

UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE PATENT TRIAL AND APPEAL BOARD

NETNUT LTD.,

v.

LUMINATI NETWORKS LTD.,

Patent Owner

U.S. Patent No. 9,742,866

DECLARATION OF KEITH J. TERUYA

Mail Stop "PATENT BOARD"
Patent Trial and Appeal Board
U.S. Patent & Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

I, KEITH J. TERUYA, declare the following:

1. I have been retained by counsel for NetNut Ltd. to provide my opinions on certain issues in connection with claims 15-20, 23, 24, 27, and 28 of U.S. Patent No. 9,742,866 (the '866 patent).

2. I am being compensated at a normal hourly rate for my time in preparing this declaration, with reimbursement for actual expenses. My compensation is not tied to the outcome of this matter, and not based on the substance of the opinions rendered here.

3. In preparing my opinions expressed herein, I have reviewed the '866 patent and its prosecution history and all of the patents and printed publications listed in Exhibit 1010. I may rely upon these materials, my knowledge and experience, and/or additional materials to rebut arguments raised by the patent owner. Further, I may also consider additional documents and information in forming any necessary opinions, including documents that may not yet have been provided to me.

4. I am familiar with the technology described in the '866 patent as of its earliest claimed priority date of August 28, 2013. I have been asked to provide my technical review, analysis, insights, and opinions regarding the '866 patent. I have used my experience and insight along with the above-noted references as the basis

for my opinions, which support the grounds of unpatentability set forth in the Petition for *Inter Partes* Review of the '866 Patent.

QUALIFICATIONS

5. I have a long professional background in information technology and network engineering, and for the past 20 years I have been the chief executive officer of a specialized Internet hosting company that I founded, which among other things provides Tier I Network Operations Center capabilities for corporations, local municipalities, Federal Government programs, and regulatory agencies in and around southern California.

6. My day-to-day work in my present capacity involves direct hands-on as well as strategic involvement in the issues of networked data distribution and access addressed by the '866 patent, including without limitation architecting and configuring high-capacity content servers, proxy servers, content distribution networks (CDNs), edge and origin servers, peer-to-peer communications, as well as the lower-level routing and switching infrastructure and communications protocols and standards underlying such systems.

7. In prior positions, I was the Chief Technology Architect (in addition to being the CEO) for 15 years for a company I founded that was the original communications technology “skunk-works” for Novell Inc. In this capacity, I designed basic and advanced telecommunications and network interfaces for

Novell and other companies and developed a mastery of the standards and protocols underlying the Internet. I authored the Network Communications Services Interface (“NCSI”) that became a de-facto communications software LAN/WAN standard, with more than 3 million deployments of software. I also developed protocol adjustments in Novell IPX Protocol for adaptive packet buffering required by LAN/WAN communication (Asynchronous and LAPB X.25) gateways, receiving Industry Product of the Year awards for successive years (1988, 1990, 1991 and 1996).

8. I previously served for 10 years architecting network information processing technologies for Goldman Sachs as a senior consultant. In this capacity, I was the architect, designer, development manager, and developer in Goldman’s Network Workstation Technologies Department. I was also the architect of Goldman’s product strategy and deployment of online delivery of consolidated live market data information into local and wide area network-based workstations for mission critical securities trading operations in the worldwide trading rooms of the firm. In particular, I developed proprietary adaptive buffering protocols to mitigate stream delays when terrestrial transatlantic data links were routed through backup satellite connections affecting the flow of steaming market data feeds used for program trading operations.

9. As Chief Technology Officer of ShowBizData Inc. between 2000 and 2002 as an “early adopter” pioneered the online Internet streaming of various lived events of the Cannes Film Festival, BFTA Awards and the Sundance Film Festival using both commercial and proprietary systems we architected and developed.

10. I am also a Network Computing Paradigm Award recipient.

11. A copy of my curriculum vitae is attached as Exhibit 1004.

12. I believe that I am qualified to provide reliable opinions in the field of the '866 Patent, including what a person of ordinary skill in the art (“POSITA”) would have understood from the prior art in this field at that time and his or her view of obviousness therein.

My Understanding of Claim Construction

13. My opinions herein address only technical issues that I have been asked to assess, relative to interpreting the claims of the patent. I understand that the ultimate claim interpretation is a question of law, which is subject to determination by the PTAB panel.

My Understanding of Anticipation

14. I understand that for purposes of an *inter partes* review (IPR) proceeding, the only bases for invalidity that may be considered are patents and printed publications. In particular whether the claims in question are anticipated

under 35 U.S.C. § 102 or rendered obvious under 35 U.S.C. § 103, by patents or printed publications.

15. I understand that a claim is considered anticipated (under 35 U.S.C. § 102) if each and every limitation of the claim may be found, expressly or inherently, within the disclosure of a single prior art reference (patent or printed publication), with the features arranged in the reference in the manner claimed.

16. My understanding of inherent disclosure of a claimed feature by a prior art reference is that it means the prior art reference must necessarily reflect the presence or operation of the feature (and that possibility or even strong likelihood that the feature is present is not sufficient).

My Understanding of Obviousness

17. I understand a patent claim is invalid if the claimed invention would have been obvious to a POSITA at the time the invention was made. 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 based on a combination of references that would have been obvious to a POSITA at the time of the invention.

18. I understand that an obviousness analysis requires an understanding of the scope and content of the prior art, any differences between the alleged

invention and the prior art, and the level of ordinary skill in evaluating the pertinent art.

19. I further understand that to prove prior art or a combination of prior art renders a patent obvious, it is necessary to: (1) identify the particular references that, singly or in combination, make the patent 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 challenged claim, and why it would have been obvious to a POSITA at the time of the alleged invention to have made that combination.

20. I understand that exemplary rationales that may support a conclusion of obviousness include:

- Combining prior art elements according to known methods to yield predictable results;
- Simple substitution of one known element for another to obtain predictable results;
- Use of known technique to improve similar devices (methods, or products) in the same way;
- Applying a known technique to a known device (method, or product) ready for improvement to yield predictable results;
- “Obvious to try” – choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success;
- Known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations are predictable to one of ordinary skill in the art; and

- Some teaching, suggestion, or motivation in the prior art that would have led a POSITA to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention.

