

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

LIMELIGHT NETWORKS, INC.,
Petitioner,

v.

AKAMAI TECHNOLOGIES, INC.,
Patent Owner.

Case IPR2018-00310
U.S. Patent 8,194,538

**PATENT OWNER'S PRELIMINARY RESPONSE PURSUANT TO
37 C.F.R. § 42.107**

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. BACKGROUND TECHNOLOGY	7
A. Retrieving Content Directly from Origin Servers	7
B. Retrieving Content from CDNs	8
III. THE ’538 PATENT	13
IV. THE PRIOR ART CITED IN THE PETITION	16
A. The Speedera Reference	16
B. The Ricciulli Reference	21
V. CLAIM CONSTRUCTION	23
A. “alternate path”	24
B. “server node”	27
VI. THE PRIOR ART CITED IN THE PETITION FAILS TO DISCLOSE THE LIMITATIONS OF THE ’538 PATENT	28
A. Speedera and Ricciulli Do Not Disclose the Claimed “Alternate Path”	29
B. Petitioner’s Arguments Regarding the Claimed “Alternate Path” Fail	35
VII. CONCLUSION	37

TABLE OF AUTHORITIES

	Page(s)
Federal Cases	
<i>Allergan, Inc. v. Sandoz Inc.</i> , 796 F.3d 1293 (Fed. Cir. 2015)	33
<i>Arendi S.A.R.L. v. Apple Inc.</i> , 832 F.3d 1355 (Fed. Cir. 2016)	35
<i>In re Fulton</i> , 391 F.3d 1195 (Fed. Cir. 2004)	6, 33
<i>In re Gordon</i> , 733 F.2d 900 (Fed. Cir. 1984)	33
<i>In re Magnum Oil Tools Int'l</i> , 829 F.3d 1364 (Fed. Cir. 2016)	36, 37
<i>Perfect Web Techs., Inc. v. InfoUSA, Inc.</i> , 587 F.3d 1324 (Fed. Cir. 2009)	35
<i>Personal Web Techs., LLC v. Apple, Inc.</i> , 848 F.3d 987 (Fed. Cir. 2017)	35
<i>Polaris Indus., Inc. v. Arctic Cat, Inc.</i> , 882 F.3d 1056 (Fed. Cir. 2018)	6
<i>In re Zurko</i> , 258 F.3d 1379 (Fed. Cir. 2001)	35

Patent Owner Akamai Technologies, Inc. ("Patent Owner") submits the following Preliminary Patent Owner Response ("POPR") to the Petition filed by Limelight Networks, Inc. ("Petitioner") requesting *inter partes* review of claims 1-5 of U.S. Patent No. 8,194,538 ("the '538 patent").

I. INTRODUCTION

The '538 patent relates to a content delivery network ("CDN") having "server nodes" that are distributed geographically and that include content servers called "edge servers." *See* '538 patent at 1:29-32 ("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."), claim 1 ("a server node comprises a cluster of content delivery network edge servers") (Ex. 1001). The edge servers store "objects" (e.g., webpage images, text, videos, etc.) on behalf of content providers, such as NBC. *Id.* at 5:32-39. When an end user desires to retrieve a webpage, for example, he or she enters the website's name—such as www.nbc.com—into his or her computer. *Id.* at 5:39-45. Afterwards, the computer sends a request for the webpage to an edge server in one of the server nodes. *See id.* ("the content server 200 receives end user requests for content, determines whether the requested object is present ... [and] serves the requested object ... (if it is present)").

Occasionally, an edge server does not have a copy of a requested object and

therefore cannot immediately deliver the object to the end user computer. *See id.* at 5:45-48 (“a cache miss can occur when the requested object is not in cache” on the edge server). In that case, the edge server may need to retrieve the object from the content provider’s “origin server” (e.g., NBC’s server). *Id.* at 5:45-50 (when there is a “cache miss” “it may be necessary for the edge server to contact a content provider origin server to fetch the requested object”). Before the ’538 patent, if an edge server in a given server node needed to retrieve that object from the origin server, typically it would send a request over a “direct” path between the edge server and the origin server. For example, this direct path typically was determined by the Internet’s default routing protocol—known as the “Border Gateway Protocol” or “BGP.” *Id.* at 6:27-29 (“When an edge server contacts the origin server, the ‘direct’ route typically is the route obtained through the Border Gateway Protocol (BGP).”). However, “the Internet ... is far from an ideal communications medium” and the direct path is often “congested or faulty.” *See id.* at 1:60-2:1. As a result, the object retrieved along this “direct path” from the edge server to the origin server might be delayed or even lost. *See id.* (“When a path in the network is either congested or faulty, packets sent over that path can get dropped or delayed.”).

To overcome this and related problems, the ’538 patent implements an “inventive routing technique” in which the request for the object is sent from the

edge server to *another intermediate* edge server in another server node before going to the origin server. *Id.* at 6:29-33 (request is sent “from the edge server to another intermediate ... server/region before going to the origin”). Sending the request to the origin server via another edge server in a different server node—as opposed to using the “direct path” (determined via the Internet’s default routing protocol)—often is faster and more reliable, and can avoid network congestion. As the ’538 patent disclosure teaches, this is true even where the “direct path” itself is shorter in terms of network distance. *See id.* at 6:34-38 (“While one might assume that adding the intermediate step would reduce performance, it frequently improves performance, because CDN [edge] servers and regions [i.e., server nodes] typically are well-connected, and the indirect route can often bypass network congestion.”).

Each challenged claim expressly requires sending a request for an object from “a particular edge server in [a] given server node” to the origin server along an “alternate path” that includes “at least one other ... [intermediate] edge server that is external to the given server node”:

- “identifying at least one *alternate path* to be used for *edge server-to-content provider origin server communications*, the at least one *alternate path including* at least *one other* content delivery network *edge server that is external to [a] given server node ...*”; and
- “upon receipt at a *particular edge server in the given server node* of a

request for an object ..., issuing a new request for the object *from the particular edge server over the at least one alternate path...*”

Id. at claim 1 (16:1-11), claim 2 (16:24-34), claim 5 (16:50-55 and 58-62), claims 3 and 4 (each requiring “[t]he method as described in claim 2”).

