

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC,
Petitioner,

v.

SANDPIPER CDN, LLC.,
Patent Owner.

Case IPR2025-00969
Patent No. 8,478,903

**DECLARATION OF TODD C. MOWRY, PH.D.
IN SUPPORT OF PETITION FOR INTER PARTES REVIEW OF
U.S. PATENT NO. 8,478,903**

TABLE OF CONTENTS

I. INTRODUCTION11

II. SUMMARY OF OPINIONS.....12

III. QUALIFICATIONS AND BACKGROUND12

IV. MATERIALS CONSIDERED17

V. LEGAL STANDARDS18

A. Claim Construction19

B. Obviousness Under 35 U.S.C. § 103.....20

VI. THE '903 PATENT.....23

A. Overview23

B. Summary of the Prosecution History25

C. Claim Construction of Terms in the '903 patent.....26

VII. LEVEL OF ORDINARY SKILL IN THE ART26

VIII. OVERVIEW OF THE PRIOR ART27

A. *Kenner*27

1. *Kenner's* system.....27

2. *Kenner* accommodates multiple users and content providers.....29

B. *Vetter*34

C. *Rekimoto*.....35

D. *Boyles*35

IX. CLAIMS 1, 2, 22-24, AND 26-46 ARE UNPATENTABLE.....36

A. Ground 1: *Kenner* Renders Obvious Claims 1, 2, 22-24, 26, 28-32, 35, 37, 38, 40, 41, and 43-45.....36

1. Independent Claim 136

a. [1.pre]: “A content delivery system operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines”36

b. [1.a.i]: “at least one shared repeater server constructed and adapted to replicate at least some of the resources associated with a first content provider of said plurality of content providers”37

c. [1.a.ii]: “at least one shared repeater server constructed and adapted ... to replicate at least some of the resources associated with a second content provider of said plurality of content providers, wherein the second content provider is distinct from the first content provider”39

d. [1.b]: “at least one table listing a plurality of alias names corresponding to content providers authorized to have resources delivered to client machines via the at least one shared repeater server, wherein the plurality of content providers comprises the first content provider and the second content provider”42

e. [1.c]: “wherein at least a first resource associated with the first content provider is associated with a first alias name of the plurality of alias names, said first alias name being associated with said at least one shared repeater server”44

f. [1.d]: “wherein, requests for the first resource from said client machines are directed to the at least one shared repeater server, based at least in part on said first alias name”46

g.	[1.e]: “wherein at least a second resource associated with the second content provider is associated with a second alias name of the plurality of alias names, said second alias name being associated with said at least one shared repeater server”	47
h.	[1.f]: “wherein requests for the second resource made from said client machines are directed to the at least one shared repeater server, based at least in part on said second alias name”	48
i.	[1.g]: “wherein the at least one shared repeater server is further constructed and adapted to analyze, using the table, an alias name received with a client request for a particular resource to determine a content provider associated with the particular resource”	49
2.	Claim 2: “wherein the least one shared repeater server is addressable using the first alias name and the second alias name”	50
3.	Claim 22: “wherein at least one of the first resource and the second resource is an embedded object”	51
4.	Claim 23: “wherein the requests comprise host tag names that identify specific origin servers from which the resources originate”	51
5.	Claim 24: “wherein the host tag names are included in HTTP headers”	52
6.	Claim 26: “wherein the table further comprises, for each of the plurality of content providers, information from the group consisting of: a reflector name and a committed aggregate information rate (CAIR)”	52
7.	Independent Claim 28	55

a. [28.pre.i]: “A method, in a content delivery system operative in a computer network for delivering content to client machines”.....55

b. [28.pre.ii]: “the computer network comprising a plurality of origin servers, each of said origin servers having resources associated therewith”56

c. [28.pre.iii]: “the content delivery system comprising at least one shared repeater server operable to replicate resources associated with the plurality of origin servers”.....56

d. [28.a]: “associating the at least one repeater server with a first alias name, wherein requests for a first resource located on a first origin server are directed, based at least in part on said first alias name, to the at least one repeater server for delivery of the first resource from said at least one repeater server”57

e. [28.b]: “associating the at least one repeater server with a second alias name, wherein requests for a second resource located on a second origin server are directed, based at least in part on said second alias name, to the at least one repeater server for delivery of the second resource from said at least one repeater server, wherein the second origin server is distinct from the first origin server”.....58

f. [28.c]: “providing a table listing origin servers having content located thereon, wherein said content is authorized for delivery to client machines via the at least one shared repeater server, and wherein the origin servers comprise the first origin server and the second origin server”59

g. [28.d]: “wherein the at least one repeater server is further constructed and adapted to analyze, using the table, an alias name received with a client

	request for a particular resource to determine an origin server associated with the particular resource”	60
8.	Claim 29: “further comprising: in response to a request for the first resource issued by a particular client machine, identifying the at least one shared repeater server from a set of shared repeater servers based at least in part on a network cost of delivering resources to the particular client machine; and associating the first resource with the first alias name”	60
9.	Claim 30: “further comprising: in response to a request for the first resource issued by a particular client machine, identifying the at least one repeater server from a set of shared repeater servers based at least in part on load characteristics of the set of shared repeater servers; and associating the first resource with the first alias name”	62
10.	Claim 31: “further comprising: in response to a request for the first resource issued by a particular client machine, identifying the at least one repeater server from a set of shared repeater servers based at least in part on a location of the particular client machine relative to the at least one shared repeater server; and associating the first resource with the first alias name”	63
11.	Claim 32: “wherein the resources comprise an HTTP-based resource”	63
12.	Claim 35: “wherein the table further comprises, for each of the listed origin servers, information from the group consisting of: a reflector name and a committed aggregate information rate (CAIR)”	64
13.	Independent Claim 37	64

- a. [37.pre.i]: “A method, in a content delivery system operative in a computer network for delivering content to client machines”64
- b. [37.pre.ii]: “comprising at least one shared repeater server operable to replicate resources stored on a plurality of origin servers”64
- c. [37.a.i]: “associating a plurality of alias names with the at least one shared repeater server, each of said plurality of alias names being associated with an origin server, wherein a first alias name of said plurality of alias names is associated with a first origin server”64
- d. [37.a.ii]: “at least a second alias name of said plurality of alias names is associated with a second origin server distinct from said first origin server”65
- e. [37.b]: “providing at least one table associating alias names with origin servers having content located thereon, wherein said content is authorized for delivery to client machines via the at least one shared repeater server, wherein the origin servers comprise the first origin server and the second origin server”65
- f. [37.c]: “wherein requests for a first resource originating on the first origin server are directed, based at least in part on said first alias name, to the at least one shared repeater server for delivery of the first resource from said at least one repeater server”66
- g. [37.d]: “wherein requests for a second resource originating on the second origin server are directed, based at least in part on said second alias name, to the at least one shared repeater server for delivery of the second resource from said at least one repeater server”66

h. [37.e]: “wherein the at least one repeater server uses the at least one table and an alias name received with a client request for a resource to determine an origin server associated with the requested resource”67

14. Claim 38: “wherein a copy of the at least one table is located on each of a plurality of repeater servers”67

15. Independent Claim 4068

a. [40.pre.i]: “A server operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines”68

b. [40.pre.ii]: “wherein the plurality of content providers comprise at least a first content provider and a second content provider distinct from the first content provider”68

c. [40.a]: “the server comprising: ... cache storage”68

d. [40.b]: “a table listing content providers of said plurality of content providers having content that is authorized for delivery to client machines via the server”70

e. [40.c]: “wherein the server is associated with a first alias name corresponding to said first content provider”71

f. [40.d]: “wherein the server is associated with a second alias name corresponding to said second content provider, said second alias name being distinct from the first alias name”71

g. [40.e]: “wherein the server is operable to replicate at least some of the resources associated with the first content provider, and at least some of the

	resources associated with the second content provider”	74
h.	[40.f]: “wherein, in response to a request for a resource, the server is operable to use at least the table to analyze an alias name associated with the request to determine a content provider associated with the resource”	74
16.	Claim 41: “wherein the table further comprises, for each of the listed content providers, information from the group consisting of: a reflector name and a committed aggregate information rate (CAIR)”	74
17.	Independent Claim 43	75
a.	[43.pre]: “A server operative in a computer network for delivering resources to multiple client machines”	75
b.	[43.a]: “the server comprising: ... cache storage”	75
c.	[43.b]: “a table listing a plurality of origin servers having content that is authorized for delivery to client machines via the server”	75
d.	[43.c]: “wherein the server is associated with a first alias name corresponding to a first origin server of the plurality of origin servers”	75
e.	[43.d]: “wherein the server is associated with a second alias name corresponding to a second origin server of the plurality of origin servers, said second origin server being distinct from the first origin server and said second alias name being distinct from the first alias name”	75
f.	[43.e]: “wherein the server is operable to replicate at least some of the resources located on the first origin server, and at least some of the resources located on the second origin server”	76

- g. [43.f]: “wherein, in response to a request for a resource, the server uses at least the table to analyze an alias name associated with the request to determine an origin server associated with the particular resource”76
 - 18. Claim 44: “wherein the first origin server and the second origin server are associated with distinct content providers located at distinct physical locations”76
 - 19. Claim 45: “wherein the table further comprises, for each of the listed origin servers, information from the group consisting of: a reflector name and a committed aggregate information rate (CAIR)”77
- B. Ground 2: The Combination of *Kenner* and *Vetter* Renders Obvious Claim 33.....77
 - 1. Rationale for Combining *Kenner* and *Vetter*77
 - 2. Claim 33: “wherein the resources comprise a File Transfer Protocol (FTP)-based resource”78
- C. Ground 3: The Combination of *Kenner* and *Rekimoto* Renders Obvious Claim 34.....79
 - 1. Rationale for Combining *Kenner* and *Rekimoto*.....79
 - 2. Claim 3481
 - a. [34.a]: “wherein at least the first alias name is a domain name and the computer network further comprises at least one name server that provides domain name service (DNS) resolution”81
 - b. [34.b]: “the method further comprising: receiving, at the at least one name server, the first alias name as part of a client request; and resolving, at the at least one name server, the received first alias name to identify the at least one shared repeater server to

	deliver the first resource in response to the client request”	83
D.	Ground 4: The Combination of <i>Kenner</i> and <i>Boyles</i> Renders Obvious Claims 27, 36, 39, 42, and 46	85
1.	Rationale for Combining <i>Kenner</i> and <i>Boyles</i>	85
2.	Claim 27: “wherein the at least one repeater server is configured and adapted to reject the client request for the particular resource if the alias name received with the client request is not included the table”	86
3.	Claim 36: “further comprising: rejecting the client request for the particular resource if the alias name received with the client request is not included [in] the table”	87
4.	Claim 39: “further comprising: rejecting the client request for the resource if the alias name received with the client request is not included [in] the table”	88
5.	Claim 42: “wherein in response to a request for a resource, the server is further configured and adapted to reject the request for the resource if the alias name received with the request is not included [in] the table”	88
6.	Claim 46: “wherein in response to a request for a resource, the server is further configured and adapted to reject the request for the resource if the alias name received with the request is not included [in] the table”	88
X.	CONCLUSION.....	89

I. INTRODUCTION

1. I have been retained by Google LLC (“Google”), as an independent expert for the above-identified *inter partes* review proceeding involving U.S. Patent No. 8,478,903 (“the ’903 patent”).

2. I am over 21 years of age and, if I am called upon to do so, I would be competent to testify as to the matters set forth herein.

3. Although I am being compensated for my time in connection with this proceeding at the rate of \$800 per hour for time that I spend on this case and reimbursed for reasonable out-of-pocket expenses, no part of my compensation depends on the outcome of this proceeding, and I have no other interest in this proceeding.

4. I have been asked to consider the validity of certain claims of the ’903 patent based on certain prior art references. I have also been asked to consider the state of the art and prior art available as of February 10, 1998, which is the earliest claimed priority date on the face of the ’903 patent. Based on the combinations of prior art references discussed in this declaration, it is my opinion that claims 1, 2, 22-24, and 26-46 of the ’903 patent are unpatentable.

5. Throughout this declaration, I refer to specific pages, figures, and/or line numbers of various exhibits. These citations are illustrative and are not intended to suggest that they are the only support for the propositions for which they are cited.

II. SUMMARY OF OPINIONS

6. This declaration analyzes claims 1, 2, 22-24, and 26-46 of the '903 patent. Below, I set forth the opinions I have formed, the conclusions I have reached, and the bases for my opinions and conclusions. I believe the statements contained in this declaration to be true and correct to the best of my knowledge.

7. Based on my experience and my knowledge of the art at the time of the applicable priority date, it is my opinion that claims 1, 2, 22-24, and 26-46 of the '903 patent would have been obvious based on the asserted grounds discussed below.

III. QUALIFICATIONS AND BACKGROUND

8. Below is a summary of my education, work experience, and other qualifications. My qualifications are stated more fully in my curriculum vitae, which I understand is provided in EX-1004.

9. I am a Professor in the Department of Computer Science at Carnegie Mellon University. I also have a courtesy appointment in the Department of Electrical and Computer Engineering. I have served on the faculty of Carnegie Mellon University for 28 years, starting in 1997 through the present (2025).

10. I served on the faculty of the University of Toronto for four years between 1993 and 1997, in the Department of Electrical and Computer Engineering and by courtesy appointment in the Department of Computer Science. Prior to that

appointment, I served as a Graduate Research Assistant in the Department of Electrical Engineering at Stanford University for four years between 1989 and 1993.

11. As a faculty member, I have taught and continue to teach courses and direct research in computer systems and software, operating systems, distributed and network systems, parallel processing, and compiler optimization.

12. I received a B.S. degree in Electrical Engineering with Highest Distinction from the University of Virginia in May 1988. I received an M.S. in Electrical Engineering from Stanford University in June 1989, and a Ph.D. in Electrical Engineering from Stanford University in March 1994.

13. I have worked in the computer industry in various capacities. I was a part-time Computer Architect and then Computer Architecture Consultant at Silicon Graphics, Inc., in Mountain View, California (formerly MIPS Computer Systems in Sunnyvale, California), from 1989 to 1993 and 1993 to 1996, respectively. I was a Visiting Scientist at IBM in Toronto from 1996 to 2004. During that same time period (1996 to 2004), I was also a Member of the Technical Advisory Board of SandCraft, Inc. in Santa Clara, California. I was the Director of the Intel Research Pittsburgh Lab at Intel Corporation in Pittsburgh, Pennsylvania from 2004 to 2007.