21. I further understand that the above list of rationales provided is not all-inclusive. Other rationales may also be sufficient, so long as there is a link between the factual underpinnings relied on and an articulated rationale for why those factors render the invention obvious.

22. I understand certain objective indicia can also be important evidence regarding 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. I understand that to have legal weight, there must be a nexus—a connection—between any such secondary considerations and the alleged invention.

23. 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

field of the prior 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.

24. Rather, my testimony addresses the underlying facts and provides factual analysis that would support a legal conclusion of obviousness or non-obviousness. When I use the term obvious, it is with reference to the perspective of one of ordinary skill at the time of claimed invention.

Level of Skill in the Art

25. As part of this inquiry, I have been asked to consider the level of ordinary skill in the field that someone would have had at the time the claimed invention was made. In deciding the level of ordinary skill, I considered the following:

- 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.

26. The '866 patent presumes familiarity with Internet standards and practices, including Internet Requests for Comment (RFCs), which are cited repeatedly throughout the specification. Persons who have the skill to operate in this field will have a working familiarity with the relevant RFCs, as well as the basic standards applicable to distribution content over the Internet.

27. The education and/or working experience necessary to acquire the requisite familiarity with the subject matter to qualify as a POSITA would have

included either (1) at least a bachelor's degree or equivalent in a field such as Electrical Engineering, Computer Engineering, or Computer Science, or an equivalent field that includes network engineering as a topic of study, plus at least one year of practical academic or industry technical experience in the computer network field, such as serving as an engineer for an Internet service provider performing network design, development, or configuration tasks, or as a software developer for network communications software or related utility software, or (2) or at least three years' full-time technical experience as stated (or an equivalent combination of academic study and work experience).

OVERVIEW OF THE '866 PATENT

28. The '866 patent concerns retrieving Internet content through intermediary devices. The following description is made with reference to Fig. 5 of the '866 patent (depicting the first embodiment disclosed in the Detailed Description). *See* Ex. 1001, 83:25 *ff.*

29. Initially, tunnel device 33a signs into acceleration server 32, providing to acceleration server 32 tunnel device 33a's "identification on the Internet 113, such as its IP address," which acceleration




FIG. 5

server 32 then records in a database of such tunneling devices. *Id.*, 84:6-19. Client device 31 similarly signs in with acceleration server 32. *Id.*, 84:29-44. Thereafter, client device 31 sends a first request to acceleration server 32 for a list of available tunnels that can be used as proxies. *See id.*, 84:63-65. Acceleration server 32 responds with a list of tunnel devices. *Id.*, 85:1-5. This list includes "the identification (*e.g.*, IP address) of the tunnel devices to be used as intermediary devices to the client device #1 31a." *Id.*, 91:31-33. Tunnel devices 33a can be selected by client 31a or acceleration server 32 based upon various criteria. (*Id.*, 91:35-60.) These criteria could be, for example, the geographical location of tunnel device 33a. *See id.*, 91:55-25.

30. Client device 31a uses the Internet address of selected tunnel device 33a to establish a connection with tunnel device 33a. *Id.* 83:19-22. When client device 31a requires content from data server 22a, client device 31a sends a

“Content Request” message to tunnel device 33a. *Id.*, 83:46-51. This “Content Request” message identifies both the data server 22a and the desired content. *See, e.g., id.*, 95:47-66. The content can be partitioned in multiple parts or slices. *Id.*, 94:63-64. Different tunnels may be used to fetch different slices and to provide redundant access for the same slice. *Id.*, 95:10-92:6. *See generally id.*, 92:7-93:62.

31. Tunnel device 33a then fulfills its role as a proxy: “In response [to the “Content Request” message], the tunnel device 33a sends a ‘Content Request’ message [] to the data server #1 22a . . . requesting the content that was requested by the client device #1 31a. The data server #1 22a receives the request and . . . sends the requested content back to the tunnel device #1 33a” (*id.*, 85:53-60), with tunnel device 33a thereafter forwarding this content to client device 33a. *Id.* 83:60-65.

32. According to the ’866 patent, the “content requested by the client device #131a may be partitioned into multiple parts or ‘slices’.” *Id.*, 93:63-65. In particular, “the content may be partitioned into multiple slices (overlapping or non-overlapping), where each slice is requested and fetched using a distinct tunnel device (or via the client device serving as its own tunnel).” (*Id.*, 94:60-63.) The ’866 patent also refers to these content “slices” as “chunks,” and explains that chunks may be identified “by the content (e.g., URL, web-site, or web-page), and a number such as the number,” (*id.*, 105:21-23), a CRC value (*id.* 105:25-28), or be

based on a hash function of the chunk content (*id.*, 105:29-30). However, the '866 patent also discloses that a “chunk” can consist of an entire webpage, for a site comprising a plurality of webpages. *Id.*, 104:22-24.

33. The '866 patent discloses that client device 31a may also serve as a tunnel device and *vice versa*. *Id.*, 95:22-25. Moreover, the tunnel devices may store content obtained from earlier fetches from a data server, in which case rather than re-requesting such content from, *e.g.*, data server 22a, the tunnel device instead provides its locally-stored copy of the requested information to client 31a. *See id.*, 87:38-57.

34. Another embodiment is disclosed at length in Figs 26-26d for a peer-to-peer system 260, with related discussion from columns 101-109.

OVERVIEW OF THE CHALLENGED CLAIMS

35. A Claims Appendix, containing a listing of all challenged claims, is provided in Exhibit 1005. All claims are method claims. Only claim 15 of the fourteen challenged claims is independent. Claim 15 is reproduced below:

15. A method for fetching a content over the Internet from a first server identified in the Internet by a second identifier via a group of multiple devices, each identified in the Internet by an associated group device identifier, the method comprising the step of partitioning the content into a plurality of content slices, each content slice containing at least part of the content, and identified using a content slice identifier, and for each of the content slices, comprising the steps of:

(a) selecting a device from the group;

(b) sending over the Internet a first request to the selected device using the group device identifier of the selected device, the first request including the content slice identifier and the second identifier;

(c) in response to receiving the sent first request by the selected device, receiving over the Internet the content slice from the selected device; and

wherein the method further comprising the step of constructing the content from the received plurality of content slices,

and wherein each of the devices in the group is a client device.

