

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC,
Petitioner

v.

SANDPIPER CDN, LLC,
Patent Owner

Case IPR2025-01010
U.S. Patent No. 10,057,322

**DECLARATION OF SAMUEL H. RUSS, PH.D.,
IN SUPPORT OF PATENT OWNER'S PRELIMINARY 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 EX2010
Google v. Sandpiper
IPR2025-01010

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	MY BACKGROUND AND QUALIFICATIONS	2
III.	LIST OF DOCUMENTS CONSIDERED.....	5
IV.	LEGAL UNDERSTANDING.....	6
	A. My Understanding of Claim Construction.....	6
	B. My Understanding of Obviousness	7
	C. My Understanding of a Person Having Ordinary Skill in the Art	9
V.	PERSON OF ORDINARY SKILL IN THE ART	10
VI.	OVERVIEW OF THE STATE OF THE ART	10
VII.	THE '322 PATENT	12
	A. Overview of the '322 Patent	12
	B. Overview of the Prosecution History of the '322 Patent.....	14
VIII.	OVERVIEW OF THE PRIMARY REFERENCES	17
	A. Verma	17
	B. Raciborski.....	19
	C. Lewin.....	20
IX.	DR. MOWRY'S PROPOSED OBVIOUSNESS GROUNDS	21
	A. Dr. Mowry does not establish that the cited art discloses or renders obvious the claimed "application programming interface (API)" (all Grounds 1A/1B, 2, 4; Claims 11 and 15)	21
	B. Dr. Mowry does not establish that Verma, alone or in combination with Lewin, renders obvious the "tag" limitations of independent claim 1 (Grounds 3-5).....	27
	C. Dr. Mowry does not establish that the cited art discloses or renders obvious the claimed "modified embedded resource" limitations (all Grounds, all independent claims).	32
X.	CONCLUSION	35

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 litigation. I have no other interest in this proceeding.

2. I understand that this declaration accompanies Patent Owner’s Preliminary Response filed in this IPR involving U.S. Patent No. 10,057,322 (“the ’322 patent”) (EX1001), which resulted from U.S. Application No. 14/985,968 (“the ’968 application”), filed December 31, 2015, which claims priority to U.S. Provisional Application No. 62/098,930 (“the ’930 provisional”), filed December 31, 2014. My opinions regarding the knowledge in the art at the time of the ’322 patent are based on this date of December 31, 2014. Although I express no opinion as to whether the ’322 patent is entitled to the priority date of the ’930 provisional, my opinions regarding the knowledge in the art would not materially differ if based on the date of December 31, 2015.

3. In preparing this declaration, I have reviewed the ’322 patent and each of the documents cited herein, in light of general knowledge in the art before the time of the ’322 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 ’322 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 EX2011, 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. 10,057,322

Exhibit No.	Description
1001	U.S. Patent No. 10,057,322 (“the ’322 Patent”)
1002	Prosecution History of the ’322 patent
1003	Declaration of Todd C. Mowry, Ph.D.
1004	<i>Curriculum Vitae</i> of Todd C. Mowry, Ph.D.
1005	U.S. Patent Application Publication No. 2009/0132640 to Verma (“Verma”)
1006	U.S. Patent Application Publication No. 2007/0067424 to Raciborski (“Raciborski”)

Exhibit No.	Description
1007	U.S. Patent Application Publication No. 2014/0052811 to Lewin (“Lewin”)
2001	“Cloud CDN: content delivery network,” <i>Google Cloud</i> , accessed at https://cloud.google.com/cdn (last accessed July 2, 2025)
2011	<i>Curriculum Vitae</i> of Samuel H. Russ, Ph.D.
2012	Michael Goodwin, “What is an API (application programming interface)?” accessed at https://www.ibm.com/think/topics/api (last accessed August 5, 2025)
2013	“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 (last accessed August 1, 2025)
2014	“What is a CDN (Content Delivery Network)?” <i>Amazon Web Services</i> , accessed at https://aws.amazon.com/what-is/cdn/ (last accessed August 1, 2025)

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 '322 patent.

C. My Understanding of a Person Having 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 '322 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 '322 patent.

