

UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE PATENT TRIAL AND APPEAL BOARD

Limelight Networks, Inc.,
Petitioner

v.

Akamai Technologies, Inc.,
Patent Owner.

Patent No. 8,194,538 to Bornstein et al.

IPR Case No.: IPR2018-00310

**PETITION FOR *INTER PARTES* REVIEW OF CLAIMS 1-5
OF U.S. PATENT 8,194,538 UNDER 35 U.S.C. §§ 311-319
AND 37 C.F.R. §§ 42.100 ET SEQ.**

TABLE OF CONTENTS

TABLE OF EXHIBITS	iv
I. MANDATORY NOTICES UNDER 37 C.F.R. § 42.8(a)(1) FOR INTER PARTES REVIEW.....	1
A. Real Party in Interest Under 37 C.F.R. § 42.8(b)(1).....	1
B. Related Matters Under 37 C.F.R. § 42.8(b)(2)	1
C. Lead and Backup Counsel Under 37 C.F.R. § 42.8(b)(3) and Service Information under 37 C.F.R. § 42.8(b)(4).....	1
II. PAYMENT OF FEES UNDER 37 C.F.R. § 42.15.....	3
III. Certification of Word Count Under 37 C.F.R. § 42.24(d)	3
IV. REQUIREMENTS FOR IPR UNDER 37 C.F.R. § 42.104.....	3
A. Grounds for Standing Under 37 C.F.R. § 42.104(a).....	3
B. Identification of Challenge Under 37 C.F.R. § 42.104(b) and Relief Requested.....	4
V. OVERVIEW OF THE '538 PATENT	4
A. Admitted Prior Art Content Delivery Networks and the Alleged Problem Addressed by the '538 Patent	4
B. Priority Date	9
C. The Prosecution History of the '538 Patent	9
D. The Challenged Claims	10
E. Person of Ordinary Skill in the Art (“POSITA”).....	10
F. Claim Construction.....	10
1. “alternate path”	10
2. “server node”.....	13
G. Overview of Prior Art.....	17
1. Overview of Speedera (Ex. 1004-Speedera).....	17
2. Overview of Ricciulli-470 (Ex. 1005-Ricciulli).....	19
3. Overview of HTTP/1.1 (Ex. 1006-HTTP).....	21
VI. SPECIFIC GROUNDS FOR PETITION.....	24

A.	GROUND 1: Speedera in view of Ricciulli-470 and HTTP/1.1 renders the Challenged Claims obvious.....	24
1.	The scope and content of the prior art	24
2.	Motivation to Combine and Rationale for Combining Speedera, Ricciulli-470, and HTTP/1.1.....	24
3.	Challenged Claims	32
VII.	CONCLUSION.....	59
	CERTIFICATE OF SERVICE	61

TABLE OF EXHIBITS

Exhibit	Description
1001	U.S. Patent No. 8,194,538 (“’538 patent”)
1002	File history of U.S. Patent No. 8,194,538 (“’538 FH”)
1003	Declaration of Kevin Almeroth, Ph.D., Regarding Invalidity Of U.S. Patent No. 8,194,538 (“Almeroth”)
1004	U.S. Patent No. 6,405,252 (“Speedera”)
1005	U.S. Patent No. 6,275,470 (“Ricciulli-470”)
1006	Fielding, R. et al., Hypertext Transfer Protocol - HTTP/1.1, RFC 2616, dated June 1999 (“HTTP/1.1”)
1007	U.S. Provisional Application No. 60/166,906
1008	Pete Loshin, ed., <i>Big Book of World Wide Web RFCs</i> , Morgan Koffman Publishing (2000)
1009	Library of Congress Online Catalog record of publication of <i>Big Book of World Wide Web RFCs</i>
1010	WorldCat record of publication of <i>Big Book of World Wide Web RFCs</i>
1011	Limelight Networks, Inc.’s Preliminary Claim Construction Brief in <i>Akamai Techs., Inc. v. Limelight Networks, Inc.</i> , No. 1:16-cv-12624-GAO (D. Mass. Dec. 27, 2016), <i>consolidated with</i> 1:16-cv-10253 (D. Mass. Feb. 16, 2016)
1012	Akamai Technologies, Inc.’s Opening Claim Construction Brief in <i>Akamai Techs., Inc. v. Limelight Networks, Inc.</i> , No. 1:16-cv-12624-GAO (D. Mass. Dec. 27, 2016), <i>consolidated with</i> 1:16-cv-10253 (D. Mass. Feb. 16, 2016)
1013	U.S. Patent No. 6,356,907
1014	U.S. Patent No. 6,567,857

1015	W3 Press Release, available at http://www.w3.org/1999/07/HTTP-PressRelease
------	---

Limelight Networks, Inc. (“Petitioner”) hereby seeks *inter partes* review of Claims 1-5 (“the Challenged Claims”) of U.S. Patent No. 8,194,538. (Ex. 1001 (the “’538 patent”).) The Challenged Claims of the ’538 patent do not claim anything inventive; they claim a previously-known way of hosting and delivering content on the Internet. The Challenged Claims in the patent should, therefore, be canceled for the reasons described in this Petition.

I. MANDATORY NOTICES UNDER 37 C.F.R. § 42.8(a)(1) FOR INTER PARTES REVIEW

A. Real Party in Interest Under 37 C.F.R. § 42.8(b)(1)

The real party-in-interest in this petition is Limelight Networks, Inc., 222 South Mill Ave., Suite 800, Tempe, Arizona 85281.

B. Related Matters Under 37 C.F.R. § 42.8(b)(2)

A pending federal district court litigation may affect or be affected by the decision in this proceeding: *Akamai Techs., Inc. v. Limelight Networks, Inc.*, No. 1:16-cv-12624-GAO (D. Mass. Dec. 27, 2016), *consolidated with* 1:16-cv-10253 (D. Mass. Feb. 16, 2016) (“the Litigation”). The undersigned is unaware of any other any other judicial or administrative matter that would affect, or be affected by, a decision in the proceeding.

C. Lead and Backup Counsel Under 37 C.F.R. § 42.8(b)(3) and Service Information under 37 C.F.R. § 42.8(b)(4)

Petitioner designates the following lead and backup counsel:

Lead Counsel:	Barry J. Schindler (Reg. No. 32,938) Greenberg Traurig, LLP 500 Campus Drive, Suite 400 Florham Park, NJ 07932 Telephone: 973-360-7900 Facsimile: 973-301-8410 SchindlerB@gtlaw.com
Backup Counsel:	Heath J. Briggs (Reg. No. 54,919) Greenberg Traurig, LLP 1200 17th St., Suite 2400 Denver, CO 80202 Telephone: 303-685-7418 Facsimile: 720-904-6118 BriggsH@gtlaw.com
Backup Counsel:	Stephen M. Ullmer (<i>pro hac vice</i> forthcoming) Greenberg Traurig, LLP 4 Embarcadero Center, Suite 3000 San Francisco, CA 94111 Telephone: 415-655-1261 Facsimile: 415-707-2010 UllmerS@gtlaw.com
Backup Counsel:	Douglas R. Weider (<i>pro hac vice</i> forthcoming) Greenberg Traurig, LLP 500 Campus Drive, Suite 400 Florham Park, NJ 07932 Telephone: 973-360-7900 Facsimile: 973-301-8410 WeiderD@gtlaw.com
Backup Counsel:	Vimal M. Kapadia (Reg. No. 73,310) Greenberg Traurig, LLP MetLife Building 200 Park Ave. New York, NY 10166 Telephone: 212-801-2241 Facsimile: 212-801-6400 KapadiaM@gtlaw.com

Service on Petitioner may be made by mail or hand delivery to: Greenberg Traurig, LLP, 1200 17th St., Suite 2400, Denver, CO, 80202. Petitioner also consents to, and prefers, electronic service by emailing counsel of record at schindlerb@gtlaw.com, briggsh@gtlaw.com, and limelight-iprs@gtlaw.com.

II. PAYMENT OF FEES UNDER 37 C.F.R. § 42.15

Petitioner authorizes the U.S. Patent & Trademark Office to charge Deposit Account No. 50-2638 for the fee set in 37 C.F.R. § 42.15(a) for this Petition and further authorizes for any additional fees to be charged to this Deposit Account.

III. CERTIFICATION OF WORD COUNT UNDER 37 C.F.R. § 42.24(d)

Petitioner certifies that the word count in this Petition is 11,902 words, as counted by the word-processing program (Microsoft Word 2010) used to generate this Petition, where such word count excludes the table of contents, table of authorities, mandatory notices, certificate of service, appendix of exhibits, and this certificate of word count. This Petition is in compliance with the 14,000 word limit set forth in 37 C.F.R. § 42.24(a)(1)(i).

IV. REQUIREMENTS FOR IPR UNDER 37 C.F.R. § 42.104

A. Grounds for Standing Under 37 C.F.R. § 42.104(a)

Petitioner certifies that the '538 patent is available for *inter partes* review, and that Petitioner is not barred or estopped from requesting an *inter partes* review on the grounds identified in the petition.