TECHNICAL ISSUES CONCERNING CLAIM CONSTRUCTION

36. I will address claim construction to the extent it presents technical issues.

“Content Slice Identifier”

37. The '866 patent appears to treat the terms “chunk” and “slice” to be interchangeable and mean a part or the entirety of a content. Ex. 1001, 169:48-50. The '866 patent does not define a “slice identifier,” but it does explain that “the content in the chunks is identified by a chunk identifier, where each chunk identifier is associated with one, and only one, chunk.” *Id.*, 105:16-18; *see also id.*, 142:18-27.

38. To make technical sense of the claims, the term “content slice identifier” need only unambiguously identify the corresponding content slice to the responding device. Global identification of the content slice on the Internet may be

resolved through the separately specified device identifier (*e.g.*, “second identifier” (of the first server), as in claim 15).

39. To the extent this term requires construction, I would submit that a POSITA would understand this term, in light of the specification, to mean “a machine-readable expression in any form sufficient for the device that will retrieve the corresponding content slice to uniquely identify that content slice.”

“Identifier” (With Respect to a Device)

40. Within the challenged claims, a “first server” is “identified in the Internet by a second identifier,” and similarly each device in the “group of multiple devices” is also “identified in the Internet by an associated group device identifier ...” Ex. 1001, 173:42-44. Hence, the term “identifier” when used with respect to a device, is used to reference a machine-readable expression that a computing device uses to identify the subject device in the Internet. The challenged claims do not impose any further restrictions on “identifier,” and the specification makes clear that this term is intended to have broad scope.

41. According to the specification, “[e]ach of the identifiers herein may be a URL or an IP address in IPv4 or IPv6 form.” *Id.*, 55:31-32. Hence, “identifier” at least means a URL, an IP address, or both. But, the ’866 patent clearly intended “identifier” to be even broader than this, noting that “peer and agent devices may be identified by their respective IP address, *or by any other*

mechanism allowing addressing over the Internet.” Id., 107:53-55 (emphasis added).

42. Hence, when used with respect to a device, “identifier” would be understood by a POSITA as “any machine-readable expression that can be used by a computing device to address the device over the Internet.

“Partitioning the content into a plurality of content slices”

43. Claim 15 is silent as to what instrumentality performs the partitioning and where it is performed. I see no basis to impose further limitations in this regard.

44. In general, the recited partitioning step may be viewed as either (i) partitioning the content in advance on the device during retrieval, or (ii) selecting elements responsive to piece-wise requests, whether it has been pre-sliced or not. Either understanding is workable for purposes of accomplishing the object of fetching the content. Therefore, I believe a POSITA would see no reason to impose any further limitation as to which instrumentality does the partitioning, or when, or on which device it is carried out.

“Constructing the Content”

45. The District Court construed “constructing the content” to mean “assembling or combining content slices so as to reconstruct the original content.” Ex. 1023 at 48. I believe this is technically reasonable.

46. Claim 15 first recites the step of “partitioning *the content* into a plurality of content slices,” and later recites “constructing *the content* from the received plurality of content slices ... ” Ex. 1001, 173:45-174:11. In each case, the same content that is being partitioned is also being constructed. Further, this content is not being constructed from anything, but is specifically being constructed from the received plurality of content slices. Hence, the plain meaning of “constructing the content” within the context of claim 15 is that the original content is reconstructed from the received slices. This understanding is confirmed by the specification, which explains that when the slices are overlapping,

the content may be fully reconstructed from any two of the slices, hence providing a degree of redundancy. For example, in case of carrying the three slices over the Internet and a failure to receive one of the slices, the remaining two slices may be used to fully reconstruct the whole content.

Id., 95:52-57; *see also id.*, 105:57-58 (“multiple chunks may be combined to reconstruct the original content, such as website or content” (emphasis added)).

47. Based on the foregoing, I believe it is technically reasonable to interpret “constructing the content” to mean “assembling or combining content slices so as to reconstruct the original content.”

“wherein each of the devices in the group is a client device”


48. The '866 patent defines “client” as follows: “the term ‘client’ is used herein to include, but not limited to, a program or to a device or a computer (or a series of computers) executing this program, which accesses a server over the Internet for a service or a resource.” Ex. 1001 at 4:50-53. The District Court construed the subject limitation to mean “a device that is operating in the role of a client by requesting services, functionalities, or resources from other devices” (Ex. 1023 at 51), rejecting the interpretation that the other devices as to which the “client device” serves as a client must specifically be the tunnel, agent, or peer devices that otherwise participate in the claim. I believe a POSITA would consider the Court’s construction reasonable, since it follows the usual understanding per Internet RFCs of what a “client” is (see RFC 2616 (Ex. 1007) at 8 (defining “client” as “[a] program that establishes connections for the purpose of sending requests”), and closely follows the language of the specification in this regard.

Overview of the State of the Art at the Time of Filing

Sharp KK, EP Pat. Pub. No. 2597869 A1, Ex. 1018

49. Sharp KK published on May 29, 2013, which is about three months before the earliest priority date of August 28, 2013 for the '866 patent and I understand is thus prior art to the '866 patent under 35 U.S.C. § 102(a)(1).

50. Sharp KK discloses an online content distribution system, and of particular relevance is Embodiment 3 of Sharp KK, extending from columns 48 to 65. Fig. 20 of Sharp KK is reproduced below concerning this third embodiment:




51. In Sharp KK, server 2 specifies a plurality of relaying devices 3 from which content playing device 4 can acquire content. *See* Ex. 1018 ¶¶ 0217, 0229, 0232, 0257. Content playing device 4 acquires content by selecting the address of relaying device 3 from the plurality of relaying devices specified by server 2. *Id.* Content is in the form of a media file, which is broken into fragments. *Id.* ¶ 0240; Fig. 21. After selecting a relaying device 3, content playing device 4 sends a media request to the selected relaying device 3, requesting a fragment of the media by way of a media segment number, which is thereafter incremented. *Id.* ¶¶ 0272,

0302. This process continues until all media content has been received. *Id.* ¶ 0269; Fig. 23, step S632.

Shribman, U.S. Pat. Publ. No. 2011/0087733, Ex. 1017

52. Shribman published on April 14, 2011, which is more than one year before the earliest priority date of August 28, 2013 for the '866 patent, and I understand it is thus prior art to the '866 patent under 35 U.S.C. § 102(a)(1). Shribman subsequently issued as U.S. Pat. No. 8,560,604 (the "'604 patent").

