

**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.

ROCKETHUB INC., a New York corporation, and
ELEQT GROUP LTD. a United Kingdom limited
company,

Defendants.

18 Civ. 271

COMPLAINT

JURY TRIAL DEMANDED

Plaintiff PersonalWeb Technologies, LLC ("Plaintiff" or "PersonalWeb") files this Complaint for patent infringement against Defendants RocketHub Inc. and ELEQT Group Ltd. (collectively "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. 5,978,791; 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) 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 ELEQT Group Ltd. is, upon information and belief, a United Kingdom limited company having a principal place of business or regular and established place of business at 142 Cromwell Road, London, SW7 4EF, United Kingdom.

8. Defendant RocketHub Inc. is, upon information and belief, a New York corporation having a principal place of business or regular and established place of business at 340 W. 42nd Street, Suite 880, New York, NY, 10008. Upon information and belief, RocketHub Inc. is a subsidiary of ELEQT Group Ltd.

JURISDICTION AND VENUE

9. 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.*

10. Venue is proper in this federal district pursuant to 28 U.S.C. §§ 1391(b)-(c) and 1400(b) because Defendant RocketHub Inc. is incorporated in the State of New York and has a regular and established place of business in this District. Defendant has done business in this District and has committed acts of infringement in this District, entitling PersonalWeb to relief in this District.

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: by replacing conventional naming and storing conventions with system-wide “substantially unique,” content-based identifiers. Their approach assigned substantially unique identifiers to all “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.

DEFENDANT'S BACKGROUND

20. On information and belief, Defendant operates or has operated a website located at **rockethub.com**, and has done so since before expiration of the last to expire of the Patents-in-Suit, which has operated to provide webpage content to its authorized users in the manner herein described.¹ On information and belief, Defendant's webpage servers utilize a system of notifications and authorizations to control the distribution of content, *e.g.*, what webpage content may be served from webpage servers and intermediate caches and what webpage content a user's browser is authorized to use to render Defendant's webpage(s). On information and belief, Defendant's system and its associated method of providing webpage content, use CONDITIONAL GET requests with IF-NONE-MATCH headers and associated E-Tag values for each file required to render a webpage of the Defendant, including the index file for that webpage. In this manner, Defendant's system and associated method force both intermediate cache servers and end point caches to check whether it is still authorized to access the previously cached webpage files of Defendant, or whether it must access new content in rendering Defendant's webpage.

¹ While the complaint is sometimes written in the present tense, and though it is believed that Defendant's system operates in substantially the same manner currently, all specific allegations are focused on the system's operations in the relevant time period.

21. On information and belief, Defendant has thereby reduced the bandwidth required and the amount of data to be served from origination servers or intermediate cache servers to field user requests to render Defendant's webpages, because such servers only need to serve files whose content has changed. This has allowed for the efficient update of cached information only when such content has changed, thereby reducing transaction overhead and allowing the authorized content to be served from the nearest cache.

22. On information and belief, Defendant's website uses a Ruby on Rails architecture to develop and compile its webpage files that are required to render a webpage, and to generate a fingerprint of the content of each of the files when compiled. On information and belief, the fingerprint of each file that is part of the webpage's content is appended to Defendants Uniform Resource Locator ("URL") to make it a Uniform Resource Identifier ("URI") used to access the file; wherein when the file's content changes, a new fingerprint is generated and appended to its URL. On information and belief, the file fingerprint has been generated with a message digest hash function.

23. On information and belief, once Defendant's webpage files have been compiled and are complete, Defendant uploads them to an Amazon S3 host system as objects. On information and belief, Defendant has contracted with, directed and/or controlled the uploading of its files and subsequent actions that occur on the S3 host system due to Defendant's contractual choice of using content-based identifiers, e.g., fingerprints of content of files necessary to render webpages, as well as Defendant's relationship with Amazon, so that it may control its content distribution in an infringement of the Patents-In-Suit in the manner specified herein.

24. On information and belief, the object's value comprises a sequence of bits, and the object's associated E-Tag value generated, on Defendant's behalf, upon upload by applying a

hash function to the sequence of bits; wherein any two objects comprising identical content have identical associated E-Tag values. Thus, when the object's content changed, such as where the overall webpage to be rendered required updated content, and a new associated E-Tag value was generated, on Defendant's behalf, to authorize or disallow the respective service or use of the object's content by intermediate cache servers and end point caches such as browser caches.