V. PERSON OF ORDINARY SKILL IN THE ART

25. For purposes of this declaration, I have applied Dr. Mowry's articulation of a POSITA. EX1003, ¶62. It is my opinion that the definition of a POSITA would not materially affect the limited analysis required for deciding institution and presented in this declaration. If called to give further opinions in this proceeding, I reserve the right to reconsider the definition of a POSITA and arrive at my own conclusion.

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. EX2013, 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. EX2013, 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. EX2013, 3-6; EX2014, 2-3. This service and its architecture were quickly adopted by many others in the industry. EX2013, 3-6; EX2014, 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). EX2013, 2; EX2014, 2-3. Propagating data from origin servers to edge servers is known as "caching." EX2013, 3-6; EX2014, 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. EX2013, 3-6; EX2014, 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. EX2014, 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. EX2013, 3; EX2014, 1.

30. In the early-to-mid 2000s, demand for CDNs exploded. EX2013, 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.

VII. THE '322 PATENT

A. Overview of the '322 Patent

31. The '322 patent is directed to novel techniques for improving the performance of content delivery networks (CDNs), in some cases using an application programming interface (API) “to embed link resolution information within an HTML document or other content returned to a requesting client.” EX1001, 1:16-18. Aspects of the '322 patent “involve an alternative to using conventional domain name system (DNS) servers for edge selection by instead calling an API or otherwise querying the CDN when constructing a uniform resource locator (URL) for the resource to be delivered from the CDN.” EX1001, 3:10-14. Aspects of the '322 patent also involve “[a]n origin server [for] maintaining original copies of content for distribution to clients and/or distribution

of copies of content to other nodes.” EX1001, 4:15-18. The origin server is connected with a database that includes “the HTML document with an embedded resource,” in which the embedded resource may include “a tag indicating that a resource address compute engine should be contacted via a call to an API”—“to retrieve the address (only) or complete URL.” EX1001, 4:22-26, 32-34, 41-42. Specifically, the presence of the tagged embedded resource prompts the API call, the result of which “causes the HTML document returned to the requesting client to include a modified embedded resource with an address or URL to the CDN.” EX1001, 4:42-47. It is worth noting that, in the ’322 patent, an HTML document contains the modified URL. That is, embodiments of the ’322 patent expressly disclose modifying an HTML document in order to modify the URLs found inside the document. EX1001, 4:22-26, 32-34, 41-47. The solutions described in the ’322 patent were unconventional at the time of the invention of the ’322 patent and continue to offer numerous technical benefits.

32. For example, as users were able to access more content (e.g., television, movies, news, software updates) on networks, network providers sought alternatives to routing requests through DNS to select an edge server. EX1001, 1:39-50. The inventors of the ’322 patent devised techniques for “bypassing DNS,” thereby reducing the number of components involved in servicing requests while improving response speeds while delivering content. EX1001, 2:45-49, 5:29-32,

6:60-7:1. Additionally, when the origin server receives the request, the origin server calls an API in order to communicate with the compute engine to format the address for the embedded resource. EX1001, 4:49-52. In particular, the address is not static, instead, the address is optimally suited to serve the content to the requesting client. EX1001, 4:54-5:15. The solutions described in the '322 patent improve the functioning of a CDN, including, for example, offloading computational resource consumption away from a client-side DNS infrastructure, and instead having content retrieved from an optimal node in the network.

B. Overview of the Prosecution History of the '322 Patent

33. I understand that the '322 patent received a thorough prosecution on the merits, including overcoming both §§ 101 and 102 rejections. EX1002, 15-18, 60-74, 95-108.

34. After receiving a non-final rejection (*see* EX1002, 95-108), Patent Owner amended the independent claims to clarify the architecture of the claimed systems and methods. EX1002, 56-59. For example, I understand that Patent Owner amended independent claim 12 (which issued as independent claim 11) as follows:

12. (Currently Amended) A content delivery ~~device~~ network comprising:
at least one processor in communication with a non-transitory computer readable medium
including computer readable instructions for:

~~generating~~ providing an application programming interface (API);

at the API, receiving, from an origin system that is separate from the content
delivery network, an embedded resource and at least one parameter associated with the
embedded resource;

generating a modified embedded resource providing ~~information as to a location
within a~~ a direct link to a node within the content delivery network for obtaining content
associated with the embedded resource;

delivering, by the content delivery network, the modified embedded resource to
the origin system;

receiving, at the node, a request from a client device to obtain the content; and

providing the content to the client device.

EX1002, 57-58.

