

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

DISNEY ENTERTAINMENT & SPORTS LLC,
Petitioner

v.

ADEIA MEDIA HOLDINGS LLC,
Patent Owner.

Case No. IPR2026-00051
U.S. Patent No. 8,280,987

**DECLARATION OF HENRY HOUH IN SUPPORT OF PETITION FOR
INTER PARTES REVIEW OF U.S. PATENT NO. 8,290,987**

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. BACKGROUND AND QUALIFICATIONS	2
III. INFORMATION CONSIDERED	8
IV. RELEVANT LEGAL STANDARDS	8
A. Claim Interpretation	8
B. Perspective of One of Ordinary Skill in the Art.....	9
C. Obviousness.....	9
V. LEVEL OF ORDINARY SKILL IN THE ART	12
VI. SUMMARY OF MY OPINIONS	13
VII. TECHNOLOGICAL BACKGROUND	14
A. Cloud data persistence system for distribution of data	14
B. Origin servers for storing a plurality of data objects comprising a link to a plurality of media assets	15
C. Cache servers that store a plurality of data objects and request a data object from the origin server if the data object is not present at the cache server	20
D. Application servers that store a plurality of data objects and request the data object from the cache servers if the data object is not present at the application server	22
E. Data objects are at least partially determined based on frequency of requests and/or when requests for those data objects were last received at the server	25
F. Data objects at the origin server comprise all of the data objects in the cache and application servers.....	27
VIII. THE '987 PATENT AND ITS FILE HISTORY	28
IX. CLAIM CONSTRUCTION	30
A. “Data object” (Claims 1, 3-13).....	30
B. “Data persistence cache servers” (Claim 7).....	31
X. OVERVIEW OF THE APPLIED PRIOR ART REFERENCES	31

TABLE OF CONTENTS
(continued)

	Page
A. Dilley (Ex-1005).....	31
B. Russell (Ex-1006).....	34
C. Menon (Ex-1007)	37
XI. DETAILED EXPLANATION OF THE UNPATENTABILITY GROUNDS	38
A. Grounds 1-2: Claims 1-13 are obvious over Dilley alone or Dilley in view of Menon.	38
1. A POSITA would have been motivated to combine Dilley’s teachings with Menon’s teachings and would have had a reasonable expectation of success.	38
2. Independent Claim 1	40
a. Element 1[pre]: A cloud data persistence system for distribution of data, the cloud data persistence system comprising:	40
b. Element 1[a][1]: a first server which stores a first plurality of data objects and is communicatively coupled with a plurality of cache servers, wherein:	42
c. Element 1[a][2]: the first server is located at a first location;	46
d. Element 1[a][3]: each of the first plurality of data objects comprise a link to one of a plurality of media assets; and	47
e. Element 1[a][4]: the plurality of media assets are stored at a second location separate from the first location;	48
f. Element 1[b][1]: a second server that stores a second plurality of data objects and is configured to request a data object from the first server if the data object is not present in the second plurality of data objects, wherein:	51

TABLE OF CONTENTS
(continued)

	Page
g. Element 1[b][2]: the second server is one of the plurality of cache servers; and.....	53
h. Element 1[b][3]: the second server is communicatively coupled with a first plurality of application servers;	54
i. Element 1[c][1]: a third server that stores a third plurality of data objects and is configured to request the data object from the first server if the data object is not present in the third plurality of data objects, wherein:	56
j. Element 1[c][2]: the third server is one of the plurality of cache servers; and.....	58
k. Element 1[c][3]: the third server is communicatively coupled with a second plurality of application servers; and.....	58
l. Element 1[d][1]: a fourth server that stores a fourth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fourth plurality of data objects, wherein:	60
m. Element 1[d][2]: the fourth server is one of the first plurality of application server[s]; and	63
n. Element 1[d][3]: the fourth server is configured to receive a request from a first client for the data object;.....	63
o. Element 1[e][1]: a fifth server that stores a fifth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fifth plurality of data objects, wherein:	65
p. Element 1[e][2]: the fifth server is one of the first plurality of application servers; and	67

TABLE OF CONTENTS

(continued)

	Page
q. Element 1[e][3]: the fifth server is configured to receive a request from a second client for the data object;.....	67
r. Element 1[f][1]: a sixth server that stores a sixth plurality of data objects and is configured to request the data object from the third server if the data object is not present in the sixth plurality of data objects, wherein:	68
s. Element 1[f][2]: the sixth server is one of the second plurality of application servers; and	70
t. Element 1[f][3]: the sixth server is configured to receive a request from a third client for the data object; and.....	71
u. Element 1[g][1]: a seventh server that stores a seventh plurality of data objects and is configured to request the data object from the third server if the data object is not present in the seventh plurality of data objects, wherein:	72
v. Element 1[g][2]: the seventh server is one of the second plurality of application servers; and	74
w. Element 1[g][3]: the seventh server is configured to receive a request from a fourth client for the data object.....	75
3. Independent Claim 5	76
a. Element 5[pre]: A cloud data persistence system for distribution of data, the cloud data persistence system comprising:	76
b. Element 5[a][1]: a first server which stores a first plurality of data objects and is communicatively coupled with a second server and a third server, wherein:	76
c. Element 5[a][2]: the first server is located at a first location;	78

TABLE OF CONTENTS

(continued)

	Page
d. Element 5[a][3]: each of the first plurality of data objects comprises a link to a media asset of a plurality of media assets;	78
e. Element 5[a][4]: the plurality of media assets is located at a second location ; and the second location is distinct from the first location;.....	78
f. Element 5[b][1]: the second server stores a second plurality of data objects and is configured to request a data object from the first server if the data object is not present in the second plurality of data objects, ...	78
g. Element 5[b][2]: wherein the second server is communicatively coupled with a fourth server and fifth server; and.....	78
h. Element 5[c][1]: the fourth server stores a third plurality of data objects and is configured to request the data object from the second server if the data object is not present in the third plurality of data objects, ...	80
i. Element 5[c][2] wherein the fourth server is configured to receive a request from a first client for the data object.	80
4. Dependent Claims 2-4 and 6-13	80
a. Claim 2: The cloud data persistence system of claim 1, wherein the second location comprises a content delivery network.	80
b. Claim 3.....	81

TABLE OF CONTENTS
(continued)

	Page
c. Claim 4: The cloud data persistence system of claim 1, wherein the first plurality of data objects comprises all data objects in the second plurality of data objects, the third plurality of data objects, the fourth plurality of data objects, the fifth plurality of data objects, the sixth plurality of data objects, and the seventh plurality of data objects.	84
d. Claim 6: The cloud data persistence system of claim 5 wherein the first plurality of data objects comprises all data objects present in the second plurality of data objects and the third plurality of data objects.	86
e. Claim 7.....	86
f. Claim 8.....	88
g. Claim 9: The cloud data persistence system of claim 5, wherein the request from the first client for the data object comprises a universal resource indicator.	92
h. Claim 10: The cloud data persistence system of claim 5 wherein each of the first plurality of data objects comprise a universal resource locator and a title linked to a media asset of the plurality of media assets.	94
i. Claim 11: The cloud data persistence system of claim 5, wherein data objects in the second plurality of data objects and the third plurality of data objects are at least partially determined based on a frequency of requests for those data objects received from clients.	96

TABLE OF CONTENTS

(continued)

	Page
j. Claim 12: The cloud data persistence system of claim 5, wherein data objects in the second plurality of data objects and the third plurality of data objects are at least partially determined based on when requests for those data objects were last received from clients.	97
k. Claim 13: The cloud data persistence system of claim 5, wherein the first plurality of data objects comprises all data objects in the second plurality of data objects, the third plurality of data objects, the fourth plurality of data objects, and the fifth plurality of data objects.	97
B. Grounds 3-4: Claims 1-13 are obvious over Russell alone or Russell in view of Menon.....	97
1. A POSITA would have been motivated to combine Russell’s teachings with Menon’s teachings and would have had a reasonable expectation of success.	97
2. Independent Claim 1	100
a. Element 1[pre]: A cloud data persistence system for distribution of data, the cloud data persistence system comprising:	100
b. Element 1[a][1]: a first server which stores a first plurality of data objects and is communicatively coupled with a plurality of cache servers, wherein:	101
c. Element 1[a][2]: the first server is located at a first location;	105
d. Element 1[a][3]: each of the first plurality of data objects comprise a link to one of a plurality of media assets; and	106
e. Element 1[a][4]: the plurality of media assets are stored at a second location separate from the first location;	108

TABLE OF CONTENTS
(continued)

	Page
f. Element 1[b][1]: a second server that stores a second plurality of data objects and is configured to request a data object from the first server if the data object is not present in the second plurality of data objects, wherein:	109
g. Element 1[b][2]: the second server is one of the plurality of cache servers; and	111
h. Element 1[b][3]: the second server is communicatively coupled with a first plurality of application servers;	112
i. Element 1[c][1]: a third server that stores a third plurality of data objects and is configured to request the data object from the first server if the data object is not present in the third plurality of data objects, wherein:	115
j. Element 1[c][2]: the third server is one of the plurality of cache servers; and	116
k. Element 1[c][3]: the third server is communicatively coupled with a second plurality of application servers; and	117
l. Element 1[d][1]: a fourth server that stores a fourth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fourth plurality of data objects, wherein:	118
m. Element 1[d][2]: the fourth server is one of the first plurality of application server; and	121
n. Element 1[d][3]: the fourth server is configured to receive a request from a first client for the data object;	121

TABLE OF CONTENTS

(continued)

	Page
o. Element 1[e][1]: a fifth server that stores a fifth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fifth plurality of data objects, wherein:	123
p. Element 1[e][2]: the fifth server is one of the first plurality of application servers; and	125
q. Element 1[e][3]: the fifth server is configured to receive a request from a second client for the data object;.....	125
r. Element 1[f][1]: a sixth server that stores a sixth plurality of data objects and is configured to request the data object from the third server if the data object is not present in the sixth plurality of data objects, wherein:	126
s. Element 1[f][2]: the sixth server is one of the second plurality of application servers; and	128
t. Element 1[f][3]: the sixth server is configured to receive a request from a third client for the data object; and.....	128
u. Element 1[g][1]: a seventh server that stores a seventh plurality of data objects and is configured to request the data object from the third server if the data object is not present in the seventh plurality of data objects, wherein:	129
v. Element 1[g][2]: the seventh server is one of the second plurality of application servers; and	131
w. Element 1[g][3]: the seventh server is configured to receive a request from a fourth client for the data object.....	131
3. Independent Claim 5	132

TABLE OF CONTENTS

(continued)

	Page
a. Element 5[pre]: A cloud data persistence system for distribution of data, the cloud data persistence system comprising:	132
b. Element 5[a][1]: a first server which stores a first plurality of data objects and is communicatively coupled with a second server and a third server, wherein:	133
c. Element 5[a][2]: the first server is located at a first location;	134
d. Element 5[a][3]: each of the first plurality of data objects comprises a link to a media asset of a plurality of media assets;	134
e. Element 5[a][4]: the plurality of media assets is located at a second location ; and the second location is distinct from the first location;.....	134
f. Element 5[b][1]: the second server stores a second plurality of data objects and is configured to request a data object from the first server if the data object is not present in the second plurality of data objects, ...	135
g. Element 5[b][2]: wherein the second server is communicatively coupled with a fourth server and fifth server; and.....	135
h. Element 5[c][1]: the fourth server stores a third plurality of data objects and is configured to request the data object from the second server if the data object is not present in the third plurality of data objects, ...	136
i. Element 5[c][2] wherein the fourth server is configured to receive a request from a first client for the data object.	136
4. Dependent Claims 2-4 and 6-13	137

TABLE OF CONTENTS
(continued)

	Page
a. Claim 2: The cloud data persistence system of claim 1, wherein the second location comprises a content delivery network.	137
b. Claim 3.....	138
c. Claim 4: The cloud data persistence system of claim 1, wherein the first plurality of data objects comprises all data objects in the second plurality of data objects, the third plurality of data objects, the fourth plurality of data objects, the fifth plurality of data objects, the sixth plurality of data objects, and the seventh plurality of data objects.	140
d. Claim 6: The cloud data persistence system of claim 5 wherein the first plurality of data objects comprises all data objects present in the second plurality of data objects and the third plurality of data objects.	142
e. Claim 7.....	142
f. Claim 8.....	144
g. Claim 9: The cloud data persistence system of claim 5, wherein the request from the first client for the data object comprises a universal resource indicator.	147
h. Claim 10: The cloud data persistence system of claim 5 wherein each of the first plurality of data objects comprise a universal resource locator and a title linked to a media asset of the plurality of media assets.	148
i. Claim 11: The cloud data persistence system of claim 5, wherein data objects in the second plurality of data objects and the third plurality of data objects are at least partially determined based on a frequency of requests for those data objects received from clients.	150

TABLE OF CONTENTS
(continued)

	Page
j. Claim 12: The cloud data persistence system of claim 5, wherein data objects in the second plurality of data objects and the third plurality of data objects are at least partially determined based on when requests for those data objects were last received from clients.	150
k. Claim 13: The cloud data persistence system of claim 5, wherein the first plurality of data objects comprises all data objects in the second plurality of data objects, the third plurality of data objects, the fourth plurality of data objects, and the fifth plurality of data objects.	150
XII. CONCLUSION.....	150

LIST OF EXHIBITS

Ex. No.	Description
Ex-1001	U.S. Patent No. 8,280,987 (“the ’987 Patent”)
Ex-1002	Declaration of Dr. Houh
Ex-1003	Curriculum Vitae of Dr. Houh
Ex-1004	Prosecution History of U.S. Patent No. 8,280,987
Ex-1005	U.S. Patent Publication No. 2007/0055764 to Dilley (“Dilley”)
Ex-1006	U.S. Patent Publication No. 2002/0069420 to Russell (“Russell”)
Ex-1007	U.S. Patent Publication No. 2002/0152318 to Menon (“Menon”)
Ex-1008	INTENTIONALLY LEFT BLANK
Ex-1009	INTENTIONALLY LEFT BLANK
Ex-1010	INTENTIONALLY LEFT BLANK
Ex-1011	INTENTIONALLY LEFT BLANK
Ex-1012	INTENTIONALLY LEFT BLANK
Ex-1013	INTENTIONALLY LEFT BLANK
Ex-1014	INTENTIONALLY LEFT BLANK
Ex-1015	INTENTIONALLY LEFT BLANK
Ex-1016	INTENTIONALLY LEFT BLANK
Ex-1017	INTENTIONALLY LEFT BLANK
Ex-1018	INTENTIONALLY LEFT BLANK
Ex-1019	INTENTIONALLY LEFT BLANK
Ex-1020	Claim Mapping Table

Ex. No.	Description
Ex-1021	U.S. Patent No. 8,627,446 to Eaton (“Eaton”)
Ex-1022	U.S. Patent No. 8,180,720 to Kovacs (“Kovacs”)
Ex-1023	U.S. Patent Publication No. 2011/0289126 to Aikas (“Aikas”)
Ex-1024	U.S. Patent No. 7,428,540 to Coates (“Coates”)
Ex-1025	U.S. Patent Publication No. 2011/0276623 to Girbal (“Girbal”)
Ex-1026	U.S. Patent Publication No. 2003/0046239 to Geilfuss (“Geilfuss”)
Ex-1027	U.S. Patent Publication No. 2010/0057894 to Glasser (“Glasser”)
Ex-1028	U.S. Patent Publication No. 2010/0180208 to Kasten (“Kasten”)
Ex-1029	U.S. Patent No. 9,167,036 to Seed (“Seed”)
Ex-1030	T. Berners-Lee, IETF Request for Comments: 3986 (January 2005) (“IETF RFC 3986”)
Ex-1031	U.S. Patent Publication No. 2010/0250772 to Mao (“Mao”)
Ex-1032	U.S. Patent No. 5,796,952 to Davis (“Davis”)
Ex-1033	U.S. Patent No. 6,996,616 to Leighton (“Leighton”)
Ex-1034	<i>Adeia Techs., Inc. v. The Walt Disney Co.</i> , No. 1:24-cv-01231, Defendants’ Supplemental Disclosure of Proposed Claim Constructions (D. Del. Sept. 16, 2025)
Ex-1035	<i>Adeia Techs., Inc. v. The Walt Disney Co.</i> , No. 1:24-cv-01231, Plaintiffs’ Initial Proposed Claim Constructions (D. Del. Sept. 5, 2025)
Ex-1036	U.S. Patent Publication No. 2007/0299681 to Plastina (“Plastina”)

I. INTRODUCTION

1. My name is Henry Houh. I am presently an independent technical consultant with extensive experience in the field of web based systems, streaming multimedia systems, distributed systems, and network architectures.

2. I have been retained to provide this Declaration by Disney Entertainment & Sports LLC (“Petitioner”) as an independent expert consultant in this *inter partes* review (“IPR”) proceeding before the United States Patent and Trademark Office (“PTO”).

3. I have been asked by Petitioner Counsel (“Counsel”) to consider whether certain references disclose, teach and/or suggest the features recited in Claims 1-13 of U.S. Patent No. 8,280,987 (“the ’987 Patent”) (Ex-1001), which I understand is currently assigned to currently assigned to Adeia Guides Inc. (“PO”). My opinions and the bases for my opinions are set forth below.

4. I am being compensated at my ordinary and customary consulting rate for my work, which is \$820 per hour. My compensation is in no way contingent on the nature of my findings, the presentation of my findings in testimony, or the outcome of this or any other proceeding. I have no other direct financial interest in this proceeding.

II. BACKGROUND AND QUALIFICATIONS

5. I have a Bachelor of Science Degree in Electrical Engineering and Computer Science from Massachusetts Institute of Technology (MIT), which I received in June 1989. I also have a Bachelor of Science Degree in Physics from MIT which I received in February 1990. I have a Master of Science Degree in Electrical Engineering and Computer Science from MIT which I received in February 1991. I have a PhD in Electrical Engineering and Computer Science from MIT which I received in February 1998. My thesis was directed to “Designing Networks for Tomorrow’s Traffic,” and was supervised by Professor David Tennenhouse and Professor John Guttag.