53. Shribman is highly similar to the '866 patent architecturally. It discloses deploying a geographically diverse array of devices that can interchangeably function as clients, peers, or agents (caching proxies), so that those devices, when operating in the role of clients, can more efficiently retrieve content from origin servers over the Internet.



54. According to Shribman, any device in network 100 may serve as a client, peer, or agent, depending upon the requirements of network 100. Ex. 1017 ¶ 0035; *see also id.* Figs 4-6 and ¶¶ 0041-65 (disclosing general computing device 200 with software that enables peer, agent, and client functionalities). When one of the communication devices acts as a client 102, it can communicate with various peers 112, 114, 116 and one or more agents 122. *Id.* ¶ 0036.

55. To obtain content from a web server 152, client 102 of Shribman first registers with acceleration server 162, providing acceleration server 162 the IP address of client 102. *Id.* ¶ 0072. Thereafter, client 102 sends a request to acceleration server 162 that contains the IP address of web server 152, requesting a list of agents 122 that client 102 can use in relation to web server 152. *Id.* ¶ 0079. In response to this request, acceleration server 162 sends a list of agents to client 102. *Id.* ¶ 0080. This list necessarily includes the IP address of each agent. *See id.* ¶¶ 0080, 0101; Fig. 7, item 166; Fig. 13, items 560, 562, 564; claim 13.

56. Client 102 of Shribman reaches out to each agent in the received list, sending each agent the URL of the desired content. *Id.* ¶ 0081. In response, each agent sends to the client a list of peers and checksums of chunks of the desired content that these peers have. *Id.* ¶ 0083. This list of peers includes the IP address of each peer. *Id.*, Fig. 7, item 310. After deciding which agent to use, client 102 sends a request for specific chunks of content from the selected agent. *Id.* ¶ 0086.

57. In response to this chunk request from the client, if the selected agent determines that the requested information “exists and is still valid, then the agent prepares a response to the client, which includes for each of the chunks: (i) the checksum of the chunk; [and] (ii) a list of peers that according to the database of the selected agent contains these chunks” *Id.* ¶ 0087. As before, this list of peers necessarily includes the IP address of each peer. *Id.*, Fig. 7, item 310.

58. Upon receiving the response list, client 102 “sends a request to each of the peers listed for the chunk to download the chunk.” *Id.* ¶ 0090. The peers, in turn, respond to indicate whether or not they still have the requested chunks of information. *Id.* ¶ 0091. Client 102 “then selects the quickest peer that responds with a positive answer regarding the requested information, the client [letting] that peer know that it is chosen to provide the client with the chunk” *id.* ¶ 0091, and “the chosen peer then sends the chunk to the client.” *Id.* ¶ 0092.

ISO/IEC 23009-1:2012(E), MPEG-DASH standard, Jan. 5, 2012, Ex. 1027

59. This is a portion of the ISO/IEC standard for MPEG-DASH streaming, which was a draft standard at the time of this document. This document is a printed publication for purposes of 35 U.S.C. § 102, in that it was published in the ordinary course by an established and leading international standards organization, intended to be viewed as of its stated date of publication by the interested Internet engineering audience at large. MPEG-DASH is well known

streaming media standard in the technical community that includes the POSITAs herein.

RFC 2616, June 1999, Ex. 1007; RFC 3040, January 2001, Ex. 1020

60. These are Standards Track publications of the Internet Engineering Task Force. RFC 2616 was the principal standard on the HTTP protocol that underlies the World Wide Web that was current as of 2013. RFC 2616 is a printed publication for purposes of 35 U.S.C. § 102, in that it was published in the ordinary course by an established and leading standards organization and intended to be viewed as of its stated date of publication by the interested Internet engineering audience at large. It was a fundamental technical specification fully accessible to and well-known by POSITAs in this field as of 2013. The same is equally true for RFCs 791 (Ex. 1006) (defining the basic Internet Protocol), 1928 (defining Socks5 protocol, particularly relevant to other prior art cited with respect to the '044 patent), 3040 (Ex. 1020) (defining web replication and caching infrastructure), and 1945 (defining an earlier version of the HTTP protocol), each of which reflect basic, everyday knowledge of POSITAs as of the claimed priority date.

61. The Background section of the specification also cites RFCs 2616 and 791, acknowledging them as prior art. Ex. 1001 at 3:33, 4:63.

Prior Art Proxies, e.g., SQUID

62. The Background section of the '866 patent refers to “web proxies.”

One widely-known prior art example of a web (HTTP) proxy is Squid, which is extensively documented in Wessels, *Squid: The Definitive Guide* (O'Reilly Media 2004), Ex. 1012 hereto (“Wessels”). The book was published by a major publisher in the field and has the usual conventional markers (copyright, ISBN, etc.) indicating a reliable publication date in 2004. Also relevant is the 1998 book *Web Proxy Servers*, by Ari Luotonen, principal architect of the CERN proxy server, Ex. 1014 (the direct predecessor of Squid), which was similarly published (by Prentice Hall) and reliably dated in 1998 by conventional markers.

Open Proxies, e.g., ProxyList.net

63. The Background section of the '866 patent also refers to “open proxies.” One widely-known prior art example of open proxies is ProxyList.net, a web site that maintained a list of links to open proxies, available to the Internet user public. Anyone on the Internet could go to this list and click into an open public proxy server.

64. I now address the invalidity particulars for each of the cited references.

ANTICIPATION OF CLAIMS 15-17 AND 23-24 BY SHARP KK, EXHIBIT 1018

65. Sharp KK identically discloses all aspects of claims 15-17 and 23-24.

Anticipation of Claim 15 By Sharp KK

66. Concerning the preamble of claim 15, Sharp KK discloses a method for fetching content over the Internet from a first server identified in the Internet by a second identifier via a group of multiple devices. *See, e.g.*, “Embodiment 3” in Ex. 1018, Fig. 20, reproduced in the summary of Sharp KK, above.) (*See also id.*, Fig. 24.

67. As shown in Fig. 20 and its referencing disclosure, Sharp KK discloses a client 4e using multiple proxies 3a-3c (“a group of multiple devices”) to download content originating from server 2 (“a first server”). (Ex. 1018 ¶ 0220.) Server 2 is identified in the Internet by a domain name “example.com” (the “second identifier”). *See id.* Fig. 27, message “a”, “Host: example.com”). This occurs over networks including the Internet. (*Id.* ¶ 0569.

