference

85. Peleg is directed to packet classification using pattern recognition and traffic characteristics rather than traffic content recognition or port matching. *PAN-1006* [0001], [0008].

86. Peleg explains that prior art classification methods that rely on content analysis are not able to classify coded or encoded packets and require the use of

powerful network processors. *PAN-1006* [0004]. Peleg also describes that prior art methods that rely on port identification are not robust, as the port identification algorithms can be circumvented and the ports can be impersonated. *PAN-1006* [0005]-[0006].

87. Peleg describes a classifier that can use a combination of traffic characteristics to classify a packet or a session. *PAN-1006* [0027]. The traffic characteristics include:

(a) IP address and MAC address,

(b) application port number,

(c) length of packet,

(d) type of service,

(e) quality of service,

(f) time interval between packets, and

(g) throughput per second.

PAN-1006 [0028]-[0040].

88. Peleg describes an exemplary classifier as one that "classifies the packets based on their length and/or the time interval between successive packets belonging to the same session." *PAN-1006* [0022]. Peleg also discloses that a sessions table can be maintained that includes identified and potential sessions. *PAN-1006* [0041].

89. In one embodiment, the sessions table includes Source IP address, Destination IP address, Source Port number, and Destination Port number. *PAN-1006* [0069]. When a packet is received, the classifier checks to see if the packet Source IP address, Destination IP address, Source Port number, and Destination Port number appears in the sessions table. If the current session does not appear in the sessions table, the classifier creates a new entry which represents a new "potential session." *PAN-1006* [0041].

90. Peleg explains that implementing classification of IP traffic based on traffic characteristics was more efficient, cheaper, and required less processing power than content-based classification. *PAN-1006* [0009]-[00011].

B. Detailed Invalidity Analysis

a. <u>Claim 15</u>

i. A method of identifying session type, comprising:

91. Pappu discloses a method of "identifying, classifying, and controlling information packet flows in a network". *PAN-1004* Abstract, 2:34-44. Pappu discloses identifying the particular type of network traffic, such as "VOIP, gamming, streaming and P2P flows." *PAN-1004* Abstract.

> In accordance with one embodiment of the present invention, a mechanism for effectively identifying,

> classifying, and controlling information packet flows in a network is provided. This mechanism may be applied to any type of network traffic that cannot otherwise be statically classified, including, but certainly not limited to, P2P traffic, online gaming traffic, and VOIP traffic. In one embodiment of the invention, flows are identified and classified based upon their observed behavior.

PAN-1004 2:34-44.

92. The "type of network traffic" described in Pappu is the claimed "session type."

ii. obtaining passing packets of respectively unknown sessions and unknown session types;

93. Pappu discloses receiving and processing of packets (*obtaining passing packets*) including packets that "do[] not match any previously established flow". (*unknown sessions and unknown session types*). *PAN-1004* 7:4-6.

> FM 210 is assumed to be on a line card that is acting as an egress line card (i.e. the line card is receiving packets from an ingress line card and sending packets out to another router).

PAN-1004 6:63-66.

> FM 210 may establish a new flow whenever FM 210 receives and processes a packet that does not match any previously established flow (e.g., based on that packet's header information). As is described in greater

> detail below, in one embodiment of the invention, when FM 210 receives a packet, FM 210 determines whether that packet matches any existing "flow block" data structure. If the packet matches an existing flow block, then the packet is deemed to belong to the flow that corresponds to that flow block. Alternatively, if the packet does not match an existing flow block, then FM 210 creates a new flow block for a new flow to which the packet is deemed to belong.

PAN-1004 7:4-16.

iii. obtaining traffic packet characteristics of said passing packets of respectively unknown session types;

94. Pappu discloses obtaining certain header information (*traffic packet characteristics of the passing packets*):

> In one embodiment of the invention, packets are grouped into a flow if those packets share a sufficient amount of header information. More specifically, in one embodiment of the invention, packets are deemed to belong to the same flow if they have the five tuple in common. Thus, if two or more packets have the same source address, the same source port number, the same destination address, the same destination port number, and the same transport layer protocol, then those packets are deemed to belong to the same flow.

