

**UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF NEW YORK**

PERSONALWEB TECHNOLOGIES, LLC, a
Texas limited liability company, and
LEVEL 3 COMMUNICATIONS, LLC, a
Delaware limited liability company,

Plaintiffs,

v.

WEWORK COMPANIES, INC, a Delaware
corporation,

Defendant.

18 Civ. 6932

COMPLAINT

JURY TRIAL DEMANDED

Plaintiff PersonalWeb Technologies, LLC (“Plaintiff” or “PersonalWeb”) files this Complaint for patent infringement against Defendant WeWork Companies, Inc. (“Defendant”).

Plaintiff PersonalWeb Technologies, LLC alleges:

PRELIMINARY STATEMENT

1. PersonalWeb and Level 3 Communications, LLC (“Level 3”) are parties to an agreement between Kinetech, Inc. and Digital Island, Inc. dated September 1, 2000 (the “Agreement”). Pursuant to the Agreement, PersonalWeb and Level 3 each own a fifty percent (50%) undivided interest in and to the patents at issue in this action: U.S. Patent Nos. 6,928,442, 7,802,310, 7,945,544, and 8,099,420 (“Patents-in-Suit”). Level 3 has joined in this Complaint pursuant to its contractual obligations under the Agreement, at the request of PersonalWeb.

2. Pursuant to the Agreement, Level 3 has, among other rights, certain defined rights to use, practice, license, sublicense and enforce and/or litigate the Patents-in-Suit in connection with a particular field of use (“Level 3 Exclusive Field”). Pursuant to the Agreement PersonalWeb has, among other rights, certain defined rights to use, practice, license, sublicense, enforce and/or

litigate the Patents-in-Suit in fields other than the Level 3 Exclusive Field (the “PersonalWeb Patent Field”).

3. All infringement allegations, statements describing PersonalWeb, statements describing any Defendant (or any Defendant’s products), any statements made in the paragraphs set forth in the Section entitled “General Background,” and any statements made regarding jurisdiction and venue are made by PersonalWeb alone, and not by Level 3. PersonalWeb alleges that the infringements at issue in this case all occur within, and are limited to, the PersonalWeb Patent Field. Accordingly, PersonalWeb has not provided notice to Level 3—under Section 6.4.1 of the Agreement or otherwise—that PersonalWeb desires to bring suit in the Level 3 Exclusive Field in its own name on its own behalf or that PersonalWeb knows or suspects that Defendant is infringing or has infringed any of Level 3’s rights in the patents.

THE PARTIES

4. Plaintiff PersonalWeb Technologies, LLC is a limited liability company duly organized and existing under the laws of Texas with its principal place of business at 112 E. Line Street, Suite 204, Tyler, TX 75702.

5. Plaintiff Level 3 Communications, LLC is a limited liability company organized under the laws of Delaware with its principal place of business at 100 CenturyLink Drive, Monroe, Louisiana, 71203.

6. PersonalWeb’s infringement claims asserted in this case are asserted by PersonalWeb and all fall outside the Level 3 Exclusive Field. Level 3 is currently not asserting patent infringement in this case in the Level 3 Exclusive Field against any Defendant.

7. Defendant WeWork Companies, Inc. is, upon information and belief, a Delaware corporation having a principal place of business or regular and established place of 115 W. 18th St., 5th Floor, New York, New York 10011.

JURISDICTION AND VENUE

8. The court has subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a) because this action arises under the patent laws of the United States, 35 U.S.C. §§ 1 *et seq.*

9. Venue is proper in this federal district pursuant to 28 U.S.C. §§ 1391(b)–(c) and 1400(b) because Defendant is incorporated in the State of Delaware, and on information and belief, has a regular and established place of business in this District and has committed acts of infringement in this District.

10. This court has personal jurisdiction over Defendant because, in addition to the allegations in above paragraphs, on information and belief, Defendant purposefully directed activities at residents of New York, the claims herein arise out of and relate to those activities, and assertion of personal jurisdiction over Defendant would be fair.

PERSONALWEB BACKGROUND

11. The Patents-in-Suit cover fundamental aspects of cloud computing, including the identification of files or data and the efficient retrieval thereof in a manner which reduces bandwidth transmission and storage requirements.

12. The ability to reliably identify and access specific data is essential to any computer system or network. On a single computer or within a small network, the task is relatively easy: simply name the file, identify it by that name and its stored location on the computer or within the network, and access it by name and location. Early operating systems facilitated this approach with standardized naming conventions, storage device identifiers, and folder structures.

13. Ronald Lachman and David Farber, the inventors of the Patents-in-Suit, recognized that the conventional approach for naming, locating, and accessing data in computer networks could not keep pace with ever-expanding, global data processing networks. New distributed storage systems use files that are stored across different devices in dispersed geographic locations. These different locations could use dissimilar conventions for identifying storage devices and data partitions. Likewise, different users could give identical names to different files or parts of files—or unknowingly give different names to identical files. No solution existed to ensure that identical file names referred to the same data, and conversely, that different file names referred to different data. As a result, expanding networks could not only become clogged with duplicate data, they also made locating and controlling access to stored data more difficult.