25. On information and belief, Defendant's webpages have generally comprised one or more asset files and each webpage is represented by an index file. The index file lists each asset file needed to render the webpage to be loaded, where each of these files is uploaded as an individual object with its own URL.

26. On information and belief, when an intermediate cache server or an end point browser has requested a webpage of the Defendant for the first time, it has sent an Hyper Text Transfer Protocol ("HTTP") GET request with the webpage's URL and Defendant's origination server has responded by sending individual HTTP 200 messages respectively containing the index file and asset files necessary to render that webpage, along with their respective associated E-Tags. On information and belief, upon receipt of the HTTP 200 message, the intermediate cache server and end point browser have cached the index and asset files with their associated URI and associated E-Tag values and the browser has used them in rendering the requested web page of the defendant. On information and belief, the intermediate cache and browser caches have maintained a database/table which maps the URI of each asset/index file to its associated E-Tag.

27. On information and belief, by responding to a HTTP GET request for a given webpage by sending down the authorized index/asset file content with an associated E-Tag, Defendant has forced the browser cache and all intermediate cache servers to use the E-Tag in an HTTP CONDITIONAL GET with "IF-NONE-MATCH" protocol to re-verify that they are still

authorized to serve or use the content the next time that they are called to do so, or whether they are not still authorized to use that content and must use new content, in the manner as follows.

28. On information and belief, when the user has again requested the Defendant's webpage, the user's browser sends a CONDITIONAL GET 'IF-NONE-MATCH' request using the associated E-Tag value and the URI for the index file so as to be notified whether the browser still has Defendant's authority to render the webpage with its locally cached asset files for that webpage. On information and belief, a responding intermediate cache server having an unexpired E-Tag for that URL responds to the request by determining whether it has the same associated E-Tag value in its list of associated E-Tag values; (if it had no E-Tag value for that URL, the request was passed up to an upstream server capable of responding or, if none, to the Defendant's origination server which performed the response). On information and belief, if the responding server had webpage content for that URL and there was a match between the E-Tag it received in the request with the E-Tag it currently had associated for that URL, it has sent back an HTTP 304 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 again use the previously cached asset files to render the webpage. On information and belief, upon receipt of the HTTP Protocol 304 response, the browser accessed the locally cached asset files in rendering the webpage.

29. On information and belief, if the index file's associated E-Tag sent by the browser in the 'IF-NONE-MATCH' request did not match the associated E-Tag maintained at the responding server for that URI, the responding server sent back an HTTP 200 response along with the new index file along with its new E-Tag value. The HTTP 200 indicated to the downstream server and/or the browser that it was not authorized to use (or serve, as the case may be) the previously cached web page content but must acquire some newly authorized content. In response

to receiving the HTTP 200 message, the intermediate cache server and browser were forced to update their respective caches with the new index file and associated E-Tag. The browser read the new index file to identify the list of asset file URIs contained therein.

30. On information and belief, for any asset file URI for which it already had a cached associated E-Tag value, the browser likewise sent an 'IF_NONE_MATCH' CONDITIONAL GET request with the URI and associated E-Tag to the first intermediate cache server. On information and belief, if the responding server had an unexpired E-Tag value for the URL from that URI, the responding server compared the associated E-Tag value received in the CONDITIONAL GET with its list of associated E-Tag values for the URL from that URI. On information and belief, if there was a match, then the responding server sent an HTTP 304 message with the new max-age value and associated E-Tag value, which reauthorized the browser to use the previously cached content of that asset file to render the webpage. If there was not a match, the responding server sent an HTTP 200 message with the new content for that asset file and its new associated E-Tag value. The HTTP 200 message directed the downstream server or the browser that it was not authorized to access the previously cached content for that URL to serve it or to render the webpage. Rather, in response to receiving such a message, the browser accessed the new asset file content provided in the HTTP 200 message in rendering the webpage. Thusly the end cache and the intermediate caches in the network updated their respective databases to map the new URI to the new content and E-Tag value.