Petitioner relies on prior art that does not disclose this functionality. In particular, Petitioner argues that the challenged claims are invalid over the combination of U.S. Patent No. 6,405,252 (“Speedera”) (Ex. 1004), U.S. Patent No. 6,275,470 (“Ricciulli”) (Ex. 1005), and Hypertext Transfer Protocol – HTTP/1.1, RFC 2616 (“HTTP/1.1”) (Ex. 1006). Petitioner relies on Speedera and Ricciulli for the limitations requiring issuing a request from an edge server in one node along an alternate path that includes an edge server in another node. *See* Pet. at 38-45, 52, 56. But neither reference, alone or in combination, discloses an alternative path that is both (1) from “a particular edge server” to an origin server, and (2) includes another edge server that is external to the server node of the particular edge server, as the relevant claim elements require.

First, although Speedera allegedly discloses communications between an edge server and an origin server, Petitioner concedes—correctly—that Speedera “does not specify which ‘path’ through the Internet such communications would take.” Pet. at 38. Accordingly, the Petition does not establish that Speedera teaches a path that (1) is between a “particular edge server in [a] given server

node” and an origin server and (2) includes “at least one other ... edge server that is external to the given server node,” as the claims require. *See* ’538 patent at claim 1 (16:1-11), claim 2 (16:24-34), claim 5 (16:50-55 and 58-62), claims 3 and 4 (each requiring “[t]he method as described in claim 2”) (Ex. 1001).

Second, Ricciulli does not fill this gap. Ricciulli does not address a “content delivery network having a set of server nodes” or, as a consequence, “edge server-to-content provider origin server communications.” Instead, Ricciulli simply discloses a network of nodes (not a “content delivery network”) that is located between a “source node” and a “destination node” and that includes ordinary computer systems located at Internet Service Providers (“ISPs”). *See* Ricciulli at 3:65-4:9 (“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.***”), *id.* (“[A]lthough it uses new additional hardware (co-located boxes), it preferably uses the existing network and cabling infrastructure for all communications.”) (Ex. 1005). In fact, Ricciulli only briefly mentions CDNs and when it does, it criticizes them. Ricciulli contends that CDNs are “***limited or ineffective*** with respect to dynamically generated content” and “***entail significant costs*** for using large-capacity computer servers to physically store extra copies of massive volumes of data, and for maintaining synchronization among the various repositories of frequently updated content.” *Id.* at 1:58-2:6. *See*

In re Fulton, 391 F.3d 1195, 1201 (Fed. Cir. 2004) (reference “teaches away” where it “criticize[s], discredit[s], or otherwise discourage[s]” solution claimed by patent); *see also Polaris Indus., Inc. v. Arctic Cat, Inc.*, 882 F.3d 1056, 1069 (Fed. Cir. 2018) (reference “teaches away” where a POSITA “would be discouraged from following the path set out in the reference, or would be led in a direction divergent from the path that was taken by the applicant” (citation omitted)). Ricciulli therefore does not disclose an “alternate path” between a “particular edge server” and an origin server that includes “at least one ... other edge server that is external to [a] given server node [including the particular edge server].”

Third, even assuming that a person of ordinary skill in the art (“POSITA”) would have been motivated to combine Speedera and Ricciulli (which Patent Owner does not concede), the combination of these references does not disclose a request that is sent from an edge server in one node along an alternate path that includes an edge server that is external to that node—as required by all challenged claims. ’538 patent at claim 1 (16:1-11), claim 2 (16:24-34), claim 5 (16:50-55 and 58-62), claims 3 and 4 (each requiring “[t]he method as described in claim 2”) (Ex. 1001). Instead, Petitioner’s arguments reduce to the contention—without evidence or support—that a POSITA would have combined Speedera and Ricciulli and then added the missing feature. This allegation is tantamount to an impermissible reconstruction based on the words of the claims themselves, and not

evidence, and thus constitutes impermissible hindsight. It is an insufficient basis to establish obviousness. Because the Petition fails to make any showing of how the cited prior art discloses this claimed feature or how this difference would have been obvious to a POSITA at the time of the invention, the Petition should not be instituted.

II. BACKGROUND TECHNOLOGY

A. Retrieving Content Directly from Origin Servers

Content on the Internet—such as webpages, graphics, videos, and text—is created by companies called “content providers” (such as NBC) and is stored on servers owned by the content providers. *See* U.S. Patent No. 6,108,703 (“the ’703 patent”) at 1:15-24, 1:36-40 (Ex. 2001). These content provider servers are also known as “origin servers” because they are the original source of the content. *See* ’538 patent at 1:32-36 (Ex. 1001).

Traditionally, end users used computers to retrieve webpages directly from origin servers. *See* ’703 patent at 1:26-34 (“Use of an HTML-compatible browser (e.g., Netscape Navigator or Microsoft Internet Explorer) at a client machine involves specification of a link via the URL. In response, the client makes a request to the server identified in the link”) (Ex. 2001). For instance, if a user wanted to retrieve NBC’s webpage, she would enter “www.nbc.com” into her end user computer. *See id.* Afterwards, the computer would send a request for the


NBC webpage directly to NBC's origin server, which in turn would deliver the webpage back to the computer. *See id.* at 1:29-33 (“the client makes a request to the Server identified in the link and, in return, receives a document or other object formatted according to HTML”). Delivering webpages in this manner, however, was slow and unreliable because most content providers maintained a limited number of origin servers, which were often located far from end user computers. *See id.* at 3:42-49, 15:14-32 (describing deficiencies with traditional content delivery method).