B. Identification of Challenge Under 37 C.F.R. § 42.104(b) and Relief Requested

Petitioner respectfully requests that the Challenged Claims of the '538 patent (Ex. 1001) be cancelled based on the following ground of unpatentability:

Ground of Unpatentability	'538 Patent Claim(s)	Basis for Rejection under 35 U.S.C. § 103
Ground 1	1-5	Rendered obvious over Speedera (Ex. 1004) in view of Ricciulli-470 (Ex. 1005) and HTTP/1.1 (Ex. 1006)

V. OVERVIEW OF THE '538 PATENT

A. Admitted Prior Art Content Delivery Networks and the Alleged Problem Addressed by the '538 Patent

The '538 patent explains in the “Description of the Related Art” section that content delivery networks (“CDNs”) were previously known in the art:

“One way that enterprises have sought to manage and move information effectively via the Internet is through the use of content delivery networks. A content delivery network or ‘CDN’ is a network of geographically distributed content delivery nodes that are arranged for efficient delivery of content on behalf of third party content providers. ...

Typically, a CDN is implemented as a combination of a content delivery infrastructure, a request-routing mechanism, and a distribution infrastructure. The content delivery infrastructure usually comprises a set of ‘surrogate’ origin servers that are located at


strategic locations (e.g., Internet network access points, Internet Points of Presence, and the like) for delivering copies of content to requesting end users. ... The distribution infrastructure consists of on-demand or push-based mechanisms that move content from the origin server to the surrogates. An effective CDN serves frequently-accessed content from a surrogate that is optimal for a given requesting client. In a typical CDN, a single service provider operates the request-routers, the surrogates, and the content distributors. In addition, that service provider establishes business relationships with content publishers and acts on behalf of their origin server sites to provide a distributed delivery system. A well-known commercial CDN service that provides web content and media streaming is provided by Akamai Technologies, Inc. of Cambridge, Mass.” (Ex. 1001 at 1:27-59 (emphasis added).)

Consistent with the admission that CDNs were known in the art, FIG. 1 of the '538 patent illustrates a known CDN in which the purported invention of the '538 patent may be implemented:

“FIG. 1 is a block diagram of a known content delivery network in which the present invention may be implemented.” (Ex. 1001 at 3:63-65 (emphasis added).)

Below, Petitioner provides an annotated version of this FIG. 1 to indicate components that are pertinent to the claim limitations discussed below in the context of the Challenged Claims:

FIG. 1



As Petitioner understands the claims in the context of FIG. 1, the First through Nth server nodes (orange) make up a “set of server nodes” (preamble element 1.P.a) and each server node comprises a “cluster of content delivery network edge servers” (purple; preamble element 1.P.b). Each of the server nodes (orange) are organized by region (preamble element 1.P.a). The server nodes (orange) provide delivery of content on behalf of participating content providers (preamble element

1.P.b), such as by obtaining content from the content provider's origin server 115 (yellow). (Ex. 1001 at 1:32-36; 2:63-67.)

Petitioner notes that labeling ambiguities exist in FIG. 1. First, callout (106) is used twice, as a label for "region" and as a label for "content migration." This error is repeated in the specification. (Ex. 1001 at 4:42; 6:6-7.) Petitioner also notes that the "REGION 106" appears to refer to a set of CDN servers (102) within a point of presence ("POP"):

"A CDN service provider (CDNSP) may organize sets of surrogate origin servers as a 'region.' In this type of arrangement, a CDN region 106 typically comprises a set of one or more content servers that share a common backend, e.g., a LAN, and that are located at or near an Internet access point. Thus, for example, a typical CDN region may be co-located within an Internet Service Provider (ISP) Point of Presence (PoP) 108." (Ex. 1001 at 4:40-47.)

"[A] CDN service provider typically segments its servers into regions, with each region comprising a set of content servers (e.g., up to about ten (10) servers) that preferably operate in a peer-to-peer manner and share data across a common backbone such as a local area network (LAN)." (*Id.* at 6:6-11.)

"[A] server node comprises a cluster of content delivery network edge servers." (*Id.* at Claims 1, 2, and 5 (preambles).)

In other words, consistent with the specification, a POP and a Region may be co-located, and may include a “server node” (orange) that comprises a cluster of content delivery network edge servers (purple).

The ’538 patent describes the purported problem it solves as follows:

“The present invention, in particular, addresses the problem of finding the best way to route data between content delivery network (CDN) regions and content providers, while quickly adjusting to changing network conditions.” (Ex. 1001 at 2:38-41 (emphasis added).)

As the “Description of the Related Art” section of the ’538 patent admits, it was known in the art that the performance of the Internet—as affected by issues including congestion, economics, cost, speed, misconfiguration, and damage—impacts the performance of a CDN. (Ex. 1001 at 1:60-2:27.) Any or all of these issues may cause a server to communicate via a path that is not ideal. (*Id.* at 2:28-37.) The ’538 patent to solve this purported problem by “selectively routing through the CDN’s own nodes, thereby avoiding network congestion and hot spots.” (*Id.* at 2:61-63.)

As shown in detail in this Petition with respect to the Challenged Claims, this identified problem, and the purported solutions provided in the Challenged Claims, did not need solving – solutions to the purported problem were already known and obvious to a person of ordinary skill in the art.

B. Priority Date

The '538 patent issued on June 5, 2012, from Application No. 13/088,517 filed on April 18, 2011 and claims priority through other applications to U.S. Patent Provisional Application No. 60/273,241, filed on March 1, 2001. Thus, the earliest possible priority date for the '538 patent is March 1, 2001.¹

C. The Prosecution History of the '538 Patent

The '538 patent was subject to a single office action. (Ex. 1002 at 35-40.) In that action, the Examiner rejected pending claims 1, 2, 4, and 10 as being unpatentable over U.S. Patent No. 6,785,704 to McCanne in view of U.S. Patent No. 5,369,686 to Dutra. The Examiner also objected to pending claims 3, 5, and 6-9 as being dependent upon a rejected base claim, but allowable if rewritten in independent form. (*Id.*)

In response, Patent Owner amended claim 1 to include the limitations of original claims 2-3. (*Id.* at 26.) Similar amendments were made to other dependent claims (*id.*), resulting in allowance. (*Id.* at 9-16).

¹ Petitioner does not concede the '538 patent is entitled to the priority date of the provisional patent application.

D. The Challenged Claims

Claims 1, 2, and 5 are the independent Challenged Claims of the '538 patent, and are generally of the same scope, as shown in detail below. Claims 1-5 comprise the Challenged Claims.

E. Person of Ordinary Skill in the Art (“POSITA”)

A POSITA in the field of the '538 patent at the time of the earliest possible priority date (March 1, 2001) would have at least a Bachelor's degree in Computer Science, Computer Engineering, or the equivalent, and several years experience in the field of distributed systems, name services, or Internet content delivery. (Ex. 1003 [Almeroth] at ¶¶ 27-39.)

F. Claim Construction

Pursuant to 37 C.F.R. § 42.100(b), the claims in *inter partes* review are given the “broadest reasonable construction in light of the specification.” Except for the claim terms specifically discussed below, Petitioner requests that each of the various claim terms be given their plain meaning for the purposes of this IPR proceeding only.

1. “alternate path”

In the Litigation, Petitioner and Patent Owner dispute the term “alternate path.” (Ex. 1011 (Petitioner Limelight Networks, Inc.'s Preliminary Claim Construction Brief) at 19-21; Ex. 1012 (Patent Owner Akamai Technologies, Inc.'s

Opening Claim Construction Brief) at 7-10.) As shown below, Petitioner’s proposed construction from the Litigation is correct. Further, the Board need not construe “alternate path” because the prior art discloses the claimed “alternate path,” irrespective of which construction applies. *See Limelight Networks, Inc. v. Akamai Techs., Inc.*, IPR2016-01894, Paper 8 at 8 (PTAB Mar. 8, 2017) (“Only those terms that are in controversy need to be construed and only to the extent necessary to resolve the controversy. *See Vivid Techs. Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999).”).

The dispute over “alternate path” is whether the term means a “path other than a primary path already specified by the content delivery network,” as Petitioner has proposed, or whether no construction is necessary, as Patent Owner has proposed. (Ex. 1012 at 19.) Patent Owner has also proposed an alternate construction of “one of multiple paths between a given edge server and an origin server.” (*Id.*)

The ’538 patent describes a purported invention that uses a CDN’s server node locations as waypoints in what it calls “alternate paths” between an edge server and an origin server. The CDN can consider network conditions and select a new path—an alternate path—through its server nodes around the internet, to avoid congestion and outages. Petitioner’s proposed construction reflects this, while Patent Owner’s ignores it.

Specifically, the '538 patent states:

“The *present invention* identifies alternate paths from a CDN edge server to an origin server and uses these alternates to either improve the performance of content delivery, or to provide for fail-over in the event that the direct route is congested or otherwise unavailable. When an edge server contacts the origin server, the “direct” route typically is the route obtained through the Border Gateway Protocol (BGP). When the inventive routing technique is used, in contrast, alternate routes to the origin server are accessed by sending the request from the edge server to another intermediate (preferably CDN-managed) server/region before going to the origin.” (Ex. 1001 at 6:23-33 (emphasis added).)

Per the limiting language “the present invention,” the alternate paths are not just any path that happens to work, they are specifically ***identified*** as such by the CDN before they are used. (See *id.* at 7:4-16.) *Hill-Rom Services, Inc. v. Stryker Corp.*, 755 F. 3d 1367, 1372 (Fed. Cir. 2014). When the direct route cannot or should not be used, the CDN then uses this identified alternate route instead. The specification further states:

“a map of potential alternate routes is created [C]urrent ping data is used to identify the two best alternates . . . [and] each CDN edge server chooses among the direct route and the alternate routes based on the current actual performance. . . .” (*Id.* at 7:23-32.)

