

Filed: April 15, 2025

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC,

Petitioner,

v.

SANDPIPER CDN, LLC,

Patent Owner.

Case: IPR2025-00826

U.S. Patent No. 9,021,112

Issue Date: April 28, 2015

Title: Content Request Routing and Load Balancing for
Content Distribution Networks

PETITION FOR *INTER PARTES* REVIEW

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	STANDING	1
III.	RELIEF REQUESTED	1
IV.	THE '112 PATENT	2
A.	Overview	2
B.	Prosecution History	3
C.	POSITA	3
V.	CLAIM CONSTRUCTION	3
VI.	PRIOR ART	4
A.	Glines	4
B.	Gupta	5
C.	Chung	6
D.	Chiou	6
E.	Colby	6
VII.	Grounds 1A/1B: Glines or Glines-Gupta	7
A.	Motivation to Combine Glines with Gupta	7
B.	Claim 1	8
C.	Claim 2: “some ... CDN DNS servers ... located at ... (POPs)”	26
D.	Claim 3: “each ... CDN DNS server[] ... located at ... (POPs)”	30
E.	Claim 4: “grouped as a cluster”	31
F.	Claim 5: “each ... content server[] is associated with a ... (POP)”	32

G.	Claim 6: “attempting to serve the content”	33
H.	Claim 7: “if a valid version of the content is available ... serving .. otherwise obtaining”	34
I.	Claim 8: “origin server associated with said content provider”	37
J.	Claim 10: “load balancing policy”	38
K.	Claim 11: “switch, selecting the content server”	38
L.	Claim 12: “predetermined selection policy”	39
M.	Claim 13	39
N.	Claim 14: “load balancing policy”	47
O.	Claims 15-17	47
P.	Claim 18	48
Q.	Claim 19	50
R.	Claim 20: “load balancing policy”	51
S.	Claim 21	51
T.	Claim 22	52
U.	Claim 23	54
VIII.	Grounds 2A/2B: Glines, with or without Gupta, and Chung	58
A.	Motivation to Combine	58
B.	Claim 4: “grouped as a cluster”	60
IX.	Grounds 3A/3B: Glines, with or without Gupta, and Chiou	61
A.	Motivation to Combine	61
B.	Claim 7: “if a valid version of the content is available ... serving .. otherwise obtaining”	63
C.	Claim 8: “origin server associated with said content provider”	64

D.	Claim 9: “us[ing] a table”	65
E.	Claim 19	66
F.	Claim 20: “load balancing policy”	67
X.	Grounds 4A/4B and 5A/5B: Glines, with or without Gupta and Chiou, and Colby	68
A.	Motivation to Combine	68
B.	Claim 10: “load balancing policy”	69
C.	Claim 11: “switch, selecting the content server”	69
D.	Claim 12: “predetermined selection policy”	70
E.	Claim 13	71
F.	Claim 14: “load balancing policy”	73
G.	Claims 15-17	74
H.	Claim 18	74
I.	Claim 19	75
J.	Claim 20: : “load balancing policy”	77
K.	Claim 21	77
L.	Claim 22	79
XI.	35 U.S.C. § 314.....	80
XII.	35 U.S.C. § 325.....	81
XIII.	MANDATORY NOTICES UNDER 37 C.F.R. § 42.8.....	82
A.	Real Party-in-Interest	82
B.	Related Matters.....	82
C.	Lead and Backup Counsel Information.....	82

XIV. CONCLUSION.....	83
----------------------	----

TABLE OF AUTHORITIES

Cases	Page(s)
<i>Alcon Inc. v. AMO Dev., LLC</i> , IPR2021-00843, Paper 15 (P.T.A.B. Nov. 12, 2021).....	46
<i>B/E Aerospace, Inc. v. C&D Zodiac, Inc.</i> , 962 F.3d 1373 (Fed. Cir. 2020)	40
<i>Becton, Dickinson & Co. v. B. Braun Melsungen AG</i> , IPR2017-01586, Paper 8 (P.T.A.B. Dec. 15, 2017)	81
<i>Microsoft Corp. v. Parallel Networks Licensing LLC</i> , IPR2015-00486, Paper 10 (P.T.A.B. July 15, 2015).....	81
<i>Sotera Wireless, Inc. v. Masimo Corp.</i> , IPR2020-01019, Paper 12 (P.T.A.B. Dec. 1, 2020)	80, 81
<i>Target Corp. v. Proxicom Wireless, LLC</i> , IPR2020-00931, Paper 10 (P.T.A.B. Nov. 10, 2020).....	46
<i>Thorne Rsch., Inc. v. Trustees of Dartmouth Coll.</i> , IPR2021-00491, Paper 18 (P.T.A.B. Aug. 12, 2021).....	82
<i>Zillow Grp., Inc. v. Int’l Bus. Machs. Corp.</i> , IPR2020-01656, Paper 8 (P.T.A.B. Mar. 15, 2021).....	46
Statutes	
35 U.S.C. § 314.....	80
35 U.S.C. § 325.....	81
Regulations	
37 C.F.R. § 42.8	82

All emphases added unless otherwise indicated.

LIST OF EXHIBITS

Exhibit	Description
EX-1001	U.S. Patent No. 9,021,112 (“the ’112 patent”)
EX-1002	WO2001039470 A1 to Glines (“Glines”)
EX-1003	U.S. Patent No. 6,405,252 to Gupta (“Gupta”)
EX-1004	U.S. Patent No. 6,470,389 to Chung (“Chung”)
EX-1005	U.S. Patent No. 6,792,507 to Chiou (“Chiou”)
EX-1006	U.S. Patent No. 6,006,264 to Colby (“Colby”)
EX-1007	Prosecution History of the ’112 patent
EX-1008	Declaration of Todd C. Mowry, Ph.D.
EX-1009	Curriculum Vitae of Todd C. Mowry, Ph.D.
EX-1010	U.S. Patent No. 8,041,814 to Chase
EX-1011	James Kurose & Keith Ross, Computer Networking: A Top-Down Approach Featuring the Internet (1st ed. 2000)
EX-1012	Andrew Tanenbaum, Computer Networks (3rd ed. 1996)
EX-1013	U.S. Patent No. 6,829,654 to Jungck
EX-1014	James Kurose & Keith Ross, Computer Networking: A Top-Down Approach Featuring the Internet (3rd ed. 2005)
EX-1015	Mark Baker & Rajkumar Buyya, Cluster Computing at a Glance <i>in</i> High Performance Cluster Computing – Architectures and Systems, 3-47 (1999).
EX-1016	U.S. Patent No. 6,578,066 to Logan
EX-1017	U.S. Patent No. 8,041,814 to Chase
EX-1018	U.S. Patent No. 7,523,181 to Swildens
EX-1019	U.S. Patent No. 7,594,030 to Teodosiu
EX-1020	Order Granting Stipulation Regarding Amended Complaint and Case Schedule, <i>Sandpiper CDN, LLC v. Google LLC</i> , 2:24-cv-03951, Dkt. 54 (C.D. Cal. Jan. 21, 2025)
EX-1021	Sandpiper’s Disclosures of Asserted Claims & Infringement Contentions dated Jan. 14, 2025 in <i>Sandpiper CDN, LLC v. Google LLC</i> , 2:24-cv-03951.
EX-1022	Exhibit C to Sandpiper’s Disclosures of Asserted Claims & Infringement Contentions dated Jan. 14, 2025 in <i>Sandpiper CDN, LLC v. Google LLC</i> , 2:24-cv-03951.
EX-1023	Schroeder, et al., <i>Scalable Web Server Clustering Technologies</i> , IEEE Network, Vol. 14, Issue 3 (June 2000).

Exhibit	Description
EX-1024	The Wayback Machine – Yahoo! News, http://dailynews.yahoo.com/ , dated February 2, 2001 (<i>available at</i> https://web.archive.org/web/20010202075000/http://dailynews.yahoo.com/)
EX-1025	The Wayback Machine, “U.S. Senate Approves Ashcroft, 58-42,” http://dailynews.yahoo.com/h/nm/20010201/ts/congress_ashcroft_dc_71.html , dated February 2, 2001 (<i>available at</i> https://web.archive.org/web/20010202073200/http://dailynews.yahoo.com/h/nm/20010201/ts/congress_ashcroft_dc_71.html)
EX-1026	A. Bestavros et al., <i>Distributed Packet Rewriting and its Application to Scalable Server Architectures</i> , Proceedings Sixth International Conference on Network Protocols, IEEE (Oct. 1998).
EX-1027	Letter to Robert H. Reckers <i>re Sandpiper CDN LLC v. Google LLC</i> , No. 2:24-3951-AB-RAO (C.D. Cal.) (Apr. 14, 2025)

CLAIM LISTINGS

Independent Claim 1

[1p]	A method of content delivery in a network, the method comprising:
[1a]	(A) providing a plurality of Domain Name System (DNS) servers associated with a Content Delivery Network (CDN), said plurality of CDN DNS servers sharing a common anycast address, wherein each CDN DNS server is associated with a respective plurality of content servers;
[1b]	(B) causing said plurality of CDN DNS servers to be authoritative for a hostname associated with a content provider by causing said common anycast address to be associated with said hostname;
[1c]	(C) responsive to a request for content associated with the content provider and issued by a client, said request including at least said hostname,
[1d]	(c1) causing said hostname to be resolved to said common anycast address by an Internet Service Provider (ISP) DNS server; and then
[1e]	(c2) by one of said plurality of CDN DNS servers, resolving said hostname to identify an IP address for use by the client to retrieve the content from a content server.

Dependent Claim 2

2	The method as recited in claim 1 wherein at least some of said plurality of CDN DNS servers are located at network Points of Presence (POPs).
----------	---

Dependent Claim 3

3	The method as recited in claim 2 wherein each of said plurality of CDN DNS servers is located at a network POP.
----------	---

Dependent Claim 4

4	The method as recited in claim 2 wherein at least one of said plurality of CDN DNS servers and respective plurality of content servers is grouped as a cluster.
----------	---

Dependent Claim 5

5	The method as recited in claim 1 wherein each of said plurality of content servers is associated with a network Point of Presence (POP).
----------	--

Dependent Claim 6

6	The method as recited in claim 1 further comprising: (D) by the content server, attempting to serve the content.
----------	--

Dependent Claim 7

7	The method as recited in claim 6 wherein said step (D) of attempting comprises: (d1) if a valid version of the content is available on the content server, serving the content; otherwise (d2) obtaining the content from a content source and then serving the content.
----------	--

Dependent Claim 8

8	The method as recited in claim 7 wherein, said content source is an origin server associated with said content provider.
----------	--

Dependent Claim 9

9	The method as recited in claim 7 wherein, if a valid version of the content is not available on the content server, the content server uses a table to ascertain where to obtain the content.
----------	---

Dependent Claim 10

10	The method as recited in claim 1, wherein selection of the content server is based on a load balancing policy.
-----------	--

Dependent Claim 11

11	The method of claim 1 further comprising: (c3) by at least one switch, selecting the content server from said respective plurality of content servers associated with said resolving CDN DNS.
-----------	---

Dependent Claim 12

12	The method of claim 11, wherein the step of (c3) is performed in accordance with a predetermined selection policy.
-----------	--

Independent Claim 13

[13p]	A method of content delivery for delivering content on behalf of a plurality of content providers, the method comprising:
[13a]	(A) providing a plurality of Domain Name System (DNS) servers, said plurality of DNS servers sharing a common anycast address, wherein each DNS server is associated with a respective plurality of content servers;
[13b]	(B) configuring each DNS server and respective plurality of content servers to support content delivery for a plurality of content providers, wherein each of the plurality of content providers has a respective hostname associated therewith;
[13c]	(C1) responsive to a first request for first content associated with a first content provider of said plurality of content providers, said first request including at least a first hostname associated with the first content provider,
[13d]	(c11) by a first DNS server of said plurality of DNS servers, resolving said first hostname to identify a first IP address to be used for retrieval of the first content from a first content server of said content servers associated with said first DNS server;
[13e]	(c12) by a switch associated with said resolving DNS server of step (c11), selecting the first content server from said respective plurality of content servers associated with the DNS server of step (c11); and
[13f]	(c13) attempting to serve said first content from the first content server; and
[13g]	(C2) responsive to a second request for said first content associated with said first content provider of said plurality of content providers, said second request including at least the first hostname associated with the first content provider,
[13h]	(c21) by a second DNS server of said plurality of DNS servers, resolving said first hostname to identify a second IP address for retrieval of the first content from a content server associated with the second DNS server;
[13i]	(c22) by a switch associated with said resolving DNS server of step (c21), selecting the second content server from said respective plurality of content servers associated with the DNS server of step (c21); and

[13j]	(c23) attempting to serve said second content from the second content server.
--------------	---

Dependent Claim 14

14	The method as recited in claim 13, wherein selection of at least one of the first content server and the second content server is based on a load balancing policy.
-----------	---

Dependent Claim 15

15	The method as recited in claim 13, wherein the resolving DNS server of step (c11) is the same as the resolving DNS server performing of step (c21).
-----------	---

Dependent Claim 16

16	The method as recited in claim 15, wherein the switch of step (c12) is the same as the switch of step (c22).
-----------	--

Dependent Claim 17

17	The method as recited in claim 15, wherein the first content server is the same as the second content server.
-----------	---

Independent Claim 18

[18p]	A content delivery system for delivering content on behalf of a plurality of content providers, each of said plurality of content providers having a hostname associated therewith, the system comprising:
[18a]	(A) a plurality of Domain Name System (DNS) servers, said plurality of DNS servers sharing a common anycast address,
[18b]	said plurality of DNS servers being authoritative for said hostnames associated with said plurality of content providers;
[18c]	(B) a set of content servers associated with each DNS server, each set comprising at least a plurality of content servers; and
[18d]	(C) at least one switch associated with each set of content servers, wherein
[18e]	(D) responsive to a request to resolve a hostname associated with a particular content provider of said plurality of content providers, said request being associated with an end-user request for content of said

	particular content provider, one of said plurality of DNS servers is operable to resolve said hostname to identify an IP address for use by the end-user to retrieve the content from a content server,
[18f]	(E) the at least one switch associated with the set of content servers of the resolving DNS server is operable to perform load balancing to select the content server from the set of content servers associated with the resolving DNS server, and
[18g]	(F) the content server is operable to serve the content to the end-user.