31. On information and belief, the browser has repeated this process for each asset file for which it has an associated E-Tag value.

32. On information and belief, for any asset file for which it did not have cached a previously received associated E-Tag value, the browser sent an HTTP GET request with the asset

file's URI; and the responding intermediate or origination server responded to the GET request by sending the asset file for that URI and the corresponding associated E-Tag with an HTTP 200 message. On information and belief, in response to receiving the HTTP 200 message, the browser cached the asset file and its associated E-Tag and used the newly received asset files in rendering Defendant's webpage. On information and belief, when the downstream intermediate cache or the browser was later required to again 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 render the webpage via the HTTP 304 and HTTP 200 messages.

33. On information and belief, in this manner, Defendant used E-Tag values to control the behavior of in-network intermediate cache servers and end point caches to make sure that they only accessed authorized webpage content to serve or to use.

34. On information and belief, recognizing that some out of network intermediate cache servers rendered their own E-Tag by hashing the index or asset file's URI, Defendant appended to the URL a content fingerprint that was generated by applying a hash function to the file's content. On information and belief, Defendant's appendment of the fingerprint to the URL similarly controlled the behavior of such intermediate cache servers by making sure that such intermediate cache servers always revalidated whether they are still authorized to serve the cached content or had to access new authorized content to serve or use in rendering Defendant's webpages.

FIRST CLAIM FOR RELIEF

INFRINGEMENT OF U.S. PATENT NO. 5,978,791

35. PersonalWeb repeats and realleges paragraphs 1-34, as if the same were fully stated herein.

36. On November 2, 1999, United States Patent No. 5,978,791 (the "'791 patent") was duly and legally issued for an invention entitled "Data Processing System Using Substantially Unique Identifiers to Identify Data Items, Whereby Identical Data Items Have the Same Identifiers." PersonalWeb has an ownership interest in the '791 patent by assignment, including the exclusive right to enforce the '791 patent within the PersonalWeb Patent Field, and continues to hold that ownership interest in the '791 patent. A true and correct copy of the '791 patent is attached hereto as Exhibit A.

37. Defendant has infringed at least claims 38 and 42 of the '791 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 is liable for its infringement of the '791 patent pursuant to 35 U.S.C. § 271.

38. For example, claim 38 covers "a method of locating a particular data item at a location in a data processing system." On information and belief, Defendant's website has been a data processing system and has performed the claimed method by using a system of notifications and authorizations to locate and control the distribution of data items necessary to render its webpages such as various index and asset files.

39. Claim 38 then recites the act of "(A) determining a substantially unique identifier for the data item, the identifier depending on and being determined using all of the data in the data item and only the data in the data item, whereby two identical data items in the system will have the same identifier." On information and belief, Defendant's website has determined a substantially unique identifier for the data item by calculating a hash fingerprint and E-Tags of the file's contents, and only its contents; for example, each asset file has comprised a sequence of bits and the hash of any two files comprising the identical sequence of bits has had identical

substantially unique identifiers, *e.g.*, identical fingerprints and E-Tags. If either the file's content has changed, a new substantially unique identifier has been determined for the index file both during compilation of the file and its upload as an object into Defendant's chosen content distribution system.

40. Claim 38 then recites the act of "(B) requesting the particular data item by sending the data identifier of the data item from the requester location to at least one location of a plurality of provider locations in the system." On information and belief, Defendant's use of the E-Tags and fingerprints has controlled how multiple provider locations such as origin or intermediate servers have interfaced with requester locations such as users' browsers to perform this act. On information and belief, for example, by including the E-Tags in the HTTP 200 messages and by appending the fingerprint to the URL, Defendant forced intermediate cache servers and end caches (such as used by a browser) to use `CONDITIONAL GET` requests with `IF-NONE-MATCH` headers and associated E-Tag values for each file needed to render Defendant's webpages, and forced the responding upstream servers to respond to the `CONDITIONAL GET` requests with HTTP 200 and HTTP 304 messages to verify whether they were still authorized to serve or use previously cached file contents needed to render Defendant's webpages, or must access newly provided authorized content to serve or use.

