UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

MICROSOFT CORPORATION,
Petitioner,

v.

SANDPIPER CDN, LLC,
Patent Owner.

Case IPR2026-00095
Patent No. 8,478,903

PETITION FOR *INTER PARTES REVIEW*
OF U.S. PATENT NO. 8,478,903

TABLE OF CONTENTS

TABLE OF AUTHORITIES

CASES

STATUTES

LIST OF EXHIBITS

Exhibit	Description
1001	U.S. Patent No. 8,478,903 to Farber et al. ("'903 Patent")[11]
1002	Prosecution History of U.S. Application No. 11/065,412
1003	Declaration of Dr. Todd Mowry ("Mowry")
1004	Curriculum Vitae of Dr. Todd Mowry
1005	PCT International Patent Pub. No. WO 1996041285 to Kenner et al. ("*Kenner*")
1006	Ronald J. Vetter et. al, Mosaic and the World-Wide Web, Computer, vol. 27, no. 10, pp. 49-57, Oct. 1994 ("*Vetter*")
1007	European Patent Pub. No. EP 0753836 to Rekimoto et al. ("*Rekimoto*")
1008	U.S. Patent No. 5,511,208 to Boyles et al. ("*Boyles*")
1009	Andrew Tanenbaum, Computer Networks (3rd ed. 1996) ("*Tanenbaum*")
1010	Tim Berners-Lee, The World-Wide Web, Communications of the ACM, Vol. 37, No. 8, August 1994 ("*Berners-Lee*")
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 ("*Malpani*")

[1] 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.

1012	Fielding et al., Hypertext Transfer Protocol – HTTP/1.1, RFC 2068, Jan. 1997, https://www.rfc-editor.org/rfc/rfc2068 ("*HTTP/1.1*")
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 ("*Zivkov*")
1014	Declaration of Lauren Gluckman Regarding EX-1006
1015	Scheduling Order, *Sandpiper CDN, LLC v. Google LLC*, 2:24-cv-03951, Dkt. 62 (C.D. Cal. Dec. 6, 2024)
1016	Letter to Robert H. Reckers re Sandpiper CDN LLC v. Google LLC, No. 2:24-3951-AB-RAO (C.D. Cal.) (May 6, 2025)
1017	*Sandpiper CDN, LLC v. Microsoft Corp.*, No. 2-25-cv-00664 (E.D. Tex.), Docket

*All emphasis is added unless otherwise indicated.

II. INTRODUCTION

Microsoft Corp. ("Petitioner") requests review and cancellation of claims 1, 2, 22-24, and 26-46 of U.S. Patent No. 8,478,903 (the "'903 patent") (EX-1001). The patentability analysis of this Petition is substantively equivalent to the petition in IPR2025-00969, and Petitioner conditionally requests joinder to that proceeding. *See* Paper 3. More specifically, Petitioner is not time-barred from filing this Petition. Thus, at the time the Director considers institution of this Petition, Petitioner requests institution and joinder to IPR2025-00969 if that IPR has been instituted and remains pending. If that IPR has been terminated (e.g., based on settlement), Petitioner requests the Director consider this Petition on its own without joinder.

The '903 patent claims a content delivery system for delivering resources from content providers to requesting clients. EX-1001, Abstract; claim 1. The '903 patent acknowledges that serving resources to client devices in a network was known. *Id.*, 7:4-38; *see Malpani*, 107; *Rekimoto*, 17:40-49. The alleged improvement—delivering content from shared repeater servers (EX-1001, 4:60-62; 7:40-49)—was also known or at least rendered obvious in prior art networks. Claims 1, 2, 22-24, and 26-46 are unpatentable.

III. STATEMENT OF PRECISE RELIEF REQUESTED

Petitioner requests review under 35 U.S.C. §311 of claims 1, 2, 22-24, and 26-46 of the '903 patent and their cancellation in view of the following:

Prior Art References	
Ref. 1:	*Kenner* (EX-1005), published December 19, 1996, filed on June 7, 1996.
Ref. 2:	*Vetter* (EX-1006), published in October 1994.
Ref. 3:	*Rekimoto* (EX-1007), published January 15, 1997, filed on July 12, 1996.
Ref. 4:	*Boyles* (EX-1008) published April 23, 1996, filed on March 23, 1993.

Each of the above references is prior art under at least 35 U.S.C. §102(b).

Grounds of Unpatentability	
1	Claims 1, 2, 22-24, 26, 28-32, 35, 37, 38, 40, 41, and 43-45 are obvious under 35 U.S.C. §103 over *Kenner*.
2	Claim 33 is obvious under 35 U.S.C. §103 over *Kenner* and *Vetter*.
3	Claim 34 is obvious under 35 U.S.C. §103 over *Kenner* and *Rekimoto*.
4	Claims 27, 36, 39, 42, and 46 are obvious under 35 U.S.C. §103 over *Kenner* and *Boyles*.

IV. THE '903 PATENT

A. Overview

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; Mowry, ¶40. As shown in Fig. 1 (annotated below), the '903 patent's network includes an **origin server**, a **client**, and various **shared repeater servers**:

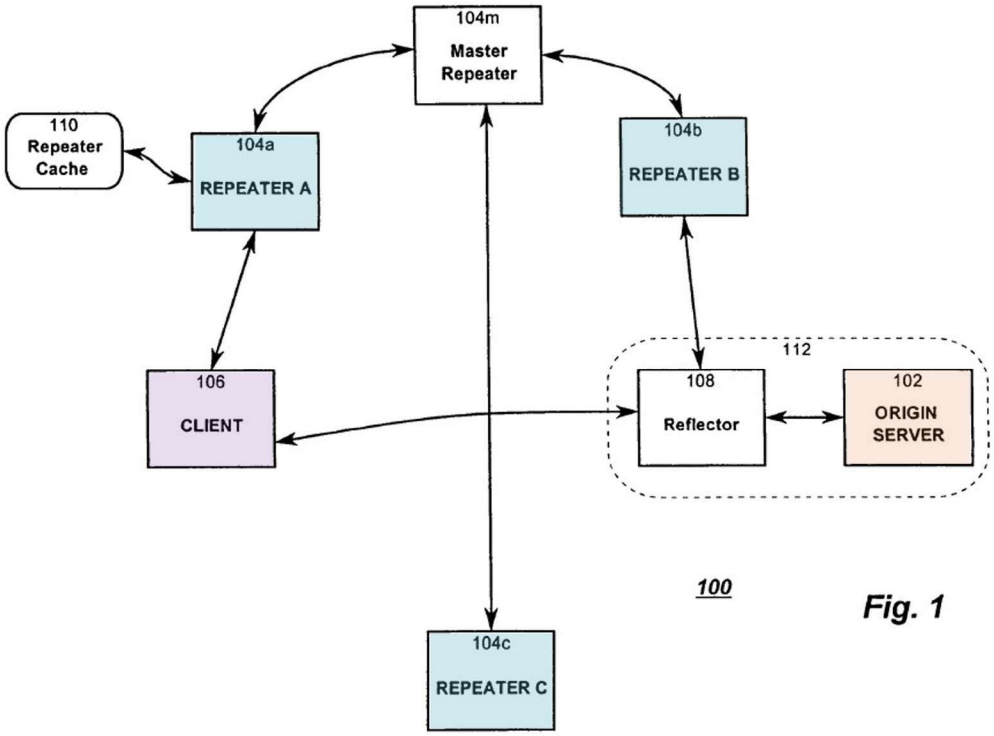

Fig. 1

EX-1001, 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; 5:9-23. The reflector selects either a repeater server or the origin server to serve the request. *Id.*, 5:9-23, 7:50-56.

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.*, 9:56-63, 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; Mowry, ¶ 50.

The purported advantages of the '903 patent arise from the reflector and repeater servers. EX-1001, 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. Mowry, ¶42.

B. **Prosecution History**

The '903 patent was filed on February 23, 2005, claiming priority to February 10, 1998. EX-1001, cover. The '903 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; Mowry, ¶43. A final rejection was issued, indicating claims 53 and 54 included allowable subject matter. EX-1002, 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; Mowry, ¶43.