6. My professional career has spanned more than 25 years. As set forth in my curriculum vitae, a copy of which is appended to this note, during these years I have gained extensive experience in web sites, communications systems, networking, distributed systems, protocols and streaming audio/video communications. I am currently employed by Houh Consulting Inc. as an independent technical consultant.

7. I first entered telecommunications in 1987 when I worked as a summer intern at AT&T Bell Laboratories as part of a five-year dual degree program at MIT. I continued to work at AT&T Bell Laboratories as part of this MIT program. While at MIT, I was a teaching assistant (“TA”) in the Electrical Engineering and Computer

Science Department's core Computer Architectures course. I first was a TA as a senior for a role typically reserved for graduate students. I later became head TA. The course covered various topics in computer architectures. As a TA, I helped write homework assignments, lab assignments, and exams. I also taught in the recitation sections.

8. Later, as part of my doctoral research at MIT from 1991-1998, I was a research assistant in the Telemedia Network Systems ("TNS") group at the Laboratory for Computer Science. The TNS group built a high-speed gigabit network and created applications that ran over the network. Example applications included ones for remote video capture, processing, and display of video on computer terminals. In addition to working on the design of core network components (which was based on FPGAs), designing and building the high-speed links (which included programmable logic devices), and designing and writing the device drivers for the interface cards (which included programmable logic devices), I also set up the group's web server. Also, I helped to maintain, install, and upgrade the networking devices used within the group, along with other graduate students. When I started with the TNS group, the networking infrastructure consisted of shared coaxial cable-based Ethernet, with bridges and repeaters interconnecting various network segments. I became familiar with the Ethernet specifications in order that we not configure our network improperly, typically by chaining too many

repeaters together. We gradually replaced the coaxial Ethernet infrastructure with twisted pair 10BASE-T wiring that were interconnect via hub repeaters, for which we also had to mind the Ethernet specifications. Eventually, many of the hub repeaters were upgrade to switches. I also studied the various parts of the Ethernet specifications both independently and through graduate course work in networking.

9. I also helped to build the web pages that initiated the above-mentioned video sessions via a web interface. Vice President Al Gore visited our group in 1996 and received a demonstration of—and remotely drove—a radio-controlled toy car with a wireless video camera mounted on it that was built by our group. This toy car device received commands transmitted over a network from a remote computer, and video data from the toy car was transmitted wirelessly then over a computer network back to the user controller. On occasion, we allowed users visiting our web site to drive the toy car from their remote computer while they watched the video on their computer. The video stream was encoded by TNS-designed hardware, streamed over the TNS-designed network, and displayed using TNS-designed software.

10. I defended and submitted my Ph.D. thesis, titled “Designing Networks for Tomorrow’s Traffic,” in January 1998. As part of my thesis research, I analyzed local area and wide area flows to show a more efficient method for routing packets in a network, based on traffic patterns at the time. The traffic flow data included

Ethernet, IP, TCP or UDP, and RTP header information, which I analyzed to come to the conclusions in my thesis.

11. While I was in graduate school, I started a company that provided web site design services and targeted advertising opportunities. We also sold targeted advertising placements. As the web became popular first at colleges and universities, we focused initially on building college recruiting web sites and sold advertisements to companies wishing to target students to recruit at specific colleges. We developed technology that utilized the web user's college affiliation to customize advertisements targeted to students at that college.

12. From 1997 to 1999, I was a Senior Scientist and Engineer at NBX Corporation, a start-up that made business telephone systems for streaming packetized audio over data networks instead of using traditional telephone lines. NBX was later acquired by 3Com Corporation and the phone system is still used today by numerous businesses.

13. As part of my work at NBX, I designed the core audio reconstruction algorithms for the telephones, as well as the packet transmission algorithms. I also designed and validated the core packet transport protocol used by the phone system. The protocol was used for all signaling in the phone system, including for the setup of conference calls. The NBX system also featured a computer interface for initiating phone calls, which could also initiate conference calls. The NBX system also

supported the Telephony Application Programming Interface (“TAPI”) that allowed other computer programs to integrate with our system telephony features. We obtained U.S. Patent No. 6,967,963, entitled “Telecommunication method for ensuring on-time delivery of packets containing time-sensitive data,” as part of this work. The first release of the phones had integrated Ethernet repeater hubs embedded in them, and I worked with the documentation and field support people to develop sets of guidelines so that customers would not build their networks to violate the Ethernet specifications, typically by chaining too many hub repeaters and NBX phones together. The first release of the phones also worked using Ethernet-layer packets (without the use of IP). I also programmed the first prototype of our phone which communicated using IP, and I demonstrated our IP phones working over the Internet when we attended a trade show in California. The phone connected over the Internet to our headquarters in Andover, MA.

14. From 1999-2004, I was employed by Empirix or its predecessor company, Teradyne. Teradyne had traditionally made products to test chips and hardware, but the founder and board chairman, Alex d’Arbeloff wanted to expand the company’s products to include software systems testing. Teradyne made a number of acquisitions to achieve this goal, and I was hired to work for Mr. d’Arbeloff to work with all the acquired software test divisions. These divisions were later spun-out of Teradyne into a separate company, Empirix. Empirix was a

leader in test tools for software systems including telecommunications protocols and systems, providing functional testing tools as well as load testing tools. From 2000-2001, I conceived and built a test platform for testing Voice-over-IP (VoIP). The first application on this new test platform was a cloud emulator for simulating the effects of transmitting VoIP over a busy network. The test platform was based on a network processor chip, which could be programmed to cut-through packets while processing packet data such as various protocol layer headers including addresses included therein and even packet data contents. I also designed a protocol analyzer built on the same platform. The application captured and performed protocol decoding at various layers in the protocol stacks of captured packets, including detailed Ethernet header decoding, IP header decoding, TCP header decoding, UDP header decoding, RTP header decoding, and many other specified protocols. The application was also designed to reconstruct entire conversations that spanned multiple packets.

15. In 2006, as part of my role at BBN Technologies, I helped found PodZinger Inc., now known as RAMP Inc. PodZinger utilized BBN's speech recognition algorithms to search through the spoken words in audio and video segments. While I was Vice President of Operations and Technology, PodZinger followed its initial prototype with a full streaming audio and video search solution. I also created a social networking web site, which BBN sold to a venture-funded

startup company. In the process of creating the web site, I designed and specified the authentication and authorization protocols.

III. INFORMATION CONSIDERED

16. In preparation for my Declaration, I have considered the materials I discuss in my Declaration, including, for example, the '987 Patent, the references cited by the '987 Patent, the prosecution history of the '987 Patent (including the references cited therein), various background articles and materials referenced in my Declaration, and the prior art references I identify in my Declaration. In addition, my opinions are further based on my education, training, experience, and knowledge in the relevant field.

IV. RELEVANT LEGAL STANDARDS

17. I am not an attorney and offer no legal opinions. For the purposes of my Declaration, I have been informed about certain aspects of the law that are relevant to my analysis, as summarized below.

A. Claim Interpretation

18. I have been informed and understand that in an IPR proceeding, claims are to be interpreted according to the *Phillips* claim construction standard. I understand that under this standard, the meaning of claim terms is considered from the viewpoint of a hypothetical “person of ordinary skill in the art” (“POSITA”) at the time of the alleged invention. I have been informed and understand that claim

construction is a matter of law and that the final claim constructions for this proceeding will be determined by the Patent Trial and Appeal Board (“PTAB”).

19. I provide my opinions regarding terms that may require explicit construction to resolve the grounds I present in my Declaration in Section IX.

B. Perspective of One of Ordinary Skill in the Art

20. I have been informed and understand that a patent is to be understood from the perspective of a hypothetical POSITA. Such an individual is considered to possess normal skills and knowledge in a particular technical field (as opposed to being a genius). I understand that in considering what the claims of a patent require, what was known prior to that patent, what a prior art reference discloses, and whether an invention is obvious or not, one must use the perspective of such a POSITA.

C. Obviousness

21. I have been informed and understand that a patent claim is obvious under 35 U.S.C. §103, and therefore invalid, if the claimed subject matter, as a whole, would have been obvious to a POSITA as of the priority date of the patent based on one or more prior art references and/or the knowledge of a POSITA.

22. I understand that an obviousness analysis must consider (1) the scope and content of the prior art, (2) the differences between the claims and the prior art, (3) the level of ordinary skill in the pertinent art, and (4) secondary considerations,

if any, of non-obviousness (such as unexpected results, commercial success, long-felt but unmet need, failure of others, copying by others, and skepticism of experts).

23. I understand that a prior art reference may be combined with other references to disclose each element of the invention under 35 U.S.C. §103. I understand that a reference may also be combined with the knowledge of a POSITA, and that this knowledge may be used to combine multiple references. I further understand that a POSITA is presumed to know the relevant prior art. I understand that the obviousness analysis may take into account the inferences and creative steps that a POSITA would employ.

24. In determining whether a prior art reference would have been combined with other prior art or other information known to a POSITA, I understand that the following principles may be considered:

- a. whether the references to be combined involve non-analogous art;
- b. whether the references to be combined are in different fields of endeavor than the alleged invention in the Patent;
- c. whether the references to be combined are reasonably pertinent to the problems to which the inventions of the Patent are directed;
- d. whether the combination is of familiar elements according to known methods that yields predictable results;

- e. whether a combination involves the substitution of one known element for another that yields predictable results;
- f. whether the combination involves the use of a known technique to improve similar items or methods in the same way that yields predictable results;
- g. whether the combination involves the application of a known technique to a prior art reference that is ready for improvement, to yield predictable results;
- h. whether the combination is “obvious to try”;
- i. whether the combination involves the known work in one field of endeavor prompting variations of it for use in either the same field or a different one based on design incentives or other market forces, where the variations are predictable to a POSITA;
- j. whether there is some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill in the art to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention;
- k. whether the combination requires modifications that render the prior art unsatisfactory for its intended use;

- l. whether the combination requires modifications that change the principle of operation of the reference;
- m. whether the combination is reasonably expected to be a success; and
- n. whether the combination possesses the requisite degree of predictability at the time the invention was made.

25. I understand that in determining whether a combination of prior art references renders a claim obvious, it is helpful to consider whether there is some teaching, suggestion, or motivation to combine the references and a reasonable expectation of success in doing so. I understand, however, that a teaching, suggestion, or motivation to combine is not required.

V. LEVEL OF ORDINARY SKILL IN THE ART

26. I am familiar with the level of ordinary skill in the art with respect to the '987 Patent around its filing date. Based on my experience working, teaching, and conducting research in the relevant field, and based on my review of the '987 Patent specification, claims, file history, and prior art, I believe one of ordinary skill in the art around the time of the alleged invention of the '987 Patent would have had at least a bachelor's degree in electrical engineering, computer science, computer engineering, or a related field, and two-to-three years of experience in the research, design, development, or testing of media streaming technologies, user interfaces, and platforms, with additional education substituting for experience and vice versa.

27. In determining the level of ordinary skill in the art, I considered, for example, the type of problems encountered in the art, prior art solutions to those problems, the rapidity with which innovations are made, the sophistication of the technology, and the educational level of active workers in the field.

28. I met or exceeded the knowledge and experience that a POSITA would have had in 2010. I have extensive research and work experience through researching and analyzing streaming media for my doctoral thesis, working at a company to develop and ship the first business telephone system that worked over a company's data network, and working at a company that sold the industry-leading web site testing and monitoring solution.

VI. SUMMARY OF MY OPINIONS

29. As I explain below in detail in my Declaration, it is my opinion that Claims 1-13 of the '987 Patent would have been obvious under 35 U.S.C. §103 on the following grounds:

- **Ground 1:** Claims 1-13 are obvious over U.S. Patent Publication No. 2007/0055764 Dilley ("Dilley") (Ex-1005) alone.
- **Ground 2:** Claims 1-13 are obvious over Dilley in view of U.S. Patent Publication No. 2002/0152318 Menon ("Menon") (Ex-1007).
- **Ground 3:** Claims 1-13 are obvious over U.S. Patent Publication No. 2002/0069420 Russell ("Russell") (Ex-1006) alone.

- **Ground 4:** Claims 1-13 are obvious over Russell in view of Menon.

VII. TECHNOLOGICAL BACKGROUND

A. Cloud data persistence system for distribution of data

30. Cloud data persistence systems for distributing data were known by 2010. For example, U.S. Patent No. 8,180,720 to Kovacs (Ex-1022, “Kovacs”) teaches that one “distributed computer system is a ‘content delivery network’ or ‘CDN’ that is operated and managed by a service provider.... ‘[C]ontent delivery’ means the storage, caching, or transmission of content, streaming media and applications on behalf of content providers.” Ex-1022, 1:16-28. Similarly, U.S. Patent Publication No. 2011/0289126 to Aikas (Ex-1023, “Aikas”) teaches that “[a] service provider may provide a hosted, or cloud-based, storage service into which users (for example, stand-alone or web-based application developers) can store objects, and an associated Content Delivery Network (CDN) that provides caching of stored objects in one or more CDN nodes.” Ex-1023 ¶25.

31. *See also* U.S. Patent No. 7,428,540 to Coates (Ex-1024, “Coates”), 20:27-28 (“Content delivery networks attempt to optimize the delivery of commonly accessed rich media objects.”); U.S. Patent Publication No. 2011/0276623 to Girbal (Ex-1025, “Girbal”), Fig. 1; U.S. Patent Publication No. 2010/0250772 to Mao (Ex-1031, “Mao”), Fig. 2; U.S. Patent No. 6,996,616 to Leighton (Ex-1033, “Leighton”), 1:18-53.





FIG. 1

B. Origin servers for storing a plurality of data objects comprising a link to a plurality of media assets


32. Origin servers for storing a plurality of data objects comprising a link to a plurality of media assets were known by 2010. For example, U.S. Patent Publication No. 2003/0046239 to Geilfuss (Ex-1026, “Geilfuss”) teaches “content origination servers 12, 14, and 16 store[] and make[] available original versions of content or references to original content that is to be distributed through the Internet 17 and server 18 to users 32 of user systems 28. Content is considered a digital asset and includes, for example, images, video, audio, text documents, software, electronic books, Hyper-Text Markup Language (HTML), Extended Markup Language (XML), software and services.” Ex-1026 ¶31; *see also id.* at Fig. 1.

Figure 1



33. It was also known that each of the data objects would comprise a URL and a title linked to the media asset. For example, Geilfuss also teaches that “[t]he content management and distribution process 56 includes a link generation process 60. When content and/or references to content is received and stored in local storage device 20, the link generation process 60 generates a unique reference that is used to refer to the location of the content where it can be found on the local storage device 20, or to the pointer where content may be found internal or external to the content management and distribution network 10. The unique reference represents a unique descriptor. The unique descriptor may be a filename, pointer, Universal Resource Locator (URL), and so forth.” Ex-1026 ¶36; *see also id.* at Fig. 2.

Figure 2




34. *See also* Ex-1021 (Eaton), 8:26–36 (“Referring to FIG. 4, a file 130 is shown as including a metadata file object 132 and a plurality of data file objects. The metadata file object 132 contains information that points to each of the data file objects 134-136. Accessing the file includes first accessing the metadata file object 132 and then using information therein to locate the appropriate one or more of the corresponding data file objects 134-136.”); Ex-1031 (Mao) ¶51 (“VOD backoffice 115 may receive asset title and content metadata from CMS 112 through AMS 114”); Ex-1036 (Plastina) ¶44 (“The metadata transmitted with the media content items describes the media content items” and “includes, but is not limited to, album, artist, title, track, and album art” as well as “information about the distribution of the media content items: a date associated with transmission of the media content item

to the computing device, the parties involved in the transmission, identification of the media content items transferred with a particular media content item, identification of a local communication region where the transmission took place, and the like.”); *see also* Ex-1033 (Leighton), 1:54-2:6 (“Typically, content objects are made up of a basic structure that includes references to additional, embedded content objects. Most web pages, for example, consist of an HTML document that contains plain text together with some embedded objects, such as .gif or .jpg images. The embedded objects are referenced using embedded HTML directives, e.g., Uniform Resource Identifiers (URIs). A similar scheme is used for some types of streaming content that may be embedded within an SMIL document. Embedded HTML or SMIL directives tell the client to fetch embedded objects from the origin server. Using a CDN content modification scheme, a content provider can modify references to embedded objects so that the client is told to fetch an embedded object from the best surrogate (instead of from the origin server).”).

35. It was also a well-known concept for the origin server to be at a first location, for the plurality of media assets to be at a second location separate from the first location, and for the second location to comprise a content delivery network. This is taught, for example, in U.S. Patent Publication No. 2007/0299681 to Plastina (Ex-1036, “Plastina”). Plastina teaches that “[t]he metadata transmitted with the media content items describes the media content items” and “includes information


about the distribution of the media content items” such as “identification of a local communication region where the transmission took place, and the like.” Ex-1036 (Plastina) ¶44. Plastina further teaches that “[t]he metadata may be stored with the media content items or stored separately (e.g., in a different folder).” Ex-1036 (Plastina) ¶44; *see also, e.g.*, Ex-1026 (Geilfuss), Figs. 1-2. Moreover, the plurality of media assets are commonly stored at parent/cache and edge/application servers in addition to the origin server, wherein the parent/cache and edge/application servers are located at a separate physical location from the origin server and comprise a content delivery network. *See, e.g.*, Ex-1031 (Mao), Fig. 2.

36. Moreover, it was known that the origin servers in a cloud data persistence system would be communicatively coupled with a plurality of cache servers. For example, U.S. Patent Publication No. 2010/0057894 to Glasser (Ex-1027, “Glasser”) teaches that an “[e]dge cache server 118 and local cache server 120 can cache the content in order to respond to additional requests for the content ... [and] [i]n this way, the content may be retrieved from the origin server 116.” Ex-1027 ¶19, Fig. 1.



C. Cache servers that store a plurality of data objects and request a data object from the origin server if the data object is not present at the cache server

37. Cache servers that store a plurality of data objects and request a data object from the origin server if the data object is not present at the cache server were known by 2010. For example, Kovacs teaches that “[i]f the second tier does not have the content the request may then be forwarded to the origin server.” Ex-1022 (Kovacs), 8:64-65, Fig. 4.



