

**IN THE UNITED STATES DISTRICT COURT
DISTRICT OF MASSACHUSETTS**

AKAMAI TECHNOLOGIES, INC.

Plaintiff,

v.

LIMELIGHT NETWORKS, INC.

Defendant.

Case No. 1:16-cv-10253-GAO

Case No. 1:16-cv-12624-GAO

LEAVE TO FILE ADDITIONAL
PAGES GRANTED ON
OCTOBER 10, 2017

**AKAMAI TECHNOLOGIES, INC.'S
OPENING CLAIM CONSTRUCTION BRIEF**

TABLE OF CONTENTS

	<u>Page</u>
I. INTRODUCTION	1
II. LEGAL STANDARD.....	1
III. TECHNOLOGY OVERVIEW	2
IV. CONSTRUCTIONS OF DISPUTED TERMS.....	4
A. Akamai’s Patents	4
1. ’319 Patent – “host configuration table (HCT)”	4
2. ’538 Patent	6
a) “alternate path” (claims 1-5)	7
b) “server node” (claims 1, 2, 5)	10
3. ’211 Patent	11
a) “global cache” (claims 1, 5-8)	12
b) “third party edge network” (claims 1, 8)	13
c) “filling the global cache” (claim 1).....	15
d) “selectively filling at least a portion of the global cache” (claim 8)	16
e) “close to end users” (claims 1, 8).....	17
f) “high-bandwidth content” (claims 1, 8).....	20
g) “popular” / “popularity” (claims 1, 3, 8).....	22
B. Limelight’s ’661 Patent.....	24
1. “[the identified] one or more [global] policies” (all claims).....	25
2. “circumstances” (all claims)	28
3. “one or more authorized transfers of [at least a portion] of the content” (all claims)	29
V. CONCLUSION.....	30

TABLE OF AUTHORITIES

	Page(s)
Cases	
<i>Adams Respiratory Thera., Inc. v. Perrigo Co.</i> , 616 F.3d 1283 (Fed. Cir. 2010).....	1
<i>Akamai Techs., Inc. v. Limelight Networks, Inc.</i> , No. CIV. A. 06-11109-RWZ, 2008 WL 697707 (D. Mass. Feb. 8, 2008)	19
<i>Akzo Nobel Coatings, Inc. v. Dow Chem. Co.</i> , 811 F.3d 1334 (Fed. Cir. 2016).....	27
<i>Epos Techs. Ltd. v. Pegasus Techs. Ltd.</i> , 766 F.3d 1338 (Fed. Cir. 2014).....	9
<i>Leggett & Platt, Inc. v. Hickory Springs Mfg. Co.</i> , 285 F.3d 1353 (Fed. Cir. 2002).....	28
<i>Lexington Luminance LLC v. Amazon.com, Inc.</i> , No. 12-cv-12216-DJC, 2016 WL 1337254 (D. Mass. Apr. 4, 2016)	15
<i>Markman v. Westview Instruments, Inc.</i> , 517 U.S. 370 (1996).....	1
<i>Merck & Co. v. Teva Pharm. USA, Inc.</i> , 395 F.3d 1364 (Fed. Cir. 2005).....	13, 27
<i>Microunity Sys. Eng’g, Inc. v. Acer, Inc.</i> , No. 2:10-CV-91-LED-RSP, 2013 WL 12134146 (E.D. Tex. Mar. 8, 2013).....	21
<i>Music Choice v. Stingray Digital Grp. Inc.</i> , No. 2:16-CV-586-JRG-RSP, 2017 WL 2896025 (E.D. Tex. July 6, 2017)	23
<i>Nautilus, Inc. v. Biosig Instruments, Inc.</i> , 134 S. Ct. 2120 (2014).....	2, 18, 22, 24
<i>Oatey Co. v. IPS Corp.</i> , 514 F.3d 1271 (Fed. Cir. 2008).....	17
<i>Phillips v. AWH Corp.</i> , 415 F.3d 1303 (Fed. Cir. 2005) (en banc).....	1, 6
<i>Phoenix Licensing, L.L.C. v. AAA Life Ins. Co.</i> , No. 2:13-cv-1081, 2015 WL 3866832 (E.D. Tex. June 22, 2015)	28

<i>Premier Int’l Assocs. v. Apple Computer, Inc.</i> , 512 F. Supp. 2d 737 (E.D. Tex. 2007)	23
<i>Shire Dev. LLC v. Amneal Pharm. LLC</i> , No. 15-2865 (RBK/JS), 2016 WL 4119940 (D.N.J. Aug. 2, 2016)	28
<i>SuperGuide Corp. v. DirecTV Enters., Inc.</i> , 358 F.3d 870 (Fed. Cir. 2004).....	6
<i>Trs. of Boston Univ. v. Everlight Elecs. Co.</i> , 23 F. Supp. 3d 50, 58 (D. Mass. 2014)	1
<i>Trustees of Columbia Univ. v. Symantec Corp.</i> , 811 F.3d 1359 (Fed. Cir. 2016).....	1
<i>Veracode, Inc. v. Appthority, Inc.</i> , No. 12-10487-DPW, 2013 WL 5587946 (D. Mass. Oct. 9, 2013)	6
Statutes	
35 U.S.C. § 112.....	1

I. INTRODUCTION

The parties disagree about the construction of claims of five of the asserted patents, and their respective constructions reflect vastly different approaches to this process. Limelight's positions are inconsistent with well-established claim construction principles, contrary to the plain meaning of the claim language, and contrary to the teachings of the patent specifications. Akamai's positions, in contrast, are well grounded in the law and the intrinsic evidence. Akamai accordingly requests that the Court adopt its proposed constructions.