PAN-1004 6:17-26.

95. As explained above, the ’629 patent specifically teaches that the “traffic packet characteristics” include source and destination addresses. For example, ’629 patent teaches that its invention “relates to IP traffic inspection by means of identifying those packets having similar traffic characteristics as probabilistically belonging to the same media session. Traffic characteristics relate to those parameters inherent to packets traffic… includes (*sic*) features such as… certain header characteristics ***such as source address and destination address***…” *PAN-1001* 6:8-17 (emphasis added). Thus, at least part of the “header information” of Pappu is included in the list of information the ’629 Patent defines as “traffic characteristics”

96. Thus, Pappu’s header information is the claimed “traffic packet characteristics.”

iv. comparing said obtained packets with each other using respectively obtained traffic packet characteristics;

97. Pappu discloses comparing certain header information of one packet with the same header information (*traffic packet characteristics*) of another packet to determine if they share sufficient header information:

> In one embodiment of the invention, *packets are grouped into a flow if those packets share a sufficient amount of*

> *header information.* More specifically, in one embodiment of the invention, packets are deemed to belong to the same flow if they have the five tuple in common. *Thus, if two or more packets have the same source address, the same source port number, the same destination address, the same destination port number, and the same transport layer protocol*, then those packets are deemed to belong to the same flow. (emphasis added)

PAN-1004 6:17-26.

98. A POSA would have understood that determining if packets "share a sufficient amount of header information" and that determining "if two or more packets have the same source address, the same source port number, the same destination address, the same destination port number, and the same transport layer protocol" are both comparisons of packets/traffic packet characteristics with each other, and, more specifically, are both comparisons that are performed by using the claimed "traffic packet characteristics." As noted above in claim limitation iii, the header information of Pappu comprises the claimed "traffic packet characteristics."

v. grouping together those packets having similar values of said traffic packet characteristics into a presumed session;

99. Pappu discloses grouping packets with the similar header information (*similar values of said traffic packet characteristics)* into previously unknown flows *(presumed sessions*):

> In one embodiment of the invention, *packets are grouped into a flow if those packets share a sufficient amount of header information.* More specifically, in one embodiment of the invention, packets *are deemed to belong to the same flow* if they have the five tuple in common. Thus, if two or more packets have the same source address, the same source port number, the same destination address, the same destination port number, and the same transport layer protocol, then *those packets are deemed to belong to the same flow*. (emphasis added)

PAN-1004 6:17-26.

100. Pappu discloses that this grouping into a flow is not necessarily sufficient to identify the specific session and therefore the grouping is only a presumed session:

> Although a packet's five tuple may determine the flow to which that packet belongs, *the packet's five tuple alone does not necessarily determine which actions router 102 will apply to that packet. This is because the five tuple lacks the information that categorizes the traffic as being P2P, VOIP, online gaming traffic, or*

> *another kind of traffic*; the five tuple lacks this information primarily because of the often random nature of transport layer port numbers that are present in five tuples of packets that belong to such flows. (emphasis added)

PAN-1004 8:16-25.

101. Because the received packet does not match an existing packet flow, and it is not yet determined what type of flow the packet belongs to, a new flow is "presumed": "the packet is assumed to be the first packet of a new flow." *PAN-1004* 14:50-51.

102. To the extent that the Patent Owner asserts that Pappu does not disclose grouping packets in a presumed session, it would have been obvious in view of Peleg. Peleg discloses specific criteria that are used to initially group packets into potential sessions (*presumed sessions*) and identify sessions:

> Moreover, in an alternative embodiment, there is looking for the appropriate IP source, IP destination, and port numbers, in a sessions table. In an embodiment of the present invention, the sessions table stores the identified and potential sessions. In case there is no entry for the specific session in the sessions table, a new entry in the sessions table is created. Actually, the new entry represents a new potential session.

PAN-1006 [0041]. Only after observing a few consecutive packets does Peleg's system decide that a session is "identified":

> The classifier has to decide when there is a session according to several consecutive packets [*sic*] that matches its rules. The classifier may identify that is it not the same session according a few consecutive packets and their time. i.e. the classifier has to be able to know when there is a streaming session and when the streaming session ends.