V. LEVEL OF ORDINARY SKILL

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. Mowry, ¶47.

VI. CLAIM CONSTRUCTION

The Board construes claims under *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc). 37 C.F.R. §42.100(b) (2018). Under this standard, terms receive their plain and ordinary meaning as understood by one of ordinary skill in the art, consistent with the disclosure and prosecution history. *Phillips*, 415 F.3d at 1314–19. Claims should only be construed to the extent necessary to resolve a controversy. *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017). No claim terms need to be construed by the Board at this time, and all should be given their ordinary meanings.

VII. Claims 1, 2, 22-24, and 26-46 Are Unpatentable Over the Prior Art

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

1. Overview of Ground 1

a. *Kenner*'s system

Like the '903 patent, *Kenner* deals with content delivery networks. Mowry, ¶50. It is thus analogous art. *In re Bigio*, 381 F.3d 1320, 1325 (Fed. Cir. 2004).

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." *Kenner*, Abstract. *Kenner*'s system includes a network of database index managers (IMs), including a **source IM**, a **user terminal**, and various **regional IMs**:

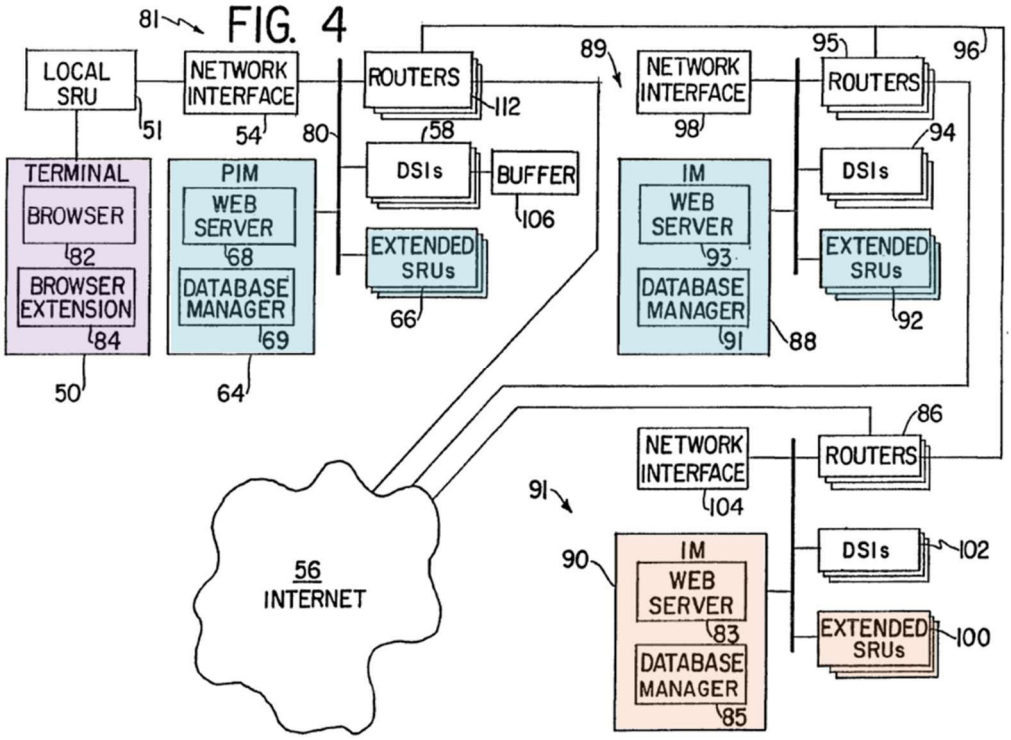

Kenner, Fig. 4. As shown in Fig. 4, each IM hosts a server and has associated extended storage and retrieval units (SRUs) for storing resources. *Id.*, 10:18-31, 31:19-21, 31:30-32, 32:14-17, 44:6-14. Though shown separately, the extended

SRUs 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.

In *Kenner*, a video clip is first uploaded by a content provider using one of the IMs. *Id*., 44:1-5; Mowry, ¶51. For a particular video clip, the originating server is referred to as the "source IM 90." *Kenner*, 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.

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). Mowry, ¶52. 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. *Kenner*, 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.

b. *Kenner* accommodates multiple users and content providers

Kenner's system is scalable and thus accommodates many user terminals and content providers. *Kenner*, 54:12-16 (describing scalability), 4:32-5:5 (contrasting

other non-scalable systems); Mowry, ¶53. For example, *Kenner* explains its system accommodates "multiple index managers for a large number of concurrent users in disparate geographical areas" (*Kenner*, 40:24-28) and supports "thousands of simultaneous users" (*id.*, 5:21-26).

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

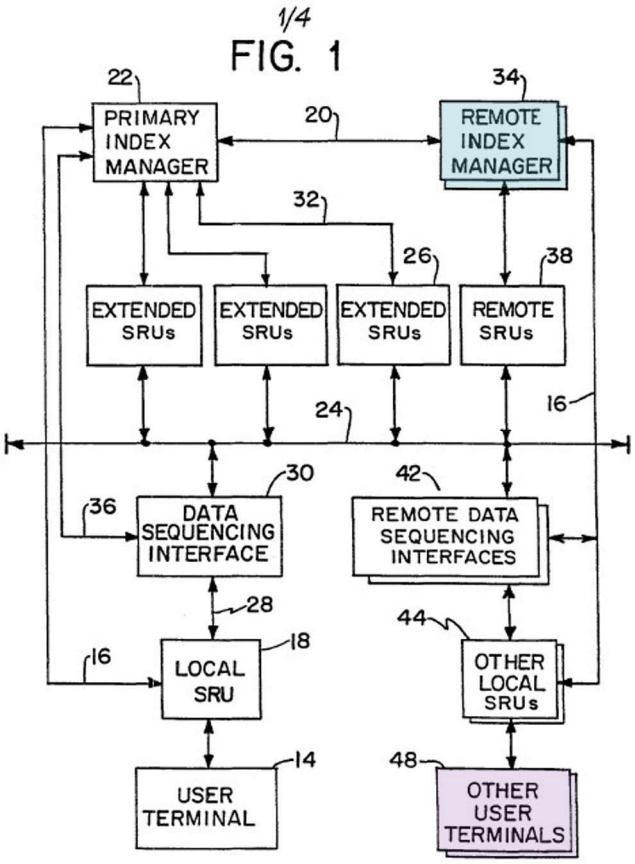

Kenner, Fig. 1, 28:3-6. Each of these remote IMs has its own associated user terminal. *Id.*, claims 45, 46. Accordingly, any of the remote IMs may be a "primary"

IM depending on where a user request is received. *Id.* In other words, there is no functional difference between regional IMs, and reference to a primary IM is in the context of a particular user terminal. Mowry, ¶54. Accordingly, the **IM for a neighboring region** may equally include an associated **user terminal**, as shown in modified Figure 4 below:

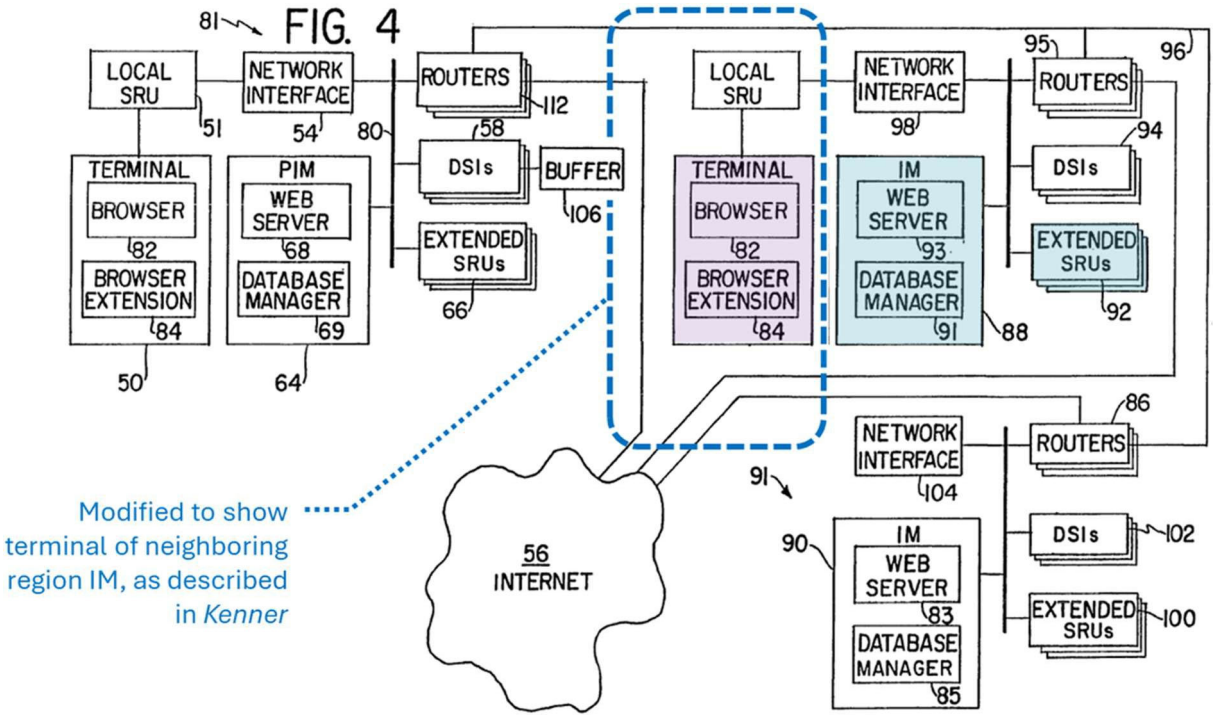

Id., ¶55 (modified *Kenner* Fig. 4).

Kenner's system also accommodates multiple content providers and source IMs. *Id.*, ¶56. *Kenner* explains a primary IM connects to "many geographically distributed video data sources." *Kenner*, 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).

Accordingly, *Kenner* is not limited to a single source IM and can include an **additional source IM**, as shown in further modified Figure 4 below:

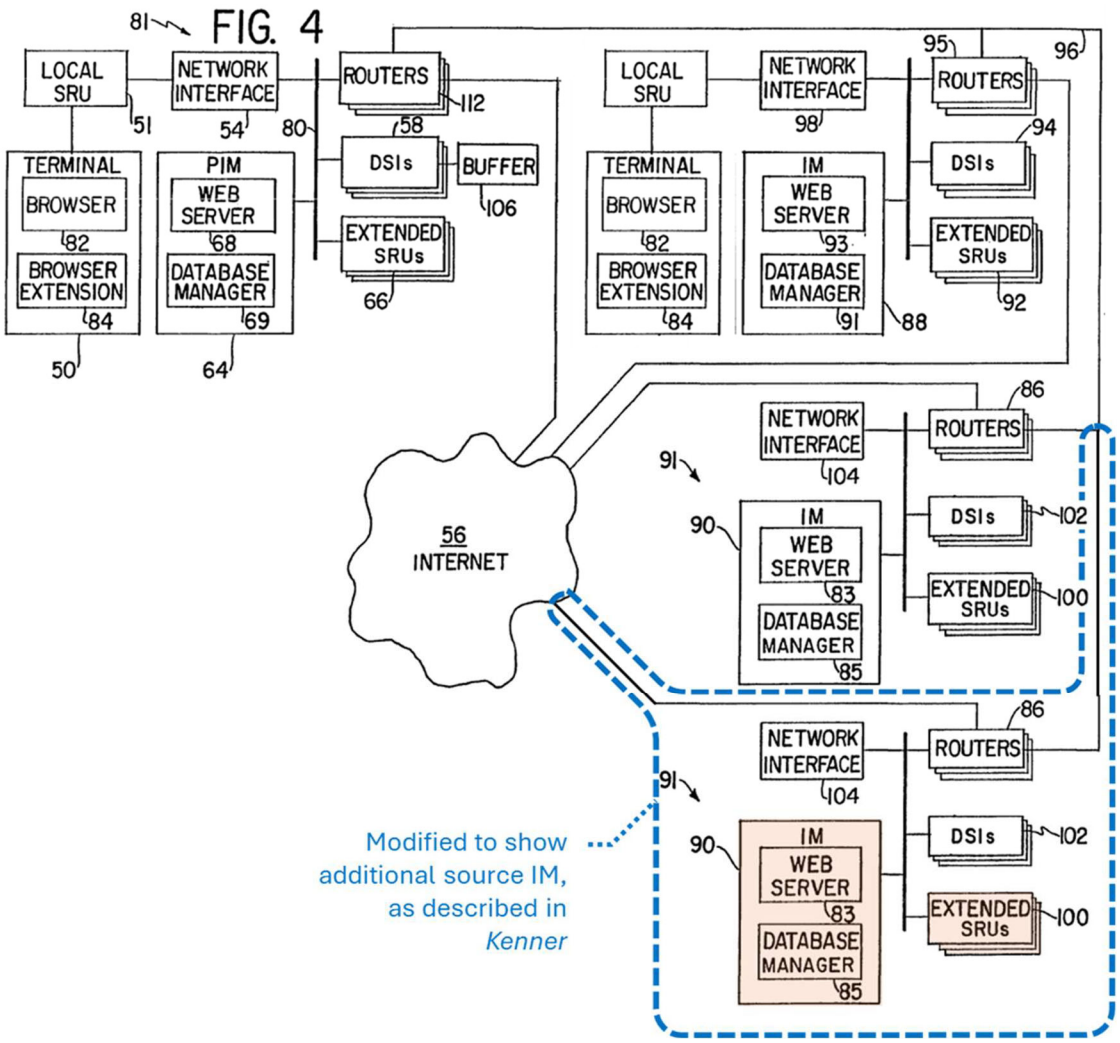

Mowry, ¶56 (modified *Kenner* Fig. 4). A client terminal may therefore be used to access content uploaded from multiple source IMs. *Id.*

With these multiple source IMs, *Kenner* also accommodates multiple content providers. *Id.*, ¶57; *Kenner*, 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 content providers, depending on which content

providers uploaded the clips. Mowry, ¶57. For example, *Kenner* envisions a subscription service, where a user must subscribe to a particular content provider to view that provider's content. *Kenner*, 33:4-11, 35:11-16, 43:17-25. Accordingly, a user may subscribe to one content provider and not another. Mowry, ¶57. And each content provider may be associated with a different source IM. *Kenner*, 45:25-31 (describing system instruction sent to source IM of a particular content provider).

Kenner is thus 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. Mowry, ¶58.

2. 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"

Kenner discloses [1.pre]. Mowry, ¶¶66-68.

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 (describing "clips" uploaded by "content providers"); §VI.A.1.b. *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:

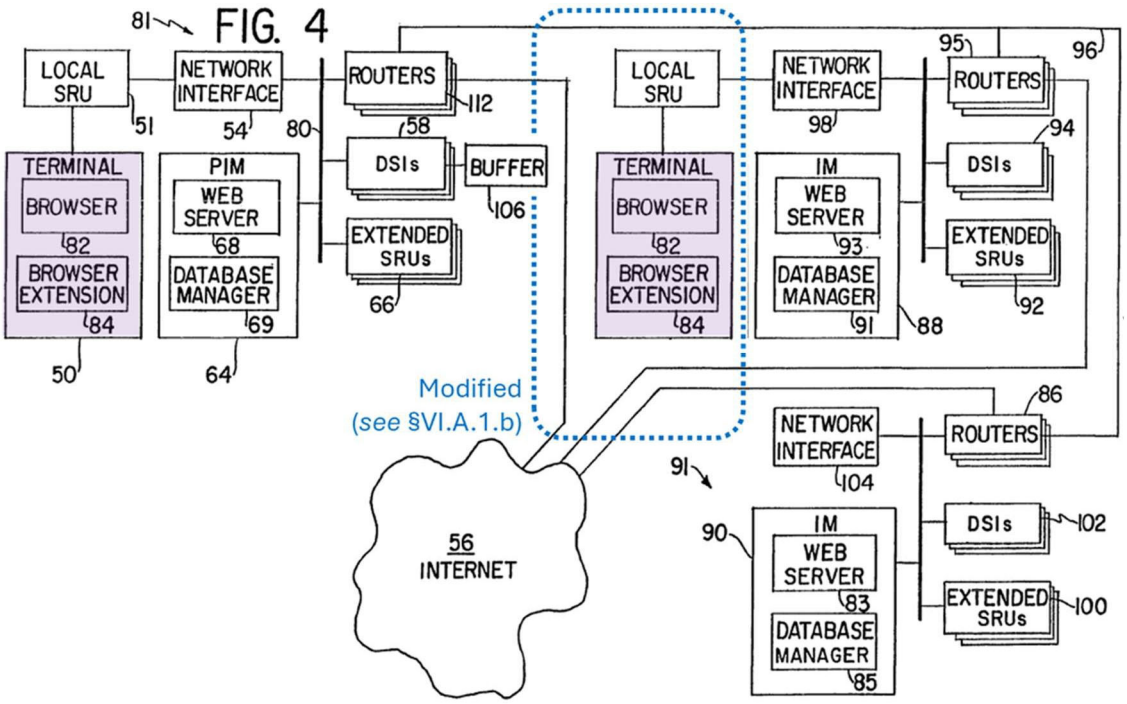

Mowry, ¶67 (modified *Kenner* Fig. 4); *Kenner*, 5:21-26.

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

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"

Kenner discloses [1.a.i]. Mowry, ¶¶69-72.

As explained, *Kenner* discloses a network of IMs. §VI.A.2.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; §VI.A.1.a; Mowry, ¶69.

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:

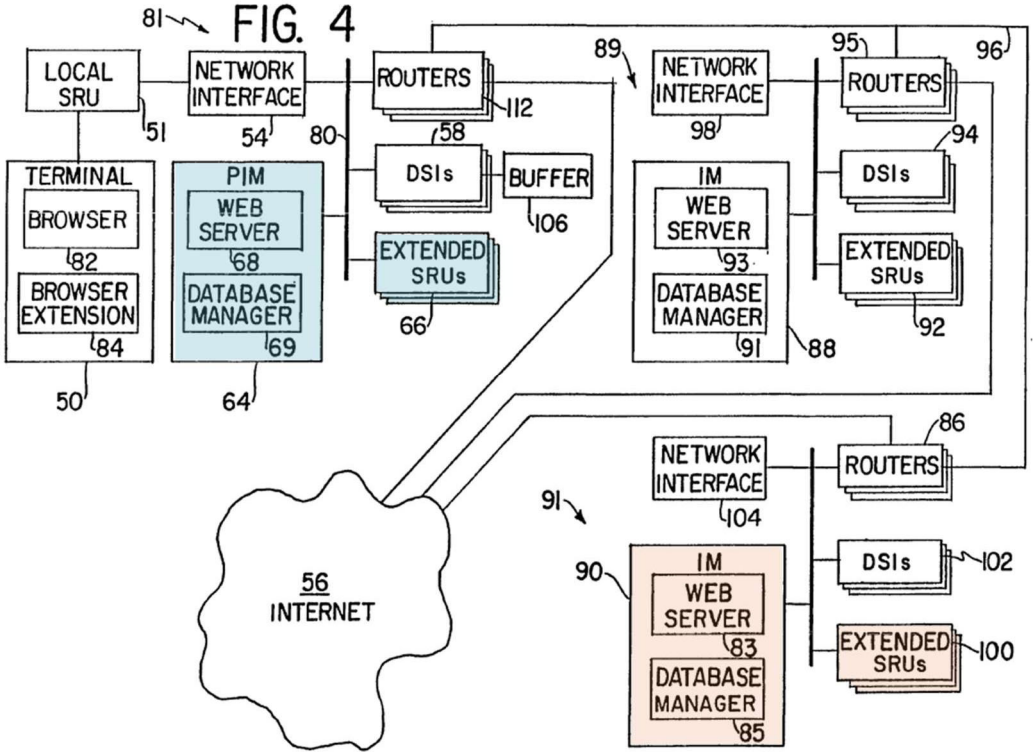

Id., Fig. 4, 41:13-16, 44:1-5. *Kenner*'s primary IM is a *repeater* server because it replicates the resources of the source IM. Mowry, ¶70; EX-1001, 4:27-32. For example, once a file is uploaded, the source IM distributes the file to other IMs throughout the network. *Kenner*, 44:15-26. The IM is *shared* because it is used for

delivering resources from multiple source IMs (and associated content providers). Mowry, ¶71; §VI.A.1.b; *see* EX-1001, 23:38-41.

Kenner's clips are *associated with a first content provider* by virtue of being uploaded by the content provider. Mowry, ¶72. 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*.; Mowry, ¶72.

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"

Kenner discloses [1.a.ii]. Mowry, ¶¶73-76.

As 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. §VI.A.2.b. *Kenner*'s primary IM equally replicates resources from other content providers, each of which may be associated with respective source IMs. Mowry, ¶73; §VI.A.1.b; *Kenner*, 19:18-19 ("connect[ing] many geographically distributed video data sources to one subscriber destination").

For example, the primary IM maintains a clip database storing information for each of the clips included in its extended SRUs. *Id*., 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. Mowry, ¶74.

Because the primary IM may store 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*. Mowry, ¶75. For example, *Kenner*'s network includes an **additional source IM**, as shown in modified Figure 4 below:

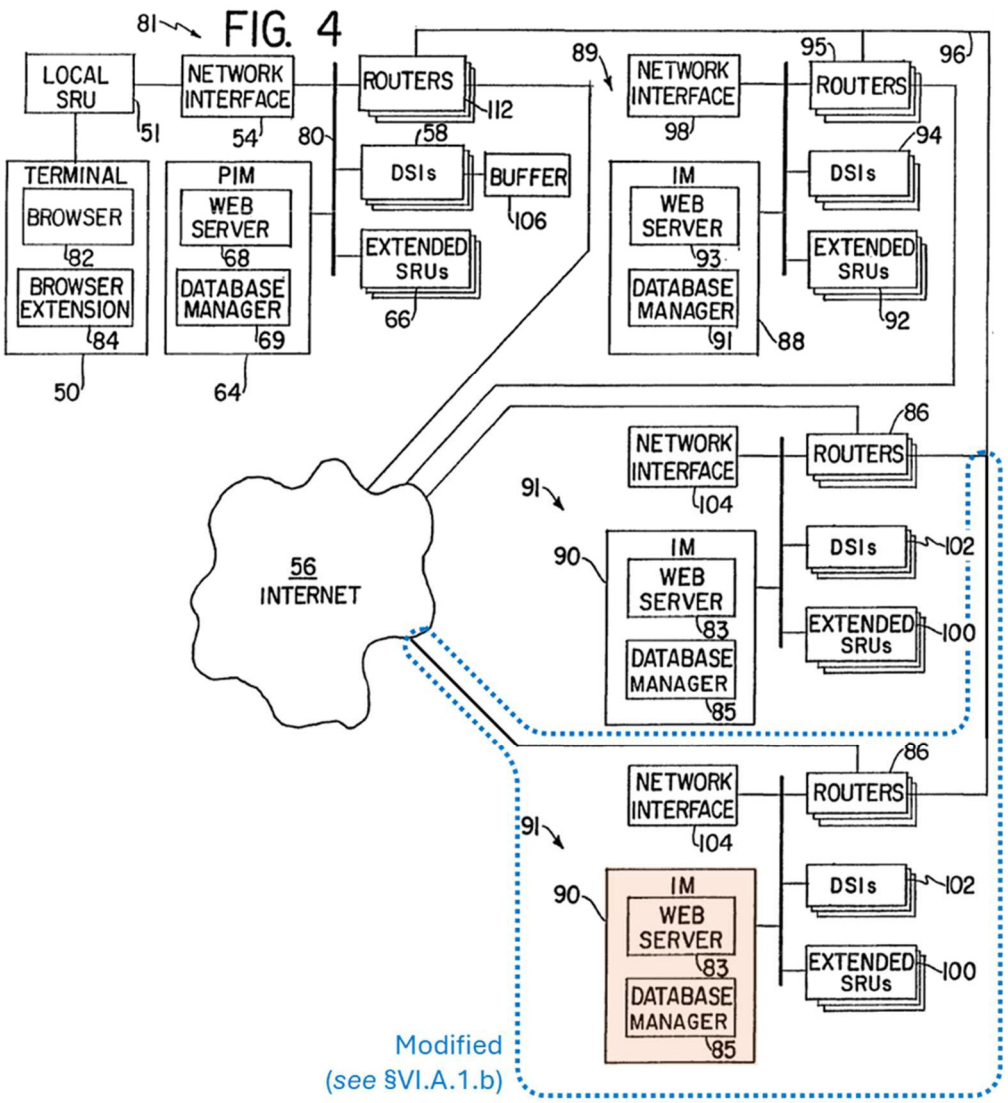

Mowry, ¶75 (modified *Kenner* Fig. 4); §VI.A.1.b.

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. Mowry, ¶76; *Kenner*, 43:17-25 (referring to "content providers"); *see* §VI.A.1.b.

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"

Kenner discloses or at least suggests [1.b]. Mowry, ¶¶77-82.

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. Mowry, ¶77; *Kenner*, 34:3-35:3. Storing data in this manner, with rows representing each clip and columns representing corresponding clip information, constitutes a table. Mowry, ¶78. *Kenner* confirms this format, indicating the database may be stored as a "flat-file database." *Kenner*, 20:21-27, 21:30-22:13 (describing database format); Mowry, ¶78. 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 (*Kenner*, 21:30-22:13, 34:3-35:3) and tables are a fundamental data structure for organizing data reflecting these relationships. Mowry, ¶78.

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."

Kenner, 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.*, Mowry, ¶79.

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. Mowry, ¶80. 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." *Kenner*, 36:32-37:6.

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. The '903 patent discloses a similar authorization. EX-1001, 23:17-20 ("Each new request for the resource must be tested at the origin server to assure that the requester is authorized to access the resource."); Mowry, ¶81.

Kenner discloses *the plurality of content providers compris[ing] the first content provider and the second content provider* as explained for [1.a.i] and [1.a.ii]. §§VI.A.2.b, VI.A.2.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; Mowry, ¶82.

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"

Kenner discloses [1.c]. Mowry, ¶¶83-87.

Kenner discloses *at least a first resource associated with the first content provider* because each clip (*resource*) is uploaded by a content provider. *Id.*, ¶83; *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.

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. Mowry, ¶84. In uploading a clip, the source IM registers the clip along with a video ID (*alias name*). *Kenner*, 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").

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. *Kenner*, 34:3-6 (database of clips stored on

extended SRUs), 41:1-6 (database indicating "which IMs are likely to have the desired clip"); Mowry, ¶85. Indeed, not every IM receives every clip. *Kenner*, 44:15-20 (updates sent only to IMs storing a particular clip). Accordingly, an IM is associated with a video ID by virtue of storing the associated clip. Mowry, ¶85.

This association is reflected, for example, in the clip database of the IM. *Id.*, ¶86. As 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" (*Kenner*, 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).

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. Mowry, ¶87; *Kenner*., 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"

Kenner discloses [1.d]. Mowry, ¶¶88-90.

Kenner discloses requests for the first resource from said client machines. Mowry, ¶88. 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.

Kenner's requests are directed to the primary IM and thus are *directed to the at least one shared repeater server*. Mowry, ¶89. 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." *Kenner*, 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.

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. Mowry, ¶90. 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. Mowry, ¶90; §VI.A.1.b.

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"

Kenner discloses [1.e]. Mowry, ¶¶91-92.

As 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*. §VI.A.2.e. A POSITA would have understood that the same applies to *a second resource associated with the second content provider*. Mowry, ¶91; §VI.A.1.b. 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* §VI.A.2.c), the given IM is associated with a second clip uploaded by the second content provider. Mowry, ¶91; §VI.A.2.e.

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. Mowry, ¶92; *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. Mowry, ¶92; *Kenner*, 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"

Kenner discloses [1.f]. Mowry, ¶¶93-94.

As 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*. §VI.A.2.f. The same would apply to requests for the second resource. Mowry, ¶93. The user similarly requests access to the second clip through the user terminal (client machine). *Kenner*, 36:10-27.

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. Mowry, ¶94; §VI.A.2.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"

Kenner discloses [1.g]. Mowry, ¶¶95-96.

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.

Mowry, ¶95. As 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; §VI.A.2.f.

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 resourc*e because it uses the clip database to identify the source IM. Mowry, ¶96. As explained for [1.d], if the primary IM determines a clip is not stored at its associated extended SRUs, it queries the closest IMs. §VI.A.2.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." *Kenner*, 41:13-15. The primary IM thus consults the clip database to determine an internet address of the source IM. *Id*. 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. Mowry, ¶96. Indeed, the "content provider's account number" is included in the video ID within the database. *Kenner*, 36:32-37:6.

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

Kenner discloses claim 2. Mowry, ¶¶97-98.

As 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*. §VI.A.2.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. The primary IM is therefore *addressable using the first alias name*, because the virtual URL addresses the primary IM and the video ID is used to generate the URL. Mowry, ¶97.

As 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*. §VI.A.2.h. The primary IM is also *addressable using the second alias name* for the same reason as the first alias name. Mowry, ¶98.

4. Claim 22: "wherein at least one of the first resource and the second resource is an embedded object"

Kenner discloses claim 22. Mowry, ¶99.

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.

5. Claim 23: "wherein the requests comprise host tag

names that identify specific origin servers from which the resources originate"

Kenner discloses claim 23. Mowry, ¶¶100-01.

Kenner discloses delivering HTTP-based clips. *Id*., ¶100; *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.'").

A POSITA would have recognized that including host tag names in requests for HTTP-based resources was standard practice. Mowry, ¶101. These HTTP requests include HTTP headers containing metainformation about the requested resource. *Id.* (citing *Berners-Lee*, ¶ 78; *HTTP/1.1*, §5). It was standard practice at least as early as January 1997 (HTTP version 1.1) for HTTP requests to include host header fields (*host tag names*) in these HTTP headers. *Id.* (citing *HTTP/1.1*, §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.* (citing *HTTP/1.1*, §§5.1.2, 14.23).

6. Claim 24: "wherein the host tag names are included in HTTP headers"

Kenner discloses claim 24. Mowry, ¶102.

As explained for claim 23, *Kenner*'s HTTP-based requests include *host tag names*. §VI.A.5. It was standard practice to include the host tags in HTTP headers

for the request. *Id.*; Mowry, ¶102 (citing *HTTP/1.1*, §§5.1.2, 14.23).

7. 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)"

Kenner discloses or at least suggests claim 26. Mowry, ¶¶103-10.

As explained for [1.b], *Kenner* discloses a table listing a plurality alias names corresponding to content providers. §VI.A.2.d.

Kenner further discloses including *a reflector name* in its clip database. Mowry, ¶104. 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.

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. Mowry, ¶105.

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. Mowry, ¶106. And each of the source IMs is associated with a corresponding content provider. *Id*.; §VI.A.1.b.

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. Mowry, ¶107. This rate of transfer of information committed to a particular subscriber is the same aggregate rate described in the '903 patent as the CAIR. Mowry, ¶108. For example, *Kenner* indicates it varies playback quality and resolution for different users (*id.*, 51:30-52:6), which would affect transmission bandwidth. Mowry, ¶108. And *Kenner* tracks the download speeds for individual users. *Kenner*, 47:8-19. Mowry, ¶108.

To the extent *Kenner* does not explicitly disclose the claimed CAIR, it would have been obvious to include this information in *Kenner*'s database. Mowry, ¶109.