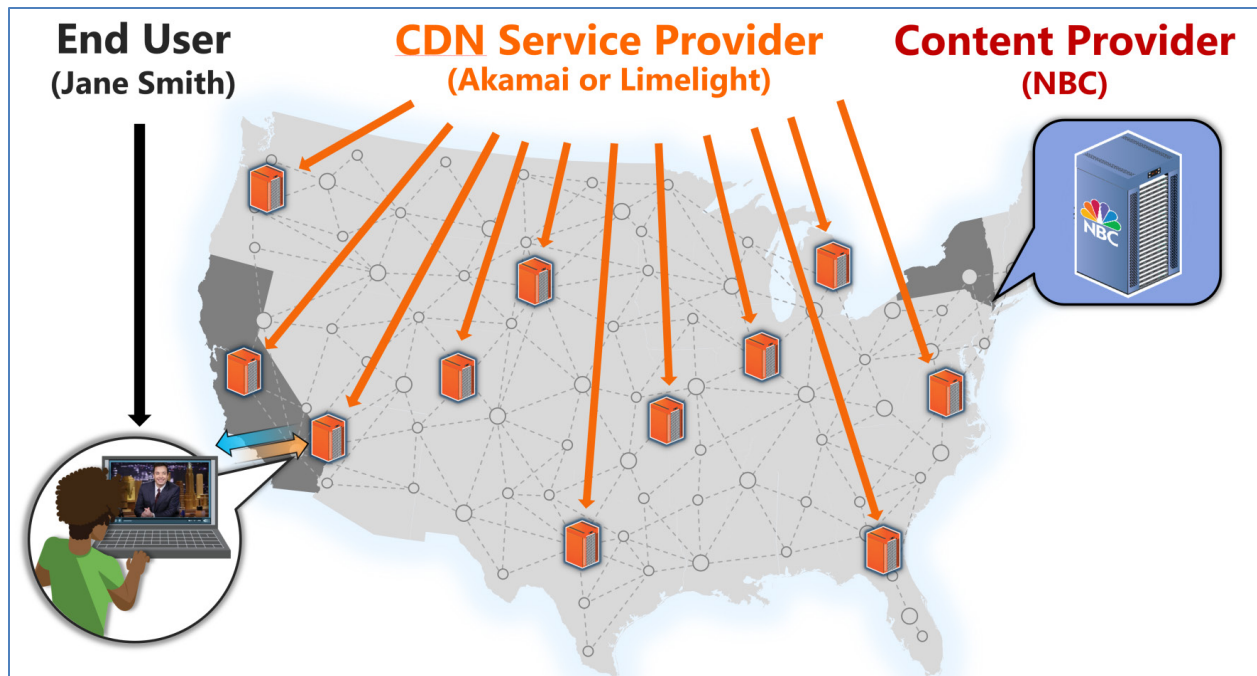
II. LEGAL STANDARD

Claim construction is an issue of law for the Court. *Markman v. Westview Instruments, Inc.*, 517 U.S. 370, 372 (1996). The Court should view the claim language through the lens of a person of ordinary skill in the art at the time of the invention. *See Phillips v. AWH Corp.*, 415 F.3d 1303, 1313 (Fed. Cir. 2005) (en banc). "The specification is the single best guide to the meaning of a disputed term, and is, thus, the primary basis for construing the claims." *Trustees of Columbia Univ. v. Symantec Corp.*, 811 F.3d 1359, 1362 (Fed. Cir. 2016). If "nothing in the specification imparts any special meaning" to a term, it "should be given its plain and ordinary meaning." *Adams Respiratory Thera., Inc. v. Perrigo Co.*, 616 F.3d 1283, 1291 (Fed. Cir. 2010). The prosecution history also can "inform the meaning of the claim language by demonstrating how the inventor understood the invention and whether the inventor limited the invention in the course of prosecution, making the claim scope narrower than it otherwise would be." *Phillips*, 415 F.3d at 1317. Extrinsic evidence, such as dictionaries and treatises, can be useful but should be read in the context of the intrinsic evidence. *Trs. of Boston Univ. v. Everlight Elecs. Co.*, 23 F. Supp. 3d 50, 58 (D. Mass. 2014) (citing *Phillips*, 415 F.3d at 1318-1319).

Under 35 U.S.C. § 112, a patent claim must particularly point out and distinctly claim the subject matter which the inventor regards as his or her invention. A party challenging a patent claim as indefinite must establish that “its claims, read in light of the specification delineating the patent, and the prosecution history, fail to inform, with reasonable certainty, those skilled in the art about the scope of the invention.” *Nautilus, Inc. v. Biosig Instruments, Inc.*, 134 S. Ct. 2120, 2124 (2014). The party challenging the claim bears the burden of proof on this issue, and must establish indefiniteness with clear and convincing evidence. *Id.* at 2130 n.10.

III. TECHNOLOGY OVERVIEW

Each of the asserted patents relates to the design and features of content delivery networks (“CDNs”). Before CDNs, if a person (*i.e.*, “end user”) wanted to access Internet content from a website such as www.NBC.com, the end user’s computer would communicate directly with the website provider’s own web servers (*i.e.*, the www.NBC.com servers). This made Internet content delivery slow and unreliable because most website providers (typically called “content providers”) maintained only a limited number of servers in limited geographical locations, far away from most of the end users requesting the content. To solve these problems, Akamai developed a new third-party network of servers—CDN content servers—to host Internet content from multiple content providers. These CDN content servers are often referred to as “edge servers” because they are located at the “edge” of the Internet (*i.e.*, close to where end users access the internet). Companies providing CDN services (such as Akamai and Limelight) are often referred to as “CDN service providers,” and the content providers (such as NBC) are often referred to as the CDN service provider’s “customers.” The diagram below shows a simplified CDN network:



A CDN service provider (such as Akamai or Limelight) stores content from multiple content providers (such as NBC in New York) in the cache memory of its edge servers. When an end user (Jane Smith in California) requests a video from NBC’s website, her computer is redirected to a nearby CDN edge server to download the video. If the video is already stored on the CDN edge server, it will be retrieved from there and delivered to Ms. Smith. If not, then the CDN will retrieve the video from the content provider’s servers (*e.g.*, the NBC.com servers in New York), deliver it to Ms. Smith, and store (*i.e.*, “cache”) a copy in the CDN edge server to use for future requests. Then, if another end user in the same area later requests that same NBC video, he or she would likely be directed to the same CDN edge server, which—following Ms. Smith’s request—now has a cached copy of that video. By distributing content closer to end users in this manner, CDNs make the Internet faster, more efficient, and more reliable.

IV. CONSTRUCTIONS OF DISPUTED TERMS

A. Akamai's Patents

1. '319 Patent – “host configuration table (HCT)”

U.S. Patent No. 8,484,319 (“the ’319 patent”) (No. 16-12624, Dkt. No. 1-1) introduces techniques for configuring the delivery of internet content according to the different preferences of particular CDN customers (*e.g.*, NBC.com). *See, e.g.*, ’319 patent, 2:46-3:35 (“Brief Summary of the Invention”). The ’319 claims recite a number of features that are part of this system, including using a “host configuration table” (HCT) to configure a host, such as a CDN server, to deliver content based on customer-specific configurations. *See* ’319 patent, claim 1.

Disputed Claim Term: “Host Configuration Table”	
Akamai’s Construction	Limelight’s Construction
“a table that includes information used to configure a host for the delivery of content”	“a table containing content delivery network service provider-wide metadata with entries following a dictionary form, including a unique version attribute corresponding to the customer for which the table applies”

The parties’ primary dispute over “host configuration table” is whether it should be construed consistent with the plain meaning of the ’319 claims (as Akamai proposes), or whether the Court should add limitations from the specification that have no basis in the claims, such as “content delivery network service provider-wide metadata,” “entries following a dictionary form,” and a “unique version attribute corresponding to [a] customer” (as Limelight proposes).

The claims, specification, and prosecution history all support Akamai’s construction. The claim language indicates that the “host configuration table” (“HCT”) is a “table” that includes information used to configure a “host” for the delivery of content. ’319 patent, claim 1 at 22:22-25; 22:47-50. The specification uses “HCT” consistent with the plain language of the claims by providing an embodiment, for example, in which a “host configuration table” is used to configure a “host” (*e.g.*, a CDN server) for the delivery of content. *See, e.g., id.* 8:47-55

(defining “CDN server” as a “global *host*”); 9:28-30; 16:55-17:10; 20:21-23 (describing an example of an HCT that can be used to configure a CDN server (*i.e.*, host) for the delivery of content); *see also* Declaration of Dr. Samrat Bhattacharjee, Ph.D. in Support of Akamai Technologies, Inc.’s Opening Claim Construction Brief (“Bhattacharjee Decl.”) ¶¶ 24-25.

Further, the patentees explained during prosecution that the “host configuration table” is “*defined in claim 1*”¹ (as described above), and thus is not limited to the preferred embodiments. Ex. 1² (U.S.A.N. 13/400,146, Oct. 23, 2012 Amendment) at 9 (“[T]he claim elements directed to the ‘host configuration table’ *as that table is defined in claim 1.*”); *see also* Bhattacharjee Decl. ¶ 26. Consistent with this statement, to the extent the patentees intended to impose any additional requirements for the “host configuration table,” those requirements are specified in other limitations of the claims. For example, claim 1 requires that the HCT “identifies content provider domains or subdomains” being managed by the CDN, that it include an “entry” for “a content provider domain or subdomain,” and that this entry “point[s] to a configuration file” having “at least one content handling instruction.” ’319 patent, 22:23-31.