“The edge server then chooses among these alternate routes and the direct route by occasionally testing the performance of each connection, e.g., with a request for content.” (*Id.* at 7:64-67; *see also id.* at 9:7-46.)

These passages confirm that the primary and alternate paths are distinct specified paths, identified in advance and tested for performance or accessibility, so that the CDN can make a selection among the specified paths.

Patent Owner’s alternative proposal in the Litigation (“one of multiple paths between a given edge server and an origin server”) is nonsensical and appears to be for the purpose of allowing Patent Owner to argue in the Litigation that any path, including the default path, could be an “alternate path” as long as there may also be other possible paths. Indeed, Patent Owner’s alternative proposal from the Litigation nullifies the claim language (and plain meaning) of “alternate path.”

For at least these reasons, Petitioner’s proposal from the related Litigation is the correct interpretation of “alternate path.” Accordingly, Petitioner requests that the term “alternate path” be construed as “a path other than a primary path already specified by the content delivery network.”

2. “server node”

In the Litigation, Petitioner and Patent Owner dispute the term “server node.” (Ex. 1011 (Petitioner Limelight Networks, Inc.’s Preliminary Claim Construction Brief) at 21-23; Ex. 1012 (Patent Owner Akamai Technologies, Inc.’s

Opening Claim Construction Brief) at 10.) As shown below, Petitioner’s proposed construction from the Litigation is correct. Further, the Board need not construe “server node” because the prior art discloses the claimed “server node,” irrespective of which construction applies. *See Limelight Networks*, IPR2016-01894, Paper 8 at 8.


The dispute over “server node” in the Litigation is whether the term means a “datacenter containing the content delivery network edge servers for a particular geographic region,” as Petitioner has proposed, or whether no construction is necessary, as Patent owner has proposed. (Ex. 1011 at 21.) Patent Owner has also proposed an alternative construction of “a set of one or more CDN servers.” (*Id.*)

The ’538 patent purports to disclose a beneficial way of routing traffic between regions, from one part of the internet through a CDN’s own nodes rather than using default Internet paths:

“The inventive routing technique enables an edge server operating within a given CDN region to retrieve content (cacheable, non-cacheable and the like) more efficiently by ***selectively routing through the CDN’s own nodes***, thereby avoiding network congestion and hot spots. ***The invention*** thus enables an edge server to fetch content from an origin server through an intermediate CDN server or, more generally, ***enables an edge server within a given first region to fetch content from the origin server through an intermediate CDN***

region. As used herein, this routing through an intermediate server, node or region” (Ex. 1001 at 6:11-22 (emphasis added).)

Server nodes are waypoints through which traffic is routed in the invention. “Node” is used in the specification, such as to explain that “[a] content delivery network or ‘CDN’ is a network of geographically distributed content delivery **nodes** that are arranged for efficient delivery of content on behalf of third party content providers.” (*Id.* at 1:29-32.) But the support for what a “server node” actually **is**, or how it actually gets used, comes in Figure 4 and the related description. Figures 3 and 4 together show the difference between default CDN routing and the alternative routing purportedly invented by the ’538 Patent. Figure 4 shows the alternative routing between two datacenters, with a third datacenter used as a waypoint:



(*Id.* at FIGS. 3; 4.) In this regard, the ’538 Patent discloses:

“FIGS. 3-4 illustrate *the basic concept of routing through intermediate CDN regions according to the present invention*. In FIG. 3 an end user request is directed (e.g., through a DNS request routing mechanism) to an edge server located at a **datacenter 300** in Texas. In this case the content provider is operating mirrored origin servers at a **datacenter 302** on the East Coast, and at a **datacenter 304** on the West Coast. In a typical scenario, the edge server connects to the **datacenter 302** over a direct connection 306 if it becomes necessary to fetch content from the origin server. *In FIG. 4, in contrast, a pair of alternative routes is made available to the edge server. One alternative route 402 uses region 406, which is located intermediate of the datacenter 402 and the datacenter 400 to which the end user was originally directed. Another alternative route 408 uses region 410, which is located intermediate of the datacenter 404 and the datacenter 400.* (*Id.* at 6:43-59 (emphasis added).)

This confirms that the “server nodes” are the datacenters containing the CDN edge servers for a particular geographic region, which datacenters are the waypoints in the alternative routing scheme of the ’538 Patent. (*See id.* at 6:39-7:3; 9:47-10:24.)

Patent Owner’s proposed “alternative” construction of “a set of one or more CDN servers” is redundant of other claim language and, by focusing on the things that a server node *has*, neglects to explain what a server node *is*. Claim 1 already requires that “a server node comprises a cluster of content delivery network edge

servers.” (Ex. 1001 at Claim 1.) Saying that a server node is “a set of one or more CDN servers” is redundant of this existing claim requirement.

Thus, Petitioner’s proposal from the related Litigation is the correct interpretation of “server node.” Accordingly, Petitioner requests that the term “server node” be construed as “a datacenter containing the content delivery network edge servers for a particular geographic region.”

G. Overview of Prior Art

1. Overview of Speedera (Ex. 1004-Speedera)

Speedera is an issued U.S. Patent, filed on August 23, 2000, claiming priority to U.S. Provisional Application No. 60/166,906, filed November 22, 1999. (Ex. 1007 [Provisional App. No. 60/166,906].) Thus, Speedera is prior art under at least pre-AIA § 102(e).

Speedera discloses a CDN that services multiple customers. The CDN is organized into regional “points of presence” or “POPs,” and each POP contains a datacenter having multiple content delivery network edge servers (WebCaches). In other words, each Speedera POP includes a server node having a cluster of content delivery network edge servers (WebCaches). The POPs are organized by geographic region (WestUSA, WestEurope, EastUSA, etc.). FIG. 1 of Speedera is annotated and highlighted below:




FIG. 1

(Ex. 1006 at FIG. 1.) As shown, just like the conventional CDN recited in the '538 patent, Speedera's CDN includes multiple server nodes (orange), each having a cluster of content delivery network edge servers (purple). The server nodes (orange) make up a set of server nodes, and the server nodes are organized by region. The server nodes (orange) provide delivery of content on behalf of

participating content providers, such as by obtaining content from the content provider's origin server (yellow). (Ex. 1004-Speedera at 4:61-62; 5:16-32; 6:66-7:7.)

2. Overview of Ricciulli-470 (Ex. 1005-Ricciulli)

Ricciulli-470 is an issued U.S. Patent, filed on June 18, 1998. (Ex. 1005-Ricciulli at 1 (22).) Thus, Ricciulli-470 is prior art under at least pre-AIA § 102(e).

Ricciulli-470 discloses methods and apparatuses for determining whether to use a default path or an alternate path for communications between a source and a destination over the Internet:

“Methods and apparatus are disclosed for dynamically discovering and utilizing an optimized network path through overlay routing for the transmission of data. A determination whether to use a default network path or to instead use an alternate data forwarding path through one or more overlay nodes is based on real-time measurement of costs associated with the alternative paths, in response to a user request for transmission of message data to a destination on the network. Cost metrics include delay, throughput, jitter, loss, and security. *The system chooses the best path among the default forwarding path and the multiple alternate forwarding paths, and implements appropriate control actions to force data transmission along the chosen path.* No modification of established network

communication protocols is required.” (Ex. 1005-Ricciulli at Abstract (emphasis added).)

This routing system enables requests and responses for resources over the Internet to be routed over (a) a default path (if appropriate), or (b) over an alternate path if performance measurements indicate the alternate path is optimal as compared to a default path. (*Id.* at 2:27-57; 3:41-50.)

Specifically, existing network infrastructure and Internet communications protocols are used in combination with “overlay path modules” (software) to determine whether to use the default communication path or an alternative communication path. (*Id.* at 3:50-4:15). The overlay software may be installed at various locations in the network to facilitate “a large set of alternate paths with measurable communication costs around the Internet.” (*Id.* at 3:65-4:13.) In determining whether to use a default Internet communication path or an alternate path, the Ricciulli-470 software may determine “one or more cost/performance metrics.” (*Id.* at 5:6-7.)

“[When a] new path is found through the overlay network nodes that has better performance than the default path... information describing this improved path is preferably sent to each of the overlay network nodes 130 that will be involved in the new path. ... [A] preferred embodiment of [the] invention employs a mechanism to route packets through the new path in a transparent manner and without

modification of any of the default Internet communication mechanisms.” (*Id.* at 4:38-46.)

An alternate path would not generally be used if it would not outperform the default path: “The default path is preferably used in step 260, in those cases where no new path is found through the overlay network nodes that has better performance than the default path.” (*Id.* at 4:48-51.)

3. Overview of HTTP/1.1 (Ex. 1006-HTTP)

HTTP (“hypertext transfer protocol”) is an application protocol that facilitates the exchange and transfer of hypertext and serves as the foundation for the web. (Ex. 1003-Almeroth at ¶ 63.) The first documented version of HTTP was released in 1991 as HTTP v0.9. (*Id.*) HTTP/1.1 was in development as RFC 2068 by early 1996, at which point the HTTP/1.1 specification was garnering rapid adoption. (*Id.*) HTTP/1.1 was ultimately released and published in January 1997 as RFC 2068, and replaced by RFC 2616 (Ex. 1006-HTTP) in June 1999. (*Id.* at ¶ 63; Ex. 1006-HTTP at 1.) Accordingly, HTTP/1.1 is prior art under at least pre-AIA § 102(b).