For example, *Kenner* discloses storing data to improve load projection, such as a usage count for a particular clip. *Kenner*, 35:1. *Kenner*'s IMs also maintain information correlating usage rates for individual users to content providers the users are subscribed to. Mowry, ¶109. For example, *Kenner*'s IMs 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." *Kenner*, 33:19-22. A POSITA would have found it obvious to use this information to allocate bandwidth usage across content providers and store such information as a committed aggregate data rate across IMs for each given content provider, which would further improve the load projections in *Kenner*, especially in view of the bandwidth limitations discussed above. Mowry, ¶110.

8. 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"

Kenner discloses [28.pre.i]. Mowry, ¶111.

As explained for [1.pre], *Kenner* discloses a content delivery system operative in a computer network. *Id.*; §VI.A.2.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. Mowry, ¶111; *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"

Kenner discloses [28.pre.ii]. Mowry, ¶112.

As explained for [1.a.ii], *Kenner*'s network includes a plurality of source IMs through which content providers upload clips. §VI.A.2.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. Mowry, ¶112; *Kenner*, 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"

Kenner discloses [28.pre.iii]. Mowry, ¶113.

As explained for [1.a.i], *Kenner*'s IM is *at least one shared repeater server constructed and adapted to replicate* resources. §VI.A.2.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 explained for [28.pre.ii], the resources replicated by the IM are associated with source IMs (i.e., origin servers). §VI.A.8.b; Mowry, ¶113.

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"

Kenner discloses [28.a]. Mowry, ¶¶114-16.

As explained for [1.c], *Kenner* discloses a first alias name associated with the at least one repeater server. §VI.A.2.e. *Kenner* likewise discloses *associating the at least one repeater server with a first alias name*. Mowry, ¶114. 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.

As 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*. §VI.A.2.f. *Kenner* discloses the first resource being *located on a first origin serve*r 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; Mowry, ¶115.

As 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*. §VI.A.2.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; Mowry, ¶116.

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"

Kenner discloses [28.b]. Mowry, ¶¶117-19.

As explained for [1.e], *Kenner* discloses a second alias name associated with the at least one repeater server. §VI.A.2.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; Mowry, ¶118.

As 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*. §VI.A.2.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* §VI.A.8.d. Indeed, *Kenner*'s IMs store clips from multiple different source IMs. §VI.A.2.c; Mowry, ¶119.

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"

Kenner discloses or at least suggests [28.c]. Mowry, ¶¶120-23.

As 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; §VI.A.2.d; Mowry, ¶121.

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. *Id*., ¶122; *Kenner*, 41:13-16. As 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*. §VI.A.8.e.

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. Mowry, ¶123; §VI.A.2.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"

Kenner discloses [28.d], as explained for [1.g].[2] Mowry, ¶124. §VI.A.2.i; *Kenner*, 41:13-16 (the IM "contact[s] the source IM 90, where the content provider first uploaded the file.").

9. 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"

Kenner discloses claim 29. Mowry, ¶¶125-29.

As explained for [28.a], *Kenner* discloses *a request for the first resource* §VI.A.8.d. *Kenner* discloses this request being *issued by a particular client machine*

[2] Claim 28 invalid. [28.pre.i] recites a method, whereas [28.d] recites an apparatus, thus mixing two statutory classes of invention. Claim 28 "does not apprise a person of ordinary skill in the art of its scope" and is invalid under 35 U.S.C. §112. *IPXL Holdings, L.L.C. v. Amazon.com, Inc.*, 430 F.3d 1377, 1384 (Fed. Cir. 2005). For purposes of this Petition, and without conceding that the statutory requirements are met for claim 28, Petitioner nonetheless addresses this limitation.

because a request for a clip in *Kenner* is issued through a user terminal (i.e., a client machine). *Kenner*, 25:7-24.

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. Mowry, ¶126. *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. *Kenner*, 41:1-6

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." *Kenner*, 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. Mowry, ¶127; 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.").

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. Mowry, ¶128.

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. *Kenner*, 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"). Mowry, ¶129; §VI.A.2.e.

10. 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"

Kenner discloses claim 30. Mowry, ¶¶130-32.

As 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*. §VI.A.9.

Kenner further discloses identifying the IM (*repeater server*) based *on load characteristics*. Mowry, ¶132. 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.

11. 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"

Kenner discloses claim 31. Mowry, ¶¶133-35.

As 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*. §VI.A.9.

Kenner further discloses identifying the IM (*repeater server*) based on location. Mowry, ¶135. 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") (emphases added).

12. Claim 32: "wherein the resources comprise an HTTP-based resource"

Kenner discloses claim 32. Mowry, ¶136.

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").

13. 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)"

Kenner discloses or at least suggests claim 35, as explained for claim 26. §VI.A.7; Mowry, ¶137.

14. 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"

Kenner discloses [37.pre.i], as explained for [28.pre.ii]. Mowry, ¶138.

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

Kenner discloses [37.pre.ii], as explained for [28.pre.iii]. §VI.A.8.c; Mowry, ¶139.

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"

Kenner discloses [37.a.i]. Mowry, ¶¶140-42.

As explained for [28.a] and [28.b], *Kenner* discloses *associating a plurality of alias names with the at least one shared repeater server*. §§VI.A.8.d, VI.A.8.e; Mowry, ¶141. 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).

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). Mowry, ¶142. 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. *Id.* Accordingly, an alias name for a first one of the clips is associated with a first source IM. *Id.*

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"

Kenner discloses [37.a.ii]. Mowry, ¶¶143-44.

As explained for [37.a.i], *Kenner* discloses *each of said plurality of alias names being associated with an origin server*. §VI.A.14.c; Mowry, ¶144. This includes *a second alias name of said plurality of alias names is associated with a second origin server*. Mowry, ¶144. 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 explained for [28.b], *Kenner* discloses *the second origin server [being] distinct from the first origin server*. §VI.A.8.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"

Kenner discloses [37.b]. Mowry, ¶¶145-46.

As 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*. §VI.A.8.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; Mowry, ¶146.

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"

Kenner discloses [37.c], as explained for [28.a]. §VI.A.8.d; Mowry, ¶147. 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"

Kenner discloses [37.d], as explained for [28.b]. §§VI.A.8.e, VI.A.14.f; Mowry, ¶148.

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"

Kenner discloses [37.e], as explained for [28.d]. §VI.A.8.g; *Kenner*, 41:13-16; Mowry, ¶149.

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

Kenner discloses claim 38. Mowry, ¶150.

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). *Id*. *Kenner* explains that each IM "will query its clip database to determine if any of its extended SRUs are storing the clip." *Kenner*, 45:25-31, *see* 44:6-14 (describing the source IM's "own clip database"), 45:32-33 ("Each IM …

performs its own maintenance on the clip database").

16. 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"

Kenner discloses [40.pre.i]. Mowry, ¶151.

As explained for [1.pre], *Kenner* discloses a *computer network for delivering resources associated with a plurality of content providers to multiple client machines*. §VI.A.2.a. *Kenner*'s IM is *a server operative* in the network. §VI.A.2.b; Mowry, ¶151.

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"

Kenner discloses [40.pre.ii]. Mowry, ¶152.

As 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*. §§VI.A.2.b, VI.A.2.c.

c. [40.a]: "the server comprising: … cache storage"

Kenner discloses [40.a]. Mowry, ¶¶153-55.