Nothing in the claim language, the specification, or the prosecution history imposes the requirements that Limelight is attempting to add to the construction of the HCT. Specifically, Limelight argues that the “host configuration table” must also contain (1) “content delivery network service provider-wide metadata,” (2) entries “following a dictionary form,” and (3) “a unique version attribute corresponding to the customer for which the table applies.” These features are *not mentioned anywhere in the ’319 claims*. *See* ’319 patent, claims 1-11. Instead,

¹ All emphasis in quotations is added unless otherwise noted.

² All references to “Ex. ___” are to the exhibits attached to the Declaration of Arthur Coviello, filed concurrently with this brief.

they are specifically described in the patent as a “*preferred embodiment*” in column 16 of the ’319 specification. *Id.* 3:53-54 (confirming features cited by Limelight in column 16 are part of the “**PREFERRED EMBODIMENT**”); 16:14-33 (describing the features of the preferred embodiment that Limelight is attempting to read into the claims); 9:26-32 (explaining that “[t]he following [including the description of the HCT at column 16] is a detailed specification of a *preferred CDN content server interface*” and that “[t]he following describes these files and identifies *preferred requirements* for their formats.”). *See* Bhattacharjee Decl. ¶ 27

Limelight’s attempt to read these features from the specification into the claims thus violates the “bedrock principle” of patent law that “the *claims of a patent* define the invention.” *Phillips*, 415 F.3d at 1312. Indeed, “reading a limitation from the written description into the claims” is “*one of the cardinal sins of patent law.*” *Id.* at 1319-20; *SuperGuide Corp. v. DirecTV Enters., Inc.*, 358 F.3d 870, 875 (Fed. Cir. 2004) (“[I]t is important not to import into a claim limitations that are not a part of the claim. For example, a particular *embodiment* appearing in the written description *may not be read into a claim when the claim language is broader than the embodiment.*”); *Veracode, Inc. v. Appthority, Inc.*, No. 12-10487-DPW, 2013 WL 5587946, at *9 (D. Mass. Oct. 9, 2013) (“[T]he Federal Circuit has repeatedly asserted that reading limitations from particular embodiments into the claims is the ‘cardinal sin’ of claim construction.”).

2. ’538 Patent

U.S. Patent No. 8,194,538 (“the ’538 patent”) (No. 16-12624, Dkt. No. 1-2) introduces an “inventive routing technique [that] enables an edge server operating within a given CDN region to retrieve content ... from an origin server more efficiently by selectively routing through the CDN’s own nodes.” ’538 patent, Abstract. When an end user makes a request for content and the CDN must retrieve the content from the customer’s origin server, the path across the Internet

to the origin server may be congested. *See, e.g., id.* 1:60-2:1. Therefore, the CDN also identifies multiple “alternate paths” to the origin server. *See, e.g., id.* 6:23-25. These “alternate paths” may be less direct paths, traveling from the CDN “server node” through **other**, intermediate CDN server nodes before arriving at the content provider’s server (*e.g.*, an NBC.com server). *See, e.g., id.* 6:11-15. Nevertheless, they may avoid congestion and improve performance. *Id.* Namely, as the inventors discovered, identifying these alternate paths can improve performance, because sometimes paths that are less direct can perform better than a more direct path. *See id.* 6:34-38 (“While one might assume that adding the intermediate step [of routing a path through another CDN server node] would reduce performance, it frequently improves performance, because CDN servers and regions typically are well-connected, and the indirect route can often bypass network congestion.”). The asserted claims describe methods for identifying and using at least one of these “alternate paths.” *See, e.g., id.* claim 1.

a) “alternate path” (claims 1-5)

Akamai’s Construction	Limelight’s Construction
Plain and ordinary meaning; no construction needed Alternatively, if the court deems construction to be necessary, “alternate path” should be construed as “one of multiple paths between a given edge server and an origin server”	“path other than a primary path already specified by the content delivery network”

Akamai contends that no construction of this term is needed because the plain meaning of “alternate path” in the context of the ’538 patent claims and specification is “one of multiple paths between a given edge server and an origin server.” On the other hand, Limelight’s construction requiring the alternate path to be different (*i.e.*, “other”) than an “already specified” primary path is contrary to the embodiments set forth in the ’538 patent.

Indeed, the patent confirms that the CDN service provider **first** determines a set of multiple paths (the “alternate paths”) and **then** selects an optimal (*i.e.*, primary) path from

among those multiple alternate paths. *See, e.g.,* '538 patent, 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...*”); Bhattacharjee Decl. ¶ 30. More specifically, the specification of the '538 patent clearly and repeatedly spells out a sequence for selecting a path, where the CDN service provider first determines a set of alternate paths, then conducts performance tests such as a download “race” to determine which path is best according to current conditions, and finally selects that best path as “primary” or “optimal.” *See id.* 8:3-8 (“These races consist of having the edge server forward a client request through the a [sic] *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.*”); 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); 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...*”). In other words, the primary path is *one of* the alternate paths.

Limelight’s proposed construction—requiring the alternate path be “other than” a primary path, and that the primary path be “already specified”—is flatly inconsistent with the specification. *First*, Limelight has the timing backwards. The “primary path” is not “*already* specified” when the multiple alternate paths are determined. Bhattacharjee Decl. ¶ 34. Instead, the “primary” path is selected from among a set of alternate paths after they have been determined. *See id.* 8:3-8 (“These 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.”); 7:23-32 (describing three phases of path selection, where first a map of “potential alternate routes is created” and is then refined into a choice of a “best” path); Bhattacharjee Decl. ¶¶ 31-32 (describing that “best,” “optimal,” and “primary” path are used as equivalents in the specification).

Second, and for similar reasons, the “alternate path” is not “*other than*” a primary path. Instead, the patent is clear that any of the alternate paths might become the “primary path” depending on various conditions. *See* ’538 patent, 8:1-11 (“[Performance] races consist of having the edge server forward a client request through ... a given number (e.g., three (3)) 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.* 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....”); Bhattacharjee Decl. ¶ 35. By way of analogy, the alternate paths are like different coats from which one can choose on a given day—a rain coat, a warm winter coat, or a light fleece. These are “alternate” coats, any of which may be chosen as the “primary” coat that day based on one or more conditions. On a rainy day, the rain coat may be selected as the “primary” coat, or on a cold day, the winter coat may be chosen as the “primary” coat. Accordingly, Limelight’s construction, which defines an “alternate” as “other than” a “primary,” is inconsistent with the preferred embodiments and should be rejected. *See, e.g., Epos Techs. Ltd. v. Pegasus Techs. Ltd.*, 766 F.3d 1338, 1347 (Fed. Cir. 2014) (“[A] claim construction that excludes a preferred embodiment ... is rarely, if ever

correct”).

b) “server node” (claims 1, 2, 5)

Akamai’s Construction	Limelight’s Construction
Plain and ordinary meaning; no construction needed Alternatively, if the court deems construction to be necessary, “server node” should be construed as “a set of one or more CDN servers”	“data center containing the content delivery network edge servers for a particular geographic region”

Akamai contends that no construction of “server node” is needed because the term can easily be understood and applied based on its ordinary meaning. As described in the preamble of claim 1, a “server node” is simply a set of one or more CDN servers. *See, e.g.*, ’538 patent, claim 1, 15:65-66 (“a server node comprises a cluster of content delivery network [CDN] edge servers”); *see also id.* 4:40-55.