HTTP/1.1 is an authentic, publicly accessible printed publication. For example, HTTP/1.1 was re-published in a compilation of RFCs in 2000. (Ex. 1008 at RFC2616-1 to 176; *see also* Exs. 1009-10 (records of the Library of Congress and WorldCat indicating publication in 2000).) Indeed, this compilation of RFCs

notes that HTTP/1.1 “is the most important RFC of this volume,” “it defines one of the most visible protocols on the Internet,” and “it is also one of the most widely used [RFCs].” (*Id.* at xviii.) Further, U.S. Patent No. 6,356,907 cites and describes RFC 2616. (Ex. 1013 at 1:28-30.) The ’907 patent was filed on July 26, 1999, one month after RFC 2616 indicates it was published. As another example, U.S. Patent No. 6,567,857 cites and describes RFC 2616. (Ex. 1014 at 5:25-35.) The ’857 patent was filed on July 26, 1999, also one month after RFC 2616 indicates it was published. As yet another example, a press release from July 7, 1999 from <http://www.ietf.org/> and <http://www.w3.org/> announced that “The World Wide Web Consortium (W3C) is pleased to recognize that HTTP/1.1, along with the accompanying authentication specification, has been approved by the Internet Engineering Steering Group (IESG) of the Internet Engineering Task Force (IETF) as a IETF Draft Standard.” (Ex. 1015 at 1.) The press release goes on to state “The RFC, based on draft 6 of the specification revision with some minor final editorial changes, is available as RFC 2616.” (*Id.* at 2.) Thus, the evidence of record indicates that HTTP/1.1 was published and publicly accessible in June 1999, as it states on its face. (Ex. 1006-HTTP at 1.) Petitioner is not aware of any circumstances that call into question the authenticity or public accessibility of HTTP/1.1, and indeed the distinctive pagination, formatting, publication date, and publication information within HTTP/1.1 all confirm that HTTP/1.1 is what it

purports to be. RFC 2616 also continues to be available online from the IETF at <https://www.ietf.org/rfc/rfc2616.txt>, the website of the organization that publishes requests for comment of this type, precisely as a POSITA would expect. (Ex. 1003-Almeroth at ¶¶ 64-65.)

HTTP/1.1 discloses cache-control techniques, including disclosing circumstances in which it might be inappropriate to cache a resource, and further explains why such controls are provided:

“Unless specifically constrained by a cache-control (section 14.9) directive, a caching system MAY always store a successful response (see section 13.8) as a cache entry, MAY return it without validation if it is fresh, and MAY return it after successful validation. If there is neither a cache validator nor an explicit expiration time associated with a response, we do not expect it to be cached, but certain caches MAY violate this expectation (for example, when little or no network connectivity is available). A client can usually detect that such a response was taken from a cache by comparing the Date header to the current time.

Note: some HTTP/1.0 caches are known to violate this expectation without providing any Warning.

However, in some cases it might be inappropriate for a cache to retain an entity, or to return it in response to a subsequent request. This might be because absolute semantic transparency is deemed necessary by the service author, or because of security or privacy considerations.

Certain cache-control directives are therefore provided so that the server can indicate that certain resource entities, or portions thereof, are not to be cached regardless of other considerations.” (Ex. 1006-HTTP at § 13.4 (Response Cacheability) (emphasis added).)

Examples of cache-control cacheability settings that are specified in HTTP/1.1 include the “private” and “no-store” control settings. (*Id.* at §§ 14.9.1; 14.9.2; Ex. 1003-Almeroth at ¶ 66.) These are described in detail in § 14.9 of HTTP/1.1, as well as below with respect to the Challenged Claims. In general, and as HTTP/1.1 indicates, a POSITA would have understood that content providers would utilize cache controls to prevent the caching of resources for a variety of reasons, including ensuring semantic transparency, for security reasons, or for privacy reasons, among others. (Ex. 1003-Almeroth at ¶ 66.)

VI. SPECIFIC GROUNDS FOR PETITION

A. GROUND 1: Speedera in view of Ricciulli-470 and HTTP/1.1 renders the Challenged Claims obvious

1. The scope and content of the prior art

For Ground 1, the prior art consists of Speedera (Ex. 1004), Ricciulli-470 (Ex. 1005), and HTTP/1.1 (Ex. 1006).

2. Motivation to Combine and Rationale for Combining Speedera, Ricciulli-470, and HTTP/1.1

Speedera discloses a CDN as outlined above in § V.G.1 (Overview of Speedera). The Speedera CDN includes edge servers called WebCaches that are

regionally distributed in POPs. (Ex. 1004-Speedera at FIG. 1.) On occasion, Speedera's WebCache will not have a resource that an end user has requested, and will need to communicate with a content provider customer's origin server to obtain the resource from its source in order to fulfill the customer's request. (*Id.* at 13:36-14:7; *see also* limitation 1.1.a.) In this regard, Speedera provides:

“Initially, WebCache server 280 determines whether it has the requested static content in memory, step 600. ... If the requested static content is not cached within WebCache server 280, WebCache server 280 sends a request (260) for the static content to WebServer server 290, step 610.

In response, WebServer server 290 determines whether it has the requested static content on disk, or the like, 620. ... If the requested static content is not stored within WebServer server 290, WebServer server 290 determines sends a request (270) for the static content to a customer web server, step 630. ...

In this example, customer web server 40 provides the static content to WebServer server 290, which in turn stores the static content to disk, or the like, step 640.

Once the data is stored within WebServer server 290, the static content is copied into the memory of WebCache server 280, step 650. Next, WebCache server 280 returns the static content to the user, step 660.” (*Id.* at 10:43-11:7.)

Thus, the WebCache communicates with the customer's origin server in order to request and receive resources that it has not cached from the customer's origin server. After it has obtained the resource, the WebCache serves the resource to the user. (*Id.* (“...WebCache server 280 returns the static content to the user, step 660.”))

Regarding communications between the POPs (10, 20, 30), customers (40, 45), and users (150, 160), Speedera provides that such communications may be implemented via the Internet. (*Id.* at 4:49-51 (“In FIG. 1, the servers 10-50, and ISP 60 are typically coupled via a computer network 70, such as the Internet.”).) (*See also* Ex. 1003-Almeroth at ¶¶ 67-71 (explaining that the purpose of the Speedera CDN, and CDN's in general, was to make content delivery faster and optimize response times between requests for content resources and delivery of those resources).) However, Speedera is silent as to how to optimize Internet communication pathways between the POPs (10, 20, 30), customers (40, 45), and users (150, 160).

Ricciulli-470 discloses techniques for optimizing communication paths over the Internet. As noted above in § V.G.2 (Overview of Ricciulli-470), Ricciulli-470 explains that default communication paths are not necessarily *optimal* communication paths within a network. Accordingly, Ricciulli-470 discloses a technique for assessing alternate paths that travel through special intermediate

nodes in a network, and further discloses selecting and using those alternate paths when performance metrics indicate that an alternate path would provide better performance:

“More particularly, the preferred embodiment of our invention creates a new forwarding path from endpoint A to endpoint B only when: (1) an end-to-end communication is requested between A and B (per step 200 of FIG. 2), and (2) a path id discovered through the overlay routing network that provides better performance than the default Internet route (per steps 210-215 of FIG. 2).” (Ex. 1005-Ricciulli at 4:65-5:4; *see also id.* at Abstract; 2:27-4; 2:49-57; FIGS. 1-2, 4A-4B.)

Ricciulli-470 further discloses that its technique may be employed with already-existing infrastructure and that computers (nodes) of a network may be used to facilitate the optimal communication:

“1) An overlay network of alternate routing mechanisms is constructed on top of the existing Internet routing mechanisms to find and exploit available resources. The overlay routing mechanism is completely transparent and separate from the Internet routing protocols and is preferably deployed throughout some small, but widely distributed, portion of the Internet as a distributed user application. FIG. 1 exemplifies the concept. ... ***Overlay network nodes 130a-n utilize existing network transmission lines and infrastructure, via network links 135a-n, to create a virtual topology.*** The overlay network preferably includes a number of computing devices such as nodes 130a-n that cooperate to provide forwarding

paths overlaid over an underlying network. Overlay network nodes preferably communicate using existing, established Internet protocols and thus do not require any modifications to current standards. Each overlay node 130 preferably includes overlay path module 150, and either the source or destination node similarly includes overlay path module 120; these components are programmed and operable to combine available IP protocols in order to provide additional functionality for exploiting overlay routing when it is advantageous to do so, as described below in detail. In our preferred embodiment, the overlay network includes a set of ordinary computer systems (e.g., Pentium-based PC's) co-located at Internet sites such as major ISP locations. These PC's each include overlay path modules 150 in the form of additional custom software modules, for purposes of the present invention, operable to measure and record connection cost information and optimal forwarding path information, as described in greater detail herein. The overlay network is a virtual network; in other words, although it uses new additional hardware (co-located boxes), it preferably uses the existing network and cabling infrastructure for all communications. Effectively, this overlay network provides a large set of alternate paths with measurable communication costs around the Internet. Clients can use these alternate paths if they are found to provide better service to their desired destinations than would a default path.” (Ex. 1005-Ricciulli at 3:34-4:14.)