38. It was also known that cache servers in a cloud data persistence system would be communicatively coupled with a plurality of application servers. For example, Glasser teaches that “the edge cache server can determine if the content is cached or if the content must be retrieved from the origin server. When the content is cached at the edge cache server, the edge cache server can provide the content to the client, as illustrated at 312. Alternatively, when the content is not cached at the edge cache server, the edge cache server can retrieve the content from the origin server, as illustrated at 314, prior to providing the content to the client, as illustrated at 312.” Ex-1027 (Glasser) ¶24. Similarly, U.S. Patent Publication No. 2010/0180208 to Kasten (Ex-1028, “Kasten”) teaches that the “application servers 110, cache servers 125, and update server 145 may all be at different locations and configured to communicate with each other. For example, each user device of the user devices 105 may communicate with any of the application servers 110, each application server of the application servers 110 may communicate with any of the cache servers 125, and each cache server may communicate with any of the databases 140.” Ex-1028 ¶16; *see also id.* at Fig. 1.




FIG. 1

D. Application servers that store a plurality of data objects and request the data object from the cache servers if the data object is not present at the application server

39. Application servers that store a plurality of data objects and request the data object from the cache servers if the data object is not present at the application server were known by 2010. For example, Kovacs teaches that “[a] multi-tier cache hierarchy approach is shown in FIG. 7 ... If the edge region does not yet have the content in cache, the request is forwarded to one of the cache hierarchy tiers ... If the content is not on the first tier parent, that parent forwards the request to the second tier.” Ex-1022 (Kovacs), 8:56-64, Fig. 7.




FIG. 7

40. Similarly, U.S. Patent No. 9,167,036 to Seed (Ex-1029, “Seed”) also teaches “[a]ccording to embodiments of the invention, a method and system for managed object replication and delivery over a network redirects, directly or indirectly, a client’s request for an object that is not available at a best or optimal handling edge server of the network to a parent server of the network that has the requested object.” Ex-1029, 2:63-3:1.

41. It was additionally known that the application servers were configured to receive requests from clients for the data objects. For example, Seed teaches that “[i]nitially, the method of managed object replication and delivery directs (at 200, 300) a client, requesting one or more objects, to an edge server in the network, whether or not the edge server has the requested object(s).” Ex-1029, 4:37-40; *see also* Ex-1033 (Leighton), 1:32-36 (“The content delivery infrastructure usually comprises a set of ‘surrogate’ origin servers that are located at strategic locations ... for delivering copies of content to requesting end users.”).

42. And it was known that the requests for the data objects would comprise a URL. For example, Coates teaches that a URL identifies an object file associated with the client’s web site. In response to the end user’s URL file request, the client site (e.g., content we server) generates a local file system request for the object file (block 2020, FIG. 20).” Ex-1024 (Coates), 22:24-27, Fig. 20; *see also* U.S. Patent No. 5,796,952 to Davis (Ex-1032, “Davis”), 7:10-13 (“Upon specification of a link,

the client makes a TCP/IP request to the server and receives information that was specified in that URL ... in return.”).




Figure 20

43. Moreover, it was known that the application servers would be geographically separated from one another. For example, as shown in Figure 1 below, Seed teaches that “[t]he network 100 comprises one or more parent server sites 120 and one or more edge server sites 130.” Ex-1029 (Seed), 3:36-37, Fig. 1; *see also* Ex-1031 (Mao) ¶50 (“multiple head ends 190 may be distributed geographically to server the various client devices 201, and may each contain the functional blocks as shown in FIG. 1”).




Fig. 1

E. Data objects are at least partially determined based on frequency of requests and/or when requests for those data objects were last received at the server

44. It was a well-known concept by 2010 that data objects are at least partially determined based on frequency of requests and/or when requests for those data objects were last received at the server. For example, Kovacs teaches that “it may be desirable to cache an object only after it is clear that the object is popular (that is, it will be requested by clients fairly frequently). The CDN server can be configured to cache a response only after it has received a few requests for the object.” Ex-1022 (Kovacs), 8:23-27.

45. See also Ex-1031 (Mao) ¶¶55 (“The actual locations at which each media content asset is stored may depend upon one or more factors, such as how popular the media content asset is to the end users, how popular the media content asset is expected to be, how often the media content assets is requested by one or more of the end users, and/or which end users have requested or are expected to

request the media content assets. The locations at which the media content assets are stored may change dynamically over time in responses to changes in these and/or other factors.”), 66 (“In addition, although a media content asset may be pre-positioned prior to any or substantial requests for that media content asset by client devices 201, the media content asset (or portion thereof) may further be replicated to one or more additional locations based on actual experienced requests by client devices 201 for that media content asset. And, once a media content asset has been pre-positioned or otherwise replicated to a location, the replicated copy of the media content asset may remain at that location for a predetermined period of time or until it is later determined that the popularity for that media content asset has dropped below a predetermined threshold, after which time the replicated copy may be deleted or moved to yet another location in the network.”), 67 (“Popularity of a media content asset may be determined in many ways, such as being based on a measured frequency of client device 201 requests for the media content asset, determining whether the media content asset has been requested by client devices 201 a sufficient number of times over a predetermined period of time, and/or based on historical or predicted future demand for the media content asset. Also, such determinations may be made on a global basis (i.e., across the entire network) and/or on a geographic regional basis, and may be made more than once over time to re-determine the popularity of the media content asset and re-replicate the entire media

content asset or a portion thereof as appropriate based on the newly-determined popularity.”).

F. Data objects at the origin server comprise all of the data objects in the cache and application servers

46. It was a well-known concept by 2010 that data objects at the origin server comprise all of the data objects in the cache and application servers. For example, Mao teaches that “[a]s a default, most content may be stored in the main content library. Then, depending upon content popularity and/or other factors, the content may be replicated and propagated ahead of time, or in response to a client request, to or near one or more of the local head end systems. Upon a subscriber’s request for content, the local system serving that subscriber may begin immediately streaming the content if the content is already cached at the local system. If the content is not cached at the local system, then the local system may pull the content from the content library or elsewhere. The pulled content may thereafter continue to be cached at the local system for a period of time to serve expected future requests from other subscribers served by that local system, and then later removed if desired.” Ex-1031 (Mao) ¶9. *See also, e.g.,* Ex-1029 (Seed), Fig. 1.




Fig. 1

VIII. THE '987 PATENT AND ITS FILE HISTORY

47. The '987 Patent relates to “distribution, storage, and retrieval of data on a cloud networking arrangement.” Ex-1001, 1:13-16. The '987 Patent states that in “a typical data hosting architecture, multiple servers may be spread over a geographical area” but storing the same data on all of the servers “may result in a substantial amount of data being stored at servers that wastes storage resources and is infrequently accessed.” Ex-1001, 1:17-37. The '987 Patent proposes a system and “method for retrieving data objects using a cloud data persistence system configured for distribution of data.” *Id.*, 2:35-37.

48. As shown in annotated Figure 5A below,¹ the '987 Patent's cloud data persistence system comprises a first origin server (red), which is communicatively coupled to a plurality of cache servers (i.e., a second server (green) and third server (yellow)). Ex-1001, Cl. 1, Fig. 5A. The cache servers are further communicatively

¹ Throughout, emphasis and annotations are added unless otherwise specified.

coupled to a plurality of application servers (i.e., a fourth server (purple), fifth server (orange), sixth server (blue), and seventh server (pink)), which are configured to receive a request for a “data object” from a client (grey). *Id.*, Cl. 1, 1:42-43.




FIG. 5A

49. The '987 Patent teaches that each server stores a plurality of data objects that are linked to media assets, “such as an audio file[,], a video file, or a flash file.” Ex-1001, 9:19-20, 9:30-37. The “media asset may be stored in a different location” from the data object such as on another server in the system. Ex-1001, 1:45-46, 9:38-45. The '987 Patent further teaches that if an application server does

not have the data object requested by the client stored locally, it may contact a cache server located in close proximity to determine if that server has it. Ex-1001, 10:43-51. If the cache server “determines that it does not have the data object requested” stored, it “may contact a data object origin server, which may have all existing data objects stored.” Ex-1001, 1:48-50, 11:48-52. The origin server “may then transmit the data object to” the cache server, which in turn, “may then transmit the data object” to the application server to transmit to the client. Ex-1001, 12:36-48.

50. The '987 Patent was filed on January 31, 2011. The '987 Patent claims priority to Australian Patent application No. 2010/202782, filed on July 1, 2010.

51. I understand from Counsel that the prior art and arguments in my Declaration were not before the PTO during prosecution. Ex-1004.

IX. CLAIM CONSTRUCTION

52. For the purposes of my Declaration, I identify the following claim construction issues. I do not currently believe that any other terms require express construction.

A. “Data object” (Claims 1, 3-13)

53. I understand from Counsel that in the Disney litigation, Petitioner has proposed that the recited “data object” in Claims 1 and 3-13 is “metadata about a particular media asset, distinct from the asset itself.” Ex-1034, 2. PO has proposed that the recited “data object” is “data structure that contains information about a

media asset.” Ex-1035, 2. As I explain below, in my opinion, there is no meaningful difference between these constructions as applied to the prior art references at issue in my Declaration.

B. “Data persistence cache servers” (Claim 7)

54. I understand from Counsel that in the Disney litigation, Petitioner has proposed that the recited “data persistence cache servers” in Claim 7 is invalid under 35 U.S.C. §112 ¶4. See Ex-1034, 2. PO has proposed that this limitation means that “the third server and second server are on the same level.” Ex-1035, 2. For the purposes of this proceeding only, I have interpreted this limitation to mean that the third and second servers are cache servers.


X. OVERVIEW OF THE APPLIED PRIOR ART REFERENCES


A. Dilley (Ex-1005)

55. Dilley (U.S. Patent Publication No. 2007/0055764), titled “Method and system for tiered distribution in a content delivery network” was filed on November 6, 2006 and published on March 8, 2007. I understand from Counsel that Dilley is §102(b) prior art to the ’987 Patent.

56. Dilley teaches a “tiered distribution service ... in a content delivery network (CDN).” Ex-1005, Abstract. As shown in annotated Figures 3-4 below, Dilley teaches that “a given content provider operates an origin server” (i.e., first server (red)). Ex-1005 ¶14. Dilley further teaches that “content may be marked for distribution from a so-called cache hierarchy that includes ... a given subset of the

origin server regions (a ‘multi-region’ cache hierarchy)’ called parent server regions (i.e., second server (yellow) and third server (green)). Ex-1005 ¶28. Dilley teaches that the cache hierarchy also comprises “given edge server region[s]” connected to the “parent server regions” (i.e., fourth server (purple), fifth server (orange), sixth server (blue), and seventh server (pink)). Ex-1005 ¶14, Figs. 3-4 (brown arrows and dashed boxes added).





57. Moreover, Dilley teaches that “a given edge server receives a client request for an object” (Fig. 3 (grey)) and “the edge server cache management process look[s] for the object in the content server’s cache.” Ex-1005 ¶44. If the requested object is not in the edge server’s cache, “the servers in the tiered distribution hubs” such as the given parent region, “will provide the edge servers with the requested content.” Ex-1005 ¶¶14, 27. The request is “only forwarded on to the origin server if it cannot be serviced by the intermediate node in the hierarchy.” Ex-1005 ¶14. Additionally, Dilley teaches that the CDN service provider (“CDNSP”) provides


“object-specific metadata” to the servers to determine how they “will handle a request for an object being served by the CDN.” Ex-1005 ¶25. The metadata includes, for example, “a modified URL ... that points an end user browser to attempt to retrieve the object from the CDN.” Ex-1005 ¶¶32-33.

B. Russell (Ex-1006)

58. Russell (U.S. Patent Publication No. 2002/0069420), titled “System and process for delivery of content over a network” was filed on April 6, 2001 and published on June 6, 2002. I understand from Counsel that Russell is §102(b) prior art to the ’987 Patent.

59. Russell teaches a system for delivering content over a network. Ex-1006, Abstract. As shown in annotated Figures 1-2 below, the system’s “network architecture includes three or more layers of servers, including a main server” (i.e., first server (red)) “supporting an interface for recipient processors [i.e., user network device (UND) (grey)] to request content items, a plurality of parent servers” (i.e., second server (green) and third server (yellow)), “and a plurality of edge servers” (i.e., fourth server (purple), fifth server (orange), sixth server (blue), and seventh server (pink)). Ex-1006, Abstract, Figs. 1-2 (brown dashed boxes added). Russell teaches that “[a] copy of each content item available through the service is distributed from the main server to each parent servers for storage,” which in turn, “distribute content to edges servers, pursuant to instructions from the main server.”

Ex-1006, Abstract. Russell further teaches that “[t]he content items (such as, movies) may be supplied” to main server 12 “by content owners or holders” from “source server 18.” Ex-1006 ¶32.






FIG. 2

60. Moreover, Russell teaches that “main server 12 publishes an XML document” to the parent and edge servers “containing a list of all content items ... available” on main server 12 with additional information such as, for example, whether a content item is marked “hot” and “geographic restrictions (so that a content item ... that is not available for download in a region will never be stored on edge servers in that region).” Ex-1006 ¶81. Russell further teaches that based on the instructions, “each edge server 16 may determine whether or not its associated storage facility is storing all of the appropriate content items and, if not, the edge server may request appropriate content items from a parent server.” Ex-1006 ¶82. Russell also teaches that “the main server provides the user’s UND with a URL (uniform resource locator) that will allow the user to connect to an edge server to

download an encrypted electronic file containing the selected content item.” Ex-1006 ¶12.

C. Menon (Ex-1007)

61. Menon (U.S. Patent Publication No. 2002/0152318), titled “Metadata enabled push-pull model for efficient low-latency video-content distribution over a network” was filed on March 4, 2002 and published on October 17, 2002. I understand from Counsel that Menon is §102(b) prior art.

62. Menon teaches a “metadata enabled push-pull model and method for efficient low-latency video-content distribution over a network.” Ex-1007, Abstract. Specifically, Menon teaches a “metadata enabled server” for distributing content objects to users over a network in response to user requests, which includes “a storage device coupled to [the] server and storing metadata describing content objects accessible to the server including a location from where a content object is stored.” Ex-1007 ¶35. For example, Menon teaches that “when a user makes a selection (such as by identifying a URL) for a video ... the metadata (MD) corresponding to the selected URL is examined to see if the corresponding video object is in the cache.” Ex-1007 ¶50. The metadata may comprise “(i) a globally unique identifier that identifies and optionally describes the content associated with the metadata, and (ii) a location or locations at which the content may be found, usually in the form of at least the origin server identifier.” Ex-1007 ¶63. The

metadata is distributed “to a plurality of locations (e.g to the origin server and edge servers, or to the origin server, edge servers, and directory servers[]).” Ex-1007 ¶84.

XI. DETAILED EXPLANATION OF THE UNPATENTABILITY GROUNDS

A. Grounds 1-2: Claims 1-13 are obvious over Dilley alone or Dilley in view of Menon.

1. A POSITA would have been motivated to combine Dilley’s teachings with Menon’s teachings and would have had a reasonable expectation of success.

63. In addition to the reasons set forth below, a POSITA would have been motivated to combine the teachings of Dilley and Menon and had a reasonable expectation of success in doing so because each relates to a content delivery network that uses metadata to enable efficient distribution of media content. Ex-1005, Abstract, ¶16; Ex-1007, Abstract, ¶42.

64. Dilley teaches a “tiered distribution service ... in a content delivery network (CDN)” comprising an origin server and a plurality of parent servers and edge servers. Ex-1005, Abstract. Dilley teaches that in addition to being stored at the origin server, the requested content may be stored at the edge servers and parent servers (Ex-1005 ¶¶27, 43), which are at separate geographical locations from the origin server that controls the metadata. Ex-1005 ¶¶15, 28, Fig. 3.

65. Menon teaches a “metadata enabled push-pull model and method for efficient low-latency video-content distribution over a network.” Ex-1007, Abstract. Menon further teaches a directory server, which is “a different machine than the

origin server or the edge server” where the content items are stored (Ex-1007 ¶¶66-67), and “provides a single location data structure that may be queried to identify all locations at which a content item is located.” Ex-1007 ¶73.

66. To the extent that a POSITA was interested in additional ways to implement Dilley’s content delivery network such that storage and distribution of the metadata and content were further separated, a POSITA would have looked to Menon, which teaches storing the metadata at a separate directory server. A POSITA would have been motivated to store the metadata at a separate location than the server that stores the content because, as Dilley itself acknowledges, it would have been “desirable to reduce wide-area network bandwidth and the load on a content provider’s origin server as much as possible.” Ex-1005 ¶6. Furthermore, in computer files systems, the data about the data is stored separately from the data itself. Indeed, Menon teaches that its directory server advantageously “provide[s] redundancy and reduce[s] loading on any particular server during peak loads.” Ex-1007 ¶67. In the Dilley-Menon combination, Dilley’s system would store the metadata and the content at separate locations such as in a directory server and origin server, as taught by Menon.

67. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), it was a well-known concept for the origin server to be at a first location and for the plurality of media assets to be at a second location

separate from the first location. Ex-1026 (Geilfuss), Figs. 1-2; Ex-1031 (Mao), Fig. 2.

68. A POSITA would have had a reasonable expectation of success in combining the teachings of Dilley and Menon. Both describe using metadata to enable efficient distribution of media content within a tiered server network infrastructure. Ex-1005, Abstract, ¶16, Figs. 3-4; Ex-1007, Abstract, ¶42, Fig. 3. A POSITA would thus have understood that these references disclose interrelated teachings based on well-understood technologies that would have been amenable to various well-understood and predictable combinations.

2. Independent Claim 1

a. Element 1[pre]: A cloud data persistence system for distribution of data, the cloud data persistence system comprising:


69. In my opinion, Dilley alone or Dilley in view of Menon (“Dilley-Menon”) teaches² the preamble, to the extent it is limiting.