Limelight’s proposed construction is unnecessary, overly narrow and, once again, attempts to read in limitations from the specification. In particular, Limelight’s construction attempts to artificially limit the servers in the “server node” to those that serve a “particular geographic region.” But nothing in the specification requires or suggests that a “server node” must *serve* a particular “geographic” region. Instead, the specification merely describes that servers in a server node typically *may be located* geographically near one another in a region (“such as” in a local area network). *See, e.g.*, ’538 patent, 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”).

Bhattacharjee Decl. ¶ 38.

3. '211 Patent

U.S. Patent No. 8,250,211 (the “’211 patent”) (No. 16-10253, Dkt. 1-2) addresses the problem of providing reliable, fast access to large multimedia files, such as videos, which require significant bandwidth to traverse networks. ’211 patent, 1:17-45. The ’211 patent provides techniques that a CDN service provider can employ in delivering this high bandwidth content. The claimed invention involves a “global cache” that is designed to provide “intelligent distribution” of multimedia files across the Internet and to “dynamically adapt to varying demands for multimedia files.” *Id.* 1:40-45; 12:29-34; 13:6-10.

In particular, the CDN service provider has a “global cache” that includes multiple clusters of “edge servers” that hold popular videos. *Id.* figs. 1, 3; 3:53-63; 4:51-58; 6:12-17; 9:52-54. Within each cluster, the edge servers are connected as part of a local area network (“LAN”). *E.g., id.* fig. 3; 10:53-59; 12:31-43. The CDN service provider also maintains, in association with the cluster, another server for storing popular content. When an end user requests a specific video, he or she is directed and connected to an appropriate edge server (*e.g.*, one that is located close to the end user) that is likely to have a copy of that video file in its local storage. If the video is not stored on that edge server, the user will be redirected to another edge server or to the service provider’s other content server, associated with the cluster, to retrieve the video. *Id.* figs. 2, 5-8, 4:18-23, 5:65-6:25, 10:42-11:23. The edge servers keep track of the number of requests for each video (thus providing data about the popularity of videos) and can also download popular videos from nearby service provider servers (thus ensuring that current popular videos are available in local storage for efficient delivery to end users). *Id.* Abstract; fig. 9; 4:65-5:5; 6:12-17; 9:16-10:39; 11:24-44.

a) “global cache” (claims 1, 5-8)

Akamai’s Construction	Limelight’s Construction
<p>Plain and ordinary meaning; no construction needed</p> <p>Alternatively, if the court deems construction to be necessary “global cache” should be construed as “geographically distributed collection of edge servers.”</p>	<p>“cache on a cluster of LAN-connected edge servers that primarily serve worldwide requests, not primarily requests from the region in which the edge server is located” (claims 1, 5-7)</p> <p>“cache on multiple clusters of LAN-connected edge servers that primarily serve worldwide requests, not primarily requests from the region in which the edge server is located” (claim 8)</p>

Akamai contends that no construction is needed because “global cache” can easily be understood and applied based on its ordinary meaning. As the parties appear to agree, the term “cache” refers to the cache memory of the CDN service provider’s edge servers. A “global cache” accordingly is a geographically distributed collection of edge servers.

The claims use the term “global cache” consistent with its ordinary meaning, specifying that that the “global cache” includes clusters of edge servers located “close to end users” (*i.e.*, they are geographically distributed in multiple locations close to end users). *Id.* 12:31-34; 13:7-9. The specification similarly confirms that the “global edge servers” are geographically distributed—they are located at the edges of the “global network” or “worldwide Internet”—to enable efficient local delivery of content. *Id.* 4:3-4; 6:66-67; *see also id.* 7:1-6. This geographical distribution enables the CDN service provider to select an “appropriate” edge server for delivery of content based on, for example, “geographic location in relation to the user.” *Id.* 5:66-6:2. To the extent that “global cache” requires construction, it accordingly should be a “geographically distributed network of edge servers.” *See* Bhattacharjee Decl. ¶¶ 41-43.

Limelight proposes two different constructions for the same term, both of which are inconsistent with the specification and the role of the global cache. **First**, Limelight proposes

that the edge servers in the global cache must “primarily serve worldwide requests.” But nothing in the specification requires or even implies that the edge servers in the global cache must serve “worldwide requests.” Instead, the specification explains that the edge servers are “global” because of their locations—they are widely distributed across the global Internet. *Id.* 4:3-4; 6:66-67; 7:1-6. In fact, Limelight’s reading is **contrary** to one of the primary purposes of the ’211 patent—locating the edge servers “close to end users” so that requests can be routed to a nearby edge server (*i.e.*, serving local requests). *Id.* 5:66-6:2; 12:31-34; 13:8-9. This means that each edge server will primarily serve **local** requests, not “worldwide” requests.

Second, Limelight’s reading incorrectly conflates the terms “cluster” and “global cache.” *See Merck & Co. v. Teva Pharm. USA, Inc.*, 395 F.3d 1364, 1372 (Fed. Cir. 2005) (“A claim construction that gives meaning to all the terms of the claim is preferred over one that does not do so.” (citations omitted)). Although claim 1 describes “organizing a cluster of LAN-connected edge servers into a global cache,” it does not follow that the “global cache” can include only servers in a single “cluster.” Instead, the cluster is “organized ... **into**” a broader global cache, which includes **multiple** clusters. As noted above, the global edge servers constituting the “global cache” are distributed geographically (’211 patent, 4:3-4; 6:66-67; 7:1-6), but the edge servers within each “cluster,” as described in the claims, must be connected as part of a **local** area network—necessarily demonstrating that a global cache cannot be a single cluster.

b) “third party edge network” (claims 1, 8)

Akamai’s Construction	Limelight’s Construction
Plain and ordinary meaning; no construction needed Alternatively, if the court deems construction to be necessary “third party edge network” should be construed as “an edge network, operated by an entity other than the content provider or end user.”	“a network of an entity other than the service provider where the entity owns and manages the network infrastructure and the network provides, administers and controls end users’ Internet access”

Akamai contends that no construction is needed because the term can be understood and

applied based on its plan and ordinary meaning. A “third party edge network” is an edge network, operated by an entity *other than* the content provider or end user—which represent the “first” and “second” parties.

As the claims and specification confirm, the “first party” in the claims is the end user, which requests and receives content. The “second party” in the claims is the content provider. For example, an end user, Jane Smith, might request a news video from a content provider such as NBC.com. *See* ’211 patent, fig. 1; 1:53-61; 4:51-58 (illustrating embodiment where “customer site [content provider]” uploads content and “user” requests content); fig 2; 5:65-6:10 (illustrating embodiment where “user” requests a “customer’s [content provider’s] file”); fig. 4 (showing “customer [content provider]” as party uploading content); fig. 5; 10:45-52 (showing “end user” as party requesting content provided by the content provider); *see also* Bhattacharjee Decl. ¶ 46.

The “third party edge network” in the claims accordingly is a network operated by a party *other than* the end user or the content provider and can include the CDN service provider’s network. Indeed, the only “third party edge network” described in any detail in the specification is the service provider’s network. *See, e.g., id.* 3:65-4:30, 4:35-50, 6:18-7:28; *see also* Bhattacharjee Decl. ¶¶ 45, 47.