Independent Claim 19

[19p]	A method of content delivery, the method comprising:
[19a]	(A) providing a plurality of CDN Domain Name System (DNS) servers, said plurality of DNS servers sharing a common anycast address, each DNS server being associated with a respective plurality of content servers,
[19b]	wherein at least some of said plurality of content servers are located at network Points of Presence (POPs);
[19c]	(B) causing said plurality of DNS servers to be authoritative for a hostname associated with a content provider by causing said common anycast address to be associated with said hostname;
[19d]	(C) responsive to a request for content associated with the content provider by a client, said request including at least said,
[19e]	(c1) by one of said plurality of DNS servers, resolving said hostname to identify an IP address for use by the client to retrieve the content from a content server; and
[19f]	(c2) by at least one switch, selecting the content server from said respective plurality of content servers associated with said resolving DNS server,
[19g]	(D) by the content server, (d1) attempting to serve the content by:
[19h]	(d11) if a valid version of the content is available on the content server, serving the content; otherwise (d12) obtaining the content from a content source and then serving the content.

Dependent Claim 20

20	The method as recited in claim 19, wherein selection of the content server is based on a load balancing policy.
-----------	---

Independent Claim 21

[21p]	A method of content delivery in a network, the method comprising:
[21a]	(A) providing a plurality of Domain Name System (DNS) servers associated with a Content Delivery Network (CDN), the plurality of DNS servers sharing a common anycast address, wherein each DNS server is associated with a set of content servers, each set comprising at least a plurality of content servers;
[21b]	(B) causing the plurality of DNS servers to be authoritative for a hostname associated with a content provider by causing the common anycast address to be associated with the hostname;
[21c]	(C) responsive to a request for content associated with the content provider and issued by a client, the request including at least the hostname,
[21d]	(c1) by one of the plurality of DNS servers, resolving the hostname to identify an IP address for use by the client to retrieve the content from a content server; and
[21e]	(c2) by at least one switch associated with the set of content servers of the resolving DNS server, selecting the content server from the set of content servers associated with the resolving DNS server.

Independent Claim 22

[22p]	A content delivery system comprising a plurality of nodes, wherein each node comprises:
[22a]	a Domain Name System (DNS) server sharing an anycast address with DNS servers in other nodes of the content delivery system,
[22b]	wherein the DNS server is authoritative for hostnames associated with a plurality of content providers;
[22c]	a plurality of content servers; and
[22d]	at least one switch, wherein
[22e]	responsive to a request by a client for content associated with a particular content provider of the plurality of content providers, the request including at least a hostname associated with the particular content provider,
[22f]	the DNS server is operable to resolve the hostname to identify an IP address for use by the client to retrieve the content from a content server,
[22g]	the at least one switch is operable to select the content server from the plurality of content servers, and

[22h]	the content server is operable to serve the content to the client.
--------------	--

Independent Claim 23

[23p]	A method comprising:
[23a]	(A) providing a plurality of Domain Name System (DNS) servers, said plurality of DNS servers sharing a common anycast address, wherein each DNS server is associated with a respective plurality of content servers,
[23b]	wherein each DNS server and respective plurality of content servers is configured to support content delivery for a plurality of content providers;
[23c]	(B) responsive to a first request for first content associated with a first content provider of said plurality of content providers, said first request including at least a first hostname associated with the first content provider,
[23d]	by a first DNS server of said plurality of DNS servers, resolving said first hostname to identify a first IP address of a content server of the plurality of content servers associated with said first DNS server; and
[23e]	(C) responsive to a second request for second content associated with a second content provider of said plurality of content providers, said second request including at least said first hostname associated with the first content provider,
[23f]	by a second DNS server of said plurality of DNS servers, resolving said first hostname to identify a second IP address associated with a content server of the plurality of content servers associated with the second DNS server.

I. INTRODUCTION

Petitioner requests IPR of claims 1-23 U.S. Patent No. 9,021,112.

II. STANDING

Petitioner is not barred or estopped from requesting IPR on these grounds.


III. RELIEF REQUESTED

Ground	Art	Basis	Claims
1A	Glines	§103	1-8, 10-23
1B	Glines-Gupta	§103	1-8, 10-23
2A	Glines-Chung	§103	4
2B	Glines-Gupta-Chung	§103	4
3A	Glines-Chiou	§103	7-9, 19-20
3B	Glines-Gupta-Chiou	§103	7-9, 19-20
4A	Glines-Colby	§103	10-22
4B	Glines-Gupta-Colby	§103	10-22
5A	Glines-Chiou-Colby	§103	19-20
5B	Glines-Gupta-Chiou-Colby	§103	19-20

IV. THE '112 PATENT

A. Overview

The '112 patent discloses “hosting and distributing content on a content delivery network such as the Internet.” EX-1001, 1:16-19. A CDN DNS system 78 of CDN DNS servers 78a-c (triangles) is authoritative for a web domain, represented by a hostname (e.g., www.abcd.com). *Id.*, 6:32-7:32, Fig. 5. To resolve the hostname, the system returns an IP address associated with the CDN DNS servers 78a-c, such as an anycast (IP) address (ovals). *Id.* An anycast address is “a unicast address that is used in multiple places” and may “specify more than one interface” on a network. *Id.*, 5:32-58.



Id., Fig. 5 (annotated); EX-1008 ¶¶41-44.

The user requests and receives the content from a cache server 80 (squares) associated with one of the CDN DNS servers 78 that responds to the anycast address. EX-1001, 6:32-7:32. This “enables distribution of the Web site among the cache servers 80.” *Id.*

B. Prosecution History

The applicant amended the claims during prosecution to require a hostname to be resolved to a common anycast address “by an Internet Service Provider (ISP) DNS server,” EX-1007, 130-145, and argued the prior art lacked “two DNS servers performing two resolves,” *id.*, 86-90, 25-49; *also id.*, 36-37. These features are taught in the prior art, as discussed below. EX-1008 ¶¶45-47.

C. POSITA

A POSITA at the time of the invention would have had at least a bachelor’s degree in computer science, electrical engineering, or a related field, and at least two years of work or research experience in the field of content delivery management or networks. *Id.* ¶¶51-53. Work experience could substitute for education, and additional education could substitute for work experience. *Id.*


V. CLAIM CONSTRUCTION

No terms require construction to resolve this Petition.

VI. PRIOR ART

A. Glines

Glines discloses redirecting client requests for content to a replica of the content using anycast messaging. Glines, Abstract, 11:1-13:14. Glines's replica content servers 50 include DNS servers (or name servers) 53 and content servers 51. *Id.*, 11:1-13:14, Fig. 1. Requests to resolve a domain name are "forwarded as an anycast message over the internetwork (e.g., Internet) 30" to DNS servers 53, and the closest one responds with the network address of an associated content server 51. *Id.*, Abstract, 11:1-13:14, 13:15-16:8, Figs. 1, 2. The user then requests the content from that content server 51. *Id.*



Id., Fig. 1; EX-1008 ¶¶54-55.

B. Gupta

Gupta discloses a network of point of presence (“POP”) servers that share a hostname and determine a web (content) server to provide content. Gupta, Abstract, 2:9-3:8, Fig. 1. The POP servers may include a name server (e.g., servers 110, 120) and multiple WebCache servers (e.g., in POP 10 and 30). *Id.*, 5:16-22, 10:16-22, 10:40-50, 11:1-7, Fig. 1. A user reaches the POP servers via an Internet Service Provider (ISP) cluster 60, which includes a client DNS server 170. *Id.*, 5:65-6:9, 4:45-53, Fig. 1. A “POP” is a point or physical location where two networks connect using, for example, a router. *E.g.*, EX-1008 ¶¶56-59; EX-1013 at 4:65-5:54.





FIG. 1

Id., Fig. 1.

C. Chung

Chung discloses “deploying a set of machines, also known as a cluster, ... to work together to host a single service.” *Id.*, 1:15-40; *also id.*, Abstract, 2:33-50.




Id., Fig. 2; EX-1008 ¶¶60-62.

D. Chiou

Chiou discloses updating a content cache to prevent stored content from becoming “stale.” Chiou, Abstract, 1:28-40, 2:56-59. Chiou determines whether the stored data is “valid data” before returning the content. *Id.*, 12:52-13:40, Fig. 4; EX-1008 ¶¶63-64.

E. Colby

Colby discloses a “content-aware flow switch” that directs content requests to a server “well-suited to serve the content request.” Colby, Abstract, 2:48-58, 5:43-6:28, Fig. 1c.



Id., Fig. 1c; EX-1008 ¶¶65-66.

VII. Grounds 1A/1B: Glines or Glines-Gupta

A. Motivation to Combine Glines with Gupta

It would have been obvious to combine Glines and Gupta. EX-1008 ¶¶68-74. Both disclose methods of delivering content over the Internet. Glines, 1:4-21, 5:24-6:7; Gupta, 1:41-2:17, 2:23-64. This includes, for example, resolving DNS requests and identifying the most efficient server to serve content. Glines, 9:9-11:12; Gupta, 9:1-13, 12:13-37.

Glines and Gupta share a similar architecture to one another—both disclose DNS servers and associated content servers that are connected to a network and function together to resolve hostnames and deliver content. EX-1008 ¶¶68-71. Given

their similarities and shared desire to increase content delivery performance, a POSITA would have understood that Glines and Gupta were analogous art. *Id.* ¶72. And given both teach well-known methods, a POSITA would have had a reasonable expectation of success in the combination. *Id.*

The combination would have involved, for example, applying Gupta’s use of multiple WebCache servers with each name server to Glines to yield a system with multiple content servers for each name server. *Id.* ¶¶73-74 The combination also would have applied Gupta’s use of an ISP DNS server to Glines to resolve URLs or hostnames for residential Internet users. *Id.*


B. Claim 1

Glines, alone or in combination with Gupta, would have rendered claim 1 obvious.

1. [1p] A method of content delivery in a network, the method comprising:

Glines discloses a method of content delivery in a network. For example, Glines “redirect[s] client computer requests for content files to the closest replica of the requested content, by using anycast messaging,” Glines, Abstract, 1:4-2:3, and it “replicat[es] copies of popular content and locat[es] the copies throughout various locations in the network” to reduce original-server demand. *Id.*, 2:4-3:3; *also id.*, 5:24-8:23, 9:9-21, 10:8-11:5, 12:1-3; EX-1008 ¶¶75-77.

Glines's content is delivered over, for example, network 10, [internetwork 30](#) (e.g., the Internet), local area network (LAN) 14, and the “common local area network segments” of replica content servers 50-1, 50-2, 50-3, each of which is a “network” as claimed. Glines, 1:4-3:3, 3:19-5:7, 6:24-6:2, 6:3-7:13, 8:14-23, 9:9-15, 11:14-28, 13:3-8, 15:15-25.



Id., Fig. 1 (annotating internetwork 30); EX-1008 ¶¶78-81.


2. [1a] (A) providing a plurality of [Domain Name System \(DNS\) servers](#) associated with a [Content Delivery Network \(CDN\)](#), the plurality of [CDN DNS servers](#) sharing a [common anycast address](#), and where each [CDN DNS server](#) is associated with a respective plurality of [content servers](#);

Glines provides “a group of [name servers](#) ... located throughout [a] [network](#).” Glines, 6:9-22. The name servers are DNS servers. EX-1008 ¶¶82-83.

Glines’s computer network 10, including internetwork 30, local area network 14, and name server local area networks, is a Content Delivery Network (CDN). EX-1008 ¶84. Computer network 10 (square, below) “redirect[s] client computer requests for content files to the closest replica ... to a name service provided by a group of name servers distributed in the network.” Glines, 9:9-15; *also id.*, 11:14-22.

Each of Glines’s name servers is “addressable via a common anycast address” as claimed. *Id.*, 6:9-22, 3:4-18, 4:10-5:8, 6:8-22, 7:9-8:23; EX-1008 ¶85. Glines further discloses that its name servers 53 “advertise themselves as reachable at an anycast address [(rectangles, below)],” such as 50.100.20.1. Glines, 11:1-12, 12:24-13:2, 13:9-14, 14:17-15:14; *also id.*, 9:9-15, 16:15-17:2.

Each name server (filled circles, below) is “peered with a nearby content server [(filled rectangles, below)] ... that contains (or can serve) desired content file replicas via unicast addressing.” *Id.*, 6:9-22; *also id.*, 7:1-18.



Id., Fig. 1 (annotated); EX-1008 ¶¶86-87.

Glines discloses providing CDN DNS servers that are associated with a respective plurality of content servers, as claimed.