68. Client 4e requests content from server 2, for example, “GET /content1/0 HTTP/1.1”. *Id.* Fig. 27, message “(a)”, line 2; ¶¶ 0290, 0293. Server 2 responds with a list of proxies (“a group of multiple devices”) from which client 4e can obtain the requested content (“example-proxy1.com,” “example-client2.com,” “example-proxy3.com”). *See id.* Fig. 27, message “b”; *id.* ¶¶ 0294-0297. Each of these devices in the group has an associated group device identifier (*i.e.*, the domain name of the group device as recited in this list). *Id.* Each proxy is a

“device,” since each proxy 3 (or client 4f, which can also serve as a proxy) is a computing device. *See id.* ¶ 0566.

partitioning the content into a plurality of content slices, each content slice containing at least part of the content, and identified using a content slice identifier

69. Sharp KK discloses this fetching method further comprising the step of partitioning the content into a plurality of content slices (disclosed as “fragments” in Sharp KK), each content slice containing at least part of the content, and identified using a content slice identifier. In particular, Sharp KK discloses a media file being fragmented by a predetermined unit, such as by a group of pictures. *Id.* ¶ 0240, Fig.21. Consecutive reference numerals are assigned to each such fragment. *Id.* ¶ 0245. A media segment includes one or more movie fragments. *Id.* ¶ 0248. A media segment is a transmission unit, and the combined segments form the content. *Id.* ¶ 0247; *id.* Fig. 21.

70. As set forth in the following, for each such segment of desired content, Sharp KK also discloses performing the steps recited in the body of claim 15 to thus fetch the overall original content – that is, the media content.

[for each of the content slices ...] (a) selecting a device from the group

71. According to Sharp KK, in order to download media content, a client first provides a content identifier of the desired content to server 2. Ex. 1018 ¶¶

0290-93; *id.* Fig. 27. Server 2 responds with a list of proxies (or clients acting as proxies).

[0229] The client 4 of Embodiment 3 ... when functioning as a content playing device, (i) receives a content storage location list which is transmitted from the server 2 and (ii) acquires a content that the client 4 requests, from one of relaying devices which exist at an address indicated by the address information contained in the content storage location list thus received.

Id. ¶ 0229; *see also id.* ¶¶ 0294-96.

72. From this group of proxy devices, Sharp KK discloses client 4 selecting one of proxies 3 from this group. *Id.* ¶ 0229 (“from one of relaying devices which exist at an address indicated by the address information contained in the content storage location list thus received”); *id.* ¶¶ 0232-35. This proxy may be used for each of the content slices. The selected proxy may also be changed during the retrieval process. Sharp KK checks quality of service for each chunk, and if necessary, changes the proxy selection. *See, e.g., id.* ¶¶ 235, 283. Although the claim need not be interpreted to require repeated proxy selection for each chunk, Sharp KK, by disclosing to check the need for a proxy change after every retrieval, meets that additional requirement. This is clearly shown in Fig. 23 of Sharp KK, which shows an inner and outer processing loop for movie fragments and segments, respectively. The outer loop performs a quality of service check for each segment (S633), and decides whether to keep the current proxy or change to a new one (S634). Thus, Sharp KK meets limitation 15(a) however it may be construed.

(b) sending over the Internet a first request to the selected device using the group device identifier of the selected device, the first request including the content slice identifier and the second identifier

73. Sharp KK discloses using the group device identifier of the selected device to send over the Internet a first request that includes the content slice identifier and the second identifier. More particularly, Sharp KK discloses client 4e transmitting an instruction to acquire content from the proxy 3 at the address indicated by the address information that was selected in step (a) above. Ex. 1018 ¶¶ 0229, 0232, 0259, 0270, 0272, 0281; *id.* Fig. 27, message “(c)”. This transmission includes both the media segment number (“content slice identifier”) and also the domain name of server 2 from which the content originates (the “second identifier”). As shown in Fig. 27, the GET request to the select proxy includes the content identifier (“example.com/content1”) with the appended fragment number of “/0” (which is incremented in subsequent requests (*id.* ¶¶ 0270, 0302)). The “0” in this example is a content slice identifier. *See also id.* ¶ 0291. The request to selected proxy 3 also includes the host name “Host: example.com” of server 2, which is the “second identifier” of the “first server.” *Id.* Fig. 27 (“Host: example.com” in message “(a)”.

(c) in response to receiving the sent first request by the selected device, receiving over the Internet the content slice from the selected device; and

74. Sharp KK discloses, in response to receiving the sent first request by the selected device, client 4 receiving over the Internet the content slice from the selected device. More specifically, Sharp KK discloses proxy 3, in response to the client's request, checking if the requested content is stored in its cache. *Id.*, ¶0281. If this content is fresh, proxy 3 transmits the requested media segment to client 4. *Id.* Figs. 21, 24, and 27, message "(f)".

wherein the method further comprising the step of constructing the content from the received plurality of content slices

75. Sharp KK discloses a content playing device 4 that plays content received in the form of a plurality of segments. Ex. 1018 ¶¶ 0302, 0364 (referencing Fig. 23 and related description), ¶¶ 0559-61; *id.* Fig. 32, steps S724 and S632. Hence, the media being played is constructed from the content of the received plurality of media segments.

and wherein each of the devices in the group is a client device

76. Embodiment 3 discussed above is based upon Sharp KK's earlier embodiment 2 and differs only in that embodiment 3 supports a plurality of proxies from which a client 4 may select in the event of slow transmission from a current proxy 3. See Ex. 1018 ¶¶ 0219, 0285. Sharp KK's underlying embodiment 2, in turn, discloses "a client 4 has a function as and acts as a proxy so that it is possible

to widely distribute (i) a processing load of a server 2 and (ii) a network load which is used to transmit data from the server 2.” *Id.* ¶ 0142. “That is, in Embodiment 2, the client 4 acts both (i) as a device that requests a content and then acquires the content and (ii) as (a) a relaying device (proxy) that stores the content thus requested and then transfers the content to another device or (b) a playing device that acquires the content thus requested and then plays the content.” *Id.* ¶ 0144.

