

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

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

**SUPPLEMENTAL DECLARATION OF TODD C. MOWRY, PH.D.
IN SUPPORT OF PETITION FOR INTER PARTES REVIEW OF
U.S. PATENT NO. 9,021,112**

TABLE OF CONTENTS

I.	Introduction	1
II.	A POSITA Would Have Interpreted Claim 1’s “causing” Resolution Consistent With Its Plain Meaning, the Written Description, and Known DNS Systems.....	1
III.	A POSITA Would Have Understood Glines Teaches Delivering Content on Behalf of a “plurality” of Content Providers (Claims 13, 18, 22, and 23).....	11
IV.	POSITAs Were Aware of News Aggregator Websites Like Yahoo! News (Claim 23)	13
V.	Conclusion	15

I. Introduction

1. I, Todd Mowry, submit this supplemental declaration to state my opinions on the matters described below.

2. As explained in my opening declaration (EX-1008), I have been retained by Google LLC (“Petitioner”) as an independent expert in this proceeding involving U.S. Patent No. 9,021,112 (“the ’112 patent,” EX-1001), before the United States Patent and Trademark Office. I am being compensated at the rate of \$800 per hour for time that I spend on this case. No part of my compensation depends on the outcome, and I have no other interest in this proceeding.


3. My education, work experience, and other qualifications are summarized in my opening declaration (EX-1008) and curriculum vitae (EX-1009).

4. In connection with this supplemental declaration, I have reviewed Patent Owner’s Response (“POR,” Paper 25), Dr. Russ’s declaration in support of the POR (EX2018), as well as the exhibits cited in those papers.

II. A POSITA Would Have Interpreted Claim 1’s “causing” Resolution Consistent With Its Plain Meaning, the Written Description, and Known DNS Systems

5. Consistent with its plain language, a POSITA would have understood claim 1 recites a DNS resolution by query. That is, the claimed ISP DNS server “causes” resolution by sending a request to another server to perform the DNS lookup to the anycast address.

6. The written description of the '112 patent confirms this plain reading of claim 1. For example, as depicted below in connection with Figure 5, the '112 patent discloses that a user first sends a request to the **local DNS server 70a**—the claimed ISP DNS server—via DNS Req 91 to resolve the hostname. EX-1001 at 7:2-9, Fig. 5. The **local DNS server 70a** then sends a request in **DNS Req 93** to **ABCD DNS server 64** to resolve the hostname to the anycast address, which is returned via **DNS Resp 95**. EX-1001 at 7:2-9, Fig. 5.¹



Id. at Fig. 5 (annotated).

¹ The written description identifies “DNS Req 94” and “DNS Resp 96” while Figure 5 depicts “DNS Req 93” and “DNS Resp 95.” I understand these are the same DNS Req and DNS Resp and either the specification or the figure includes typographic errors.

7. The plain meaning of claim 1 and its corresponding written description are consistent with how DNS resolution is typically performed in the art to “cause” resolution of a hostname. As I explained in my opening declaration (EX-1008), local DNS resolvers—including those associated with Internet Service Providers (ISPs)—typically “act[] as an *intermediary*, performing recursive queries by contacting DNS servers in a hierarchical order: root servers, top-level domain (TLD) servers (e.g., .com), and finally the domain’s authoritative DNS server (e.g., the ‘example’ company server, often referred to as the origin server), which maintains the actual IP address of the requested domain name.” EX-1008 ¶ 36 (emphasis added) (citing EX-1011, 124-132; EX-1012, 622-630).

8. Materials I cite in my opening declaration support this understanding. For example, Kurose explains that “[w]hen a local name server [associated with an ISP] cannot immediately satisfy a query from a host (because it does not have a record for the hostname being requested), the local name server *behaves as a DNS client and queries* one of the root name servers.” EX-1011 at 128 (emphasis added)); *see also* EX-1014, 128-129 (explaining that local DNS servers associated with ISPs are “central to the DNS architecture” and that “[w]hen a host makes a DNS query, the query is sent to the local DNS server, *which acts a proxy, forwarding the query into the DNS server hierarchy.*”) (emphases added). The two editions of Kurose I cited in my opening declaration (EX-1011 and EX-1014) are widely-recognized

textbooks known to POSITAs and used to teach undergraduate courses in network computing.