14. Lachman and Farber developed a solution: replacing conventional naming and storing conventions with system-wide “substantially unique,” content-based identifiers. Their approach assigned substantially unique identifiers to “data items” of any type: “the contents of a file, a portion of a file, a page in memory, an object in an object-oriented program, a digital message, a digital scanned image, a part of a video or audio signal, or any other entity which can be represented by a sequence of bits.” Applied system-wide, this invention would permit any data item to be stored, located, managed, synchronized, and accessed using its content-based identifier.

15. To create a substantially unique, content-based identifier, Lachman and Farber turned to cryptography. Cryptographic hash functions, including MD4, MD5, and SHA, had been used in computer systems to verify the integrity of retrieved data—a so-called “checksum.” Lachman and Farber recognized that these same hash functions could be devoted to a vital new purpose: if a cryptographic hash function was applied to a sequence of bits (a “data item”), it would produce a substantially unique result value, one that: (1) virtually guarantees a different result

value if the data item is changed; (2) is computationally difficult to reproduce with a different sequence of bits; and (3) cannot be used to recreate the original sequence of bits.

16. These cryptographic hash functions would thus assign any sequence of bits, based on content alone, with a substantially unique identifier. Lachman and Farber estimated that the odds of these hash functions producing the same identifier for two different sequences of bits (i.e., the “probability of collision”) be about 1 in 2 to the 29th power. Lachman and Farber dubbed their content-based identifier a “True Name.”

17. Using a True Name, Lachman and Farber conceived various data structures and methods for managing data (each data item correlated with a single True Name) within a network—no matter the complexity of the data or the network. These data structures provide a key-map organization, allowing for a rapid identification of any particular data item anywhere in a network by comparing a True Name for the data item against other True Names for data items already in the network. In operation, managing data using True Names allows a user to determine the location of any data in a network, determine whether access is authorized, and to selectively provide access to specific content not possible using the conventional naming arts.

18. On April 11, 1995, Lachman and Farber filed their patent application, describing these and other ways in which content-based “True Names” elevated data-processing systems over conventional file-naming systems. The first True Name patent issued on November 2, 1999. The last of the Patents-in-Suit has expired, and the allegations herein are directed to the time period before expiration of the last of the Patents-in-Suit.

19. PersonalWeb has successfully enforced its intellectual property rights against third party infringers, and its enforcement of the Patents-In Suit is ongoing. This enforcement has

resulted in PersonalWeb obtaining settlements and granting non-exclusive licenses regarding the Patents-in-Suit.

GENERAL BACKGROUND

20. A webpage is a type of document that is typically retrieved over the World Wide Web, made viewable and formatted (rendered) by a web browser, and displayed electronically. A “webpage” often refers to what is visible in a browser, but sometimes also refers to a computer file (“webpage file”), usually written in Hypertext Markup Language (“HTML”) or a comparable markup language. Such HTML files typically include text, formatting, and references (hyperlinks) to other web contents, such as style sheets, scripts, and images. Web contents referenced in an HTML or similar file are also called “webpage assets” or “asset files” herein. The web browser coordinates the retrieval of the various asset files of a webpage and renders the webpage for display from the webpage file and the asset files referenced in the webpage file.

21. On the World Wide Web, hyperlinks generally include Uniform Resource Identifiers (“URIs”), which each typically include an address of a server from which the asset file could be retrieved (e.g., “www.website.com”), a “path” to the location of that asset file on that server (e.g., “/directory/”), and a filename (e.g., “filename.ext”).

22. On the Internet, a web browser typically retrieves a webpage file from a remote web server and retrieves referenced asset files from the same or different servers. The web browser retrieves a webpage file and asset files by making a GET “request” to a web server using the Hypertext Transfer Protocol (“HTTP”), an industry standard. The web server may respond to such a request with a HTTP “response” that includes the requested web content.

23. A static webpage is delivered exactly as stored, as web content in the web server’s file system or memory. In contrast, a dynamic webpage is generated by a web server application,

usually driven by server-side software, upon receipt of a request from a browser (user). For example, a picture of a building might be delivered as static content (a picture) whereas the latest traffic conditions may be delivered dynamically based on real time traffic information.

24. The speed of a browser retrieving webpage files and incorporated asset files can be increased by the browser storing previously retrieved webpage files and asset files in a browser “cache” on the computer running the browser. If a browser’s user later requests a previously retrieved webpage file or requests a webpage that includes an asset file previously used by the browser in rendering the same or a different webpage (for example, by reloading a webpage or visiting the same webpage again), the browser may use the cached webpage file or asset file rather than having to download the same file repeatedly over the Internet again.

25. Two computers communicating over the Internet usually are not directly connected to each other but rather interact via chains of network appliances and other computers (e.g., “switches” and “intermediate” servers). Many intermediate servers have caches similar to and complementing the browser cache that store webpage files and assets that pass through that intermediate server. If a browser or server requests a file from the intermediate server that is present in that intermediate server’s cache, the intermediate server can use the content in its cache to respond to the request rather than send the request upstream towards the web server from which the file initially originated (also called the “origin server”).