77. Furthermore, Sharp KK expressly refers to “a process which is carried out by the client 4 functioning as a relaying device of Embodiment 3 is identical to that of Embodiment 2.” *Id.* ¶ 0251. Since, as stated the process is “identical to that of Embodiment 2,” the disclosure then states: “As such, those processes will not be discussed here.” *Id.* But ¶ 0251 tells us all we need to know - an express disclosure that the clients of Embodiment 3 can also act as proxies, just as in Embodiment 2. Moreover, Embodiment 3 itself shows a scenario in which the proxy selection logic on the server “changes proxy 3a which is currently selected as the relaying device, from which content is required, to the client 4f having second priority to be selected as the relaying device,” expressly disclosing the dual roles that clients and proxies can have in Embodiment 3. *Id.* ¶ 283.

78. Sharp KK thus discloses that the group of proxies 3 are also client devices 4.

Anticipation of Claim 16 by Sharp KK

79. Claim 16 depends from claim 15 and recites “wherein the content is composed of bits, nibbles, bytes, characters, words, or strings, and the partitioning is based on bit, nibble, byte, multi-byte, number, character, word, or string level.” Sharp KK discloses that media may be fragmented along any suitable unit. Ex. 1018 ¶ 0240. It is inherent in such computer-related inventions that data is composed of bytes. Since each segment in Sharp KK contains one or more pictures that are divided at the byte level, Sharp KK discloses partitioning at the byte level. Sharp KK thus anticipates claim 16.

Anticipation of Claim 17 by Sharp KK

80. Claim 17 depends from claim 15 and recites, “wherein the content is composed of files, or programs, and the partitioning is based on file or program level.”

81. As discussed above, Embodiment 3 of Sharp KK identically discloses all limitations of the parent claim, 15.

82. Sharp KK further discloses, with respect to Embodiment 3, a scenario in which the client device is unable to process the “multipart” MIME format of the default media fragments in Embodiment 3. Ex. 1018 ¶ 0307. In that case, Sharp KK teaches to combine the fragments into “one body.” *Id.* ¶ 308. A POSITA would understand the resulting exchange (*i.e.*, a GET request without an

“Accept:multipart/media-segment” header, and a response with the combined non-MIME body) to be a request for a file, where the file contains a segment and constitutes a portion of the content (movie). Sharp KK also states that the format of the content in one embodiment is assumed to be a MP4 file format (*id.* ¶ 0022) and refers to a Figure illustrating segments as representing “files.” *Id.* ¶ 0240.

83. Furthermore, a web page may be represented by a file, and Sharp KK provides explicit disclosure of this. See point that follows regarding claim 18.

Anticipation of Claim 23 by Sharp KK

84. Claim 23 depends from claim 15 and recites “wherein the partitioning is sequential in the content.” Sharp KK discloses sequential fragments in each media segment, with this sequential ordering continuing across segments, thus anticipating claim 23. Ex. 1018 ¶ 0250.

Anticipation of Claim 24 by Sharp KK

85. Claim 24 depends from claim 15 and recites, “wherein the partitioning is non-sequential in the content.” Sharp KK discloses that fragmentation of the content generally “has consecutive playing time.” Ex. 1018 ¶ 0248. Sharp KK further discloses, however, “two or more movie fragments each of which has no consecutive playing time.” *Id.* Sharp KK thus anticipates claim 24.

Obviousness of Claim 18 over Sharp KK + MPEG DASH

Background - Claim 17

86. Claim 17 depends from claim 15 and recites “wherein the content is composed of files, or programs, and the partitioning is based on file or program level.”

87. If not anticipated by Sharp KK itself (per Ground 1), partitioning content at the file level as recited in claim 17 is obvious over Sharp KK (which identically discloses all limitations of parent claim 15), in view of the draft MPEG-DASH streaming standard as published in June 2012 (Ex. 1027). Ex. 1027 teaches that, in MPEG-DASH streaming, media segments such as those in Sharp KK could be partitioned as files. Ex. 1027 at 121 (partitioning media into a series of “.ts” files). Sharp KK itself states in ¶ 0022 that the media content is stored as files (“format of the content is assumed to be a MP4 file format”). This is arguably anticipation in itself. In any case, however, Sharp KK addresses MPEG-DASH streaming (*see* ¶ 0410), and expressly teaches to use the then draft MPEG-DASH standard. Since that standard contemplates segmenting at the file level, and Sharp KK contemplates that the server’s content will ordinarily exist as files (*id.* ¶ 0022 (content assumed to be in “MP4 file format” and use of HTTP as a “general file transferring protocol.”)), the combination in claim 17 is at the very least obvious, if not outright anticipated by Sharp KK itself.

Claim 18

88. Claim 18 depends from claim 17 and recites “wherein the content is a website content comprising multiple webpages, and the partitioning is based webpages level.” As stated in the specification of the ’866 patent, a web page can constitute a content slice where it is part of a plurality of web pages comprising a web site.

89. As noted above with reference to claim 17, Sharp KK discloses a VOD system. Sharp KK, however, also states that its content delivery system is not limited to such a VOD system. Ex. 1018 ¶ 0022. Indeed, Sharp KK explicitly contemplates server 2 responding to web page requests. *See id.* ¶ 0462. This reflects partitioning at the web page level as in claim 18. A POSITA would understand from the disclosure of a web page request in Sharp KK that such website data could be distributed to clients in the same manner that the video data is distributed in Embodiment 2, as doing so would not require any substantial changes to the underlying distribution process disclosed in Sharp KK, by simply substituting one known type of digital data (website data) for another type of digital data (video data).

90. It has been obvious since the beginning of the World Wide Web in the early 1990s that websites can comprise multiple webpages, with each webpage in the form of one or more respective files. Webster’s Third International Dictionary

defines “Web site or website” as a noun with first known use in 1992, meaning “a group of World Wide Web pages usually containing hyperlinks to each other and made available online by an individual, company, educational institution, government, or organization.” Ex. 1011. Microsoft’s Computer Dictionary says much the same thing. Ex. 1029.

91. Having multiple web pages rather than just one changes nothing in the operations that Sharp KK already teaches to perform to retrieve a web page via its proxy system. It is obvious to use the same technique on the multiple pages of the typical web site, and there would be every expectation that if Sharp KK worked to serve one of the pages it would work equally well to serve the others.