9. Kurose provides examples of DNS queries typically used to resolve a hostname. One such query is a recursive query, in which “a host or name server A makes a . . . query to a name server B, then name server B obtains the requested mapping on behalf of A and then forwards the mapping to A.” EX-1011 at 132.

10. As shown in Figure 2.19, a requesting host “first sends a DNS query message to its local name server,” and “[t]he local name server *forwards the query message* to a root name server.” *Id.* at 129-130, Fig. 2.19. “The root name server forwards the query message to the name server that is authoritative for all the hosts in the [requested] domain.” *Id.*

11. Figure 2.20 depicts another example in which multiple “DNS messages can be sent in order to translate a single hostname—there can be two or more intermediate name servers in the chain between the root name server and the authoritative name server.” *Id.* at 130-131, Fig. 2.20.

12. In yet another example, shown in Figure 2.21, Kurose explains an “iterative” DNS query. *Id.* at 132. When a name server makes an iterative query to another name server, if the queried name server “does not have the requested mapping, it immediately sends a DNS reply to [the requesting name server] that

contains the IP address of the next name server in the chain” for the original name server to query next. *Id.* at 132, Fig. 2.21.

13. Thus, there may be several queries and intermediate name servers involved in the typical DNS resolution process, as depicted below.




Figure 2.19 + Recursive queries to obtain the mapping for `gaia.cs.umass.edu`

Id. at Fig. 2.19 (depicting recursive queries).




Figure 2.20 + Recursive queries with an intermediate name server between the root and authoritative name servers

Id. at Fig. 2.20 (depicting recursive queries with intermediate DNS servers)




Figure 2.21 + A query chain with recursive and iterative queries

Id. at Fig. 2.21 (depicting DNS resolution with iterative queries).

14. Tanenbaum, another well-known textbook on network computing, confirms this known DNS-by-query process and discusses recursive queries. For example, “[w]hen a resolver has a query about a domain name, it passes the query to one of the local name servers,” which either returns the authoritative resource records or “*sends a query message* to the top-level name server for the domain requested” if the requested information is not available locally. EX-1012 at 628-630 (emphasis added); *see also id.* at 630 (explaining that “the query method described here is known as a *recursive query*, since each server that does not have the requested information goes and finds it somewhere, then reports back.”) (emphasis added).

15. Prior art patents cited in my opening declaration also describe this query-based DNS resolution process by ISP DNS servers. For example, Jungck discloses that when a local ISP DNS server (e.g., DNS A 204) “does not know how to translate the given domain name in the request or the translation in its routing table has expired, it must go up the DNS hierarchy to complete the translation.” EX-1013 at 11:6-26, Fig. 2. To do so, the ISP DNS server will “*forward a request*, labeled ‘C’, upstream to the top-level DNS translation server 202 associated with the top-level domain in the domain address, i.e. com, gov, edu etc.” *Id.* (emphasis added).

16. In another example, Logan discloses a “local domain name server” that “checks to see if it already knows the IP-addresses for the hosts that service particular domain name and host.” EX-1016 at 3:49-4:15. “It could know this by having previously needing this information and storing the answer it discovered in a local private cache memory.” *Id.* “If the local domain name server does not know the hostname IP-address for a requested URL domain name, *it will perform an iterative query* to a domain name server higher in the DNS hierarchy.” *Id.* (emphasis added). That domain name server query will then be “answered by a higher level domain name server.” *Id.*

17. A POSITA would have recognized that Kurose, Tanenbaum, Jungck, and Logan describe the typical, query-based DNS resolution process performed by ISP DNS servers in the art to “cause” resolution of hostname, consistent with resolution-by-query described in the ’112 patent and Glines. This typically involved the local ISP DNS server sending a request to another server to resolve a hostname if the answer is not available locally. Indeed, this was common practice in the art as of the filing date of the ’112 patent.