B. Retrieving Content from CDNs

To address this problem, companies like Akamai developed content delivery networks (“CDNs”) that delivered content to end users on behalf of content providers. *See* '538 patent at 1:52-59 (a CDN “service provider establishes business relationships with content publishers and acts on behalf of their origin server sites to provide a distributed delivery system”) (Ex. 1001). As described in the '538 patent, CDNs typically include “server nodes” that are distributed geographically on a wide basis. *See id.* at 1:29-32 (“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.”). A typical CDN server node includes CDN content servers—called “edge servers”—that (as described in one embodiment) are located at the “edge” of

the Internet close to the access networks (e.g., Comcast) by which end user computers access the Internet. *See id.* at 5:32-39, 4:52-55 (“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.”). These edge servers selectively obtain, server and sometimes store content—known as “objects”—on behalf of content providers, like NBC. *Id.* at 5:32-39. Edge servers are clustered into “server nodes.” *See, e.g., id.* at claim 1 (“a server node comprises a cluster of content delivery network edge servers”); Fig. 1 (showing edge servers arranged into server nodes 102).

When an end user computer sends a request for an object (e.g., a portion of NBC’s webpage), the request is not sent to NBC’s origin server. *Id.* at 5:39-45. Instead, as illustrated in Figure 3 of the ’538 patent (reproduced below), it is routed to a CDN edge server located in a server node, preferably one that is close to the end user computer:




Id. at Fig. 3 (annotated).

Namely, when the end user computer (purple) sends a request for an object (green), such as NBC's webpage object, the request is routed to an edge server in the server node (which in one embodiment is located in a data center ("DC") 300) that is close to the end user computer and that attempts to respond to the request.

See id. at Fig. 3, 6:45-50 ("end user request is directed ... to an edge server located at a datacenter 300"); *see also id.* at 5:39-45 ("the [edge server] receives end user requests for content"). If the edge server has an up-to-date copy of the requested object, it delivers that object to the end user computer. *See id.* at 5:39-45 ("the [edge server] receives end user requests for content, determines whether the requested object is present ... [and] serves the requested object ... (if it is present)"). And, because in this exemplary embodiment the edge server is much

closer to the end user computer than the origin server 302 (such as NBC's origin server), it can deliver the requested object much more quickly.

Sometimes, however, an edge server does not have an up-to-date copy of the object requested by an end user computer. *See id.* at 5:45-48 (“a cache miss can occur when the requested object is not in cache” on the edge server). In that case, the edge server may need to retrieve the object from the content provider's origin server 302 (e.g., NBC's server). *Id.* at 2:30-34 (“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.”), 5:45-50 (when a “cache miss” occurs, “it may be necessary for the edge server to contact a provider origin server to fetch the requested object”). In those circumstances, before the '538 patent, the edge server would typically send a request for the object to the origin server 302 over the Internet via a “direct path” 306 (determined via the Internet's default routing protocol) between the edge server and the origin server 302:



Id. at 6:27-29, Fig. 3 (annotated), 6:50-52 (“In a typical scenario, the edge server connects to the datacenter [*sic*: origin server] 302 over a direct connection 306 if it becomes necessary to fetch content from the origin server.”).

However, “the Internet ... is far from an ideal communications medium.”

Id. at 1:60-2:1. As a result:


- The direct path between the edge server and the origin server may be “**congested**.” *Id.* at 1:64-65.
- The direct path may be “**faulty**” due, for instance, to “routing misconfiguration, software and hardware problems, and even fiber cuts.” *Id.* at 1:64-65, 2:11-14.

- And the direct path may be configured to use the “*cheapest*,” rather than the “*fastest*,” pathway. *Id.* at 2:1-4 (“[e]conomical issues can also interfere with the best” pathway).

The object retrieved along this direct path therefore may be delayed or even lost, despite the fact that the default Internet routing may represent a shorter network path. *Id.* at 1:60-2:1 (“When a path in the network is either congested or faulty, packets sent over that path can get dropped or delayed.”).

III. THE '538 PATENT

To overcome this problem, the '538 patent implements an “inventive routing technique” in which a request for an object is sent from an edge server to another intermediate edge server in another server node before it goes to the origin server. '538 patent at 6:29-33 (“When the inventive routing technique is used ... alternate routes to the origin server are accessed by sending the request from the edge server to another intermediate ... server/region before going to the origin.”), claim 1 (alternate path “includ[es] at least one other ... edge server that is external to the given server node”) (Ex. 1001). Figure 4 of the '538 patent, reproduced below, shows an example of this inventive technique:



Id. at Fig. 4 (annotated). As shown, when the end user computer (purple) sends a request for an object (green), the request is routed to a particular edge server in a nearby server node—in this example, a server in a node in the data center (“DC”) 400. *See id.* If that edge server does not have the requested object, it can retrieve the object from the origin server. *Id.*; *see also id.* at 5:45-50 (in the event of a “cache miss” “it may be necessary for the edge server to contact a content provider origin server to fetch the requested object”). The request for the object is sent to the origin server over an “alternate path” 403 that includes another, intermediate edge server located in another, external server node (such as the one located in the data center to the right of 400 in the figure above). *Id.* at Fig. 4, 6:54-57 (“One alternative route 402 [*sic*: 403] uses [the region to the right of datacenter 400], which is located intermediate of the datacenter [with the origin server] 402 and the datacenter 400 to which the end user was originally directed.”); *see also id.* at

6:29-33 (“alternate routes to the origin server are accessed by sending the request from the edge server to another *intermediate* ... *server*[] before going to the origin”). Then, in response to this request, the origin server sends a copy of the object back to the particular edge server along the same alternative path, and then onto the end user computer. *Id.* at 16:16 (noting the that particular edge sever server “receiv[es] the object over the at least one alternate path”), 16:4-5 (noting that “the at least one alternate path includ[es the] other [intermediate] edge server”).

By sending the request and the object along the alternative path via the intermediate edge server in the external server node, the '538 patent is able to quickly route information between the edge server and origin server by avoiding network congestion and hot spots, even where the alternative path is longer than the “direct path” that would be taken if the Internet’s default protocol (BGP) were used to determine the path. *Id.* at 2:59-63 (“The invention enables an edge server operating within a given CDN region to retrieve content ... more efficiently by selectively routing through the CDN’s own nodes, thereby avoiding network congestion and hot spots.”), 6:34-38 (“While one might assume that adding the intermediate step would reduce performance, it frequently improves performance, because CDN servers and regions [i.e., server nodes] typically are well-connected, and the indirect route can often bypass network congestion.”).