14. I have authored 21 journal articles and 71 conference papers. I am also an inventor on five patents.

15. I have published a number of papers in the top research conferences in the fields of data storage, caching, operating systems and distributed systems (including a paper that won the Best Paper Award at the USENIX 2nd Symposium on Operating Systems Design and Implementation (OSDI '96)). *See, e.g.*, Charles Garrod, Amit Manjhi, Anastasia Ailamaki, Bruce Maggs, Todd Mowry, Christopher Olston, and Anthony Tomasic, *Scalable Query Result Caching for Web Applications*, in Proceedings of the 34th International Conference on Very Large Databases (VLDB) (2008); Andrew Pavlo, Gustavo Angulo, Joy Arulraj, Haibin Lin, Jiexi Lin, Lin Ma, Prashanth Menon, Todd Mowry, Matthew Perron, Ian Quah, Siddharth Santurkar, Anthony Tomasic, Skye Toor, Dana Van Aken, Ziqi Wang, Yingjun Wu, Ran Xian, and Tieying Zhang, *Self-Driving Database Management Systems*, in Proceedings of the 2017 Conference on Innovative Data Systems Research (CIDR) (2017); Amit Manjhi, Charles Garrod, Bruce M. Maggs, Todd C. Mowry, Anthony Tomasic, *Holistic Query Transformations for Dynamic Web Applications*, in 2009 IEEE 25th International Conference on Data Engineering (ICDE) (2009); and Amit Manjhi, Phillip B. Gibbons, Anastassia Ailamaki, Charles Garrod, Bruce M. Maggs, Todd C. Mowry, Christopher Olston, Anthony Tomasic, and Haifeng Yu, *Invalidation Clues for Database Scalability Services*, in

Proceedings of the 2007 IEEE 23rd International Conference on Data Engineering (ICDE) (2007).

16. I have also been an active member of the Parallel Data Lab at Carnegie Mellon University since 1997, which is often described as “academia’s premiere data systems research center.”

17. I am the recipient of several honors and awards:

- Arthur Samuel Thesis Award (awarded by the Stanford Computer Science department to the top two Ph.D. theses in a given year);
- IBM Faculty Development Awards (1996, 1997, 1998, 2000, 2001, 2002, and 2003);
- Best Paper Awards (Second Symposium on Operating Systems Design and Implementation in 1996; 20th International Conference on Data Engineering (ICDE) in 2004);
- Alfred P. Sloan Research Fellowship (awarded to researchers in recognition of distinguished performance and a unique potential to make substantial contributions to their field);
- Most Thought-Provoking Idea Award in 2004 (awarded by the Architectural Support for Programming Languages and Operating Systems (ASPLOS));
- TR100 Award in 1999 (awarded by MIT’s Technology Review magazine to the top 100 most promising young innovators in science and technology); and
- I became an ACM Fellow in 2016.

18. I am a member of the Institute of Electrical and Electronics Engineers (IEEE) and the Association of Computing Machinery (ACM). I was the Editor-in-Chief of ACM Transactions on Computer Systems (the premier journal for computer systems research) from 2013 to 2018. I was an Associate Editor for the journal prior to that, since 2001. I was the Program Chair of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS) in 2010. I was the Co-Program Chair of the International Conference on Parallel Architectures and Compilation Techniques (PACT) in 2001. I have been on the programming committee in various years for ASPLOS, the International Symposium on Computer Architecture (ISCA), the International Symposium on Microarchitectures, and the Workshop on Architectural and System Support for Improving Format.

19. Overall, I have over 30 years of experience in the field of computer science and, specifically, computer architecture, operating systems, and document storage and retrieval.

20. In light of the foregoing, I consider myself to be an expert in the field of computer science, and I believe that I am qualified to provide an opinion as to what a person of ordinary skill in the art would have understood, known, or

concluded regarding the subject matter of the '903 patent at the time of its alleged invention.

IV. MATERIALS CONSIDERED

21. In forming my opinions, I read and considered the '903 patent and its prosecution history, the exhibits listed in the Exhibit List filed with the petition for *inter partes* review of the '903 patent, as well as any other materials referenced herein. In preparing this declaration, I have reviewed at least the following documents:

<u>Exhibit</u>	<u>Description</u>
1001	U.S. Patent No. 8,478,903 to Farber et al. ("903 Patent") ¹
1002	Prosecution History of U.S. Application No. 11/065,412
1005	PCT International Patent Pub. No. WO 1996041285 to Kenner et al. (" <i>Kenner</i> ")
1006	Ronald J. Vetter et. al, Mosaic and the World-Wide Web, Computer, vol. 27, no. 10, pp. 49-57, Oct. 1994 (" <i>Vetter</i> ")
1007	European Patent Pub. No. EP 0753836 to Rekimoto et al. (" <i>Rekimoto</i> ")
1008	U.S. Patent No. 5,511,208 to Boyles et al. (" <i>Boyles</i> ")
1009	Andrew Tanenbaum, Computer Networks (3rd ed. 1996) (" <i>Tanenbaum</i> ")

¹ I understand that, because pages 31-54 of *Kenner* do not include line numbers or paragraph numbers, Petitioner has added line numbering for these pages to facilitate review.

<u>Exhibit</u>	<u>Description</u>
1010	Tim Berners-Lee, The World-Wide Web, Communications of the ACM, Vol. 37, No. 8, August 1994 (“ <i>Berners-Lee</i> ”)
1011	Radhika Malpani, Making World Wide Web Caching Servers Cooperate, WWW4: Proceedings of the Fourth International Conference on World Wide Web, Pages 107-117, December 11, 1995 (“ <i>Malpani</i> ”)
1012	Fielding et al., Hypertext Transfer Protocol – HTTP/1.1, RFC 2068, Jan. 1997, https://www.rfc-editor.org/rfc/rfc2068 (“ <i>HTTP/1.1</i> ”)
1013	Barbara Tockey Zivkov et al., Disk Caching in Large Databases and Timeshared Systems, Proceedings Fifth International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 97), Haifa, Israel, 184–95, January 1997 (“ <i>Zivkov</i> ”)

22. I reserve the right to supplement and amend any of my opinions in this declaration based on documents, testimony, and other information that becomes available to me after the date of this declaration.

V. LEGAL STANDARDS

23. In considering the subject matter of the '903 patent and its claims in light of the prior art and forming my opinions, I am relying on certain legal principles that counsel in this case explained to me. My understanding of these concepts is summarized below.

24. I understand that the claims define the invention. I also understand that an unpatentability analysis is a two-step process. First, the claims of the patent are construed to determine their meaning and scope. Second, after the claims are construed, the content of the prior art is compared to the construed claims.

25. I understand that an invention, as defined by the claims of the patent, must not be anticipated or rendered obvious by the prior art.

A. Claim Construction

26. I understand that the United States Patent and Trademark Office interprets claim terms in an *inter partes* review proceeding under the same claim construction standard that is used in a United States federal court. I understand that under this standard, the meaning of claim terms is considered from the viewpoint of a person of ordinary skill in the art (“POSITA”) at the time of the alleged invention.

27. I have been informed that claim terms are generally given their ordinary and customary meaning as understood by a POSITA in light of the specification and the prosecution history pertaining to the patent. I understand, however, that claim terms are generally not limited by the embodiments described in the specification.

28. I understand that in addition to the claims, specification, and prosecution history, other evidence may be considered to ascertain the meaning of claim terms, including textbooks, encyclopedias, articles, and dictionaries. I have

been informed that this other evidence is often less significant and less reliable than the claims, specification, and prosecution history.

B. Obviousness Under 35 U.S.C. § 103

29. I understand that a patent claim is invalid as obvious if the claimed invention would have been obvious to a POSITA at the time the claimed invention was made. I understand that in an obviousness determination, the POSITA is presumed to have knowledge of all material prior art. I understand that whether a claim is obvious is based upon the determination of several factual issues.

30. In considering obviousness, I understand that one must determine the scope and content of the prior art. I understand that, in order to be considered prior art to a patent, a reference must be reasonably related to the claimed invention of that patent. A reference is reasonably related if it is in the same field as the claimed invention or is from another field to which a person of ordinary skill in the art would look to solve a known problem.

31. I understand that one must determine what differences, if any, existed between the claimed invention and the prior art. In evaluating whether such a claim would have been obvious, one may consider whether a reason has been identified that would have prompted a POSITA to combine the elements or concepts from the prior art in the same way as recited in the claimed invention.

32. I understand that one may consider whether (1) the change is merely the predictable result of using prior art elements according to their known functions or the result of true inventiveness; (2) there is some teaching or suggestion in the prior art to make the modification or combination of limitations claimed in the patent; (3) the innovation applies a known technique that had been used to improve a similar device or method in a similar way; (4) the claimed invention would have been obvious to try, meaning that the claimed innovation is one of a relatively small number of possible approaches to the problem that could have been used by those of ordinary skill in the art with a reasonable expectation of success; (5) the invention merely substitutes one known element for another known element in order to obtain predictable results; (6) the invention merely applies a known technique to a known device, method, or product to yield predictable results; or (7) known work in the field may have prompted variations of use of the same inventions in the same or different fields due to market forces or design incentives that would have been predictable to a person of ordinary skill in the art.

33. I understand that one must be careful not to determine obviousness using the benefit of hindsight. I understand that the determination should be based on the position of a POSITA at the time the claimed invention was made and should not reflect what is known today or what is learned from the teaching of the patent.

34. I understand that certain secondary considerations may be considered in evaluating obviousness in order to prevent hindsight bias. These secondary considerations include the commercial success of products that practice the patent, long-felt need for the patented technology, failure by others to solve the problem addressed by the patent, initial skepticism by others in the industry, industry recognition and praise of the patented products, and efforts by others to copy the patented technology.

35. I have been informed that these secondary considerations of non-obviousness must be balanced against the strength of the *prima facie* case of obviousness. Where the invention represents no more than the predictable use of prior art elements according to their established functions, secondary considerations are inadequate to establish non-obviousness.

36. I have also been informed that the patentee must establish a nexus between any secondary considerations and the claimed invention. Where a secondary consideration results from something other than what is both claimed and novel in the patent, there is no nexus to the merits of the claimed invention.

37. I have been informed that, for example, the nexus requirement as it applies to evidence of commercial success requires that the patentee must show that the sales of a product that allegedly practices the claimed invention must be a direct

result of the unique characteristics of the claimed invention, as opposed to other economic and commercial factors that are unrelated to the patented invention. I understand this to mean that if the product's commercial success is due to an element that existed in the prior art or that is not claimed by the patent, then no nexus exists.

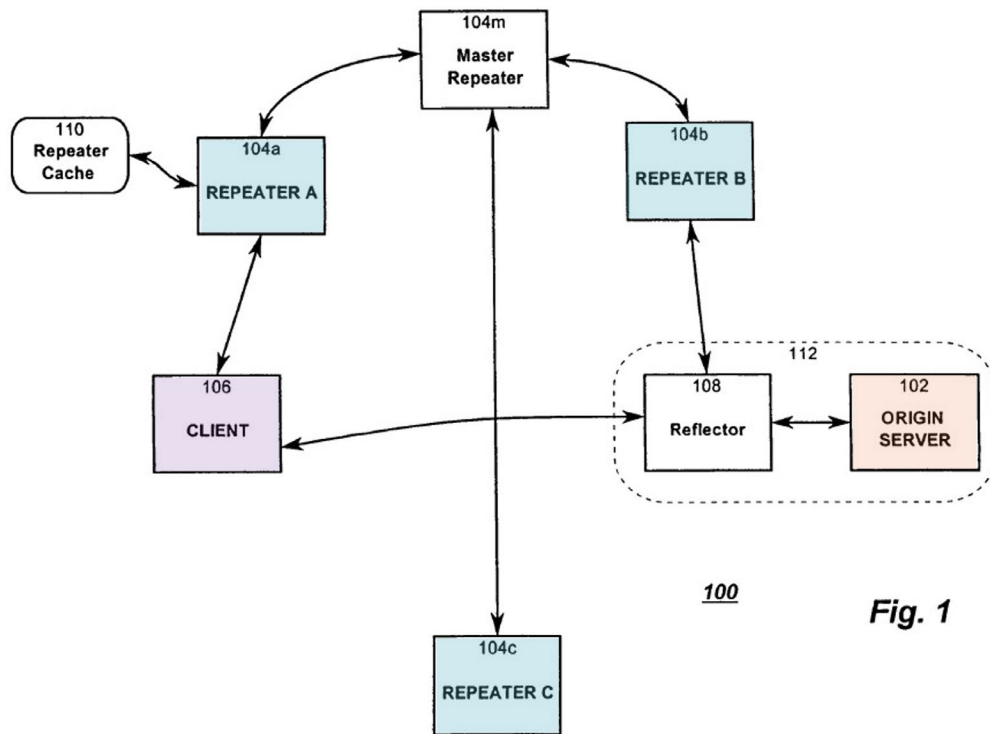
38. I have been further informed that, for example, the nexus requirement as it applies to the long-felt need element must likewise be supported by evidence. Where the differences between the prior art and the claimed invention are minimal, it cannot be said that any long-felt need was met by the patented invention or that the patent solved any technological problems that were unaddressed by others.

39. I have been informed that evidence of industry praise must also show a nexus with the claimed invention, and that self-serving statements by the patentee regarding the merits of the claimed invention are generally not given weight.

VI. THE '903 PATENT

A. Overview

40. The '903 patent discloses a system enabling servers to “off-load their processing of requests for selected resources by determining a different server (a ‘repeater’) to process those requests.” EX-1001, 2:62-65. As shown in Fig. 1 (annotated below), the '903 patent's network includes an **origin server**, a **client**, and various **shared repeater servers**:



Id., Fig. 1, 4:22-36. Network resources are originally uploaded at the origin server. *Id.*, 4:48-55. Each repeater server (104a, 104b, and 104c) replicates some or all of the origin server’s resources. *Id.*, 4:22-36. A user submits a request to access a resource through client 106, which is received at a reflector 108. *Id.*, 4:56-60. The reflector selects either a repeater server or the origin server to serve the request. *Id.*, 5:9-23, 7:50-56.

41. Each repeater server stores a table correlating alias names with names of origin servers. *Id.*, 9:56-63. When responding to a request for a resource, a repeater server uses the table to identify the origin server associated with the requested resource using an alias name included with the request. *Id.*, claims 1, 28,

37, 40, 43. If the repeater server stores a copy of the resource locally, the repeater server retrieves the resource. *Id.*, 10:35-37. Otherwise, the repeater “quer[ies] its ‘peer caches’ to see if one of them contains the resource,” in which case the resource is delivered from the peer cache. *Id.*, 10:48-53.

42. The purported advantages of the ’903 patent arise from the reflector and repeater servers. *Id.*, 7:40-49. The ’903 patent acknowledges that delivering content from origin servers to client terminals was known. *Id.*, 4:56-5:2, 7:4-38. But cache or repeater servers were also known and used before the effective date of the ’903 patent, as outlined below.

B. Summary of the Prosecution History

43. From the cover of the patent, I understand that the ’903 patent was filed on February 23, 2005, claiming priority to February 10, 1998. EX-1001, cover. I have reviewed the prosecution history of the ’903 patent and I understand the patent was initially rejected over references not at issue here. EX-1002, 536-47. In response, the Applicant amended the claims, adding new claims 53 and 54, which recite, in part, analyzing an alias name using a table to determine an origin server or content provider associated with a request. *See id.*, 512-13; 326-27. A final rejection was issued, indicating claims 53 and 54 included allowable subject matter. *Id.*, 241-50. The Applicant then amended the independent claims to include the alias

name and table of dependent claims 53 and 54 to reach an allowance. *Id.*, 209-26; 75-95. The table and alias names were thus the basis for the allowance. *Id.*, 15.