Limelight’s proposed construction is inconsistent with the plain meaning of the term and attempts to impose additional requirements that have no support in the patent. For example, there is nothing in the claim language—and nothing in the specification—describing a third party “other than the service provider” which “owns and manages the network infrastructure” and “provides, administers and controls end users’ Internet access.” Instead, the only “third party network” described in any detail in the specification is the service provider’s network. The

“Internet” is simply depicted as a cloud. *Id.* fig. 3. Further, there is nothing in the claims requiring that the third party network “control[]” the end users’ access to the Internet. Instead, the only requirement is that end users “obtain access to the Internet via [the] third party edge network.” *Id.* 13:9-10. Limelight’s proposed inclusion of these additional limitations finds no support in the patent and should be rejected.

c) “filling the global cache” (claim 1)

Akamai’s Construction	Limelight’s Construction
“adding content to edge servers in the global cache”	“filling all available storage of the global cache selected on the basis of [popularity and high-bandwidth]”

The parties’ primary dispute over the term “filling the global cache” is whether it should be construed consistent with its ordinary meaning (as Akamai proposes) or whether the Court should read in an additional requirement not supported by the claims or the specification—namely, filling “***all*** available storage” (as Limelight proposes).

The claims and specification confirm that the term should have its ordinary meaning. The edge servers in the global cache store content for delivery to end users. *See* ’211 patent, 3:53-56 (“An edge server maintains a dynamic number of popular files in its memory for the customer.”); 9:15-23. The plain meaning of “filling” the global cache accordingly is “adding content” to the edge servers that comprise the cache. *See* Bhattacharjee Decl. ¶¶ 48-49.

Limelight’s proposal not only deviates from the plain meaning of the claims and the disclosure in the patent, but to rewrites the clear claim language to add a requirement of filling “**all available storage** of the global cache.” *See Lexington Luminance LLC v. Amazon.com, Inc.*, No. 12-cv-12216-DJC, 2016 WL 1337254, at *7 (D. Mass. Apr. 4, 2016) (rejecting party’s proposed claim construction for “fail[ing] to recognize that its proposal would read in the modifier ‘all’ or ‘every’ without providing any support for that insertion” (citing *Johnson*

Worldwide Assocs., Inc. v. Zebco Corp., 175 F.3d 985, 989 (Fed. Cir. 1999))). Claim 1 simply refers to “filling the global cache”—not “filling ***all available storage*** of the global cache” —with popular content. Indeed, the specification indicates the opposite—memory reserved for popular files is ***less*** than the total memory available on an edge server. *See* ’211 patent, 3:53-59 (“An edge server maintains a dynamic number of popular files in its memory for the customer. The number of files is determined by the fixed amount of memory that the edge server reserves for the customer.”); 9:15-20 (“Each edge server is configured to have storage memory set aside for popular files”). Limelight’s modification accordingly has no basis in either the claims or the specification. Bhattacharjee Decl. ¶¶ 48-49.

d) “selectively filling at least a portion of the global cache”
(claim 8)

Akamai’s Construction	Limelight’s Construction
“adding content to edge servers in at least one selected portion of the global cache”	“filling all available storage of a portion of the global cache assigned to a specific customer with [popular, high-bandwidth] content selected by the customer”

The parties’ dispute over the term “selectively filling at least a portion of the global cache” is similar. Akamai proposes construing the term consistent with its ordinary meaning. Limelight again improperly proposes adding requirements that are inconsistent with the claim language and the specification, including filling “***all available storage***,” filling a portion of the global cache “***assigned to a specific customer***,” and filling the cache with “content ***selected by the customer***.”

Once again, the claim language and specification confirm that the term should be construed consistent with its ordinary meaning. “Selectively filling at least a portion of the global cache” is simply “adding content to edge servers ***in at least one selected portion*** of the global cache.” “Selectively” plainly refers to the “portion” of the global cache that is filled—that

is, the particular edge servers in the global cache that are filled—and not to the content that is added. This reading is consistent with the specification, which notes that “the popularity of the content is determined at the edge of the network[, which] means a particular content can vary from one edge location to another.” ’211 patent, 9:52-54; *see also* Bhattacharjee Decl. ¶¶ 50-51.

Limelight’s construction, in contrast, is inconsistent both with the claim language and with the specification. **First**, nothing in claim 8 requires that a portion of the global cache be “assigned to a specific customer.” Rather, the claim describes filling “a portion of the global cache” without reference to any “assignment” to a “specific customer.” **Second**, nothing in claim 8 requires that the content be “selected by the customer.” Instead, contrary to Limelight’s construction, the specification describes an embodiment where popular content is determined **by the edge servers** (not in advance by the customer) through an automated process based on the number of requests a content file receives. *Id.* fig. 4, 4:65-5:5; 6:12-17; 9:4-10:40. Limelight’s construction would improperly exclude this embodiment, limiting the patent to an alternative embodiment allowing a customer to select content files that are to be transferred to edge servers. *Oatey Co. v. IPS Corp.*, 514 F.3d 1271, 1276 (Fed. Cir. 2008) (courts “normally do not interpret claim terms in a way that excludes embodiments disclosed in the specification.”). **Finally**, as noted above with respect to “filling the global cache,” claim 8 also does not require that “**all available storage** of a portion of the global cache” be filled. Bhattacharjee Decl. ¶¶ 50-51.

e) “close to end users” (claims 1, 8)

Akamai’s Construction	Limelight’s Construction
Plain and ordinary meaning; no construction needed Alternatively, if the Court deems construction to be necessary “close to end users” should be construed as “proximate to end users”	<i>Indefinite</i>

The parties dispute whether the term “close to end users” is indefinite.

The Supreme Court has confirmed that a patent is invalid for indefiniteness only “if its claims, read in light of the specification delineating the patent, and the prosecution history, fail to inform, with reasonable certainty, those skilled in the art about the scope of the invention.” *Nautilus*, 134 S. Ct. at 2124. The Court has further confirmed that the party claiming a term is indefinite—in this case, Limelight—bears the burden of proof, and must establish indefiniteness by clear and convincing evidence. *Id.* at 2130 n.10.

Applying this standard, there is no basis for Limelight’s contention that the claims are indefinite. As Akamai’s expert, Dr. Bhattacharjee, has confirmed, the concept of locating an edge server “close to end users” is well understood to persons of skill in the relevant art. Bhattacharjee Decl. ¶ 53. Specifically, CDNs have long been structured around the principle that content should be accessed from a nearby server at the “edge” of the network—rather than from a distant server. *Id.* In fact, the specification confirms that the patent uses the term consistent with this understood meaning. *See, e.g.*, ’211 patent, 4:54-55 (“Edge servers 103 are **known** to be servers distributed at the edge of a network area.”). As the specification explains, the edge servers are located close to end users so that they can then “serve content to user requests within [that] region,” *id.* 4:57-58, and so that an “appropriate” edge server may be determined by “geographic location in relation to the user,” *id.* 5:66-6:2. As confirmed by Dr. Bhattacharjee (¶¶ 52-57), this term accordingly “inform[s], with reasonable certainty, those skilled in the art about the scope of the invention,” and is thus not indefinite. *Nautilus*, 134 S. Ct. at 2124.

In fact, Limelight’s suggestion that this concept of locating an edge server “close to end users” would **not** be known to persons of skill in the art is flatly inconsistent with Limelight’s own patents and public statements. For example, Limelight’s own U.S. Patent No. 8,626,876 notes that “CDNs can include geographically distributed points of presence (POPs) to locate

edge servers close to end users.” Ex. 3 (’876 patent) at 1:15-17. Similarly, in a press release and data sheet, Limelight said that its own “Cloud Storage” product can “optimize worldwide delivery and management of media and software content by positioning content *close to end-users.*” Ex. 4 (June 16, 2015 Press Release); Ex. 5 (Limelight Data Sheet); *see also* Bhattacharjee Decl. ¶ 55.