All of the challenged claims require this inventive technique. In particular, each challenged claim requires sending a request for an object from “a particular edge server” in a given server node to the origin server along an “alternate path” that includes another intermediate edge server that is external to the given node:

- “identifying at least one *alternate path* to be used for *edge server-to-content provider origin server communications*, the at least one *alternate path including* at least one other content delivery network *edge server that is external to [a] given server node ...*”; and
- “upon receipt at a *particular edge server in the given server node* of a request for an object ..., issuing a new request for the object *from the particular edge server over the at least one alternate path....*”

Id. at claim 1, claim 2, claim 5, claims 3 and 4 (each requiring “[t]he method as described in claim 2”).

IV. THE PRIOR ART CITED IN THE PETITION

A. The Speedera Reference

Petitioner relies on Speedera and Ricciulli in an attempt to show an “alternate path ... used for edge server-to-content provider origin server communications [that] includ[es] at least one other content delivery network edge

server that is external to the given server node,” as set forth in Elements [1.1.a] and [1.1.b] of the '538 patent.¹ See Pet. at 36-45.

Figure 1 of Speedera, reproduced below, shows a CDN:




FIG. 1

¹ Petitioner only relies on the HTTP/1.1 reference (Ex. 1006) to allegedly show features of Element [1.2.b] of claim 1 (“associated cacheability setting, wherein the cacheability setting indicates that the object should not be cached in the particular edge server”). See Pet. at 46-49. Petitioner does not rely on this reference to show features of the claim limitations at issue in this POPR.

Speedera at Fig. 1 (annotated) (Ex. 1004).

As shown, the CDN includes customer web servers 40 and 45, which contain web servers 130 and 135 (such as NBC's origin server), and a plurality of "Point of Presence" servers or "POPs" 10, 20, and 30 (blue). *Id.* at Fig. 1, 4:46-47 ("FIG. 1 illustrates a series of 'point of presence' (POP) servers, 10, 20, and 30"), 5:54-59 ("As shown in FIG. 1, customer web servers 40 and 45 typically include web servers 130 and 135 and servers 140 and 145. In this embodiment, web servers 130 and 135 provide static content, such as text and graphics, further servers 140 and 145 provide non-static content, or content that changes in response to a user query, or the like."). Each of the POPs includes WebCache servers (e.g., servers 80, 90, and 280 (pink)). *Id.* at Fig. 1, 4:56-57 ("each POP server will include a plurality of dedicated servers"), 4:59-63 ("In this example ... server 80 and 90 operate as a cache according to WebCache software"). Figure 1 further illustrates end user computers 150 and 160 (brown) that are located in an "ISP cluster" 60. *Id.* at Fig. 1, 4:45-49, 6:1-3.

Petitioner alleges that Speedera's customer web server 130 corresponds to an "origin server," the WebCache servers (e.g., servers 80, 90, and 280) correspond to "edge servers," and the POPs 10, 20, and 30 contain "server nodes organized into regions." *See* Pet. at 24-25, 32-33, 36.

Speedera discloses that, when an end user computer (e.g., computer 150) sends a request for content, the request is routed to a WebCache server (e.g., server 280) in an alleged region (e.g., one of the POPs, such as POP 30). If the WebCache server 280 does not have the requested content, it may need to retrieve that content from a customer web server 130 (i.e., the alleged origin server). *See* Speedera at 10:43-65 (Ex. 1004); Pet. at 36 (alleging customer web server 130 corresponds to an origin server). In particular, if the requested content “is not cached within WebCache [server] 280,” the WebCache server 280 requests the content from WebServer server 290, which is in the same POP (e.g., the same alleged server node) as the WebCache server 280. Speedera at 10:43-65 (Ex. 1004). Then, if the requested content “is not stored within WebServer server 290, WebServer server 290” sends a request “for the ... content to [the] customer web server” 130. *Id.* The customer web server 130, in turn, “provides the ... content to WebServer server 290.” *Id.* at 11:1-3. The content then is “copied into the memory of WebCache server 280,” *id.* at 11:4-6, and the “WebCache server 280 returns” the requested content “to the user,” *id.* at 11:6-7.


WebCache server 280 and WebServer server 290 sit in the same alleged “server node.” Petitioner does not contend that the WebServer server 290 corresponds to the claimed “other ... edge server” that is included in the alternate path and “external to the given server node.” *See* Pet. at 17 (asserting that “each

POP contains a datacenter having multiple content delivery network edge servers (WebCaches)”), 18 (describing the WebCaches in a given POP as a “cluster of delivery edge servers”), 37-38, 43-45. Instead, Petitioner alleges that the WebCache 280 constitutes the “particular edge server” in the claims, and that the POP 30 includes the claimed “given server node” where the “particular edge server” is located. Pet. at 17, 18, 33-34 (contending that “the POPs [in Speedera] include datacenters that contain the CDN edge servers (WebCaches)” and that “the POPs in 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.”), 37-38.

Petitioner concedes that Speedera is silent on the path by which the WebCache server 280 (i.e., the alleged particular edge server) communicates with the customer web server 130 (i.e., the alleged origin server). See Pet. at 26 (“*Speedera is silent* as to how to optimize Internet communication pathways between the POPs (10, 20, 30) [containing WebCache 280], customers, (40, 45) and users (150, 160).”), 38 (Speedera “*does not specify* which ‘path’ through the Internet such communications [i.e., communications from the alleged edge server to the alleged origin server] would take.”); see also *id.* at 24-25 (contending that Speedera’s “WebCaches” (e.g., 280) constitute “edge servers”), 36 (describing “customer web server[] 130” as an “origin server”).