PAN-1006 [0047].

103. A POSA would be motivated to modify the teachings of Pappu (grouping packets in flows) with the teachings of Peleg (initially grouping packets in potential sessions, and then deciding if there is an identified session) to thereby allow treating the identified sessions differently from potential sessions in order to allocate resources more efficiently. Specifically, Pappu explains that a router would benefit from an ability to identify different types of traffic, and to take specified appropriate actions relative to those types of traffic, because it would make the best use of available resources best control the traffic that passes through router 102. *PAN-1004* 6:38-42. A POSA would have understood that taking into consideration that packets can belong to an identified or a potential session is a

better use of available resources than treating all flows interchangeably. For example, it may be more efficient to focus resources on identified sessions since their resource demands are known and more predictable. Moreover, Pappu and Peleg are in the same field of endeavor (e.g., routing different types of packet flows), address the same problems (e.g., efficiently allocating resources network resources to packet flows), and propose the same solutions (e.g., identifying and classifying packet flows using traffic characteristics and not either packet contents or port numbers). Accordingly, modifying the teachings of Pappu with the teachings of Peleg results in a more efficient use of resources using a known technique to improve a similar system and to yield predictable results. A POSA, seeking to modify the system of Pappu to improve its efficiency, would be drawn to the teachings of Peleg, inasmuch as it is drawn to the same field of art and to solving the same problems, and would find its teachings readily adaptable, easily implemented without undue experimentation, and likely to yield predictable results.

104. Therefore, in view of the above, it would have been obvious to a POSA to group the packets of Pappu in a potential session (*presumed session*) as taught by Peleg. A POSA would have been motivated to make the modification to more efficiently allocate resources as described above.

vi. and analyzing said grouped packets of said presumed session for session characteristic;

105. Pappu discloses obtaining behavioral statistics (*session characteristics*) of the passing packets of the previously unknown flow (*presumed session*).

106. Pappu describes creating a flow block including the Flow ID and a set of behavioral statistics of the flow. *PAN-1004* 7:21-27, 14:50-56, Fig. 3. The behavioral statistics include: Total Byte Count, Life Duration, Flow Rate, Number, Average Packet Size, Average Packet Rate, Average Inter-Packet Gap, Moving Average Flow Rate, Instantaneous Flow Rate, and Flow Block Creation Timestamp. *PAN-1004* 7:28-50, 14:57-15:5, Fig. 4.