C. Claim Construction of Terms in the '903 patent

44. As discussed above, I have been informed that, for purposes of *inter partes* review, the standard for construction of terms in a patent's claims is the same as that applied in federal district court litigation.

45. I understand that a district court has yet to construe the claim terms of this patent. At this time, I see no need to construe any terms of the challenged claims for the purposes of this proceeding. Accordingly, in preparing my analysis I have applied the ordinary and customary meaning that each claim term would have had to POSITA in light of the '903 patent's specification and prosecution history.

46. If any claim terms are differently construed, I reserve the right to reassess my opinions under those constructions.

VII. LEVEL OF ORDINARY SKILL IN THE ART

47. As discussed above, I have been informed that patentability must be analyzed from the perspective of a "person of ordinary skill in the art" in the same field as the patent-in-suit at the time of the invention. I also understand that several factors are considered in assessing the level of ordinary skill in the art, including (1) the types of problems encountered in the art; (2) the prior art solutions to those

problems; (3) the rapidity with which innovations are made; (4) the sophistication of the technology; and (5) the education level of active workers in the field.

48. I refer to a person of ordinary skill in the art here as a “POSITA.” A POSITA at the time of the invention would have had at least a bachelor’s degree in computer science, electrical engineering, or a related field, and at least two years of work or research experience in the field of content delivery management or networks. Work experience can substitute for formal education and additional formal education can substitute for work experience. I at least meet this definition of a POSITA, now and as of the priority date of the ’903 patent.

49. I am also familiar with the knowledge of a POSITA as of the priority date of the ’903 patent. I am able to opine on how a POSITA would have understood the disclosure and claims of the ’903 patent, the disclosures of the prior art, the various motivations to combine the prior art, and what combinations would have been obvious to a POSITA.

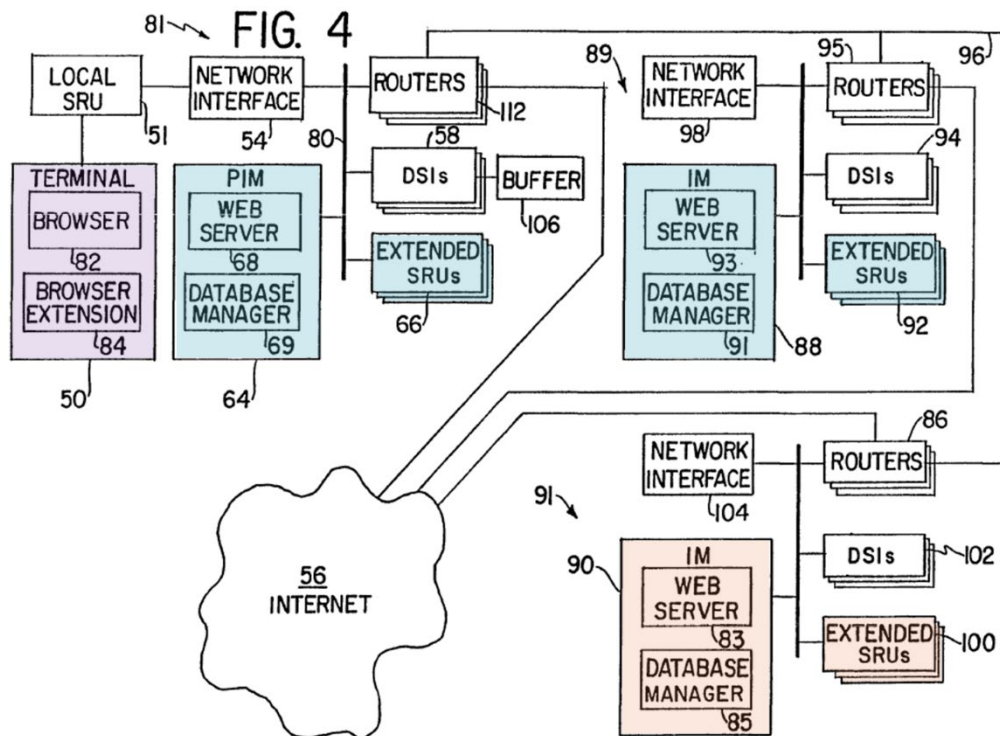
VIII. OVERVIEW OF THE PRIOR ART

A. *Kenner*

1. *Kenner’s* system

50. Like the ’903 patent, *Kenner* deals with content delivery networks. *Kenner*, Abstract (describing a “distributed digital video clip delivery system”). *Kenner* discloses a “storage and retrieval system whereby video clips, stored locally

and/or at a more remote location, can be requested and retrieved by a user at the user's multimedia terminal." *Id.* Kenner's system includes a network of database index managers (IMs), including a **source IM**, a **user terminal**, and various **regional IMs**:



Kenner, Fig. 4 (annotated). As shown in Fig. 4, each IM hosts a server and has associated extended storage and retrieval units (SRUs) for storing resources. *Id.*, 10:18-31, 31:19-21, 31:30-32, 32:14-17, 44:6-14. Though *Kenner* shows the extended SRUs separate from the IM, the extended SRU's can be "located on the same computer" as the IM. *Id.*, 27:18-20, 54:7-12. Each IM maintains a "clip database" indicating which clips it stores in its associated extended SRUs. *Id.*,

34:3-6. The IMs are geographically dispersed such that each IM may be associated with a different region. *Id.*, 31:30-32:6, 40:24-28, 54:12-16.

51. In *Kenner*, a video clip is first uploaded by a content provider using one of the IMs. *Id.*, 44:1-5. For a particular video clip, the originating server is referred to as the “source IM 90.” *Id.*, 32:14-21. The source IM stores the uploaded file on one of its extended SRUs and then distributes copies of the clip to other IMs in the system. *Id.*, 44:6-14.

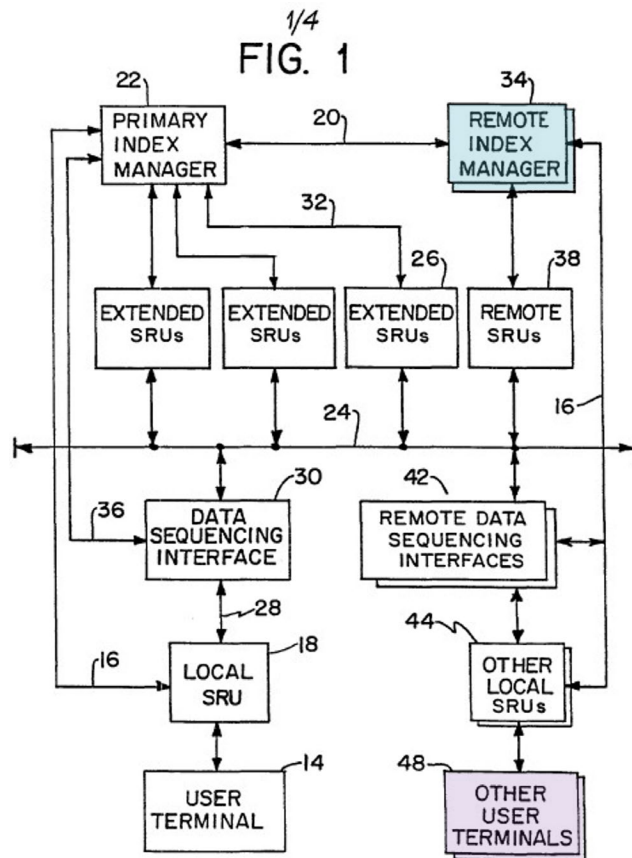
52. To access a clip, the user generates a request through the user terminal (*id.*, 6:20-23, 35:11-16, 36:6-9), which is transmitted to a primary IM (“PIM”) associated with the user terminal (*id.*, 6:32-7:6, 25:7-11, 37:7-17). The primary IM consults its clip database to determine whether its SRUs store a copy of the clip using a unique “video ID” for the requested clip included with the request. *Id.*, 37:20-22. If the clip is not stored locally, the primary IM queries other remote IMs to locate a copy of the clip. *Id.*, 40:29-41:6. If the primary IM is still unable to locate the clip, it consults its clip database again to contact the source IM. *Id.*, 41:13-19.

2. *Kenner* accommodates multiple users and content providers

53. *Kenner*’s system is scalable and thus accommodates many user terminals and content providers. *Kenner*, 54:12-16. For example, *Kenner* specifically distinguishes its system from other non-scalable systems. *Id.*, 4:32-5:5. Accordingly, the system may include more IMs than those shown in Fig. 4. For

example, *Kenner* explains its system accommodates “multiple index managers for a large number of concurrent users in disparate geographical areas” (*id.*, 40:24-28) and supports “thousands of simultaneous users” (*id.*, 5:21-26).

54. To support multiple users, each IM includes its own associated user terminal. As shown in Fig. 1, for example, *Kenner* includes **remote IMs** with associated **user terminals 48**:

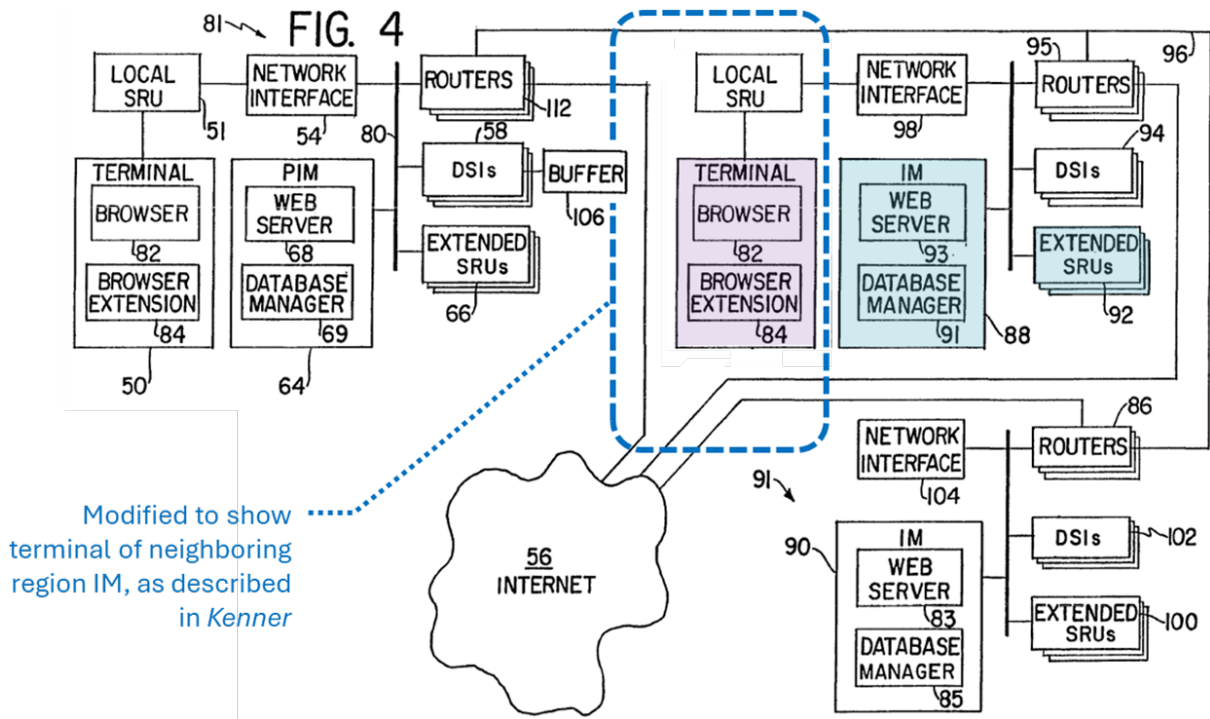


Id., Fig. 1, 28:3-6. Each of these remote IMs has its own associated user terminal.

Id., 32:4-6, claims 45, 46. Accordingly, any of the remote IMs may be a “primary” IM depending on where a user request is received. *Id.*, 32:4-6. In other

words, there is no functional difference between regional IMs. When *Kenner* refers to a primary IM, this is in the context of a particular user terminal. A neighboring region’s IM would likewise be a “primary” IM for requests received from clients in that region.

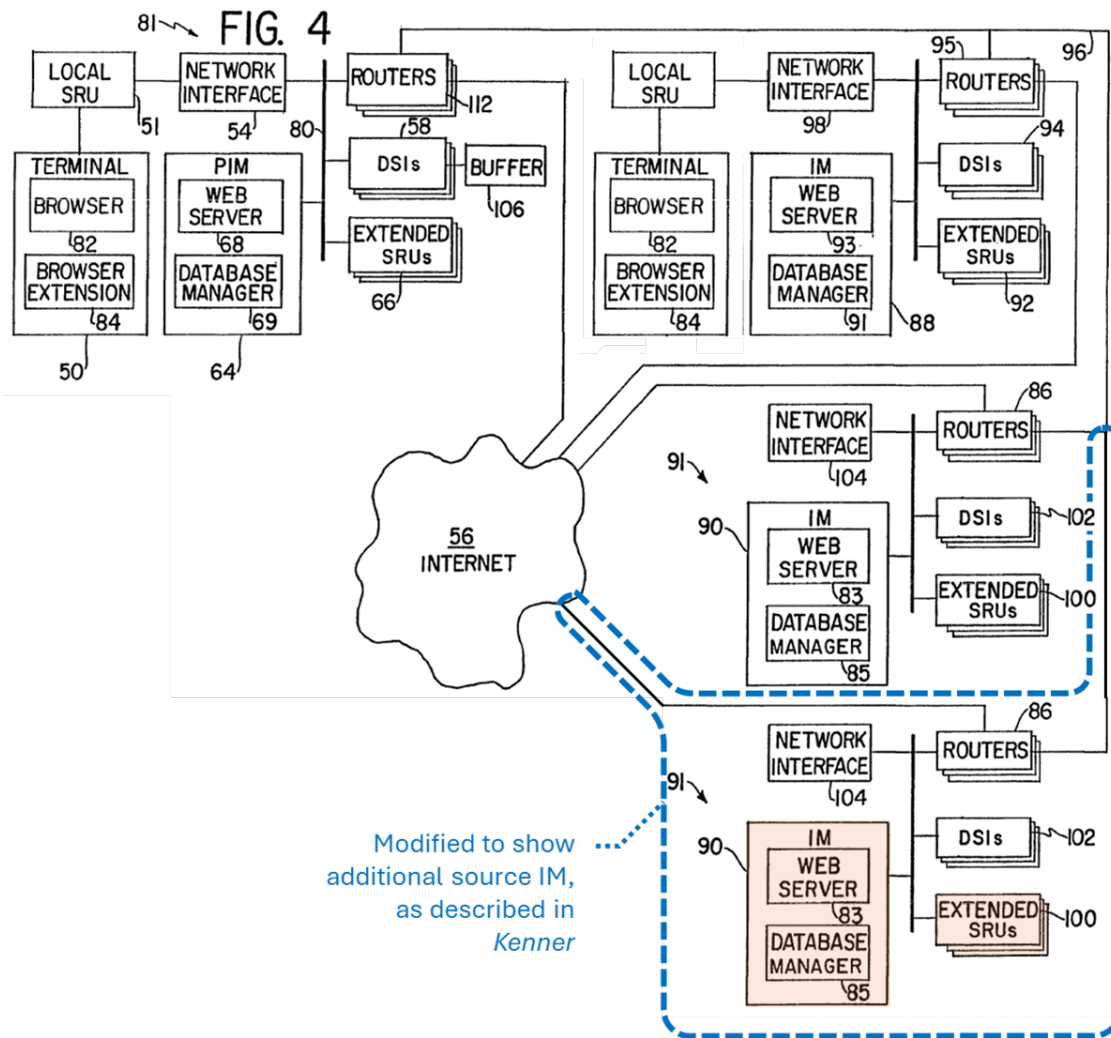
55. Although Fig. 4 does not show the terminals associated with other IMs, the **IM for a neighboring region** may equally include an associated **user terminal**. *Id.*, 32:4-6. Modified Figure 4 below replicates the local SRU and terminal from PIM 64 to show the terminal local to IM 88:



Id., Fig. 4 (modified as indicated).