26. Responses to HTTP requests may include header elements (control elements) and a body (the “object” that was requested). Under HTTP, web servers can include a “cache-control” header with a response including a webpage or asset file. A “cache-control” header includes one or more directives that instruct browsers and intermediate server caches (“intermediate caches”) as to whether and for how long the file (object) included in the response may be cached. HTTP

also provides for including other headers in responses that provide similar types of instructions to browsers and intermediate caches. Collectively, these other headers and directives in a “cache-control” header are referred to herein as “cache-control headers.”

27. Given that webpage content changes, sometimes rather quickly and regularly, a problem that website owners face is effectively instructing a browser that is re-rendering a previously cached webpage that one or more of its cached files for that webpage are no longer the correct and authorized content (the content of those files has changed) and similarly reauthorizing the use of those cached files whose content has not changed.

28. On one hand, website owners want to encourage the browsers that render their web pages to use cached files so as to reduce the number of requests for these files that are being made to their webpage servers. Therefore, they frequently will set cache-control headers that authorize the browser to cache their webpage and asset files so they are on hand when the browser needs to render that webpage again. On the other hand, website owners want the browsers to use the latest authorized files so that their users do not see the wrong content when viewing their webpage.

DEFENDANT’S BACKGROUND

29. On information and belief, Defendant has operated a website located at **wework.com**, and has done so since before expiration of the last to expire of the Patents-in-Suit, which has operated to provide authorized webpage content to its users in the manner herein described.¹

30. On information and belief, Defendant’s web servers utilized a system of notifications and authorizations to control the distribution of content, *e.g.*, what webpage content

¹ While the complaint is sometimes written in the present or present perfect tense, all specific allegations are directed to the system’s operations and the method’s performance in the relevant time period.

may be served from web servers and intermediate caches and what cached webpage content a browser is re-authorized to use to render Defendant's webpage(s).

31. On information and belief, Defendant's system and its associated method of providing webpage content, used "conditional" HTTP GET requests with If-None-Match headers and associated content-based ETag values for various webpage and/or asset files required to render various webpages of the Defendant.

32. On information and belief, Defendant's system and its associated method of providing webpage content, also inserted fingerprints generated based on the content of asset files into the filenames of asset files required to render various webpages of the Defendant.

33. On information and belief, Defendant's system and associated method used these ETags and fingerprints to instruct both the intermediate cache servers and the endpoint caches at browsers to verify whether they were still authorized to reuse he previously cached webpage files of Defendant and to instruct them to obtain newly authorized content in rendering Defendant's webpage when that content had changed.

34. On information and belief, Defendant thereby reduced the bandwidth and computation required by their origin servers and any intermediate cache servers to field user requests to render Defendant's webpages as those servers only need to serve files whose content has changed. On information and belief, this has allowed for the efficient update of cached information only when such content has changed, thereby reducing transaction overhead and bandwidth and allowing the authorized content to be served from the nearest cache.

35. More particularly, on information and belief, each of Defendant's webpages included a webpage file (e.g., a main or initial HTML file) and one or more asset files referenced in the webpage file. On information and belief, the references in the webpage file to the asset files

needed to render the webpage were typically Uniform Resource Identifiers (“URIs”), which each typically included a filename, the address of a server from which the asset file could be retrieved, and a “path” to the location of that asset file on that server.

36. On information and belief, Defendant’s website used a web application framework to develop and compile various webpages of the Defendant, including asset files that were used in rendering the webpages, and to generate fingerprints of the contents of asset files when the webpages were compiled. On information and belief, the fingerprint of individual asset files that were part of the webpage’s content were included in the respective filenames of the individual asset files. On information and belief, the modified filenames were then used as part of the URI used to access the individual asset files over the Internet. On information and belief, when an asset file’s content was changed, a new fingerprint was generated and included in the filename, its URI thus being changed accordingly.

37. On information and belief, the asset file fingerprint was generated with a hash function and used to identify content changes. Furthermore, on information and belief, asset file URIs (with respective fingerprints) were included in webpage files. On information and belief, static webpage files, if any, were recompiled when any URI of a referenced asset file was changed (due to the fingerprint of the referenced asset file changing). Thus, a content change in an asset file for a given webpage would result in a change to its fingerprint, its URI, and a subsequent change to the content of any static webpage files referencing that changed asset file for that webpage.

38. On information and belief, a dynamic webpage file generated for a webpage of Defendant webpages in response to one request from a user could be the same as it was when it was generated in response to a prior request from that or another user. On information and belief,

this would not be the case if any of the asset files referenced in the webpage file had changed between the time of the two requests and the URIs of the changed assets included fingerprints as described above.

39. On information and belief, when Defendant created a webpage file for a webpage, whether dynamic or static, that webpage file included a sequence of bits and an associated ETag value was generated by Defendant by applying a hash function to the sequence of bits; wherein any two webpage files comprising identical sequences of bits had identical associated ETag values. Thus, on information and belief, when a webpage file's content was changed and a new associated ETag value was generated by Defendant, it thereafter instructed the respective service by intermediate cache servers or use by endpoint caches such as browser caches to no longer use the previous webpage file's content.