35. Patent Owner also added independent claim 22 (which issued as
independent claim 15) with similar limitations:

22. (New) A method comprising:

providing, at a content delivery network, an application programming interface (API);

at the API, receiving, from an origin system that is separate from the content delivery network, an embedded resource and at least one parameter associated with the embedded resource;

generating a modified embedded resource providing a direct link to a node within the content delivery network for obtaining content associated with the embedded resource;

delivering, by the content delivery network, the modified embedded resource to the origin system;

receiving, at the node, a request from a client device to obtain the content; and

providing the content to the client device.

EX1002, 59.

36. Patent Owner argued that the claimed architecture differed from the prior art (U.S. Patent Pub. No. 2014/0289319 to Richardson) because “Richardson requires the client to either (a) communicate with the DNS server of the CDN to determine the node that will be used to deliver the content; or (b) execute translation code to ask the CDN for translation information, which translation information may include the address of the node that will be used to deliver the content.” EX1002, 72-73. Based on these distinctions, Patent Owner argued that Richardson did not disclose “generating modified embedded resource providing a direct link to a node within the content delivery network for obtaining content associated with the embedded resource.” EX1002, 73.

37. I understand that the Examiner allowed the claims after considering

this and other arguments. EX1002, 15-18. The '322 patent issued on August 21, 2018. EX1001, (45).


VIII. OVERVIEW OF THE PRIMARY REFERENCES

38. I understand that Petitioner's expert, Dr. Mowry, relies on numerous references in his proposed obviousness grounds challenging the claims of the '322 patent. Three of those references are particularly relevant to my opinions provided in this declaration: Verma (EX1005), Raciborski (EX1006), and Lewin (EX1007). I provide brief overviews of these primary references below.

A. Verma

39. Verma discloses a “computer method and related apparatus for delivering content files to a client computer.” EX1005, Abstract, ¶¶10-11. The method includes modifying a resource link—based on client location and content availability—to direct requests to a cache server node in a CDN.

40. Verma's exemplary “content delivery network” is illustrated as shown below.




EX1005, FIG. 1.

41. In Verma, the content delivery network includes a client that runs a “browser 21,” a “content provider 22,” a “content delivery network broker 25,” and “a set of cache servers 23/24.” EX1005, FIG. 1, ¶¶22, 25. The client “browser 21 operates to request a content file . . . by sending a URL 28 to a content provider 22.” EX1005, ¶25, FIG. 1. “The request from the client” can “comprise delay tolerance information.” EX1005, ¶18. This request is then “forward[ed]” from content provider 22 to the content delivery network broker 25 that “manages at least one, and preferably, a set of cache servers.” EX1005, claim 9, ¶25. Also, content provider 22 can “dynamically request the broker 25 for a content file delivery to a location close to the client.” EX1005, ¶26. The broker sends availability inquiries to the cache servers to “find out whether they currently store

the requested” file. EX1005, ¶25. The broker 25 can select the appropriate cache server to serve the content and return a “re-written URL to the content server [that] received the [client] request.” EX1005, ¶¶12, 26, 27, 37. When the client receives the prepended (re-written) URL, “it requests the content file from the appropriate cache server when the content is available.” EX1005, ¶37.

B. Raciborski

42. Raciborski discloses a “content delivery network (CDN) for delivering content objects from a content provider to a recipient.” EX1006, Abstract. Raciborski’s exemplary network is illustrated as shown below.



EX1006, FIG. 1C.

43. Within this network, Raciborski discloses a “content provider” that

“sends various information to the one or more CDNs” in the network to facilitate content delivery. EX1006, ¶¶26, 55-56. The information can include the “IP address of the recipient computer,” the “full object location (i.e., server, path and name),” “if delivery can be delayed,” and a “protocol used to deliver the content object,” among other things. EX1006, ¶¶26, 55-56. This information may be included in a “bid message,” like the one illustrated in Table 1 as shown below. EX1006, ¶¶55-56, Table 1. Such “bid message” parameters, as illustrated in Table 1, provide the CDN with information it can use to select an optimal delivery path. EX1006, ¶¶55-56, Table 1.