A POSITA would have recognized the benefit of using the path optimization techniques of Ricciulli-470 to improve the communication paths used within the

Speedera CDN. (Ex. 1003-Almeroth at ¶¶ 72-74 (citing Ex. 1005-Ricciulli at 1:58-2:12, which shows that the routing techniques of Ricciulli-470 is particular useful in CDNs).) Ricciulli-470 provides a detailed explanation of how to implement its routing technique in an Internet environment using existing network architecture, such as Speedera's CDN, and explains how its technique can improve performance in a network. (Ex. 1003-Almeroth at ¶¶ 75-76.) A POSITA would have found it natural, predictable, and obvious to implement the Ricciulli-470 routing technique in the Speedera CDN to ensure that the optimal communication pathway between the POPs (10, 20, 30), customers (40, 45), and users (150, 160) was utilized. In doing so, a POSITA would have found it obvious to include/use Ricciulli-470's software, called "overlay network module" (120, 150), at least at each of the Speedera POPs (10, 20, 30), Speedera's Operation Center (50), and the customer POP servers (40, 45, unlabeled). A POSITA would have found this obvious because Ricciulli-470 teaches to employ such overlay network modules (120, 150) at high-traffic areas (e.g., ISP's) and at the source or destination. (*Id.* at 3:65-4:1.) Doing so would provide "a large set of alternate paths with measurable communication costs around the Internet." (*Id.* at 4:10-12.) As Ricciulli-470 notes, "ordinary computer systems" and preexisting "custom software modules" are used to implement the "overlay network nodes." (*Id.* at 3:65-4:10.) Thus, it would be a matter of routine skill to implement the overlay network at the

Speedera POPs (10, 20, 30), the Speedera Operation Center (50), and the customer POP servers (40, 45, unlabeled), among other locations. (Ex. 1003-Almeroth at ¶¶ 77-78.)

Riccuilli-470 discloses that an alternate communication path, rather than a default path, is used when a path through the overlay nodes provides better performance:

“If a new path is found through the overlay network nodes that has better performance than the default path, at steps 230 and 235, then at step 240 information describing this improved path is preferably sent to each of the overlay network nodes 130 that will be involved in the new path. At step 250, a preferred embodiment of our invention employs a mechanism to route packets through the new path in a transparent manner and without modification of any of the default Internet communication mechanisms. Preferably, a form of IP encapsulation is used, as described below in more detail. The default path is preferably used in step 260, in those cases where no new path is found through the overlay network nodes that has better performance than the default path.” (*Id.* at 4:38-52 (emphasis added).)

Accordingly, in the combined system of Speedera and Riccuilli-470, when a network communication is required, e.g., in the event of a “cache miss” at WebCache (e.g., 80, 90, 280, 290), the overlay path module located at the applicable Speedera POP would use the overlay network to determine whether to

use a default pathway or an alternate pathway for requesting and receiving content from the customer's origin server. (Ex. 1003-Almeroth at ¶¶ 79-80.)

A POSITA would also know that the WebCaches of Speedera would not always have a requested resource. (Ex. 1003-Almeroth at ¶ 81.) This is expressly noted in Speedera. (Ex. 1004-Speedera at 10:40-50.) Cache controls are built into the hypertext transfer protocol, and they restrict the ability of caching servers to cache certain resources. (Ex. 1006-HTTP at § 14.9 (disclosing cache-control mechanisms that prevent caching in several ways).) For instance, cache controls were used for reasons related to privacy, security, semantic transparency, or other design considerations. (*Id.* at § 13.4 (Response Cacheability).) In view of these well-known cache-controls described in the HTTP/1.1 specification, a POSITA would have understood that a WebCache in the Speedera CDN would sometimes need to request content from a customer's origin server. One conventional reason this request would be needed was due to the use of cache control mechanisms by content providers that prevent caching of certain resources. (*Id.* at ¶ 82.) For example, a content provider customer would commonly have used cache controls like the "private" and "no-store" controls to prevent the Speedera CDN from caching the applicable static content of Speedera, such as any of Speedera's "movies, audio, images, text, plug-in program data, programs, and the like" that would normally be cached by the CDN. (Ex. 1003-Almeroth at ¶¶ 82-83 (citing

Ex. 1004-Speedera at 15:14-19).) A POSITA would also have recognized that the optimized communications path disclosed by Speedera in view of Ricciulli-470 (described above) would be used to retrieve content from a customer's origin server, such as when content is not cached due to an HTTP/1.1 cache control. (*Id.* at ¶¶ 82-83.)

3. Challenged Claims

[CLAIM 1.P.a] 1. A method operative in a content delivery network having a set of server nodes organized into regions,

In the '538 patent, a region comprises one or more content servers and may be co-located with a POP. (Ex. 1001 at FIG. 1; 4:40-47; 6:6-11; *see also* § V.A.) As shown in the Claim Construction section (§ V.F), above, “server nodes” should be construed as “datacenters containing the content delivery network edge servers for a particular geographic region.” Further, per part 1.P.b of the preamble, addressed below, server nodes comprise a “cluster of content delivery network edge servers.”

Speedera discloses this part of the preamble, whether or not it is a limitation. Specifically, Speedera discloses a CDN having a set of server nodes organized into regions, as illustrated in the annotated version of FIG. 1:




FIG. 1

As shown, the Speedera CDN includes a set of server nodes (orange) that are organized by region. Speedera discloses the claimed “set of server nodes,” i.e., it discloses a set of “datacenters containing the content delivery network edge servers for a particular geographic region.” In particular, the POPs shown above include datacenters that contain the CDN edge servers (WebCaches) for particular

geographical regions (e.g., WestUSA, EastUSA, and the like). Likewise, under Akamai's proposed construction from the Litigation, the POPs of Speedera contain the claimed "server nodes" because they contain "a set of one or more CDN servers," i.e., the sets of WebCaches in each POP.

The POPs of the Speedera system are organized into regions, e.g., a West USA region, an East USA region, and a West Europe region. (*Id.* at FIG. 1.) These regions comprise multiple CDN edge servers, as detailed below with respect to part 1.P.b of the preamble. Hence, Speedera discloses the claimed "a set of server nodes organized into regions" under any reasonable construction.

Thus, Speedera discloses part 1.P.a of the preamble.

[1.P.b] wherein the server nodes provide delivery of content on behalf of participating content providers, wherein a server node comprises a cluster of content delivery network edge servers, the method comprising:

In the '538 patent, "content delivery network edge servers" include servers comprised of basic computer components that serve content from locations around the Internet:

"A representative CDN content server is a Pentium-based caching appliance running an operating system (e.g., Linux, Windows NT, Windows 2000) and having suitable RAM and disk storage for CDN applications and content delivery network content (e.g., HTTP content, streaming media and applications). Such content servers are sometimes referred to as 'edge' servers as they are located at or near

the so-called outer reach or ‘edges’ of the Internet.” (Ex. 1001 at 4:47-55.)

Speedera discloses this part of the preamble, whether or not it is a limitation. In particular, FIG. 1 shows that the regions include clusters of WebCache servers. In the annotated version of FIG. 1 above, the individual WebCache servers are highlighted purple. As shown, each region (e.g., WestUSA, WestEurope, EastUSA, etc.) includes multiple WebCache servers, which constitute a “cluster of content delivery network edge servers” and are part of the server node. (See Ex. 1004-Speedera at 4:45-53 (describing the embodiment of FIG. 1).) The WebCache servers of the server node provide delivery of content on behalf of participating content providers:

“Customers using the POP network, according to one embodiment, can store content such as HTML, images, video, sound, software, or the like on the network for fast and highly available access by clients (end users). The network also provides load balancing and high availability for servers outside the network. Customers with generated content, such as search engines, auctions and shopping carts, can use the latter feature to add their own content servers to the network.” (*Id.* at 2:9-17 (emphasis added); *see also id.* at 13:48-52.)

“Initially, WebCache server 280 determines whether it has the requested static content in memory, step 600. In one embodiment of the present invention, this step may include determining whether the file requested is resident in memory. If the requested static content is

not cached within WebCache server 280, WebCache server 280 sends a request (260) for the static content to WebServer server 290...

...

Once the data is stored within WebServer server 290, the static content is copied into the memory of WebCache server 280, step 650. Next, WebCache server 280 returns the static content to the user, step 660.” (Ex. 1004-Speedera at 10:43-11:7 (emphasis added); *see also id.* at 2:1-8, 4:45-53; 13:7-10; 13:19-21.)

As shown above, the WebCache servers deliver content from geographically distributed locations around the Internet. Speedera also explains that the servers, like the claimed “edge” servers, are comprised of basic computer components. (See, *e.g.*, *id.* at 7:21-64 & FIG. 2 (explaining that typical servers include a processor, such as an Intel processor, as well as memory and an operating system).)

Thus, Speedera discloses part 1.P.b of the preamble.

[1.1.a] for each given server node and a content provider origin server, identifying at least one alternate path to be used for edge server-to-content provider origin server communications,

Speedera in view of Ricciulli-470 discloses this limitation. As explained above, Speedera discloses the claimed server nodes and a content provider origin server. (See, *e.g.*, Ex. 1004-Speedera at FIG. 1; 5:54-64 (describing “content provider origin servers,” i.e., customer web servers 130 and 135); 9:42-49 (same).)