Glines “associate[es] multiple replica content servers 50-1, 50-2, ... 50-3, with the origin server 40” to serve content associated with the origin server. Glines, 10:13-18. Each replica content server includes a name server, including “name servers 53-1, 53-2, 53-3.” *Id.*, 10:13-25. The name servers 53-1, 53-2, 53-3 are “peered with content file servers 51-1, 51-2, 51-3 that contain replica copies of files stored at the origin server 40.” *Id.*; *also id.*, 12:11-23, 12:24-13:8, 6:9-22, 7:1-18. Glines discloses and depicts content servers 51-1, 51-2, and 51-3 as being addressable by unicast addresses “40.100.20.x,” “30.100.20.x,” and “20.100.20.x,”

respectively. *Id.*, Fig. 1; *also id.*, 6:8-22, 7:18-8:2, 13:9-13, 15:26-16:2. This discloses more than one machine associated with the content servers (e.g., “40.100” “30.100,” and “20.100.”), and at least suggests each name server 53-3 has a plurality of associated content servers. EX-1008 ¶¶88-92; EX-1012, 416-418; *also id.*, 11; Glines, 8:19-23 (“return[ing] the IP address of *one or many . . . content servers*”).

Further, Glines refers to each name server-content server group as a “service group” and discloses functionality that would be relevant only if more than one content server was in each group. Glines, 16:15-17:2, 17:20-19:16. For example, Glines discloses fallover or fallback policies to avoid issues with content server outages or overloads. *Id.*, 17:20-19:16. This includes “provid[ing] redundancy in the service group” and “return[ing] *two or more IP* address[es] in the service group” in response to the anycast message. *Id.*, 18:20-19:3. Load-balancing techniques “within the service group” are also described using “layer 4 switches.” *Id.*, 19:10-16. Glines further discloses “map[ping] the lookup of [a] name into a local IP address of a *member* of the service group,” *id.*, 16:28-17:2, which teaches more than one content server for each name server. EX-1008 ¶¶93-96.

Glines therefore teaches *each* name server in its network being paired with *multiple* content servers. *Id.*

To the extent Glines does not expressly disclose this feature (it does), it would have been obvious to implement and/or modify Glines in this manner, just as Glines suggests.

First, a POSITA would have understood that having multiple content servers associated with each name server would have improved performance, reliability, and scalability. EX-1008 ¶¶97-100. This would have increased storage capacity at each name server location, allowing more content to be stored locally and improving system performance. *Id.* Using multiple servers in this way was well known and would have been obvious to a POSITA. *Id.*; Chung at 1:16-40.

Second, employing multiple content servers would have enabled higher request volumes because each name server-content server group would have been able to load balance across multiple content servers. Glines discloses this technique expressly and a POSITA would have recognized its benefits. Glines, 19:13-15; EX-1008 ¶¶101-102.


Third, using multiple content servers would have enabled each group to better manage faults. If one content server fails, others in the group would continue to handle content requests. EX-1008 ¶¶103-104; Chung, 4:3-28.

A POSITA therefore would have been motivated to modify Glines for these reasons (e.g., to improve performance, load distribution, and fault management). EX-1008 ¶¶105-108. Doing so would have been nothing more than combining

known replication techniques in the prior art to yield predictable results. *Id.* This would have been well within the capabilities of a POSITA and the result would have been more than one content server for each name server, which was a common practice to increase storage, performance, and reliability. *Id.* A POSITA therefore would have had a reasonable expectation of success in the combination. *Id.*

It also would have been obvious to combine Glines and Gupta to arrive at this feature.

Gupta discloses multiple “point of presence” (POP) servers that each include a name server and WebCache servers. Gupta, 4:45-67, 5:16-53, 6:10-21, 10:4-22, Fig. 1. For example, Gupta’s POP servers 10 and 30 each include multiple **WebCache servers** (ovals) with an associated **name server** (rectangles). *Id.*



Id., Fig. 1 (annotated); EX-1008 ¶¶109-110.

It would have been obvious to combine and apply Gupta's teaching of multiple WebCache servers with each name server to Glines. Glines's content servers 51 may be or include "cache server[s]." Glines, 6:8-15, 8:4-6. Gupta's WebCache servers were therefore compatible with Glines. EX-1008 ¶¶111-114.

Combining Gupta's use of multiple WebCache servers with Glines's name servers would have been nothing more than combining known elements according to understood principles to yield predictable results. *Id.* ¶¶111-118. As discussed above, a POSITA would have understood that Glines's system would have benefited from multiple content/WebCache servers at each name server location by enabling load balancing, redundancy, and failure tolerance at these locations. *Id.* A POSITA thus would have been motivated to combine Gupta's multiple servers per name server in view of Glines's suggestion in order to achieve these same benefits. *Id.*; Glines, Fig. 1, 8:19-23, 18:20-19:3, 19:10-16; Gupta, 10:4-22, Fig. 1; *see also* Section VII.A. Given their similarities and overlap, a POSITA would have had a high expectation of success in combining Glines and Gupta. EX-1008 ¶¶117-118.


Combining Glines with Gupta in this way would not alter the features of Glines discussed elsewhere in this petition. *Id.*

3. [1b] (B) causing said plurality of CDN DNS servers to be authoritative for a hostname associated with a content provider by causing said common anycast address to be associated with said hostname;

Glines discloses these features. A client in Glines enters a URL into a browser that "will typically contain a domain name" (e.g., hostname). Glines, 9:16-10:4. Glines discloses "www.example.com" as an exemplary hostname associated with a content provider. *Id.*, 6:23-7:13, 13:15-15:10; *also* EX-1001, 3:24-54 (listing hostname as www.abcd.com); EX-1008 ¶¶119-124.

Glines teaches that each of its name servers 53 (the claimed CDN DNS servers) is “authoritative” for any domain name (hostname) for which the name server’s associated content servers have replica content. Glines, 13:9-14 (“[e]ach name server 53 is considered to be an authoritative”); *also id.*, 6:8-15, 12:24-13:2, 14:17-25. For example, where content servers include a replica of the “example.com” content, the associated name servers 53 are “defined as being authoritative for the ‘example.com’ domain.” *Id.*, 7:1-8; EX-1008 ¶¶125-127.

Glines further discloses that, to define the name servers 53 as authoritative, the anycast address is associated with the hostname. EX-1008 ¶¶128-133. When a user enters “www.example.com” into a browser, the local resolver 18 “consult[s] public name servers that have been defined as being authoritative for the ‘example.com’ domain.” Glines, 7:1-8. Rather than returning the address of the origin server associated with the “example.com” domain, the resolver “instead return[s] the anycast address 50.100.20.1,” which is associated with the name servers. *Id.*, 14:13-25; *also id.*, 7:1-8, 12:24-13:2, 13:9-11. Glines’s Figure 2, annotated below, illustrates this process:



Id. at Fig. 2 (annotated); EX-1008 ¶¶128-133.

4. [1c] (C) responsive to a request for content associated with the content provider and issued by a client, said request including at least said hostname,

Glines explains that a user operating client computer 12 “runs a web browser program that enables a user to submit requests for content files that are nominally located at an origin file server 40.” Glines, 9:9-22; *also id.*, 9:9-15, 11:14-12:10. The user’s request for a content file nominally located at an origin server is a request for content associated with the content provider issued by a client as claimed. EX-1008 ¶¶134-135.

Glines’s client issues the request in the form of a Uniform Resource Locator (URL) that “will typically contain a domain name [e.g., a hostname], a file name for the content file, and an access method.” Glines, 9:9-10:4, 13:15-14:25. “The client browser 12 then makes a request to discover the correct Internet Protocol (IP) address of the origin server 40 that contains the requested content file.” *Id.*, 9:9-10:4. For example, Glines discloses requesting content from www.example.com (a hostname) and specifying the URL “http://www.example/homepage.html.” *Id.*, 13:15-14:25. This discloses a request that includes a hostname. EX-1008 ¶136.

Glines also discloses that its system is “responsive” to the user’s request. For example, in response to the user’s request, the root DNS name server 38 is “programmed to . . . return the anycast address 50.100.20.1” of name servers 53-1, 53-2, 53-3 that “have been designated as being the authoritative name servers for ‘example.com.’” Glines, 14:4-25; *also infra* [1d]; EX-1008 ¶137. The name server then provides the browser with a content server IP address so that the user may retrieve the requested content. Glines, 7:18-8:2, *also infra* [1e]; EX-1008 ¶137-140.

5. [1d] (c1) causing said hostname to be resolved to said common anycast address by an Internet Service Provider (ISP) DNS server; and then

Glines, alone or combined with Gupta, discloses this limitation. EX-1008 ¶¶141-169.


For example, Glines “caus[es] said hostname to be resolved to said common anycast address” by returning an anycast address associated with DNS servers name servers 53-1, 53-2, 53-3 that “have been designated as being the authoritative name servers for ‘example.com.’” Glines at 13:21-15:4, 7:1-7. By returning an anycast address (e.g., 50.100.20.1) associated with name servers 53-1, 53-2, and 53-3 in response to a user entering a request into a browser (e.g., via URL or hostname) for a website (e.g., content), Glines discloses causing a hostname to be resolved to a common anycast address, as claimed. EX-1008 ¶¶141-146.

Glines alone discloses that the anycast address is resolved by an ISP DNS Server.

Glines discloses local DNS resolver 18, which is an ISP DNS server. Glines, 13:24-14:12. A POSITA would have understood that residential users’ Internet access is usually provided by an Internet service provider (ISP). EX-1008 ¶¶147-153; EX-1011, 3-4, 128; EX-1012, 229-230; EX-1013, 4:65-5:53. With an ISP, users typically access the internet through an ISP’s DNS server. EX-1014, 128-129, 338-339; *also* EX-1011, 307, 128 (“DHCP is used extensively in LANs and in residential Internet access.”). A POSITA would have accordingly understood that when Glines describes the browser of a “user associated with an Internet Service Provider” interacting with local DNS resolver 18, local DNS resolver 18 is an ISP DNS server, as claimed. EX-1008 ¶¶147-153. This is consistent with the ’112 patent,

which describes a “local DNS server 70a” or an “end user DNS 70a” but does not discuss a separate “ISP DNS server.” EX-1001, 4:56-5:31, 6:65-7:32, Fig. 5.

Glines further teaches that, when a user “associated with an Internet Service Provider” enters a URL (which includes a hostname, *supra* Limitation [1c]) in a web browser, local DNS resolver 18 sends a request to a root DNS server 38. Glines, 13:24-14:12; *also id.*, 11:24-28; EX-1008 ¶¶154-156. In turn, root DNS name server 38 is “programmed to ... return the anycast address 50.100.20.1” of name servers 53-1, 53-2, 53-3 that “have been designated as being the authoritative name servers for ‘example.com.’” Glines at 14:4-15:4; *also id.*, 7:1-7, 9:9-15. The browser or local resolver can then “send[] out a DNS request as a UDP datagram to [the] anycast message” to resolve the content. *Id.*, 7:9-13; *also id.*, 6:16-17.



Id., Fig. 1 (annotating local DNS resolver 18).

Glines's local DNS resolver, an ISP DNS server, thus causes the hostname entered as part of the URL to be resolved to the name servers' common anycast address, as claimed. EX-1008 ¶¶154-156.

To the extent Glines does not expressly disclose this feature (it does), it would have been obvious to implement and/or modify Glines's local resolver 18 to be an ISP DNS server.


As noted above, a POSITA would have understood that Glines's local DNS resolver 18/local resolver is an ISP DNS server, but even if not, it would have been obvious to substitute an ISP DNS server for Glines's local resolver 18. EX-1008 ¶¶157-163.

ISP-provided DNS servers are typically located in close proximity to a user's physical location, making them faster and more efficient than more distant DNS servers. *Id.*; EX-1011, 128; EX-1013, 5:35-54. Modifying Glines to implement local DNS resolver 18 as an ISP DNS server would accordingly have been nothing more than the simple substitution of known elements to yield predictable results, consistent with Glines's description of the DNS resolver as "local." Glines, 13:24-14:3; EX-1008 ¶¶159-161. A POSITA would have had an expectation of success in the modification because using local ISP DNS servers was known and understood in the art for accessing the Internet to resolve content in systems such as Glines. EX-1008 ¶¶162-163.

The modified Glines's ISP DNS server would perform the same functions as Glines ascribes to local DNS resolver 18, and would thus cause the hostname to be resolved to the common anycast address by an ISP DNS server, as claimed. *Id.* ¶163.

It also would have been obvious to combine Glines and Gupta to arrive at this feature.

Gupta discloses that “users 150 and 160 may be home users using dial-up lines, cable modems, DSL, or the like, to connect with **ISP 60**.” Gupta, 6:1-9, Fig. 1; *also id.*, 4:45-53. Gupta further discloses that “client DNS server 170 is typically provided within **ISP 60**.” *Id.*, 6:1-9. And Gupta further discloses that “client DNS server 170 is typically used to resolve a host name into an IP address.” *Id.*; *also id.*, Fig. 1.



Id., Fig 1. (cropped and annotating ISP 60).