56. *Kenner*’s system also accommodates multiple content providers and source IMs. *Kenner* explains a primary IM connects to “many geographically

distributed video data sources.” *Id.*, 19:18-19. The primary IM thus maintains a database listing information for each clip that it uses to identify the original source IM. *Id.*, 41:13-19 (listing “[t]he Internet address of the source IM 90” for each clip). The inclusion of this database indicates that the IM stores clips from many different source IMs. Accordingly, although not shown in Fig. 4, *Kenner* is not limited to a single source IM and can include additional source IMs. For purposes of illustration, modified Figure 4 below replicates source IM 90 and associated components to show this **additional source IM**:



Id., Fig. 4 (modified as indicated). A client terminal may therefore be used to access content uploaded from multiple source IMs.

57. With these multiple source IMs, *Kenner* also accommodates multiple content providers. *Id.*, 43:17-25 (describing multiple content providers). Accordingly, two video clips could be associated with the same content provider, or could be associated with different contentment providers, depending on which content providers uploaded the clips. For example, *Kenner* envisions a subscription

service, where a user must subscribe to a particular content provider to view that provider's content. *Id.*, 33:4-11, 35:11-16, 43:17-25. Accordingly, a user may subscribe to one content provider and not another. And each content provider may be associated with a different source IM. For example, in describing how copies of a specified clip can be removed from the system, *Kenner* indicates a content provider is associated with a particular source IM 90. *Id.*, 45:25-31.

58. In summary, *Kenner* is configured to accommodate multiple users, each accessing the system through a different user terminal associated with a primary IM for that particular region, and each capable of accessing content uploaded by multiple content providers via their associated source IMs.

B. *Vetter*

59. *Vetter* is a publication promoting Mosaic, a graphics-oriented browser for accessing web-based content. *Vetter*, Abstract. *Vetter* is in the same field of endeavor as the '903 patent: web-based content delivery. *Id.*, 1.

60. *Vetter* explains that browsers, like Mosaic, allow users to access content using various protocols, including a file-transfer protocol (FTP). *Id.*, 29; see *Berners-Lee*, 78. This content is specified using a uniform resource locator (URL), which identifies the type of resource being accessed and the path for the file. *Vetter*, 29. *Vetter* explains for example that FTP and HTTP are common schemes or resource types. *Id.*

C. *Rekimoto*

61. *Rekimoto* is also in the same field of endeavor as the '903 patent. *Rekimoto*, 2:22-29.

62. Like *Kenner*, *Rekimoto* discloses a system for delivering content to user terminals within a network. *Id.* In *Rekimoto*'s system, this content is a shared virtual reality environment provided by a shared server terminal. *Id.*, 2:30-34, 12:9-12. Accordingly, users in different parts of the world can interact through the same virtual environment. *Id.*, 10:29-36. To avoid delay caused by remote locations, *Rekimoto*'s system includes a mapping server that identifies a shared server local to the users' terminals managing the virtual environment. *Id.*, 17:40-57. This mapping server uses domain name system (DNS) resolution to identify the closest server. *Id.*, 3:3-27.

D. *Boyles*

63. *Boyles* is likewise in the same field of endeavor as the '903 patent: web-based content delivery. *Boyles*, 1:7-10.

64. *Boyles* discloses a network of cache server nodes "responsible for caching or storing information about all or a subset of the network's resources." *Id.*, 2:55-3:3. A cache server node receives a "LOCATE request" and, in response, "searches its local resource directories for an entry defining the location and characteristics of the target resource." *Id.* "If no suitable local entry is found, the

origin cache server node directs the LOCATE requests to one or more alternate cache server nodes” and then “to other nodes in the network.” *Id.* Various other checks may be performed and “[i]f ... the target resource is either not known or not available in the network[,] ... [a] negative reply is returned ... to the requesting node.” *Id.*, 8:58-65.

IX. CLAIMS 1, 2, 22-24, AND 26-46 ARE UNPATENTABLE

65. As I outline below, the '903 claims recite a combination of well-known prior art elements that perform their known functions to produce predictable results.

A. Ground 1: *Kenner* Renders Obvious Claims 1, 2, 22-24, 26, 28-32, 35, 37, 38, 40, 41, and 43-45


1. Independent Claim 1

- a. [1.pre]: “A content delivery system operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines”**

66. In my opinion, *Kenner* discloses [1.pre]. *Kenner* discloses a “distributed digital video clip delivery system” (*content delivery system*) in which video clips (*resources*) “can be requested ... and retrieved at the user’s multimedia terminal” (i.e., delivered to *client machines*). *Kenner*, Abstract. The video clips are uploaded by and thus are *associated with a plurality of content providers*. *Id.*, 43:17-25; § VIII.A.2. A POSITA would have understood that a clip uploaded by a content provider is associated with that content provider. *Kenner* confirms this

association, explaining the video ID for a clip includes a content provider’s account number. *Kenner*, 36:32-37:6.

67. *Kenner*’s architecture is suitable for supporting “thousands of simultaneously attached users” and thus delivers resources to **multiple terminals** (*multiple client machines*), as shown in modified Figure 4:



Kenner, Fig. 4 (modified as indicated), 5:21-26.


68. *Kenner*’s system is operative in a computer network because it delivers content over the “Internet or another general purpose network.” *Id.*, 4:5-8, 1:10-11; *see also id.*, 32:2-6, 5:24-26.

- b. [1.a.i]: “at least one shared repeater server constructed and adapted to replicate at least some of

the resources associated with a first content provider of said plurality of content providers”

69. In my opinion, *Kenner* discloses [1.a.i]. As I explained, *Kenner* discloses a network of IMs. § IX.A.1.a. When a clip is requested, the request is directed to a primary IM (*shared repeater server*) associated with the user’s terminal. *Kenner*, 10:24-31, 16:2-9; § VIII.A.1.

70. The primary IM is *constructed and adapted to replicate at least some of the resources associated with a first content provider* because it stores copies of resources, such as video files, on an associated extended search and retrieval unit (SRU). *Kenner*, 10:24-31. For example, a video file (*resource*) stored at the **primary IM** is initially uploaded by a content provider at a **source IM**, as shown in Fig. 4:



Id., Fig. 4, 41:13-16, 44:1-5.

71. Despite the different terminology, *Kenner's* primary IM is a shared repeater server. *Kenner's* IM is a *server* because it includes web server 83. *Id.*, 19-23. As the '903 patent explains, a repeater is a server that “replicates some or all of the information available on the origin server.” EX-1001, 4:22-25. *Kenner's* IM performs this same function, replicating some or all of the resources of a source IM. For example, once a file is uploaded, the source IM distributes the file to other IMs throughout the network. *Kenner*, 44:15-26. The IM is a *shared* server because it is used for delivering resources from multiple source IMs. In other words, content providers and source IMs share the IMs in the system for distributing content, consistent with the meaning in the '903 patent. § VIII.A.2 (explaining *Kenner* includes multiple source IMs); EX-1001, 23:38-41.

72. *Kenner's* clips are *associated with a first content provider* by virtue of being uploaded by the content provider. § IX.A.1.a. Once uploaded, the source IM transfers the file to other IMs throughout the system, including the primary IM and/or neighboring region IMs. *Kenner*, 44:15-26. These IMs then “accept the new clip and subsequently transfer it to their respective SRUs,” thus replicating the resource of the content provider. *Id.*

- c. **[1.a.ii]: “at least one shared repeater server constructed and adapted ... to replicate at least some**

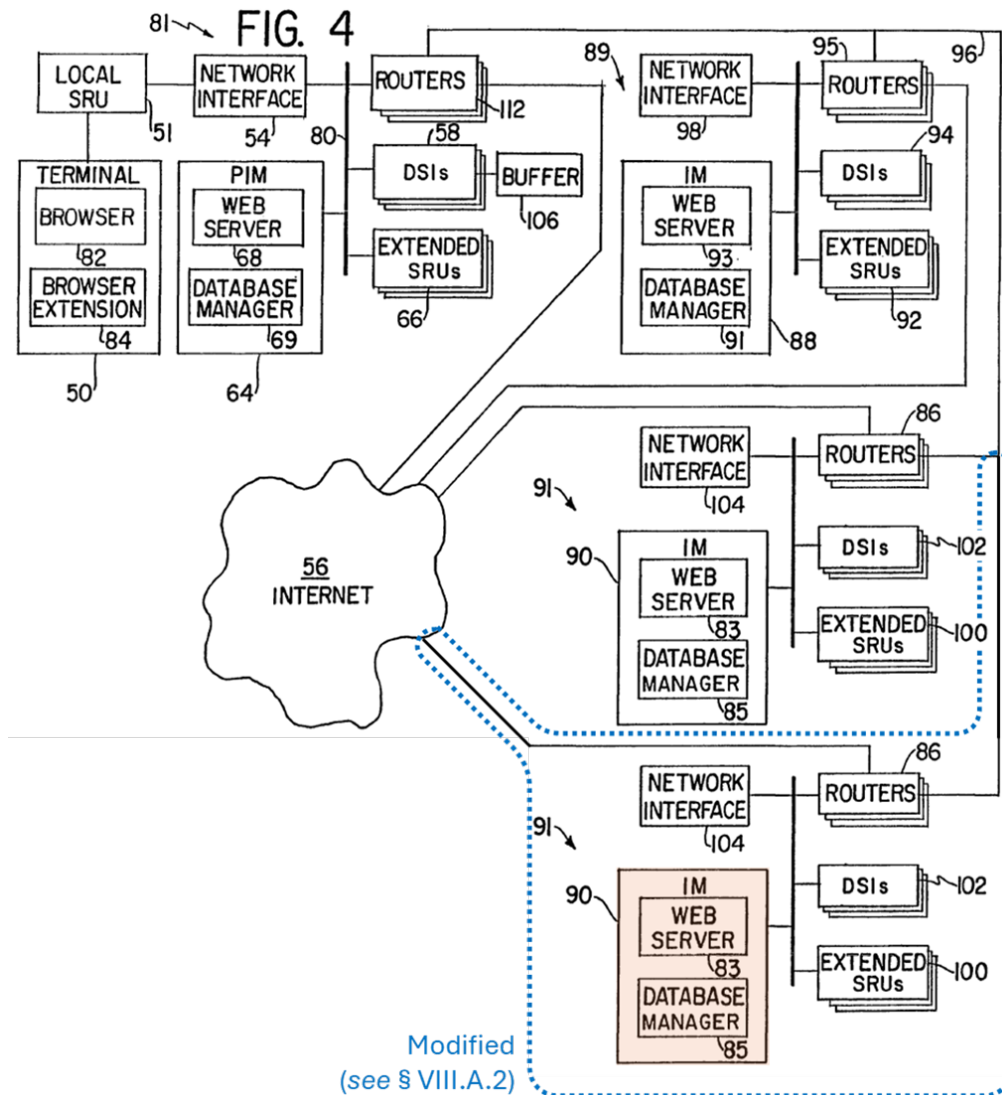
of the resources associated with a second content provider of said plurality of content providers, wherein the second content provider is distinct from the first content provider”

73. In my opinion, *Kenner* discloses [1.a.ii]. As I explained for [1.a.i], *Kenner* discloses *replicat[ing] at least some of the resources associated with a first content provider* by a primary IM. § IX.A.1.b. *Kenner*'s primary IM equally replicates resources from other content providers, each of which may be associated with respective source IMs. § VIII.A.2; *Kenner*, 19:18-19 (“connect[ing] many geographically distributed video data sources to one subscriber destination”).

74. For example, the primary IM maintains a clip database storing information for each of the clips included in its extended SRUs. *Kenner*, 34:3-35:3. The database further correlates each of the clips with the source IM that uploaded the clip, which the primary IM references to identify the source IM. *Id.*, 41:13-19. Accordingly, a POSITA would have understood that the primary IM stores clips from multiple source IMs.

75. As I explained, a primary IM stores clips from multiple source IMs. § VIII.A.2. Because the primary IM stores multiple clips, each associated with a different source IM, the primary IM is *constructed and adapted ... to replicate at least some of the resources associated with a second content provider of said*

plurality of content providers. For example, Kenner's network includes an **additional source IM**, as shown in modified Figure 4 below:



Kenner Fig. 4; § VIII.A.2.

76. A POSITA would further have understood that the second content provider associated with this additional source IM would be *distinct from the first*

content provider because *Kenner's* system accommodates multiple content providers. *Kenner*, 43:17-25 (referring to “content providers”); *see* § VIII.A.2.

- d. **[1.b]: “at least one table listing a plurality of alias names corresponding to content providers authorized to have resources delivered to client machines via the at least one shared repeater server, wherein the plurality of content providers comprises the first content provider and the second content provider”**

77. In my opinion, *Kenner* discloses or at least suggests [1.b]. *Kenner* discloses *at least one table* because the primary IM maintains a clip database (*table*) storing information on the video clips stored on its extended SRUs. *Kenner*, 34:3-6, 37:32-38:3, 38:32-39:4, 41:13-19, 43:17-25. A POSITA would have recognized that this database is a table because the database includes information for multiple clips and with several corresponding data fields for each clip. *Id.*, 34:3-35:3. In the database, each of these fields represents columns in the table, with each row of the table corresponding to a different clip. For example, *Kenner's* database includes a “[c]ounter” field representing “[p]rimary index for the records” where “[e]ach record represents one video clip.” *Id.*, 34:8. A POSITA would have understood that one column of the table includes these counters for different video clips and each other column includes corresponding properties of the clip.