18. Caching can occur at every level in the DNS hierarchy, so it was typical for the ISP DNS server to first check locally for requested content. EX-1008 ¶¶ 36-40. If the content is not available or is stale (e.g., its Time-to-Live has expired), the ISP DNS server would query an authoritative DNS server for the requested


content. *Id.* ¶ 36-38. The authoritative DNS server would then “cause” resolution of the hostname by the techniques described above, in the ’112 patent, and in Glines, to return an IP address of the requested content. *Id.*

19. In my opinion, Glines’s local DNS resolver 18 sending a request to root DNS server 38 for resolution of the hostname to the anycast address discloses claim 1’s “causing” resolution of the hostname, consistent with the plain meaning of the claim, the ’112 patent’s written description, and a POSITA’s understanding of how query-based DNS resolution typically occurred in the art.

20. For example, as I explained in my opening declaration, Glines discloses that when a user enters a URL in a web browser, “[i]n the case where the user is associated with an Internet Service Provider (ISP) operating the local area network 20, this first name request is made to a local DNS resolver 18, to determine the location of the domain.” EX-1008 ¶ 143 (citing Glines at 13:21-14:3). The local DNS resolver 18 is the “Internet Service Provider (ISP) DNS server,” as claimed. *Id.* ¶¶ 148-156. Glines further discloses that, “[i]f the ‘example.com’ domain is not located in a local domain, the local resolver will then proceed **to consult** public name servers that have been defined as being authoritative for the ‘example.com’ domain.” *Id.* ¶ 144 (citing Glines 7:1-7) (emphasis added). Then, “[i]n response, 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.’” *Id.* ¶ 144 (citing Glines at 14:4-15:4, 7:1-7, 9:9-15); *see also id.* ¶¶ 154-156.

21. This process is depicted in annotated Figure 1 of Glines below. The user sends a request to resolve a hostname to **local DNS resolver 18**—the claimed ISP DNS server. Glines at 13:15-14:5; EX-1008 ¶¶ 141-146. If the content is unavailable locally, **local DNS resolver 18** “consults” **root DNS server 38** by sending it a request to resolve the domain. Glines at 14:8-25; EX-1008 ¶¶ 143-144, 154-156. In response, **root DNS server 38** is programmed to return an anycast address to the **local DNS resolver 18**. *Id.*



Id. at Fig. 1 (annotated and cropped).

22. In my opinion, Glines’s resolution-by-request is indistinguishable from claim 1’s “causing” resolution of the hostname to the anycast address based on how a POSITA would understand the plain meaning of this term.

23. Claim 1 therefore would have been obvious over Glines and Glines in view of Gupta.

III. A POSITA Would Have Understood Glines Teaches Delivering Content on Behalf of a “plurality” of Content Providers (Claims 13, 18, 22, and 23)

24. As I stated in my opening declaration, 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). EX-1008 ¶ 262 (citing Glines at 12:11-22, 13:21-14:3 (requesting “example.com”)). This is clear from the description of Glines’s system and its disclosure of delivering content on behalf of exemplary content providers in “www.example.com.” Glines at 13:21-14:3. Indeed, Glines places no limitation on the source of its content or how many content providers its system services. *See id.*

25. A POSITA instead would have understood that content delivery networks like those of Glines were designed and intended to store and deliver content on behalf of one or more content providers because that is their business model. The prior art cited in my opening declaration confirms this. Swildens—a patent assigned on its fact to Akamai, a well-known and exemplary CDN provider in the art—

discloses what a POSITA would have understood at the time of CDN networks like Glines: that a “customer” (e.g., with a website “www.customer.com”) of a CDN typically pays the CDN provider a “service fee” to use its network. EX-1018 at 3:46-4:60. This “allows a customer to scale up its Web site, without a need for expensive and difficult to use hardware and/or software.” *Id.*; *see also* EX-1011 at 126 (referencing Akamai as a known DNS provider).

26. Glines itself further confirms a POSITA’s understanding that its system delivers content on behalf of one or more content providers. For example, Glines discloses that “[e]ach name server [] advertises itself as being authoritative for *the domains* associated with the content files stored [on] (or to be served from) the associated content server.” Glines at 6:11-15 (emphasis added). To do so, each name server 53 in Glines may be “an authoritative DNS resolver *for domain names* associated with the replica content files stored in its associated replica content server 51.” *Id.* at 13:11-13 (emphasis added).