Gupta's ISP 60 thus performs the same functions as Glines's local DNS resolver 18 (e.g., resolving a host name into an IP address). In view of Glines's teaching that local DNS resolver 18 is used "where the user is associated with an Internet Service Provider (ISP)," Glines, 13:24-27, a POSITA would have understood Gupta's ISP 60 to be an obvious substitute for Glines's local resolver 18. EX-1008 ¶¶164-169. A POSITA would have been motivated to make the modification to ensure compatibility with "typical" Internet service architecture, e.g., Gupta, 4:45-51, 5:65-6:9, and to achieve the speed and efficiency associated with ISP-provided DNS servers, EX-1011, 128; EX-1013, 5:35-54; EX-1008

¶¶164-169. Given the ubiquity of ISP-provided DNS servers and the disclosures of Gupta and Glines, a POSITA would have viewed the combination as nothing more than the simple substitution of known elements according to understood principles to yield predictable results. EX-1008 ¶169. A POSITA also would have had an expectation of success in the substitution. *Id.*


Combining Glines with Gupta in this way would not alter the features of Glines discussed elsewhere in this petition. *Id.*

6. **[1e] (c2) by one of said plurality of CDN DNS servers, resolving said hostname to identify an IP address for use by the client to retrieve the content from a content server.**

Glines discloses that, after receiving the common anycast address, the user sends a request to the anycast address associated with the name servers. Glines, 15:6-16:2. The closest name server 53, one of the plurality of CDN DNS servers, “will then respond by reporting the unique IP address of its associated replica content server 51,” thus resolving the hostname to identify an IP address for use by the client to retrieve the content from a content server as claimed. *Id.*; *also id.*, 11:1-12; EX-1008 ¶¶170-172.

For example, the user may send a “DNS request for the IP address of “www.example.com” to the anycast address of the name servers. Glines, 15:15-25. Where name server 53-3 is the “closest in terms of network hops,” *id.*, 15:15-25, it responds and returns the IP address of its associated replica content server 51-2, e.g.,

unicast address 30.100.20.x, *id.*, Fig. 1. The “x” in this IP address identifies the host address of a machine within content server 51-2, allowing the client to use that address to receive the requested content. *Id.*, 15:26-16:8; EX-1008 ¶¶173-176.



Glines at Fig. 1 (annotated).

C. Claim 2: “some ... CDN DNS servers ... located at ... (POPs)”


Glines, alone or combined with Gupta, would have rendered this limitation obvious.

Glines alone teaches this limitation.

Glines discloses name servers 53 and content servers 51 that are physically co-located and share access points to the internet network 30. Glines, 13:3-8, 8:19-23,

6:8-15, 7:6-7, 10:23-25, 13:3-9, 16:6-8, Figs. 1, 2. Glines explains that the name and content servers may share “the same physical machine, if desired.” *Id.*, 8:11-13.

Glines also discloses that each of the peered name and content servers may access Internetwork 30 (e.g., Internet) via “internetwork devices 16.” *Id.*, 11:14-22, Fig. 1; *also id.*, Abstract, 6:3-7, 8:21-23, 9:9-12, 12:24-13:2, 15:11-16:8, Fig. 2 (Step 116). Glines’s internetwork devices 16 may include a “modem, network interface card, router, switch, bridge, gateway, or the like.” *Id.*, 11:14-22.



Id. at Fig. 1 (annotated).

Glines collocating its name server 53 and content server 51 service groups and causing them to share the same entry into the greater network (e.g., Internetwork 30), discloses at least some of the CDN DNS servers (and in fact, all of the CDN

DNS servers) being located at network Points of Presence (POPs), as claimed. EX-1008 ¶¶177-183. As discussed, a “POP” is a point or physical location where two networks connect using, for example, a router. *Id.*; EX-1013 at 4:65-5:54.

Glines’s disclosure corresponds with the ’112 patent, which discloses nodes 76 comprised of DNS and content servers “located so as to be as close as possible to various network entry points.” EX-1001, 5:11-18, Fig. 5.

Alternatively, it also would have been obvious to combine Glines and Gupta to arrive at this feature.

Gupta expressly discloses “a series of ‘point of presence’ (POP) servers,” including POP servers 10, 20, and 30. Gupta, 4:45-67. These “are typically coupled via a computer network 70, such as the Internet.” *Id.* Each POP server “include[s] a plurality of dedicated servers,” including for example both WebCache servers (content servers) and “server[s] ... that include[] Domain Name Server (DNS) software,” (CDN DNS servers). *Id.*




FIG. 1

Id. at Fig. 1 (annotated); EX-1008 ¶¶184-185.

A POSITA would have combined Gupta and Glines to implement and/or modify Glines's name server and content server groups to be located at network POPs. EX-1008 ¶¶186-188. A POSITA would have been motivated to do so to distribute Glines's service groups to different access points on the network, including different geographic and topological locations. *Id.* By distributing content servers at different POPs, content and services would be delivered from locations closer to users, thereby reducing latency. *Id.*; also Gupta 2:46-3:25. By localizing content delivery through, for example regional POPs located closer to many users, Glines's

system would be improved by reducing congestion and decreasing bandwidth costs associated with long-distance data transmission. EX-1008 ¶¶186-188; EX-1011, 126-127.

Combining Glines with Gupta in this way would not alter the features of Glines discussed elsewhere in this petition. EX-1008 ¶188.

D. Claim 3: “each ... CDN DNS server[] ... located at ... (POPs)”

Glines, alone or combined with Gupta, would have rendered this limitation obvious.

As explained in claim 2, Glines, alone or combined with Gupta, renders obvious placing “at least some of said plurality of CDN DNS servers” at POPs. *Supra* claim 2. For the same reasons, Glines, alone or combined with Gupta, renders obvious placing each of the CDN DNS servers at a network POP. EX-1008 ¶¶189-192. Glines’s disclosures expressly apply to “each” name server and content server 51. *E.g.*, Glines, 13:3-8, Fig. 1; *also id.*, 8:19-23, 6:8-15, 7:6-7, 10:23-25, 13:3-9, 16:6-8, Fig. 2. To the extent Glines alone does not disclose this limitation, the benefits described in claim 2 of modifying Glines to position the servers at POPs, as expressly disclosed by Gupta, *e.g.*, Gupta, 4:45-67, Fig. 1, would be heightened by placing each name server at a POP, *supra* Claim 2; EX-1008 ¶¶189-192.

E. Claim 4: “grouped as a cluster”

Glines, alone or combined with Gupta, would have rendered this limitation obvious.

As discussed in claim 1, Glines, alone or combined with Gupta, discloses that each name server 53 (CDN DNS server) is associated with a plurality of content servers. *Supra* Limitation [1a].

Glines further discloses that associated name and content servers should be arranged together and may even share a “cluster address.” Glines, 13:3-8, 6:8-15, 7:6-7, 7:14-24, 8:14-23, 10:23-25, 13:3-9, 16:6-8, Figs. 1, 2, 17:15-19:16. A POSITA therefore would have understood that Glines discloses its name servers (the claimed CDN DNS servers) and associated content servers being grouped as a cluster, and a POSITA would have implemented Glines, whether or not modified by Gupta, in this way. EX-1008 ¶¶193-197.

Cluster computing was well known in the art. EX-1008 ¶¶198-201; EX-1015, 9-11. A POSITA would have grouped Glines’s name and associated content servers as a cluster based on Glines’s teaching and suggestion to do so, and because of the benefits of cluster computing, including improved performance, scalability, throughput, and high availability over existing systems. EX-1008 ¶¶198-201; EX-1015, 9-11, EX-1011, 11, 102-103; EX-1023, 38, 39. Indeed, Glines discloses using Layer 4 switches to select a content server and to perform load balancing,

which was a common practice in cluster computing. EX-1008 ¶201; EX-1023, 39, 40, Glines, 19:13-16.

Implementing and/or modifying Glines's name server/content server pairs to be grouped as clusters, as Glines suggests, Glines, 7:14-24, 19:13-16, to the extent not expressly disclosed, would have been nothing more than combining prior art elements to yield predictable results, EX-1008 ¶¶202-204. This would have improved scalability, fault tolerance, and load-balancing in Glines. *Id.*; EX-1015, 9-10; EX-1023, 38.

F. Claim 5: “each ... content server[] is associated with a ... (POP)”

Glines, alone or combined with Gupta, would have rendered this limitation obvious. EX-1008 ¶¶205-209.

Glines, with or without Gupta, teaches each CDN DNS server being located at network Points of Presence (POP). *Supra* Claim 3.

Further, Glines, with or without Gupta, teaches that each name server (CDN DNS server) is clustered with its plurality of content servers, and a POSITA would have been motivated to group Glines's name and associated content servers as a cluster. *Supra* Claim 4.


For the same reasons, Glines, alone or combined with Gupta, discloses that each content server is associated with a network POP. Each of the content servers would be associated with and clustered with a name server, *supra* claim 4, and each

of those name servers would be associated with a POP, *supra* claim 3, such that each of the content servers would also be associated with a POP as claimed.

G. Claim 6: “attempting to serve the content”

Glines, alone or combined with Gupta, would have rendered this limitation obvious.

Glines explains that after receiving the IP address of the content server, *see* Limitation [1e], the user’s browser makes an HTTP level request for the content (e.g., “homepage.html”) using “standard TCP/IP and HTTP network protocols” via the returned “unicast address for the content server 51.” Glines, 15:21-16:8; *also id.*, 11:1-12. In response, “[t]he requested file is then returned from the content server 51-2.” *Id.* at 16:3-8.



Id. at Fig. 2 (cropped and annotated).


Because the content server returns the requested file, Glines discloses the content server attempting to serve the content. EX-1008 ¶¶210-214. To successfully serve content, the content server must “attempt” to serve it. *Id.* This is consistent with the ’112 patent, which “serves the object to the end-user system that requested it.” EX-1001, 7:24-32. The specification does not otherwise use the term “attempt.” This is also consistent with Patent Owner’s allegations. EX-1022, 18-20.

H. Claim 7: “if a valid version of the content is available ... serving .. otherwise obtaining”

Glines, alone or combined with Gupta, would have rendered this limitation obvious.

Glines, alone or combined with Gupta, discloses claim 6. *Supra* Claim 6.

Glines further discloses that “replicas of content files that originate at the origin server 40 [(red rectangle)] are distributed [(yellow lines)] and stored in the replica content servers 50.” Glines, 12:11-23.



Id., Fig. 1 (annotated).

A POSITA would have understood that, for Glines’s content servers 50-1, 50-2, 50-3 to serve the requested content, the system would have confirmed the content is stored or otherwise available and would have understood the benefits of obtaining unavailable content at a user’s request, including distributing content across Glines’s system to fulfill current and future requests for content. EX-1008 ¶¶215-222. Accordingly, a POSITA would have found it obvious to implement Glines in this manner. *Id.*

Gupta also discloses these features. Gupta, 5:16-40, 13:23-30, Fig. 4. Gupta’s WebCache servers (content servers), like Glines’s content servers, “respond[] to requests for static web content from clients.” *Id.*, 5:16-40. Gupta teaches, however, first determining “[i]f the requested content does not exist” and thus determining if

a valid version is available on the content server or not. *Id.* Where available, the server responds to the request “from memory,” serving the content. *Id.*

If the requested content is not cached, Gupta teaches that its WebCache server software “will generate a request to a WebServer, that is known to include the content,” which provides the content to the WebCache, which in turn “returns the static content to the user.” *Id.*, 10:40-11:7; *also id.*, 5:16-40, 13:23-30, Fig. 4. Gupta thus teaches obtaining unavailable content from a content source and then serving the content, as claimed. This is consistent with Patent Owner’s allegations, which include “serv[ing] content from the content server (such as cached sources) *if available* on the content server.” EX-1022, 20-21.

A POSITA would have implemented Gupta’s teachings in Glines to facilitate content distribution across Glines’s network and to improve future service of that content at the closest available content server. EX-1008 ¶¶223-231. Doing so would have been nothing more than combining known elements to yield predictable results. *Id.* A POSITA would have had an expectation of success because Gupta and Glines disclose similar content caching techniques and it is foundational to first check that content is available when retrieving and/or providing it. *Id.*

A POSITA would have implemented Gupta’s content-checking and cache miss signal in Glines either before a name server responds to the anycast address or after Glines’s name server provides the IP address of a content server. *Id.* ¶232. This

includes requesting the content from origin server 40 or from another content server. *Id.* In addition to fulfilling a user's request, this would have added otherwise missing content to that content server, thereby distributing the content across Glines's system to efficiently fulfill future requests. *Id.* It would have reduced latency and improved the usefulness of Glines. *Id.* Requesting the content from the source and delivering it via the content server would have avoided the user having to initiate another request for the content. *Id.*

I. Claim 8: “origin server associated with said content provider”

Glines, alone or combined with Gupta, would have rendered this limitation obvious.

Glines, alone or combined with Gupta, discloses claim 7. *Supra* claim 7. As discussed in claim 7, Glines discloses a content source that includes origin file server 40. *Id.*; Glines, 2:9-13, 9:16-10:25, 12:11-23, 14:13-25. The origin server is associated with the content provider—Glines's content provider “example” provides content from “www.example.com.” Glines, 13:15-14:25; EX-1008 ¶¶233-236.

Gupta similarly discloses customer web server 130, which is the content source from where unavailable content is obtained. *Supra* claim 7. Gupta teaches that the “customer” in Gupta is the content provider, and a POSITA would accordingly understand the customer web server 130 to be an origin server as claimed. Gupta, 5:23-33, 2:10-17, 8:53-59, Fig. 1; EX-1008 ¶¶237-239. As

discussed in claim 7, a POSITA would have combined Glines and Gupta, and the combination would have rendered obvious an origin server in Glines's system that is "associated with said content provider," as claimed. *Supra* claim 7; EX-1008 ¶¶237-239.

J. Claim 10: "load balancing policy"

Glines, alone or combined with Gupta, would have rendered this limitation obvious.