78. Storing data in this manner, with rows representing each clip and columns representing corresponding clip information, constitutes a table. *Kenner*

confirms this format, indicating the database may be stored as a “flat-file database.” *Id.*, 20:21-27, 21:30-22:13 (describing database format). To the extent *Kenner* does not explicitly disclose a table, it would have been obvious to implement the clip database as a table because *Kenner*’s database is used to identify information related to a particular video ID. *Id.*, 21:30-22:13, 34:3-35:3. Tables are a fundamental data structure for organizing data reflecting these relationships.

79. *Kenner*’s database *list[s] a plurality of alias names* because it includes a video ID (*alias name*) field representing a “globally unique name of the video clip.” *Id.*, 34:9. The video IDs are alias names because they are identifiers (i.e., names) used to refer to a particular resource or entity. *Id.*

80. *Kenner*’s video IDs (*alias names*) *correspond[] to content providers* because the video IDs refer to the particular content provider that uploaded the particular video clip. For example, *Kenner* explains that the video ID includes “a text name of the file as defined by the content provider” and “the content provider’s account number.” *Id.*, 36:32-37:6.

81. *Kenner*’s content providers are *authorized to have resources delivered to client machines via the at least one shared repeater server* because prior to downloading a clip to a user’s terminal the primary IM “checks the user’s subscription rights in its user database” to determine if the download is “authorized

and necessary.” *Id.*, 37:18-31, 38:17-22, 52:7-12. This is the same form of authorization described in the ’903 patent. For example, the ’903 patent explains that “[e]ach new request for the resource must be tested at the origin server to assure that the requester is authorized to access the resource.” EX-1001, 23:17-20.

82. *Kenner* discloses *the plurality of content providers compris[ing] the first content provider and the second content provider* as I explained for [1.a.i] and [1.a.ii]. §§ IX.A.1.b, IX.A.1.c. Because *Kenner*’s database can include multiple clips, each associated with its own respective content providers, a POSITA would have understood *Kenner*’s database includes first and second content providers. *Kenner*, 41:13-16.

- e. **[1.c]: “wherein at least a first resource associated with the first content provider is associated with a first alias name of the plurality of alias names, said first alias name being associated with said at least one shared repeater server”**

83. In my opinion, *Kenner* discloses [1.c]. *Kenner* discloses *at least a first resource associated with the first content provider* because each clip (*resource*) is uploaded by a content provider. *Kenner*, 44:1-5. A first content provider uploads a first clip (*first resource*) via a first source IM and this clip would be distributed to IMs within the system. *Id.*, 44:6-26 The clip is also associated with the content provider because the “content provider’s account number” is included in the video ID. *Id.*, 36:32-37:6.

84. *Kenner's first clip (first resource) is associated with a first alias name of the plurality of alias names* because each clip is associated with a corresponding video ID. In uploading a clip, the source IM registers the clip along with a video ID (*alias name*). *Id.*, 44:6-14. And the video ID “is a unique reference identifier for [the clip.” *Id.*, 22:1-2, 43:27-29 (“[E]ach clip ... ha[s] a unique video ID”).

85. *Kenner also discloses the first alias name being associated with said at least one shared repeater server* because *Kenner's* system tracks which video clips are stored in which IMs using the video ID. For example, *Kenner's* IMs maintain a database of the clips stored on its SRUs, reflecting an association between the clips and the IM. *Id.*, 34:3-6. And *Kenner's* IMs store a database indicating “which IMs are likely to have the desired clip,” also reflecting this association. *Id.*, 41:1-6. Indeed, not every IM receives every clip. For example, when updating a clip, updates are only sent to IMs storing the clip. *Id.*, 44:15-20. Accordingly, an IM is associated with a video ID by virtue of storing the associated clip.

86. This association is reflected, for example, in the clip database of the IM. *Id.* As I explained for [1.a.i], when accepting a new clip from a source IM, the primary IMs “subsequently transfer [the clip] to their respective SRUs” (*Id.*, 44:15-20) and register the video ID in its clip database (*id.*, 34:3-6 (the IM

“maintains information on the audio-visual clips stored on its extended SRUs 66”), 37:32-38:3, 38:32-39:4).

87. This association is also reflected in another database maintained at each IM, which lists the “the Internet address of [every] IM and the content coordinates of all audio-video files that it maintains.” *Id.*, 38:3-8. A POSITA would have recognized that IMs identified as “likely to have the desired clip” are associated with that video ID. *Id.*, 41:1-5.

- f. **[1.d]: “wherein, requests for the first resource from said client machines are directed to the at least one shared repeater server, based at least in part on said first alias name”**

88. In my opinion, *Kenner* discloses [1.d]. *Kenner* discloses *requests for the first resource from said client machines* because, to access a particular clip, a user requests access to the clip through a browser application on a user terminal (client machine). *Kenner*, 32:25-33, 36:10-27.

89. *Kenner*'s requests are directed to the primary IM and thus are *directed to the at least one shared repeater server*. A local SRU within the client's network “intercepts [the] video ID ... specifying the selected clip” in the request and “passes the video ID to the PIM 64 associated with the user's terminal.” *Id.*, 36:22-27. Because the video ID is “embedded within the HTML of the Web page,” the request is directed to the IM based on the video ID (alias name). *Id.*, 36:15-16.

90. Alternatively, *Kenner* also discloses requests *directed to the at least one shared repeater server* because a request may be directed to the primary IM from a neighboring IM. If a given IM (i.e., a neighboring IM) does not identify a desired clip on its extended SRUs, it “will query the closest IMs (e.g., IM 88) to determine if any of the remote SRUs 92 ... have the desired clip.” *Id.*, 40:29-33. To do so, the IM uses content coordinate data included in the video ID (*alias name*) (*id.*, 36:28-31) to “determine which IMs are likely to have the desired clip, and query only those IMs” (*id.*, 41:1-5). Accordingly, when a neighboring IM does not store a clip on its own SRUs, a request for the clip will be directed to the primary IM (*at least one shared repeater server*) based on the video ID. *See* VIII.A.2.

- g. [1.e]: “wherein at least a second resource associated with the second content provider is associated with a second alias name of the plurality of alias names, said second alias name being associated with said at least one shared repeater server”**

91. In my opinion, *Kenner* discloses [1.e]. As I explained for [1.c], *Kenner* discloses *at least a first resource associated with the first content provider [being] associated with a first alias name of the plurality of alias names.* § IX.A.1.e. A POSITA would have understood that the same applies to *a second resource associated with the second content provider.* In *Kenner*, a given IM stores multiple clips in its associated SRUs, each being associated with their own video IDs (*Kenner*, 34:3-9, 43:27-29) and respective content providers (*id.*, 36:32-37:6).

Accordingly, where the system includes a second content provider (*see* § IX.A.1.c), the given IM is associated with a second clip uploaded by the second content provider. § IX.A.1.e.

92. *Kenner's* second video clip (*second resource*) is associated with a *second alias name of the plurality of alias names* because this second clip would be assigned its own unique video ID. *Kenner*, 22:1-2, 43:27-29. And *Kenner* discloses the *second alias name being associated with said at least one shared repeater server* because the IM records the associated video ID in its clip database. *Id.*, 34:3-6, 37:32-38:3, 38:32-39:4.

h. [1.f]: “wherein requests for the second resource made from said client machines are directed to the at least one shared repeater server, based at least in part on said second alias name”

93. In my opinion, *Kenner* discloses [1.f]. As I explained for [1.d], *Kenner* discloses *requests for the first resource made from said client machines [being] directed to the at least one shared repeater server*. § IX.A.1.f. The same would apply to requests for the second resource as the user requests access to any resource in the same way. The user thus requests access to the second clip through the user terminal (client machine). *Kenner*, 36:10-27.

94. *Kenner's* requests are *directed to the at least one shared repeater server* for the same reasons the requests for the first clip are directed to the at least one

shared repeater server. § IX.A.1.f. Indeed, the IM connects “many geographically distributed video data sources to once subscriber destination.” *Kenner*, 19:18-19.

- i. **[1.g]: “wherein the at least one shared repeater server is further constructed and adapted to analyze, using the table, an alias name received with a client request for a particular resource to determine a content provider associated with the particular resource”**

95. In my opinion, *Kenner* discloses [1.g]. *Kenner* discloses *an alias name received with a client request for a particular resource* because *Kenner*’s video ID is provided to the primary IM with the request. As I explained for [1.d], in *Kenner*, when a user requests a particular video, a local SRU for the user’s terminal “passes the video ID to the PIM 64 associated with the user’s terminal.” *Kenner*, 36:25-27; § IX.A.1.f.

96. *Kenner* discloses *the at least one shared repeater server [being] constructed and adapted to analyze the alias name using the table ... to determine a content provider associated with the particular resource* because it uses the clip database to identify the source IM. As I explained for [1.d], if the primary IM determines a clip is not stored at its associated extended SRUs, it queries the closest IMs. § IX.A.1.f; *Kenner*, 40:29-41:5. If “after having queried the neighboring remote IMs, [PIM 64] is still unable to locate the desired clip on an SRU 66 or 92, the PIM 64 will then contact the source IM 90, where the content provider first uploaded the file.” *Id.*, 41:13-15. The primary IM thus consults the clip database to determine an

internet address of the source IM. *Id.* In view of this disclosure, a POSITA would have understood that identifying the source IM also identifies the content provider because a content provider is associated with a particular source IM it used to upload the content. *Kenner* confirms this association because the “content provider’s account number” is included in the video ID within the database. *Id.*, 36:32-37:6.

2. Claim 2: “wherein the least one shared repeater server is addressable using the first alias name and the second alias name”

97. In my opinion, *Kenner* discloses claim 2. As I explained for [1.d], *Kenner* discloses *requests for the first resource ... [being] directed to the at least one shared repeater server, based at least in part on said first alias name.* § IX.A.1.f. The SRU constructs a “‘virtual URL’ ... in the form ‘http:’ plus the Internet address of the PIM 64 ... plus the video ID,” and directs the request to the primary IM using the virtual URL through an NPN_GetURLNotify procedure. *Kenner*, 37:7-15. Accordingly, the URL is used to address the primary IM. And the primary IM is *addressable using the first alias name*, because the video ID is used to generate the URL.

98. As I explained for [1.f], *Kenner* further discloses *requests for the second resource ... [being] directed to the at least one shared repeater server, based at least in part on said second alias name.* § IX.A.1.h. The primary IM is also *addressable using the second alias name* for the same reason as the first alias name.

In other words, the process described above for addressing the IM using the first alias name would be the same as when a second alias name is used, although the second alias name would be used in place of the first alias name.

3. Claim 22: “wherein at least one of the first resource and the second resource is an embedded object”

99. In my opinion, *Kenner* discloses claim 22. *Kenner* expressly discloses accessing video clips as an *embedded object*. For example, *Kenner* describes accessing clips formatted as an “embedded file.” *Kenner*, 36:19-27. For these clips, a “reference to [the] desired clip is embedded within the HTML of the Web page” such that a “request is made of the Web server 83 to transmit the embedded file.” *Id.*, 36:15-18.

4. Claim 23: “wherein the requests comprise host tag names that identify specific origin servers from which the resources originate”

100. In my opinion, *Kenner* discloses claim 23. *Kenner* discloses delivering HTTP-based clips. *Kenner*, 36:15-18 (embedding a “reference to a desired clip ... within the HTML of the Web page”), 2:16-21 (the “precise location” of the clip is specified “in the format ‘http://internet.address/directory/filename.html.’”).

101. A POSITA would have recognized that including host tag names in requests for HTTP-based resources. This was standard practice at the time. These HTTP requests include HTTP headers containing metainformation about the

requested resource. *Berners-Lee*, ¶ 78. The Request for Comments (RFC) 2068 published in January 1997 by the Internet Engineering Task Force (IETF) (EX-1012) confirms that HTTP requests included request headers at the time. *HTTP/1.1*, § 5. And it was standard practice for HTTP requests to include host header fields (*host tag names*) in these HTTP headers. *Id.*, § 5.3. The host header fields *identify specific origin servers from which the resources originate* because they identify the host of the origin server for a requested resource. *Id.*, §§ 5.1.2, 14.23.

5. Claim 24: “wherein the host tag names are included in HTTP headers”

102. In my opinion, *Kenner* discloses claim 24. As I explained for claim 23, *Kenner*’s HTTP-based requests include *host tag names*. § IX.A.4. It was standard practice to include the host tags in HTTP headers for the request. *Id.*; *HTTP/1.1*, §§ 5.1.2, 14.23.

6. Claim 26: “wherein the table further comprises, for each of the plurality of content providers, information from the group consisting of: a reflector name and a committed aggregate information rate (CAIR)”

103. In my opinion, *Kenner* discloses or at least suggests claim 26. As I explained for [1.b], *Kenner* discloses a table listing a plurality alias names corresponding to content providers. § IX.A.1.d.

104. *Kenner* further discloses including *a reflector name* in its clip database. As the ’903 patent explains, a reflector is “preferably a software program” that

“intercepts requests that would normally be sent directly to the origin server.” EX-1001, 5:32-39; 3:4-20. The “reflector 108 and the origin server 102 are typically co-located, e.g., on a particular system such as data server 112.” *Id.*, 5:32-39.

105. *Kenner*'s IMs include various software modules for receiving requests and determining whether to handle them locally in its extended SRUs, or through remote SRUs. *Kenner*, 10:32-12:35, 16:10-17:27, 37:7-31, 40:29-41:6. For example, *Kenner*'s IM creates local DSIs executing at the IM (*id.*, 19:4-17) that query the extended and remote SRUs (*id.*, 25:25-26:7, 38:32-39:4). This software executed at the IM is a *reflector* because it performs the same function as the reflector software described in the '903 patent. EX-1001, 5:32-39; 3:4-20.

106. *Kenner*'s clip database includes the “Internet address of the source IM.” *Kenner*, 41:13-16. A POSITA would have understood that this address is a reflector name because it identifies the IM of the server where the clip was uploaded and the software executing on this IM performs the same functions as the reflector software in the '903 patent. EX-1001, 5:32-39; 3:4-20. And each of the source IMs is associated with a corresponding content provider. § VIII.A.2.

107. *Kenner* also at least suggests including a *committed aggregate information rate (CAIR)* in its table. The '903 patent explains that CAIR is an aggregate rate at which data is served on behalf of a given subscriber by all repeaters.

EX-1001, 17:51-57, 19:57-59. *Kenner* explains that its IM “maintains information on the subscriber in a user database” including “types of content subscribed to, user preferences, **limitations on service**, and billing information.” *Kenner*, 33:4-11. A POSITA would have understood that these “limitations on service” include bandwidth limitations for a particular subscriber.

108. This rate of transfer of information committed to a particular subscriber is the same aggregate rate described in the '903 patent as the “CAIR.” EX-1001, 17:51-57, 19:57-59. For example, *Kenner* indicates it varies playback quality and resolution for different users (*Kenner*, 51:30-52:6), which would affect transmission bandwidth. And *Kenner* tracks the download speeds for individual users. *Id.*, 47:8-19.

