

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

**DECLARATION OF SAMUEL H. RUSS, PH.D.,
IN SUPPORT OF THE PATENT OWNER'S RESPONSE**

Mail Stop "PATENT BOARD"

Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

Sandpiper EX2018
Google v. Sandpiper
IPR2025-00826

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	MY BACKGROUND AND QUALIFICATIONS	2
III.	LIST OF DOCUMENTS CONSIDERED	5
IV.	LEGAL UNDERSTANDING	7
	A. My Understanding of Claim Construction	7
	B. My Understanding of Obviousness	8
	C. My Understanding of a Person of Ordinary Skill in the Art	10
V.	PERSON OF ORDINARY SKILL IN THE ART	11
VI.	OVERVIEW OF THE STATE OF THE ART	11
VII.	THE '112 PATENT	14
	A. Overview of the '112 Patent	14
	B. Overview of the Prosecution History of the '112 Patent	17
VIII.	OVERVIEW OF THE PRIMARY REFERENCES	19
	A. Glines	19
	B. Gupta	22
IX.	DR. MOWRY'S PROPOSED OBVIOUSNESS GROUNDS	23
	A. Dr. Mowry does not establish that the cited art teaches "causing said hostname to be resolved to said common anycast address by an Internet Service Provider (ISP) DNS server" (Claim 1, Element [1d]).	23
	B. Dr. Mowry does not establish that the cited art teaches "a second request for second content associated with a second content provider" (Claim 23, Element 23[e])	28
	C. Dr. Mowry does not establish that the cited art renders obvious a system configured to support "a plurality of content providers" (Independent claims 13, 18, 22, and 23)	33
	1. Obviousness in view of Glines	34
	2. Obviousness in view of Glines and Gupta	37

D.	Dr. Mowry does not establish that the cited art teaches the “at least one switch” limitations (Independent claims 13, 18, 19, 21, and 22 and dependent claim 11)	39
E.	A POSITA would not have been motivated to combine Glines and Gupta in the manner suggested by Dr. Mowry for any Ground.	42
X.	CONCLUSION.....	49

I, Samuel H. Russ, hereby declare as follows.

I. INTRODUCTION

1. I have been retained as an expert witness on behalf of Patent Owner Sandpiper CDN, LLC (“Patent Owner”) for the above-captioned *inter partes* review (“IPR”). I am being compensated for my time in connection with this IPR at my normal hourly rate. My compensation is not contingent on either my expert opinion or the outcome of this proceeding. I have no other interest in this proceeding.

2. I understand that this declaration accompanies Patent Owner’s Response filed in the IPR involving U.S. Patent No. 9,021,112 (“the ’112 patent”) (EX1001), which resulted from U.S. Application No. 12/050,648 (“the ’648 application”), filed March 18, 2008, which claims priority as a divisional application to U.S. Application No. 09/982,721, filed October 18, 2001. My opinions regarding the knowledge in the art at the time of the ’112 patent are based on this date of October 18, 2001.

3. In preparing this declaration, I have reviewed the ’112 patent and each of the documents cited herein, in light of general knowledge in the art before the time of the ’112 patent. In formulating my opinions, I have relied upon my experience, education, and knowledge in the relevant art. In formulating my opinions, I have also considered the viewpoint of a person of ordinary skill in the

art (“POSITA”) (as discussed further below in Section V) at the time of the ’112 patent.

II. MY BACKGROUND AND QUALIFICATIONS

4. I received a Bachelor’s degree in Electrical Engineering from the Georgia Institute of Technology (“Georgia Tech”) in 1986 and a Ph.D. in Electrical Engineering from Georgia Tech in 1991.

5. From 2007 to the present, I have been a member of the faculty of the University of South Alabama as an Assistant and Associate Professor in the Department of Electrical and Computer Engineering. During that time, I have won awards for excellent teaching and have been actively publishing research in home networking and digital video recording (DVR) technologies. I am active in the Institute of Electrical and Electronic Engineers (IEEE) and was a Distinguished Lecturer for the IEEE Consumer Electronics Society. As an example of the type of work I perform at South Alabama, I managed the design, construction, and launch of a fully functional satellite from the International Space Station into space in 2022.

6. From 2000 to 2007, I worked for Scientific-Atlanta (now Cisco’s Service Provider Video Tech. Group), where I managed a cable set-top box (STB) design group that designed four STB models, including the Explorer 4200 (non-DVR) and 8300 (DVR) models. Both models sold several million units. As design-

group manager, I was responsible for managing the design and prototyping activities of the group and for interfacing with other groups (especially integrated-circuit design, procurement, software developers, the factory where prototypes were built, and product managers) and for maintaining the hardware and mechanical development schedule. Since the products were produced in extremely high volumes, the projects had very high visibility in the company, and therefore carried a great deal of responsibility.

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

8. Both my work with home networking and coaxial DVR video distribution were technologies that improve content distribution networks (CDNs). In order to implement these systems, I had to develop a detailed understanding of, *e.g.*, bandwidth and resource management, video distribution, and the requirements

for storage media to support multiple video streams. For example, because of my work on multi-room DVR and hard drive reliability, I identified the level of performance needed for a hard disk drive to support seven simultaneous streams of video when functioning as a home video server. The set-top boxes I designed supported video on demand and switched digital video, both forms of content distribution.

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

10. I also served on the faculty of Mississippi State University from 1994 to 1999 as an Assistant Professor in the Department of Electrical & Computer Engineering where I taught circuit board design and two-way interactive video classes, among other things.

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

12. I understand that a copy of my complete *curriculum vitae* has been submitted as Exhibit 2014, which includes a list of my publications and the cases

in which I have testified as an expert in deposition or trial in the past four years,
and other details regarding my expertise and qualifications.

III. LIST OF DOCUMENTS CONSIDERED

13. I have considered the following documents listed below in addition to any other documents cited in this declaration. To the best of my knowledge, these documents are true and accurate copies of what they purport to be. An expert in the field would reasonably rely on them to formulate opinions such as those set forth in this declaration.

Paper No.	Description
1	Petition for <i>Inter Partes</i> Review of U.S. Patent No. 9,021,112
15	Institution Decision

Exhibit No.	Description
1001	U.S. Patent No. 9,021,112 (“the ’112 Patent”)
1002	WO2001039470 A1 to Glines (“Glines”)
1003	U.S. Patent No. 6,405,252 to Gupta (“Gupta”)
1004	U.S. Patent No. 6,470,389 to Chung (“Chung”)
1005	U.S. Patent No. 6,792,507 to Chiou (“Chiou”)
1006	U.S. Patent No. 6,006,264 to Colby (“Colby”)
1007	Prosecution History of the ’112 patent
1008	Declaration of Todd C. Mowry, Ph.D.