Glines discloses each name server and content server "service group" may include fallover/fallback policies to avoid issues with content server outages or overloads. Glines, 18:20-19:16. These policies include, "within the service group, standard load balancing" using "layer 4 switches," which is selection of a content server based on a load balancing policy. *Id.*, 19:10-16; EX-1008 ¶¶240-244; EX-1016, 1:28-53; EX-1017, 1:19-48.

Glines further teaches that when the "local load reache[s] some watermark, it can load shed to content servers in other service groups by returning their IP address to name lookups coming at the anycast address," which is also selection of a content server based on a load balancing policy. Glines, 19:17-22; EX-1008 ¶¶245-246.

K. Claim 11: "switch, selecting the content server"

Glines, alone or combined with Gupta, would have rendered this limitation obvious.

As explained in claim 10, Glines describes a Layer 4 switch used for “standard load balancing” among multiple content servers to apply a load balancing policy and select a content server. *Supra* Claim 10. For the same reasons, Glines discloses selecting the content server by at least one switch—Glines’s switch is used for the load balancing that determines which content server should be selected. EX-1008 ¶¶247-253.

L. Claim 12: “predetermined selection policy”

Glines, alone or combined with Gupta, would have rendered this limitation obvious.

As explained in claims 10 and 11, Glines discloses that its Layer 4 switch is used for “standard load balancing” across content servers. *Supra* Claims 10, 11. Glines accordingly describes selection of a content server by a switch in accordance with a predetermined selection policy (here, a load balancing policy). EX-1008 ¶¶254-257. This corresponds with the ’112 patent, which discloses “a predetermined selection policy, for example, . . . , load balancing.” EX-1001, 8:52-58.

M. Claim 13

Glines, alone or combined with Gupta, would have rendered claim 13 obvious.

1. [13p]: “plurality of content providers”

As described above, Glines discloses a method of content delivery for delivering content. Limitation [1p].

Glines further discloses a method of content delivery for delivering content on behalf of a plurality of content providers because Glines expressly discloses that its system may be used with either a “single provider” or “multiple providers,” and thus discloses content delivery on behalf of a *plurality* of content providers, as claimed. Glines, 17:3-9; EX-1008 ¶¶258-261.

A POSITA would have understood Glines’s disclosure of providing content from an origin source would have included supporting more than one origin source (e.g., more than one website or content provider). Glines, 12:11-22, 13:21-14:3; EX-1008 ¶262. Adding origin servers and content providers to Glines’s system, as expressly contemplated in Glines, would have been nothing more than replicating Glines’s teaching of storing and providing content associated with one origin server, with expected results. EX-1008 ¶263; *B/E Aerospace, Inc. v. C&D Zodiac, Inc.*, 962 F.3d 1373, 1380 (Fed. Cir. 2020) (“repetition of an existing element” is obvious). Doing so would have enabled Glines’s system to support more (and different) content to fulfill user requests. EX-1008 ¶263.

Gupta likewise discloses content delivery on behalf of a plurality of content providers, and for this additional reason, it would have been obvious to provide content delivery on behalf of a plurality of content providers in a combined Glines-Gupta system. *Id.* ¶264-265. For example, Gupta teaches that its POP server network provides “traffic management, as well as content hosting to more than one

specific customer,” including “customer1 and customer2” that share the server network.” Gupta, 7:8-20, Fig. 1 (depicting customer1 and customer2 web servers); *also id.*, 8:53-59 (“customer1.com.”); EX-1008 ¶¶264-266.

It would have been obvious to apply these teachings to Glines’s system. EX-1008 ¶267. The combined system would have stored and provided content from multiple content providers (e.g., multiple websites, etc.). *Id.* This would have been nothing more than combining known elements in the prior art according to understood principles, with an expectation of success. *Id.* Doing so would have been expected and well within the capabilities of a POSITA to, for example, maximize the usefulness and value of Glines’s system by providing Glines’s service to more than one customer/content provider. *Id.*

2. [13a]-[13f]

The following limitations are substantively identical to limitations addressed above and would have been obvious for the same reasons. EX-1008 ¶¶268-289.

<u>Claim Element</u>	<u>Corresponding Disclosure</u>
[13a]	[1a] (§VII.B.2)
[13b]	[1b] (§VII.B.3)
[13c]	[1c] (§VII.B.4)

[13d]	[1e] (§VII.B.6)
[13e]	Claim 11 (§VII.K)
[13f]	Claim 6 (§VII.G)

3. [13g]: “responsive to a second request”

Limitation [13g] differs from Limitation [13c] in that it recites a “second” request for the first content. Glines, with or without Gupta, would have rendered this limitation obvious.

Glines discloses providing content distribution “as a service where client *requests* are automatically routed to the closest available content server.” Glines 6:3-15; *id.*, Abstract, 9:9-17 (“requests” for content), 10:5-12 (handling “requests”). Similar or repeated processing would occur to fulfill those “requests,” whether for the same or different content between requests. EX-1008 ¶¶290-295. As discussed, repeating a process—for the same or different content—is an obvious and expected implementation of Glines’s system. *See* Limitation [13p]. This includes Glines providing a user content at different times or to different users. EX-1008 ¶¶294-295.

4. [13h]: “by a second DNS server ... identify[ing] a second IP address for retrieval of the first content from a content server associated with the second DNS server”

Limitation [13h] differs from Limitation [13d] in that a “second” DNS server resolves the first hostname to identify a “second” IP address for retrieval of the first content.

Dependent claim 15 recites that the first DNS server and second DNS server are the same server. EX-1001, claim 15. Dependent claim 17 further recites that the second content server (and thus the “second IP address”) and the first content server are the same server. EX-1001, claim 17.

Element [13h] therefore encompasses repeating a DNS resolution for the first content based on a second user request.

Glines, with or without Gupta, would have rendered this limitation obvious for the same reasons provided for Limitation [13d] (citing [1e], and [13p]).

As explained, it would have been obvious for Glines to manage more than one request for content, including multiple requests for the same or different content. Limitations [13p], [13g]. Glines’s resolution process would have been handled in the same way for each of the requests, and Glines, with or without Gupta, would have rendered this limitation obvious as explained in Limitations [1e] and [13d]. EX-1008 ¶¶296-302.

If the first and “second” DNS servers are different servers, it would also have been obvious to resolve a second request for particular content via a different name server in Glines (e.g., via a different DNS server 52 and different associated content server 51). Glines, 14:17-16:8, Fig. 1. For example, if a first name server is overloaded or offline, or if a second name server is available that is “closer” to the user sending the second request, it would have been obvious for a second name server in Glines’s network 10 to resolve the content. *Id.*, 14:17-16:8, 18:13-19:20; EX-1008 ¶¶303-305. The second name server in Glines would share the same anycast address as the first name server and respond by identifying its associated content server 51, and thus identify a “second IP address” that is different from the “first IP address” provided in response to the first request. EX-1008 ¶306. Indeed, Glines contemplates this with respect to its system in Figure 1. *See* Glines, 14:17-16:8, Fig. 1.

5. [13i]: “by a switch”

Limitation [13i] differs from Limitation [13e] in that a switch associated with the “second” DNS server selects the “second” content server.

Dependent claim 15 recites that the first DNS server and second DNS server are the same server. EX-1001, claim 15. Dependent claim 17 further recites that the second content server (and thus the “second IP address”) and the first content server are the same server. EX-1001, claim 17. Dependent claim 16 additionally recites the

switch of (c22), limitation [13e], and the switch of (c12), limitation [13i], are the same switch.

Element [13i] therefore encompasses a DNS resolution of the second request using the same switch.

Glines, with or without Gupta, would have rendered this limitation obvious for the reasons explained in Claim 11 and Limitation [1e]. As discussed, it would have been obvious to manage more than one request in Glines, including managing multiple requests for the same or different content using different DNS servers. *See* Limitation [13h]; EX-1008 ¶¶306-314.

Even if the first and second DNS server are understood to be different servers, this feature would have been obvious over Glines and/or Gupta. EX-1008 ¶¶315-316. Glines and Gupta both disclose DNS server-content server pairs with similar configurations. *See* Glines, 12:11-13:14, 14:17-16:8, Fig. 1; Gupta, 4:45-67, 5:41-6:21, Fig. 1. Limitation [13i] would have obvious in either Glines or Glines-Gupta because it would have involved mere replication of steps (e.g., replication of limitation [13e]) in a different CDN node (e.g., different DNS server-content server pair). EX-1008 ¶¶315-316.

6. [13j]: “attempting to serve”

Whether [13j] is understood to be attempting to serve the same or different content in response to the second request, Glines, alone or combined with Gupta,

would have rendered this limitation obvious for the reasons explained in claim 6.¹
Supra Claim 6; EX-1008 ¶¶317-319.

It would have been obvious to manage more than one request for content, Limitation [13h]; EX-1008 ¶318, and a POSITA would have understood that each request would be resolved in the same way such that, in response to a second request, the system would attempt to serve the requested content from the identified content server in the same manner described in claim 6. *Supra* Claim 6. As discussed in [13g], Glines discloses serving client “requests” and similar or repeated processing would occur to fulfill those requests, including serving the same content or different content at different times. This would have been an obvious and expected implementation of Glines’s system. EX-1008 ¶¶318-319.

¹ Petitioner’s alternative indefiniteness assertion in district court does not preclude review here. *Zillow Grp., Inc. v. Int’l Bus. Machs. Corp.*, IPR2020-01656, Paper 8 at 11 (P.T.A.B. Mar. 15, 2021); *Target Corp. v. Proxicom Wireless, LLC*, IPR2020-00931, Paper 10 at 9 (P.T.A.B. Nov. 10, 2020) (noting “[a]lternative pleading before a court is common practice”); *Alcon Inc. v. AMO Dev., LLC*, IPR2021-00843, Paper 15 at 16 (P.T.A.B. Nov. 12, 2021) (stating Board is “able to construe the disputed claim term with sufficient precision to address the obviousness arguments”).

N. Claim 14: “load balancing policy”

Glines, alone or combined with Gupta, would have rendered this limitation obvious. EX-1008 ¶¶320-322.

As explained, in connection with claims 10 and 11, Glines, alone or combined with Gupta, would have rendered obvious selection of a content server based on a load balancing policy. *Supra* Claims 10, 11.

Further, it would have been obvious for Glines to manage more than one request for content, including multiple requests for the same or different content, and content server selection would proceed in the same way for each request. *See* Limitations [13g], [13h]. Repeating the same process twice, for example selecting a content server by a load balancing policy across the same or different service groups would have been an obvious and expected implementation of Glines, and this limitation would have been obvious for the reasons explained in claims 10 and 11. EX-1008 ¶¶320-322.

O. Claims 15-17

These limitations are substantively identical to limitations addressed above and would have been obvious for the same reasons. EX-1008 ¶¶323-331.

<u>Claim</u>	<u>Corresponding Disclosure</u>
15	[13h] (§VII.M.4)

16	[13i] (§VII.M.5)
17	[13h] (§VII.M.4)

P. Claim 18

1. [18p]-[18c]

These limitations are substantively identical to limitations addressed above and would have been obvious for the same reasons. EX-1008 ¶¶332-342.

<u>Claim Element</u>	<u>Corresponding Disclosure</u>
[18p]	[13p] (§VII.M.1)
[18a]	[13a] (§VII.M.2)
[18b]	[13b] (§VII.M.2)
[18c]	[13a] (§VII.M.2)

2. [18d]: “switch associated with each set of content servers”

For the reasons explained in claim 11 and limitations [13e] and [13i], Glines, alone or combined with Gupta, would have rendered this limitation obvious.

Each of claim 11 and limitations [13e] and [13i] require a switch used to select among a set of content servers. Glines, alone or combined with Gupta, thus would

have rendered obvious at least one switch associated with each set of content servers as claimed. *Supra* Claim 11, Limitations [13e], [13i]; EX-1008 ¶¶343-347.

3. [18e]: “responsive to a request”

For the reasons explained in Limitations [13c] and [13d], Glines, alone or combined with Gupta, would have rendered this limitation obvious.

Each of Limitations [13c] and [13d] require one of the DNS servers to resolve the hostname by identifying an IP address in response to a user’s request. Glines, alone or combined with Gupta, accordingly would have rendered obvious Limitation [18e]. *Supra* Limitations [13c], [13d]; EX-1008 ¶¶348-353.

4. [18f]: “switch ... operable to perform load balancing”

For the reasons explained in Claims 10, 11 and Limitation [13e], Glines, alone or combined with Gupta, would have rendered this limitation obvious.

Each of Claims 10, 11 and Limitation [13e] require selection of a content server based on load balancing and doing so by a switch. Glines, alone or combined with Gupta, accordingly would have rendered obvious Limitation [18f]. *Supra* Claims 10, 11 and Limitation [13e]; EX-1008 ¶¶354-360.

5. [18g]: “serve the content”

This limitation is substantively identical to claim 6 addressed above and would have been obvious for the same reasons. *Supra* Claim 6; EX-1008 ¶361.

Q. Claim 19

1. [19p]-[19h]

These limitations are substantively identical to limitations addressed above and would have been obvious for the same reasons. EX-1008 ¶¶362-386.