Speedera also discloses that a WebCache server may not have a requested resource. Speedera states that in such instances, the WebCache communicates with a customer's origin server via a WebServer, which Speedera discloses may be part of the same software/program and on the same server as the WebCache. (Ex. 1004-Speedera at 13:36-40.) Speedera explains:

“Initially, WebCache server 280 determines whether it has the requested static content in memory, step 600. ... If the requested static content is not cached within WebCache server 280, WebCache server 280 sends a request (260) for the static content to WebServer server 290, step 610.

In response, WebServer server 290 determines whether it has the requested static content on disk, or the like, 620. ... If the requested static content is not stored within WebServer server 290, WebServer server 290 determines sends a request (270) for the static content to a customer web server, step 630. ...

In this example, customer web server 40 provides the static content to WebServer server 290, which in turn stores the static content to disk, or the like, step 640.

Once the data is stored within WebServer server 290, the static content is copied into the memory of WebCache server 280, step 650. Next, WebCache server 280 returns the static content to the user, step 660.” (*Id.* at 10:43-11:7.)

Thus, the WebCache communicates with the customer's origin server in order to request and receive resources that it has not cached from the customer's origin server. In turn, the WebCache serves the resource to the user. (*Id.* (“...WebCache server 280 returns the static content to the user, step 660.”))

Although Speedera discloses communications between edge servers and origin servers, it does not specify which “path” through the Internet such communications would take. As a POSITA would have recognized, there are a vast number of available paths across the global Internet, and it was obvious to utilize existing techniques to identify non-default routes between two servers in a manner that optimizes performance. (Ex. 1003-Almeroth at ¶¶ 89-92.) Indeed, this is expressly disclosed and obvious in view of Ricciulli-470.

As explained above in the Motivation to Combine and Rationale for Combining section, Ricciulli-470 discloses an “overlay network” that includes a default path and at least one alternate path for communications between nodes in a network:

“The invention further provides steps and means for measuring the costs of transmitting the message from the source to the destination along one or more non-default, alternative paths passing through a special group of intermediate nodes. In a preferred embodiment, those intermediate nodes are referred to as an ‘overlay network’. An optimized path for sending the requested transmission is ultimately

selected by comparing the default cost against the alternative costs.” (Ex. 1005-Riccillui at 2:33-41 (emphasis added); *see also id.* at Abstract; 2:27-33; 2:49-57.)

The reference further explains:

“FIG. 1 exemplifies the concept. Nodes 100 and 160 are, respectively, source and destination nodes for an intended communication on a network such as the Internet. These nodes are connected to the underlying network via transmission links 110 and 170, respectively. ... Overlay network nodes 130a-n utilize existing network transmission lines and infrastructure, via network links 135a-n, to create a virtual topology. The overlay network preferably includes a number of computing devices such as nodes 130a-n that cooperate to provide forwarding paths overlaid over an underlying network. Overlay network nodes preferably communicate using existing, established Internet protocols and thus do not require any modifications to current standards. Each overlay node 130 preferably includes overlay path module 150, and either the source or destination node similarly includes overlay path module 120; these components are programmed and operable to combine available IP protocols in order to provide additional functionality for exploiting overlay routing when it is advantageous to do so, as described below in detail. In our preferred embodiment, the overlay network includes a set of ordinary computer systems (e.g., Pentium-based PC's) co-located at Internet sites such as major ISP locations. These PC's each include overlay path modules 150 in the form of additional custom software modules, for purposes of the present invention, operable to measure and record

connection cost information and optimal forwarding path information, as described in greater detail herein. The overlay network is a virtual network; in other words, although it uses new additional hardware (co-located boxes), it preferably uses the existing network and cabling infrastructure for all communications. Effectively, this overlay network provides a large set of alternate paths with measurable communication costs around the Internet. Clients can use these alternate paths if they are found to provide better service to their desired destinations than would a default path.” (*Id.* at 3:41-4:14.)

Thus, Ricciulli-470 discloses a technique for utilizing existing infrastructure with corresponding special software (overlay path module) to optimize a network path between a source and a destination.

As explained in the Motivation to Combine and Rationale for Combining section, a POSITA also would have considered it obvious to improve communications over the Speedera CDN using the communication path optimization techniques of Ricciulli-470. (Ex. 1003-Almeroth at ¶¶ 93-94.) Ricciulli-470 provides a detailed explanation of how to implement its routing technique in an Internet environment using existing network architecture, such as Speedera’s CDN, and explains how its technique can improve performance in a network. (*Id.*) A POSITA would have found it natural, predictable, and obvious to implement the Ricciulli-470 routing technique in the Speedera CDN to ensure that the optimal communication pathway between the POPs (10, 20, 30), customers

(40, 45), and users (150, 160) was utilized. In doing so, a POSITA would have found it obvious to include Riccuilli-470's software, called "overlay network module" (120, 150), at each of the Speedera POPs (10, 20, 30), Speedera's Operation Center (50), and the customer POP servers (40, 45, unlabeled) because Riccuilli-470 teaches to employ such overlay network modules (120, 150) at high-traffic areas, such as ISP locations (*id.* at 3:65-4:1), and at the source or destination. Doing so would help provide "a large set of alternate paths with measurable communication costs around the Internet." (*Id.* at 4:10-12.) As Riccuilli-470 notes, "ordinary computer systems" and preexisting "custom software modules" are used to implement the "overlay network nodes" (*id.* at 3:65-4:10), so it would be a matter of routine skill to implement such overlay network at the Speedera POPs (10, 20, 30), the Speedera Operation Center (50), and the customer POP servers (40, 45, unlabeled), among other locations. (Ex. 1003-Almeroth at ¶¶ 94-95.)

Regarding communication, Riccuilli-470 provides:

"If a new path is found through the overlay network nodes that has better performance than the default path, at steps 230 and 235, then at step 240 information describing this improved path is preferably sent to each of the overlay network nodes 130 that will be involved in the new path. At step 250, a preferred embodiment of our invention employs a mechanism to route packets through the new path in a

transparent manner and without modification of any of the default Internet communication mechanisms. Preferably, a form of IP encapsulation is used, as described below in more detail. The default path is preferably used in step 260, in those cases where no new path is found through the overlay network nodes that has better performance than the default path.” (*Id.* at 4:38-52.)

Accordingly, in the combined system of Speedera and Ricciulli-470, when a network communication is required, such as when there is a “cache miss” at WebCache (e.g., 80, 90, 280, 290), the overlay path module located at the applicable Speedera POP would determine, using the overlay network, whether to use a default pathway or an alternate pathway when requesting the content from the customer’s origin server. (Ex. 1003-Almeroth at ¶ 95.)


As noted above, in the Litigation, there is a claim construction dispute regarding “alternate path.” Under Petitioner’s proposed construction of “alternate path” (i.e., “a path other than a primary path already specified by the content delivery network”), the alternate paths of the combined system of Speedera in view of Ricciulli-470 are the claimed “alternate path” because they are predetermined overlay network paths that are evaluated in order to determine if they would perform better than a default path. (Ex. 1005-Ricciulli at 4:38-51.) Thus, they are “paths other than a primary path already specified by the CDN.” Likewise, under Patent Owner’s proposed construction from the Litigation, these alternate paths are

paths between a WebCache and a customer's origin server that are not a default path. Hence, they are "one of multiple paths between a given edge server and an origin server." Thus, under any reasonable construction, the alternate paths of the combined system of Speedera in view of Ricciulli-470 meet the requirements of limitation 1.1.b. (Ex. 1003-Almeroth at ¶ 96.)

Thus, Speedera in view of Ricciulli-470 discloses this limitation.


[1.1.b] the at least one alternate path including at least one other content delivery network edge server that is external to the given server node;

The '538 patent illustrates the difference between using a default path and an alternate path that is routed through another edge server in FIGS. 3 and 4:



As shown above, the route in FIG. 3 provides a default route between a datacenter 300 and an origin server 302; in FIG. 4, alternate routes are presented that route through other datacenters in the network. (Ex. 1001 at 6:41-65.)

Speedera in view of Ricciulli-470 discloses this limitation. As described above with respect to limitation 1.1.a, Speedera in view of Ricciulli-470 discloses using alternate paths that route through overlay network nodes in an “overlay network.” This is illustrated in FIGS. 4A-4B, showing that alternate paths involving nodes in an overlay network are examined, and an optimized path is selected for routing communications between a source node “A” and a destination node “B:”



(See also Ex. 1005-Ricciulli at 6:66-7:19 (explaining the process for assessing alternate paths through the overlay network).) As explained previously, a POSITA would have found it obvious to include Ricciulli-470’s software called “overlay network module” (120, 150), at each of the Speedera POPs (10, 20, 30), Speedera’s Operation Center (50), and the customer POP servers (40, 45,

unlabeled) because Riccuilli-470 teaches to employ such overlay network modules (120, 150) at high-traffic areas, such as ISP locations. (*Id.* at 3:65-4:1.) A POSITA would find it obvious to install and use Riccuilli-470's software on the Speedera WebCaches because Riccuilli-470 discloses using the overlay software at the source and destination computers, and the WebCaches are both sources and destinations. (Ex-1003 Almeroth at ¶¶ 97-99.) Thus, the alternate paths provided for by the combination of Speedera and Riccuilli-470 meet the requirement of "the at least one alternate path including at least one other content delivery network edge server." Further, given that Riccuilli-470 discloses to use a large number of alternative nodes, it would be obvious to use the overlay software at the many POP's of Speedera. In doing so, communications between sources and destinations in the combined Speedera-Riccuilli-470 system would include at least one Speedera edge server that is "external to the given source node." (Ex-1003 Almeroth at ¶ 99.)