Other sources confirm that “close to end users” is well understood by persons of skill in the art, and far from indefinite. In previous litigation between these parties in this District, for example, the court found that when the term “close to end users” appeared in the specification, it shed light on the meaning of the meaning of the claims—confirming that the term is far from ambiguous. *Akamai Techs., Inc. v. Limelight Networks, Inc.*, No. CIV. A. 06-11109-RWZ, 2008 WL 697707, at *2-3 (D. Mass. Feb. 8, 2008). As Dr. Bhattacharjee confirms, textbooks in the field—including those published before the ’211 patent’s claimed priority date—have also used this (or materially identical) language to describe the design of CDNs. Bhattacharjee Decl. ¶ 54; *see also* Ex. 6 (Rabinovich (2002)) at 248 (“A CDN can achieve scalable content delivery by distributing load among its servers, by serving client requests from servers that are *close to requesters*”); Ex. 7 (Tanenbaum (2002)) at 694 (“CDNs allow documents to be dynamically replicated to servers *close to users.*”); Ex. 8 (Peterson (2003)) at 817 (“The goal of widely distributing the surrogate servers is to have a surrogate *close to the client*, making it possible to respond to requests more quickly.”); Ex. 9 (Gourley (2009)) (“[C]aches ... keep copies of popular web pages *close to clients.*”); Ex. 10 (Molina (2012)) at 2 (CDNs, including Limelight’s, “deliver[] content *close to end users*”). Other publications predating the ’211 patent’s priority date used similar language. Ex. 11 (Besen (2001)) at 292 (“[C]ontent distribution services (CDS) ... replicate web pages at locations *close to end users* in order to deliver information in

the most effective way.”); Ex. 12 (Kapitskaia (2000)) at 202 (“To achieve fast performance in such systems, it is desirable to cache information *close to the (client) applications* that access the (server) directory.”); Ex. 13 (*InformationWeek* (1999)) at 79 (describing product that “allows users to store frequently used content *close to end users*”).

To the extent a construction is required, it should be “proximate to end users,” which reflects the understanding of a person of ordinary skill in the art, who would understand that content delivery is faster when the requested content is proximate to the user that requested it. *See* Bhattacharjee Decl. ¶ 57.

f) “high-bandwidth content” (claims 1, 8)

Akamai’s Construction	Limelight’s Construction
“content requiring high-bandwidth delivery, including, for example, video and other multimedia files”	<i>Indefinite</i>

Limelight also contends that the term “high-bandwidth” is indefinite. Like the concept of locating an edge server “close to end users,” however, the term “high bandwidth content” is well understood by persons of skill in the art. As Akamai’s expert Dr. Bhattacharjee explains, this concept would be clear to a person of ordinary skill in the art, who would understand it to reflect the conditions necessary for satisfactory delivery of requested content, such as smooth, interruption-free playback. Bhattacharjee Decl. ¶¶ 58-59. The specification uses the term consistent with this well understood meaning, describing the claimed invention as a solution to known constraints in bandwidth for “large multimedia files” and noting that the “bottleneck” for delivery of these files was “the bandwidth that the [content providers’] servers could handle. ’211 patent, 1:25-27; Bhattacharjee Decl. ¶ 60. Akamai’s proposed construction accurately reflects the use of the term in the claims and specification, including that certain content—especially multimedia files—must be delivered with higher bandwidth than other

content (such as text).

Additionally, claims 1 and 8 both include the further limitation that “the high-bandwidth content is ***video content***” (*id.* 12:45-46; 14:10-11)—further illuminating the meaning of “high-bandwidth content” while also providing clear guidance to a person of ordinary skill in the art regarding the scope of the claimed invention. Bhattacharjee Decl. ¶ 61. As other courts have found, identification of a specific threshold for “high bandwidth” is unnecessary where the rest of the claim provides further detail about the scope of the claimed invention. *See, e.g., Microunity Sys. Eng’g, Inc. v. Acer, Inc.*, No. 2:10-CV-91-LED-RSP, 2013 WL 12134146, at *5 (E.D. Tex. Mar. 8, 2013) (construing “***high bandwidth*** external interface” to mean “an interface between the media processor and external sources of data,” and finding “that it [was] not necessary to specify a rate constraint within the construction of this term because surrounding limitations in the claims where the ‘high bandwidth external interface’ term appears specify the rate constraint”).

Limelight’s position that the term is indefinite is further belied by its own contentions in this case, which show that Limelight clearly understands the scope of “high-bandwidth content.” For example, in its Invalidity Contentions, Limelight noted that “[t]o the extent that the 703 Patent does not expressly state that the content served from the global cache is high-bandwidth, one of ordinary skill in the art would have understood from 703 Patent’s disclosure that the content served from the edge server can be video content that ***therefore the content served is indeed high-bandwidth.***” Ex. 14 (8/11/17 Invalidity Contentions, A.1) at 5; *see also* Ex. 15 (8/11/17 Invalidity Contentions, A.2) at 26 (“[O]ne of ordinary skill would have known that the higher-level cache server or servers fill the lower level global cache edge servers with popular, high-bandwidth content (***including video content***).”).

The '211 patent is focused on solving known issues with delivery of high-bandwidth content, and its claims are explicitly restricted to video content—which Limelight itself has treated as *per se* high-bandwidth. As the specification confirms, and as confirmed by the declaration of Dr. Bhattacharjee (¶¶ 58-62), Limelight cannot meet its burden of establishing that this term does not “fail to inform, with reasonable certainty, those skilled in the art about the scope of the invention,” and the term is thus not indefinite. *Nautilus*, 134 S. Ct. at 2124.

g) “popular” / “popularity” (claims 1, 3, 8)

Akamai’s Construction	Limelight’s Construction
“content receiving a relatively higher number of requests than other content”	<i>Indefinite</i>

Limelight also contends that the terms “popular” and “popularity” are indefinite. Once again, however, these terms have a reasonably certain meaning to persons of skill in the art, particularly when “read in light of the specification delineating the patent,” as the Supreme Court’s *Nautilus* decision requires. 134 S. Ct. at 2124. As Dr. Bhattacharjee confirms, a person of skill in the art would understand that the “popularity” of content is based on the number of requests for content, relative to requests for other content. Bhattacharjee Decl. ¶ 64. Akamai’s proposed construction is not only consistent with this well understood meaning, but also firmly grounded in the specification, which contains ample description of what constitutes “popular” content:

- “The popularity of the content in the system is based on how many requests an edge server gets for that content within a specified period of time.” ’211 patent, 9:21-23.
- “The files are ranked from most popular to least popular (among the stored files). The edge server records each time a file is requested. It periodically reviews the requests and, when a file has been requested from an edge server a sufficient number of times to become more popular than the lowest popular stored file, the file is obtained from an origin site.” *Id.* 3:57-63.
- “The edge servers 103 operate using a file popularity profile. An edge server

maintains a dynamic number of popular files in its local memory for the customer. The files are ranked from most popular to least popular (among the stored files). When a file has been requested from an edge server a sufficient number of times to become more popular than the lowest popular stored file, then the file is obtained from the origin site 102.” *Id.* 4:65-5:5.