41. Claim 38 then recites the act of "(C) on at least some of the provider locations, (a) for each data item of a plurality of data items at the provider locations, (i) determining a substantially unique identifier for the data item, the identifier depending on and being determined using all of the data in the data item and only on the data in the data item, whereby two identical data items in the system will have the same identifier; and (ii) making and maintaining a set of identifiers of data items." On information and belief, Defendant's origination servers stored URI's

(that include appended content fingerprints) mapped to the authorized content and its E-Tag; and by sending the URI and the E-Tag in each HTTP 200 message containing their website content, Defendant forced intermediate cache servers and end-point caches to do the same.

42. Claim 38 then recites “(b) determining, based on the set of identifiers, whether the data item corresponding to the requested data identifier is present at the provider location.” On information and belief, by doing so, Defendant has also forced the intermediate cache servers and end point caches to send the URI and E-Tag back in CONDITIONAL GET requests with IF-NONE-MATCH headers; and thereby forced a responding server (origination or intermediate cache server) that received such a CONDITIONAL GET request from a downstream cache server or end point cache, to determine whether the file content corresponding to the received E-Tag, is present on the responding server by comparing it to the E-Tag values identifiers it has in its database to determine whether there is a match. On information and belief, this same process has been used for out-of-network intermediate cache servers that generate their own E-Tag value by hashing the URI.

43. Claim 38 then recites “(c) based on the determining, when the provider location determines that the particular data item is present at the provider location, notifying the requestor that the provider has a copy of the given data item.” On information and belief, by using this system, Defendant forced the responding server to issue an HTTP 304 message to the requesting downstream cache when there has been a match between the E-Tag in the CONDITIONAL GET request and the E-Tag in the database thereby notifying the requesting location that the same file content is present both at the responding and requesting locations and that it was therefore re-authorized to serve/use the existing content corresponding to that E-Tag value.

44. Defendant's acts of infringement have caused damage to PersonalWeb, including impairment of the value of the '791 patent, 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. 6,928,442

45. PersonalWeb repeats and realleges paragraphs 1-34, as if the same were fully stated herein.

46. 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. A true and correct copy of the '442 patent is attached hereto as Exhibit B.

47. 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 is liable for its infringement of the '442 patent pursuant to 35 U.S.C. § 271.

48. 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 origin

servers, intermediate cache servers and end point caches used by browsers rendering Defendant's webpages.

49. 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 obtained E-Tags and fingerprints for its index and asset files used to render its webpages using a hash function, wherein the E-Tag and fingerprint has been based on the contents of the particular file. Moreover, Defendant caused the intermediate caches servers and end point caches to obtain the E-tags and URIs (which contain the fingerprint) in HTTP Protocol 200 messages sent from Defendants origination servers. On information and belief, Defendant caused intermediate cache servers and its origination servers to obtain E-tags and URIs (which the fingerprint) in CONDITIONAL GET messages from end point and intermediate caches, as described *supra*. On information and belief, by also inserting the fingerprint into the URI for the file, Defendant caused certain out-of-network intermediate cache servers (that obtain their own E-Tag by hashing the URI) to make content based E-Tags, so that when the content of the file has changed, these out-of-network caches were caused to verify that they already had or needed Defendant's latest authorized content in the same manner outlined *supra* for in-network servers via the HTTP 200 and HTTP 304 message system or to notify such caches that they already had and were still authorized to access the previously cached content or to provide such latest authorized content.

50. 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 origination servers and the intermediate cache servers

in-between an end point cache and one of its origination servers, in response to receiving a CONDITIONAL GET request with the IF-NONE-MATCH header, to compare the E-Tag in the CONDITIONAL GET to the E-Tags of files it has present and determine whether a copy of the content having that E-Tag is present.

51. Claim 10 then recites the act of “determining whether a copy of the data file that is present on a 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 origination or intermediate cache server determined that the copy of the file present at the downstream intermediate cache server and/or the end-point 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 end-point cache was an unauthorized or unlicensed copy of the data file.

52. Defendant's acts of infringement caused damage to PersonalWeb, including impairment of the value of the '442 patent, 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. 7,802,310

53. PersonalWeb repeats and realleges paragraphs 1-34, as if the same were fully stated herein.

54. 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. A true and correct copy of the '310 patent is attached hereto as Exhibit C.