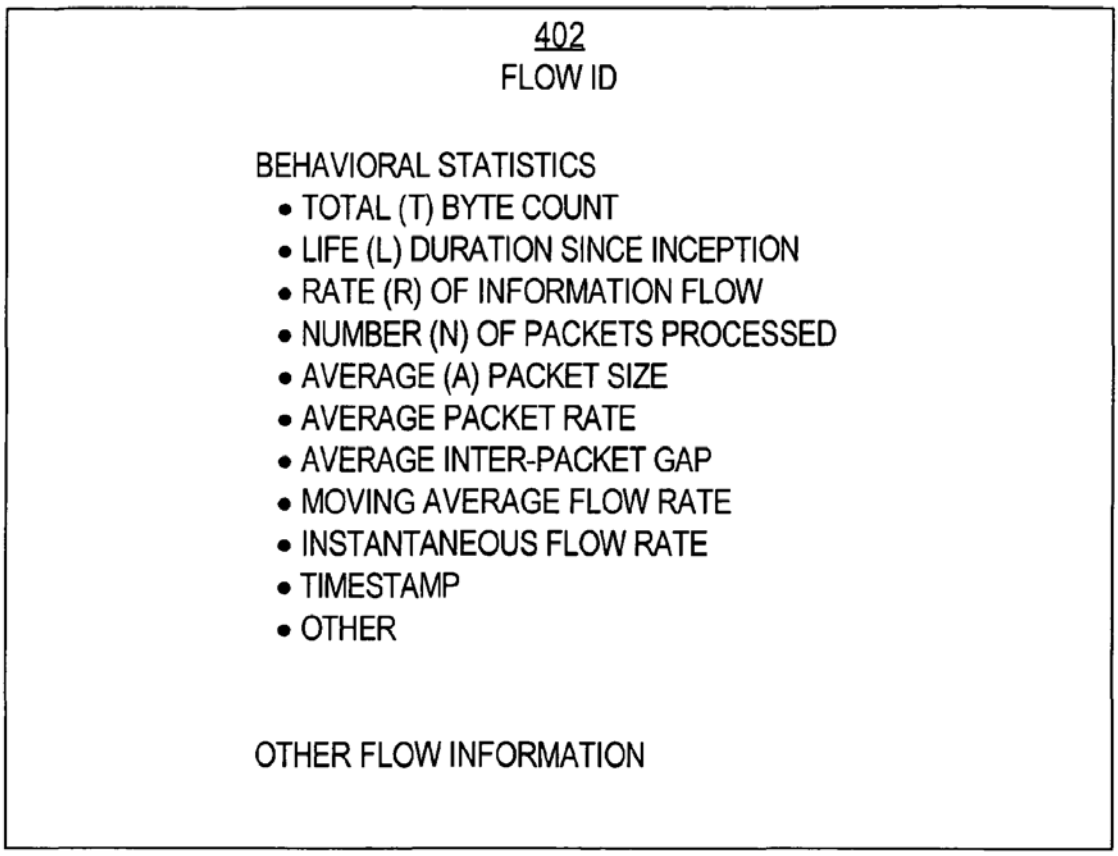

FIG. 4

107. Pappu discloses that the behavioral characteristics of the flow can include averaging the individual packet characteristics including packet size, packet rate and inter-packet gap:

> an average packet size (derived by dividing the total byte count of the flow by the total number of packets in the flow that have been processed). In one embodiment of the invention, the behavioral statistics include an average packet rate (derived by dividing (a) the total

> number of packets in the flow that have been processed by (b) the life duration of the flow)
>
> …
>
> an average inter-packet gap (derived by averaging the amounts of time that pass between the arrivals of packets that belong to the flow),

PAN-1004 7:34-42.

108. Pappu discloses that the grouped packets are analyzed by the flow manager to determine what behaviors those flows are exhibiting. *PAN-1004* 6:44-49.

vii. using said session characteristics to identify a session type of said presumed session.

109. Pappu discloses using behavioral statistics (*session characteristics*) to identify the type of traffic (*session type*) represented by the flow (*presumed session*):

> The type of traffic that a particular flow represents can be estimated from that flow's behavioral statistics.

PAN-1004 8:26-27.

110. Pappu provides several examples of determining the type of traffic represented by the flow by using behavioral statistics, including VOIP and online gaming:

> In one embodiment of the invention, the VOIP policy that is applied to UDP flows, to determine whether those

UDP flows represent VOIP traffic, contains the following criteria: the flow's average bit rate needs to be lower than 100 Kbps, the flow's life duration needs to be greater than 10 seconds, and the flow's average packet size needs to be less than 230 bytes. Flows that satisfy these criteria may be identified as flows that represent VOIP traffic, and actions appropriate to VOIP traffic may be applied to these flows.

PAN-1004 9:39-48.

In one embodiment of the invention, the online gaming policy that is applied to UDP flows, to determine whether those UDP flows represent online gaming traffic, contains the following criteria: the flaw's average bit rate needs to be both greater than 8 Kbps and less than 32 Kbps, the flow's life duration needs to be greater than 10 seconds, and the flow's average packet size needs to be greater than 100 bytes and less than 500 bytes. Flows that satisfy these criteria may be identified as flows that represent online gaming traffic, and actions appropriate to online gaming traffic may be applied to these flows.

PAN-1004 10:17-27.

b. <u>Claim 17</u>

The method of claim 15 wherein said identifying a session type is followed by at least one of the group comprising: providing provisioning and providing control actions to the session.

111. Pappu describes applying user specified actions (*control actions*) to the flow (*presumed session*) based on the prior step of identifying of flow type (*session type*).