109. To the extent *Kenner* does not explicitly disclose the claimed CAIR, it would have been obvious to include this information in *Kenner*'s database. For example, *Kenner* discloses storing data to improve load projection, such as a usage count for a particular clip. *Id.*, 35:1. *Kenner*'s IMs also maintain information correlating usage rates for individual users to content providers the users are subscribed to. For example, *Kenner*'s IM's maintain information such as a “list of services to which the user's account has subscribed” and the “maximum monthly expense that can be incurred by the account.” *Id.*, 33:19-22. The system thus tracks

which users are subscribed to each content provider and the approximate download rates for those users.

110. A POSITA would have found it obvious to use this information to allocate bandwidth usage across content providers. *Kenner* already relies on load projections (*id.*, 35:1) and emphasizes the need for dynamic load management based on predicted usage (*id.*, 46:7-14). Allocating bandwidth across content providers would further improve the load projections and bandwidth management. A POSITA would have further found it obvious to store information to allocate bandwidth usage across content providers as a committed aggregate data rate across IMs for each given content provider as the IMs dynamically assess the load for their respective regions. *Id.*, 48:1-3. A POSITA would thus have stored committed aggregate information rate data to further improve the load projections in *Kenner*, especially in view of the bandwidth limitations discussed above.

7. Independent Claim 28

- a. **[28.pre.i]: “A method, in a content delivery system operative in a computer network for delivering content to client machines”**

111. In my opinion, *Kenner* discloses [28.pre.i]. As I explained for [1.pre], *Kenner* discloses a content delivery system operative in a computer network. § IX.A.1.a. *Kenner* discloses a *method ... for delivering content to client machines*

in this network because it describes the process for accessing content via user terminals. *Kenner*, 35:11-16, claim 55.

b. [28.pre.ii]: “the computer network comprising a plurality of origin servers, each of said origin servers having resources associated therewith”

112. In my opinion, *Kenner* discloses [28.pre.ii]. As I explained for [1.a.ii], *Kenner*'s network includes a plurality of source IMs through which content providers upload clips. § IX.A.1.c. These source IMs (or “originating” IMs) (*Kenner*, 45:8-14) are *origin servers* because they include a web server and are the origin of uploaded clips (*id.*, 32:14-20). *Kenner*'s source IMs have *resources associated therewith* at least because the source IMs store copies of uploaded clips to their associated SRUs. *Id.*, 44:6-14, *see also id.*, 41:13-16 (describing clip database correlating files to source IMs).

c. [28.pre.iii]: “the content delivery system comprising at least one shared repeater server operable to replicate resources associated with the plurality of origin servers”

113. In my opinion, *Kenner* discloses [28.pre.iii]. As I explained for [1.a.i], *Kenner*'s IM is *at least one shared repeater server constructed and adapted to replicate* resources. § IX.A.1.b. *Kenner*'s IM is *operable to replicate* these resources because it performs the required operations to do so. *Kenner*, 44:17-18. And as I

explained for [28.pre.ii], the resources replicated by the IM are associated with source IMs (i.e., origin servers). § IX.A.7.b.

- d. [28.a]: “associating the at least one repeater server with a first alias name, wherein requests for a first resource located on a first origin server are directed, based at least in part on said first alias name, to the at least one repeater server for delivery of the first resource from said at least one repeater server”**

114. In my opinion, *Kenner* discloses [28.a]. As I explained for [1.c], *Kenner* discloses a first alias name associated with the at least one repeater server. § IX.A.1.e. *Kenner* likewise discloses *associating the at least one repeater server with a first alias name*. This association is formed through the IM recording the video ID (i.e., alias name) for a clip stored in its extended SRUs in a clip database. *Kenner*, 34:3-6, 37:32-38:3, 38:32-39:4. IMs also form this association in another database listing “the Internet address of [each other] IM and the content coordinates of all audio-video files that it maintains.” *Id.*, 38:3-8.

115. As I explained for [1.d], *Kenner* discloses *requests for [a] first resource ... [being] directed to the at least one shared repeater server, based at least in part on said first alias name*. § IX.A.1.f. *Kenner* discloses the first resource being *located on a first origin server* because when uploading a clip, *Kenner*’s source IM “cop[ies] the file to at least one of its extended SRUs.” *Kenner*, 44:6-14.

116. As I explained for [1.c], *Kenner* discloses requests being *directed, based at least in part on said first alias name, to the at least one repeater server.*

§ IX.A.1.e. The requests are directed “for delivery of the first resource from said at least one repeater server” because in response to the request, *Kenner*’s IM downloads the clip to the user’s terminal. *Kenner*, 37:18-31.

- e. **[28.b]: “associating the at least one repeater server with a second alias name, wherein requests for a second resource located on a second origin server are directed, based at least in part on said second alias name, to the at least one repeater server for delivery of the second resource from said at least one repeater server, wherein the second origin server is distinct from the first origin server”**

117. In my opinion, *Kenner* discloses [28.b].

118. As I explained for [1.e], *Kenner* discloses a second alias name associated with the at least one repeater server. § IX.A.1.g. As with the first alias name, the IM forms the association through recording the video ID (alias name) for the second clip in the clip database. *Kenner*, 34:3-6, 37:32-38:3, 38:32-39:4.

119. As I explained for [1.f], *Kenner* discloses *requests for a second resource located on a second origin server are directed, based at least in part on said second alias name, to the at least one repeater server.* § IX.A.1.c. *Kenner* discloses the second resource being *located on a second origin server* and the requests being directed *for delivery of the second resource from said at least one*

repeater server for the same reasons as the first resource. *See* § IX.A.7.d. Indeed, *Kenner's* IMs store clips from multiple different source IMs. § IX.A.1.c.

- f. **[28.c]: “providing a table listing origin servers having content located thereon, wherein said content is authorized for delivery to client machines via the at least one shared repeater server, and wherein the origin servers comprise the first origin server and the second origin server”**

120. In my opinion, *Kenner* discloses or at least suggests [28.c].

121. As I explained for [1.b], an IM in *Kenner's* network maintains a clip database storing information on the video clips stored on its extended SRUs, which would be represented as a table. *Kenner*, 34:3-6, 37:32-38:3, 38:32-39:4, 41:13-19, 43:17-25; § IX.A.1.d.

122. *Kenner's* table *list[s] origin servers having content located thereon* because the clip database includes the address of the source IM where the content provider first uploaded the file. *Kenner*, 41:13-16. As I explained for [28.b], *Kenner's* IM includes clips from first and second source IMs, and its clip database would thus list *the first origin server and the second origin server*. § IX.A.7.e.

123. *Kenner* discloses *content ... authorized for delivery to client machines via the at least one shared repeater server* because users must be authorized to access the clips. § IX.A.1.d. Before delivering content to the user, the IM checks “that the user is authorized to receive the desired clip.” *Kenner*, 38:23-31.

-
- g. **[28.d]: “wherein the at least one repeater server is further constructed and adapted to analyze, using the table, an alias name received with a client request for a particular resource to determine an origin server associated with the particular resource”**

124. *Kenner* discloses [28.d], as I explained for [1.g]. § IX.A.1.i; *Kenner*, 41:13-16 (the IM “contact[s] the source IM 90, where the content provider first uploaded the file.”).

8. **Claim 29: “further comprising: in response to a request for the first resource issued by a particular client machine, identifying the at least one shared repeater server from a set of shared repeater servers based at least in part on a network cost of delivering resources to the particular client machine; and associating the first resource with the first alias name”**

125. In my opinion, *Kenner* discloses claim 29. As I explained for [28.a], *Kenner* discloses a request for the first resource § IX.A.7.d. *Kenner* discloses this request being issued by a particular client machine because a request for a clip in *Kenner* is issued through a user terminal (i.e., a client machine). *Kenner*, 25:7-24.

126. *Kenner* discloses identifying the at least one shared repeater server in response to the request because *Kenner*’s IM determines whether the clip is stored on its own extended SRUs or on remote SRUs. *Id.*, 25:25-26:7, 38:32-39:4, 39:29-41:5. The primary IM thus identifies either the primary IM (if the clip is retrieved from its extended SRUs) or a remote IM (if the clip is retrieved from remote SRUs) for delivering the clip to the user’s client machine. *Kenner* discloses

identifying the server *from a set of shared repeater servers* because *Kenner's* IM “maintains a database of all other IMs connected to the system,” which it uses to determine IMs likely to have the desired clip. *Id.*, 41:1-6

127. *Kenner* further discloses *identifying the at least one shared repeater server ... based at least in part on a network cost of delivering resources to the particular client machine* because in determining which SRU should deliver the clip, the IM (via a data sequencing interface (“DSI”) process) “actively determines which computing systems and communication paths to the user should be used for each download.” *Id.*, 19:29-31. For example, *Kenner's* IM queries “succeedingly more remote servers” such that “the fastest possible path will be selected, and traffic will be minimized on the network.” *Id.*, 9:26-29, 8:11-15 (“minimizing the number of network nodes traversed”). This is the same form of cost described in the '903 patent. EX-1001, 13:16-21 (describing a “statically determined cost for transmitting data” based on “sums of the costs of the links along each path.”).

128. *Kenner* also discloses this network cost by accounting for “parameters such as the *cost* of transmitting and storing duplicate information” when identifying an SRU for storing a clip. *Kenner*, 16:32-17:2 (emphasis added). These IMs are identified based on this network cost because the primary IM identifies which SRUs store the clip.

129. *Kenner* discloses *associating the first resource with the first alias name* because the video ID (*alias name*) for a clip (i.e., *first resource*) is assigned when the source IM uploads the clip. *Id.*, 43:2-7, 44:6-14, 22:1-2 (video IDs are “unique reference identifier[s] for each video clip”), 34:9 (the video ID is a “globally unique name of the video clip”); § IX.A.1.e.

9. **Claim 30: “further comprising: in response to a request for the first resource issued by a particular client machine, identifying the at least one repeater server from a set of shared repeater servers based at least in part on load characteristics of the set of shared repeater servers; and associating the first resource with the first alias name”**

130. In my opinion, *Kenner* discloses claim 30.

131. As I explained for claim 29, *Kenner* discloses *identifying the at least one repeater server from a set of shared repeater servers and associating the first resource with the first alias name*. § IX.A.8.

132. *Kenner* further discloses identifying the IM (*repeater server*) based on *load characteristics*. For example, when identifying an SRU to deliver the clip, *Kenner*'s SRUs “are prioritized according to apparent load” (i.e., *load characteristics*). *Kenner*, 39:1-4, 41:7-12 (considering apparent load of remote SRUs); *see* EX-1001, 11:32-39.

10. Claim 31: “further comprising: in response to a request for the first resource issued by a particular client machine, identifying the at least one repeater server from a set of shared repeater servers based at least in part on a location of the particular client machine relative to the at least one shared repeater server; and associating the first resource with the first alias name”

133. In my opinion, *Kenner* discloses claim 31.

134. As I explained for claim 29, *Kenner* discloses *identifying the at least one repeater server from a set of shared repeater servers and associating the first resource with the first alias name.* § IX.A.8.

135. *Kenner* further discloses identifying the IM (*repeater server*) based on location. For example, *Kenner*'s IM queries “succeedingly more remote servers” to identify an SRU and thus selects an IM based on location. *Kenner*, 9:26-29, 8:10-13 (“locat[ing] audio/video content on servers close to those users expected to request it”).

11. Claim 32: “wherein the resources comprise an HTTP-based resource”

136. In my opinion, *Kenner* discloses claim 32. *Kenner* expressly discloses *HTTP-based resource[s]*. For example, *Kenner* discloses constructing a “virtual URL” for a clip “in the form ‘http://’ plus the Internet address of the PIM 64, plus the user’s subscriber ID number, plus the video ID.” *Kenner*, 37:7-17, 2:16-21 (describing URLs specifying “the precise location of the HTML file”).

12. Claim 35: “wherein the table further comprises, for each of the listed origin servers, information from the group consisting of: a reflector name and a committed aggregate information rate (CAIR)”

137. *Kenner* discloses or at least suggests claim 35, as I explained for claim 26. § IX.A.6.

13. Independent Claim 37

a. [37.pre.i]: “A method, in a content delivery system operative in a computer network for delivering content to client machines”

138. *Kenner* discloses [37.pre.i], as I explained for [28.pre.ii]. § IX.A.7.b.

b. [37.pre.ii]: “comprising at least one shared repeater server operable to replicate resources stored on a plurality of origin servers”

139. *Kenner* discloses [37.pre.ii], as I explained for [28.pre.iii]. § IX.A.7.c.

c. [37.a.i]: “associating a plurality of alias names with the at least one shared repeater server, each of said plurality of alias names being associated with an origin server, wherein a first alias name of said plurality of alias names is associated with a first origin server”

140. In my opinion, *Kenner* discloses [37.a.i].

141. As I explained for [28.a] and [28.b], *Kenner* discloses *associating a plurality of alias names with the at least one shared repeater server*. §§ IX.A.7.d, IX.A.7.e. Indeed, *Kenner*’s IM (shared repeater server) stores multiple clips (*Kenner*, 34:3-6), each associated with their own video IDs (*id.*, 43:27-29).

142. *Kenner* further discloses *each of said plurality of alias names being associated with an origin server* at least because the video ID includes “the content provider's account number” (*id.*, 36:32-37:6) and the clip database associates the video ID with the source IM (*id.*, 41:13-16). This would include *a first alias name of said plurality of alias names ... associated with a first origin server* because each clip is associated with the source IM that uploaded the clip. Accordingly, an alias name for a first one of the clips is associated with a first source IM.

- d. **[37.a.ii]: “at least a second alias name of said plurality of alias names is associated with a second origin server distinct from said first origin server”**

143. In my opinion, *Kenner* discloses [37.a.ii].

144. As I explained for [37.a.i], *Kenner* discloses *each of said plurality of alias names being associated with an origin server*. § IX.A.13.c. This includes *a second alias name of said plurality of alias names is associated with a second origin server*. Indeed, *Kenner*'s IM (*shared repeater server*) stores multiple clips (*Kenner*, 34:3-6), each associated with their own video IDs (*alias name*) (*id.*, 43:27-29). As I explained for [28.b], *Kenner* discloses *the second origin server [being] distinct from the first origin server*. § IX.A.7.e.

- e. **[37.b]: “providing at least one table associating alias names with origin servers having content located thereon, wherein said content is authorized for delivery to client machines via the at least one shared**

repeater server, wherein the origin servers comprise the first origin server and the second origin server”

145. In my opinion, *Kenner* discloses [37.b].

146. As I explained for [28.c], *Kenner* discloses *providing at least one table ... [listing] origin servers having content located thereon, wherein said content is authorized for delivery to client machines via the at least one shared repeater server, wherein the origin servers comprise the first origin server and the second origin server.* § IX.A.7.f. The clip database (*table*) *associat[es] alias names with origin servers* because the clip database includes the address of the source IM where the content provider first uploaded the file. *Kenner*, 41:13-16.