Kenner discloses *cache storage* via a buffer at each IM. *Id*., ¶153. *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:

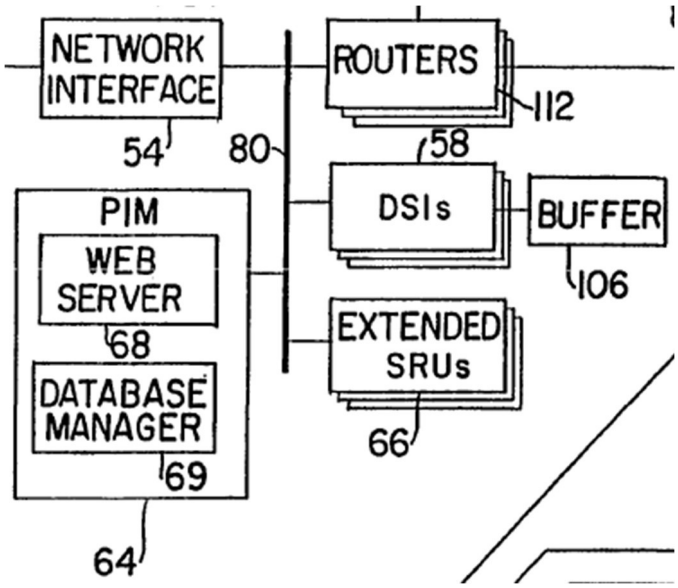

Id., Fig. 4, 48:9-15. "[A]ny computer connected to the backbone 80," can create these DSIs, including the web server of the primary IM. *Id*., 31:26-29.

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. Mowry, ¶154. Caching refers to storing data temporarily in an accessible manner, allowing information that is used frequently to be stored in close proximity to where the information is likely to be used. *Id.*, (citing *Zivkov*, §2). *Kenner*'s buffer 106 operates in the same way. *Id*. 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." *Kenner*, 48:9-15. Accordingly, by storing data buffer 106, *Kenner*'s IM includes cache storage. Mowry, ¶154.

Kenner's IMs also include *cache storage* by virtue of storing copies of content replicated from origin servers, as described for [1.a.1]. §VI.A.2.b. Consistent with the general principles of caching explained above, *Kenner*'s IM acts as a web cache by storing copies of clips local to a client requesting access to the clips. Mowry, ¶155 (citing *Kenner*, 1:13-16; *Zivkov*, §2). In web-based content distribution networks, web caches store copies of documents throughout the Web *Id*. (citing *Malpani*, Introduction). A POSITA would have understood the primary IM acts as a Web cache for the network and thus includes *cache storage*. *Id*.

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"

Kenner discloses or at least suggests [40.b]. Mowry, ¶¶156-59.

As 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; §VI.A.2.d.

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.

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." *Kenner*, 38:23-31; Mowry, ¶159; §VI.A.2.d.

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

Kenner discloses [40.c]. Mowry, ¶160.

As explained for [1.c], *Kenner* discloses a first alias name associated with the server in the form of a first video ID. §VI.A.2.e. And as explained for [1.b], the video IDs *correspond[] to content providers* because the video IDs include "the content provider's account number." §VI.A.2.d; *Kenner*, 36:32-37:6. Accordingly, a first video clip in the database would correspond to a first content provider. Mowry, ¶160.

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"

Kenner discloses [40.d]. Mowry, ¶¶161-63.

As explained for [1.e], *Kenner* discloses a second alias name associated with the at least one repeater server. §VI.A.2.g. And as explained for [1.b], the video IDs *correspond[] to content providers* because the video IDs include "the content provider's account number." §VI.A.2.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. Mowry, ¶162. 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:

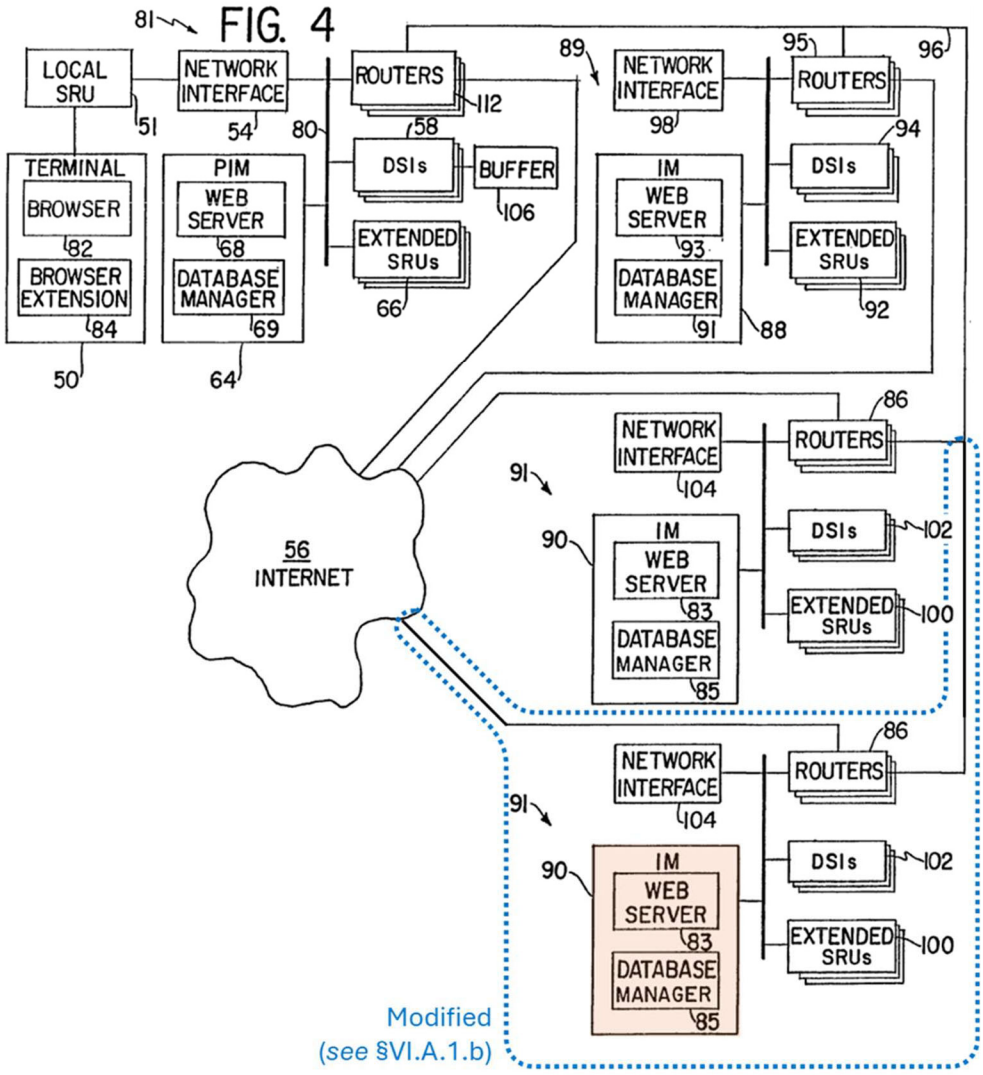

Id. (modified *Kenner* Fig. 4).

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."

Kenner, 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; Mowry, ¶163.

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"

Kenner discloses [40.e]. Mowry, ¶164.

As 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. §§VI.A.2.b, VI.A.2.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; Mowry, ¶164.

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"

Kenner discloses [40.f], as explained for [1.g]. Mowry, ¶165; §VI.A.2.i; *Kenner*, 41:13-16 (explaining the IM "contact[s] the source IM 90, where the content provider first uploaded the file.").

17. 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)"

Kenner discloses claim 41, as explained for claim 35. §VI.A.13; Mowry, ¶166.

18. Independent Claim 43

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

Kenner discloses [43.pre], as explained for [40.pre.i]. §VI.A.16.a; Mowry, ¶167.

b. [43.a]: "the server comprising: … cache storage"

Kenner discloses [43.a], as explained for [40.a]. §VI.A.16.c; Mowry, ¶168.

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

Kenner discloses [43.b], as explained for [40.b]. §VI.A.16.d; Mowry, ¶169. *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; §VI.A.2.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"

Kenner discloses [43.c], as explained for [40.c]. §VI.A.16.e; Mowry, ¶170. *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"

Kenner discloses [43.d], as explained for [40.d]. §VI.A.16.f; Mowry, ¶171; *see* §VI.A.18.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"

Kenner discloses [43.e], as explained for [40.e]. §VI.A.16.g; Mowry, ¶172.

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"

Kenner discloses [43.f], as explained for [40.f]. §VI.A.16.h; Mowry, ¶173.