<u>Claim Element</u>	<u>Corresponding Disclosure</u>
[19p]	[1p] (§VII.B.1)
[19a]	[1a] (§VII.B.2)
[19b]	Claim 2 (§VII.C)
[19c]	[1b] ² (§VII.B.3)
[19d]	[1c] (§VII.B.4)
[19e]	[1e] (§VII.B.6)
[19f]	Claim 11 (§VII.K)
[19g]	Claim 6 (§VII.G)

² The prior art teaches and renders obvious “said request including at least said” in Limitation [19c] for the same reasons it does so for “at least said hostname” in Limitation [1b]. EX-1008 ¶¶373-376.

[19h]	Claim 7 (§VII.H)
-------	------------------

R. Claim 20: “load balancing policy”

For the reasons explained in claims 10, 12, and 14, Glines, alone or combined with Gupta, would have rendered this limitation obvious.

Each of claims 10, 12, and 14 require selection of a content server based on a load balancing policy. Glines, alone or combined with Gupta, accordingly would have rendered obvious selection of the content server based on a load balancing policy as claimed. *Supra* Claims 10, 12, and 14; EX-1008 ¶¶387-388.

S. Claim 21

1. [21p]-[21e]

These limitations are substantively identical to limitations addressed above and would have been obvious for the same reasons. EX-1008 ¶¶389-404.

<u>Claim Element</u>	<u>Corresponding Disclosure</u>
[21p]	[1p] (§VII.B.1)
[21a]	[1a] (§VII.B.2)
[21b]	[1b] (§VII.B.3)
[21c]	[1c] (§VII.B.4)

[21d]	[1e] (§VII.B.6)
[21e]	Claim 11 (§VII.K)

T. Claim 22

1. [22p]: “plurality of nodes”

For the reasons explained in Limitations [1p] and [1a], Glines, alone or combined with Gupta, would have rendered this limitation obvious.

The ’112 patent discloses that a “node 76” includes a DNS server 78 and cache server 80. EX-1001, 6:65-32, Fig. 5. Each of Glines’s server groups, which likewise include name and content servers, therefore discloses a “node.” Glines, 10:13-18, 12:11-23, 13:3-7; EX-1008 ¶¶405-409.

Limitations [1p] and [1b] also require a plurality of DNS servers and associated content servers, which discloses a “plurality of nodes.” EX-1008 ¶408. Therefore, Glines, alone or combined with Gupta, would have rendered obvious selection a plurality of nodes, as claimed. *Supra* Limitations [1p] and [1b].

2. [22a]: “Domain Name System (DNS) server sharing an anycast”

This limitation is substantively identical to Limitation [1a] addressed above and would have been obvious for the same reasons. *Supra* §VII.B.2; EX-1008 ¶¶410-412.

3. [22b]: “plurality of content providers”

As discussed in Limitation [13p], Glines expressly discloses that its system may be used with “multiple providers” and adding additional content providers would have been obvious. *Supra* Limitation [13p]. Gupta likewise discloses multiple content providers, and it would have been obvious to apply these teachings to Glines’s system. *Supra* Limitation [13p]; EX-1008 ¶¶413-416.

Limitation [13p] requires delivering content on behalf of a plurality of content providers. Glines, alone or combined with Gupta, accordingly would have rendered obvious Limitation [22b]. *Supra* Limitations [13p]; EX-1008 ¶¶413-416.

4. [22c]-[22h]

These limitations are substantively identical to limitations addressed above and would have been obvious for the same reasons. EX-1008 ¶¶417-436.

<u>Claim Element</u>	<u>Corresponding Disclosure</u>
[22c]	[1a] (§VII.B.2)
[22d]	Claim 11 (§VII.K)
[22e]	[13c] (§VII.M.2)
[22f]	[13d] (§VII.M.2)
[22g]	Claim 11 (§VII.K)

[22h]	[13f] (§VII.M.2)
-------	------------------

U. Claim 23

1. [23p]-[23b]

These limitations are substantively identical to limitations addressed above and would have been obvious for the same reasons. EX-1008 ¶¶437-445.

<u>Claim Element</u>	<u>Corresponding Disclosure</u>
[23p]	[13p] (§VII.M.1)
[23a]	[13a] (§VII.M.2)
[23b]	[13b] (§VII.M.2)

2. [23c]: “responsive to a first request”

Glines, alone or combined with Gupta, would have rendered this limitation obvious as explained in Limitation [13c]. *See* Section VII.M.2.

Glines, alone or combined with Gupta, additionally renders Limitation [23c] obvious based on Internet content as of the priority date. For example, it would have been obvious for Glines’s name servers to respond to an anycast message based on a request for a news aggregator website, such as Yahoo! News (e.g., <http://dailynews.yahoo.com>), as of this date. EX-1024; EX-1008 ¶¶446-448.

News aggregator websites, such as Yahoo! News, were known in the art. EX-1008 ¶449. In EX-1024, for example, the first story is associated with the URL http://dailynews.yahoo.com/h/nm/20010201/ts/congress_ashcroft_dc_71.html, where the hostname is “http://dailynews.yahoo.com,” and the path is “h/nm/20010201/ts/congress_ashcroft_dc_71.html.” *See* EX-1024; EX-1025. The hostname is associated with Yahoo! News, while the path is associated with news content from Reuters. *See id.*; EX-1008 ¶449.

Applying Glines, alone or combined with Gupta, to respond to a first request for the URL <http://dailynews.yahoo.com> would have been an obvious application of Glines’s system. EX-1008 ¶¶449-450. The first content would have been the Yahoo! News webpage (e.g., a first content provider of said plurality of content providers), and the first request would have included a first hostname associated with the first content provider (e.g. <http://dailynews.yahoo.com>). *Id.*

This would have been nothing more than applying the art in understood ways to resolve particular content and combining known prior art techniques with predictable results. EX-1008 ¶¶451-453; Glines, 1:4-20, 9:9-21. This would have required no modifications of Glines or Gupta to implement, EX-1008 ¶¶451-453, and is indistinguishable from Glines’s request for content via the URL “<http://www.example.com/homepage.html>.” *Id.*; Glines, 6:23-7:24, 13:15-15:10.

3. [23d]: “resolving said first hostname”

Limitation [23d] is substantively identical to Limitation [13d] addressed above and would have been obvious for the same reasons. *Supra* §VII.M.2.

It would have been obvious to resolve the first request discussed in Limitation [23c] using Glines’s content delivery network. EX-1008 ¶¶454-458. This would have been nothing more than applying the art in understood ways for resolving particular content, and combining known prior art resolve techniques with known content available via the Internet, with predictable results. *Id.*

4. [23e]: “responsive to a second request”

It would have been obvious for Glines’s name servers, alone or combined with Gupta, to respond to an anycast message based on a second request that includes a request for a particular news article via a news aggregator website, such as Yahoo! News. *See* EX-1024; *see* Limitation [23c]; EX-1008 ¶¶459-461. This includes, as shown in EX-1025, a request for the article “U.S. Senate Approves Ashcroft, 58-42” at http://dailynews.yahoo.com/h/nm/20010201/ts/congress_ashcroft_dc_71.html, where the hostname is “http://dailynews.yahoo.com,” and the second content is a news article. *See* EX-1025. The hostname is associated with the first content provider (Yahoo! News), and the second content is associated with a second content provider (e.g., Reuters). *See id.*; EX-1008 ¶461.

It would have been obvious for Glines, with or without Gupta, to respond to a first request to Yahoo! News or similar new aggregator, as discussed, and a particular article via the aggregator. *See* EX-1024; EX-1025. Both requests would include the first hostname associated with the first content provider (e.g., <http://dailynews.yahoo.com> associated with Yahoo! News). EX-1008 ¶462.

Repeating Glines's resolution process for the same or different content was an obvious and expected implementation of Glines. *See* Limitation [13g]. This includes providing a user content at different times, or content to different users. Indeed, Glines discloses providing content to "client requests." Glines, 6:3-15; *id.*, Abstract, 9:9-17, 10:5-12. Similar or repeated processing would occur to fulfill those "requests." EX-1008 ¶463.

Resolving a second request such as this would have been nothing more than applying the art in understood ways to resolve particular content and combining known prior art techniques with predictable results. *Id.* ¶464. This would have involved combining known content available via the Internet with Glines's content delivery techniques (with or without Gupta), which resolves and delivers Internet content. *Id.*; Glines, 1:4-20, 9:9-21. This would have required no modifications of Glines or Gupta to implement. *Id.* ¶¶464-465. Resolving this particular content request is indistinguishable from Glines's exemplary request for content via "<http://www.example.com/homepage.html>." *See* Glines, 6:23-7:24, 13:15-15:10.


5. [23f]: “by a second DNS server”

Limitation [23f] is substantively identical to Limitation [13h] addressed above and would have been obvious for the same reasons. *Supra* §VII.M.4; EX-1008 ¶¶466-474.

VIII. Grounds 2A/2B: Glines, with or without Gupta, and Chung

A. Motivation to Combine

Chung discloses “distributing HTTP client requests to a cluster of servers,” including “publicizing the address ... of the server-side router 32 [associated with a cluster of servers] to the clients via the DNS.” EX-1004 (Chung), 2:33-3:10, Fig. 2. A client request “is directed to the router 32, which dispatches the request to a selected server 14-k of server cluster 34 based on load characteristics.” *Id.* The selected server then responds, “as indicated by the dashed line connecting server 14-k to client 12,” making the transaction “totally transparent.” *Id.*



Id., Fig. 2.

A POSITA would have found it obvious to combine Glines, with or without Gupta, with Chung. EX-1008 ¶¶475-481. Chung explains the benefits of its server cluster, including increasing performance, allowing server failures and maintenance to not interfere with system operation, and increasing capacity and scale. *Id.* ¶¶482-484; Chung, 1:16-40.

A POSITA would have understood that Glines and Glines-Gupta would have been compatible with Chung. EX-1008 ¶¶485-487. Chung teaches resolving a DNS request to a common cluster address and then responding via a particular server in the cluster (e.g., one of server 14-1 through 14-N). Chung, 2:33-3:10, Fig. 2. Glines and Glines-Gupta similarly include name servers that are addressable by an anycast address, where one name server responds and provides the IP address of a content server from which the user can directly retrieve content. *Supra* Limitations [1a], [1b], [1c]. Chung also contemplates the use of a variety of “dispatching techniques” to distribute requests among its cluster of servers, and a POSITA would have understood these are compatible with the routing described in Glines. *See, e.g.*, Chung, 2:51-61; Glines, 19:13-16.

Combining Chung’s use of a cluster with Glines’s/Glines’s-Gupta’s replica content servers 50-1, 50-2, and 50-3 would have been nothing more than combining known features in the art according to understood principles to yield predictable results. The result would have been a cluster of servers responding to the request for

content after one of the DNS servers resolves the hostname to the cluster address, and then providing content in Glines (or Glines-Gupta) via Chung's server cluster 34. A POSITA would have had an expectation of success in the combination because deploying multiple servers as a server cluster was known and understood in the art. EX-1008 ¶¶488-491; Chung, 1:16-40.

B. Claim 4: “grouped as a cluster”

Glines, alone or combined with Gupta, discloses claim 2. *Supra* VII.C.

Glines, with or without Gupta, in further combination with Chung, would have rendered this limitation obvious. EX-1008 ¶¶492-494.

Hosting content on multiple servers, including server clusters, was known and understood in the prior art. *Id.*; Chung, 1:15-40, 2:33-3:10; *also id.*, Abstract, 3:65-4:17, 5:39-6:15. For example, Chung discloses a server cluster that is “a group or set of servers interconnected or otherwise configured to host a network service.” Chung, 5:47-52. Chung's server cluster may advertise a single cluster address for its servers in the cluster to support the service, including serving content. *Id.*, 1:29-33, 2:33-3:10, Fig. 2.

Glines itself suggests deploying its servers as a cluster by disclosing they are arranged in a peering arrangement and “operate quite closely together . . . , such as on a common local area network segment sharing the same internetwork device 16-50-1.” Glines, 13:3-8, Fig. 1; EX-1008 ¶¶495-496. Glines also discloses

reporting a “cluster address” in response to the anycast message and the internetwork device may be a “router,” similar to Chung. Chung, 7:18-20, 15:23-25, 20:1-2, 7:14-18, 11:18-20.

This at least suggests deploying Glines (and Glines-Gupta) as Chung expressly discloses, including using a router (e.g., internetwork device 16-50-1) to connect users with a server cluster. EX-1008 ¶¶495-496. This would have improved the capacity, scalability, and performance of Glines’s system, as discussed above. *Id.*; *also supra* §VII.E (claim 4 (Grounds 1A/1B)), §VIII.A.

A POSITA would accordingly have combined Chung’s use of a server cluster with Glines (or Glines and Gupta) so that each name server includes a plurality of content servers grouped as a cluster, as claimed. EX-1008 ¶¶497-498. This would not alter the features of Glines or Gupta discussed elsewhere in this petition. *Id.*


IX. Grounds 3A/3B: Glines, with or without Gupta, and Chiou

A. Motivation to Combine

Chiou discloses methods of improving cache coherency. Chiou, 1:4-41; *also id.*, 2:15-46, 2:56-67. Caches must refresh themselves frequently “so that applications reading from the cache will not get ‘stale’ (e.g., invalid) data from the cache.” *Id.*, 1:4-41. Chiou receives requests for content, determines if content is cached, and either serves the content or requests it from an origin or “target” server. *Id.*, 12:52-13:40, Fig. 4.

A cache manager checks a “cache tag table,” which contains “pointers to the data stored in the cache.” *Id.*, 11:33-51, 12:61-67. Chiou discloses that, “[i]f a cache hit occurs (e.g., the requested data is located in the SSC cache *and is valid data*),” the read request is forwarded and the cache manager retrieves the requested data from the cache location and sends it to the requester. *Id.*, 13:5-17, Fig. 4.