40. On information and belief, when an intermediate cache server or a browser requested a webpage from the Defendant for the first time, it sent an HTTP GET request with the webpage's URI and Defendant's origin server or an upstream cache server responded by sending an HTTP 200 (OK) response message containing the webpage file, along with its respective associated ETag. On information and belief, the intermediate cache server or a browser then sent individual HTTP Get requests, each with an asset URI that was referenced in the webpage file, and Defendant's origin server or an upstream cache server responded by sending individual HTTP 200 responses containing the requested asset files, along with their respective associated ETags. On information and belief, upon receipt of the HTTP 200 responses, the intermediate cache server or browser cached the webpage and asset files with their associated URI and associated ETag values and the browser used them in rendering the requested web page of the Defendant. On information and belief, the intermediate cache servers and browser caches were caused to maintain

databases/tables which mapped the URIs of asset/webpage files to their respective responses and, if applicable, associated cache-control headers and ETags.

41. On information and belief, by responding to an HTTP GET request for a given webpage by transmitting content of a webpage/asset file with an associated ETag, Defendant instructed the browser cache and all intermediate cache servers, to use an HTTP conditional GET request the next time that webpage/asset file is requested. More specifically, on information and belief, the browser or intermediate cache is instructed to include the ETag in the HTTP conditional GET request with an “If-None-Match” header to re-verify that they are still authorized to serve or use that content or determine that they are no longer authorized to use that content and therefore must use new content.

42. On information and belief, Defendant did this, for example, by causing cache-control headers to be included in HTTP responses containing its webpage/asset files. On information and belief, Defendant benefits from using the ETags to control the distribution of its webpage content by communicating to a downstream cache and to a browser which of Defendant’s cached webpage files it is reauthorized to serve/use and what newly authorized files it must first obtain in serving/rendering Defendant’s webpages.

43. More particularly, on information and belief, when a browser again requested the Defendant’s webpage, the browser, based on the cache-control headers received in the original response, sent a conditional GET request with an If-None-Match header using the associated ETag value and the URI for the webpage file so as to be notified whether the browser still had Defendant’s authority to render the webpage with its locally cached webpage file.

44. On information and belief, under most circumstances, a responding intermediate cache server, having an ETag for that URI responded to the request by determining whether it had

the same associated ETag value in its list of associated ETag values for that URI. If it had no ETag value for that URI, the request was passed up to an upstream intermediate cache server capable of responding or, if none, to the Defendant's origin server which performed the response.

45. On information and belief, if the responding server had the webpage content for that URI and there was a match between the ETag it received in the request with the ETag it currently had associated for that URI, it sent back an HTTP 304 (Not Modified) response message; this message notifying the browser that the same webpage content was present at the responding server and that the browser was still authorized to use that previously cached webpage file to render the webpage. On information and belief, upon receipt of the HTTP 304 response, the browser accessed the locally cached webpage file in rendering the webpage.

46. On information and belief, if the webpage file's associated ETag sent by the browser in the conditional GET If-None-Match request did not match the associated ETag maintained at the responding server (or other intermediate cache servers further upstream or the origin server) for that URI, the responding server sent back an HTTP 200 response along with the new webpage file and its new ETag value. The HTTP 200 response indicated to the browser that it was not authorized to use (or serve, in the case of an intermediate cache server receiving the HTTP 200 response) the previously cached webpage file. In response to receiving the HTTP 200 response, the browser (or intermediate cache server) was instructed to update its respective cache with the new webpage file and associated ETag. The browser subsequently read the new webpage file and the asset file URIs contained therein to render the webpage.

47. On information and belief, for a particular asset file URI in the webpage file for which there was a matching entry in the browser's cache and that cache entry did not include an associated ETag value, the browser was allowed to use the cached content of the previously

received asset file, subject to the asset file's cache-control headers. On information and belief, if the cache-control headers did not allow the cached asset file content to be re-used or if there was no entry in the browser's cache for the asset file's URI, the browser sent an HTTP GET request with the asset file's URI; and the responding intermediate or origin server responded to the GET request by sending the asset file for that URI and, if any, the corresponding cache-control headers and/or associated ETag with an HTTP 200 response. On information and belief, in response to receiving the HTTP 200 response, the browser cached the asset file, if any, and its cache-control headers and/or associated ETag and used the newly received asset files in rendering Defendant's webpage. On information and belief, if the downstream intermediate cache or the browser was later required to again serve or render the webpage, it went through the above process to determine which file content it still had authority to access or whether it needed to access different authorized content to serve or render the webpage via the HTTP 304 and HTTP 200 responses.

48. On information and belief, the browser repeated this process for various asset files for which it had an associated ETag value and asset files with URIs in the webpage file.

49. On information and belief, in this manner, Defendant used (1) ETag values and (2) asset files referenced by URIs with fingerprints based on the asset files' content to control the behavior of downstream intermediate cache servers and endpoint caches to assure that they only accessed and used Defendant's latest authorized webpage content to serve or to render their webpages.