- “The edge server checks the frequency of requests over a certain amount of time (configurable by an administrator) to see if any new files have gained enough popularity to replace another file in the local storage. If so, it then requests the file from the origin site 102.” *Id.* 6:12-17.

“Popularity” thus depends on the number of requests for content, relative to requests for other content. *E.g., id.* 5:1-5 (“When a file has been *requested* from an edge server a *sufficient number of times* to become *more popular* than the lowest popular stored file”); *see also* Bhattacharjee Decl. ¶¶ 65-70

As multiple courts have confirmed, “popularity” is not an indefinite term, particularly where, as here, the specification provides ample guidance on the use of the term. For example, one court recently found that “popular music videos” was entitled to its plain and ordinary meaning where the specification was clear that “popular videos are those videos that have a relatively high frequency of download, such as the ‘local top 10 videos.’” *Music Choice v. Stingray Digital Grp. Inc.*, No. 2:16-CV-586-JRG-RSP, 2017 WL 2896025, at *23 (E.D. Tex. July 6, 2017); *see also Premier Int’l Assocs. v. Apple Computer, Inc.*, 512 F. Supp. 2d 737, 746 (E.D. Tex. 2007) (construing “popularity” in light of specification).

Akamai’s proposed construction—“content receiving a relatively higher number of requests than other content”—faithfully reflects the disclosure of the specification, as it would be understood by a person of ordinary skill in the art. Limelight’s indefiniteness claim, by contrast, is undermined by the extensive discussion of popularity in the patent. Therefore, as confirmed by the declaration of Dr. Bhattacharjee, ¶¶ 63-70, Limelight has not met its burden of establishing that these terms “fail to inform, with reasonable certainty, those skilled in the art

about the scope of the invention,” and they are thus not indefinite. *Nautilus*, 134 S. Ct. at 2124.

B. Limelight’s ’661 Patent

U.S. Patent No. 8,775,661 (the “’661 patent”) (No. 16-10253, Dkt. No. 22-2) relates to combining CDN technology with peer-to-peer content delivery (“P2P”). Generally speaking, P2P networks deliver content from one end user to another, rather than delivering content from a server (such as a CDN server) to an end user. P2P networks became popular in the late 1990s as a way for end users to share content directly with one another. ’661 patent, 1:37-46; 1:62-2:2; 2:7-9. As the patent acknowledges, both CDNs and P2P networks were well known before the ’661 patent. *Id.* 1:27-46.

The ’661 patent claims **combining** a CDN with a P2P network. *Id.* 2:20-26. Specifically, the patent proposes a purportedly novel hybrid CDN/P2P system that supports **both** “traditional origin download” (from a CDN) and “user originated download” (from a “peer” end user in a P2P network). *Id.* 1:28-41; 2:21-26; 4:42-64.

When a user (or “client”) requests content, this hybrid network must determine whether the content should be delivered from the CDN or from another end user (peer), and if the latter, which peer should deliver the content. To achieve this, the patent discloses a “global resource authority infrastructure” that determines global policies, which may, for example, identify the circumstances under which clients are authorized to receive requested content from peers instead of from CDN servers, and which also may identify the peers eligible to provide the requested content. *See id.* claim 1; 2:30-37; 11:19-25.

Figure 1 depicts the “global resource authority” (item 104), which receives status information and settings from clients, content providers, and/or ISPs and uses this information to determine global policies. *Id.* 4:43-46; 4:53-61; 5:58-67; 11:15-25; fig. 1. In contrast, Figure 2 depicts a “local resource authority” that performs a similar role locally, for example by receiving

status information and settings from clients within a local network and transmitting the resulting policies to the global resource authority. *Id.* 8:62-9:11; fig. 2.

1. “[the identified] one or more [global] policies” (all claims)

Akamai’s Construction	Limelight’s Construction
“global policies” = “policies determined by a global resource authority from settings received from one or more entities”	“policies that can be used to manage content delivery to, from and between clients”

The parties dispute whether “global policies” should be construed consistent with the plain meaning of the claim language and specification (as Akamai proposes), or whether the Court should read out of the claims the requirement that the policies be “**global**” (as Limelight proposes).³

Akamai’s construction is consistent both with the plain language of the claims and with the specification. The claims plainly require “**global** policies,” not “**local** policies.” Consistent with the claim language, the specification expressly distinguishes between “**global** policies” and “**local** policies.” *Compare* ’661 patent Abstract (“**global** policies”), *with* 16:33-41 (“**local** policies”). Likewise, the patent distinguishes between a “**global** resource authority” and a “**local** resource authority.” *Compare id.* 5:62-65 (“**global** resource authority”), *with* 8:64-9:12 (“**local** resource authority”).

As the specification confirms, these “policies” are determined by “resource authorities”—“global policies” are determined by a “global resource authority,” and “local

³ As mentioned above, “global” is also used in Akamai’s ’211 patent claims. While the two specifications are different, and do not require the same construction, Akamai’s construction here is not inconsistent with the use of “global” in the ’211 patent. Both patents distinguish between “global” entities and “local” or “regional” ones.

policies” by a “local resource authority.” *Id.* 5:62-65 (“The **global resource authority** 104 uses the status information and control layer settings to determine **a set of policies**....”); 11:15-33 (describing multiple aspects in which different settings are sent to the **global resource authority** which “then determines the **policies**”); 7:41-44 (“the **global resource authority** 104 determines the **policies**”); 8:64-9:12 (“The **local resource authority** 208 operates within a limited network and performs locally a similar role as the global resource authority.... In one aspect, the **local resource authority** 208 processes the various control layer settings and **provides the resultant setting policies**....”); *see also id.* 8:25-27; fig. 1 (showing a “global resource authority”); fig. 2 (showing a “local resource authority”). Because policies are determined by resource authorities, and the claims require global, not local, policies, it follows that the claimed “**global** policies” are policies that are determined by the “**global** resource authority.” *See id.* 5:62-65; 11:15-33; 4:41-44; 8:64-9:12; *see also id.* 8:25-27; figs. 1 & 2. Akamai’s construction thus makes clear that the “**global** policies” are different than “**local** policies” and are created by a “**global** resource authority” instead of a “**local** resource authority.”

The specification also repeatedly explains that these “policies” are determined from “settings.” *Id.* 2:39-44; 2:63-3:7; 3:25-37; 5:62-65 (“The global resource authority 104 uses the status information and control layer settings to determine a set of policies....”); 7:41-44 (“the global resource authority 104 determines the policies from the various control layer settings, status information and the like”); 11:15-20; 14:50-60; fig. 5, 15:60-65; 16:1-6; 16:15-18 (“control layer policies are determined based on ... control layer settings”).

Moreover, the specification confirms that the global resource authority determines the global policies based on “settings received from one or more entities.” *Id.* 2:39-44 (“A first setting from a first entity ... is received,” “[a] second setting from a second entity ... is

received,” and “[a] policy based on the first and second settings is determined.”); 2:63-3:7; 3:25-37; *see also id.* 11:61-13:48 (describing Tables II-V containing settings received from multiple entities: applications, clients, ISPs, and the system); fig. 5 (Step 510 (“Receiving a First Setting from a First Entity”), Step 520 (“Receiving a Second Setting from a Second Entity”), Step 540 (“Determining a Policy Based on the First and Second Settings”)); 15:60-16:33 (explaining flow of Figure 5); fig. 4; 11:15-20; 14:31-60.

The claims and specification thus make clear that “policies” are determined from “settings,” and that “**global** policies” are different from other policies—including “**local** policies.”