55. Defendant has infringed at least claims 20, 69 and 71 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 is liable for its infringement of the '310 patent pursuant to 35 U.S.C. § 271.

56. For 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 network of computers, such as its origin servers, intermediate cache servers and end-point caches, comprising hardware including a processor. On information and belief, Defendant has utilized the Ruby on Rails code, software utilized in implementing the HTTP web protocol, and both hardware and software hosted on the Amazon S3 hosting system that Defendant uses to serve its content.

57. Claim 69 then recites the system "(a)...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 end-point caches to send CONDITIONAL GET requests with IF-NONE-MATCH headers containing E-Tags that are fielded by upstream cache or origination servers. On information and belief, the E-Tags have been content-dependent names for a data item calculated by hashing the file's contents; and when the file's content has changed a new content-dependent name has been determined. On information and belief, in Defendant's system, a first computer, such as the intermediate cache server or origination server, received CONDITIONAL GET requests from a second computer, such as a user browser, regarding data items, such as index or asset files, using content-dependent names (E-tags) associated with the data items.

58. Claim 69 then recites "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 origination server, has maintained a plurality of E-tag values associated with Defendant's asset and index files; has compared the E-tag received in the CONDITIONAL GET request from the second (downstream) computer to that plurality of values; that comparison having allowed the first computer to determine whether the content-dependent name in the request corresponded to one of the plurality of stored 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 received E-tag was not still unauthorized to be used, the first computer has sent back an HTTP 304 message 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 has been determined that the data item corresponding to received E-tag was unauthorized, the first computer has sent back an HTTP 200 message 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 message to serve it or to use it to render the webpage.

59. Defendant's acts of infringement have caused damage to PersonalWeb, including impairment of the '310 patent, 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. 7,945,544

60. PersonalWeb repeats and realleges paragraphs 1-34, as if the same were fully stated herein.

61. 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. A true and correct copy of the '544 patent is attached hereto as Exhibit D.

62. Defendant has infringed at least claims 46, 48, 49, 52, 55 and 56 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 is liable for its infringement of the '791 patent pursuant to 35 U.S.C. § 271.

63. 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 index and asset files, necessary to render its webpages.

64. 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 an associated index file, the index file lists the URI’s of a plurality of asset files comprising the webpage, and once the asset files are compiled and complete, Defendant uploads them to the S3 host system as objects. On information and belief, the object’s associated E-Tag value is generated by applying a hash algorithm to the object’s contents, wherein any two objects comprising the identical content will have identical associated E-Tag values. On information and belief, whenever a new object is uploaded to an S3 server or the object’s content changes, Defendant determines and associates an E-Tag for the object by receiving or identifying the associated E-Tag value generated at the time of upload. On information and belief, this applies also to webpage’s E-tag, which is generated when its index file is uploaded, and this E-Tag value is a search key to contents of the webpage.

65. 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, the webpage’s E-Tag value is generated by applying a second hash function to the index file’s contents, which consist of the URI’s of one or more of the asset files which comprise the webpage’s contents. On information and belief, because the respective asset file’s URI’s include the fingerprints of their content, the webpage’s E-Tag value will change and a new associated E-Tag value is generated to represent the webpage’s content, when the content changes and two identical webpage’s having the identical content represented by their index file will have the same E-Tag value.

66. 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, the origination server, intermediate caches and browser caches maintain a database/table which maps the E-Tag of each webpage’s index file to its URI, storage location and information about the corresponding file, and whenever a new index file is uploaded to an S3 server for that webpage (e.g. when the webpage’s content changes and therefore it’s index file’s content changes), Defendant determines and associates a new E-Tag for the index file by receiving or identifying the associated E-Tag value generated at the time of upload. On information and belief, this associated E-Tag is added to the database/table and maps to the corresponding file information.

67. 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 requests a webpage of Defendant, it sends a CONDITIONAL GET request with IF-NONE-MATCH with the webpage’s associated E-Tag value. On information and belief, the receiving server will determine the received E-Tag value and use it as a search key to check whether the webpage’s content has changed.