> In one embodiment of the invention, whenever a packet belonging to the flow is processed, a set of zero or more user-specified policies (within a set of user-specified policies) that the flow's behavioral statistics satisfy is determined. For each policy that the flow's behavioral statistics satisfy, one or more user-specified actions that are associated with that policy are applied relative to the packet. The actions may be designed to cause a router to handle, in a user-specified manner, packets that are likely to represent a particular kind of traffic. For example, actions may include penalizing or rewarding a particular flow by changing that flow's "drop priority," changing a particular flaw's "flow type," changing 15 an "aggregate class" to which a particular flow belongs, rerouting packets that belong to a particular flow, mirroring packets that belong to a particular flow, capturing and/or logging information about packets that belong to a particular flow, and other actions that are described in greater detail 20 below.

PAN-1004 3:4-21.

112. Pappu teaches that these actions are applied to flows that have been identified as having the same flow type. For example, flows identified as VOIP traffic are each treated in a similar manner, as are flows identified as online gaming.

> Thus, different flows that are associated with similar behavioral statistics may be handled in similar ways. For example, multiple flows that exhibit behaviors that are characteristic of P2P traffic all may be handled in one way that is appropriate for P2P traffic. Multiple flows that exhibit characteristics of VOIP traffic all may be handled in another way that is appropriate for VOIP traffic. Multiple flows that exhibit characteristics of online gaming traffic all may be handled in yet another way that is appropriate for online gaming traffic.

PAN-1004 3:22-30. See also *PAN-1004* 11:60-13:16.

c. <u>Claim 18</u>

Method according to claim 15, further comprising thresholding said traffic characteristics, and identifying those packets whose characteristics are mutually within said thresholds for said grouping.

113. Pappu discloses that a flow is a set of packets that are related in some manner. In one embodiment of the invention, packets are grouped into a flow "if those packets share a sufficient amount of header information." *PAN-1001* 6:15-19. A POSA understood that in order to determine whether packets share a sufficient amount of header information, a threshold can be set. The threshold condition is whether or not the shared amount of header information is "sufficient." For example, in one embodiment, Pappu set the threshold that packets are deemed to belong to the same flow if they have the five tuple in common, e.g., if two or more packets have the same source address, the same source port number, the same destination address, the same destination port number, and the same transport layer protocol. *PAN-1001* 6:19-26. Thus, Pappu disclosed using thresholding of the traffic characteristics to group the packets.

d. <u>Claim 19</u>

The method of claim 15, wherein said characteristics comprise at least one member of the group consisting of packet length, inter-arrival period, and average bandwidth.

114. At least one antecedent basis for "said characteristics" are the "session characteristics" of claim 15. Pappu discloses that behavioral statistics (*session characteristics*) include average packet length, average inter-arrival period, and

average bandwidth as discussed with reference to Claim 15, above. *See* Section X.B.a.vi.

> an average packet size (derived by dividing the total byte count of the flow by the total number of packets in the flow that have been processed). In one embodiment of the invention, the behavioral statistics include an average packet rate (derived by dividing (a) the total number of packets in the flow that have been processed by (b) the life duration of the flow)
>
> …
>
> an average inter-packet gap (derived by averaging the amounts of time that pass between the arrivals of packets that belong to the flow)

PAN-1004 7:34-42.

> *The behavioral statistics include* a total (T) byte count (sum of all of the bytes in all of the packets of the flow that have been processed up to the current time), a life duration (L) (how long the flow has been in existence since inception), a flow rate (R) (derived by dividing T by L), a number (N) of the flow's packets processed up to the current time, an average (A) packet size (derived by dividing T by N), an average packet rate (derived by dividing N by L), an average inter-packet gap (derived by averaging the amounts of time that pass between the arrival of packets that belong to the flow), a moving

average flow rate (discussed above), and an instantaneous flow rate (derived by dividing (a) the size of the most recently arrived packet of the flow by (b) the *inter-arrival time gap* between the two most recently arrived packets of that flow)

PAN-1004 14:57-15:4.

e. Claim 20

The method of claim 15, further comprising inserting into a header of a respective packet an indication of said presumed session with which said respective packet has been identified.

115. Pappu discloses inserting into the header of a packet a code (*indication*) associated with the flow (*presumed session with which said respective packet has been identified*) of each respective packet belonging to the identified flow.