Thus, Speedera in view of Ricciulli-470 discloses this limitation.

[1.2.a] upon receipt at a particular edge server in the given server node of a request for an object that is not available for delivery from the particular edge server, issuing a new request for the object from the particular edge server over the at least one alternate path,

Speedera in view of Ricciulli-470 discloses this limitation. Specifically, as explained above with respect to limitation 1.1.a, Speedera discloses receiving

requests for content that is not cached at a WebCache. (Ex. 1004-Speedera at 10:43-11:7.) This discloses “**upon receipt at a particular edge server in the given server node of a request for an object that is not available for delivery from the particular edge server.**”

As further explained with respect to the same limitation (1.1.a), the WebCache then communicates a request for the content that was not cached to the customer’s origin server. This discloses “**issuing a new request for the object from the particular edge server.**”

Finally, as explained above with respect to limitations 1.1.a and 1.1.b, it was obvious to improve Speedera in view of Ricciulli-470 by utilizing an optimized, alternate network path. Accordingly, in the combined system of Speedera and Ricciulli-470, when a cache miss occurs, and when the default path is not optimal, the system issues “**a new request for the object from the particular edge server over the at least one alternate path.**”

Thus, Speedera in view of Ricciulli-470 discloses this limitation.

[1.2.b] wherein the object has an associated cacheability setting, wherein the cacheability setting indicates that the object should not be cached in the particular edge server;

The background “Description of the Related Art” section of the ’538 patent explains that there is often a need in a CDN for an edge server to retrieve content from a customer’s origin server:

“Edge servers operating within a content delivery network often have need to return to a content provider’s origin server, e.g., when requested data is not available at the server or is otherwise stale, to obtain non-cacheable content, and the like.” (Ex. 1001 at 2:30-34.)

Thus, the ’538 patent admits that it was known and often the case that an edge server needed to request content from a content provider’s origin server, such as “non-cacheable contact,” and that there were multiple reasons for doing so.

Speedera in view of HTTP/1.1 discloses this limitation. As explained above with respect to limitation 1.1.a, Speedera discloses the need for a WebCache server to retrieve content from a customer’s origin server because it is not stored in the WebCache or at the WebCache’s POP location. (Ex. 1004-Speedera at 10:43-11:7.)

As shown by HTTP/1.1, and as admitted by the ’538 patent, it was commonly known to a POSITA that one reason an edge server, such as Speedera’s WebCaches, may not have a stored copy of a resource is because of a cacheability setting, i.e., a “cache control:”

“Certain cache-control directives are therefore provided so that the server can indicate that certain resource entities, or portions thereof, are not to be cached regardless of other considerations.” (Ex. 1006-HTTP at § 13.4 (Response Cacheability) (emphasis added).)

HTTP/1.1 further discloses several specific cache-control directives that expressly indicate that an object cannot be cached on an edge server. For example, the private directive prevents an object from being cached by a shared cache:

“private

Indicates that all or part of the response message is intended for a single user and MUST NOT be cached by a shared cache. This allows an origin server to state that the specified parts of the response are intended for only one user and are not a valid response for requests by other users. A private (non- shared) cache MAY cache the response.”

(*Id.* at § 14.9.1 (What is Cacheable).)

Likewise, the “no-store” directive indicates that an origin server may prevent caching entirely:

“no-store

The purpose of the no-store directive is to prevent the inadvertent release or retention of sensitive information (for example, on backup tapes). The no-store directive applies to the entire message, and MAY be sent either in a response or in a request. If sent in a request, a cache MUST NOT store any part of either this request or any response to it. If sent in a response, a cache MUST NOT store any part of either this response or the request that elicited it. This directive applies to both nonshared and shared caches. ‘MUST NOT store’ in this context means that the cache MUST NOT intentionally store the information in nonvolatile storage, and MUST make a best-effort attempt to remove the information from volatile storage as promptly

as possible after forwarding it.” (*Id.* at § 14.9.2 (What May be Stored by Caches).)

Thus, it was commonly known to a POSITA that cache-control directives would cause an edge server (the WebCaches of Speedera, in particular) to avoid caching a resource, such as the “static content stored on the web caches” of Speedera, which “may include movies, audio, images, text, plug-in program data, programs, and the like” that would normally be cached in the absence of a cache control. (Ex. 1003-Almeroth at ¶¶ 105. This would in turn require the WebCache to request the resource from a corresponding customer’s origin server as described in Speedera. (Ex. 1003-Almeroth at ¶¶ 104-06.)

Thus, Speedera in view of HTTP/1.1 discloses this limitation.

[1.3] receiving the object over the at least one alternate path; and

Speedera in view of Ricciulli-470 discloses this limitation. As explained above with respect to limitation 1.2.a, in the combined system of Speedera and Ricciulli-470, a WebCache server would issue a request for a resource that was not cached to a customer’s origin server over either a default or an alternate path. Ricciulli-470 further explains that, when an alternative pathway was used to issue the request, the same alternative pathway would be used to respond to the request:

“As a third example, we will now consider the case of a message that requests a return reply (such as an http request to get a file), once

again in the context of the multi-hop forwarding path through overlay nodes 130a and 130b as in the previous example. ...

In our preferred embodiment, the last overlay node on a forwarding path performs the task of masquerading, in order to allow bi-directional use of the overlay forwarding path. ...

Returning to step 510, module 150b forwards the encapsulated message to the next overlay node on an optimized path to node 100, by accessing path information previously stored at step 240 (in this case, the path information is of course just the inverse of the optimized overlay path for communications being sent from source 100 to destination 160)....” (Ex. 1005-Ricciulli at 9:9-62 (emphasis added).)

Thus, in the combined system of Speedera and Ricciulli-470, when a WebCache requests content from a customer’s origin server over an optimized alternate path, the response to that request is returned over the same alternate path. (Ex. 1003-Almeroth at ¶¶ 107-08.)

Thus, Speedera in view of Ricciulli-470 discloses this limitation.

[1.4] serving the object in response to the request.

Speedera discloses this limitation. Specifically, as shown above with respect to limitation 1.1.a, after the WebCache receives the requested resource from the customer’s origin server, the WebCache servers the content to the end user. (Ex.

1004-Speedera at 10:43-11:7 (“...WebCache server 280 returns the static content to the user, step 660.”).)

Thus, Speedera discloses this limitation. Further, Speedera in view of Riciulli-470 and HTTP/1.1 discloses every limitation of claim 1 and renders obvious claim 1.

[CLAIM 2]

Claim 2 is substantially identical to claim 1. To facilitate the Board’s review, Petitioner provides a chart below, comparing claim 1 to claim 2, with emphasis added to highlight the minor differences:

<u>[1.P.a]</u> A method operative in a content delivery network having a set of server nodes organized into regions,	<u>[2.P.a]</u> A method operative in a content delivery network having a set of server nodes organized into regions,
<u>[1.P.b]</u> wherein the server nodes provide delivery of content on behalf of participating content providers, wherein a server node comprises a cluster of content delivery network edge servers, the method comprising:	<u>[2.P.b]</u> wherein the server nodes provide delivery of content on behalf of participating content providers, wherein a server node comprises a cluster of content delivery network edge servers, the method comprising:
<u>[1.1.a]</u> for each given server node and a content provider origin server, identifying at least one alternate path to be used for edge server-to-content provider origin server communications,	<u>[2.1.a]</u> for each given server node and a content provider origin server, identifying at least one alternate path to be used for edge server-to-content provider origin server communications,

[1.1.b] the at least one alternate path including at least one other content delivery network edge server that is external to the given server node;	[2.1.b] the at least one alternate path including at least one other content delivery network edge server that is external to the given server node;
[1.2.a] upon receipt at a particular edge server in the given server node of a request for an object that is not available for delivery from the particular edge server, issuing a new request for the object from the particular edge server over the at least one alternate path;	[2.2.a] upon receipt at a particular edge server in the given server node of a request for an object that is not available for delivery from the particular edge server, issuing a new request for the object from the particular edge server over the at least one alternate path;
[1.2.b] <i>wherein the object has an associated cacheability setting, wherein the cacheability setting indicates that the object should not be cached in the particular edge server;</i>	
[1.3] receiving the object over the at least one alternate path; <i>and</i>	[2.3] receiving the object over the at least one alternate path;
[1.4] serving the object in response to the request.	[2.4] serving the object in response to the request; <i>and</i>
	[2.5] <i>identifying a new alternate path as a result of changing Internet conditions.</i>

In view of the chart above, the only limitation that is in claim 2 that was not previously described above with respect to claim 1 is limitation 2.5. Limitation 2.5 is analyzed below.

[2.5] identifying a new alternate path as a result of changing Internet conditions.