Exhibit No.	Description
1009	<i>Curriculum Vitae</i> of Todd C. Mowry, Ph.D.
1010	U.S. Patent No. 8,041,814 to Chase
1011	James Kurose & Keith Ross, Computer Networking: A Top-Down Approach Featuring the Internet (1st ed. 2000)
1012	Andrew Tanenbaum, Computer Networks (3rd ed. 1996)
1013	U.S. Patent No. 6,829,654 to Jungck
1014	James Kurose & Keith Ross, Computer Networking: A Top-Down Approach Featuring the Internet (3rd ed. 2005)
1015	Mark Baker & Rajkumar Buyya, Cluster Computing at a Glance in High Performance Cluster Computing – Architectures and Systems, 3-47 (1999).
1016	U.S. Patent No. 6,578,066 to Logan
1017	U.S. Patent No. 8,041,814 to Chase
1018	U.S. Patent No. 7,523,181 to Swildens
1019	U.S. Patent No. 7,594,030 to Teodosiu
1023	Schroeder, et al., Scalable Web Server Clustering Technologies, IEEE Network, Vol. 14, Issue 3 (June 2000).
1024	The Wayback Machine – Yahoo! News, http://dailynews.yahoo.com/ , dated February 2, 2001 (available at https://web.archive.org/web/20010202075000/http://dailynews.yahoo.com/)
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 (available at https://web.archive.org/web/20010202073200/http://dailynews.yahoo.com/h/nm/20010201/ts/congress_ashcroft_dc_71.html)

Exhibit No.	Description
1026	A. Bestavros et al., Distributed Packet Rewriting and its Application to Scalable Server Architectures, Proceedings Sixth International Conference on Network Protocols, IEEE (Oct. 1998).
2001	“Cloud CDN: content delivery network,” <i>Google Cloud</i> , accessed at https://cloud.google.com/cdn
2015	“CDN Stands For Content Delivery Network: History and Evolution,” <i>BlazingCDN</i> , accessed at https://blog.blazingcdn.com/en-us/cdn-stands-for-content-delivery-network-history-and-evolution
2016	“What is a CDN (Content Delivery Network)?” <i>Amazon Web Services</i> , accessed at https://aws.amazon.com/what-is/cdn/
2017	“How Anycast Works to Bring Content Closer to Your Visitors,” <i>Imperva</i> , accessed at https://www.imperva.com/blog/how-anycast-works/
2019	“DNS root server,” Cloudflare, accessed at https://www.cloudflare.com/learning/dns/glossary/dns-root-server/
2020	Lee <i>et al.</i> , “Availability and Effectiveness of Root DNS servers: A long term study,” IEEE (2010).
2021	Deposition Transcript of Todd C. Mowry

IV. LEGAL UNDERSTANDING

A. My Understanding of Claim Construction

14. I understand that during an *inter partes* review proceeding, claims are to be construed in light of the specification as would be read by a person of ordinary skill in the relevant art at the time of the invention. I understand that claim terms are given their plain and ordinary meaning as would be understood by a

person of ordinary skill in the relevant art in the context of the patent's specification and prosecution history. A claim term, however, will not receive its ordinary meaning if the patentee acted as his or her own lexicographer and clearly set forth a definition of the claim term in the specification. In this case, the claim term will receive the definition set forth in the patent.

B. My Understanding of Obviousness

15. I understand that a patent claim is invalid if the claimed invention would have been obvious to a person of ordinary skill in the art at the time the application was filed. I understand that this means that even if all of the requirements of the claim cannot be found in a single prior-art reference that would anticipate the claim, the claim can still be invalid.

16. To obtain a patent, a claimed invention must have been, as of its priority date, nonobvious in view of the prior art in the field. I understand that a patent claim is obvious when the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art.

17. I understand that to prove that prior art or a combination of prior art renders a patent claim obvious, it is necessary to: (1) identify the particular references that, singly or in combination, render the patent claim obvious; (2)

specifically identify which elements of the patent claim appear in each of the asserted references; and (3) explain how the prior art references could have been combined in order to create the inventions claimed in the patent claim.

18. I also understand that when considering the obviousness of a patent claim, one should consider whether a teaching, suggestion, or motivation to combine the references exists so as to avoid impermissibly applying hindsight when considering the prior art. I understand this test should not be rigidly applied, but that the test can be important to avoiding such hindsight.

19. I understand that certain objective indicia can be important evidence as to whether a patent is obvious or nonobvious. Such indicia include: (1) commercial success of products covered by the patent claims; (2) a long-felt need for the invention; (3) failed attempts by others to make the invention; (4) copying of the invention by others in the field; (5) unexpected results achieved by the invention as compared to the closest prior art; (6) praise of the invention by the infringer or others in the field; (7) the taking of licenses under the patent by others; (8) expressions of surprise by experts and those skilled in the art at the making of the invention; and (9) the patentee proceeded contrary to the accepted wisdom of the prior art.

20. I also understand that “obviousness” is a legal conclusion based on the underlying factual issues of the scope and content of the prior art, the differences

between the claimed invention and the prior art, the level of ordinary skill in the pertinent art, and any objective indicia of non-obviousness. For that reason, I am not rendering a legal opinion on the ultimate legal question of obviousness. Rather, my testimony addresses the underlying facts and factual analysis that would support a legal conclusion of obviousness or non-obviousness, and when I use the term obvious, I am referring to the perspective of one of ordinary skill at the time of the '112 patent.

C. My Understanding of a Person of Ordinary Skill in the Art

21. I understand that a POSITA is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity.

22. I understand that the following should be considered when deciding the level of ordinary skill in the art that someone would have had at the time the claimed invention was made:

- the levels of education and experience of persons working in the field;
- the types of problems encountered in the field; and
- the sophistication of the technology.

23. My opinion below explains how a POSITA would have understood the technology described in the documents I have identified herein around the time of the '112 patent.

24. Regardless of whether I use “I” or “POSITA” during my technical analysis below, all of my statements and opinions are always to be understood to be based on how a POSITA would have understood or read a document at the time of the ’112 patent.

V. PERSON OF ORDINARY SKILL IN THE ART

25. In my opinion, a POSITA at the time of the ’112 patent would have had at least a bachelor’s degree in computer science, electrical engineering, or other suitable field, and at least two years of work or research experience related to the field of content delivery networks. Further professional experience could substitute for education, and vice versa. It is also my opinion that the definition of a POSITA would not materially affect the analysis required for upholding the patentability of the challenged claims.