70. Dilley teaches a CDN that provides a tiered distribution service using “a set of surrogate origin (namely, ‘edge’) servers organized into regions,” which deliver content “on behalf of participating content providers,” each of which “operates an origin server.” Ex-1005, Abstract. Dilley further teaches that “a cache

² I understand from Counsel that the term “teaches” includes both express teachings and those fairly suggested to a person of ordinary skill in the art.

hierarchy is established,” which includes “a given edge server region and ... a subset of the edge server regions.” Ex-1005, Abstract. If it is determined that a user’s “request cannot be serviced in the given edge region, instead of contacting the origin server, the request is provided to ... a given one of the subset of edge server regions for handling, preferably as a function of metadata associated with the” request. Ex-1005, Abstract. Only “if the request cannot be serviced by an intermediate node” is it forwarded “to the origin server.” Ex-1005, Abstract.



71. Dilly’s cloud data persistence system is depicted in Figure 3, annotated below. The coloration will be discussed in further details in the subsequent element-by-element analysis. As also I explained in the Technological Background section of my Declaration (*see* Section VII.A), cloud data persistence systems for distributing data were known by 2010. Ex-1022 (Kovacs), 1:16-28; Ex-1023 (Aikas) ¶25; Ex-1024 (Coates), 20:27-28; Ex-1025 (Girbal), Fig. 1; Ex-1031 (Mao), Fig. 2; Ex-1033 (Leighton), 1:18-53.



b. Element 1[a][1]: a first server which stores a first plurality of data objects and is communicatively coupled with a plurality of cache servers, wherein:

72. In my opinion, Dilley or Dilley-Menon teaches Element 1[a][1].

73. As shown in annotated Figures 3-4 below, Dilley teaches a first server (i.e., origin server (red)), which is communicatively coupled with a plurality of cache servers called parent region servers (green, yellow). Ex-1005, Figs. 3-4 (brown arrow added).



Specifically, Dilley teaches that “a given content provider operates *an origin server*.” Ex-1005 ¶14. Dilley further teaches that “*a cache hierarchy* is established in the CDN comprising a given edge server region and ... *a subset of well-connected ‘core’ parent server regions*.” Ex-1005 ¶14.

74. Dilley also teaches that the origin server (i.e., first server) “stores a first plurality of data objects.” The ’987 Patent does not explicitly define the term “data object” but states that a “data object may contain a link to a media asset and other information linked to the video file, such as a title and duration.” Ex-1001, 1:43-45. A POSITA would thus have understood that a “data object” includes, at least, a metadata (such as title, keywords, creation date, duration, file formats, and permissions, etc.) about a particular media asset, as distinct from the media asset itself.

75. Dilley teaches that “[p]referably, only a relatively small, specific subset of objects are serviced using the hierarchy, while most objects are fetched using a default mechanism where the edge server contacts the origin server to handle a request that cannot be serviced in the edge region,” with a selectivity process that is accomplished “using object ‘metadata.’” Ex-1005 ¶29. Specifically, the CDNSP provides “*object-specific metadata* to the CDN content servers to determine how” they “will handle a request for an object being served by the CDN.” Ex-1005 ¶25. Dilley defines metadata as “the set of all control options and parameters for the


object (e.g., coherence information, origin server identity information, load balancing information, customer code, other control codes, etc.),” which may be provided to the servers “via a configuration file, in HTTP headers, or in other ways.” Ex-1005 ¶25. In one example, metadata may be provided to the origin server by the “content provider (e.g., using a secure extranet application) and periodically delivered to the” other servers, as necessary. Ex-1005 ¶25. Dilley also teaches that “a given metadata tag” is preferably used to “enable the tiered distribution service or feature and to specify a policy or type of hierarchy parent to be used.” Ex-1005 ¶16; *see also id.*, ¶45 (providing an example using a “metadata tag”). Moreover, content may be “tagged for delivery from the CDN using a content migrator or rewrite tool 106 operated ... at a participating content provider server,” which “rewrites embedded object URLs to point to the CDNSP domain.” Ex-1005 ¶25. The object-specific metadata stored at the origin server (and other servers), for example, via a configuration file and/or tag corresponds to the recited “plurality of data objects” because it links the client to the requested content. Dilley or Dilley-Menon thus teaches the “data objects” limitation under either parties’ construction because it teaches “metadata about a particular media asset, distinct from the asset itself” and more broadly teaches a “data structure that contains information about a media asset.”

76. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), origin servers for storing a plurality of data objects comprising a link to a plurality of media assets were known by 2010. Ex-1026 (Geilfuss) ¶31, Fig. 1. Additionally, as I explained in the Technological Background section of my Declaration (*see* Section VII.B), it was known that the origin servers in a cloud data persistence system would be communicatively coupled with a plurality of cache servers. Ex-1027 (Glasser) ¶19, Fig. 1.

c. Element 1[a][2]: the first server is located at a first location;

77. In my opinion, Dilley or Dilley-Menon teaches Element 1[a][2].

78. Dilley teaches that “[a] CDN is a network of geographically-distributed content delivery nodes that are arranged for efficient delivery of content on behalf of third party content providers.” Ex-1005 ¶5. Figure 3 illustrates that the first server (i.e., origin server (red)) is located at a first location relative to the other servers in the CDN. Ex-1005, Fig. 3 (brown arrow added).



- d. **Element 1[a][3]: each of the first plurality of data objects comprise a link to one of a plurality of media assets; and**

79. In my opinion, Dilley or Dilley-Menon teaches Element 1[a][3].

80. Dilley teaches that each of the first plurality of data objects (i.e., metadata) comprise a link to one of a plurality of media assets (i.e., content) for the reasons described in Section XI.A.2.b. Ex-1005 ¶¶25 (URLs are embedded to point to the content), 29, 45. Moreover, Dilley teaches that “representative” CDN metadata may include the “metadata tag ‘use–hierarchy,’” which “turns on the cache hierarchy feature, specifies the policy or type of hierarchy parent, and lists the base domain to be used.” Ex-1005 ¶31. In one example, the “default setting for cache hierarchy is OA: <md name=“use–hierarchy”>OA</md>,” where “OA indicates that the edge server is to go direct to the *(O)origin-server* specified *in the (A)RL*),

where the ARL is *a modified URL (or portion thereof) that points an end user browser to attempt to retrieve the object from the CDN.*” Ex-1005 ¶¶32-33; *see also id.*, ¶6.


81. As also I explained in the Technological Background section of my Declaration (*see* Section VII.B), it was known that each of the data objects would comprise a URL linked to the media asset. Ex-1026 (Geilfuss) ¶36, Fig. 2; Ex-1021 (Eaton), 8:26–36; Ex-1031 (Mao) ¶51; Ex-1033 (Leighton), 1:54-2:6.

e. Element 1[a][4]: the plurality of media assets are stored at a second location separate from the first location;

82. In my opinion, Dilley or Dilley-Menon teaches Element 1[a][4].

83. Ground 1: Dilley teaches that the plurality of media assets may be stored at the edge servers and parent region servers. Specifically, Dilley teaches that each edge server “includes a cache for storing content provider content.” Ex-1005 ¶43. Moreover, “[t]he tiered distribution hubs ensure that the edge servers have the information requested by the end users” and, if they “do not have the appropriate file when it is requested, the servers in the tiered distribution hubs will provide the edge servers with the requested content.” Ex-1005 ¶27. The edge servers and parent region servers are at separate geographical locations from the origin server (i.e., first location) as shown in Figure 3. Ex-1005 ¶¶15 (“edge server regions typically are co-located in large hosting data centers in the well-connected ‘core’ of the Internet web

hosting infrastructure”), 28 (“a cluster of ‘parent’ servers ... may be co-located in large hosting data centers in well-connected ‘core’ portions of the Internet web hosting infrastructure”), Fig. 3. To the extent PO argues that the first location where the metadata is stored may not also store the plurality of media assets, this would have been obvious to a POSITA, and is taught by Menon.



84. Ground 2: Menon teaches that the plurality of media assets are stored at a second location separate from the first location. Specifically, Menon teaches a directory server, which is “a different machine than the origin server or the edge server.” Ex-1007 ¶¶66-67. Menon teaches that when the system “utilizes one or more servers other than the origin server and the edge server, such as one or more directory servers, the content metadata may also be pushed to the directory server” to “provide[] a single location data structure that may be queried to identify all


locations at which a content item is located.” Ex-1007 ¶73. The directory server essentially acts as a metadata server in that it finds “the optimal source location for [a] requested video” and “receives metadata updates from the system.” Ex-1007 ¶91. Thus, a POSITA would have understood that Menon teaches that metadata may be stored at a first location (i.e., directory server) and the content may be stored at a separate, second location (i.e., origin server). Menon teaches that the metadata is stored in a metadata database, and a POSITA would have understood that a database would be separate from a content server such as the origin server. In the Dilley-Menon combination, Dilley’s system would store the metadata and the content at separate locations such as in a directory server and origin server, as taught by Menon. A POSITA would have been motivated to combine Dilley-Menon and had a reasonable expectation of success in doing so for the reasons described in Section XI.A.1.


85. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), it was a well-known concept for the origin server to be at a first location and for the plurality of media assets to be at a second location separate from the first location. Ex-1026 (Geilfuss), Figs. 1-2; Ex-1031 (Mao), Fig. 2.

- f. **Element 1[b][1]: a second server that stores a second plurality of data objects and is configured to request a data object from the first server if the data object is not present in the second plurality of data objects, wherein:**

86. In my opinion, Dilley or Dilley-Menon teaches Element 1[b][1].

87. As shown in annotated Figures 3-4 below, Dilley teaches a second server (green). Ex-1005, Figs. 3-4. Dilley teaches that “a cache hierarchy is established in the CDN” comprising a “subset of well-connected ‘core’ *parent server* regions.” Ex-1005 ¶¶14, 27 (“a hierarchical set of ‘parent’ or hub regions”).





88. In my opinion, Dilley teaches that the second server stores a second plurality of data objects and is configured to request a data object from the first server if the data object is not present in the second plurality of data objects. Specifically, Dilley teaches that “in response to a determination that a given object request cannot be serviced in the given edge region” it is provided to “a given one of the subset of server regions for handling, preferably *as a function of metadata* associated with the given object request.” Ex-1005 ¶14. “The given object request is then serviced, if possible, by a given CDN server in ... *the given parent region*” and “is only forwarded on to the origin server if it cannot be serviced by the intermediate node in


the hierarchy.” Ex-1005 ¶¶14, 51 (“parent regions get the content (if needed) from the origin server(s)”).

89. As I also explained in the Technological Background section of my Declaration (*see* Section VII.C), cache servers that store a plurality of data objects and request a data object from the origin server if the data object is not present at the cache server were known by 2010. Ex-1022 (Kovacs), 8:64-65, Fig. 4.

g. Element 1[b][2]: the second server is one of the plurality of cache servers; and

90. In my opinion, Dilley or Dilley-Menon teaches Element 1[b][2].



91. Dilley teaches that the second server is one of the plurality of cache servers. Specifically, Dilley teaches that “content may be marked for distribution from a so-called cache hierarchy that includes ... a given subset of the origin server regions (*a ‘multi-region’ cache hierarchy*)” such as “the given *subset of parent regions*” 408 a-n in annotated Fig. 4 below. Ex-1005 ¶28, Fig. 4. This region “includes a cluster of ‘parent’ servers,” which “may be co-located in large hosting data centers in well-connected ‘core’ portions of the Internet web hosting infrastructure.” Ex-1005 ¶28.



h. Element 1[b][3]: the second server is communicatively coupled with a first plurality of application servers;

92. In my opinion, Dilley or Dilley-Menon teaches Element 1[b][3].

93. Dilley teaches that “a cache hierarchy is established in the CDN comprising a given *edge server region* and ... a subset of *well-connected ‘core’ parent server regions*.” Ex-1005 ¶14. As shown in annotated Figures 3-4 below, Dilley’s second server (green) is communicatively coupled to a plurality of edge servers such as a fourth server (purple) and fifth server (orange). Ex-1005, Figs. 3-4 (brown arrows added).



94. Dilley's edge servers are application servers. Dilley teaches that edge servers are typically content servers "located at or near the so-called outer reach or



‘edge’ of the Internet.” Ex-1005 ¶24. Moreover, Dilley teaches that a representative edge server “is a Pentium-based caching appliance running an operating system (e.g., Linux, Windows NT, Windows 2000) and having suitable RAM and disk storage for CDN applications and content delivery network content (e.g., HTTP content, streaming media and applications).” Ex-1005 ¶24. The edge servers include “an application programming interface (API) for the guide process to facilitate location of a parent region (from which the object can be fetched).” Ex-1005 ¶48.

95. As I also explained in the Technological Background section of my Declaration (*see* Section VII.C), it was known that cache servers in a cloud data persistence system would be communicatively coupled with a plurality of application servers. Ex-1027 (Glasser) ¶24; Ex-1028 (Kasten) ¶16, Fig. 1.

- i. Element 1[c][1]: a third server that stores a third plurality of data objects and is configured to request the data object from the first server if the data object is not present in the third plurality of data objects, wherein:**

96. In my opinion, Dilley or Dilley-Menon teaches Element 1[c][1].

97. As shown in annotated Figures 3-4 below, Dilley teaches a third server (yellow), which is a server in the parent region. Ex-1005, Figs. 3-4 (brown arrow added).



98. Dilley teaches that the third server stores a third plurality of data objects and is configured to request a data object from the first server if the data object is not present in the third plurality of data objects for the reasons described in Section XI.A.2.f. Specifically, Dilley teaches that a given object request is handled as a function of metadata and “is then serviced, if possible, by a given CDN server in ... *the given parent region*” and “only forwarded on to the origin server if it cannot be serviced by the intermediate node in the hierarchy.” Ex-1005 ¶¶14, 51.

99. As I explained in the Technological Background section of my Declaration (*see* Section VII.C), cache servers that store a plurality of data objects and request a data object from the origin server if the data object is not present at the cache server were known by 2010. Ex-1022 (Kovacs), 8:64-65, Fig. 4.

j. Element 1[c][2]: the third server is one of the plurality of cache servers; and

100. In my opinion, Dilley or Dilley-Menon teaches Element 1[c][2].

101. Dilley teaches that the third server is one of the plurality of cache servers for the reasons described in Section XI.A.2.g. Ex-1005 ¶28, Fig. 4.


k. Element 1[c][3]: the third server is communicatively coupled with a second plurality of application servers; and


102. In my opinion, Dilley or Dilley-Menon teaches Element 1[c][3].

103. As shown in annotated Figures 3-4 below, Dilley’s third server (yellow) is communicatively coupled to a plurality of edge servers such as a sixth server

(blue) and seventh server (pink). Ex-1005, Figs. 3-4 (brown arrows and dashed box added), ¶14. As explained in Section XI.A.2.h, the edge servers are application servers. Ex-1005 ¶¶24, 53.

104. As I also explained in the Technological Background section of my Declaration (*see* Section VII.C), it was known that cache servers in a cloud data persistence system would be communicatively coupled with a plurality of application servers. Ex-1027 (Glasser) ¶24; Ex-1028 (Kasten) ¶16, Fig. 1.







- I. **Element 1[d][1]: a fourth server that stores a fourth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fourth plurality of data objects, wherein:**


105. In my opinion, Dilley or Dilley-Menon teaches Element 1[d][1].

106. As shown in annotated Figures 3-4 below, Dilley teaches a fourth server (purple). Ex-1005, Figs. 3-4. Dilley teaches that the CDN “employs a tiered infrastructure using two or more tiers: *a set of edge servers* and a hierarchical set of ‘parent’ or hub regions that service the edge servers.” Ex-1005 ¶27.



107. Dilley teaches that the fourth server stores a fourth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fourth plurality of data objects. Specifically, as shown in Figure 5 below, Dilley teaches that each edge server runs software “implemented as part of the CDN software” and “includes a cache for storing content provider content,” a guide process component “that directs the edge server to use another CDN server (i.e., an intermediate) instead of the customer’s origin server for given content requests,” a “C map,” which maps “from a client network address to a nearest CDN region ... to direct client requests,” and “a more global G map ... used by the CDN service provider as part of the request routing mechanism.” Ex-1005 ¶43, Fig. 5. Moreover, Dilley teaches that “[i]f the edge servers do not have *the appropriate file when it is requested*, the servers in *the tiered distribution hubs will provide the edge servers with the requested content*.” Ex-1005 ¶27. For example, “in response to a determination that a given object request cannot be serviced in the given edge region” it is provided to “a given one of the subset of server regions for handling, preferably *as a function of metadata* associated with the given object request.” Ex-1005 ¶14. “The given object request is then serviced, if possible, by a given CDN server in ... *the given parent region*” (i.e., second server). Ex-1005 ¶14, Fig. 4 (showing second server connected to the fourth server).

108. As I explained in the Technological Background section of my Declaration (*see* Section VII.D), application servers that store a plurality of data objects and request the data object from the cache servers if the data object is not present at the application server were known by 2010. Ex-1022 (Kovacs), 8:56-64, Fig. 7; Ex-1029 (Seed), 2:63-3:1.



m. Element 1[d][2]: the fourth server is one of the first plurality of application server[s]; and

109. In my opinion, Dilley or Dilley-Menon teaches Element 1[d][2].

110. Dilley teaches that the fourth server is one of the first plurality of application servers for the reasons described in the previous section and Section XI.A.2.h, namely, because the second server is communicatively coupled to the fourth server and Dilley teaches that its edge servers are application servers.


n. Element 1[d][3]: the fourth server is configured to receive a request from a first client for the data object;

111. In my opinion, Dilley or Dilley-Menon teaches Element 1[d][3].

112. Dilley teaches that the fourth server is configured to receive a request from a first client for the data object. As shown in annotated Figure 3 below, the

fourth server (purple) directly communicates with the client (grey) to receive requests for data objects and provide the corresponding content. Ex-1005, Fig. 3.

113. As I explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the application servers were configured to receive requests from clients for the data objects. Ex-1029 (Seed), 4:37-40; Ex-1033 (Leighton), 1:32-36.





Dilley further teaches that “a given edge server receives a client request for an object” and “the edge server cache management process look[s] for the object in the content server’s cache.” Ex-1005 ¶44. “If the requested object is in cache ... the object is served back to the requesting client” and if it is not “the edge server makes a request (e.g., via a cache sharing protocol) to other edge servers in the region” so that it can “return the object to the requesting client.” Ex-1005 ¶44.