Stale data may be determined based on, for example, Time-to-Live (TTL) expiration. *Id.*; *also id.*, 1:29-40. Using this time duration technique, a cache stores content for a predefined duration, after which it is stale.



Id., Fig. 4.

It would have been obvious to combine Glines, with or without Gupta, with Chiou so only valid data stored in content servers. EX-1008 ¶¶499-504. Determining whether data is present and valid was common in cache systems. *Id.* ¶¶505-506; EX-1019, 13:44-15:28; EX-1013, 24:7-51. A POSITA would have been motivated to implement Chiou's teachings in Glines or Glines-Gupta to improve user experience and provide the correct data in response to requests. EX-1008 ¶¶507-508.

This would have been nothing more than combining known features to yield predictable results. *Id.* ¶¶509-513. Chiou contemplates employing its methods on distributed cache systems, including those serving Internet content from distributed servers. Chiou, 1:5-2:53. Chiou's techniques were therefore compatible with Glines and Glines-Gupta, and a POSITA would have had a reasonable expectation of success in the combination. EX-1008 ¶¶509-513.

B. Claim 7: “if a valid version of the content is available ... serving .. otherwise obtaining”

Glines, alone or combined with Gupta, discloses claim 6. *Supra* §VII.G.

Glines, with or without Gupta, in further combination with Chiou, would have rendered this limitation obvious. EX-1008 ¶¶514-519.

Chiou discloses confirming the validity of available content and either serving valid available content or obtaining valid content from a “target” server and serving that content. *E.g.*, Chiou, 12:52-13:40, Fig. 4.

A POSITA would have been motivated to implement these teachings in Glines or Glines-Gupta to ensure only valid content is served and improve user experience. *Supra* §IX.A; EX-1008 ¶¶514-517.

In the combined Glines-Chiou and Glines-Gupta-Chiou systems, the content server would determine whether a valid version of the requested content is available on the content server (e.g., whether there was a “cache hit” or “miss”) as taught by Chiou. Chiou, 12:52-13:40; *see also id.*, 13:5-17 (determining if “the requested data is located in the [] cache and is valid data”). As Chiou explains, this may include verifying whether the data is stale using, for example, “pre-set time limits.” *Id.*, 12:52-13:40, 1:29-40.

If the data is either not available or is stale, the content server would request the content from the origin server according to Chiou’s teaching of reading requested data from a “target disk” and Glines’s and Gupta’s teachings regarding receiving content from an origin server. Chiou, 13:17-40; *supra* §VII.H (claim 7, regarding teachings of Glines and Gupta); EX-1008 ¶¶517-518.

Combining Glines or Glines-Gupta with Chiou in this way would not alter the features of Glines or Gupta discussed elsewhere in this petition. EX-1008 ¶519.

C. Claim 8: “origin server associated with said content provider”

Glines, alone or combined with Gupta and/or Chiou, discloses claim 7. *Supra* §§VII.H, IX.B.

Glines, with or without Gupta, in further combination with Chiou, would have rendered this limitation obvious. EX-1008 ¶¶520-525

As explained in claim 7, in the combined Glines-Chiou and Glines-Gupta-Chiou systems, the content server would request the content from the origin server. *Supra* §§VII.H, IX.B. For those reasons, Glines-Chiou or Glines-Gupta-Chiou would have rendered this limitation obvious.

Combining Glines or Glines-Gupta with Chung in this way would not alter the features of Glines or Gupta discussed elsewhere in this petition. EX-1008 ¶525.

D. Claim 9: “us[ing] a table”

Glines, alone or combined with Gupta and/or Chiou, discloses claim 7. *Supra* §§VII.H, IX.B.

Glines, with or without Gupta, and Chiou, would have rendered this limitation obvious. EX-1008 ¶¶526-532.

Chiou discloses a “cache tag table” that includes pointers to cached data. Chiou, 12:52-13:40. For each piece of data stored in cache, the cache tag table includes an entry that provides a “target device address” or “target IP address” associated with the data. *Id.* The “target” in Chiou is the server or location where the data is requested if the data is not located (or is stale) in cache, e.g., the origin server. *Id.*; *also id.*, 2:4-45. Chiou discloses that the cache tag table is checked and if there

is a cache miss (e.g., the content is stale), the system uses the table to identify the target device from which to obtain the content, as claimed. *Id.*, 12:52-13:40.

A POSITA would have found it obvious to combine Chiou's cache tag table with Glines, with or without Gupta, for example, because using the target device from the table would have reduced lookup time by determining the source of data directly and without additional DNS resolution, which can add latency. EX-1008 ¶¶529-532. This also would have made the system more robust against DNS outages, misconfigurations, or delays. *Id.* The combination would have been nothing more than combining known elements to yield predictable results, and there would have been a reasonable expectation of success in the combination. *Id.*

Combining Glines or Glines-Gupta with Chung in this way would not alter the features of Glines or Gupta discussed elsewhere in this petition. *Id.*

E. Claim 19

1. [19p]-[19h]

The following limitations are substantively identical to limitations addressed above and would have been obvious for the same reasons. EX-1008 ¶¶533-557.

<u>Claim Element</u>	<u>Corresponding Disclosure</u>
[19p]	[1p] (§VII.M.1)
[19a]	[1a] (§VII.B.2)

[19b]	Claim 2 (§VII.C)
[19c]	[1b] (§VII.B.3.)
[19d]	[1c] (§VII.B.4.)
[19e]	[1e] (§VII.B.6.)
[19f]	Claim 11 (§VII.K)
[19g]	Claim 6 (§VII.G)
[19h]	Claim 7 (§IX.B)

F. Claim 20: “load balancing policy”

Glines, alone or combined with Gupta and/or Chiou, discloses claim 19. *Supra* §§VII.Q, IX.E.

Glines, with or without Gupta, and Chiou, would have rendered this limitation obvious. EX-1008 ¶¶558-561.

As explained above, *supra* §VII.R, Glines alone or combined with Gupta, would have rendered claim 20 obvious. The addition of Chiou to the combination would not have disturbed Glines’s (or Glines-Gupta’s) teaching that content server selection may be based on a load balancing policy. Accordingly, Glines, with or

without Gupta, and Chiou, would have rendered this limitation obvious. EX-1008 ¶¶560-561.

X. Grounds 4A/4B and 5A/5B: Glines, with or without Gupta and Chiou, and Colby

A. Motivation to Combine

Colby discloses a content-aware flow switch for a “set of local web servers 100a-c, constituting a web server farm 150.” Colby, 5:43-45. Switch 110 may switch packets based on a number of factors, including server loading and network path optimization. *Id.*, 6:7-28, Abstract, 20:1-33.

A POSITA would have found it obvious to combine Glines, Glines-Gupta, Glines-Chiou, or Glines-Gupta-Chiou with Colby to utilize a switch such as Colby’s flow switch 100 at each name server group to select a content server based on a policy. Colby, Abstract, 5:23-6:28, 6:35-67. Doing so would have distributed load across multiple content servers at each location. Based on load, connection count, and/or capacity, the “best-fit” server would be selected based on these metrics. Colby, Abstract, 6:7-28, 20:1-33. This would maintain server stability across Glines’s content servers. EX-1008 ¶¶562-570.

This would have been nothing more than combining known elements to yield predictable results. *Id.* ¶¶571-576. A POSITA would have had a reasonable expectation of success in the combination because it was known to apply switches

to server farms to respond to content requests, including DNS requests. *Id.*; Glines, 19:13-22; Chase, 1:19-36; Colby, 5:23-41; EX-1011, 126.

B. Claim 10: “load balancing policy”

Glines, alone or combined with Gupta, discloses claim 1. *Supra* §VII.B.

Glines, with or without Gupta, in further combination with Colby, would have rendered this limitation obvious. EX-1008 ¶¶577-581.

Colby discloses a switch 100 for selecting a server from a plurality of servers to provide content to a user using, e.g., a load-balancing algorithm. Colby, Abstract, 6:7-28, 20:1-33, Fig. 1c. Colby also discloses rules that identify a load balance algorithm to use when accessing particular content. *Id.*, 6:35-67, Fig. 2. A POSITA would have been motivated to combine Colby’s switch and load-balancing algorithm (a policy) to select a content server as claimed for the reasons explained in §X.A. EX-1008 ¶¶577-581.

Combining Glines or Glines-Gupta with Colby in this way would not alter the features of Glines or Gupta discussed elsewhere in this petition. *Id.*

C. Claim 11: “switch, selecting the content server”

Glines, alone or combined with Gupta, discloses claim 1. *Supra* §VII.B.

Glines, with or without Gupta, in further combination with Colby, would have rendered this limitation obvious. EX-1008 ¶¶582-586.

As discussed in claim 10, Colby's switch 110 applies an algorithm, such as a load-balancing algorithm, to select a server from multiple servers. *Supra* §X.B; Colby, Abstract, 6:7-28, 20:1-33, Fig. 1c. Colby also discloses content rules that provide a load-balancing algorithm to use when accessing particular content. *Id.*, 6:35-67, Fig. 2.

A POSITA would have implemented these teachings in Glines-Colby or Glines-Gupta-Colby to apply a load-balancing algorithm (a policy) to select a content server using a switch as claimed for the reasons explained in §X.A. EX-1008 ¶¶582-586.

Combining Glines or Glines-Gupta with Colby in this way would not alter the features of Glines or Gupta discussed elsewhere in this petition. *Id.*

D. Claim 12: “predetermined selection policy”

Glines, alone or combined with Gupta and/or Colby, discloses claim 11. *Supra* §§VII.K, X.C.

Glines, with or without Gupta, in further combination with Colby, would have rendered this limitation obvious. EX-1008 ¶¶587-591.

As discussed in claims 10 and 11, Colby discloses a switch selecting a server based on, for example, a load-balancing algorithm. *Supra* §§X.B, X.C. Colby's load-balancing algorithm is a “predetermined selection policy,” as claimed. Colby's switch may consider factors such as source and destination IP addresses, transport

layer protocol, and transport layer source and destination port numbers when selecting a content server. Colby, 6:7-28. Consideration of these factors would likewise be a “predetermined selection policy,” as claimed. EX-1008 ¶¶587-591.

A POSITA would have implemented these teachings in Glines-Colby or Glines-Gupta-Colby to select a content server using a switch based on a predetermined selection policy as claimed for the reasons explained in Section X.A. *Id.*

Combining Glines or Glines-Gupta with Colby in this way would not alter the features of Glines or Gupta discussed elsewhere in this petition. *Id.*

E. Claim 13

Glines, alone or combined with Gupta, and Colby discloses claim 13. *Supra* §VII.M.

Limitations [13a]-[13d], [13f]-[13h], and [13j] are disclosed by Glines, alone or combined with Gupta for the reasons discussed above. *Supra* §VII.M. Colby discloses limitations [13e], and [13i] as discussed below, and a POSITA would have been motivated to combine Glines, with or without Gupta, and Colby for the reasons discussed. *Supra* §X.A. A POSITA would have understood that this combination would not have altered the features of Glines and Gupta discussed in connection with Limitations [13a]-[13d], [13f]-[13h], and [13j]. EX-1008 ¶¶592-595.

1. [13e]

Glines, with or without Gupta, in further combination with Colby, would have rendered the “switch” of Limitation [13e] obvious. EX-1008 ¶¶596-599.

As discussed in claims 10 and 11, Colby discloses a switch for selecting a server. *Supra* §§X.B, X.C. Colby applies a load-balancing algorithm, and its switch may consider factors such as source and destination IP addresses, transport layer protocol, and transport layer source and destination port numbers when selecting a content server. Colby, 6:7-28. A POSITA would have implemented these teachings in Glines-Colby or Glines-Gupta-Colby to select a content server using a switch for the reasons explained in §X.A. EX-1008 ¶¶596-599.

2. [13i]

Glines, with or without Gupta, in further combination with Colby, would have rendered this limitation obvious. EX-1008 ¶¶600-605.

Limitation [13i] differs from Limitation [13e] in that a switch associated with the “second” DNS server selects the “second” content server.

Dependent claim 16 recites the switch of (c22), limitation [13e], and the switch of (c12), limitation [13i], are the same switch. EX-1001, claim 16.

Element [13i] therefore encompasses a DNS resolution of the second request using the same switch.

Glines and Colby, with or without Gupta, would have rendered this limitation obvious for the same reasons provided for [13e]. *Supra* §§X.E.1, VII.M.2.

It would have been obvious to manage more than one request for content, including multiple requests for the same or different content. *Supra* §VII.M.4 (Limitation [13h]). Similar or repeated processing would occur to fulfill those requests, including using the same switch (or a second switch) to select a particular content server to provide the content to the user. *Id.*; EX-1008 ¶¶604-605.

F. Claim 14: “load balancing policy”

Glines, alone or combined with Gupta, and/or Colby, discloses claim 13. *Supra* §§VII.M, X.E.

Glines, with or without Gupta, in further combination with Colby, would have rendered this limitation obvious. EX-1008 ¶¶606-610.

As explained, in connection with claims 10 and 11, Glines, alone or combined with Gupta, in further combination with Colby, would have rendered obvious selection of a content server based on a load balancing policy. *Supra* §§X.B, X.C.