VI. OVERVIEW OF THE STATE OF THE ART

26. Content delivery networks (CDNs) provide critical services that enable content providers to quickly deliver online content to millions of consumers simultaneously over the Internet. In the early 1990s, the World Wide Web saw increasing adoption, becoming a household staple. This mass adoption led to data congestion issues from the ever-growing number of users seeking to simultaneously access Internet content. EX2015, 2.

27. A typical computer server in the 1990s could only handle a limited number of simultaneous connections before becoming overloaded. Moreover, signals take time to traverse routers and move through content-congested internet cables. Consumers living far from the physical server(s) hosting content experienced sluggish load times and high latency due to problems such as overloaded servers, congested network segments, and geographic separation. EX2015, 2-3.

28. The availability of CDN servers around the world allowed service providers to replicate appropriate content from customers' origin servers to appropriate CDN servers, and transparently rendezvous end-users requesting that content to the "best" CDN server to deliver that content, all while providing their customers with control over their content and user experience. EX2015, 3-6; EX2016, 2-3. This service and its architecture were quickly adopted by many others in the industry. EX2015, 3-6; EX2016, 2-3.

29. In the CDN model, duplicate copies of content are stored in servers located over a wide geographic area (*e.g.*, across the country or around the world). EX2015, 2; EX2016, 2-3. Propagating data from origin servers to edge servers is known as "caching." EX2015, 3-6; EX2016, 2-3. In other words, when data is propagated from its origin to one of the geographically distributed CDN servers, the data is cached at the CDN server. EX2015, 3-6; EX2016, 2-3. Using a CDN,

end-users connect to an optimal edge server—one that that is closer to them or that has available capacity. EX2016, 2-3. Distributing content across a network of servers alleviated data congestion issues and allowing consumers to connect to edge servers located near them reduced latency. Importantly, the CDN model is invisible to the user—they simply see the content they requested. Note that the closest server may not necessarily refer to one that is geographically closest but rather has the lowest latency. EX2015, 3; EX2016, 1.

30. In the early-to-mid 2000s, demand for CDNs exploded. EX2015, 7-8. Content providers like Google and YouTube use CDN technology to meet their enormous data streaming needs, as well as for video search results. EX2001, 1-3. As a result, techniques for controlling how users access content on a CDN became increasingly important.

31. One technology that fueled the rise in CDNs was anycast addressing. EX2017, 1. When browsing the Internet, one typical way that a session works is that a user types in a hostname (*e.g.*, “JohnSmithRoofing.com”), a system called Domain Name Service (DNS) maps the hostname to an IP address (the network address of the server hosting the JohnSmithRoofing.com website), and the browser uses the IP address to access the host (and therefore the website). This is called unicast addressing because it delivers a single address to the Web browser. EX2017, 2.

32. Anycast addressing works differently. To make a long story short, multiple servers are all assigned the same network address. EX2017, 2-3. The process works similarly—the user enters a hostname and DNS responds with an IP address—but, in this case, multiple servers all have the same IP address. EX2017, 2-3. When the Web browser then uses the anycast address, the routing system used by the Internet (*e.g.*, the Border Gateway Protocol, or BGP) is employed. Each of the multiple servers appears to the local router as a different route to the same server, and the Internet is optimized such that the local router chooses the “route” with the fewest numbers of routers in between. By choosing a “route,” the local router is actually choosing the nearest host. EX2017, 2-3.

33. This was and is a very powerful way to spread servers across an area and let existing Internet protocols provide a way to reach the nearest one. It promotes distribution of content without having to develop new protocols. One interesting side effect of this is that the nearest server may not be the fastest to respond—it may be the case that the nearest server is quite busy and one farther away is idle. As a result, it is not the case that anycast addressing is necessarily the best possible protocol. However, it does have the advantage of compatibility.

VII. THE '112 PATENT

A. Overview of the '112 Patent

34. The '112 patent discloses “[a] content distribution mechanism that

distributes content of a content provider at various sites across a network and selects the site that is nearest a content requestor using an anycast address that resides at each of the sites.” EX1001, Abstract. As the ’112 patent explains, “[a]s the Web server content provided by World Wide Web has continued to grow over time, so too has the number of users demanding access to such content.” EX1001, 1:36-38. “Unfortunately, the ever-increasing number of end users requesting Web content from Web sites has resulted in serious bandwidth and latency issues, which manifest themselves in delay to the end user.” EX1001, 1:38-41.

35. Therefore, when offering CDN services, there is a need to respond to requests by providing content quickly and without unnecessary network traffic to more distant sources. Providers of CDN services, however, face challenges related to serving content in a timely manner from multiple servers while maintaining a positive user experience. EX1001, 1:36-59.

36. The ’112 patent sought to provide solutions using DNS technology and anycast addressing in a CDN context in which the DNS infrastructure routes content requests to “a content site nearest the requestor” that is able to serve the requested content. EX1001, 1:63-2:24. Such techniques allow “the load from many end users [to be] distributed among multiple systems and across different parts of the network,” and “[b]ecause DNS uses a stateless protocol (UDP) for routing, the solution can handle anycast addressable caching without the problems associated

with anycast service.” EX1001, 2:11-24.

37. Figure 3 of the '112 patent illustrates embodiments of “a content distribution network (‘CDN’) 60.” EX1001, 4:18.




FIG. 3

EX1001, FIG. 3.

38. As the '112 patent explains, “[t]he CDN 60 includes at least one Web server, shown as a Web server 62, a DNS system 64 for Web server 62, and end-user stations 66a and 66b, all connected to a network 72.” EX1001, 4:19-21. The end-user clients “execute Web browsers 68a and 68b, respectively.” EX1001, 4:22-23. “Also connected to the network 72 are multiple content distribution nodes 76a, 76b and 76c, which support distributed content delivery services for one or more

Web sites on the network 72.” EX1001, 4:27-30.

39. As shown in Figure 3 and described by the ’112 patent, “[e]ach of the CDN nodes 76 includes a DNS system 78 coupled to and associated with a Web content server system or site 80.” EX1001, 4:34-36. Nodes 76 are “geographically dispersed ... so as to be as close as possible to various network entry points, exchanges or both.” EX1001, 5:11-13. “The caching servers 80 have unique IP addresses,” and “[t]he DNS systems 78 share a common IP address as well as have unique IP addresses.” EX1001, 5:19-21. As the patent explains, “[o]ne way to implement this content distribution configuration is to use an anycast address as the common address,” which is “a unicast address that is used in multiple places.” EX1001, 5:32-35. This anycast address is advertised to end-user clients “using a dynamic routing protocol, the Border Gateway Protocol,” and when a content request is received, “the selected DNS system resolves to the address of the cache server in the DNS system’s node.” EX1001, 5:59-6:32. In this manner, “the DNS anycast routing, in essence, serves to select the local content site (cache) from which content will be served.” EX1001, 6:27-31.