- o. **Element 1[e][1]: a fifth server that stores a fifth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fifth plurality of data objects, wherein:**

114. In my opinion, Dilley or Dilley-Menon teaches Element 1[e][1].

115. As shown in annotated Figures 3-4 below, Dilley teaches a fifth server (orange). Ex-1005, Figs. 3-4 (brown arrows added).





116. Dilley teaches that the fifth server stores a fifth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fifth plurality of data objects for the reasons described in Section XI.A.2.1. Specifically, Dilley teaches that “[i]f the edge servers do not have the appropriate file when it is requested, the servers in *the tiered distribution hubs will provide the edge servers with the requested content.*” Ex-1005 ¶27; *see also id.*, ¶¶14 (edge servers send requests to parent regions), 43, Figs. 4 (showing second server connected to the fifth server), 5 (showing edge server components).

117. As I explained in the Technological Background section of my Declaration (*see* Section VII.D), application servers that store a plurality of data objects and request the data object from the cache servers if the data object is not present at the application server were known by 2010. Ex-1022 (Kovacs), 8:56-64, Fig. 7; Ex-1029 (Seed), 2:63-3:1.

p. Element 1[e][2]: the fifth server is one of the first plurality of application servers; and

118. In my opinion, Dilley or Dilley-Menon teaches Element 1[e][2].


119. Dilley teaches that the fifth server is one of the first plurality of application servers for the reasons described in the previous section and Section XI.A.2.h, namely, because the second server is communicatively coupled to the fifth server and Dilley teaches that its edge servers are application servers.

q. Element 1[e][3]: the fifth server is configured to receive a request from a second client for the data object;

120. In my opinion, Dilley or Dilley-Menon teaches Element 1[e][3].

121. Dilley teaches that the fifth server is configured to receive a request from a second client for the data object. As shown in annotated Figure 3 below, the fifth server (orange) directly communicates with the client (grey) to receive requests for data objects and provide the corresponding content. Ex-1005, Fig. 3 (brown arrows added). *See also* Section XI.A.2.n; Ex-1005 ¶44.



122. As I explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the application servers were configured to receive requests from clients for the data objects. Ex-1029 (Seed), 4:37-40; Ex-1033 (Leighton), 1:32-36.



- r. **Element 1[f][1]: a sixth server that stores a sixth plurality of data objects and is configured to request the data object from the third server if the data object is not present in the sixth plurality of data objects, wherein:**

123. In my opinion, Dilley or Dilley-Menon teaches Element 1[f][1].

124. As shown in annotated Figures 3-4 below, Dilley teaches a sixth server (blue). Ex-1005, Figs. 3-4 (brown arrows added).



125. Dilley teaches that the sixth server stores a sixth plurality of data objects and is configured to request the data object from the third server if the data object is not present in the sixth plurality of data objects for the reasons described in Section XI.A.2.l. Specifically, Dilley teaches that “[i]f the edge servers do not have the appropriate file when it is requested, the servers in *the tiered distribution hubs will provide the edge servers with the requested content.*” Ex-1005 ¶27; *see also id.*, ¶¶14 (edge servers send requests to parent regions), 43, Figs. 4 (showing third server connected to the sixth server), 5 (showing edge server components).

126. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), application servers that store a plurality of data objects and request the data object from the cache servers if the data object is not present at the application server were known by 2010. Ex-1022 (Kovacs), 8:56-64, Fig. 7; Ex-1029 (Seed), 2:63-3:1.

s. Element 1[f][2]: the sixth server is one of the second plurality of application servers; and

127. In my opinion, Dilley or Dilley-Menon teaches Element 1[f][2].


128. Dilley teaches that the sixth server is one of the second plurality of application servers for the reasons described in the previous section and Sections XI.A.2.h and XI.A.2.k, namely, because the third server is communicatively coupled to the sixth server and Dilley teaches that its edge servers are application servers.

- t. Element 1[f][3]: the sixth server is configured to receive a request from a third client for the data object; and**

129. In my opinion, Dilley or Dilley-Menon teaches Element 1[f][3].

130. Dilley teaches that the sixth server is configured to receive a request from a third client for the data object. As shown in annotated Figure 3 below, the sixth server (blue) directly communicates with the client (grey) to receive requests for data objects and provide the corresponding content. Ex-1005, Fig. 3 (brown arrows added). *See also* Section XI.A.2.n; Ex-1005 ¶44.



131. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the requests for the data objects would comprise a URL. Ex-1024 (Coates), 22:24-27, Fig. 20; Ex-1032 (Davis), 7:10-13.



- u. **Element 1[g][1]: a seventh server that stores a seventh plurality of data objects and is configured to request the data object from the third server if the data object is not present in the seventh plurality of data objects, wherein:**

132. In my opinion, Dilley or Dilley-Menon teaches Element 1[g][1].

133. As shown in annotated Figures 3-4 below, Dilley teaches a seventh server (pink). Ex-1005, Figs. 3-4 (brown arrows and dashed box added).



134. Dilley teaches that the seventh server stores a seventh plurality of data objects and is configured to request the data object from the third server if the data object is not present in the seventh plurality of data objects for the reasons described in Section XI.A.2.l. Specifically, Dilley teaches that “[i]f the edge servers do not have the appropriate file when it is requested, the servers in *the tiered distribution hubs will provide the edge servers with the requested content.*” Ex-1005 ¶27; *see also id.*, ¶¶14 (edge servers send requests to parent regions), 43, Figs. 4 (showing third server connected to the seventh server), 5 (showing edge server components).

135. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the application servers were configured to receive requests from clients for the data objects. Ex-1029 (Seed), 4:37-40; Ex-1033 (Leighton), 1:32-36.

v. Element 1[g][2]: the seventh server is one of the second plurality of application servers; and

136. In my opinion, Dilley or Dilley-Menon teaches Element 1[g][2].


137. Dilley teaches that the seventh server is one of the second plurality of application servers for the reasons described in the previous section and Sections XI.A.2.h and XI.A.2.k, namely, because the third server is communicatively coupled to the seventh server and Dilley teaches that its edge servers are application servers.

- w. **Element 1[g][3]: the seventh server is configured to receive a request from a fourth client for the data object.**

138. In my opinion, Dilley or Dilley-Menon teaches Element 1[g][3].

139. Dilley teaches that the seventh server is configured to receive a request from a fourth client for the data object. As shown in annotated Figure 3 below, the seventh server (pink) directly communicates with the client (grey) to receive requests for data objects and provide the corresponding content. Ex-1005, Fig. 3 (brown arrows added). *See also* Section XI.A.2.n; Ex-1005 ¶44.

140. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the requests for the data objects would comprise a URL. Ex-1024 (Coates), 22:24-27, Fig. 20; Ex-1032 (Davis), 7:10-13.



141. Thus, Dilley alone or Dilley-Menon renders obvious Claim 1.



3. Independent Claim 5

- a. **Element 5[pre]: A cloud data persistence system for distribution of data, the cloud data persistence system comprising:**

142. In my opinion, Dilley or Dilley-Menon teaches the preamble. *See* Section XI.A.2.a.

- b. **Element 5[a][1]: a first server which stores a first plurality of data objects and is communicatively coupled with a second server and a third server, wherein:**

143. In my opinion, Dilley or Dilley-Menon teaches Element 5[a][1]. *See* Sections XI.A.2.b and XI.A.2.f-XI.A.2.j; Ex-1005, Figs. 3-4 (brown arrow added).



- c. **Element 5[a][2]: the first server is located at a first location;**

144. In my opinion, Dilley or Dilley-Menon teaches Element 5[a][2]. *See* Section XI.A.2.c.

- d. **Element 5[a][3]: each of the first plurality of data objects comprises a link to a media asset of a plurality of media assets;**

145. In my opinion, Dilley or Dilley-Menon teaches Element 5[a][3]. *See* Section XI.A.2.d.

- e. **Element 5[a][4]: the plurality of media assets is located at a second location ; and the second location is distinct from the first location;**



146. In my opinion, Dilley or Dilley-Menon teaches Element 5[a][4]. *See* Section XI.A.2.e.

- f. **Element 5[b][1]: the second server stores a second plurality of data objects and is configured to request a data object from the first server if the data object is not present in the second plurality of data objects,**

147. In my opinion, Dilley or Dilley-Menon teaches Element 5[b][1]. *See* Sections XI.A.2.f.

- g. **Element 5[b][2]: wherein the second server is communicatively coupled with a fourth server and fifth server; and**

148. In my opinion, Dilley or Dilley-Menon teaches Element 5[b][2]. *See* Sections XI.A.2.h and XI.A.2.l-XI.A.2.q; Ex-1005, Figs. 3-4 (brown arrows added).



- h. Element 5[c][1]: the fourth server stores a third plurality of data objects and is configured to request the data object from the second server if the data object is not present in the third plurality of data objects,**

149. In my opinion, Dilley or Dilley-Menon teaches Element 5[c][1]. *See* Section XI.A.2.1.

- i. Element 5[c][2] wherein the fourth server is configured to receive a request from a first client for the data object.**

150. In my opinion, Dilley or Dilley-Menon teaches Element 5[c][2]. *See* Section XI.A.2.n.

151. In my opinion, Thus, Dilley alone or Dilley-Menon renders obvious Claim 5.

4. Dependent Claims 2-4 and 6-13

- a. Claim 2: The cloud data persistence system of claim 1, wherein the second location comprises a content delivery network.**

152. In my opinion, Dilley or Dilley-Menon teaches Claim 2.

153. Ground 1: Dilley teaches Claim 2 for the reasons described in Section XI.A.2.e, namely, because content is stored at parent region servers and edge servers within a content delivery network, which are at separate geographical locations from the origin server. Ex-1005 ¶¶15, 27 (“preferably the CDN ... employs a tiered infrastructure using two or more tiers: a set of edge servers and a hierarchical set of ‘parent’ or hub regions that service the edge servers”), 28, 43 (each edge server

“includes a cache for storing content provider content”), Fig. 3. Moreover, Dilley teaches that “the original request is only forwarded on to the origin server if it cannot be serviced by the intermediate node in the hierarchy.” Ex-1005 ¶14.

154. Ground 2: Dilley-Menon also teaches Claim 2 for the reasons described in Section XI.A.2.e, namely, because the metadata is stored at a directory server separate from the origin and edge servers that store the content in the content delivery network.

155. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), it was a well-known concept for the origin server to be at a first location, for the plurality of media assets to be at a second location separate from the first location, and for the second location to comprise a content delivery network. Ex-1026 (Geilfuss), Figs. 1-2; Ex-1031 (Mao), Fig. 2.

b. Claim 3

(1) Element 3[pre]: The cloud data persistence system of claim 1, wherein:

156. In my opinion, Dilley or Dilley-Menon teaches Element 3[pre]. *See* Section XI.A.2.a.

(2) Element 3[a]: data objects in the second plurality of data objects are at least partially determined based on when requests for those data objects were last received at the second server;

157. In my opinion, Dilley or Dilley-Menon teaches Element 3[a].

158. In my opinion, Dilley teaches that one of the factors that the CDN uses in determining which server to service a request from is whether the content is “stale” or “infrequently accessed.” Ex-1005 ¶27. For example, Dilley teaches that “given content is associated with the cache hierarchy functionality using metadata” and a parent server “typically sends a child (an edge server running the CDN software) with information about how long (time-to-live or TTL) the content may stay in the child cache and still be deemed fresh.” Ex-1005 ¶51. Parameters such as TTL are determined by the server as to how long cached content in edge servers are deemed to be fresh, or current, and represents data about the content (i.e., metadata) after the request that caused it to be stored in the cache. For example, breaking news content may be considered fresh for only a short period of time, since the news may update frequently while breaking, and a newly released movie trailer may be considered fresh for a long time after it is released, since the trailer content is not expected to change. Dilley further teaches that the metadata includes Age information on the object indicating how long “the object has been in the CDN (replacing the Age header received from an upstream proxy if necessary),” and a Date header, which is “updated to the current time” of the request. Ex-1005 ¶50. Thus, Dilley’s data objects in the second plurality of data objects (i.e., metadata) are at least partially determined based on when requests for those data objects were last received at the second server.

159. As I also explained in the Technological Background section of my Declaration (*see* Section VII.E), it was a well-known concept by 2010 that data objects are at least partially determined based on frequency of requests and/or when requests for those data objects were last received at the server. Ex-1022 (Kovacs), 8:23-27; Ex-1031 (Mao) ¶¶55, 66-67.

- (3) Element 3[b]: data objects in the third plurality of data objects are at least partially determined based on when requests for those data objects were last received at the third server;**

160. In my opinion, Dilley or Dilley-Menon teaches Element 3[b]. *See* Section XI.A.4.b.2.

- (4) Element 3[c]: data objects in the fourth plurality of data objects are at least partially determined based on when requests for those data objects were last received at the fourth server;**

161. In my opinion, Dilley or Dilley-Menon teaches Element 3[c]. *See* Section XI.A.4.b.2.

- (5) Element 3[d]: data objects in the fifth plurality of data objects are at least partially determined based on when requests for those data objects were last received at the fifth server;**

162. In my opinion, Dilley or Dilley-Menon teaches Element 3[d]. *See* Section XI.A.4.b.2.

- (6) Element 3[e]: data objects in the sixth plurality of data objects are at least partially determined based on when requests for those data objects were last received at the sixth server; and**

163. In my opinion, Dilley or Dilley-Menon teaches Element 3[e]. *See* Section XI.A.4.b.2.

- (7) Element 3[f]: data objects in the seventh plurality of data objects are at least partially determined based on when requests for those data objects were last received at the seventh server.**

164. In my opinion, Dilley or Dilley-Menon teaches Element 3[f]. *See* Section XI.A.4.b.2.

- c. Claim 4: The cloud data persistence system of claim 1, wherein the first plurality of data objects comprises all data objects in the second plurality of data objects, the third plurality of data objects, the fourth plurality of data objects, the fifth plurality of data objects, the sixth plurality of data objects, and the seventh plurality of data objects.**

165. In my opinion, Dilley or Dilley-Menon teaches Claim 4.

166. Dilley teaches that the first plurality of data objects (i.e., metadata stored at the origin server) comprises all data objects in the second, third, fourth, fifth, sixth, and seventh plurality of data objects because the origin server is the central hub that all servers revert to if their location and the other server locations in the hierarchy do not have the requested data objects. Specifically, Dilley teaches that “a given content provider operates an origin server” and “a cache hierarchy is

established in the CDN comprising a given edge server region and ... a subset of well-connected ‘core’ parent server regions” that respond to requests for content. Ex-1005 ¶14. If “a given object request cannot be serviced in the given edge region, instead of contacting the origin server, the request is provided to ... a given one of the subset of server regions for handling, preferably as a function of metadata associated with the given object request.” Ex-1005 ¶14. The request is only forwarded “to the origin server if it cannot be serviced by the intermediate node in the hierarchy.” Ex-1005 ¶14; *see also id.*, ¶¶24 (“The distribution infrastructure consists of on demand or push-based mechanisms that move content from the origin server to the surrogates.”), 51 (“[P]referably only edge-of-network content servers (running CDN software) get such content from the parent regions (if they do not already have it), and parent regions get the content (if needed) from the origin server(s).”).

167. As I also explained in the Technological Background section of my Declaration (*see* Section VII.F), it was a well-known concept by 2010 that data objects at the origin server comprise all of the data objects in the cache and application servers. Ex-1031 (Mao) ¶9; Ex-1029 (Seed), Fig. 1.

- d. Claim 6: The cloud data persistence system of claim 5 wherein the first plurality of data objects comprises all data objects present in the second plurality of data objects and the third plurality of data objects.**

168. In my opinion, Dilley or Dilley-Menon teaches Claim 6. *See* Section XI.A.4.c.

- e. Claim 7**

- (1) Element 7[pre]: The cloud data persistence system of claim 5,**

169. In my opinion, Dilley or Dilley-Menon teaches Element 7[pre]. *See* Section XI.A.2.a.

- (2) Element 7[a]: wherein the third server stores a fourth plurality of data objects and is configured to request the data object from the first server if the data object is not present in the fourth plurality of data objects, wherein:**


170. In my opinion, Dilley or Dilley-Menon teaches Element 7[a]. *See* Section XI.A.2.i.


- (3) Element 7[b]: the third server and the second server are data persistence cache servers; and**

171. In my opinion, Dilley or Dilley-Menon teaches Element 7[b]. *See* Sections XI.A.2.g and XI.A.2.j.

- (4) **Element 7[c]: the third server is communicatively coupled with a sixth and seventh server.**

172. In my opinion, Dilley or Dilley-Menon teaches Element 7[c]. *See* Sections XI.A.2.k and XI.A.2.r-XI.A.2.w; Ex-1005, Figs. 3-4 (brown arrows and dashed box added).





f. Claim 8

**(1) Element 8

```
[pre]
```

: The cloud data persistence system of claim 5,**

173. In my opinion, Dilley or Dilley-Menon teaches Element 8

```
[pre]
```

. See Section XI.A.2.a.

(2) Element 8[a] wherein the fifth server stores a fifth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fifth plurality of data objects, wherein:

174. In my opinion, Dilley or Dilley-Menon teaches Element 8[a]. See Section XI.A.2.o.

(3) Element 8[b]: the fifth server and the fourth servers are application servers; and

175. In my opinion, Dilley or Dilley-Menon teaches Element 8[b]. *See* Sections XI.A.2.m and XI.A.2.p.

(4) Element 8[c]: the fifth server is configured to receive requests from a second client for the data object; and


176. In my opinion, Dilley or Dilley-Menon teaches Element 8[c]. *See* Sections XI.A.2.q.

(5) Element 8[d]: the fifth server is geographically separated from the fourth server.