Speedera in view of Ricciulli-470 discloses this limitation. Identifying an initial alternate path is described above with respect to limitation 1.1.a. Ricciulli-470 also explains that the alternate paths are determined based on “real-time measurement of costs associated with the alternative paths:”

“Methods and apparatus are disclosed for dynamically discovering and utilizing an optimized network path through overlay routing for the transmission of data. A determination whether to use a default network path or to instead use an alternate data forwarding path through one or more overlay nodes is based on real-time measurement of costs associated with the alternative paths, in response to a user request for transmission of message data to a destination on the network.” (Ex. 1005-Ricciulli at Abstract (emphasis added).)

In addition, Ricciulli-470 discloses that new alternative paths will be calculated, e.g., after a period of time passes and the initial alternative path becomes too old, in order to account for changing Internet conditions:

“Referring again to FIG. 2, following the detection of an improved overlay forwarding path at steps 230 and 235, information describing the improved forwarding path is stored at step 240 by the originator of the path query (e.g., typically module 110 of source node 100) and by each of the overlay nodes involved in the improved path. The path information is preferably stored at each overlay node 130 in a table or

the like, so that when given a specified destination endpoint, an overlay node on the forwarding path can retrieve the address of the next node on the non-default path to whom the message should be forwarded. In a preferred embodiment of our invention, this information may be deleted by the overlay nodes if no end-to-end communication happens between A and B for more than a predefined amount of time, or if the overlay forwarding path's performance is believed to become worse than the default Internet path—or simply after some specified amount of time passes, so that subsequent communication requests will result in discovery and measurement of overlay paths based on the new, current state of the network.” (*Id.* at 7:29-48 (emphasis added).)

Thus, Speedera in view of Ricciulli-470 discloses this limitation. Further, Speedera in view of Ricciulli-470 and HTTP/1.1 discloses every limitation of claim 2 and renders obvious claim 2.

[CLAIM 3] The method as described in claim 2 wherein the alternate path is valid for a given time period.

Speedera in view of Ricciulli-470 discloses this limitation. In particular, Ricciulli-470 discloses (as reproduced above with respect to claim 2) that an alternate path may only be valid for a given period of time. (*Id.* at 7:29-48.)

Thus, Speedera in view of Ricciulli-470 discloses this limitation. Further, Further, Speedera in view of Ricciulli-470 and HTTP/1.1 discloses every limitation of claim 3 and renders obvious claim 3.

[CLAIM 4] The method as described in claim 2 wherein the at least one alternate path is other than a default Border Gateway Routing (BGP) path.

Speedera in view of Ricciulli-470 discloses this limitation. Regarding BGP, the '538 patent states:

“When an edge server contacts the origin server, the ‘direct’ route typically is the route obtained through the Border Gateway Protocol (BGP). When the inventive routing technique is used, in contrast, alternate routes to the origin server are accessed by sending the request from the edge server to another intermediate (preferably CDN-managed) server/region before going to the origin.” (Ex. 1001 at 6:27-33.)

As explained above with respect to claims 1 and 2, the Ricciulli-470 alternate paths are identified via “custom software modules” (Ex. 1005-Riciulli at 4:3) rather than default techniques, of which “Border Gateway Routing (BGP)” was one example known to a POSITA. (Ex. 1003-Almeroth at ¶ 115.) Thus, the alternate paths of Ricciulli-470 are not “a default Border Gateway Routing (BGP) path.”

Thus, Speedera in view of Ricciulli-470 discloses this limitation. Further, Further, Speedera in view of Riciulli-470 and HTTP/1.1 discloses every limitation of claim 4 and renders obvious claim 4.

[CLAIM 5]

Claim 5 is substantially identical to claim 1. To facilitate the Board's review, Petitioner provides a chart below, comparing claim 1 to claim 5, with emphasis added to highlight the minor differences:

<u>[1.P.a]</u> A method operative in a content delivery network having a set of server nodes organized into regions,	<u>[5.P.a]</u> A method operative in a content delivery network having a set of server nodes organized into regions,
<u>[1.P.b]</u> wherein the server nodes provide delivery of content on behalf of participating content providers, wherein a server node comprises a cluster of content delivery network edge servers, the method comprising:	<u>[5.P.b]</u> wherein the server nodes provide delivery of content on behalf of participating content providers, wherein a server node comprises a cluster of content delivery network edge servers, the method comprising:
<u>[1.1.a]</u> for each given server node and a content provider origin server, identifying at least one alternate path to be used for edge server-to-content provider origin server communications,	<u>[5.1.a]</u> for each given server node and a content provider origin server, identifying at least one alternate path to be used for edge server-to-content provider origin server communications,
<u>[1.1.b]</u> the at least one alternate path including at least one other content delivery network edge server that is external to the given server node;	<u>[5.1.b]</u> the at least one alternate path including at least one other content delivery network edge server that is external to the given server node,
	<u>[5.1.c]</u> <i>wherein the alternate path is determined by the particular edge server as a result of performing a</i>

	<u>performance test;</u>
[1.2.a] upon receipt at a particular edge server in the given server node of a request for an object that is not available for delivery from the particular edge server, issuing a new request for the object from the particular edge server over the at least one alternate path,	[5.2.a] upon receipt at a particular edge server in the given server node of a request for an object that is not available for delivery from the particular edge server, issuing a new request for the object from the particular edge server over the at least one alternate path;
[1.2.b] <u>wherein the object has an associated cacheability setting, wherein the cacheability setting indicates that the object should not be cached in the particular edge server;</u>	
[1.3] receiving the object over the at least one alternate path; and	[5.3] receiving the object over the at least one alternate path; and
[1.4] serving the object in response to the request.	[5.4] serving the object in response to the request.

In view of the chart above, the only limitation that is in claim 5 that was not previously described above with respect to claim 1 is limitation 5.1.c. Limitation 5.1.c is analyzed below.

[5.1.c] wherein the alternate path is determined by the particular edge server as a result of performing a performance test;

Speedera in view of Ricciulli-470 discloses this limitation. As explained above relative to limitation 1.1.b, it would be obvious to use Riccuilli's overlay

software module with the Speedera WebCaches (edge servers). Ricciulli-470 explains that the alternate paths are determined as a result of performing performance tests:

“Each overlay node 130 preferably includes overlay path module 150, and either the source or destination node similarly includes overlay path module 120; these components are programmed and operable to combine available IP protocols in order to provide additional functionality for exploiting overlay routing when it is advantageous to do so, as described below in detail. In our preferred embodiment, the overlay network includes a set of ordinary computer systems (e.g., Pentium-based PC’s) co-located at Internet sites such as major ISP locations. These PC’s each include overlay path modules 150 in the form of additional custom software modules, for purposes of the present invention, operable to measure and record connection cost information and optimal forwarding path information, as described in greater detail herein.” (Ex. 1005-Ricciulli at 3:59-4:6 (emphasis added).)

“2) Alternate, improved forwarding paths through the overlay network nodes are discovered on demand. The process is outlined in the flow diagram of FIG. 2. ... In response, at step 210, if the threshold cost is exceeded by the default path for the requested communication, then at steps 220-225 (as described in more detail below) overlay path module 120 broadcasts queries which are received and processed by one or more of overlay nodes 130a-n; overlay path modules 150 cooperate to discover alternative paths through overlay nodes 130 and

to measure the costs of such alternate paths. Measuring the cost of data transmission along a given path (per steps 210, 220) is generally performed using conventional techniques, depending on the cost metric to be measured. For example, connection delay time can easily be measured by ‘pinging’ the destination of interest from the source node of interest. At steps 230-235 these alternative paths are compared to the existing Internet route or previously chosen overlay routes.” (Ex. 1005-Ricciulli at 4:16-37 (emphasis added).)

Thus, in the combined system of Speedera and Ricciulli-470, when the WebCache (“the particular edge server”) uses the Ricciulli-470 overlay software module, “the alternate path” would be determined “by the particular edge server as a result of performing a performance test.” (*Id.*; Ex. 1003-Almeroth at ¶¶ 118-20.)

Thus, Speedera in view of Ricciulli-470 discloses this limitation. Further, Further, Speedera in view of Ricciulli-470 and HTTP/1.1 discloses every limitation of claim 5 and renders obvious claim 5.

VII. CONCLUSION

For the foregoing reasons, Petitioner respectfully request that the *inter partes* review of the ’538 Patent be instituted as the Petition establishes a reasonable likelihood of prevailing with respect to the Challenged Claims. Petitioner further respectfully requests that Claims 1-5 be cancelled as unpatentable.

Respectfully submitted,
GREENBERG TRAURIG, LLP

Date: December 19, 2017

By: /s/ Heath J. Briggs
Heath J. Briggs
Registration No. 54,919
1200 17th Street, Suite 2400
Denver, CO 80202
Phone: (303) 572-6500
Fax: (303) 572-6540
Counsel for Petitioner

CERTIFICATE OF SERVICE

The undersigned certifies that a true and correct copy of the Petition together with all exhibits identified in the above Table of Exhibits and Petitioner's Power of Attorney, have been served on the Patent Owner via Priority Mail Express or by means at least as fast and reliable as Priority Mail Express on the below date, at the following addresses:

Law Office of David H. Judson
15950 Dallas Parkway
Suite 225
Dallas, TX 75248

Respectfully submitted,
GREENBERG TRAURIG, LLP

Date: December 19, 2017

By: /s/ Heath J. Briggs
Heath J. Briggs
Registration No. 54,919
1200 17th Street, Suite 2400
Denver, CO 80202
Phone: (303) 572-6500
Fax: (303) 572-6540
Counsel for Petitioner