B. Overview of the Prosecution History of the ’112 Patent

40. I understand that the ’112 patent was filed as U.S. Application No. 12/050,648 (“the ’648 application) on March 18, 2008. I also understand that the ’112 patent is a divisional application of U.S. Application No. 09/982,721, filed

October 18, 2001.

41. During prosecution, I understand that the Office issued seven Office Actions before allowing the claims. EX1007, 64-348. Of note, the applicant amended independent claim 1 during prosecution to clarify its operation:

1. (Currently Amended) A method of content delivery in a network, the method comprising:
- (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;
 - (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;
 - (C) responsive to a request for content associated with the content provider and issued by a client, said request including at least said hostname,
 - (c1) causing said hostname to be resolved to said common anycast address by an Internet Service Provider (ISP) DNS server ~~associated with an Internet Service Provider (ISP)~~; and then
 - (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.

EX1007, 132.

42. I understand that the applicant explained that independent claim 1 “includes: (c1) causing the hostname to be resolved to the common anycast address by an ISP DNS server; and then (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.” EX1007, 139. As the applicant further explained, “in response to a request associated with a content provider (the request

including a hostname), the hostname is resolved to the common anycast address by an ISP DNS and then to an IP address for one of the content servers by one of the CDN DNS servers.” EX1007, 140 (emphasis in original). In other words, claim 1 requires “two DNS servers performing two resolves,” which none of the prior art taught. EX1007, 36.

43. Upon consideration of this and other arguments, I understand that the Examiner allowed the claims. EX1007, 10. The ’112 patent issued on April 28, 2015. EX1001, (45). The prosecution history of the ’112 patent is discussed in further detail in my analysis throughout this Declaration.

VIII. OVERVIEW OF THE PRIMARY REFERENCES

44. I understand that Petitioner’s expert, Dr. Mowry, relies on numerous references in his proposed obviousness grounds challenging the claims of the ’112 patent. Two of those references are particularly relevant to my opinions provided in this declaration: Glines (EX1002) and Gupta (EX1003). I provide brief overviews of these primary references below.

A. Glines

45. Glines discloses a “technique for redirecting client computer requests for content files to the closest replica of the requested content, by using anycast messaging.” EX1002, Abstract. In Glines, name servers are distributed throughout the network and are “addressable via a common anycast address, as well as being

individually addressable via a unicast address unique to each such name server.”

EX1002, 6:8-11. Each name server is paired with a content server “that contains (or can serve) desired content file replicas via unicast addressing.” EX1002, 6:11-13. “Each name server also advertises itself as being authoritative for the domains associated with the content files stored (or to be served from) the associated content server.” EX1002, 6:13-15.

46. Glines's general architecture is illustrated in Figure 1, as shown below. I have annotated the below figure for clarity.




FIG. 1

EX1002, FIG. 1 (annotated).

47. As illustrated in Figure 1, “a 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.” EX1002, 9:16-18. “For example, as the user types a Uniform Resource Locator (URL) into a web browser and/or uses a pointing device, such as a mouse, to click on a hyperlink embedded in a previously viewed web page, the URL for a particular web site is specified to the browser 12.” EX1002, 9:18-21. Glines’s purported “invention uses a particular addressing scheme to advertise the availability of alternative or replica name servers 53-1, 53-2, 53-3 for the domain located at the origin server 40,” and “[t]he 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.” EX1002, 10:19-25.

48. When a user enters a URL into the web browser, “[t]he DNS resolver 18 determines whether or not the requested content file is available locally.” EX1002, 13:27-14:1. “For example, it determines if ‘example.com’ is located in the local web server 20.” EX1002, 14:1-2. If the content is not available locally, “a request to resolve the location of ‘example.com’ is then sent to a root DNS server 38,” which “will be recursively worked through multiple root servers associated with the Internet 30 (not shown in Fig. 1) to resolve the IP address for a DNS server authoritative for the requested domain name.” EX1002, 14:8-12. The root DNS server returns the common anycast address for the authoritative name server, and “the browser then sends out a DNS request for the IP address of

‘www.example.com,’” which is “formulated as a UDP datagram specifying the common anycast address.” EX1002, 15:6-10; *see also* EX1002, 7:6-12. “The server pair that receives the request then responds by reporting the unique (unicast) address of the associated content replica server” EX1002, 7:18-19, 15:26-16:2. The browser then sends a request to the reported unicast address of the content replica server to retrieve the requested content. *See* EX1002, 7:20-22, 16:3-8.

B. Gupta

49. Gupta discloses “a network of point of presence (POP) computers that provides a solution for integrated content hosting and traffic management.”

EX1003, 2:5-8. In contrast to Glines, which resolves client requests to anycast IP addresses, Gupta uses only unicast addressing in its DNS resolution. EX1003, 6:4-6. When a user requests a webpage with a domain within the POP network, “[t]he user’s computer sends a request for name resolution of the domain name they are requesting content from to the configured local DNS server (LDNS) for that users computer.” EX1003, 2:24-29. The LDNS resolves the IP address by sending a query to an authoritative DNS server for the domain name. EX1003, 2:29-33.

50. The authoritative DNS server identifies a subset of servers on the POP network addressable by the requested domain name, and “selects the optimal server/servers for the user based on rules defined in the POP network configuration.” EX1003, 2:37-47. As Gupta explains, “[t]hese rules define how

criteria such as POP/customer server health, latency from POP/customer servers to the user making the request, packet loss experienced at each POP, cost associated with each POP/customer server, and the like.” EX1003, 2:47-51. “These factors are then combined to define the optimal server for a given user making a request for a given domain name.” EX1003, 2:51-53. In other words, Gupta does not use anycast addressing to enable a client to access the nearest server; rather, Gupta monitors network conditions and server loading to identify the best server for each client’s request and then routes the client’s request to the identified server.

IX. DR. MOWRY’S PROPOSED OBVIOUSNESS GROUNDS

51. For purposes of this declaration, I have been asked to opine on certain positions taken by Petitioner and Dr. Mowry. That I omit certain claims, claim elements, or presented grounds from my analysis should not be viewed as agreement that any claims or claim elements are obvious.


A. Dr. Mowry does not establish that the cited art teaches “causing said hostname to be resolved to said common anycast address by an Internet Service Provider (ISP) DNS server” (Claim 1, Element [1d]).