TABLE I

<u>Real Time Bid Message Example</u>	
Field	Example
Bid Identifier	34sd89732k
Recipient IP Address	69.93.18.222
Domain of Content Object	LLNW.com
Path of Content Object	/press/2005releases/
File Name of Content Object	routemark.mp4
Format	MPEG4
Bitrate	200 Kbps
Latency Options	3
QoS Level(s)	3
Packet Loss Options	2
Buffer Duration	24 hours
Max. Volume	1000

EX1006, Table 1

C. Lewin

44. Lewin is generally directed to a “dynamic content assembly

mechanism” for constructing an “HTML page.” EX1007, ¶¶8, 12, 28, 49. As Lewin explains, “[t]raditionally,” “directives [] are placed in HTML pages” “that are evaluated on a server before the page is served.” EX1007, ¶28. These “directives” or “dynamic assembly instructions” may be “URLs for HTML chunks to be incorporated in[to] the final HTML page.” EX1007, ¶¶49, 58 (“dynamic content assembly page directives”). These “URLs” are “specif[ied]” by “tags” that can be identified when a server “parses the HTML page.” EX1007, ¶49. Of relevance, Lewin’s “tags,” such as the “<esi:include> tag” relied on by Petitioner, are used to instruct the browser to retrieve content at the URL—not to retrieve a link to the content.

IX. DR. MOWRY’S PROPOSED OBVIOUSNESS GROUNDS

45. 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 discloses or renders obvious the claimed “application programming interface (API)” (all Grounds 1A/1B, 2, 4; Claims 11 and 15)

46. Dr. Mowry relies on prior art that fails to teach or suggest the “application programming interface (API)” limitations of the claims. The unique API features described and claimed in the ’322 patent that allows the client to

communicate with an application executing on the server is not found in the prior art cited for Grounds 1A/1B, 2, 4; Claims 11 and 15.

47. Independent claims 11 and 15 recite various features involving an “application programming interface (API),” specifically:

[11b], [15a]: providing, at a content delivery network, **an application programming interface (API)**;

[11c], [15b]: **at the API**, receiving, from an origin system that is separate from the content delivery network, an embedded resource and at least one parameter associated with the embedded resource; ...

EX1001, 10:63-67 (Claim 11), 12:6-11 (Claim 15).

Dependent claims 2 and 9 also recite various features involving an “API,” specifically:

[2a]: passing information based on the client IP address to **an application programming interface (API)** of the content delivery network;

[2b]: **in association with the API**, obtaining a parameter associated with the embedded resource; ...

[9b]: passing, by the origin system, the location identifier to **an application programming interface (API)** of the content delivery network to generate the modified embedded resource relating to the node.

EX1001, 10:23-27 (Claim 2), 10:52-55 (Claim 9).

48. Dr. Mowry relies on Verma alone for the recited “application programming interface (API)” across all Grounds. EX1003, ¶83. Although Dr. Mowry further relies on Raciborski and Lewin for certain limitations of claims 2,

9, 11, and 15 in Grounds 2 and 4, Dr. Mowry does not rely on either reference for the recited “application programming interface (API).” EX1003, ¶164 (relying on Raciborski to allegedly render obvious “receiv[ing] the embedded resource and parameter associated with the embedded resource at Verma’s ‘broker’ ... to achieve Verma’s expressly stated goals”—not for the recited API), ¶233 (relying on Lewin to allegedly render obvious “using a delivery parameter to obtain, by the origin system, a modified embedded resource from the CDN ‘based upon identification of the tag.’”), ¶¶267-270, 290.

49. In particular, Dr. Mowry alleges that “a POSITA would have [] understood Verma’s ‘broker’ is the claimed ‘application programming interface (API)’ because it performs the same logical functions as the API in the ’322 patent.” EX1003, ¶87. Dr. Mowry continues, “the ’322 patent explains ‘[w]hen the origin server receives the request, . . . the origin server calls an API’ to ‘communicate with the compute engine’ in the content delivery network.” EX1003, ¶87 (citing EX1001, 4:49-52, 7:16-25). Verma’s “broker,” however, is not an API.

50. A POSITA would have understood an application programming interface (API) to be “a set of rules or protocols that enables software applications to communicate with each other to exchange data, features and functionality.” EX2012, 2. For example, “in terms of a request and response between a client and


server,” “[t]he API is the bridge establishing the connection between them.”

EX2012, 2. In other words, as its name implies, an API acts as an interface that allows the client to communicate with an application executing on the server—it is *not* the application itself. EX2012, 2. The term “application programming interface” makes this apparent—it is an interface to an application program, not the application itself. EX2012, 2.