In contrast, Limelight’s proposed construction improperly seeks to read the word “global” out of the term entirely, so that it would encompass virtually any type of policy—including “**local** policies” that the patent expressly distinguishes from “global policies.” *Compare* ’661 patent Abstract (“**global** policies”), *with* 16:33-41 (“**local** policies”). Limelight’s construction accordingly violates the well-established principle that claims should be construed to give effect to *every* claim term. *See Akzo Nobel Coatings, Inc. v. Dow Chem. Co.*, 811 F.3d 1334, 1340 (Fed. Cir. 2016) (rejecting construction that “would render [claim term] ‘collection’ entirely superfluous”); *Merck*, 395 F.3d at 1372 (“A claim construction that gives meaning to all the terms of the claim is preferred over one that does not do so.”). Moreover, Limelight’s construction—that the “global policies” “**can** be used to manage content delivery to, from and between clients”—is vague and unhelpful to the trier of fact. Limelight’s construction describes only what “global policies” **can** be used to do, rather than stating what the global policies actually **are** or actually **do**. Instead, the key points to clarify for the jury are that the “global policies” are different from other policies (such as “local policies”) because they are determined

by a **global** resource authority, and that “policies” are distinct from, and indeed determined from, “settings.” Limelight’s construction provides no assistance to, and will likely confuse, the jury, while Akamai’s construction clarifies these points.

2. “circumstances” (all claims)

Akamai’s Construction	Limelight’s Construction
“more than one condition”	“one or more condition(s)”

The “global policies” described in the claims identify, among other things, “circumstances” under which a client (*i.e.*, a computer requesting content) is authorized to receive requested content from “other clients” (*i.e.*, a P2P network) rather than from a CDN. The parties agree that “circumstances” encompass situations in which the global policies identify “more than one condition” under which content can be received from other clients. However, the parties dispute whether “circumstances”s—a plural form of the noun—is plural, meaning more than one condition (as Akamai proposes), or may also include just one condition (as Limelight proposes).

Basic grammar dictates that Akamai is correct. Akamai’s construction makes clear that the claims’ requirement of “circumstancess” under which clients are authorized to receive content requires that there be “more than one” circumstance—in other words, a plural claim term requires plurality. *See Leggett & Platt, Inc. v. Hickory Springs Mfg. Co.*, 285 F.3d 1353, 1357 (Fed. Cir. 2002) (“[T]he claim recites ‘support wires’ in the plural, thus requiring more than one welded ‘support wire.’”); *Shire Dev. LLC v. Amneal Pharm. LLC*, No. 15-2865 (RBK/JS), 2016 WL 4119940, at *11 (D.N.J. Aug. 2, 2016) (“consisting of substances” and “consists of compounds” requires “at least two substances/compounds”); *Phoenix Licensing, L.L.C. v. AAA Life Ins. Co.*, No. 2:13-cv-1081, 2015 WL 3866832, at *10 (E.D. Tex. June 22, 2015) (as to

claim term “distinct choices”: “The Court agrees with Defendants that the plain language of plural ‘choices,’ as well as the adjective ‘distinct,’ requires more than just one choice.”).

The specification further confirms that “circumstances” requires “more than one condition.” The specification does not ever use the word “circumstances,” but it does use the word “conditions”—plural—to describe how global policies (determined by a global resource authority) identify when one client is authorized to receive content sourced from other clients (in home or business networks). *See* ’661 patent, 5:29-33 (“[T]he global resource authority 104 regulates which download managers or clients in the home and business networks 132, 136 can request and source a particular content object and under what *conditions* those clients may do so.”).

The parties agree that a “circumstance” is a “condition.” Thus, “circumstances” must mean “more than one condition.”

3. “one or more authorized transfers of [at least a portion] of the content” (all claims)

Akamai’s Construction	Limelight’s Construction
“one or more transfers that satisfy the global policies determined by the global resource authority”	“content (or a portion thereof) permitted to be obtained by a data transfer”

The primary dispute between the parties is whether the term “*authorized* transfers” means, consistent with the plain language of the claims and the specification as a whole, transfers authorized by a “global resource *authority*” (as Akamai contends), or whether “authorized transfers” are merely “permitted transfers” (as Limelight argues).

Akamai’s construction is consistent with the plain meaning of the claim language and the specification. The claim language confirms that the “authorized transfers” are “authorized” based on the “global policies.” For example, claim 1 requires that “one or more *authorized transfers* of at least a portion of the content” are “identif[ied] using the one or more *global*

policies and the content identifier.” ’661 patent, 18:5-7; *see also id.* claim 5, 18:42-49; claim 13, 19:25-30. As discussed above, these “global policies” are determined by a “**global** resource authority”—and not by a “local resource authority.” *See supra* section IV.B.1, pp. 25-28. The “**authorized** transfers” are accordingly transfers that satisfy global policies determined by a global resource **authority**. *See, e.g.*, Ex. 16 (Merriam-Webster’s (2007)) (“authorize”—“to establish by or as if by **authority**”).

Akamai’s construction is also consistent with the specification. Although the specification never uses the term “authorize” or “authorized transfer,” it does repeatedly refer to “resource authorities” such as the “global resource **authority**,” which “determine[s] a set of policies which, for example, control content **transfer** over the network.” ’661 patent, 5:62-65. This is consistent with claims’ description of the “authorized transfers” as being determined based on the “global policies.” *Id.* claims 1, 5, 13. Indeed, the only “authorities” mentioned in the specification are the global resource authority and the local resource authority.

Limelight’s construction, by contrast, simply substitutes the word “permitted” for “authorized,” providing no clarification for the trier of fact. Indeed, neither “permit” nor “permitted” appears anywhere in the specification. Even if the word “permitted” were consistent with the word “authorized” outside of the context of the patent, Limelight’s construction adds no clarity to the jury’s understanding of the term “authorized transfer.” Additionally, Limelight’s construction is so broad that it would allow permission by any authority—even a local one—despite the clear claim language that authorization must be based on global policies.

V. CONCLUSION

For the reasons set forth above, Akamai respectfully requests that the Court adopt its proposed constructions.

Dated: October 20, 2017

Respectfully submitted,

/s/ Michael J. Summersgill

Michael J. Summersgill (BBO No. 632816)

Cynthia D. Vreeland (BBO No. 635143)

Timothy D. Syrett (BBO No. 663676)

WILMER CUTLER PICKERING

HALE AND DORR LLP

60 State Street

Boston, Massachusetts 02109

Tel: (617) 526-5000

michael.summersgill@wilmerhale.com

cindy.vreeland@wilmerhale.com

timothy.syrett@wilmerhale.com

Mary V. Sooter (CO Bar No. 35136)

(*pro hac vice*)

WILMER CUTLER PICKERING

HALE AND DORR LLP

1225 17th Street, Suite 2600

Denver, CO 80201

Tel: (720) 274-3135

mindy.sooter@wilmerhale.com

Arthur W. Coviello (BBO No. 670152)

(*pro hac vice*)

WILMER CUTLER PICKERING

HALE AND DORR LLP

950 Page Mill Road

Palo Alto, CA 94303

Tel: (650) 858-6000

arthur.coviello@wilmerhale.com

Attorneys for Plaintiff

Akamai Technologies, Inc.

CERTIFICATE OF SERVICE

I hereby certify that on October 20, 2017, a true copy of the foregoing will be electronically filed with the Clerk of Court using the CM/ECF system, which will then send a notification of such filing (NEF).

/s/ Michael J. Summersgill
Michael J. Summersgill