52. Claim 1 recites, “(c1) causing said hostname to be resolved to said common anycast address by an Internet Service Provider (ISP) DNS server.” EX1001, 9:11-13 (element [1d]¹). Dr. Mowry alleges that “Glines, alone or in

¹ For clarity, I reference Petitioner’s element identifiers that are set forth in

combination with Gupta, discloses this limitation.” EX1008, ¶141. But Dr. Mowry does not identify “an Internet Service Provider (ISP) DNS server” that resolves “said hostname ... to said anycast address,” which is required by element [1d]. EX1001, 9:11-13.

53. In particular, Dr. Mowry points to Glines’s “**local DNS resolver 18, which is an ISP DNS server**” according to Dr. Mowry. EX1008, ¶148 (emphasis added). Dr. Mowry alleges that when a user enters a URL, “local DNS resolver 18 ... sends a request to a root DNS server 38,” which, in turn, “return[s] the anycast address” of authoritative name servers. EX1008, ¶¶154-155.



EX1002, FIG. 1 (annotated).

54. But Dr. Mowry’s mapping does not meet the claim language—which requires “causing said hostname **to be resolved** to said common anycast address **by an Internet Service Provider (ISP) DNS server,**” EX1001, 9:11-13 (emphasis added)—because Glines’s “local DNS resolver 18” does not perform a resolution of the hostname to an anycast address. In other words, claim 1 expressly requires the ISP DNS server to resolve the hostname to an anycast address, and nowhere does Dr. Mowry explain how Glines’s “local DNS resolver 18” performs that resolution. EX1008, ¶¶154-155. Indeed, Dr. Mowry relies on “root DNS server 38” for that purpose—not Glines’s “local DNS resolver 18.” EX1008, ¶¶154-155 (“**[R]oot DNS name server 38** is ‘programmed to ... **return the anycast address**’” (emphasis added)). I agree in this respect—This understanding is consistent with Glines’s disclosure, which explains that **root DNS server 38** receives “a request to resolve the location of ‘example.com’” and “**resolve[s] the IP address for a DNS server authoritative for the requested domain name.**” EX1002, 14:8-12 (emphasis added).

55. A POSITA would have also understood that a root domain name server is not an “*Internet Service Provider (ISP) DNS Server*,” nor does Dr. Mowry allege that it is. “Root servers are DNS nameservers that operate in the root zone.” EX2019, 1. “These servers can directly answer queries for records stored or cached within the root zone, and they can also refer other requests to the appropriate Top

Level Domain (TLD) server.” *Id.*; see also EX2020, 862 (“Domain Name Systems (DNS) servers provide a critical service to users and application on the internet. At the top of the DNS hierarchy is the Root DNS servers providing pointers to Top Level Domain servers on the Internet.”). Originally, there were only 13 root servers in the world. EX2019, 2. While there remain only 13 IP addresses, those IP addresses today act as anycast addresses to allow for more root servers to exist. *Id.* But, importantly, these root servers are not part of or specific to any ISP—they operate in the root zone and are cooperatively managed by various organizations. *Id.*, 1, 3. Thus, in Glines’s architecture, a POSITA would have understood that “root DNS server 38,” not an “*Internet Service Provider (ISP) DNS Server*,” is programmed to resolve and “return the anycast address” of Glines’s name servers 53-1, 53-2, and 53-3. EX1002, 14:8-21.

56. Dr. Mowry further alleges that, to the extent Glines does not teach this element, “a POSITA also would have found it obvious to combine Glines and Gupta to cause a hostname to be resolved to a common anycast address by an Internet Service Provider (ISP) DNS server, as taught by Gupta.” EX1008, ¶165. In particular, Dr. Mowry contends that Gupta’s “client DNS server 170” provides the claimed “ISP DNS server” because “users 150 and 160 ... connect with ISP 60” and that “client DNS server 170 is typically provided within ISP 60.” EX1008, ¶166. Dr. Mowry then alleges that a “POSITA would have understood Gupta’s ISP

60 to be an obvious and known substitute for Glines’s local resolver 18” because it supposedly “performs the same functions as Glines’s local DNS resolver 18.”

EX1008, ¶167. Dr. Mowry, however, does not rely on Glines’s “local DNS resolver 18” for performing the resolution of the hostname to an anycast address, as I note above, so it is unclear how substituting Gupta’s ISP 60 for Glines’s “local DNS resolver 18” would meet the claim language. *See* EX1008, ¶¶154-155.

57. Indeed, Gupta’s “ISP 60” (and its included “client DNS server 170”) would not be “an obvious substitute for Glines’s local resolver 18” because Glines relies on a fundamentally different architecture than Gupta. Pet., 24. Glines’s “local resolver 18” functions as a cache, rather than DNS server. Specifically, Glines states that when “local resolver 18” receives a request from a browser, it “determines **whether or not the requested content file is available locally.**” EX1002, 13:24-14:1 (emphasis added). “For example, it determines if ‘example.com’ is located in the **local** web server 20.” EX1002, 14:1-2 (emphasis added). That is the purpose of Glines’s local resolver—to determine if the content is already stored locally, and if not, to reach out to a different root DNS server to perform DNS resolution. EX1002, 13:24-14:1.

58. Therefore, unlike Gupta’s “client DNS server 170” that is used “to resolve a host name into an IP address,” EX1003, 6:4-9, Glines’s “local resolver 18” is designed to determine whether requested content is available locally—not to

resolve host names to anycast addresses. EX1002, 13:24-14:1. This architecture allows Glines's "local resolver 18" to function as a cache server, rather than a DNS server, obviating the need to perform any resolution for locally available content beyond pointing to the local web server. I therefore disagree with Dr. Mowry's contention that "Gupta's ISP 60 [] performs the same functions as Glines's local DNS resolver 18," EX1008, ¶¶154-155, and I see no other reason in Dr. Mowry's declaration why a POSITA would have sought to combine Glines and Gupta in the way he suggests.

B. Dr. Mowry does not establish that the cited art teaches "a second request for second content associated with a second content provider" (Claim 23, Element 23[e])

59. Claim 23 requires performing actions in response to both "*a first request for first content associated with a first content provider of said plurality of content providers*" and "*a second request for second content associated with a second content provider of said plurality of content providers.*" EX1001, 12:33-49 (element [23c]-[23f]). The relevant limitations of claim 23 are reproduced below:

(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, 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

(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, 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.

EX1001, 12:33-49 (Claim 23, elements [23c]-[23f]) (emphasis added).

60. For element [23c], Dr. Mowry alleges that “it would have been obvious for Glines’s name servers to respond to an anycast message based on a first request for first content associated with a first content provider that includes a request for a news aggregator website, such as Yahoo! News (e.g., <http://dailynews.yahoo.com>)).” EX1008, ¶448. Dr. Mowry further contends that “[t]he 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>)).” EX1008, ¶450.