177. In my opinion, Dilley or Dilley-Menon teaches Element 8[d].

178. Ground 1: Dilley teaches that edge servers “are located at strategic locations (e.g., Internet network access points, and the like) for delivering copies of content to requesting end users.” Ex-1005 ¶24. Dilley also teaches “*a set of regions 402 a-n*, each of which comprises *a set of content (or ‘edge’) servers 404 a-n*” and a “given subset of parent regions ... 408 a-n,” each of which “includes a cluster of ‘parent’ servers, for example, servers 410 a-n that, ... may be co-located in large hosting data centers in well-connected ‘core’ portions of the Internet web hosting infrastructure.” Ex-1005 ¶28. Dilley further teaches that the “[t]he number of parent regions in the given subset is variable and *their particular location* and configuration ... *is selectable*.” Ex-1005 ¶28. As shown in annotated Figure 4 below,

the fourth server (purple) and fifth server (orange) are geographically separated in different “edge region[s]” of the CDN. Ex-1005, Fig. 3 (brown arrow added).



179. Ground 2: This limitation is also taught by Menon. Menon teaches that “[w]hen a user requests the content either explicitly (such as by making a selection from a video play list) or implicitly by accessing such a web page or link within a web page ... *the edge server closer to the user* is directed to serve the replicated content to the user.” Ex-1007 ¶15. Menon further teaches that “[e]dge server ‘closeness’ may be defined in a number of ways” including “*geographic proximity*.” Ex-1007 ¶15. Thus, a POSITA would have understood that Menon’s edge servers are geographically separated from one another, as they would serve different clusters

of users in different locations, such as different cities or nations. Ex-1007 ¶¶86 (“Intelligence embodied in the network (such as proximity of the user to an edge server near his/her geographic area) directs the request to that server.”), 58 (“Rights information may for example provide an indicator identifying either when, where, and/or to whom the content may be delivered (or when, where, and/or to whom delivery is prohibited). This permission or prohibition may also or alternatively for example be associated with the geographic or nation state location of the edge server, particular domains, be in the form of a list or database, or according to other rules or policies.”). In the Dilley-Menon combination, Dilley’s fourth and fifth servers would be geographically separated such that the edge server in closest proximity to the user is used to service the request. In addition to the reasons described in Section XI.A.1, a POSITA would have been motivated to implement Menon’s teachings in Dilley because “[s]erving content from a cache closer to the end user [] improves performance.” Ex-1007 ¶18.

180. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the application servers would be geographically separated from one another. Ex-1029 (Seed), 3:36-37, Fig. 1; Ex-1031 (Mao) ¶50.

- g. Claim 9: The cloud data persistence system of claim 5, wherein the request from the first client for the data object comprises a universal resource indicator.**

181. In my opinion, Dilley or Dilley-Menon teaches Claim 9.

182. Ground 1: Dilley teaches that the request from the first client for the data object comprises a universal resource indicator (“URL”). Dilley teaches that “content is tagged for delivery from the CDN using a content migrator or rewrite tool 106 operated ... at a participating content provider server,” which “*rewrites embedded object URLs to point to the CDNSP domain.*” Ex-1005 ¶25. Moreover, “[a] request for tagged content is resolved through a CDNSP-managed DNS to identify a ‘best’ region, and then to identify an edge server within the region that is not overloaded and that is likely to host the requested content.” Ex-1005 ¶25. Dilley further teaches that “[m]etadata is delivered to the CDN edge servers using a convenient transport mechanism” and “*[o]bject metadata may be specified by customer*, by directory, by object type, or in any other convenient manner.” Ex-1005 ¶30. An example of “representative” CDN metadata may include: “<md name=“use—hierarchy”>OA</md>,” where “OA indicates that the edge server is to go direct to the *(O)rigin-server specified in the (A)RL*), where the ARL is *a modified URL (or portion thereof) that points an end user browser to attempt to retrieve the object from the CDN.*” Ex-1005 ¶¶31-33; *see also id.*, ¶¶41 (“while the actual format of the edge server request to its parent varies, it may [be] of the form

aSerial.rRegion.cdns.net”), 49-50 (describing the edge server’s process for retrieving content based on a user request). A POSITA would thus have understood that the client request is in the form of a URL, and that a URL is a specific form of a universal resource indicator. Ex-1030 (IETF RFC 3986), 7 (“A URI can be further classified as a locator, a name, or both. The term ‘Uniform Resource Locator’ (URL) refers to the subset of URIs that, in addition to identifying a resource, provide a means of locating the resource by describing its primary access mechanism (e.g., its network ‘location’).”). It also would have been obvious to a POSITA that a client request for metadata could include a URL, as taught by Dilley. This is also taught by Menon.

183. Ground 2: Menon teaches that edge servers “maintain a directory that organizes metadata for individual content items typically in the form of a” URL “for the item.” Ex-1007 ¶47. For example, Menon teaches that “[u]ser selection of particular video content whether by URL or by other designator or identifier *may be recorded in the metadata database*, regardless of whether the video object is cached or not.” Ex-1007 ¶48. Moreover, “when *a user makes a selection (such as by identifying a URL) for a video on an edge server (ES)*, the metadata (MD) corresponding to the selected URL is examined to see if the corresponding video object is in the cache.” Ex-1007 ¶50; *see also id.*, ¶86 (“When a user selects a video title displayed on his computer, ... the user essentially transmits a URL or other

globally unique identification for the video (or other content item) to the network.”). In the Dilley-Menon combination, the client request for the metadata in Dilley’s system would include a URL, as taught by Menon. In addition the reasons described in Section XI.A.1, a POSITA would have been motivated to implement Dilley’s client request as a URL, as taught by Menon, because client requests on a webpage over the Internet are typically in the form of a URL.

184. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the requests for the data objects would comprise a URL. Ex-1024 (Coates), 22:24-27, Fig. 20; Ex-1032 (Davis), 7:10-13.

h. Claim 10: The cloud data persistence system of claim 5 wherein each of the first plurality of data objects comprise a universal resource locator and a title linked to a media asset of the plurality of media assets.

185. In my opinion, Dilley or Dilley-Menon teaches Claim 10.

186. Ground 1: Dilley teaches that the metadata stored at the first server comprises a URL linked to the media asset for the reasons described in the previous section and Sections XI.A.2.b and XI.A.2.d. Specifically, Dilley defines metadata as “the set of all control options and parameters for the object (e.g., coherence information, origin server identity information, load balancing information, customer code, other control codes, etc.).” Ex-1005 ¶25. Dilley also teaches that the metadata may include embedded URLs that point to the content (Ex-1005 ¶¶25, 31-

40) and “an object’s serial number.” Ex-1005 ¶36. It would have been obvious to a POSITA that the metadata would also include some sort of title or identifier for the content, otherwise the correct requested media would not be delivered to the requestor. This is also taught by Menon.

187. Ground 2: Menon teaches that metadata could include a URL for the reasons described in the previous section. Specifically, Menon teaches that “when a user makes a selection (such as by identifying a URL) for a video ... *the metadata (MD) corresponding to the selected URL* is examined to see if the corresponding video object is in the cache.” Ex-1007 ¶50; *see also id.*, ¶¶47 (“a directory [] *organizes metadata* for individual content items typically *in the form of a ... (URL)* for the item”), 86 (“the user essentially transmits a URL or other globally unique identification for the video (or other content item) to the network”). Menon further teaches that “the *metadata may identify* and describe *the video object or asset*, such as its physical characteristics (formats, size, bit rate, or the like physical or logical characteristics).” Ex-1007 ¶57. The metadata may also “provide documentary information, such as *textual information about the video.*” Ex-1007 ¶61. Moreover, Menon teaches that “[w]hen a *user selects a video title displayed* on his computer, TV set via a set-top box, or other information appliance or access method, the user essentially transmits a URL or other globally unique identification for the video (or other content item) to the network.” Ex-1007 ¶86. A POSITA would have

understood that the textual information about the video and/or the video title displayed correspond to the recited “title linked to a media asset.” In the Dilley-Menon combination, Dilley’s metadata would comprise a title as part of an identifier that uniquely identifies the content, as taught by Menon. Furthermore, a POSITA would have understood that each movie would also have a unique identifier, so that even movies with identical titles can be properly distinguished from each other and so that the media of those movies do not become mixed up. In addition to the reasons described in Section XI.A.1, a POSITA would have been motivated to include a title in Dilley’s metadata, as taught by Menon, because this is a typical parameter included in a metadata file.

188. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), it was known that each of the data objects would comprise a URL and a title linked to the media asset. Ex-1026 (Geilfuss) ¶36, Fig. 2; Ex-1021 (Eaton), 8:26–36; Ex-1031 (Mao) ¶51; Ex-1033 (Leighton), 1:54-2:6.

- i. **Claim 11: The cloud data persistence system of claim 5, wherein data objects in the second plurality of data objects and the third plurality of data objects are at least partially determined based on a frequency of requests for those data objects received from clients.**

189. In my opinion, Dilley or Dilley-Menon teaches Claim 11. *See* Section XI.A.4.b.

- j. **Claim 12: The cloud data persistence system of claim 5, wherein data objects in the second plurality of data objects and the third plurality of data objects are at least partially determined based on when requests for those data objects were last received from clients.**

190. In my opinion, Dilley or Dilley-Menon teaches Claim 12. *See* Section XI.A.4.b.

- k. **Claim 13: The cloud data persistence system of claim 5, wherein the first plurality of data objects comprises all data objects in the second plurality of data objects, the third plurality of data objects, the fourth plurality of data objects, and the fifth plurality of data objects.**

191. In my opinion, Dilley or Dilley-Menon teaches Claim 13. *See* Sections XI.A.4.c and XI.A.4.d.

B. Grounds 3-4: Claims 1-13 are obvious over Russell alone or Russell in view of Menon.

- 1. **A POSITA would have been motivated to combine Russell’s teachings with Menon’s teachings and would have had a reasonable expectation of success.**

192. In addition to the reasons set forth below, a POSITA would have been motivated to combine the teachings of Russell and Menon and had a reasonable expectation of success in doing so because each relates to a content delivery network that uses a tiered server architecture to efficiently distribute media content to users. Ex-1006, Abstract, ¶6; Ex-1007, Abstract, ¶42, Fig. 3.

193. Russell teaches a system for delivering content over a network. Ex-1006, Abstract. Russell’s network architecture includes “a main server layer having

at least one main server system, a parent server layer having a plurality of parent server systems, and an edge server layer having a plurality of edge server systems.” Ex-1006 ¶9. Russell further teaches that the main server “publishes instructions” to the parent servers and edge servers in the form of “an XML document containing a list of all content items” available on the main server with additional information such as whether a content item is “hot” and “geographic restrictions.” Ex-1006 ¶81. Based on the instructions, the edge server “may determine whether or not” it is “storing all of the appropriate content items and, if not,” request them “from a parent server.” Ex-1006 ¶82. Russell also teaches that “the main server provides the user’s UND with a” URL that allows “the user to connect to an edge server to download an encrypted electronic file containing the selected content item.” Ex-1006 ¶12.

194. Menon teaches a metadata enabled server for distributing content objects to users over a network in response to user requests. Ex-1007 ¶35. Menon further teaches that a directory “organizes metadata for individual content items typically in the form of a ... (URL) for the item.” Ex-1007 ¶47. Menon also teaches that the metadata is distributed “to a plurality of locations (e.g to the origin server and edge servers, or to the origin server, edge servers, and directory servers[]).” Ex-1007 ¶84.

195. To the extent PO argues that the claims’ recited “data objects” must comprise metadata, a POSITA would have been motivated to implement Russell’s

instructions to include metadata because it was known that metadata may be “used as a vehicle and mechanism to enable intelligent decisions to be made on content distribution system operation.” Ex-1007, Abstract. A POSITA would have also been motivated to implement Russell’s instructions with metadata in the form of a URL that links to the content, as taught by Menon, because URLs were a well-known mechanism for easily locating and accessing web resources and simplifying communication between applications. URLs may also carry data in key-value parameters in the “query string” portion of the URL. In the Russell-Menon combination, the instructions stored at Russell’s main server and distributed to the other servers would include metadata in the form of a URL that links to the content, as taught by Menon.

196. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), it was known that each of the data objects would comprise a URL linked to the media asset. Ex-1026 (Geilfuss) ¶36, Fig. 2; Ex-1021 (Eaton), 8:26–36; Ex-1031 (Mao) ¶51; Ex-1033 (Leighton), 1:54-2:6.

197. A POSITA would have had a reasonable expectation of success in combining the teachings of Russell and Menon. Both describe methods for enabling efficient distribution of media content within a tiered server network infrastructure. Ex-1006, Abstract, ¶6; Ex-1007, Abstract, ¶42, Fig. 3. A POSITA would thus have understood that these references disclose interrelated teachings based on well-

understood technologies that would have been amenable to various well-understood and predictable combinations.

2. Independent Claim 1

a. Element 1[pre]: A cloud data persistence system for distribution of data, the cloud data persistence system comprising:

198. In my opinion, Russell alone or Russell in view of Menon (“Russell-Menon”) teaches the preamble.

199. Russell teaches a system for delivering content over a network. Ex-1006, Abstract. The system’s “network architecture includes three or more layers of servers, including a main server supporting an interface for recipient processors to request content items, a plurality of parent servers, and a plurality of edge servers.” Ex-1006, Abstract. Russell further teaches that “[a] copy of each content item available through the service is distributed from the main server to each parent servers for storage” and “[p]arent servers distribute content to edges servers, pursuant to instructions from the main server or requests from the edge servers.” Ex-1006, Abstract.

200. As I also explained in the Technological Background section of my Declaration (*see* Section VII.A), cloud data persistence systems for distributing data were known by 2010. Ex-1022 (Kovacs), 1:16-28; Ex-1023 (Aikas) ¶25; Ex-1024

(Coates), 20:27-28; Ex-1025 (Girbal), Fig. 1; Ex-1031 (Mao), Fig. 2; Ex-1033 (Leighton), 1:18-53.

b. Element 1[a][1]: a first server which stores a first plurality of data objects and is communicatively coupled with a plurality of cache servers, wherein:

201. In my opinion, Russell or Russell-Menon teaches Element 1[a][1].

202. Ground 3: As shown in annotated Figures 1-2 below, Russell teaches a first server (i.e., main server (red)), which is communicatively coupled with a plurality of cache servers called parent servers (green, yellow). Ex-1006, Figs. 1-2.




FIG. 1




FIG. 2

Russell’s network architecture includes “*a main server layer* having at least one main server system, *a parent server layer* having a plurality of parent server systems, and an edge server layer having a plurality of edge server systems.” Ex-1006 ¶9. Russell teaches that “[*a copy of each content item* available through the server *may be* stored by the main server and distributed to and *stored by each parent server.*” Ex-1006 ¶10. Russell further teaches that “parent servers 14 provide three centralized locations for data storage” and “may be operated as *cache servers*, for cache storage and distribution of content items to the edge servers 16.” Ex-1006 ¶64; *see also id.*, 61 (“In embodiments as described above in which the main server system 12 stores content items that are made available through the service, the main server system 12 may include or operate with a content storage device or facility 36. In one embodiment, the content storage facility 36 stores all of the content items made available through the service.”).

203. Russell also teaches that the main server (i.e., first server) “stores a first plurality of data objects.” As explained in Section XI.A.2.b, a POSITA would have understood that a “data object” is a structure comprising metadata about a particular media asset, distinct from the media asset itself. As shown in Figure 3 below, Russell teaches that main server system 12 includes CMS server 38 “for controlling the distribution of content and *instructions to parent servers 14*.” Ex-1006 ¶63, Fig. 3; *see also id.* ¶73 (“*instructions* and content items *from the CMS server subsystem 38* are communicated to the parent servers 14”). Russell further teaches that “main server 12 publishes an XML document containing a list of all content items ... available” on main server 12 with additional information such as, for example, whether a content item is marked “hot” and “geographic restrictions (so that a content item, such as a movie, that is not available for download in a region will never be stored on edge servers in that region).” Ex-1006 ¶81. The instructions stored at main server system 12 and distributed to the other servers, for example, as an XML document describing the content and how it should be distributed correspond to the recited “plurality of data objects” because they are essentially metadata that provide information about the content and point to its location. To the extent PO argues that “data objects” must comprise metadata, it would have been obvious that Russell’s instructions could also include metadata, as taught by Menon.




FIG. 3

204. Ground 4: Menon teaches a “*metadata enabled server* for distributing content object to user over network in response to user request” which includes “a storage device coupled to server and *storing metadata* describing content objects accessible to the server *including a location* from where a content object is stored and may be directed to the user; and a controller for distributing the content object to the user using the metadata.” Ex-1007 ¶35. Russell or Russell-Menon thus teaches the “data objects” limitation under either parties’ construction because it teaches “metadata about a particular media asset, distinct from the asset itself” and more broadly teaches a “data structure that contains information about a media asset.”

205. In the Russell-Menon combination, the instructions stored at Russell’s CMS server subsystem and distributed to the other servers would include metadata describing the content and a location from where the content object may be retrieved,

as taught by Menon. A POSITA would have been motivated to combine Russell-Menon and had a reasonable expectation of success in doing so for the reasons described in Section XI.B.1.

206. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), origin servers for storing a plurality of data objects comprising a link to a plurality of media assets were known by 2010. Ex-1026 (Geilfuss) ¶31, Fig. 1. Moreover, as I explained in the Technological Background section of my Declaration (*see* Section VII.B), it was known that the origin servers in a cloud data persistence system would be communicatively coupled with a plurality of cache servers. Ex-1027 (Glasser) ¶19, Fig. 1.

c. Element 1[a][2]: the first server is located at a first location;

207. In my opinion, Russell or Russell-Menon teaches Element 1[a][2].

208. Russell teaches that “[e]dge servers may be located in populated areas, such as in or near large cities,” “[p]arent servers may be spaced across the service region to provide support for edge servers,” and “content items are distributed to and stored on the parent and edge servers, from the main server.” Ex-1006 ¶10. As shown in annotated Figure 2 below, the first server (i.e., main server (red)) is located at a first location relative to the other servers in the system. Ex-1006, Fig. 2.