B. The Ricciulli Reference

As shown in Figure 1 (reproduced below), Ricciulli discloses a system comprising a “source node” 100, a “destination node” 160, and a network of nodes that connects the source and destination nodes. Ricciulli at 3:41-44 (“Nodes 100 and 160 are, respectively, source and destination nodes for an intended communication on a network such as the Internet.”), 3:34-59 (network of nodes connects source and destination nodes) (Ex. 1005).



Id. at Fig. 1. The network of nodes comprises “overlay network nodes” 130*a-n*. *See id.* at 3:53-56 (“network preferably includes a number of computing devices such as nodes 130*a-n* that cooperate to provide forwarding paths overlaid over an

underlying network”). These “overlay network nodes” 130a-n include “ordinary computer systems (e.g., Pentium-based PC’s) co-located at Internet sites such as major ISP [Internet Service Provider] locations,” *id.* at 3:65-4:1; use “existing network transmission lines and infrastructure,” *id.* at 3:50-53; and “preferably communicate using existing, established Internet protocols,” *id.* at 3:56-59. *See also id.* at 4:8-10 (“[A]lthough [the “overlay network”] uses new additional hardware (co-located boxes), it preferably uses the existing network and cabling infrastructure for all communications.”).

When the source node 100 sends a message to the destination node 160, the message can be routed through the overlay network nodes 130a-n. *See id.* at 8:2-41. For instance, to route a message from the source node 100 to the destination node 160 via overlay node 130a, the message is first sent from the source node 100 to the overlay node “130a using the default network path for 100 [to] 130a,” preferably “by means of default network routing mechanisms.” *Id.* at 8:12-23. The message is then forwarded “from [the overlay node] 130a to [the destination node] 160 using the default network path for 130a [to] 160.” *Id.* at 8:12-16. In this way, “the original message gets from client ... 100 [i.e., source node 100] to destination node 160, along an optimized non-default path passing through overlay node 130a.” *Id.* at 8:34-37. In addition, the responsive message from the destination node 160 can be routed through the nodes 130a-n to the source node

100. *See id.* at 9:17-50 (“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.”).

Ricciulli does not address a “content delivery network having a set of server nodes” or “edge server-to-content provider origin server communications.” In fact, Ricciulli only briefly mentions CDNs, and when it does, it criticizes them. *Id.* at 1:58-2:6 (CDNs are “*limited or ineffective* with respect to dynamically generated content” and “*entail significant costs* for using large-capacity computer servers to physically store extra copies of massive volumes of data, and for maintaining synchronization among the various repositories of frequently updated content”).

V. CLAIM CONSTRUCTION

Petitioner asks the Board to construe two terms—“alternate path” and “server node”—and proposes the same constructions of those terms that it advanced in the related litigation.² *See* Pet. at 10-17. However, as noted below, construing these terms is not necessary to resolve the issues disputed in this POPR and deny institution of this Petition. In any event, Petitioner’s proposed

² The related litigation is *Akamai Technologies, 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).

constructions are incorrect. But under either Petitioner or Patent Owner's construction for both terms, the '538 patent is not obvious based on Speedera in light of Ricciulli-470 and HTTP/1.1.

A. “alternate path”

As Patent Owner explained in the related litigation, no construction of “alternate path” is needed, because the plain meaning of that term in the context of the '538 patent's claims and specification is “one of multiple paths between a given edge server and an origin server.” Akamai Technologies, Inc.'s Opening Claim Construction Brief at 7 (“Claim Construction Br.”) (Ex. 1012). Petitioner's proposed construction—a “path other than a primary path already specified by the content delivery network” (Pet. at 11)—by contrast, is contrary to the embodiments set forth in the '538 patent.

The '538 patent confirms that an “alternate path” is “one of the multiple paths between a server and an origin server.” In particular, the '538 patent describes that the CDN service provider *first* determines a set of multiple “alternate paths” and *then* selects an optimal path from among those alternate paths. *See, e.g.*, '538 patent at 6:39-7:3 (“Preferably, given CDN edge servers are configured with a plurality (e.g., at least two) alternate routes to use in addition to the direct route to the origin. ... According to the invention, one of these routes is selected as optimal....”), 7:23-32 (describing selecting paths in three phases,

beginning with creating “a map of potential alternate routes” which are then narrowed down to a “best” path), 8:3-8 (the “races” described in the ’538 patent “consist of having the edge server forward a client request through ... a given number (e.g., three (3)) [of] available routes simultaneously and observing which route performed the best in serving the request. The information is then used to choose the primary route for future requests.”) (Ex. 1001). The ’538 patent thus teaches that the “primary” path is *one of* the alternate paths. *See id.*; Claim Construction Br. at 7-8 (Ex. 1012).

Limelight’s proposed construction, which requires the alternate path to be “other than” a primary path, and that the primary path be “already specified” (*see* Pet. at 11), suffers from three fundamental flaws. *First*, by requiring the “alternate path” to be a path “other than a primary path *already specified*,” Petitioner requires that the primary path be specified *prior to* the claimed “alternate path.” *See* Pet. at 12, 42 (suggesting that the “alternate path” is “predetermined”). Yet, as explained above, the “primary” path is selected from *among* a set of alternate paths and, therefore, the primary path is not “already specified.” *See* ’538 patent at 8:3-8, 7:23-32 (Ex. 1001).

Second, for similar reasons, the claimed “alternate path” is not “other than” a primary path, as Petitioner argues. Instead, the ’538 patent makes clear that any “alternate path” may *become* the “primary path” depending on network conditions.

See id. at 8:1-11 (“[Performance] races consist of having the edge server forward a client request through ... a given number (e.g., three (3)) [of] available routes simultaneously and observing which route performed the best in serving the request. The information is then used to choose the primary route for future requests. Moreover, that route preferably remains the primary route until the next race is conducted.”); *see also id.* at 3:3-42 (“It is yet another more specific object of the invention to provide a routing service that predicts a best path for a data transfer between a source location ... and a target location ... by analyzing some performance metric common to a set of possible routes.”). Accordingly, Petitioner’s construction, which defines “alternate path” as “other than” a “primary path” is inconsistent with the specification, and should be rejected. *See Claim Construction Br.* at 9 (Ex. 1012).