27. By delivering content on behalf of multiple “domains,” Glines’s disclosure supports my testimony that a POSITA would have understood that Glines serves one or more content providers. *See* EX-1008 ¶ 262. At a minimum, Glines renders this feature obvious, as I explained in my opening declaration. *See id.* ¶ 263. So does Glines in view of Gupta. *See id.* ¶¶ 264-267.

28. Glines, and Glines in view of Gupta, therefore render this feature obvious in claims 13, 18, 22, and 23.

IV. POSITAs Were Aware of News Aggregator Websites Like Yahoo! News (Claim 23)

29. POSITAs were aware of news aggregator websites like Yahoo! News and news sources like Reuters as of October 18, 2001, the claimed priority date of the '112 patent.

30. As was known, Yahoo! News was launched in 1996 as an “aggregated news page” and by 2001 was in widespread use. EX-1042 (Yahoo! company timeline stating that in August 1996 “Yahoo! News is birthed by aggregated news page created by Yahoo! co-founder”); EX-1046 (NY Times discussing Yahoo! News). Yahoo! News “cull[ed] material from news organizations in exchange for a portion of Yahoo’s ad revenues” and “include[d] hyperlinking to external information sources.” EX-1046 at 3; EX-1043 at 122.

31. In addition to its widespread use, Yahoo! News was also described in Yahoo! patents before the '112 patent. For example, Nazem (EX-1044) discloses a customizable Yahoo! front page that includes links to major news stories. EX-1044 at 2:52-64 (“FIG. 1 shows a client-server system 100 which is used to display custom news pages.”), 5:50-6:20 (describing Yahoo! News “user front page 218”), Fig. 5A. “Since each of the news stories is essentially a static page linked to a headline shown

in the news section, these can simply be linked to static pages on a news server.” *Id.*
at 5:50-6:20.




FIG. 5A.

Id. at Fig. 5A.

32. POSITAs were therefore thus well-aware of news aggregator sites like Yahoo! News before the '112 patent, as I explained in my opening declaration. *See* EX-1008 ¶ 449. The Yahoo! News website I cite in my analysis of claim 23 is thus representative of known content available over the Internet at the time. *See id.* It would have been obvious for Glines's system to have delivered this content as Glines represents as "www.example.com." *See* Glines at 6:24-7:13, 13:21-15:10. Indeed, a POSITA would have understood that this was a placeholder for specific content in Glines that would have represented known hostnames at the time. *See* Glines at

6:24-7:13 (describing a search for placeholder www.example.com), 13:21-15:10 (same); EX-1008 ¶ 451 (“Resolving this particular content request is indistinguishable from Glines’s exemplary request for content via the URL “http://www.example.com/homepage.html.”) (citing Glines at 6:23-7:24, 13:15-15:10).

33. As I explained, a POSITA would have found that delivering this exemplary content nothing more than an obvious application of Glines’s system and a combination of known content available at the time. *See id.* ¶¶ 450-451. Among other things, it would have been obvious for Glines’s name servers to respond with the anycast address based on requests for content associated with different content providers known at the time. *Id.* ¶¶ 448-453, 460-465. Content at the time includes content associated with content providers Yahoo! News and Reuters. EX-1008 ¶¶ 446-453, 459-465.

34. Glines, and Glines in view of Gupta, therefore render these features obvious in claim 23.

V. Conclusion

35. For the reasons set forth above, and in my opening declaration, it is my opinion that claims 1-23 of the ’112 patent are unpatentable in view of the prior art.

36. In signing this declaration, I understand that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the

United States Patent and Trademark Office. I acknowledge that I may be subject to cross-examination in this case and that cross-examination will take place within the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

37. I declare that all statements made herein of my knowledge are true, that all statements made on information and belief are believed to be true, and 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.

Dated: May 20, 2026

By: 
Todd C. Mowry, Ph.D.