FIG. 2

- d. **Element 1[a][3]: each of the first plurality of data objects comprise a link to one of a plurality of media assets; and**

209. In my opinion, Russell or Russell-Menon teaches Element 1[a][3].

210. Ground 3: Russell teaches that each of the first plurality of data objects (i.e., instructions) comprise a link to one of a plurality of media assets (i.e., content) because, as explained in Section XI.B.2.b, they are essentially metadata that point to the location of the content. Moreover, Russell teaches that “[w]hen a user purchases a license for a selected content item *the main server provides the user’s UND with a URL* (uniform resource locator) that will allow the user *to connect to an edge server to download* an encrypted electronic file containing the selected content item.” Ex-1006 ¶12. A POSITA would have also understood that using links such as URLs to identify and fetch content was well-known, as taught by Menon. URLs were well known since the beginning of the World Wide Web and first published as

an Internet Engineering Task Force Request for Comments in December 1994 as RFC 1738.

211. Ground 4: Menon teaches that the “metadata consists of (i) a globally unique identifier that identifies and optionally describes the content associated with the metadata, and (ii) a location or locations at which the content may be found, usually in the form of at least the origin server identifier.” Ex-1007 ¶63. For example, “when a user makes a selection (such as by identifying a URL) for a video ... *the metadata (MD) corresponding to the selected URL* is examined to see if the corresponding video object is in the cache.” Ex-1007 ¶50; *see also id.*, ¶¶47 (“a directory [] *organizes metadata* for individual content items typically *in the form of a ... (URL)* for the item”), 86 (“the user essentially transmits a URL or other globally unique identification for the video (or other content item) to the network”). The metadata is distributed “to a plurality of locations (e.g to the origin server and edge servers, or to the origin server, edge servers, and directory servers[]).” Ex-1007 ¶84. In the Russell-Menon combination, Russell’s first plurality of data objects (i.e., Russell’s instructions and/or Menon’s metadata) would include a URL that links to the media asset.


212. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), it was known that each of the data objects would

comprise a URL linked to the media asset. Ex-1026 (Geilfuss) ¶36, Fig. 2; Ex-1021 (Eaton), 8:26–36; Ex-1031 (Mao) ¶51; Ex-1033 (Leighton), 1:54-2:6.

- e. **Element 1[a][4]: the plurality of media assets are stored at a second location separate from the first location;**

213. In my opinion, Russell or Russell-Menon teaches Element 1[a][4].

214. Ground 3: As shown in annotated Figure 2 below, Russell teaches that the plurality of media assets may be stored at a separate source server (maroon) location separate from the main server (red) location. Ex-1006, Fig. 2.



Russell teaches that “[t]he *content items* (such as, movies) *may be supplied to the main server 12* by content owners or holders (such as movie studios or other suitable sources), for example, *from a source server 18.*” Ex-1006 ¶32. Specifically, “main server system 12 may connect to” source server 18 to receive content items “and/or instructions relating to distribution of content items from the source server(s) ... through the network 10 or through any other suitable means, including, but not

limited to, direct connection or connection through a LAN or other network.” Ex-1006 ¶57. Additionally, Russell’s parent servers are located at a different location from the main server. Therefore, the plurality of media assets are stored at a second location separate from the first location.

215. Ground 4: Menon also teaches that the plurality of media assets are stored at a second location separate from the first location for the reasons described in Section XI.A.2.e.

216. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), it was a well-known concept for the origin server to be at a first location and for the plurality of media assets to be at a second location separate from the first location. Ex-1026 (Geilfuss), Figs. 1-2; Ex-1031 (Mao), Fig. 2.

f. Element 1[b][1]: a second server that stores a second plurality of data objects and is configured to request a data object from the first server if the data object is not present in the second plurality of data objects, wherein:

217. In my opinion, Russell or Russell-Menon teaches Element 1[b][1].

218. Ground 3: As shown in annotated Figures 1-2 below, Russell teaches a second server (green). Ex-1006, Figs. 1-2.




FIG. 1




FIG. 2

219. Russell teaches that the second server stores a second plurality of data objects (i.e., instructions) and is configured to request a data object from the first server if the data object is not present in the second plurality of data objects. Specifically, Russell teaches that “[a] copy of each content item available through the server may be stored by the main server and distributed to and stored by each parent server.” Ex-1006 ¶10. Russell further teaches that main server system 12

includes CMS server 38 “for controlling the distribution of content and *instructions to parent servers 14*.” Ex-1006 ¶¶63, 73 (“instructions and content items from the CMS server subsystem 38 are communicated to the parent servers 14”). Moreover, Russell teaches that parent servers 14 “determine which content items (such as movies) to store, for example, using a combination of instructions originating from the main server 12 and/or a least recently used (‘LRU’) algorithm.” Ex-1006 ¶68.

220. Ground 4: It would have been obvious to implement Russell’s instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.


221. As I also explained in the Technological Background section of my Declaration (*see* Section VII.C), cache servers that store a plurality of data objects and request a data object from the origin server if the data object is not present at the cache server were known by 2010. Ex-1022 (Kovacs), 8:64-65, Fig. 4.

g. Element 1[b][2]: the second server is one of the plurality of cache servers; and

222. In my opinion, Russell or Russell-Menon teaches Element 1[b][2].

223. Russell teaches that the second server is one of the plurality of cache servers. Specifically, Russell teaches that “*parent servers 14* may be operated as *cache servers*, for cache storage and distribution of content items to the edge servers 16.” Ex-1006 ¶64, Fig. 2. Russell further teaches that “one parent cache server may be located near the west coast, another parent cache server may be located near the

east coast and a further parent cache server may be located in the central portion of the country.” Ex-1006 ¶64.





h. Element 1[b][3]: the second server is communicatively coupled with a first plurality of application servers;

224. In my opinion, Russell or Russell-Menon teaches Element 1[b][3].

225. As shown in annotated Figures 1-2 below, Russell’s second server (green) is communicatively coupled to a plurality of edge servers such as a fourth server (purple) and fifth server (orange). Ex-1006, Figs. 1-2. Russell teaches that its network of servers are “*connected for communication* on a wide area network, such as the Internet, *with recipient processors*, for example, network enabled user network devices (UNDs).” Ex-1006 ¶9. The server architecture includes “a main server layer having at least one main server system, *a parent server layer* having a

plurality of parent server systems, and *an edge server layer* having a plurality of edge server systems.” Ex-1006 ¶9, Cl. 1.



226. Russell’s edge servers are application servers. Russell teaches that “[t]he *edge servers* are distributed throughout a region in which the system *provides content delivery services to users* on UNDs.” Ex-1006 ¶10. Moreover, Russell

teaches that “[a] user of a UND interacts with the main server for browsing and conducting commercial transactions” but “the main server redirects the UNDs (transparently to the users, in some embodiments) to edge servers 16 for downloading content items.” Ex-1006 ¶33; *see also id.*, ¶¶30 (“functions and operations of the UNDs and servers described herein may be implemented in any suitable manner, including, but not limited to, software running on the associated server or UND”), 55 (“the network system implements a division of functions of business logic and storage between the main server and ... edge servers” where “downloads to user’s are handled by the edge servers”). Each edge server receives instructions from the parent and/or main servers and “determine[s] whether or not its associated storage facility is storing all of the appropriate content items and, if not,” it requests “appropriate content items from a parent server.” Ex-1006 ¶82. Russell further teaches that “[o]nce the content items are sufficiently distributed to edge servers 16, users may access the system through the interface ... supported by the main server system 12” to “search and select content items and submit a request to download content items, and purchase (or otherwise obtain) a license for content items.” Ex-1006 ¶87. A POSITA would have understood that the edge servers are application servers because they are implemented by software and act as a middle layer between the client interface and the cache servers to respond to client requests for content.

227. As I also explained in the Technological Background section of my Declaration (*see* Section VII.C), it was known that cache servers in a cloud data persistence system would be communicatively coupled with a plurality of application servers. Ex-1027 (Glasser) ¶¶24; Ex-1028 (Kasten) ¶16, Fig. 1.

- i. **Element 1[c][1]: a third server that stores a third plurality of data objects and is configured to request the data object from the first server if the data object is not present in the third plurality of data objects, wherein:**

228. In my opinion, Russell or Russell-Menon teaches Element 1[c][1].

229. Ground 3: As shown in annotated Figures 1-2 below, Russell teaches a third server (yellow), which is a server in the parent region. Ex-1006, Figs. 1-2.




FIG. 1




FIG. 2

230. Russell teaches that the third server stores a third plurality of data objects (i.e., instructions) and is configured to request a data object from the first server if the data object is not present in the third plurality of data objects for the reasons described in Section XI.B.2.f.

231. Ground 4: It would have been obvious to implement Russell's instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

232. As I also explained in the Technological Background section of my Declaration (*see* Section VII.C), cache servers that store a plurality of data objects and request a data object from the origin server if the data object is not present at the cache server were known by 2010. Ex-1022 (Kovacs), 8:64-65, Fig. 4.

j. Element 1[c][2]: the third server is one of the plurality of cache servers; and

233. In my opinion, Russell or Russell-Menon teaches Element 1[c][2].



234. Russell teaches that the third server is one of the plurality of cache servers for the reasons described in Section XI.B.2.g.

k. Element 1[c][3]: the third server is communicatively coupled with a second plurality of application servers; and

235. In my opinion, Russell or Russell-Menon teaches Element 1[c][3].

236. As shown in annotated Figures 1-2 below, Russell's third server (yellow) is communicatively coupled to a plurality of edge servers such as a sixth server (blue) and seventh server (pink). Ex-1006, Figs. 1-2 (brown dashed boxes added). As explained in Section XI.B.2.h, the edge servers are application servers.



237. As I also explained in the Technological Background section of my Declaration (*see* Section VII.C), it was known that cache servers in a cloud data persistence system would be communicatively coupled with a plurality of application servers. Ex-1027 (Glasser) ¶24; Ex-1028 (Kasten) ¶16, Fig. 1.



- I. **Element 1[d][1]: a fourth server that stores a fourth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fourth plurality of data objects, wherein:**

238. In my opinion, Russell or Russell-Menon teaches Element 1[d][1].

239. Ground 3: As shown in annotated Figures 1-2 below, Russell teaches a fourth server (purple). Ex-1006, Figs. 1-2. Russell teaches that “edge servers are distributed throughout a region in which the system provides content delivery services to users on UNDs” and “content items are distributed to and stored on the parent and edge servers, from the main server.” Ex-1006 ¶10.



240. Russell teaches that the fourth server stores a fourth plurality of data objects (i.e., instructions) and is configured to request the data object from the second server if the data object is not present in the fourth plurality of data objects. Specifically, Russell teaches that edge servers 16 “determine which content items (such as movies) to store, for example, *using* a combination of *instructions* originating from the main server 12 *and/or* a least recently used (*‘LRU’*) *algorithm*.” Ex-1006 ¶68. Specifically, “*main server 12 publishes instructions to* the parent servers 14 and the *edge servers 16* and the parent servers republish instructions to edge servers.” Ex-1006 ¶81. Russell further teaches “[b]ased on such instructions, each edge server 16 may determine whether or not its associated storage facility is storing all of the appropriate content items and, if not, the *edge server may request appropriate content items from a parent server*.” Ex-1006 ¶82.

241. Ground 4: It would have been obvious to implement Russell’s instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

242. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), application servers that store a plurality of data objects and request the data object from the cache servers if the data object is not present at the application server were known by 2010. Ex-1022 (Kovacs), 8:56-64, Fig. 7; Ex-1029 (Seed), 2:63-3:1.

m. Element 1[d][2]: the fourth server is one of the first plurality of application server; and


243. In my opinion, Russell or Russell-Menon teaches Element 1[d][2].

244. Russell teaches that the fourth server is one of the first plurality of application servers for the reasons described in the previous section and Section XI.B.2.h, namely, because the second server is communicatively coupled to the fourth server and Russell teaches that its edge servers are application servers.

n. Element 1[d][3]: the fourth server is configured to receive a request from a first client for the data object;

245. In my opinion, Russell or Russell-Menon teaches Element 1[d][3].

246. Ground 3: Russell teaches that the fourth server is configured to receive a request from a first client for the data object. As shown in annotated Figure 1 below, the fourth server (purple) communicates with the client (grey) to receive requests for data objects and provide the corresponding content. Ex-1006, Fig. 1.



247. Russell teaches that “edge servers 16 provide distributed storage which, preferably, is comparatively local to users” and “may be strategically located near areas in which a high volume of usage is expected, such as on or near a college campus or associated with a college server.” Ex-1006 ¶67. Russell further teaches that “when the user attempts to access the content item, the URL provided to the user will be resolved by a DNS (domain name server) to connect the user to the appropriate edge server 16 in the system network for the download request.” Ex-1006 ¶89. For example, “the URL may include address information for a specific edge server 16 to which the user’s request is to be sent.” Ex-1006 ¶89; *see also id.*, ¶¶33 (“the main server redirects the UNDS (transparently to the users, in some embodiments) to edge servers 16 for downloading content items”).


248. Ground 4: It would have been obvious to implement Russell’s instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

249. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the application servers were configured to receive requests from clients for the data objects. Ex-1029 (Seed), 4:37-40; Ex-1033 (Leighton), 1:32-36.

- o. **Element 1[e][1]: a fifth server that stores a fifth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fifth plurality of data objects, wherein:**

250. In my opinion, Russell or Russell-Menon teaches Element 1[e][1].

251. Ground 3: As shown in annotated Figures 1-2 below, Russell teaches a fifth server (orange). Ex-1006, Figs. 1-2.






FIG. 2

252. Russell teaches that the fifth server stores a fifth plurality of data objects (i.e., instructions) and is configured to request the data object from the second server if the data object is not present in the fifth plurality of data objects for the reasons described in Section XI.B.2.1.

253. Ground 4: It would have been obvious to implement Russell's instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

254. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), application servers that store a plurality of data objects and request the data object from the cache servers if the data object is not present at the application server were known by 2010. Ex-1022 (Kovacs), 8:56-64, Fig. 7; Ex-1029 (Seed), 2:63-3:1.

- p. **Element 1[e][2]: the fifth server is one of the first plurality of application servers; and**

255. In my opinion, Russell or Russell-Menon teaches Element 1[e][2].

256. Russell teaches that the fifth server is one of the first plurality of application servers for the reasons described in the previous section and Section XI.B.2.h, namely, because the second server is communicatively coupled to the fifth server and Russell teaches that its edge servers are application servers.

- q. **Element 1[e][3]: the fifth server is configured to receive a request from a second client for the data object;**

257. In my opinion, Russell or Russell-Menon teaches Element 1[e][3].

258. Ground 3: Russell teaches that the fifth server is configured to receive a request from a second client for the data object. As shown in annotated Figure 1 below, the fifth server (orange) communicates with the client (grey) to receive requests for data objects and provide the corresponding content. Ex-1006, Fig. 1. *See also* Section XI.B.2.n.




FIG. 1


259. Ground 4: It would have been obvious to implement Russell's instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

260. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the application servers were configured to receive requests from clients for the data objects. Ex-1029 (Seed), 4:37-40; Ex-1033 (Leighton), 1:32-36.

- r. **Element 1[f][1]: a sixth server that stores a sixth plurality of data objects and is configured to request the data object from the third server if the data object is not present in the sixth plurality of data objects, wherein:**

261. In my opinion, Russell or Russell-Menon teaches Element 1[f][1].

262. Ground 3: As shown in annotated Figures 1-2 below, Russell teaches a sixth server (blue). Ex-1006, Figs. 1-2.






FIG. 2

263. Russell teaches that the sixth server stores a sixth plurality of data objects (i.e., instructions) and is configured to request the data object from the third server if the data object is not present in the sixth plurality of data objects for the reasons described in Section XI.B.2.1.

264. Ground 4: It would have been obvious to implement Russell's instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

265. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), application servers that store a plurality of data objects and request the data object from the cache servers if the data object is not present at the application server were known by 2010. Ex-1022 (Kovacs), 8:56-64, Fig. 7; Ex-1029 (Seed), 2:63-3:1.

s. Element 1[f][2]: the sixth server is one of the second plurality of application servers; and

266. In my opinion, Russell or Russell-Menon teaches Element 1[f][2].

267. Russell teaches that the sixth server is one of the second plurality of application servers for the reasons described in the previous section and Sections XI.B.2.h and XI.B.2.k, namely, because the third server is communicatively coupled to the sixth server and Russell teaches that its edge servers are application servers.

t. Element 1[f][3]: the sixth server is configured to receive a request from a third client for the data object; and

268. In my opinion, Russell or Russell-Menon teaches Element 1[f][3].

269. Ground 3: Russell teaches that the sixth server is configured to receive a request from a third client for the data object. As shown in annotated Figure 1 below, the sixth server (blue) communicates with the client (grey) to receive requests for data objects and provide the corresponding content. Ex-1006, Fig. 1. *See also* Section XI.B.2.n.




FIG. 1



270. Ground 4: It would have been obvious to implement Russell's instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

271. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the application servers were configured to receive requests from clients for the data objects. Ex-1029 (Seed), 4:37-40; Ex-1033 (Leighton), 1:32-36.

- u. **Element 1[g][1]: a seventh server that stores a seventh plurality of data objects and is configured to request the data object from the third server if the data object is not present in the seventh plurality of data objects, wherein:**

272. In my opinion, Russell or Russell-Menon teaches Element 1[g][1].

273. Ground 3: As shown in annotated Figures 1-2 below, Russell teaches a seventh server (pink). Ex-1006, Figs. 1-2 (brown dashed boxes added).



274. Russell teaches that the seventh server stores a seventh plurality of data objects (i.e., instruction) and is configured to request the data object from the third server if the data object is not present in the seventh plurality of data objects for the reasons described in Section XI.B.2.1.

275. Ground 4: It would have been obvious to implement Russell's instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

276. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), application servers that store a plurality of data objects and request the data object from the cache servers if the data object is not present at the application server were known by 2010. Ex-1022 (Kovacs), 8:56-64, Fig. 7; Ex-1029 (Seed), 2:63-3:1.

v. Element 1[g][2]: the seventh server is one of the second plurality of application servers; and

277. In my opinion, Russell or Russell-Menon teaches Element 1[g][2].


278. Russell teaches that the seventh server is one of the second plurality of application servers for the reasons described in the previous section and Sections XI.B.2.h and XI.B.2.k, namely, because the third server is communicatively coupled to the seventh server and Russell teaches that its edge servers are application servers.

w. Element 1[g][3]: the seventh server is configured to receive a request from a fourth client for the data object.