FIRST CLAIM FOR RELIEF

INFRINGEMENT OF U.S. PATENT NO. 6,928,442

50. PersonalWeb repeats and realleges paragraphs 1–49, as if the same were fully stated herein.

51. On August 9, 2005, United States Patent No. 6,928,442 (the “‘442 patent”) was duly and legally issued for an invention entitled “Enforcement and Policing of Licensed Content Using Content-Based Identifiers.” PersonalWeb has an ownership interest in the ‘442 patent by assignment, including the exclusive right to enforce the ‘442 patent within the PersonalWeb Patent Field, and continues to hold that ownership interest in the ‘442 patent.

52. Defendant has infringed at least claims 10 and 11 of the ‘442 patent by its manufacture, use, sale, importation, and/or offer for sale of products or services, and/or controlling the distribution of its webpage content in the manner described herein. Defendant’s infringement is literal and/or under the doctrine of equivalents and Defendant is liable for its infringement of the ‘442 patent pursuant to 35 U.S.C. § 271.

53. For example, claim 10 covers “a method, in a system in which a plurality of files are distributed across a plurality of computers.” On information and belief, Defendant has used a system of notifications and authorizations to distribute a plurality of files, *e.g.*, Defendant’s files containing content necessary to render its webpages, across a plurality of computers such as production servers, origin servers, intermediate cache servers and endpoint caches used by browsers rendering Defendant’s webpages.

54. Claim 10 then recites the act of “obtaining a name for a data file, the name being based at least in part on a given function of the data, wherein the data used by the function comprises the contents of the particular file.” As set forth above, on information and belief, Defendant generated or otherwise obtained ETags for its webpage and asset files used to render its webpages using a hash function, wherein the ETags were based on the contents of the particular files. Moreover, Defendant caused the intermediate caches servers and endpoint caches to obtain the ETags in HTTP 200 responses sent from Defendant’s origin servers. On information and

belief, Defendant caused intermediate cache servers and its origin servers to obtain ETags in conditional GET messages from endpoint and intermediate caches, as described *supra*.

55. Claim 10 then recites the act of “determining, using at least the name, whether a copy of the data file is present on at least one of said computers.” On information and belief, as set forth above, Defendant has caused its origin servers and the intermediate cache servers between an endpoint cache and one of its origin servers to, in response to receiving a conditional GET request with an If-None-Match header, determine whether it has a file present that matches the URI in the conditional GET and to compare the ETag in the conditional GET to the ETag for that URI and determine whether a copy of the content having that ETag is present.

56. Claim 10 then recites the act of “determining whether a copy of the data file that is present on at least one of said computers is an unauthorized copy or an unlicensed copy of the data file.” On information and belief, as set forth above, if there was a match, the origin or intermediate cache server determined that the copy of the file present at the downstream intermediate cache server and/or the endpoint cache was an authorized or licensed copy of the data file. Conversely, if there was no match, it determined that the copy of the file present at the downstream intermediate cache server and/or the endpoint cache was an unauthorized copy of the data file. Likewise, if the browser determined that it had a file with a matching URI, the browser determined that it was still authorized to use that file.

57. Defendant’s acts of infringement caused damage to PersonalWeb and PersonalWeb is entitled to recover from Defendant the damages sustained by PersonalWeb as a result of Defendant’s wrongful acts in an amount subject to proof at trial.

SECOND CLAIM FOR RELIEF

INFRINGEMENT OF U.S. PATENT NO. 7,802,310

58. PersonalWeb repeats and realleges paragraphs 1–49, as if the same were fully stated herein.

59. On September 21, 2010, United States Patent No. 7,802,310 (the “‘310 patent”) was duly and legally issued for an invention entitled “Controlling Access to Data in a Data Processing System.” PersonalWeb has an ownership interest in the ‘310 patent by assignment, including the exclusive right to enforce the ‘310 patent within the PersonalWeb Patent Field, and continues to hold that ownership interest in the ‘310 patent.

60. Defendant has infringed at least claims 20 and 69 of the ‘310 patent by its manufacture, use, sale, importation, and/or offer for sale of products or services, and/or controlling the distribution of its webpage content in the manner described herein. Defendant’s infringement is literal and/or under the doctrine of equivalents and Defendant is liable for its infringement of the ‘310 patent pursuant to 35 U.S.C. § 271.

61. For example, claim 20 covers a “computer-implemented method operable in a system which includes a plurality of computers.” On information and belief, Defendant used the claimed computer implemented method by using a system of notifications and authorizations to control the distribution of data items, such as various webpage and asset files, necessary to render its webpages, across a plurality of computers such as production servers, origin servers, intermediate cache servers, and endpoint caches.