Third, Petitioner’s argument improperly conflates the “primary” path—i.e., the best path—with the wholly separate concept of the “direct” or “default” path. *See Pet.* at 12-13. The “direct” path is the path from the edge server to the origin server that the Internet’s routing protocol would choose by default, unrelated to the ’538 patent’s invention. *See* ’538 patent at 6:27-34 (“When an edge server contacts the origin server, the ‘direct’ route typically is the route obtained through the Border Gateway Protocol (BGP).”); *see also id.* at 6:39-41 (“Preferably, given CDN edge servers are configured with a plurality (e.g., at least two) alternate

routes to use *in addition to the direct route* to the origin.”), 7:23-32, 7:64-67 (Ex. 1001). Yet the direct or default path is entirely distinct from the “primary” path, described above, that is chosen from among several alternate paths. *See id.*

Accordingly, Patent Owner's proposed construction of “alternate path” should be rejected, and the term should be given its plain and ordinary meaning.

B. “server node”

No construction of “server node” is needed because the term can be easily understood and applied based on its ordinary meaning—that is, “a set of one or more CDN servers.” *See* '538 patent at claim 1, 15:65-66 (“a server node comprises a cluster of content delivery network [CDN] edge servers”), 4:40-55 (Ex. 1001); *see also* Claim Construction Br. at 10 (Ex. 1012). Petitioner's proposed construction—“datacenter containing the content delivery network edge servers for a particular geographic region” (Pet. at 14)—is unnecessary, overly narrow, and should be rejected. *See* Claim Construction Br. at 10 (Ex. 1012).

In particular, Petitioner's proposed construction would artificially limit the servers in the “server node” to those that *serve* a “particular geographic region.” Nothing in the specification, however, requires or suggests that a “server node” must serve a particular “geographic” region. *Id.* The specification instead merely describes that servers in a server node *may be located* geographically near one another in a region (“such as” in a local area network). *See, e.g.*, '538 patent at

6:6-11 (“As noted above, 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).”), claim 1, 15:63-66 (“a set of server nodes organized into regions [and] compris[ing] ... servers”) (Ex. 1001).

Petitioner's reliance on Figures 3 and 4 of the '538 patent (*see* Pet. at 15-16) is misplaced. In fact, Figures 3 and 4 show that the servers in a given node may handle traffic traveling to or from *various* geographic regions, not a “*particular* geographic region,” as Petitioner's construction would require. *See* '538 patent at Figs. 3 & 4 (Ex. 1001). Indeed, sending traffic along an “alternate path” through intermediate server nodes is the primary purpose of the '538 patent. *See supra*.

In sum, neither “alternate path” nor “server node” requires construction, and Petitioner's proposed construction for both terms should be rejected. Under either Petitioner's or Patent Owner's construction for these terms, however, the Petition should be denied.

VI. THE PRIOR ART CITED IN THE PETITION FAILS TO DISCLOSE THE LIMITATIONS OF THE '538 PATENT

Each of the challenged claims requires an “alternate path ... used for edge server-to-content provider origin server communications [that] *includ[es] at least*

one other content delivery network edge server that is external to the given server node [of the particular edge server].”³ ’538 patent at claim 1, claim 2, claim 5, claims 3 and 4 (each requiring “[t]he method as described in claim 2”) (Ex. 1001). But neither Speedera nor Ricciulli—alone or on combination—disclose this feature, and Petitioner fails to explain why the differences between the combined teachings of Speedera and Ricciulli and the claimed feature would be obvious to a POSITA. This flaw is fatal to the Petition.

A. Speedera and Ricciulli Do Not Disclose the Claimed “Alternate Path”

First, Speedera does not teach an “alternate path” that “includ[es] at least one other content delivery network edge server.” As explained above in Section IV(A), Speedera discloses a CDN that includes customer web servers 40 and 45, which contain web servers 130 and 135, and POPs 10, 20, and 30 that contain WebCache servers (e.g., server 280). Speedera at Fig. 1, 4:46-49, 4:56-64, 5:54-59 (Ex. 1004). Petitioner alleges that the customer web server 130 corresponds to an “origin server,” that the POP 30 includes the “given server node,” and that the WebCache server 280 corresponds to the “particular edge server in the given

³ For these limitations, Petitioner only relies on Speedera and Ricciulli and does not rely on HTTP/1.1. *See* Pet. at 36-45.

server node” as set forth in the claims. *See* Pet. at 24-25, 32-33, 36. According to Petitioner, Speedera discloses that WebCache 280 (i.e., the alleged edge server) may communicate with customer web server 130 (i.e., the alleged origin server) to retrieve content requested by an end user. Pet. at 36-38.

Petitioner admits, however, that Speedera “*does not specify*” the path by which the alleged particular edge server 280 communicates with the alleged origin server. *See* Pet. at 38 (Speedera “does not specify which ‘path’ through the Internet such communications [i.e., communications from the alleged edge server to the alleged origin server] would take”); *see also id.* at 26 (“*Speedera is silent* as to how to optimize Internet communication pathways between the POPs (10, 20, 30) [which contain WebCache server 280], customers, (40, 45) and users (150, 160).”). Accordingly, even assuming that Speedera discloses communications between a particular edge server and an origin server, *it does not provide any details about that communication path*, let alone describe a path that includes *another edge server that is external to the given server node*.

Moreover, as noted above in Section IV(A), even though the WebCache server 280 (i.e., the alleged particular edge server) communicates with WebServer server 290, the WebServer server 290 does not correspond to “[an]other content delivery network edge server” included in an “alternative path.” In particular, Petitioner does not assert that WebServer server 290 is a CDN “edge server,” as

required by the claims. *See, e.g.*, Pet. at 17 (asserting that “each POP contains a datacenter having multiple content delivery network edge servers (**WebCaches**).”), 18 (identifying the **WebCaches**, but *not the WebServer*, in a given POP as a “cluster of delivery edge servers”), 37-38, 43-45. Petitioner instead alleges that the WebCaches (e.g., WebCache 280) constitutes the claimed “particular edge server.”⁴ Pet. at 17, 18, 33-34 (contending that “the POPs [in Speedera] include datacenters that contain the CDN edge servers (**WebCaches**)” and that “the POPs in 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.”), 37-38.