279. In my opinion Russell or Russell-Menon teaches Element 1[g][3].

280. Ground 3: Russell teaches that the seventh server is configured to receive a request from a fourth client for the data object. As shown in annotated Figure 1 below, the seventh server (pink) communicates with the client (grey) to

receive requests for data objects and provide the corresponding content. Ex-1006, Fig. 1 (brown dashed boxes added). *See also* Section XI.B.2.n.



281. Ground 4: It would have been obvious to implement Russell's instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

282. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the application servers were configured to receive requests from clients for the data objects. Ex-1029 (Seed), 4:37-40; Ex-1033 (Leighton), 1:32-36.

283. Thus, Russell alone or Russell-Menon renders obvious Claim 1.

3. Independent Claim 5

- a. **Element 5[pre]: A cloud data persistence system for distribution of data, the cloud data persistence system**

comprising:

284. In my opinion, Russell or Russell-Menon teaches the preamble. *See* Section XI.B.2.a.

- b. Element 5[a][1]: a first server which stores a first plurality of data objects and is communicatively coupled with a second server and a third server, wherein:**

285. In my opinion, Russell or Russell-Menon teaches Element 5[a][1] *See* Sections XI.B.2.b and XI.B.2.f-XI.B.2.j; Ex-1006, Figs. 1-2.




FIG. 1




FIG. 2

- c. **Element 5[a][2]: the first server is located at a first location;**

286. In my opinion, Russell or Russell-Menon teaches Element 5[a][2]. *See* Section XI.B.2.c.

- d. **Element 5[a][3]: each of the first plurality of data objects comprises a link to a media asset of a plurality of media assets;**

287. In my opinion, Russell or Russell-Menon teaches Element 5[a][3]. *See* Section XI.B.2.d.

- e. **Element 5[a][4]: the plurality of media assets is located at a second location ; and the second location is distinct from the first location;**

288. In my opinion, Russell or Russell-Menon teaches Element 5[a][4]. *See* Section XI.B.2.e.

- f. **Element 5[b][1]: the second server stores a second plurality of data objects and is configured to request a data object from the first server if the data object is not present in the second plurality of data objects,**

289. In my opinion, Russell or Russell-Menon teaches Element 5[b][1]. *See* Section XI.B.2.f.

- g. **Element 5[b][2]: wherein the second server is communicatively coupled with a fourth server and fifth server; and**

290. In my opinion, Russell or Russell-Menon teaches Element 5[b][2]. *See* Sections XI.B.2.h and XI.B.2.l-XI.B.2.q; Ex-1006, Figs. 1-2.




FIG. 1




FIG. 2

- h. **Element 5[c][1]: the fourth server stores a third plurality of data objects and is configured to request the data object from the second server if the data object is not present in the third plurality of data objects,**

291. In my opinion, Russell or Russell-Menon teaches Element 5[c][1]. *See* Section XI.B.2.1.

- i. **Element 5[c][2] wherein the fourth server is configured to receive a request from a first client for the data object.**

292. In my opinion, Russell or Russell-Menon teaches Element 5[c][2]. *See* Section XI.B.2.n.

293. Thus, Russell alone or Russell-Menon renders obvious Claim 5.

4. Dependent Claims 2-4 and 6-13

a. Claim 2: The cloud data persistence system of claim 1, wherein the second location comprises a content delivery network.

294. In my opinion, Russell or Russell-Menon teaches Claim 2.

295. Russell teaches Claim 2 for the reasons described in Section XI.B.2.e, namely, because the media assets (i.e., content) are stored at a source server communicatively connected to but at a separate location from the main server. Ex-1006 ¶¶32, 57. Russell also teaches that the content is stored at parent and edge servers in the network, at locations separate from the main server. Specifically, Russell teaches that “downloads to user’s are handled by the edge servers 16 and distribution to the edge servers is handled by the parent servers 14.” Ex-1006 ¶55. Moreover, CMS server 38 at the main server controls “the distribution of content and instructions to parent servers 14 ... based on instructions received from ... source server 18.” Ex-1006 ¶63; *see also id.*, ¶9.

296. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), it was a well-known concept for the origin server to be at a first location, for the plurality of media assets to be at a second location separate from the first location, and for the second location comprise a content delivery network. Ex-1026 (Geilfuss), Figs. 1-2; Ex-1031 (Mao), Fig. 2.

b. Claim 3

(1) Element 3[pre]: The cloud data persistence system of claim 1, wherein:

297. In my opinion, Russell or Russell-Menon teaches Element 3[pre]. *See* Section XI.B.2.a.

(2) Element 3[a]: data objects in the second plurality of data objects are at least partially determined based on when requests for those data objects were last received at the second server;

298. In my opinion, Russell or Russell-Menon teaches Element 3[a].

299. Ground 3: Russell teaches that edge servers 16, parent servers 14, and/or main server 12 “determine which content items (such as movies) to store ... using a *combination of instructions* originating from the main server 12 *and/or a least recently used (‘LRU’) algorithm*.” Ex-1006 ¶68. For example, Russell teaches that “parent cache servers 14 communicate some or all of the content items to various edge servers 16, initially based on instructions provided from the main server 12 and may later be *based on a combination of such instructions and other information*” such as “information obtained *from an LRU algorithm*, records of user activity (such as historical data relating to on-line activities of a user or groups of users), user input (such as user-entered preferences), or the like.” Ex-1006 ¶85; *see also id.*, ¶¶82 (“Algorithms or other suitable means, such as LRU analysis, may be used in the determination of which content items to store at the various edge servers.”), 83 (“[c]aching by edge servers 16 ... follows a combination of server ... and LRU

analysis”). Thus, Russell’s data objects in the second plurality of data objects (i.e., instructions) are at least partially determined based on when requests for those data objects were last received at the second server.

300. Ground 4: It would have been obvious to implement Russell’s instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

301. As I also explained in the Technological Background section of my Declaration (*see* Section VII.E), it was a well-known concept by 2010 that data objects are at least partially determined based on frequency of requests and/or when requests for those data objects were last received at the server. Ex-1022 (Kovacs), 8:23-27; Ex-1031 (Mao) ¶¶55, 66-67.

(3) Element 3[b]: data objects in the third plurality of data objects are at least partially determined based on when requests for those data objects were last received at the third server;

302. In my opinion, Russell or Russell-Menon teaches Element 3[b]. *See* Section XI.B.4.b.2.

(4) Element 3[c]: data objects in the fourth plurality of data objects are at least partially determined based on when requests for those data objects were last received at the fourth server;

303. In my opinion, Russell or Russell-Menon teaches Element 3[c]. *See* Section XI.B.4.b.2.

- (5) **Element 3[d]: data objects in the fifth plurality of data objects are at least partially determined based on when requests for those data objects were last received at the fifth server;**

304. In my opinion, Russell or Russell-Menon teaches Element 3[d]. *See* Section XI.B.4.b.2.

- (6) **Element 3[e]: data objects in the sixth plurality of data objects are at least partially determined based on when requests for those data objects were last received at the sixth server; and**

305. In my opinion, Russell or Russell-Menon teaches Element 3[e]. *See* Section XI.B.4.b.2.

- (7) **Element 3[f]: data objects in the seventh plurality of data objects are at least partially determined based on when requests for those data objects were last received at the seventh server.**

306. In my opinion, Russell or Russell-Menon teaches Element 3[f]. *See* Section XI.B.4.b.2.

- c. **Claim 4: The cloud data persistence system of claim 1, wherein the first plurality of data objects comprises all data objects in the second plurality of data objects, the third plurality of data objects, the fourth plurality of data objects, the fifth plurality of data objects, the sixth plurality of data objects, and the seventh plurality of data objects.**

307. In my opinion, Russell or Russell-Menon teaches Claim 4.

308. Ground 3: Russell teaches that the first plurality of data objects (i.e., instructions) comprises all data objects in the second, third, fourth, fifth, sixth, and

seventh plurality of data objects because the main server is the central hub that transmits the instructions that control content distribution to all of the other servers. For example, Russell teaches that “[a] copy of each content item available through the server may be stored by the main server and distributed to and stored by each parent server,” which “in turn, distribute content items to edge servers, *pursuant to instructions from the main server* or requests from the edge servers.” Ex-1006 ¶10; *see also id.*, ¶¶66, 73, 81 (“the main server 12 publishes instructions to the parent servers 14 and the edge servers 16 and the parent servers republish instructions to edge servers”). Russell also teaches that main server 12 “may store all of the content accessible through the main website” (Ex-1006 ¶55) and “each edge server 16 stores less than all of the content items stored by the parent servers.” Ex-1006 ¶78.

309. Ground 4: It would have been obvious to implement Russell’s instructions to include metadata, if required, as taught by Menon. *See* Section XI.B.2.b.

310. As I also explained in the Technological Background section of my Declaration (*see* Section VII.F), it was a well-known concept by 2010 that data objects at the origin server comprise all of the data objects in the cache and application servers. Ex-1031 (Mao) ¶9; Ex-1029 (Seed), Fig. 1.

- d. **Claim 6: The cloud data persistence system of claim 5 wherein the first plurality of data objects comprises all data objects present in the second plurality of data objects and the third plurality of data objects.**

311. In my opinion, Russell or Russell-Menon teaches Claim 6. *See* Section XI.B.4.c.

- e. **Claim 7**

- (1) **Element 7[pre]: The cloud data persistence system of claim 5,**

312. In my opinion, Russell or Russell-Menon teaches Element 7[pre]. *See* Section XI.B.2.a.

- (2) **Element 7[a]: wherein the third server stores a fourth plurality of data objects and is configured to request the data object from the first server if the data object is not present in the fourth plurality of data objects, wherein:**


313. In my opinion, Russell or Russell-Menon teaches Element 7[a]. *See* Section XI.B.2.i.

- (3) **Element 7[b]: the third server and the second server are data persistence cache servers; and**

314. In my opinion, Russell or Russell-Menon teaches Element 7[b]. *See* Sections XI.B.2.g and XI.B.2.j.

- (4) **Element 7[c]: the third server is communicatively coupled with a sixth and seventh server.**

315. In my opinion, Russell or Russell-Menon teaches Element 7[c]. *See* Sections XI.B.2.k and XI.B.2.r-XI.B.2.w. Ex-1006, Figs. 1-2 (brown dashed boxes added).






FIG. 2

f. Claim 8

- (1) Element 8[pre]: The cloud data persistence system of claim 5,**

316. In my opinion, Russell or Russell-Menon teaches Element 8[pre]. *See* Section XI.B.2.a.

- (2) Element 8[a] wherein the fifth server stores a fifth plurality of data objects and is configured to request the data object from the second server if the data object is not present in the fifth plurality of data objects, wherein:**

317. In my opinion, Russell or Russell-Menon teaches Element 8[a]. *See* Section XI.B.2.o.

- (3) Element 8[b]: the fifth server and the fourth servers are application servers; and**

318. In my opinion, Russell or Russell-Menon teaches Element 8[b]. *See* Sections XI.B.2.m and XI.B.2.p.

- (4) **Element 8[c]: the fifth server is configured to receive requests from a second client for the data object; and**


319. In my opinion, Russell or Russell-Menon teaches Element 8[c]. *See* Sections XI.B.2.q.

- (5) **Element 8[d]: the fifth server is geographically separated from the fourth server.**

320. In my opinion, Russell or Russell-Menon teaches Element 8[d].

321. Ground 3: Russell teaches that “one parent cache server may be located near the west coast, another parent cache server may be located near the east coast and a further parent cache server may be located in the central portion of the country.” Ex-1006 ¶64. Russell further teaches that “*edge servers 16* provide distributed storage which, preferably, is comparatively local to users” and “*may be physically spread throughout the service region*” with “[a]bout twenty to twenty-five (for example, twenty-three) edge servers 16 ... located at selected locations across the region.” Ex-1006 ¶67. Moreover, Russell teaches that “[e]dge servers 16 may be *strategically located near areas in which a high volume of usage is expected*, such as on or near a college campus or associated with a college server.” Ex-1006 ¶67; *see also id.*, ¶10 (“Edge servers may be located in populated areas, such as in or near large cities.”). As shown in annotated Figure 2 below, if the second server is located on the west coast, the fourth and fifth servers could be

geographically separated in large cities in one or multiple states in the region. Ex-1006, Fig. 2.



322. Ground 4: Menon also teaches this claim for the reasons described in Section XI.A.4.f. In addition to the reasons described in Section XI.B.1, a POSITA would have been motivated to implement Menon’s teachings in Russell because “[s]erving content from a cache closer to the end user [] improves performance.” Ex-1007 ¶18. A POSITA would understand that serving content from a cache to many nearby clients also reduces congestion in the core of the network as cached content need not be delivered numerous times across the entire network from the origin content server. A POSITA would have also understood that using a physically closer cache to retrieve content reduces the round trip time thus allowing TCP to run faster (as TCP can deliver one “window” of data each round trip time).

323. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the application servers would be geographically separated from one another. Ex-1029 (Seed), 3:36-37, Fig. 1; Ex-1031 (Mao) ¶50.

g. Claim 9: The cloud data persistence system of claim 5, wherein the request from the first client for the data object comprises a universal resource indicator.

324. In my opinion, Russell or Russell-Menon teaches Claim 9.

325. Ground 3: Russell teaches that the request from the first client for the data object comprises a universal resource indicator (“URL”). Specifically, Russell teaches that “[w]hen a user purchases a license for a selected content item the main server provides the user’s UND with *a URL (uniform resource locator) that will allow the user to connect to an edge server to download* an encrypted electronic file containing *the selected content item.*” Ex-1006 ¶12. The user inputs “the URL to the user’s web browser immediately or wait[s] and download[s] the content item at a later time.” Ex-1006 ¶12; *see also id.*, ¶88, Cl. 8. Russell further teaches that “when the user attempts to access the content item, *the URL provided to the user will be resolved by a DNS (domain name server) to connect the user to the appropriate edge server 16* in the system network for the download request.” Ex-1006 ¶89. “Alternatively, *the URL may include address information for a specific edge server 16* to which the user’s request is to be sent.” Ex-1006 ¶89. A POSITA would

understand that a URL is a specific form of a universal resource indicator. Ex-1030 (IETF RFC 3986), 7 (“A URI can be further classified as a locator, a name, or both. The term ‘Uniform Resource Locator’ (URL) refers to the subset of URIs that, in addition to identifying a resource, provide a means of locating the resource by describing its primary access mechanism (e.g., its network ‘location’).”).

326. Ground 4: Menon also teaches this claim for the reasons described in Sections XI.B.2.b and XI.A.4.g.

327. As I also explained in the Technological Background section of my Declaration (*see* Section VII.D), it was known that the requests for the data objects would comprise a URL. Ex-1024 (Coates), 22:24-27, Fig. 20; Ex-1032 (Davis), 7:10-13.

h. Claim 10: The cloud data persistence system of claim 5 wherein each of the first plurality of data objects comprise a universal resource locator and a title linked to a media asset of the plurality of media assets.

328. In my opinion, Russell or Russell-Menon teaches Claim 10.

329. Ground 3: Russell teaches that the instructions stored at the first server comprise a URL linked to the media asset for the reasons described in the previous section and Sections XI.B.2.b and XI.B.2.d. Moreover, Russell teaches that “the user may browse through *website text and images describing content items*, such as movies, available for downloading.” Ex-1006 ¶35. And “[t]he website may provide the user with *search and other functions to assist the user in locating* and selecting

content items.” Ex-1006 ¶35; *see also id.*, ¶87 (“users may search and select content items”). Russell further teaches that after the user selects a content item and pays a fee “the user may obtain a license and receive a” URL “from the main website, where the URL provides the user’s UND with an address or the like, to direct the user to an edge server 16 from which the user may download the selected movie.” Ex-1006 ¶36; *see also id.*, ¶¶88-89. The edge server then sends a query to the main server to authenticate the request, which “may include *information identifying the requested content item (for example, movie)*, the requesting user, and the user’s current IP address.” Ex-1006 ¶91. A POSITA would have understood that information identifying the requested content would include, for example, a title of the movie. Indeed, it would have been obvious to a POSITA that the instructions would include some sort of title or identifier for the content because the most obvious way to search for a movie is by its “title.” This is also taught by Menon. Furthermore, a POSITA would have understood that each movie would also have a unique identifiers, so that even movies with identical titles can be properly distinguished from each other and so that the media of those movies do not become mixed up.

330. Ground 4: Menon teaches this claim for the reasons described in Sections XI.B.2.b and XI.A.4.h.

331. As I also explained in the Technological Background section of my Declaration (*see* Section VII.B), it was known that each of the data objects would

comprise a URL and a title linked to the media asset. Ex-1026 (Geilfuss) ¶36, Fig. 2; Ex-1021 (Eaton), 8:26–36; Ex-1031 (Mao) ¶51; Ex-1033 (Leighton), 1:54-2:6.

- i. **Claim 11: The cloud data persistence system of claim 5, wherein data objects in the second plurality of data objects and the third plurality of data objects are at least partially determined based on a frequency of requests for those data objects received from clients.**

332. In my opinion, Russell or Russell-Menon teaches Claim 11. *See* Section XI.B.4.b.

- j. **Claim 12: The cloud data persistence system of claim 5, wherein data objects in the second plurality of data objects and the third plurality of data objects are at least partially determined based on when requests for those data objects were last received from clients.**

333. In my opinion, Russell or Russell-Menon teaches Claim 12. *See* Section XI.B.4.b.

- k. **Claim 13: The cloud data persistence system of claim 5, wherein the first plurality of data objects comprises all data objects in the second plurality of data objects, the third plurality of data objects, the fourth plurality of data objects, and the fifth plurality of data objects.**


334. In my opinion, Russell or Russell-Menon teaches Claim 13. *See* Sections XI.B.4.c and XI.B.4.d.

XII. CONCLUSION

335. For the reasons above, it is my opinion that Claims 1-13 of the '987 Patent are obvious based on each of the grounds specified in my Declaration.

336. I declare that all statements made in my Declaration of my knowledge are true, and that all statements made on information and belief are believed to be true, and that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Date: October 24, 2025



Dr. Henry Houh