51. In the context of the ’322 patent, the “API” acts as a bridge between the origin server and the compute engine. EX1001, 6:24-30 (“[T]he API 26 provides the interface to a computing engine 30 that uses the information, passed to the engine through the API...”), 4:32-34 (“[C]ompute engine 30 should be **contacted via a call to an API 26** (operation 220)” (emphasis added)). The ’322 patent further explains that “[w]hen the origin server receives the request (or at some time prior to receiving the request), the origin server **calls an API** in order to communicate with the compute engine 30 to format the address for the embedded resource.” EX1001, 4:49-52 (emphasis added). Then, the “computing engine 30,” not the API, “determine[s] one or more nodes 32 in the CDN to serve the content...” EX1001, 6:24-30.

52. Verma does not disclose or teach an API. In contrast, Verma describes “a content broker which receives information from the cache servers in a content delivery network as to traffic, availability of particular files and various

locations, and then returns the re-written URL to the content server which received the request from the client.” EX1005, ¶12. For example, as depicted in Verma’s FIG. 1, a “browser 21 operates to request a content file, for example the movie entitled ‘Matrix Reloaded,’ by sending a URL 28 to a content provider 22, Movies-Inc.Com in this example.” EX1005, ¶25.



EX1005, FIG. 1 (annotated).

53. As Verma explains, “Movies-Inc.Com content provider 22 sends a request 34 for information to a content delivery network broker 25, which manages at least one, and preferably, a set of cache servers, exemplified by cache server 23 and cache server 24.” EX1005, ¶25. The broker sends availability inquiries to the cache servers, respectively, “to find out whether they currently store the requested movie file.” EX1005, ¶25. Verma explains that “[t]he broker 25 provides content

current and/or future availability information to Movies-Inc.Com content provider 22 where the content file URLs are prepended (Akamai-ed) for redirecting to cache servers.” EX1005, ¶26. But nowhere does Verma disclose that its “content broker” is or makes use of an API to communicate with either the content provider or cache servers.

54. As Dr. Mowry acknowledges, the ’322 patent describes that an “API is a convenient and efficient way for disparate systems to uniformly and predictably **pass information to the computing engine and for the computing engine to return information.**” EX1003, ¶87 (emphasis added). Verma’s purported “API,” however, performs entirely different logical functions. For example, Verma explains “[i]f the content file [to be delivered] is not immediately available ... [Verma’s] content server will negotiate with the content delivery network (CDN) through the broker in order to determine the start time, end time, and the cache server where the content file will next be available during step 160 of FIG. 2.” EX1005, ¶38. Determining the start time, end time, and the availability of the content file within each cache server—which are the logical functions performed by Verma’s broker—is different than an API that merely provides an interface for passing and returning information. Instead, Verma’s determination requires that its broker “schedule the distribution of content files to edge [or cache] servers.” EX1005, ¶22. In other words, Verma has replaced a set of rules or

protocols for communication with an entirely separate software entity, the “broker,” that performs scheduling steps.

55. Therefore, it is my opinion that Verma does not teach that its broker is or provides the function of the claimed “application programming interface (API),” as required by claims 2, 9, 11, and 15. This element is wholly missing from Verma.


B. Dr. Mowry does not establish that Verma, alone or in combination with Lewin, renders obvious the “tag” limitations of independent claim 1 (Grounds 3-5).

56. Independent claim 1 recites, “**receiving**, by an origin system, a request from a client device for a content resource, the content resource including **at least one embedded resource with a tag**,” and “**upon identification of the tag**, using at least one delivery parameter to **obtain, by the origin system, a modified embedded resource** from a content delivery network that is separate from the origin system.” EX1001, 10:6-12 (elements [1a] and [1b], emphasis added). Claim 1 further requires that “the **modified embedded resource includes a direct link to a node within the content delivery network** that can deliver content associated with the embedded resource.” EX1001, 10:13-18 (element [1c], emphasis added).

57. Dr. Mowry acknowledges that “Verma does not expressly disclose its origin system receiving ‘a request from a client device for a content resource, the content resource including at least one embedded resource with a tag,’” instead

relying on Lewin for the claimed “embedded resource with a tag.” EX1003, ¶¶191-193. But Lewin does not supply this missing “tag” limitations.