62. Claim 20 then recites “controlling distribution of content from a first computer to at least one other computer, in response to a request obtained by a first device in the system from a second device in the system, the first device comprising hardware including at least one

processor, the request including at least a content-dependent name of a particular data item, the content-dependent name being based at least in part on a function of at least some of the data comprising the particular data item, wherein the function comprises a message digest function or a hash function, and wherein two identical data items will have the same content-dependent name.” On information and belief, as set forth above, Defendant has caused downstream intermediate cache servers and endpoint caches to send conditional GET requests with If-None-Match headers containing ETags that are fielded by upstream cache or origin servers. On information and belief, the ETags were content-dependent names for a data item based on hashing the data item’s contents; and when the file’s content changed a new content-dependent name was determined. On information and belief, in Defendant’s method, a first computer, such as the intermediate cache server or origin server, received such conditional GET requests from a second computer, such as a user browser or other intermediate cache server, regarding data items, such as webpage or asset files, the requests including ETags associated with the respective data items.

63. Claim 20 then recites “based at least in part on said content-dependent name of said particular data item, the first device (A) permitting the content to be provided to or accessed by the at least one other computer if it is not determined that the content is unauthorized or unlicensed, otherwise, (B) if it is determined that the content is unauthorized or unlicensed, not permitting the content to be provided to or accessed by the at least one other computer.” On information and belief, the first computer, such as an upstream intermediate cache server or origin server, maintained a plurality of ETags associated with Defendant’s asset and webpage files. On information and belief, the ETag in a request and the ETag maintained by the first computer for the particular data item sought by the request were compared to determine whether the associated content present at the downstream computer was still authorized to be used/served or whether new

authorized content must be provided thereto. If it was determined that the data item corresponding to the received ETag was still authorized to be used, the first computer sent back an HTTP 304 response authorizing the downstream cache server or end-user cache to access the file content already present in order to serve it or to use it to render the webpage. On information and belief, if it had been determined that the data item corresponding to received E-tag was no longer authorized, the first computer sent back an HTTP 200 response which indicated to the downstream cache server or end-user cache that was not authorized to access the old content and must access the new authorized file content contained in the HTTP 200 response to serve it or to use it to render the webpage.

64. For further example, claim 69 covers a “system operable in a network of computers, the system comprising hardware including at least a processor, and software, in combination with said hardware.” On information and belief, Defendant has controlled the distribution of its website content across a system that included a network of computers, such as its production servers as well as origin servers, intermediate cache servers, and endpoint caches, all comprising hardware including a processor. On information and belief, Defendant has utilized software, in combination with such hardware, such as a web development framework, software utilized in implementing the HTTP web protocol, and software used on host servers that Defendant used to serve its content.

65. Claim 69 then recites the system “(a) to receive at a first computer, from a second computer, a request regarding a data item, said request including at least a content-dependent name for the data item, the content-dependent name being based at least in part on a function of the data in the data item, wherein the data used by the function to determine the content-dependent name comprises at least some of the contents of the data item, wherein the function that was used is a message digest function or a hash function, and wherein two identical data items will have the

same content-dependent name.” On information and belief, as set forth above, Defendant has caused downstream intermediate cache servers and endpoint caches to send conditional GET requests with URIs including fingerprints that are fielded by upstream cache or origin servers. On information and belief, the URIs including fingerprints were content-dependent names for a data item calculated by hashing the file’s contents; and when the file’s content changed a new content-dependent name was determined. On information and belief, in Defendant’s system, a first computer, such as the intermediate cache server or origin server, received such conditional GET requests from a second computer, such as a user browser, regarding data items, such as asset files, using content-dependent names such as URIs including fingerprints associated with the data items.

66. Claim 69 then recites “(b) in response to said request: (i) to cause the content-dependent name of the data item to be compared to a plurality of values; and (ii) to determine if access to the data item is authorized or unauthorized based on whether or not the content-dependent name corresponds to at least one of said plurality of values, and (iii) based on whether or not it is determined that access to the data item is authorized or unauthorized, to allow the data item to be provided to or accessed by the second computer if it is not determined that access to the data item is unauthorized.” On information and belief, the first computer, such as an upstream intermediate cache server or origin server, maintained a plurality of URI values associated with Defendant’s asset and webpage files; compared the URI value received in a conditional GET request from the second (downstream) computer to that plurality of URI values; that comparison allowed the first computer to determine whether the content-dependent name in the request corresponded to one of the plurality of stored URI values and to determine whether access to the data item was still authorized or not. On information and belief, in particular when there was a match, the first computer determined the associated content present at the downstream computer was still

authorized to be used/served or whether new authorized content must be provided thereto. If it was determined that the data item corresponding to the received URI including a fingerprint was still authorized to be used, the first computer has sent back an HTTP 304 response authorizing the downstream cache server or end-user cache to access the file content already present in order to serve it or to use it to render the webpage.

67. Defendant's acts of infringement have caused damage to PersonalWeb and PersonalWeb is entitled to recover from Defendant the damages sustained by PersonalWeb as a result of Defendant's wrongful acts in an amount subject to proof at trial.

THIRD CLAIM FOR RELIEF

INFRINGEMENT OF U.S. PATENT NO. 7,945,544

68. PersonalWeb repeats and realleges paragraphs 1–49, as if the same were fully stated herein.

69. On May 17, 2011, United States Patent No. 7,945,544 (the “‘544 patent”) was duly and legally issued for an invention entitled “Similarity-Based Access Control of Data in a Data Processing System.” PersonalWeb has an ownership interest in the ‘544 patent by assignment, including the exclusive right to enforce the ‘544 patent within the PersonalWeb Patent Field, and continues to hold that ownership interest in the ‘544 patent.