68. 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 receives the webpage’s E-Tag value in a CONDITIONAL GET request with IF-NONE-MATCH header, it compares the received E-Tag with the current list of associated E-Tags it has maintained in a database/table to determine if there is matching value for that webpage.

69. 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 has a matching unexpired E-Tag value for the webpage, the responding server sends an HTTP 304 message, which informs the downstream server and/or browser that the content of the webpage has not changed, and that the downstream server and/or browser is reauthorized to access all the previously cached content necessary to render the webpage. On information and belief, if there is not a match, the responding server sends an HTTP 200 (Modified) message with the new index file for that webpage and its new associated E-Tag value, and the HTTP 200 message informs the downstream server and/or browser that it is not authorized to access all the previously cached asset files need to render that webpage. On information and belief, the receipt of the HTTP 200 message with the webpage’s new index file

and E-Tag informs the downstream server and/or browser that it is authorized to use the new index file provided in the HTTP 200 message in determining what parts of the webpage it already has cached that it can use and which new parts it needs to render the webpage. On information and belief, the end cache and the intermediate caches in the content delivery chain also update their respective databases to map the new index file URI and contents to the new index content and E-Tag value.

70. On information and belief, in this manner, the webpage's E-tag value informs the downstream cache server or end point cache via the HTTP 304 and HTTP 200 messages whether it is authorized to serve/use all of the previously cached parts of the webpage, or must use CONDITIONAL GET request(s) with IF-NONE-MATCH header(s) and E-Tags at the asset file level to determine which parts of the webpage it is re-authorized to use/serve, and what newly authorized parts of the webpage it must first obtain.

71. Defendant's acts of infringement have caused damage to PersonalWeb, including impairment of the value of the '544 patent, 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.

FIFTH CLAIM FOR RELIEF

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

72. PersonalWeb repeats and realleges paragraphs 1-34, as if the same were fully stated herein.

73. 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. A true and correct copy of the '420 patent is attached hereto as Exhibit E.

74. Defendant has infringed claims 25, 26, 27, 29, 30, 32-36 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 is liable for its infringement of the '420 patent pursuant to 35 U.S.C. § 271.

75. 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’s system has comprised hardware including a processor, such as its webpage servers; and software including the Ruby on Rails web code used in making its webpages and the Amazon S3 hosting system which have been used in combination with its hardware.

76. 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 index files necessary to render its webpage’s which represent particular data items, and each of these files comprise a corresponding sequence of bits.

77. 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 the Defendant’s webpage

files to all of the bits of the file's content to determine both a fingerprint and an E-tag for the file's content; whereby two identical data items have the same E-tag and fingerprint values. On information and belief, the fingerprint was appended to the file's URL (herein, the URL plus the appended fingerprint is referred to as the URI) and the E-Tag value was associated with the file's URL.

78. 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."

79. On information and belief, Defendant's system has included one or more webpage servers with databases containing E-tag values associated with the various URL's and/or URI's for all of the asset and manifest/index files necessary to render its webpages; moreover, Defendant's system has used a system of CONDITIONAL GET with IF-NONE-MATCH header, HTTP 304 and HTTP 200 messages containing the E-Tags, 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 E-Tag received in a

given CONDITIONAL GET message with the E-Tags 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.

80. Defendant's acts of infringement have caused damage to PersonalWeb, including impairment of the '420 patent, 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. 5,978,791, 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. 5,978,791, 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
January 11, 2018

Respectfully submitted,

KENT, BEATTY & GORDON, LLP

/s/ Joshua B. Katz

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

IP LAW GROUP, LLP

Michael A. Sherman (pro hac vice
application to be submitted)

Jeffrey F. Gersh (pro hac vice
application to be submitted)

Wesley W. Monroe (pro hac vice
application to be submitted)

15030 Ventura Blvd., #166

Sherman Oaks, CA 91403

Telephone: (818) 444-9270

michaelsherman@iplawllp.com

jeffgersh@iplawllp.com

wesleymonroe@iplawllp.com

Sandeep Seth (pro hac vice
application to be submitted)

SETHLAW

Two Allen Center

1200 Smith Street, Suite 1600

Houston, Texas 77002

Telephone: (713) 244-5017

ss@sethlaw.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