61. Then, for element [23e], Dr. Mowry alleges that “it would have been obvious for Glines’s name servers to respond to an anycast message based on a second request for second content associated with a second content provider that includes a request for a particular news article via a news aggregator website, such as Yahoo! News.” EX1008, ¶461. Specifically, Dr. Mowry identifies “a news

article” located at

“http://dailynews.yahoo.com/h/nm/20010201/ts/congress_ashcroft_dc_71.html,”

which I understand is illustrated in EX1025, as the claimed “*second content*.”


EX1008, ¶461. Dr. Mowry then contends that “[t]he hostname is associated with the first content provider (Yahoo! News), and the second content is associated with a second content provider (e.g., Reuters).” EX1008, ¶461.

62. But, from review of EX1024 and EX1025, Dr. Mowry does not identify a “*first content provider*” and a different “*second content provider*,” which is required by claim 23. EX1001, 12:24-48. Specifically, contrary to Dr. Mowry’s position, both the “*first content*” and “*second content*” identified by Dr. Mowry are from the **same** “*content provider*”: Yahoo News!

63. Initially, I note that a POSITA would have understood the plain language of claim 23 to require a request for “*first content*” associated with “*a first content provider*” and a request for different “*second content*” associated with a different “*second content provider*.” EX1001, 12:33-49. Indeed, I understand that Dr. Mowry agrees that this is the natural interpretation. Specifically, when asked about whether claim 23 “would be interpreted to require the first content to be different than the second content,” I understand that Dr. Mowry agreed that “the most natural interpretation, I think, would be that this is different content.” EX2021, 86:3-21.

64. Aside from claim 23 itself, a comparison to claim 13 makes clear that claim 23 requires two different content providers. For example, claim 13 recites “*a first request for first content associated with a first content provider of said plurality of content providers*” and “*a second request **for said first content associated with said first content provider** of said plurality of content providers.*” EX1001, 9:66-10:2, 10:15-17 (emphasis added). That is, claim 13 requires the first and second requests to be for the same content from the same content provider. Claim 23, on the other hand, specifically recites two different content providers: “*a first content provider*” and “*a second content provider.*” EX1001, 12:33-49.

65. In its mapping to Exhibits 1024 and 1025, Dr. Mowry incorrectly interprets the “*content*” of Exhibit 1024 as being associated with “Yahoo News!” but the “*content*” associated with Exhibit 1025 as being associated with “Reuters.” EX1008, ¶461. I understand that during deposition, Dr. Mowry clarified that the alleged “*second content*” is the Web page located at “http://dailynews.yahoo.com/h/nm/20010201/ts/congress_ashcroft_dc_71.html,” provided by Yahoo! News, not simply the article within the page. EX2021, 92:20-93:15 (referencing EX1008, ¶461).



EX1024

EX1025

EX1024 (annotated); EX1025 (annotated).

66. As shown above, a POSITA would have understood that the two URLs provided by Dr. Mowry are provided by the same content provider: Yahoo News!. In particular, a POSITA would have understood that while the second URL provided by Dr. Mowry appears to display an article written by Reuters, it is stored and provided by Yahoo! News, which is the same “*content provider*” that provides the first URL described by Dr. Mowry. Therefore, it is my understanding that Yahoo! News is the “*content provider*” of both Web pages, even if the article may have been written by Reuters. Such an interpretation is consistent with the description of “content provider” disclosed in the ’112 patent, which states that “[c]ontent provider” refers to an entity having a Web site on a network.” EX1001,

2:60-61. In the case of Dr. Mowry’s mapping to Exhibits 1024 and 1025, a POSITA would have understood that Yahoo! News is the entity having the Web site on the network—not Reuters. In other words, Dr. Mowry provides no explanation as to how Reuters would be the “*content provider*” for the Web page—not simply the article displayed within the Web page—located at http://dailynews.yahoo.com/h/nm/20010201/ts/congress_ashcroft_dc_71.html. For at least these reasons, it is my opinion that Dr. Mowry fails to establish that the cited art renders obvious “*a second request for second content associated with a second content provider of said plurality of content providers,*” as required by element [23e] of claim 23.

67. Accordingly, Dr. Mowry fails to establish that the cited art renders obvious element [23e] of claim 23.

C. Dr. Mowry does not establish that the cited art renders obvious a system configured to support “a plurality of content providers” (Independent claims 13, 18, 22, and 23)

Claims 13, 18, 22, and 23 require a “*plurality of content providers.*” EX1001, claim 13 (element [13b] and [13c]), claim 18 (elements [18p] and [18b]), claim 22 (element [22b] and [22e]), claim 23 (element [23c] and [23e]). For example, independent claim 13 recites: “*(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.” EX1001, 9:61-65 (element [13b])

(emphasis added). Independent claims 18, 22, and 23 recite similar features.

1. Obviousness in view of Glines

68. From Dr. Mowry’s analysis, it is unclear how Glines teaches support for “*a plurality of content providers*,” as required by claims 13, 18, 22, and 23. Dr. Mowry appears to address this limitation in his analysis of element [13p], stating that “Glines expressly discloses that its system may be used with either a ‘single provider’ or ‘multiple providers.’” EX1008, ¶260. However, a POSITA would have understood that the cited portions of Glines refer to network providers, not the claimed “*content providers*.” A POSITA would have understood that network providers are “local, regional and centralized Internet Service Providers (ISPs),” while content providers are “an entity having a Web site on a network.” EX1001, 2:60-63, 4:56-67. The full paragraph of Glines cited by Dr. Mowry is reproduced below:

The selection of routing protocol may have profound effects on the propagation [*sic*] and convergence of group membership changes. “Membership” in the group is contingent [*sic*] upon distributed routing state. In the case of deployment **within a single provider, where the anycast routing is internal to that network** (and transparent to the outside - the Internet), and an internal routing protocol like iBGP or OSPF is used, propagation [*sic*] of changes should be fast. In the case of deployment across **multiple providers**, full fledged external BGP

(“eBGP”) preferably would be used. Membership changes would be effected in such a scenario when the state change propagates at least half the way towards each other member in the anycast group, and that should cover all of external routers which would tag a particular anycast advertisement as “closest.” State changes would possibly take longer, but would not require flooding over the entire Internet.

EX1002, 17:3-14 (emphasis added).

69. Here, a POSITA would have understood that Glines’s “provider” refers to a network provider, not the claimed “*content provider*.” That is clear because the passage discloses, for example: (1) “deployment **within a single provider, where the anycast routing is internal to that network**” (e.g., deployment on a network); (2) “deployment **across multiple providers**,” (e.g., deployment across multiple networks; and (3) routing protocols (e.g., how messages are communicated within and across networks). EX1002, 17:3-14 (emphasis added). In other words, Glines specifically refers to “that network” of a particular network “provider.” EX1002, 17:3-14. Therefore, a POSITA would have understood that this passage of Glines does not in any way teach support for “*a plurality of content providers*,” but rather is discussing multiple network providers.