- f. **[37.c]: “wherein requests for a first resource originating on the first origin server are directed, based at least in part on said first alias name, to the at least one shared repeater server for delivery of the first resource from said at least one repeater server”**

147. *Kenner* discloses [37.c], as I explained for [28.a]. § IX.A.7.d. The first clip *originat[es] on the first origin server* because the clip is initially uploaded at the first source IM (*origin server*) and is distributed to other IMs in the network. *Kenner*, 44:6-26.

- g. **[37.d]: “wherein requests for a second resource originating on the second origin server are directed, based at least in part on said second alias name, to the at least one shared repeater server for delivery of the**

second resource from said at least one repeater server”

148. *Kenner* discloses [37.d], as I explained for [28.b]. §§ IX.A.7.e, IX.A.13.f.

h. [37.e]: “wherein the at least one repeater server uses the at least one table and an alias name received with a client request for a resource to determine an origin server associated with the requested resource”

149. *Kenner* discloses [37.e], as I explained for [28.d]. § IX.A.7.g; *Kenner*, 41:13-16.

14. Claim 38: “wherein a copy of the at least one table is located on each of a plurality of repeater servers”

150. In my opinion, *Kenner* discloses claim 38. *Kenner* discloses *a copy of the at least one table [being] located on each of a plurality of repeater servers* because *Kenner*’s clip database is stored on each IM (repeater server). *Kenner*, 34:3-5. *Kenner* explains that each IM “will query its clip database to determine if any of its extended SRUs are storing the clip.” *Id.*, 45:25-31; *see id.*, 44:6-14 (describing the source IM’s “own clip database”), 45:32-33 (“Each IM ... performs its own maintenance on the clip database”).

15. Independent Claim 40

- a. **[40.pre.i]: “A server operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines”**


151. In my opinion, *Kenner* discloses [40.pre.i]. As I explained for [1.pre], *Kenner* discloses a *computer network for delivering resources associated with a plurality of content providers to multiple client machines*. § IX.A.1.a. *Kenner*'s IM is a *server operative* in the network. § IX.A.1.b.

- b. **[40.pre.ii]: “wherein the plurality of content providers comprise at least a first content provider and a second content provider distinct from the first content provider”**

152. In my opinion, *Kenner* discloses [40.pre.ii]. As I explained for [1.a.i] and [1.a.ii], *Kenner* discloses multiple content providers (*plurality of content providers*) *compris[ing] at least a first content provider and a second content provider distinct from the first content provider*. §§ IX.A.1.b, IX.A.1.c.

- c. **[40.a]: “the server comprising: ... cache storage”**

153. In my opinion, *Kenner* discloses [40.a]. *Kenner* discloses *cache storage* via a buffer at each IM. For example, *Kenner*'s IMs (e.g., the primary IM) create a data sequencing interface (DSI) to “direct the efficient download of the video clips to the user terminal.” *Kenner*, 10:24-31, 31:26-29. As shown in Fig. 4, *Kenner*'s DSI includes a buffer 106 accessible to the primary IM via the DSI:



Id., Fig. 4, 48:9-15. “[A]ny computer connected to the backbone 80,” can create these DSIs, including the web server of the primary IM. *Id.*, 31:26-29.

154. A POSITA would have understood that *Kenner*’s buffer 106 is a cache storage because the buffer includes a Random Access Memory (RAM) (*id.*, 48:9-15) allowing for temporary storage of clips in memory. This form of temporary storage is cache storage. A POSITA would have understood that caching refers to temporarily holding data. *Zivkov*, § 2. Caching is beneficial because information that is used frequently is stored in close proximity to where the information is likely to be used. *Id.* *Kenner*’s buffer 106 operates in the same way. For example, *Kenner* explains that for clips in high demand, “the PIM will attempt to move ... clips into a RAM ... buffer 106 accessible by the DSI 58” to “reduce the quantity of disk

accesses required to retrieve the highest demand clips,” consistent with a cache storage device. *Kenner*, 48:9-15.

155. *Kenner*’s IMs also include *cache storage* in a broader sense, by virtue of storing copies of content replicated from origin servers, as described for [1.a.1]. § IX.A.1.b. Consistent with the general principles of caching explained above, *Kenner*’s IM acts as *cache storage* by storing copies of clips local to a client requesting access to the clips. *Kenner*, 1:13-16; *see Zivkov*, § 2 (describing locality by space benefits of caching). For example, *Malpani* explains that in web distribution systems, caching is used to store copies of documents throughout the Web. *Malpani*, Introduction. A POSITA would have understood the primary IM acts as a Web cache for the network by storing copies of clips locally, and thus the local SRUs of an IM provide *cache storage*.

d. [40.b]: “a table listing content providers of said plurality of content providers having content that is authorized for delivery to client machines via the server”

156. In my opinion, *Kenner* discloses or at least suggests [40.b].

157. As I explained for [1.b], an IM in *Kenner*’s network maintains a clip database (*table*) storing information on the video clips stored on its extended SRUs. *Kenner*, 34:3-6, 37:32-38:3, 38:32-39:4, 41:13-19, 43:17-25; § IX.A.1.d.

158. *Kenner's* table *list[s]* content providers of said plurality of content providers because the clip database includes the address of the source IM of the content providers that uploaded the file. *Kenner*, 41:13-16.

159. *Kenner* discloses the content providers being *authorized for delivery to client machines via the server* because before delivering content to the user, the IM checks "that the user is authorized to receive the desired clip." *Id.*, 38:23-31; § IX.A.1.d.

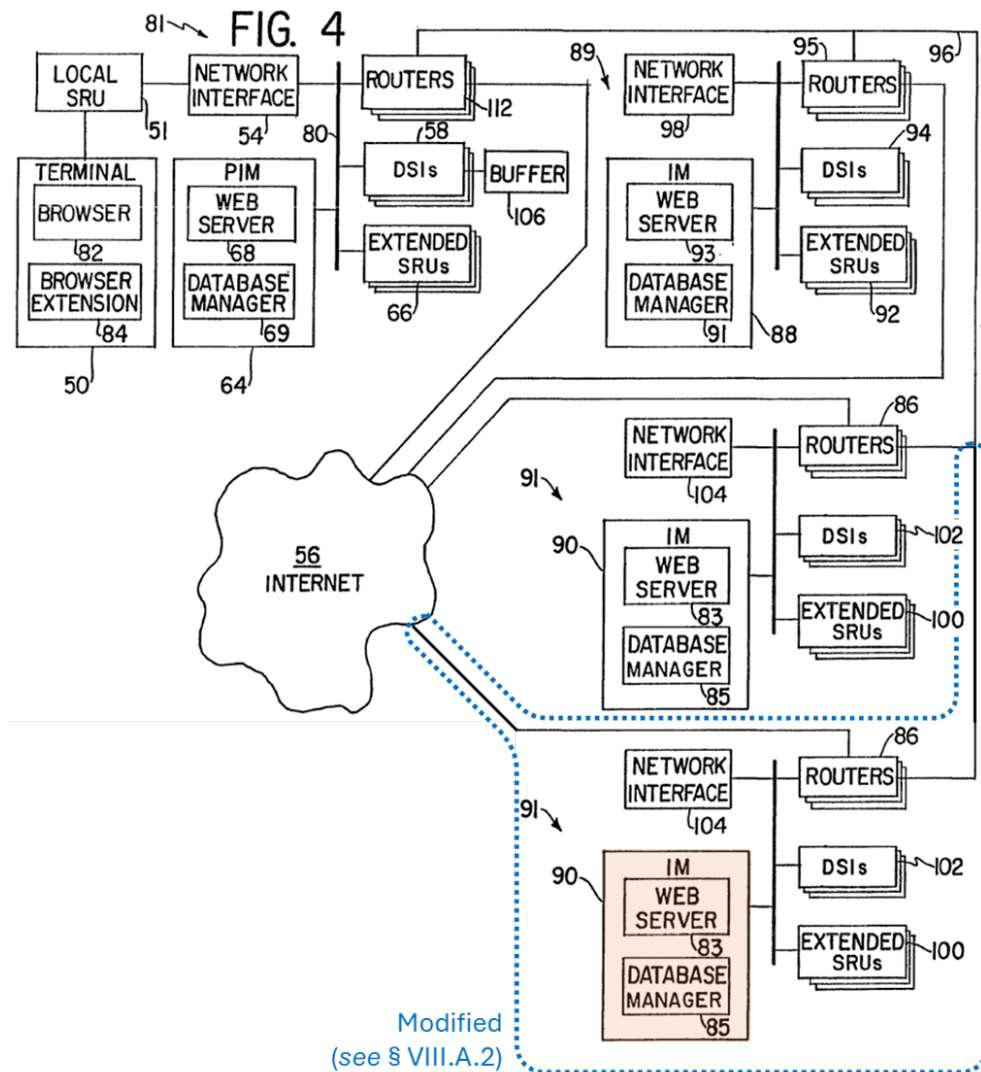
- e. **[40.c]: "wherein the server is associated with a first alias name corresponding to said first content provider"**

160. In my opinion, *Kenner* discloses [40.c]. As I explained for [1.c], *Kenner* discloses a first alias name associated with the server in the form of a first video ID. § IX.A.1.e. And as I explained for [1.b], the video IDs *correspond[] to content providers* because the video IDs include "the content provider's account number." § IX.A.1.d; *Kenner*, 36:32-37:6. Accordingly, a first video clip in the database would correspond to a first content provider.

- f. **[40.d]: "wherein the server is associated with a second alias name corresponding to said second content provider, said second alias name being distinct from the first alias name"**

161. In my opinion, *Kenner* discloses [40.d].

162. As I explained for [1.e], *Kenner* discloses a second alias name associated with the at least one repeater server. § IX.A.1.g. And as I explained for [1.b], the video IDs *correspond[] to content providers* because the video IDs include “the content provider’s account number.” § IX.A.1.d; *Kenner*, 36:32-37:6. Accordingly, a second video clip in the database corresponds to a second content provider when uploaded by a different content provider. For example, the second video clip would be uploaded by a second content provider at a **second source IM**, as shown in modified Figure 4 below:



Kenner Fig. 4 (modified as indicated).

163. Kenner discloses the *second alias name being distinct from the first alias name* because the video IDs are “unique reference identifier[s] for each video clip.” *Id.*, 22:1-2, 34:9 (the video ID is a “globally unique name of the video clip”). The system maintains a “master database of video ID numbers to ensure that each

clip uploaded, no matter from where” (e.g., from different content providers) “will have a unique video ID.” *Id.*, 43:27-29.

- g. [40.e]: “wherein the server is operable to replicate at least some of the resources associated with the first content provider, and at least some of the resources associated with the second content provider”**

164. In my opinion, *Kenner* discloses [40.e]. As I explained for [1.a.i], *Kenner*’s IM is *at least one shared repeater server constructed and adapted to replicate* resources associated with first and second content providers. §§ IX.A.1.b, IX.A.1.c. The IM (*server*) is *operable to replicate* these resources because it receives and stores clips uploaded by the content providers and thus operates to do so. *Kenner*, 44:17-18.

- h. [40.f]: “wherein, in response to a request for a resource, the server is operable to use at least the table to analyze an alias name associated with the request to determine a content provider associated with the resource”**

165. *Kenner* discloses [40.f], as I explained for [1.g]. § IX.A.1.i. For example, *Kenner* explains that the IM “contact[s] the source IM 90, where the content provider first uploaded the file.” *Kenner*, 41:13-16.

- 16. Claim 41: “wherein the table further comprises, for each of the listed content providers, information from the group consisting of: a reflector name and a committed aggregate information rate (CAIR)”**

166. *Kenner* discloses claim 41, as I explained for claim 35. § IX.A.12.

17. Independent Claim 43

- a. **[43.pre]: “A server operative in a computer network for delivering resources to multiple client machines”**

167. *Kenner* discloses [43.pre], as I explained for [40.pre.i]. § IX.A.15.a.

- b. **[43.a]: “the server comprising: ... cache storage”**

168. *Kenner* discloses [43.a], as I explained for [40.a]. § IX.A.15.c.

- c. **[43.b]: “a table listing a plurality of origin servers having content that is authorized for delivery to client machines via the server”**

169. *Kenner* discloses [43.b], as I explained for [40.b]. § IX.A.15.d.
Kenner’s clip database includes the address of the source IM where the content provider first uploaded the file and thus *list[s] ... origin servers having content that is authorized for delivery to client machines via the server. Kenner, 41:13-16; § IX.A.1.d.*

- d. **[43.c]: “wherein the server is associated with a first alias name corresponding to a first origin server of the plurality of origin servers”**

170. *Kenner* discloses [43.c], as I explained for [40.c]. § IX.A.15.e.
Kenner’s video IDs (*alias name[s]*) correspond to origin servers because the clip database includes the address of the source IM (*origin sever*). *Kenner, 41:13-16.*

- e. **[43.d]: “wherein the server is associated with a second alias name corresponding to a second origin server of the plurality of origin servers, said second origin server being distinct from the first origin server and**

said second alias name being distinct from the first alias name”

171. *Kenner* discloses [43.d], as I explained for [40.d]. § IX.A.15.f; *see* § IX.A.17.d.

- f. **[43.e]: “wherein the server is operable to replicate at least some of the resources located on the first origin server, and at least some of the resources located on the second origin server”**

172. *Kenner* discloses [43.e], as I explained for [40.e]. § IX.A.15.g.

- g. **[43.f]: “wherein, in response to a request for a resource, the server uses at least the table to analyze an alias name associated with the request to determine an origin server associated with the particular resource”**

173. *Kenner* discloses [43.f], as I explained for [40.f]. § IX.A.15.h.

18. **Claim 44: “wherein the first origin server and the second origin server are associated with distinct content providers located at distinct physical locations”**

174. In my opinion, *Kenner* discloses claim 44. As I explained for [1.a.ii], *Kenner* discloses a plurality of distinct content providers that upload content at source IMs in different regions. § IX.A.1.c.

175. *Kenner* discloses the content providers being *located at distinct physical locations* because *Kenner*’s system accommodates a “large number” of IMs in “disparate geographical areas.” *Kenner*, 40:24-28. Each of these IMs may be

associated with different content providers. *Id.*, Abstract (describing “uploading and distributing clips to geographically diverse servers”), 19:18-19.

19. Claim 45: “wherein the table further comprises, for each of the listed origin servers, information from the group consisting of: a reflector name and a committed aggregate information rate (CAIR)”

176. *Kenner* discloses or at least suggests claim 41, as I explained for claim 35. § IX.A.12.

B. Ground 2: The Combination of *Kenner* and *Vetter* Renders Obvious Claim 33

1. Rationale for Combining *Kenner* and *Vetter*

177. A POSITA would have been motivated to implement *Kenner*’s system to deliver FTP-based content, as disclosed by *Vetter*.