92. Finally, in the context of claim 18, the parent claim requires that the client reconstruct the content from the received slices. Where partitioning is at the webpage level, such reconstruction may occur in the browser cache, after a plurality of webpages from a website having multiple pages have been received. Such a cache is reflected, for example, in “client storage section 8” in Sharp KK. Ex. 1018 ¶ 0177.

93. It would therefore be obvious to a POSITA that an ordinary website would comprise multiple webpages and their related files, stored in the content storage of the website, such as section 5 of server 2, since such website data is also digital data that is partitioned into files, just as is the video data. Sharp KK further

discloses that the benefits of distributing a website in this manner is that it would reduce loading on both the network and the website server, thus motivating a POSITA to make such a substitution. *See* Ex. 1018 ¶¶ 0014, 0485.

OBVIOUSNESS OF CLAIMS 19, 20, AND 27-28 OVER SHARP KK IN VIEW OF SHRIBMAN

Claim 19

94. 93:63-94:54 in the '866 patent provide examples of the plain meaning of “wherein all of the parts of the content are included in all of the content slices.” One representative example from this citation is where the entire content is partitioned into three contiguous uninterrupted sequences of bytes which together span the entire range of bytes in the content. Ex. 1001, 94:16-23.

95. This claim is obvious over Sharp KK as above, in view of Shribman’s disclosure that “[c]hunks in the present description are defined as equally sized pieces of data that together form the whole content of the URL.” Ex. 1017 ¶ 0055. It would have been obvious to a POSITA to have the slices cover the entire content, because retrieving the content is the object of references such as Shribman and is itself an obvious objective for a content retrieval system.

Claim 20

96. Claim 20 recites “wherein all of the content slices have the same size.” While Sharp KK does not disclose that each of the media fragments is the same size, Sharp KK does disclose that each media fragment may have the same

length or the same number of pictures, each of which is a close analog to size. Ex. 1018 ¶ 0240. Indeed, Sharp KK contemplates the media file as being “fragmented by a predetermined unit,” of which length and number of pictures are presumably but two examples. *Id.*

97. Shribman discloses content being partitioned into fragments of equal size. Ex. 1017 ¶ 0055. It would therefore be obvious to a POSITA to partition the media fragments into equal sizes as taught by Shribman, as doing so would ensure a more even distribution of loading on the proxies across the system and making such a change would be a trivial variation of Sharp KK’s partitioning algorithm that would not otherwise affect the distribution of media segments so partitioned.

Claim 27

98. Claim 27 recites “wherein the number of content slices is higher than the number of devices in the group.” As discussed, Shribman discloses an embodiment in which an agent device retrieves content and thereby becomes the only agent with the requested content. Ex. 1017 ¶ 0089 (“selected agent ... provides itself as the only peer for these chunks”). Since Shribman discloses a plurality of slices, the limitations of claim 27 would be met by the combination of Sharp KK and Shribman. In view of the fact that the number of slices may be arbitrarily high, depending on the size of the content, and the number of devices is inherently a limited resource, the combination of circumstances contemplated by

claim 27 would obviously arise and, requiring no change in methodology, would be obvious to handle with completely predictable results.

Claim 28

99. Claim 28 recites “wherein the number of content slices is lower than the number of devices in the group.” As discussed, per the ’866 patent, content may be partitioned as coarsely as at the “webpage level.” Ex. 1001, 57:65-67. It is obvious that, given what may be of interest to users, many web sites will be encountered in the ordinary course, that only have a few pages. In view of this foreseeable situation, and the fact that the referenced disclosure in Shribman (of equally sized pieces of data) does not require more than two equal slices, it would have been obvious to envision scenarios in which the number of slices would be lower than the number of devices in the group. It would require no change in methodology to handle such a circumstance, again with completely predictable results.

OBVIOUSNESS OF CLAIMS 15, 17, AND 18 OVER LUOTONEN AND RFC 2616

Claim 15

100. Chunking of content was routine long before the August 28, 2013 priority date of the ’866 patent. For example, the HTTP standard for chunking data in HTTP requests is described in RFC 2616. Ex. 1007 at 18. Any web proxy wishing to be compliant, such as those disclosed in Luotonen (Ex. 1014) with such

Web standards would thus be compliant with RFC 2616 and would thus support the chunking of data in HTTP requests as described in RFC 2616. Content chunking via “byte-range” retrieval is a standard use of HTTP, which many websites employ. Web proxies themselves were also in wide use, for example as commonly incorporated in corporate firewalls. Consequently, the combination of web proxies with the chunking of data would have been obvious before August 28, 2013.

101. For example, concerning the preamble of claim 15, Luotonen discloses (throughout) methods for fetching a content over the Internet from a first server (any origin server), identified in the Internet by a second identifier (which the target server, when connected to the Internet, inherently has). *See, e.g.*, Ex. 1014 Fig. 1.4 on page 10. Luotonen further discloses having a pool of proxy servers, with a failover mechanism on the client, enabling the client, via auto proxy configuration, to select a different proxy. *Id.* at 307-08. It is inherent to the way that the Internet works that each proxy is (and must be) identified in the Internet by an associated group device identifier, that is, an IP address.

partitioning the content into a plurality of content slices, each content slice containing at least part of the content, and identified using a content slice identifier

102. The HTTP protocol is a fundamental basis for Luotonen, in that Web Proxies (the subject to the book) work over HTTP. RFC 2616, which specifies the

HTTP 1.1 protocol as of 2013, also discloses, in Section 14.3.1, that HTTP requests may be partitioned into byte ranges, enabling an HTTP client to partition requested content by byte ranges. Ex. 1007 at 85. Such partitioning may be effected with a “Range of Content-Range” request header, which includes a “content slice identifier”, identifying “at least a part of the content.” The format for this header is “(first-byte-pos ‘-’ last-byte-pos)” as disclosed in Sections 14.16 and 14.35.1 of RFC 2616. Such partitioning is a standard and widely used facility of HTTP, and it would have been obvious to use it with proxies as described in Luotonen, and indeed a standards-compliant proxy would be required and expected to support it. There is no need in this scenario to modify the proxy as envisioned by Luotonen – it simply handles this type of piece-wise request as any other HTTP request, because byte range requests are part of the HTTP 1.1 standard that Luotonen’s proxies support. *See* Ex. 1014 at 102 (internal page 84) (explicitly discussing byte range HTTP requests in the context of Web proxies).