70. Defendant has infringed at least claims 46, 48, 52, and 55 of the ‘544 patent by its manufacture, use, sale, importation, and/or offer for sale of products or services, and/or controlling the distribution of its webpage content in the manner described herein. Defendant's infringement is literal and/or under the doctrine of equivalents and Defendant is liable for its infringement of the ‘544 patent pursuant to 35 U.S.C. § 271.

71. For example, claim 46 covers a claimed “computer-implemented method.” On information and belief, Defendant uses the claimed computer implemented method by using a system of notifications and authorizations to locate and control the distribution of data items, such as various webpage and asset files, necessary to render its webpages.

72. Claim 46 then recites the act of “(A) for each particular file of a plurality of files: (a2) determining a particular digital key for the particular file, wherein the particular file comprises a first one or more parts.” On information and belief, each of Defendant’s webpages comprises one or more asset files and has an associated webpage file, the webpage file containing the URIs having fingerprints of a plurality of asset files comprising the webpage, and once the webpage and asset files are compiled and complete, Defendant stores them on a host system. On information and belief, the webpage file’s associated ETag value is generated by applying a hash algorithm to the webpage file’s contents. On information and belief, whenever a new webpage file is generated or the webpage file’s content changes, Defendant caused an ETag to be determined and associated to the webpage file.

73. Claim 46 then recites “each part of said first one or more parts having a corresponding part value, the part value of each specific part of said first one or more parts being based on a first function of the contents of the specific part, wherein two identical parts will have the same part value as determined by the first function, and wherein the particular digital key for the particular file is determined using a second function of the one or more of part values of said first one or more parts.” On information and belief, prior to various asset files being stored on a host system, a fingerprint is generated for each of these asset files by applying a hash function to the asset file’s contents and the fingerprints are inserted into the URIs for the respective asset files. On information and belief, the webpage’s ETag value is generated by applying a second hash

function to the webpage file's contents, which include the URIs of one or more of the asset files which comprise the webpage's contents. On information and belief, because the respective asset files' URIs include the fingerprints of their content, the webpage's ETag value will change and a new associated ETag value is generated to represent the webpage's content, when the content changes and two identical webpages having the identical content represented by their webpage file will have the same ETag value.

74. Claim 46 then recites the act of "(a2) adding the particular digital key of the particular file to a database, the database including a mapping from digital keys of files to information about the corresponding files." On information and belief, Defendant caused the origin server, intermediate caches and browser caches to maintain databases/tables which mapped the ETag of each webpage's webpage file to its URI, and information about the corresponding webpage, such as, for example, information from cache-control headers for the webpage.

75. Claim 46 then recites "(B) determining a search key based on search criteria, wherein the search criteria comprise a second one or more parts, each of said second one or more parts of said search criteria having a corresponding part value, the part value of each specific part of said second one or more parts being based on the first function of the contents of the specific part, and wherein the search key is determined using the second function of the one or more of part values of said second one or more parts." On information and belief, when a downstream intermediate cache server or a browser again requested a webpage of Defendant, Defendant caused it to send a conditional GET request with an If-None-Match header with the webpage's associated ETag value. On information and belief, the received ETag value was determined using the second hash function of the webpage's webpage file, which included URIs including fingerprints for one or more of the asset files which comprised the webpage's contents.

76. Claim 46 then recites “(C) attempting to match the search key with a digital key in the database.” On information and belief, when the responding server received the webpage’s ETag value in a conditional GET request with an If-None-Match header, it compared the received ETag with the ETag it has maintained in a database/table corresponding to the URI of the webpage’s webpage file to determine if there is matching value for that webpage.

77. Claim 46 then recites “(D) if the search key matches a particular digital key in the database, providing information about the file corresponding to the particular digital key.” On information and belief, if the responding server had a matching ETag value for the webpage’s webpage file, the responding server sent an HTTP 304 response, which included information about the corresponding webpage, such as, for example, information from cache-control headers for the webpage.

78. Defendant’s acts of infringement have caused damage to PersonalWeb and PersonalWeb is entitled to recover from Defendant the damages sustained by PersonalWeb as a result of Defendant’s wrongful acts in an amount subject to proof at trial.

FOURTH CLAIM FOR RELIEF

INFRINGEMENT OF U.S. PATENT NO. 8,099,420

79. PersonalWeb repeats and realleges paragraphs 1–49, as if the same were fully stated herein.

80. On January 17, 2012, United States Patent No. 8,099,420 (the “‘420 patent”) was duly and legally issued for an invention entitled “Accessing Data in a Data Processing System.” PersonalWeb has an ownership interest in the ‘420 patent by assignment, including the exclusive right to enforce the ‘420 patent within the PersonalWeb Patent Field, and continues to hold that ownership interest in the ‘420 patent.