Further, it would have been obvious to manage more than one request for content, including multiple requests for the same or different content, and content server selection would proceed in the same way for each request. §VII.M.3 (Limitation [13g]). Repeating the same process twice, for example selecting a content server by a load balancing policy across the same or different service groups

would have been an obvious and expected implementation of the system as explained in claims 10 and 11. *Id.*; EX-1008 ¶¶606-610.

Combining Glines or Glines-Gupta with Colby in this way would not alter the features of Glines, Gupta, or Colby discussed elsewhere in this petition. EX-1008 ¶610.

G. Claims 15-17

The following limitations are substantively identical to limitations addressed above, and regardless of the addition of Colby to claim 13, would have been obvious for the same reasons. *Id.* ¶¶611-622.

<u>Claim</u>	<u>Corresponding Disclosure</u>
15	[13h] (§VII.M.4.)
16	[13i] (§VII.M.5.)
17	[13h] (§VII.M.4.)

H. Claim 18

Glines, alone or combined with Gupta, and Colby, discloses claim 18. *Supra* §VII.P.

Limitations [18p]-[18c], [18e], and [18g] are disclosed by Glines, alone or combined with Gupta for the reasons discussed above. *Supra* §VII.P. Colby

discloses limitations [18d] and [18f] as discussed below, and a POSITA would have been motivated to combine Glines, with or without Gupta, and Colby for the reasons discussed. *Supra* §X.A. A POSITA would have understood that this combination would not have altered the features of Glines and Gupta discussed in connection with Limitations [18p]-[18c], [18e], and [18g]. EX-1008 ¶¶623-626.

1. [18d] and [18f]

Glines, with or without Gupta, in further combination with Colby, would have rendered the “switch” of Limitation [18d] and the “switch ... operable to perform load balancing” of Limitation [18f] obvious. EX-1008 ¶¶627-628.

As discussed in claims 10 and 11, Colby discloses a switch for selecting a server. *Supra* §§X.B, X.C. Colby applies a load-balancing algorithm, and its switch may consider factors such as source and destination IP addresses, transport layer protocol, and transport layer source and destination port numbers when selecting a content server. Colby, 6:7-28. A POSITA would have implemented these teachings in Glines-Colby or Glines-Gupta-Colby to select a content server using a switch for the reasons explained in Section X.A. EX-1008 ¶¶627-628.

I. Claim 19

Glines, alone or combined with Gupta and/or Chiou, and Colby, discloses claim 19. *Supra* §VII.Q.

Limitations [19p]-[19e], [19g], and [19h] are disclosed by Glines, alone or combined with Gupta for the reasons discussed above. *Supra* §VII.Q.

Chiou alternatively discloses Limitation [19h] as discussed above, and a POSITA would have been motivated to combine Glines, alone or combined with Gupta, with Chiou for the reasons discussed. *Supra* §IX.A.

Colby discloses Limitation [19f] as discussed below, and a POSITA would have been motivated to combine Glines, with or without Gupta, and Colby for the reasons discussed. *Supra* §X.A.

A POSITA would have understood that the combination of Colby would not have altered the features of Glines, Gupta, and Chiou discussed in connection with Limitations [19p]-[19e], [19g], and [19h]. EX-1008 ¶¶629-633.

1. [19f]

Glines, with or without Gupta and/or Chiou, in further combination with Colby, would have rendered the “switch” of Limitation [19f] obvious. EX-1008 ¶¶634-635.

As discussed in claims 10 and 11, Colby discloses a switch for selecting a server. *Supra* §§X.B, X.C. Colby applies a load-balancing algorithm, and its switch may consider factors such as source and destination IP addresses, transport layer protocol, and transport layer source and destination port numbers when selecting a content server. Colby, 6:7-28.

A POSITA would have implemented these teachings with Glines, with or without Gupta and/or Chiou, to select a content server using a switch for the reasons explained in §X.A. EX-1008 ¶¶634-636.

J. Claim 20: : “load balancing policy”

Glines, alone or combined with Gupta, Chiou, and/or Colby, discloses claim 19. *Supra* §§VII.Q, IX.E, X.I.

Glines, with or without Gupta and/or Chiou, in further combination with Colby, would have rendered this limitation obvious. EX-1008 ¶¶637-640.

As explained above, *supra* §VII.R, Glines alone or combined with Gupta, would have rendered claim 20 obvious. The addition of Colby to the combination would not have disturbed Glines’s (or Glines-Gupta’s, Glines-Colby’s, or Glines-Gupta-Colby’s) teaching that content server selection may be based on a load balancing policy. EX-1008 ¶¶637-640. Accordingly, Glines, with or without Gupta, Chiou, and Colby would have rendered this limitation obvious. *Id.*

Combining Glines, Glines-Gupta, Glines-Chiou, or Glines-Gupta-Chiou with Colby in this way would not alter the features of Glines, Gupta, Colby, or Chiou discussed elsewhere in this petition. *Id.*

K. Claim 21

Glines, alone or combined with Gupta, and Colby discloses claim 21. *Supra* §VII.S.

Limitations [21p]-[21d] are disclosed by Glines, alone or combined with Gupta for the reasons discussed above. *Supra* §VII.S.

Colby discloses Limitation [21e] as discussed below, and a POSTA would have been motivated to combine Glines, with or without Gupta, and Colby for the reasons discussed. *Supra* §X.A.

A POSITA would have understood that this combination would not have altered the features of Glines and Gupta discussed in connection with Limitations [21p]-[21d]. EX-1008 ¶¶641-644.

1. [21e]

Glines, with or without Gupta, in further combination with Colby, would have rendered the “switch” of Limitation [21e] obvious. EX-1008 ¶¶645-646.

As discussed in claims 10 and 11, Colby discloses a switch for selecting a server. *Supra* §§X.B, X.C. Colby applies a load-balancing algorithm, and its switch may consider factors such as source and destination IP addresses, transport layer protocol, and transport layer source and destination port numbers when selecting a content server. Colby, 6:7-28. A POSITA would have implemented these teachings in Glines-Colby or Glines-Gupta-Colby to select a content server using a switch for the reasons explained in §X.A. EX-1008 ¶¶645-646.

L. Claim 22

Glines, alone or combined with Gupta, and Colby, discloses claim 22. *Supra* §VII.T.

Limitations [22p]-[22c], [22e], and [22f] are disclosed by Glines, alone or combined with Gupta, for the reasons discussed above. *Supra* §VII.T.

Colby discloses Limitations [22d] and [22g] as discussed below, and a POSITA would have been motivated to combine Glines, with or without Gupta, and Colby for the reasons discussed. *Supra* §X.A.

A POSITA would have understood that this combination would not have altered the features of Glines and Gupta discussed in connection with Limitations [22p]-[22c], [22e], and [22f]. EX-1008 ¶¶647-650.

1. [22d] and [22g]

Glines, with or without Gupta, in further combination with Colby, would have rendered the “switch” of Limitation [22d] and the switch “operable to select the content server” of Limitation [22g] obvious. EX-1008 ¶¶651-652.

As discussed in claims 10 and 11, Colby discloses a switch for selecting a server. *Supra* §§X.B, X.C. Colby applies a load-balancing algorithm, and its switch may consider factors such as source and destination IP addresses, transport layer protocol, and transport layer source and destination port numbers when selecting a content server. Colby, 6:7-28. A POSITA would have implemented these teachings

in Glines-Colby or Glines-Gupta-Colby to select a content server using a switch for the reasons explained in §X.A. EX-1008 ¶¶651-652.

XI. 35 U.S.C. § 314

Fintiv Factor 1 is neutral because, although Petitioner will move to stay the litigation, the district court has not considered or ruled on any stay motion. *Sotera Wireless, Inc. v. Masimo Corp.*, IPR2020-01019, Paper 12 at 14 (P.T.A.B. Dec. 1, 2020) (precedential).

Factor 2 is neutral. The trial is presently scheduled for September 28, 2026. EX-1020, 3. Assuming Patent Owner files a preliminary response three months from this Petition (July 2025) and the Board institutes review three months later (October 2025), the Board’s anticipated Final Written Decision will be October 2026—a matter of days from current trial date. District court trial dates are subject to change, while the Board’s Final Written Decision is statutorily-defined.

Factor 3 favors institution. Fact discovery just opened on March 4, 2025, and the Markman hearing is not until September 12, 2025. EX-1020, 1-2. Petitioner diligently brought this challenge approximately three months after receiving infringement contentions (and notice of the asserted claims). EX-1021; *Sotera*, Paper 12 at 16-17.

Factor 4 favors institution. This Petition challenges twenty-three claims, while Patent Owner has asserted only eleven. *See* EX-1021, 3. Resolving invalidity

questions here would mitigate duplicative efforts in this litigation and in future cases. Petitioner also stipulates that, if this Petition is instituted, it will not pursue in district court the grounds identified in this Petition or that could have been reasonably raised in the IPR (i.e., any ground that reasonably could have been raised under 35 U.S.C. §§ 102 or 103 on the basis of prior art patents or printed publications) with respect to the claims subject to institution. EX-1027; *Sotera*, Paper 12 at 18-19.

Factor 5 is neutral because Petitioner is the defendant in the litigation.

Factor 6 favors institution. No other party has sought review of the '112 patent and Petitioner is aware of no other party likely to do so, minimizing serial or parallel petitions. Further, the compelling merits of this petition outweigh Factors 1-5. Petitioner relies on prior art that the Office never applied, presents different grounds, and relies on expert testimony.

XII. 35 U.S.C. § 325

Petitioner's grounds were not previously considered. *Advanced Bionics*, step 1, and *Becton, Dickinson* factors (a), (b), and (d) favor institution because Glines, Chung, and Chiou, were not identified during prosecution. Gupta and Colby were among 500+ other references and neither was discussed by the Examiner. *See Microsoft Corp. v. Parallel Networks Licensing LLC*, IPR2015-00486, Paper 10, 15–20 (P.T.A.B. July 15, 2015). The references are not cumulative because they teach and render obvious all challenged claims. The Board need not reach the second

Advanced Bionics prong. See *Thorne Rsch., Inc. v. Trustees of Dartmouth Coll.*, IPR2021-00491, Paper 18 at 8 (P.T.A.B. Aug. 12, 2021).

XIII. MANDATORY NOTICES UNDER 37 C.F.R. § 42.8

A. Real Party-in-Interest

Google LLC³ is the real party-in-interest for this petition.

B. Related Matters

The '112 patent is, or has been, involved in the following proceedings:

Name	Number	Forum	Filed
<i>Sandpiper CDN, LLC v. Google LLC</i>	2-24-cv-03951	C.D. Cal.	May 10, 2024

C. Lead and Backup Counsel Information

Petitioner provides the following designation of counsel:

Lead Counsel	Back-Up Counsel
Erika H. Arner (Reg. No. 57,540) erika.arners@finnegan.com Finnegan, Henderson, Farabow, Garrett & Dunner, LLP 1875 Explorer Street Suite 800 Reston, VA 20190-6023 Tel: 571-203-2700; Fax: 202-408-4400	Kara Specht (Reg. No. 69,560) Sydney Kestle (Reg. No. 78,725) Cara E. Regan (Reg. No. 70,209) cara.regan@finnegan.com Finnegan, Henderson, Farabow, Garrett & Dunner, LLP 901 New York Ave. NW Washington DC, 20001 Tel: 202-408-4000; Fax: 202-408-4400

³ Google LLC is a subsidiary of XXVI Holdings Inc., which is a subsidiary of Alphabet Inc. XXVI Holdings Inc. and Alphabet Inc. are not real parties-in-interest to this proceeding.

	Daniel C. Tucker (Reg. No. 62,781) Alexander M. Boyer (Reg. No. 66,599) alexander.boyer@finnegan.com Finnegan, Henderson, Farabow, Garrett & Dunner, LLP 1875 Explorer Street Suite 800 Reston, VA 20190-6023 Tel: 571-203-2700; Fax: 202-408-4400
--	--

In a concurrently filed Power of Attorney, Google LLC has granted Power of Attorney to Practitioners at Finnegan, Henderson, Farabow, Garrett & Dunner, LLP.

Petitioner consents to electronic service at the following email address:
Google-Sandpiper-IPRs@finnegan.com.

XIV. CONCLUSION

Petitioner respectfully requests that the Board grant IPR and find all challenged claims unpatentable.

Dated: April 15, 2025

By: /Erika H. Arner/
Erika H. Arner (Reg. No. 57,540)

CERTIFICATE OF COMPLIANCE

Pursuant to 37 C.F.R. § 42.24(a)(1)(i), the undersigned hereby certifies that the foregoing Petition for *Inter Partes* Review of U.S. Patent No. 9,021,112 contains 13,851 words, excluding parts of this Petition exempted under § 42.24(a), as measured by the word-processing system used to prepare this paper.

Dated: April 15, 2025

/Erika H. Arner/

Erika H. Arner (Reg. No. 57,540)
Lead Counsel

CERTIFICATE OF SERVICE

The undersigned certifies that the foregoing Petition for *Inter Partes* Review of U.S. Patent No. 9,021,112, the associated Power of Attorney, and Exhibits 1001 through 1027 were served on April 15, 2025, by FedEx Priority Overnight® on the correspondence address of record indicated in the Patent Office's Patent Center system for U.S. Patent No. 9,021,112.

ATTN: Patent Docketing
Level 3 Communications, LLC
931 14th St.
Denver, CO 80202

Dated: April 15, 2025

/Daniel E. Doku/
Daniel E. Doku
Senior Litigation Legal Assistant
Finnegan, Henderson, Farabow,
Garrett & Dunner, LLP