58. Figure 1 of the '322 patent illustrates the claimed “embedded resource with a tag.” EX1001, FIG. 1.



EX1001, FIG. 1 (annotated).

59. As the '322 patent explains, “the origin server is connected with a database 22, which organizes and manages stored resources in some form of storage device (or devices), such as a storage appliance, that includes the HTML document 20 with an embedded resource 24.” EX1001, 4:22-26. “In this example, the embedded resource 24A is a link to some form of content that the user can

obtain by selecting the link when the Web page is displayed.” EX1001, 4:28-31.

60. The ’322 patent further explains that “[t]he embedded resource may include **a tag indicating that a resource address compute engine 30 should be contacted** via a call to an API 26.” EX1001, 4:32-34 (emphasis added). In an example, “the tag is ‘\$’ which causes the origin server 18 to call the API 26 to retrieve the address (only) or the complete URL,” and “[h]ence, **the presence of the tagged resource 24A** in the HTML document 20 at the origin servers **causes the origin server to call the API, the result of which causes the HTML document returned to the requesting client to include a modified embedded resource 24B.**” EX1001, 4:40-48 (emphasis added).

61. This operation is expressly recited in claim 1, requiring “at least one embedded resource with a tag,” and “**upon identification of the tag**, using at least one delivery parameter to **obtain, by the origin system, a modified embedded resource** from a content delivery network that is separate from the origin system.” EX1001, 10:6-12. In other words, the claim 1 requires that the recited “tag” be used to initiate the operation of obtaining “the modified embedded resource [that] includes a direct link to a node within the content delivery network that can deliver content associated with the embedded resource.” EX1001, 10:6-18.

62. Lewin’s purported “tag,” in contrast, is not used to obtain a “modified embedded resource [that] includes a direct link to a node within the content

delivery network that can deliver content associated with the embedded resource,” as required by claim 1. In particular, Dr. Mowry alleges that in Lewin, “‘directives [] are placed in HTML pages’ and ‘evaluated on a server before the page is served,’ which ‘enable[s] the server to add dynamically-generated content to an existing HTML page.’” EX1003, ¶197 (citing EX1007, ¶28). According to Dr. Mowry, the “‘dynamic assembly instructions’ may be ‘specif[ied]’ by ‘tags,’” for example, “the ‘<esi:include> tag.’” EX1003, ¶¶197-198 (citing EX1007, ¶¶28, 30, 49).

63. Dr. Mowry, however, does not explain how Lewin’s “<esi:include>” tags are used (alone or in combination with Verma) to “obtain ... a modified embedded resource” that “includes a direct link to a node within the content delivery network that can deliver content associated with the embedded resource,” as required by claim 1. *See* EX1003, ¶¶195-224, 233-237. Instead, as Dr. Mowry notes, Lewin’s tags instruct the server to add “content to an existing HTML page,” not a new link to the content. EX1003, ¶197 (citing EX1007, ¶28). In other words, Lewin’s tags are used to insert actual content into the page, rather than insert a link to the content that can then be used to retrieve the content.

64. Moreover, Dr. Mowry merely alleges that a “POSITA would have understood Lewin’s tags advantageously allow for specification and identification of URLs, and would have been motivated to incorporate those tags into Verma’s

system to provide a standard, uniform, and more efficient and uniform way to identify URLs,” and that such tags would “enable Verma’s system to utilize and identify advantageous embedded content.” EX1003, ¶¶214, 236. But Dr. Mowry’s analysis does not address the claim language, which requires obtaining a “modified embedded resource [that] includes a direct link to a node within the content delivery network that can deliver content associated with the embedded resource” upon “identification of the tag.” EX1001, 10:6-18. Put another way, that Lewin’s tags could be used to identify embedded content has no bearing on whether a POSITA would have used those tags to obtain a “modified embedded resource [that] includes a direct link to a node within the content delivery network that can deliver content associated with the embedded resource.” EX1001, 10:6-18.

65. Additionally, Dr. Mowry provides no reason why a POSITA would seek to modify Verma’s existing markup scheme. As Dr. Mowry acknowledges, Verma already includes standard HTML tags (e.g., “a href” tags) to identify URLs. EX1003, ¶223. Dr. Mowry does not explain how replacement of these tags with those disclosed by Lewin would improve Verma’s operation, or how they would even function with Verma’s existing content broker and cache servers. EX1003, ¶¶195-224, 233-237.