178. *Kenner* and *Vetter* both relate to delivering content via a web browser. *Kenner*, 2:16-21, 31:10-18; *Vetter*, 49. Like *Kenner*, *Vetter* specifically addresses delivering “full-motion video” content. *Kenner*, 3:1-2, 8:26-9:4, 31:10-18; *Vetter*, 49. And both describe URLs for delivering HTML-based content. *Kenner*, 2:16-21, 37; *Vetter*, 52.

179. *Kenner* describes delivering HTTP-based content, but does not explicitly disclose other protocols. *Kenner*, 2:16-21, 37:7-17. *Vetter* on the other hand describes several others, including FTP. *Vetter*, 52. A POSITA would have

been motivated to include FTP-based resources as disclosed in *Vetter* to expand the types of resources available through *Kenner*'s system.

180. Combining the teachings of *Kenner* and *Vetter* would have been straightforward. *Kenner* describes URLs for delivering content under an “https://” scheme. *Kenner*, 2:16-21, 37:7-17. As *Vetter* explains, this scheme can easily be swapped out for other schemes, such as FTP, within a URL. *Vetter*, 52. This would allow *Kenner*'s network to deliver content from additional server types—namely, FTP servers. Applying *Vetter*'s teaching to *Kenner* would have been a straightforward use of a known technique to enhance a similar system in the same way.

2. Claim 33: “wherein the resources comprise a File Transfer Protocol (FTP)-based resource”

181. In my opinion, *Kenner* in view of *Vetter* teaches claim 33. *Kenner* describes delivering resources in HTML format using a uniform resource locator (URL). *Kenner*, 2:16-21, 37:7-17, 50:22-31. For example, *Kenner* explains that the precise location of an HTML file delivered in its system is provided in the format ‘http://internet.address/directory/filename.html.’” *Id.*, 2:16-21. In this example, the file is an http-based resource based on the scheme identified in the URL, as indicated by the “http” scheme. *Vetter*, 52. A POSITA would have recognized that *Kenner*'s system could equally be used for other protocols including ftp-based resources.

Kenner, 14:5-7 (using “conventional network protocols and topologies”); *Berners-Lee*, 78, 81.

182. Nonetheless, *Vetter* explicitly describes “access[ing] files using FTP” which would be ftp-based resources. *Vetter*, 49. Indeed, *Vetter* describes the same URL format “scheme://host.domain[:port]/path/filename” where “[t]he first part of the URL, the scheme, specifies the access method.” *Id.*, 52. In addition to *Kenner*’s exemplary HTTP-based scheme, *Vetter* also identifies an FTP-based scheme that “retrieves a file on [a] local system or on an anonymous FTP server.” *Id.* Accordingly, a POSITA would have easily used this FTP scheme as another form of HTML file delivered through its system. A POSITA would have been motivated to include FTP-based resources as disclosed in *Vetter* to accommodate a wide range of files and storage locations. § IX.B.1. Moreover, the ’903 patent itself admits the FTP scheme was known in the art. EX-1001, 6:22-41.

C. Ground 3: The Combination of *Kenner* and *Rekimoto* Renders Obvious Claim 34

1. Rationale for Combining *Kenner* and *Rekimoto*

183. A POSITA would have been motivated to implement *Rekimoto*’s DNS capabilities in *Kenner*’s system to reduce system latency, which would include adding *Rekimoto*’s DNS server.

184. *Kenner* and *Rekimoto* are both concerned with delivering content via a network of servers. *Kenner*, 8:3-15; *Rekimoto*, 17:46-49. Like *Kenner*, *Rekimoto* describes delivering content based on requests from users at end terminals. *Kenner*, 13:2-10; *Rekimoto*, 16:8-13. And both systems are adapted to handle HTML-based requests via URLs to deliver content via a browser. *Kenner*, 31:10-18, 37:7-17; *Rekimoto*, 16:45-57.

185. Like *Rekimoto*'s, *Kenner*'s network is designed to deliver content between servers across the world. *Kenner*, 49:1-4, 54:12-16. A POSITA would have been motivated to include *Rekimoto*'s DNS capabilities within *Kenner*'s system to achieve the advantages *Rekimoto* advertises: to reduce a latency in the network. *Rekimoto*, 17:46-49. *Kenner* shares this goal. *Kenner*, 5:21-26, 7:25-31.

186. Combining *Rekimoto*'s DNS with *Kenner*'s system would have been straightforward. *Kenner* acknowledges the advantages of prioritizing servers based on proximity. *Id.*, 40:29-33. A POSITA would have easily incorporated DNS capabilities, as disclosed by *Rekimoto*, into *Kenner*'s network to select a closest server for handling a client request. Indeed, DNS was well-known at the time. *Tanenbaum*, 622-630. Applying *Rekimoto*'s teaching to *Kenner* would have been a straightforward use of a known technique to enhance a similar system in the same way.

2. Claim 34

- a. **[34.a]: “wherein at least the first alias name is a domain name and the computer network further comprises at least one name server that provides domain name service (DNS) resolution”**

187. In my opinion, *Kenner* in view of *Rekimoto* teaches [34.a]. As the '903 patent admits, DNS resolution was known at the time of filing. EX-1001, 7:4-29. When accessing an HTML file specified by the URL “http://www.uspto.gov/A/B/C/F,” a browser would access a DNS server to look up the network IP address of a host (origin server) and establish a connection with that server. *Id.*, 6:22-61, 7:4-29 (describing “conventional[]” process); 2:4-8. In the '903 patent, a reflector “uses DNS to identify a set of candidate repeaters given a domain name that represents the repeater network” and “tries each repeater in this set until it makes a successful contact.” *Id.*, 21:56-64.

188. *Kenner* discloses similar techniques. *Kenner*'s system delivers HTML-based resources in the same format: “http://internet.address/directory/filename.html.” *Kenner*, 2:16-21. When a user requests access to a clip, the browser provides a virtual URL to the IM, which includes the video ID of the clip. *Id.*, 37:7-17. If the clip is not found locally, the IM identifies a set of IMs likely to have the clip and queries the closest IMs. *Id.*, 40:29-41:6. The IM thus maps the URL it receives to the IP addresses of its own

SRUs that may store the clip (*id.*, 34:10) or the SRUs of neighboring IMs (*id.*, 40:29-33, 42:23-36). *Kenner*'s IM thus provides the same functionality as the DNS resolution described in the '903 patent and recited in claim 34.

189. To the extent *Kenner* does not expressly disclose “a domain name” or “at least one name server that provides domain name service (DNS) resolution,” *Rekimoto* does. *Rekimoto*'s system includes a domain name system (DNS) server 130 used to “answer[] a domain name for a corresponding IP address and vice versa.” *Rekimoto*, 18:3-12. The DNS server maintains a “table listing the relationship between the source IP addresses of the requesting client terminal and the domain names assigned with the IP addresses for the corresponding domain name.” *Id.*, 18:15-19.

190. It would have been obvious to incorporate *Rekimoto*'s DNS server and DNS resolution capabilities into *Kenner*'s network to improve the response time for delivering resources. As *Kenner* explains, its network includes IMs distributed across different geographic regions. *Kenner*, 49:1-4, 54:12-16, Fig. 4. For example, *Kenner*'s clip database includes the IP addresses for its own local SRUs (*id.*, 34:10), but queries neighboring IMs to identify remote SRUs storing the clip (*id.*, 40:29-33). In the combined *Kenner-Rekimoto* system, when the clip is not available on the local SRUs, the IM will direct the client request to the DNS server described in *Rekimoto*

to identify the closest IMs (and associated remote SRUs) for handling the request. For example, the DNS server would map the URL containing the video ID to the remote IMs in the same manner described in the '903 patent. EX-1001, 21:56-64. The '903 patent admits DNS was known in the art. *Id.*, 6:45-61, 7:4-15.

191. In the combined system, a POSITA would have recognized that *Kenner's* video ID (alias name) may be used as a domain name for resolution through *Rekimoto's* DNS server. For example, *Rekimoto* describes using a domain name in the format “hanaya@ipd.sony.co.jp” for example, which denotes a user name, a host name, an organization name, an organization attribute, and country name.” *Rekimoto*, 18:3-12. *Kenner's* video ID includes similar information, such as a “provider's account number ... [,] a category coordinate, ... [and] a geographic coordinate used to determine where the file is relevant.” *Kenner*, 36:32-37:6. And *Kenner* describes formatting its video ID as a virtual URL including the internet address of the primary IM. *Id.*, 37:7-17. Accordingly, a POSITA would have found it obvious to structure *Kenner's* video ID as a domain name to provide the DNS resolution advantages described above.

- b. **[34.b]: “the method further comprising: receiving, at the at least one name server, the first alias name as part of a client request; and resolving, at the at least one name server, the received first alias name to identify the at least one shared repeater server to**

deliver the first resource in response to the client request”

192. In my opinion, *Kenner* in view of *Rekimoto* teaches [34.b]. As I explained for [1.g], *Kenner* discloses receiving an alias name as part of a client request for a particular resource. § IX.A.1.i; *Kenner*, 36:25-27 (user terminal “passes the video ID to the PIM 64 associated with the user's terminal.”). Under the combined *Kenner-Rekimoto* system, this functionality would be preserved, and a user would thus submit a request via the user terminal to be resolved by *Rekimoto*'s DNS server. *Kenner*, 25:7-24; *Rekimoto*, 18:3-12.

193. *Kenner-Rekimoto* teaches *receiving at the at least one name server, the first alias name as part of a client request* because in the combined system, *Kenner*'s video ID, formatted as a virtual URL (§ IX.C.2.a), would be provided to the DNS server 130, as described in *Rekimoto* (*Rekimoto*, 18:3-12).

194. *Kenner-Rekimoto* teaches *resolving, at the at least one name server, the received first alias name to identify the at least one shared repeater server to deliver the first resource in response to the client request* because, in combination, the DNS server would then use the domain name included in the request to identify an IP address of an IM for responding to the request. *Id.*, 18: 3-19. A POSITA would have been motivated to incorporate the DNS server into *Kenner* to reduce the response time for responding to the requests. § IX.C.2.a.

D. Ground 4: The Combination of *Kenner* and *Boyles* Renders Obvious Claims 27, 36, 39, 42, and 46

1. Rationale for Combining *Kenner* and *Boyles*

195. A POSITA would have been motivated to implement *Boyles*'s negative reply in *Kenner*'s system.

196. *Kenner* and *Boyles* are both concerned with delivering content via a network of servers. *Kenner*, 8:3-15; *Boyles*, 1:7-10. Like *Kenner*, *Boyles* describes delivering content based on requests from users at end terminals. *Kenner*, 13:2-10; *Boyles*, 4:26-41. And both networks handle requests in a similar way, searching for content locally and then expanding to additional servers. *Kenner*, 25:25-26:7, 40:29-33; *Boyles*, 2:55-3:3.

197. *Kenner* describes transmitting a "reason for the denial" to a user, although not explicitly described in response to a clip being absent from the clip database. *Kenner*, 38:17-20. However, *Boyles* expressly describes returning a "negative reply" in this scenario. *Boyles*, 8:58-65. A POSITA would have been motivated to include this form of response to improve user experience. Indeed, *Kenner* acknowledges the need to "advise the user" why a request may fail (*Kenner*, 38:17-20) and *Boyles* describes this form of notification in greater detail.

198. Implementing *Boyles*'s express failure notification in *Kenner* would have been straightforward for a POSITA. *Kenner* describes the same form of user

terminals for end user input/output described in *Boyles* (compare *Kenner*, 13:2-10 with *Boyles*, 4:26-42) and already includes mechanisms for generating user notifications in response to a request (*Kenner*, 38:17-20). To the extent not already included in *Kenner*, a POSITA would have easily displayed the “negative replies” described in *Boyles* to provide a more intuitive user experience. Implementing *Boyles*’s “negative replies” in *Kenner*, which is already designed to provide user notifications, would have been a straightforward use of a known technique to enhance a similar system in the same way.

2. Claim 27: “wherein the at least one repeater server is configured and adapted to reject the client request for the particular resource if the alias name received with the client request is not included the table”

199. In my opinion, *Kenner* in view of *Boyles* teaches claim 27. *Kenner*’s IM (e.g., the primary IM) receives a request for a clip along with associated video ID for a clip. *Kenner*, 36:25-27, 37:1-7. The IM then queries its clip database (table) to determine if any of its SRUs contain the desired clip. *Id.*, 38:32-39:4. If the IM itself does not store the clip, it queries neighboring remote IMs. *Id.*, 41:1-5. And if still unable to locate the clip, the IM uses the clip database to contact the source IM. *Id.*, 44:13-16.

200. *Kenner* at least suggests that if the video ID is not in the database, it will reject the client request. For example, *Kenner* explains that “[i]f any of the

foregoing database checks fail, ... the reason for the denial will be transmitted from ... to the ... user interface via the local SRU 51 to advise the user” (i.e., a rejection). *Id.*, 38:17-20. One of these “foregoing database checks” includes consulting the clip database to identify SRUs containing the clip (*id.*, 37:18-38:16) and a POSITA would thus understand that failing to identify the video ID in the database would trigger this rejection. In that case, the “reason for the denial” would be failure to locate the clip. *Id.*, 38:17-22.

201. To the extent Sandpiper argues further disclosure is required, *Boyles* discloses claim 38. When a target resource in *Boyles*’s network “is either not known or not available in the network[,] ... [a] negative reply is returned ... to the requesting node.” *Id.*, 8:58-65. When incorporated into *Kenner*’s network, this functionality would result in a rejection of the client request if the clip is not found in the clip database or in the clip database of other IMs in the system. And a POSITA would have been motivated to include this functionality to provide an indication to the user of why the clip is not returned. § IX.D.1.

3. Claim 36: “further comprising: rejecting the client request for the particular resource if the alias name received with the client request is not included [in] the table”

202. *Kenner* discloses claim 36, as I explained for claim 27. § IX.D.2.

-
4. **Claim 39: “further comprising: rejecting the client request for the resource if the alias name received with the client request is not included [in] the table”**

203. *Kenner* discloses claim 39, as I explained for claim 27. § IX.D.2.

5. **Claim 42: “wherein in response to a request for a resource, the server is further configured and adapted to reject the request for the resource if the alias name received with the request is not included [in] the table”**

204. *Kenner* discloses claim 42, as I explained for claim 27. § IX.D.2.

6. **Claim 46: “wherein in response to a request for a resource, the server is further configured and adapted to reject the request for the resource if the alias name received with the request is not included [in] the table”**

205. *Kenner* discloses claim 46, as I explained for claim 27. § IX.D.2.


X. CONCLUSION

206. For the reasons set forth above, in my opinion, a person of ordinary skill in the art would have found claims 1, 2, 22-24, and 26-46 of the '903 patent obvious.

207. In signing this declaration, I understand that the declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I acknowledge that I may be subject to cross-examination in this case and that cross-examination will take place within the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

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

Dated: May 8, 2025

By: 
Todd C. Mowry, Ph.D.