81. Defendant has infringed claims 25, 26, 27, 29, 30, 32, 34–36, and 166 of the ‘420 patent by its manufacture, use, sale, importation, and/or offer for sale of products or services, and/or controlling the distribution of its webpage content in the manner recited herein. Defendant’s infringement is literal and/or under the doctrine of equivalents and Defendant is liable for its infringement of the ‘420 patent pursuant to 35 U.S.C. § 271.

82. For example, claim 166 covers a “system comprising hardware, including at least a processor, and software, in combination with said hardware.” On information and belief, Defendant has controlled the distribution of its website content across a system that included hardware including a processor, such as its production servers as well as origin servers, intermediate cache servers, and endpoint caches; and software, in combination with such hardware, such as a web development framework, software utilized in implementing the HTTP web protocol, and the software used on host servers that Defendant used to serve its webpages.

83. Claim 166 then recites “(A) for a particular data item in a set of data items, said particular data item comprising a corresponding particular sequence of bits.” On information and belief, Defendant’s system has controlled the distribution of asset files and webpage files necessary to render its webpages which represent particular data items, and each of these files comprise a corresponding sequence of bits.

84. Claim 166 then recites that for the particular data item to “(a1) determine one or more content-dependent digital identifiers for said particular data item, each said content-dependent digital identifier being based at least in part on a given function of at least some of the bits in the particular sequence of bits of the particular data item, wherein two identical data items will have the same digital identifiers as determined using said given function.” On information and belief, Defendant’s system has applied hash functions to each of various Defendant’s webpage

files to all of the bits of the file's content to determine a fingerprint, an ETag, or both for the file's content; whereby two identical data items have the same ETag values and the same fingerprint values. On information and belief, fingerprints were included in files' URI and ETag values were associated with files' URIs.

85. Claim 166 then recites that for the particular data item "(a2) selectively permits the particular data item to be made available for access and to be provided to or accessed by or from at least some of the computers in a network of computers, wherein the data item is not to be made available for access or provided without authorization, as resolved based, at least in part, on whether or not at least one of said one or more content-dependent digital identifiers for said particular data item corresponds to an entry in one or more databases, each of said one or more databases comprising a plurality of identifiers, each of said identifiers in each said database corresponding to at least one data item of a plurality of data items, and each of said identifiers in each said database being based, at least in part, on at least some of the data in a corresponding data item."

86. On information and belief, Defendant's system has included one or more web servers with databases containing ETag values associated with the URIs for various of the asset and webpage files necessary to render its webpages; moreover, Defendant's system has used a system of conditional GET requests with If-None-Match headers and HTTP 304 and HTTP 200 responses containing the ETags, as described more particularly *supra*, to ensure that downstream caches only access authorized file content to either serve that file content further downstream or to use it to render Defendant's webpages. On information and belief, in particular, as more fully described *supra*, the system compared the ETag received in a given conditional GET request with

the ETags contained in the database to selectively determine whether the requesting computer could access the file content it already had or must access newly received authorized content.

87. Defendant's acts of infringement have caused damage to PersonalWeb and PersonalWeb is entitled to recover from Defendant the damages sustained by PersonalWeb as a result of Defendant's wrongful acts in an amount subject to proof at trial.

DEMAND FOR JURY TRIAL

Pursuant to Fed.R.Civ.P. 38(b), Plaintiff PersonalWeb Technologies, LLC hereby demands a trial by jury on all issues triable in this action.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff PersonalWeb requests entry of judgment in its favor and against Defendant as follows:

- a) Declaration that Defendant has infringed U.S. Patent Nos. 6,928,442, 7,802,310, 7,945,544, and 8,099,420 as described in this action;
- b) Awarding the damages arising out of Defendant's infringement of U.S. Patent Nos. 6,928,442, 7,802,310, 7,945,544, and 8,099,420, together with pre-judgment and post-judgment interest, in an amount according to proof;
- c) An award of attorneys' fees pursuant to 35 U.S.C. § 285 or as otherwise permitted by law; and
- d) For costs incurred and such other and further relief as the Court may deem just and proper.

Dated: New York, New York
August 1, 2018

Respectfully submitted,

Kent, Beatty & Gordon, LLP
/s/ Jack A. Gordon
Jack A. Gordon
Joshua B. Katz
Eleven Times Square, 10th Floor
New York, New York 10036
(212) 421-4300
jag@kbg-law.com
jbk@kbg-law.com

STUBBS ALDERTON & MARKILES, LLP
Michael A. Sherman (pro hac vice)
Jeffrey F. Gersh (pro hac vice)
Sandeep Seth (pro hac vice)
Wesley W. Monroe (pro hac vice)
Vivana Borero Hedrick
15260 Ventura Blvd., 20th Floor
Sherman Oaks, CA 91403
Telephone: (818) 444-4500
masherman@stubbsalderton.com
jgersh@stubbsalderton.com
sseth@stubbsalderton.com
wmonroe@ stubbsalderton.com
vhedrick@stubbsalderton.com

David D. Wier (pro hac vice
application to be submitted)
Vice President and Assistant General
Counsel Level 3 Communications, LLC
1025 Eldorado Boulevard
Broomfield, CO 80021
Telephone: (720) 888-3539
david.wier@level3.com