In addition, Petitioner correctly concedes that WebCache 280 and

⁴ In places, Petitioner appears to suggest that WebServer server 290 is a “WebCache.” *See, e.g.*, Pet. at 42 (“WebCache (e.g., 80, 90, 280, **290**)”). This seems to be in error, as elsewhere in the Petition, Petitioner expressly distinguishes between the “WebCache” and the “WebServer.” *See id.* at 18, 33 (identifying the WebCaches, but not the WebServer, in a given POP as a “cluster of delivery edge servers”). In any event, Petitioner’s suggestion that WebServer server 290 is a WebCache is incorrect—Speedera makes clear that the “WebCache server 280” is distinct from the “WebServer server 290.” *See, e.g.*, Speedera at 10:40-50 (Ex. 1004).

WebServer server 290 reside in *the same* alleged server node. *See* Speedera at Fig. 1 (Ex. 1004); Pet. at 33 (Fig. 1 annotated by Petitioner to show WebCache 280 and WebServer server 290 in the same alleged “server node”). The WebServer server 290 therefore is not “*external to the given server node*” of the alleged particular edge server (i.e., WebCache 280). As such, the WebServer server 290 cannot correspond to the claimed “other content delivery network edge server [that is in the alternate path and] *that is external to the given server node*” of the alleged particular edge server (i.e., the WebCache 280).

Second, Petitioner's reliance on Ricciulli fares no better. Ricciulli describes a network of nodes consisting of “overlay network nodes” 130*a-n* that allegedly provides a path from a “source node” 100 to a “destination node” 160. *See* Ricciulli at Fig. 1, 3:41-44, 3:53-56 (Ex. 1005). But Ricciulli does not address a “content delivery network having a set of server nodes” or, as a result, “edge server-to-content provider origin server communications.” In fact, Ricciulli discloses that its “overlay network nodes” 130*a-n*: include “*ordinary computer systems* (e.g., Pentium-based PC's) *co-located at Internet sites such as major ISP locations*,” *id.* at 3:65-4:1; use “*existing network transmission lines and infrastructure*,” *id.* at 3:50-53; and “preferably communicate using existing, established Internet protocols,” *id.* at 3:56-59. And the only time that Ricciulli mentions CDNs at all, it criticizes CDNs and teaches away from the claimed

invention, because CDNs are reportedly “*limited or ineffective*” with respect to dynamically generated content” and “*entail significant costs*” for using large-capacity computer servers to physically store extra copies of massive volumes of data, and for maintaining synchronization among the various repositories of frequently updated content.” *Id.* at 1:58-2:6. *See In re Fulton*, 391 F.3d at 1201 (reference “teaches away” where it “criticize[s], discredit[s], or otherwise discourage[s]” investigation into claimed invention); *see also Allergan, Inc. v. Sandoz Inc.*, 796 F.3d 1293, 1306 (Fed. Cir. 2015) (fact that the prior art “criticized, discredited, or otherwise discouraged” use of claimed invention supported finding of nonobviousness (internal alterations omitted)); *In re Gordon*, 733 F.2d 900, 902 (Fed. Cir. 1984). Ricciulli therefore does not disclose an “alternate path” that “includ[es] at least one other content delivery network edge server,” let alone another edge server that is “external to [a] given [CDN] server node.” ’538 patent at claim 1 (16:1-11), claim 2 (16:24-34), claim 5 (16:50-55 and 58-62), claims 3 and 4 (each requiring “[t]he method as described in claim 2”) (Ex. 1001).

Third, even if one were motivated to combine Speedera and Ricciulli (which Patent Owner does not concede), the combination would not disclose the claimed “alternate path.” In fact, if the teachings of Speedera and Ricciulli were combined, the resultant combination at most would simply integrate Ricciulli’s network of

nodes in Speedera's ISP location. As explained above in Section IV(B), Ricciulli discloses that its network of nodes comprises "ordinary computer systems ... co-located at Internet sites such *as major ISP locations*," Ricciulli at 3:65-4:1 (Ex. 1005), and "overlay network routers," *id.* at 5:28-30, 5:44-47. Speedera likewise describes a network that includes *an ISP location* coupled to the Internet via a router. Speedera at Fig. 1 (item 60), 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."), 5:41-42 ("In FIG. 1, the dedicated servers are typically coupled to the computer network 70 via a router."), 5:65-67 ("*ISP 60* is illustrated including users 150-160 *ISP 60 may be a conventional ISP* such as Earthlink, AOL, or the like.") (Ex. 1004). Thus, if the teachings of Speedera and Ricciulli were combined, the resultant combination would integrate Ricciulli's "overlay network nodes," which are "co-located at ... major ISP locations," at Speedera's ISP location (such as ISP 60). As a result, the resultant "alternative path" would include the "ordinary computer systems" in the ISP locations, and would *not* "includ[e] at least one other content delivery network edge server that is external to a given CDN server node."

Petitioner is therefore left to argue, without support, that a POSITA would combine Speedera and Ricciulli *and then modify* that combination to add an edge server to the alleged alternate path. Yet, as explained below, Petitioner fails to cite evidence to support this argument. And Petitioner cannot show invalidity with two

cited references that fail to disclose a limitation (either alone or in a combination) and an unsupported assertion that the missing limitation would have been obvious to a POSITA—the Federal Circuit has repeatedly rejected this kind of analysis. *See Arendi S.A.R.L. v. Apple Inc.*, 832 F.3d 1355, 1366-67 (Fed. Cir. 2016) (reversing obviousness finding due to failure to provide evidence that POSITA would modify prior art to “supply[] the missing limitation,” particularly as “the missing limitation is not a ‘peripheral’ one”); *In re Zurko*, 258 F.3d 1379, 1383 (Fed. Cir. 2001) (reversing Board due to lack of evidence that POSITA would modify cited reference to include missing limitation); *see also Perfect Web Techs., Inc. v. InfoUSA, Inc.*, 587 F.3d 1324, 1330 (Fed. Cir. 2009) (“[T]o invoke ... any ... basis for extrapolating from prior art to a conclusion of obviousness, a district court must articulate its reasoning with sufficient clarity for review.”).