70. Additionally, in discussing element [13p], Dr. Mowry alleges that 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)” because “[a]dding additional 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.”

EX1008, ¶¶262-63. I understand that, in its Institution Decision, the Board explained, “[t]hat statement [] is unsupported, and moreover, does not explain how such replication would occur or what results would be expected.” DI, 14-15. I agree.

71. Dr. Mowry provides no explanation of how Glines’s teaching of storing and providing content associated with one origin server would be replicated. For example, it is unclear if Dr. Mowry is suggesting that Glines’s existing DNS servers and content servers would need to be replicated to account for additional content providers, how the data would need to be stored at the various content servers, or whether configuration updates to the DNS servers would be needed to account for additional content providers. At bottom, I do not understand what changes Dr. Mowry is suggesting need to be made and what results he expects. Of note, Dr. Mowry’s citations to Glines in support refer only to Glines’s existing “replica content servers” that “store replicas of one or more of the content files that originate at the” single “origin server 40”—not to multiple origin servers or content providers. EX1008, ¶262 (citing EX1002, 12:11-12).

72. For these reasons, it is my opinion that Glines does not render obvious independent claims 13, 18, 22, and 23.

2. Obviousness in view of Glines and Gupta

73. Dr. Mowry alleges that it also would have been obvious to combine Glines and Gupta to arrive at “a plurality of content providers” because Gupta allegedly discloses that *customers* (plural) can store content on the network and its “POP server network may provide traffic management and content hosting ‘to *more than one specific customer,*’ including ‘customer1 and customer2’ that share the server network.” EX1008, ¶265-66 (emphasis in original). Dr. Mowry does not provide any explanation as to how the combination of Glines and Gupta would operate or why a POSITA would have expected a reasonable expectation of success in combining the teachings of the references, as the only reasoning provided by Dr. Mowry is that “[t]he combination would have been nothing more than combining known elements in the prior art according to understood principles, with an expectation of success.” EX1008, ¶267.

74. I do not see any support for Dr. Mowry’s assertions that the combined system of Glines and Gupta establishes support for multiple content providers. For example, in his section discussing the motivation to combine Glines and Gupta, Dr. Mowry asserts that “applying Gupta’s use of multiple WebCache servers with each name server to Glines’s system to yield multiple content servers associated with

each name server in Glines,” and that “doing so would have allowed more content to be stored in association with each name server, improving Glines’s ability to deliver content to a user via the closest name server.” EX1008, ¶73 (emphasis in original). However, Dr. Mowry does not explain how Glines’s system would support multiple content providers simply by increasing storage capacity.

75. A POSITA would have understood that Glines already discloses the use of multiple file servers with each name server, for example as illustrated by file servers 51-3 of Figure 1. EX1002, 12:11-23; FIG. 1. Glines explains that the file servers of Glines contain “replicas of content files that originate at the origin server 40 [and] are distributed and stored in the replica content servers 50.” EX1002, 12:14-16; EX1001, 2:60-61. Considering this already existing architecture of Glines, Dr. Mowry does not explain how merely adding more storage capacity facilitates support for multiple content providers, as required by claims 13, 18, 22, and 23, nor does he even contend that it does.

76. For these reasons, as well as reasons I discuss further in Section IX.E, it is my opinion that Dr. Mowry does not establish that a POSITA would have combined Glines and Gupta to yield a system configured to support “*a plurality of content providers*,” as required by claims 13, 18, 22, and 23.


D. Dr. Mowry does not establish that the cited art teaches the “at least one switch” limitations (Independent claims 13, 18, 19, 21, and 22 and dependent claim 11)

77. Claims 11, 13, 18, 19, 21, and 22 recite “*at least one switch*” associated with a DNS server or set of content servers. *See* EX1001, claim 13 (element [13e]), claim 18 (element [18f]), claim 19 (element [19f]), claim 21 (element [21e]), claim 22 (element [22g]). The claimed “*at least one switch*” selects a content server from the plurality of content servers—e.g., claim 19 recites, “*by at least one switch, selecting the content server from said respective plurality of content servers associated with said resolving DNS server.*” It is my opinion that Dr. Mowry did not provide sufficient evidence that Glines, alone or in combination with Gupta, teaches this limitation.

78. Specifically, Dr. Mowry alleges that “Glines teaches or at least suggests that its Layer 4 switch is associated with its name server/content server peer groups, and therefore satisfies a switch ‘associated with said resolving DNS server.’” EX1008, ¶287. In Dr. Mowry’s analysis of claim 11, he asserts that “Glines’s switch is used for the load balancing and determines which content server should be selected.” EX1008, ¶249.

79. However, it is my opinion that Dr. Mowry does not provide sufficient evidence of teachings of Glines with respect to any function performed by the layer 4 switch. Glines mentions a switch in the context of service groups (as shown in

annotated Figure 1 below), stating only: “However, within the service group,
standard load balancing and replication techniques can also apply, such as multiple
content servers (returning multiple local IP addresses to name lookup), **layer 4
switches**, etc.” EX1002, 19:13-15 (emphasis added).



EX1002, FIG. 1 (annotated).

80. Thus, Glines, therefore, generally refers to a Layer 4 switch, but Glines does not provide any detail of the functionality performed by the Layer 4 switch, let alone teach that it is used to “*select the content server from said respective plurality of content servers*” as the claims require. In contrast, the ‘112 patent specifically provides a switch that uses “content-aware selection” to select an appropriate server “based on a predetermined selection policy.” EX1001, 8:52-

58.

81. Glines instead discloses that “standard load balancing” can be accomplished by “returning multiple local IP addresses to a name lookup” that correspond to “multiple content servers,” rather than the switch selecting the appropriate content server as described in the ‘112 patent. EX1002, 19:13-16. In Glines, the DNS server returns the IP addresses for the content servers to the browser, which then requests content from one of the servers. *Id.*, 15:26-16:8 (“[T]he name server 53 associated with this closest rate will then respond by reporting the unique IP address of its associated replica content server 51.”). With such a technique in place, it is my opinion that Glines has no need to cause the disclosed Layer 4 switch to “*select*” the content server, as the claims require.

82. Additionally, Dr. Mowry does not state how or why it would have been obvious to a POSITA to have Glines’s layer 4 switch perform this function. Indeed, Dr. Mowry simply asserts that “Glines’s applying ‘standard load balancing’ techniques within a service group, including use of ‘layer 4 switches,’ discloses or at least renders obvious using a Layer 4 switch to load balance across the content servers in a service group.” EX1008, ¶252. But, as I just discussed, this overreads the passage of Glines excerpted above, which does not state the role of the Layer 4 switch in performing the “standard load balancing and replication techniques.” EX1002, 19:13-15.