116. Pappu discloses "each flow is associated with a separate QOS," and that "[a] flow's QOS generally indicates to router 102, how router 102 should buffer and forward packets belonging to that flow." *PAN*-1004 11:63-67. Pappu further discloses changing the QOS associated with a flow by changing the contents of the DSCP field in the header of each packet associated with the particular flow. A POSA would understand that the DSCP field is found inside the TOS field that is located inside an IP header, and is used as an indication of the Class Selector (*i.e.* QOS level or priority level) associated with that packet. Pappu

discloses the use of this field for its intended purpose, namely to identify the priority of the packet. Pappu discloses:

> For another example, a policy might be associated with an action that causes the contents of a packet to be altered in a specified manner. More specifically, a policy might be associated with an action that causes the contents of a "diffusely[*sic*] code point" (DSCP) field of a packet's IP header to be modified (A part of the Type of Service (TOS) field is renamed as the DSCP field in the IP header that routers and other network elements may use to determine the priority and/or the quality of service that the containing packet should be given).

PAN-1004 11:21-29. In other words, Pappu is disclosing inserting a value into the DSCP field that indicates the priority accorded to that packet, which is the claimed step of "inserting… an indication."

f. <u>Claim 21</u>

The method of claim 15, comprising obtaining said characteristics belonging to said presumed session from parameter patterns common to said packets.

117. Pappu discloses behavioral statistics are obtained from the observed behavior of the flow:

> As the packets of a flow are processed, a set of behavioral statistics are maintained (block 302 of FIG. 3)

> for the flow. These behavioral statistics reflect the empirical behavior of the flow. In one embodiment of the invention, FM 210 stores information about a flow's behavioral statistics in the flow block for that flow.

PAN-1004 7:21-27.

118. Pappu explains:

> A policy may be defined in a manner such that flows tend to satisfy (by their behavioral statistics) that policy's criteria if those flows represent a particular kind of traffic (e.g., P2P, VOIP, gaming, etc.) …

PAN-1004 9:18-21.

119. In a specific example, Pappu describes using behavioral statistics to determine if the flow represents online gaming traffic. One of the conditions that must be satisfied is that "the flow's packet size needs to be at least as great as a third threshold value" but "no greater than a fourth threshold value." *PAN-1004* 10:11-14. Thus the characteristic of packet size within threshold values is one example of a parameter pattern common to the packets of the flow. Another example of a pattern that is "common to said packets" is average packet size. See, e.g., *PAN-1004* 9:44-45, 10:22-25.

g. Claim 22

The method of claim 15, comprising using at least one member of the group consisting of:

packet length

inter-arrival period

variance of inter-arrival period

variance of length

covariance of inter-arrival period

covariance of length

a statistical derivative obtained from multi records of packet length

a statistical derivative obtained from multi records of inter-arrival period

a histogram built of at least one of the following; i. packet length ii. inter-arrival period iii. variance of inter-arrival period iv. variance of length v. covariance of inter-arrival period vi. covariance of length vii. a statistical derivative obtained from multi records of packet length viii. a statistical derivative obtained from multi records of inter-arrival period

a matrix histogram built of at least one of the following; i. packet length ii. inter-arrival period iii. variance of inter-arrival period iv. variance of length v. covariance of inter-arrival period vi. covariance of length vii. a statistical derivative obtained from multi records of packet length viii. a statistical derivative obtained from multi records of inter-arrival period.

120. Pappu, as described above with respect to Claims 15 and 19, discloses the use of header information and behavioral statistics that include packet length and inter-arrival period. *PAN-1004* 7:28-50, 14:57-15:5, Fig. 4.

Conclusion and Declaration

The findings and opinions set forth in this declaration are based on my work and examinations to date. I may continue my examinations. I may also receive additional documentation and other factual evidence over the course of this contested case that will allow me to supplement and/or refine my opinions. I reserve the right to add to, alter, or delete my opinions and my declaration upon discovery of any additional information. I reserve the right to make such changes as may be deemed necessary. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

I declare that all statements made herein on 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.

Dated: 2/7/2018 By: Samuel H. Russ

Samuel H. Russ