B. Petitioner's Arguments Regarding the Claimed “Alternate Path” Fail

Petitioner's arguments do not support a different result. Petitioner claims that it would be obvious to include Ricciulli's “overlay network module” software at each Speedera POP, Speedera's Operation Center, and the customer POP servers because, according to Petitioner, Ricciulli “teaches to employ such overlay network modules ... at high-traffic areas, such as ISP locations.” Pet. at 44-45. Yet, as explained above, combining Speedera and Ricciulli would not disclose an

“alternative path” that “includ[es] at least one other content delivery network edge server that external to [a] given server node,” as claimed. Moreover, Petitioner does not explain why a POSITA looking to improve Speedera’s CDN would look to Ricciulli, which does not disclose CDN edge servers, content provider origin servers, or communications or communication paths between edge servers and origin servers. And Petitioner’s argument is entirely conclusory, as neither Petitioner nor its expert explain how the sole cited portion of Ricciulli—which simply notes that the “network includes a set of *ordinary computer systems ... co-located at* Internet sites such as *major ISP locations*”—supports Petitioner’s argument. *See* Pet. at 44-45 (citing Ricciulli at 3:65-4:1 (Ex. 1005)). It is well-settled that a party challenging validity cannot meet its burden by simply asserting that the prior art meets the limitations of the challenged claims without citation or support. *See In re Magnum Oil Tools Int’l*, 829 F.3d 1364, 1380 (Fed. Cir. 2016) (“To satisfy its burden of proving obviousness, a petitioner cannot employ mere conclusory statements. The petitioner must instead articulate specific reasoning, based on evidence of record, to support the legal conclusion of obviousness.”); *Personal Web Techs., LLC v. Apple, Inc.*, 848 F.3d 987, 991 (Fed. Cir. 2017) (to prove obviousness, petitioner must show that a POSITA “would have been motivated to combine the prior art in the way claimed ... and had a reasonable expectation of success in doing do” and provide a “clear, evidence-supported

account of the contemplated workings of the combination”).

Petitioner's reliance on the “large number of alternative nodes” disclosed in Ricciulli (*see* Pet. at 45) is equally misplaced. Petitioner claims that, because Ricciulli discloses a plurality of nodes, it “would be obvious to use the overlay software at the many POP's of Speedera” such that the “the combined Speedera-Ricciulli[] system would include at least one Speedera edge server that is “external to the given source node.”” Pet. at 45. But, as explained above, Speedera only discloses communication between a WebCache 280 and WebServer 290 located *within the same server node*, or communication between a WebServer 290 and the alleged content provider origin server. *See* Speedera at Fig. 1, 10:43-65 (Ex. 1004). And again, Petitioner's argument is entirely conclusory. Neither Petitioner nor its expert cite any evidence to support their claim that the number of nodes in the system demonstrates that an alternative path would include an edge server external to the given server node. As such, Petitioner's argument should be rejected. *See In re Magnum Oil Tools Int'l*, 829 F.3d at 1380.

VII. CONCLUSION

Each of the challenged claims of the '538 patent requires an “alternate path” that “includ[es] at least one other content delivery network edge server that is external to the given server node.” None of the references Petitioner relies on, either alone or in combination, disclose these elements of the challenged claims.

Accordingly, the Board should decline to institute *inter partes* review of challenged claims 1-5 of the '538 patent.

Date: March 28, 2018

Respectfully submitted,
By: /Grant K. Rowan/
Grant K. Rowan, Reg. No. 41,278
Wilmer Cutler Pickering Hale and Dorr LLP
1875 Pennsylvania Ave., NW
Washington, DC 20006

PATENT OWNER'S LIST OF EXHIBITS
IPR2018-00310
U.S. PATENT 8,194,538

Exhibit No.	Description
2001	U.S. Patent No. 6,108,703 to Leighton et al. ("the '703 patent")

CERTIFICATE OF SERVICE

I hereby certify that, on March 28, 2018, I caused a true and correct copy of:

- PATENT OWNER'S PRELIMINARY RESPONSE PURSUANT TO 37 C.F.R. §42.107
- EXHIBIT 2001
- LIST OF EXHIBITS FOR PATENT OWNER'S PRELIMINARY RESPONSE PURSUANT TO 37 C.F.R. §42.107 (EXHIBIT 2001)
- CERTIFICATE OF WORD COUNT UNDER 37 C.F.R. § 42.24(d)

to be served via electronic mail on the following attorneys of record:

Barry J. Schindler
Greenberg Traurig, LLP
500 Campus Drive, Suite 400
Florham Park, NJ 07932
SchindlerB@gtlaw.com

Heath J. Briggs
Greenberg Traurig, LLP
1200 17th St., Suite 2400
Denver, CO 80202
BriggsH@gtlaw.com
limelightips@gtlaw.com
limelight-ips@gtlaw.com

By: /Scott Bertulli/
Scott Bertulli, Reg. No. 75,886
Wilmer Cutler Pickering Hale and Dorr LLP
60 State Street
Boston, MA 02109
(617) 526-6706

Date: March 28, 2018

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

I hereby certify that the foregoing Patent Owner's Preliminary Response contains 8,099 words as measured by the word-processing software used to prepare the document, in compliance with 37 C.F.R. § 42.24(d).

By: /Scott Bertulli/
Scott Bertulli, Reg. No. 75,886
Wilmer Cutler Pickering Hale and Dorr LLP
60 State Street
Boston, MA 02109
(617) 526-6706

Date: March 28, 2018