66. Therefore, it is my opinion that a POSITA would not have had reason to replace or augment Verma’s standard HTML tags with Lewin’s tag.

C. Dr. Mowry does not establish that the cited art discloses or renders obvious the claimed “modified embedded resource” limitations (all Grounds, all independent claims).

67. Independent claims 11 and 15 each recite, “generating **a modified embedded resource** providing a direct link to a node within the content delivery network for obtaining content associated with the embedded resource.” EX1001, 11:1-4 (element [11d]), 12:12-15 (element [15c]) (emphasis added). Claim 1 similarly recites, “obtain[ing], by the origin system, **a modified embedded resource** from a content delivery network.” EX1001, 10:9-12 (element [1b]) (emphasis added).

68. As the '322 patent describes, “when a user enters, into the browser Web page, an address <http://www.example.com>,” the origin server returns an “HTML document 20 with an embedded resource 24.” EX1001, 4:5-26. “The HTML document is what the browser uses to display the Web page [with an address] <http://www.example.com>” and “[t]he embedded resource 24A is a link to some form of content that the user can obtain by selecting the link when the Web page is displayed.” EX1001, 4:26-31. The '322 patent further explains that “the HTML document may include an embedded resource” with the tag “\$” which, in turn, “**causes the HTML document returned to the requesting client to include a modified embedded resource** 24B with an address or URL to the CDN whereby the content may be retrieved.” EX1001, 4:32-48 (emphasis added). Thus, as the

'322 patent describes and claims, the “modified **embedded** resource” is a resource embedded within another resource, such as an HTML document—the URL alone is not an *embedded* resource. *See, e.g.*, EX1001, 10:6-8, 56-58 (Claims 1 and 10: claiming a “content resource including at least one embedded resource, ... wherein the content resource is a hypertext markup language (HTML) document and the embedded resource is a link”), 11:10-16 (Claims 12 and 13 (“replacing the embedded resource with the modified embedded resource in a resource ... wherein the resource is an HTML document and the modified embedded resource is a link”).

69. Dr. Mowry relies on Verma for the recited “modified embedded resource” across all Grounds. EX1003, ¶¶118-120, 225-237. Specifically, Dr. Mowry alleges, the broker’s “‘generat[ed] uniform resource locator’ is a ‘prepended URL’ (or ‘rewritten URL’) because the ‘original URL (e.g., http://contentURL)’ is modified into an “‘Akamai-ed”’ URL’ identifying the ‘local_cache_server’ in the path ‘(e.g., http://local_cache_server/contentURL)’ ... [that] ‘redirects to a local cache server in the content delivery network.’” EX1003, ¶120 (citing EX1005, ¶¶12-13, 16, 26). Dr. Mowry fails to show, however, that Verma’s “‘prepended URL” or “‘rewritten URL” is a “modified **embedded** resource” as claimed—e.g., as a modified resource embedded within the client’s HTML page.

70. In contrast to the '322 patent, Verma does not disclose embedding its “prepended” or “rewritten URL” within an HTML document returned to the client. Rather, Verma discloses automatically “redirecting to the cache server” when the URL is received, without embedding the URL within any other resource. EX1005, ¶37. Verma describes that the client “browser 21 operates to request a content file, for example the movie entitled ‘Matrix Reloaded,’ by sending a URL 28 to a content provider 22.” EX1005, ¶25. Then, as Dr. Mowry recognizes, the “prepend URL redirecting to [a] cache server ... is returned to the client,” and “[w]hen the client receives the prepend URL, it requests the content file from the appropriate cache server.” EX1005, ¶37; EX1003, ¶128 (emphasis added). In fact, Dr. Mowry does not even allege that the “prepend URL” or “rewritten URL” is embedded within an HTML document or other resource. EX1003, ¶¶118-120. While Verma appears to describe altering a URL that is subsequently used to access content, Verma does not describe altering any resource, such as an HTML file, to embed the altered URL.


71. Therefore, although Verma appears to describe altering a URL, it is my opinion that Dr. Mowry does not establish that Verma discloses or renders obvious “obtain[ing]” or “generating a modified **embedded** resource,” as required by independent claims 1, 11, and 15.

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 September 15, 2025 in Mobile, AL.



Samuel H. Russ, Ph.D.