Thus, it is my opinion that Dr. Mowry fails to demonstrate that independent claims 13, 18, 19, 21, and 22 and dependent claim 11 are obvious over Glines.

E. A POSITA would not have been motivated to combine Glines and Gupta in the manner suggested by Dr. Mowry for any Ground.

83. Dr. Mowry includes a section in his Declaration titled “Motivation to Combine Glines and Gupta.” EX1008, ¶¶68-74. In my opinion, a POSITA would not have sought to combine Glines and Gupta for the reasons stated by Dr. Mowry. Specifically, Dr. Mowry alleges that a POSITA would have combined Glines and Gupta because “[b]oth Glines and Gupta disclose, for example, resolving DNS requests and identifying the most efficient server from which to serve content.” EX1008, ¶70. Dr. Mowry also alleges that “Glines and Gupta also share a similar architecture to one another—both disclose DNS servers and associated content servers that are connected by a network and function together to resolve hostnames and deliver content.” EX1008, ¶71.

84. In my opinion, Dr. Mowry does not establish a motivation to combine Glines and Gupta. Although both references discuss DNS resolution, any architectural similarity ends there. Glines and Gupta employ vastly disparate operational techniques, and the simple fact that these architectures separately existed in the prior art does not explain why a POSITA would be motivated to combine Glines and Gupta.

85. Dr. Mowry vastly oversimplifies the architectures of Glines and Gupta

in an attempt to establish a motivation to combine. A POSITA would have recognized, in fact, that Glines and Gupta disclose fundamentally disparate architectures and techniques for accomplishing content delivery. For example, Glines uses anycast addressing to direct requests to one of multiple “alternative or replica name servers 53-1, 53-2, 53-3” that “are peered with content file servers 51-1, 51-2, 51-3 that contain replica copies of files stored at the origin server 40,” as illustrated in Figure 1 below. EX1002, 10:19-25, 13:3-8.





FIG. 1

EX1002, FIG. 1 (annotated).

86. This architecture allows requests to be routed “to the closest available DNS server,” and, in turn, to the file server 51 located near the DNS server. EX1002, 4:30-5:1, 13:3-8 (“Each name server 53 and file server 51 associated with

particular content server 50 are considered to be connected in a peering arrangement. That is, they operate quite closely together and, in fact, are preferably located physically near one another ...”). In other words, when the name server receives a DNS request at its anycast address, it responds with the IP address of its peered content server. EX1002, 15:6-16:2.

87. In contrast to Glines’s anycast addressing and peering arrangement between DNS servers and file servers, Gupta uses only unicast addressing schemes, offloading its content server selection to a separate “network probe server” coupled to the DNS server. EX1003, Abstract. Gupta discloses a “client DNS server 170 [blue below] ... used to resolve a host name into an IP address.” EX1003, 6:4-6. “If client DNS server 170 does not know the IP address, client DNS server 170 may pass the request to other DNS servers, such as DNS servers 110-120 [red below], or the like.” EX1003, 6:6-9. That is, rather than sending an anycast UDP message like Glines, Gupta’s DNS servers are directly contacted via unicast.



EX1003, FIG. 1 (annotated).

88. Specifically, in Gupta’s architecture, a “local DNS server (LDNS)” “resolves the IP address of the domain name the user’s computer has queried for by sending a query to the Authoritative DNS server (ADNS) for that domain name.” EX1003, 2:30-33. The ADNS is “coupled to [a] first network probe server,” which is “configured to determine traffic loads of a plurality of customer web servers” that “stor[e] a web page.” EX1003, Abstract, 3:26-33.

89. To select the appropriate web server, the network probe server determines the “customer web server having a traffic load lower than traffic loads of remaining customer web servers from the plurality of customer web servers” and “provide[s] the IP address of the customer Web Server to the user.” EX1003, 3:35-42; *see also* EX1003, 9:1-5 (“NameServer server 110 typically uses the latency and traffic load information associated with the customer’s web sites that was determined in step 390, to select an appropriate customer web server to return to the user, step 440.”). For example, when “NameServer server 110 determines that the west coast customer web server 45 is very heavily loaded with requests, but the east coast customer web server 40 has light traffic,” the user’s request is directed to the east coast web server. EX1003, 9:5-10. Thus, rather than selecting a content server paired with the DNS server, as disclosed in Glines, Gupta uses service metrics such as traffic load to select a content server.

90. Put simply, Glines uses anycast addressing to select an appropriate replica DNS server and peered file server based on proximity to the user, while Gupta uses a different technique that offloads the selection of a web server to a network probe server coupled to the DNS server and selects the web server using standard load-balancing techniques. Glines discloses that its stated goal is to provide “a solution to the problem of route discovery by avoiding it altogether.” EX1002, 8:14-15. Gupta, in contrast, “uses the latency and traffic load information

associated with the customer's web site[] ... to select an appropriate customer web server to return to the user." EX1003, 9:1-5. Gupta's use of latency and traffic information directly contravenes Glines's stated goal because it involves route discovery—leveraging latency and traffic information to route a network packet requires knowledge of the route because it routes the packet to a server with low latency. Thus, if Glines were modified to account for Gupta's latency and load information, any efficiencies gained by using anycast addressing would be diminished because the routing system would need to have awareness of latency and load information at each DNS server, and route DNS requests based on the latency and load information. Conversely, Gupta teaches away from the use of anycast addressing, relying instead on load monitoring and load balancing. A POSITA implementing Glines may well have regarded Gupta's abandonment of anycast addressing as an improvement over Glines's architecture, supporting load leveling across a network of servers rather than relying on network proximity.

91. Therefore, it is my opinion that a POSITA would not have had reason to replace or augment Glines's anycast addressing and peered DNS and file server techniques with Gupta's disparate load allocation techniques. As I discuss above, a POSITA would have recognized that Glines and Gupta disclose fundamentally different architectures and techniques for delivering content. Glines uses anycast addressing to select an appropriate replica DNS server and peered file server based

on proximity to the user. In contrast, Gupta offloads the selection of a web server to a network probe server coupled to the DNS server and selects the web server using standard load balancing techniques. For these reasons, it is my opinion that any combination could only be made using hindsight.

X. CONCLUSION

In signing this declaration, I recognize that the declaration will be filed as evidence in an *inter partes* review before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required, I will appear for cross-examination within the United States during the time allotted.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Executed on this 11th day of February, 2026, in Mobile, AL.



Samuel H. Russ, PhD