[for each of the content slices ...] (a) selecting a device from the group

103. It was common before August 28, 2013 to have a plurality of proxies service an organization, as discussed in Luotonen. *See* Ex. 1014 at 373 (internal page 355) (proxy “pools”). The user of a proxy as described in Luotonen selects a proxy from the pool by configuring the user’s browser and/or computer. For example, Luotonen discloses that the client can have proxy autoconfiguration

features. *Id.*, Appendix A (“Proxy Auto-Configuration Support in Clients”), at 395-96 (internal pages 377-78). The proxy configuration on the client selects a proxy from the pool (“group”), which (until changed by an auto-configuration rule), applies for each of the content slices (byte-range requests).

(b) sending over the Internet a first request to the selected device using the group device identifier of the selected device, the first request including the content slice identifier and the second identifier

104. RFC 2616 discloses sending over the Internet a request to a selected device, such as a proxy, using the group device identifier (IP address) of the selected device, in which the request includes a content slice identifier and an identifier of a target server. *See* Ex. 1007 at 24-26. As discussed in RFC 2616, an HTTP request contains at least a “path” (which, in the case of a byte range request, corresponds to the “content slice identifier”). The request may contain “authority”, which corresponds to the domain name or IP address of the target server, i.e., the “second identifier,” and in any case must have a “Host” header also corresponding to the “second identifier” of claim 15. The request is directed to the “selected device,” which (as in many examples in Luotonen) may be over the Internet, in any case requiring the communication to be addressed to the selected device by its IP address.

(c) in response to receiving the sent first request by the selected device, receiving over the Internet the content slice from the selected device; and

105. This limitation is met in the combination of Luotonen and RFC 2616 when the selected device (the proxy) responds to the “first request” from the client.

wherein the method further comprising the step of constructing the content from the received plurality of content slices

106. It is inherent and well-understood in the field that a browser makes a byte range request, particularly one in which it specifies multiple ranges in a request Ex. 1007, RFC 2616 at 76 with the expectation that it will construct the desired content from responses to successive such requests.

wherein each of the devices in the group is a client device

107. Luotonen discloses that each of the proxies is also a client device. Ex. 1014 at 22 (internal page 4) (“[a] proxy acts both as a server as well as a client.”).

Claim 17:

108. According to claim 17, “the content is composed of files, or programs, and the partitioning is based on file or program level,” but this was a well-known feature long before the ’866 patent. *See, e.g.*, Ex. 1029, Microsoft Computer Dictionary (5th Ed. 2002) at 4 (internal page 564) (defining a Web site as: “a group of related HTML documents and associated files, scripts and databases”).

Claim 18:

109. The above Microsoft Computer Dictionary citation above reflects that web sites comprise multiple webpages and so are partitioned at the webpages level. Claim 18 is an obvious application of claim 15 when the content slices are taken as webpages.

OBVIOUSNESS OF CLAIMS 15, 17, AND 18 OVER LUOTONEN, RFC 2616, AND RFC 3040

110. Claim 15, when construed as suggested by the '866 patent specification, to cover as “partitioning” even delivery of content at the web page level, is so broad as to be obvious over the mere use of prior art proxies in browsing multi-page websites. It has been routine since the beginning of the World-Wide Web to use browsers such as Netscape or Firefox to browse multi-page websites. Since web proxies are ubiquitous (*e.g.*, in workplace firewall configurations, among many others, this is an obvious combination.

111. This ground is identical to the previous ground, except that the partitioning step is met simply by content retrieval at the webpage level from any web site comprising a plurality of pages. As addressed above, Luotonen already addresses the other aspects of the preamble, step (a) (selecting a device from the group), (b) the request format, (c) receiving the content slice (webpage), and the second “wherein” (proxy also acting as a client).

112. As addressed above, a web page may be a “content slice” as understood in the context of the ’866 patent, where it is one of a plurality of web pages comprising a web site. *See* Ex. 1001, claim 18. As further addressed above, it is common knowledge, even to the most casual Internet users, that web sites then (and now) typically comprised a plurality of web pages. *See* Ex. 1029, Microsoft Computer Dictionary (5th Ed. 2002) at 4 (internal page 564) (“Web site n. A group of related HTML documents, scripts, and associated files that is served up by an HTTP server on the World Wide Web.”)

113. With regard to the first “wherein” clause,

wherein the method further comprising the step of constructing the content from the received plurality of content slices

as addressed above, this limitation may be understood as “assembling or combining content slices so as to reconstruct the original content.” Where the units of “partitioning” (content slices) are at the webpage level itself, this claim step is satisfied when a user browses a plurality of different web pages from the same site. A POSITA would recognize that this reflects the routine use of the Internet at the time by ordinary Internet users. The content thus reconstructed is not just transitory. Rather, it is also noted as a matter of background in the ’866 patent that “[w]eb browser cache previously accessed web resources and reuse them when possible.” Ex. 1001, 5:22-24. Such caching was indeed well known, and is reflected, for example, in Ex. 1020, RFC 3040 at 7 (defining “user agent cache”

and “The cache within the user agent program” (browser)). A user performing the well-known and routine operation of browsing a web site thus performs “constructing the content from the received plurality of content slices,” in the browser cache.

114. As for claim 17, again as noted in the Background section of the ’866 patent and well known in the art, a web page is also a “file.” *See* Ex. 1001, 5:5-11. *See also* Ex. 1029, Microsoft Computer Dictionary (5th Ed. 2002) at 4 (internal page 564).

115. As for claim 18, the above Microsoft Computer Dictionary citation above reflects that web sites as therein defined comprise multiple webpages, partitioned based on webpages level. This is an obvious application of claim 15 when the content slices are webpages as in the reference disclosure in the ’866 patent.

* * * * *


116. In signing this declaration, I recognize 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, and that I may accordingly be subject to cross-examination. If such cross-examination is required of me, I will

appear for cross-examination within the United States during the time allotted for cross-examination.

117. This declaration represents only those opinions I have formed to date. I reserve the right to revise, supplement, and/or amend my opinions stated herein based on new information and on my continuing analysis of the materials produced in this proceeding.

I declare under penalty of perjury that the foregoing is true and correct.

Executed at McMinnville, Oregon on January 24, 2021.



KEITH J. TERUYA