19. Claim 44: "wherein the first origin server and the second origin server are associated with distinct content providers located at distinct physical locations"

Kenner discloses claim 44. Mowry, ¶¶174-75.

As explained for [1.a.ii], *Kenner* discloses a plurality of distinct content providers that upload content at source IMs in different regions. §VI.A.2.c.

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. Mowry, ¶175; *Kenner*, Abstract (describing "uploading and distributing clips to geographically diverse servers"), 19:18-19.

20. 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)"

Kenner discloses claim 41, as explained for claim 35. §VI.A.13; Mowry, ¶176.

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

1. Overview of Ground 2

a. *Vetter*

Vetter is a publication promoting Mosaic, a graphics-oriented browser for accessing web-based content. *Vetter*, Abstract. *Vetter* is analogous to the '903 patent

because it is in the same field of endeavor: web-based content delivery. Mowry, ¶59; *Bigio*, 381 F.3d at 1325. *Vetter* was published by a well-known publisher (IEEE) and was cataloged at the Library of Congress in October 1994, confirming public availability. EX-1014, ¶¶21-23.

Vetter explains that browsers, like Mosaic, allow users to access content using various protocols, including a file-transfer protocol (FTP). Mowry, ¶60; *Vetter*, 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*.

b. Motivation to Combine *Kenner* and *Vetter*

A POSITA would have been motivated to implement *Kenner*'s system to deliver FTP-based content, as disclosed by *Vetter*. Mowry, ¶¶177-80.

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.

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. Mowry, ¶179.

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. Mowry, ¶180. 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. *See KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 417-18 (2017); Mowry, ¶180.

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

Kenner in view of *Vetter* teaches claim 33. Mowry, ¶¶181-82.

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. Mowry, ¶181. A POSITA would have recognized that *Kenner*'s system could equally be used for other protocols including ftp-based resources. *Id.*; *Kenner*, 14:5-7 (using "conventional network protocols and topologies"); *Berners-Lee*, 78, 81.

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. Mowry, ¶182. 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. *Id.*; §VI.B.1.b. Moreover, the '903 patent admits the FTP scheme was known in the art. EX-1001, 6:22-41. Mowry, ¶182.

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

1. Overview of Ground 3

a. *Rekimoto*

Rekimoto is analogous to the '903 patent because they are in the same field of endeavor: web-based content delivery. Mowry, ¶61; *Bigio*, 381 F.3d at 1325.

Like *Kenner*, *Rekimoto* discloses a system for delivering content to user terminals within a network. Mowry, ¶62; *Rekimoto*, 2:22-29. 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.

b. Motivation to Combine *Kenner* and *Rekimoto*

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. Mowry, ¶183.

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.

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 latency in the network. *Rekimoto*, 17:46-49; Mowry, ¶185. *Kenner* shares this goal. *Kenner*, 5:21-26,

7:25-31.

Combining *Rekimoto*'s DNS with *Kenner*'s system would have been straightforward. Mowry, ¶186. *Kenner* acknowledges the advantages of prioritizing servers based on proximity. *Kenner*, 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. Mowry, ¶186. Indeed, DNS was well-known at the time. *Id.*; *Tanenbaum*, 622-30. 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. *See KSR*, 550 U.S. at 417-18; Mowry, ¶186.

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"

Kenner in view of *Rekimoto* teaches [34.a]. Mowry, ¶¶187-91.

As the '903 patent admits, DNS resolution was known at the time of filing. EX-1001, 7:4-29; Mowry, ¶187. 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. EX-1001, 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.

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. Mowry, ¶188.

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. Mowry, ¶189. *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.

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. Mowry, ¶190. 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 (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. Mowry, ¶190. 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. *Id*.; EX-1001, 21:56-64. The '903 patent admits DNS was known in the art. EX-1001, 6:45-61, 7:4-15. Mowry, ¶190.

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. Mowry, ¶191. *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, it would have been obvious to structure *Kenner*'s video ID as a domain name to provide the DNS resolution advantages described above. Mowry, ¶191.

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"

Kenner in view of *Rekimoto* teaches [34.b]. Mowry, ¶¶192-94.

As explained for [1.g], *Kenner* discloses receiving an alias name as part of a client request for a particular resource. §VI.A.2.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.

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 (§VI.C.2.a), would be provided to the DNS server 130, as described in *Rekimoto* (*Rekimoto*¸ 18:3-12).

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. Mowry, ¶194; *Rekimoto*, 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. §VI.C.2.a.

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

1. Overview of Ground 4

a. *Boyles*

Boyles is analogous to the '903 patent because they are in the same field of endeavor: web-based content delivery. Mowry, ¶63; *Bigio*, 381 F.3d at 1325.

Boyles discloses a network of cache server nodes "responsible for caching or storing information about all or a subset of the network's resources." *Boyles*, 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; Mowry, ¶64.

b. Motivation to Combine *Kenner* and *Boyles*

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

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; Mowry, ¶196.

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. Mowry, ¶197. 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. Mowry, ¶197.

Implementing *Boyles*'s express failure notification in *Kenner* would have been straightforward. Mowry, ¶198. *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. *See KSR*, 550 U.S. at 417-18; Mowry, ¶198.

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"

Kenner in view of *Boyles* teaches claim 27. Mowry, ¶¶199-201.

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.

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. Mowry, ¶200. In that case, the "reason for the denial" would be failure to locate the clip. *Kenner*, 38:17-22.

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. Mowry, ¶201. 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. *Id*.; §VI.D.1.b.

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"

Kenner discloses claim 36, as explained for claim 27. §VI.D.2; Mowry, ¶202.

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"

Kenner discloses claim 39, as explained for claim 27. §VI.D.2; Mowry, ¶203.

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"

Kenner discloses claim 42, as explained for claim 27. §VI.D.2; Mowry, ¶204.

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"

Kenner discloses claim 46, as explained for claim 27. §VI.D.2; Mowry, ¶205.[3]

VIII. Mandatory Notices

A. Real Party-in-Interest

The Petitioner and real party-in-interest is Microsoft Corporation.

[3] The Director's designee already determined not to discretionarily deny institution of the Petition in IPR2025-00969. *Google LLC v. Sandpiper CDN, LLC*, IPR2025-00969, Paper 13 (Oct. 10, 2025). Because the current Interim Director Discretion Process specifically states "[t]he petition should not address discretionary issues," this Petition does not repeat the discretion sections in the petition in IPR2025-00969.

B. Related Matters

Sandpiper asserted the '903 patent in the following litigations:

- *Sandpiper CDN, LLC v. Microsoft Corp.*, No. 2-25-cv-00664 (E.D. Tex., June 26, 2025);
- *Sandpiper CDN, LLC v. Google LLC*, No. 2:24-cv-03951 (N.D. Cal., May 10, 2024); and
- *Sandpiper CDN, LLC v. Comcast Cable Commc'ns Mgmt. LLC*, No. 2:24-cv-00886 (E.D. Tex., Nov. 1, 2024).

C. Lead and Back-Up Counsel, and Service Information

IX. LEAD COUNSEL
Jessica Kaiser (Reg. No. 58,937) kaiser-ptab@perkinscoie.com PERKINS COIE LLP 1900 Sixteenth Street, Suite 1400 Denver, CO 80202-5255 Telephone: (303) 291-2300
BACK-UP COUNSEL
Christopher Marando (Reg. No. 67,898) Marando-ptab@perkinscoie.com PERKINS COIE LLP 700 Thirteenth Street, N.W. Suite 800 Washington, DC 20005-3960 Telephone: (202) 654-6227 Matthew Lembo (Reg. No. 75,633) Lembo-ptab@perkinscoie.com PERKINS COIE LLP 1155 Avenue of the Americas

22nd Floor
New York, NY 10036-2711
Phone: (212) 262-6900

Steven Beigelmacher (Reg. No. 75,050)
beigelmacher-ptab@perkinscoie.com
PERKINS COIE LLP
1301 Second Avenue,
Suite 4200
Seattle, WA 98101-3804
Telephone: (206) 359-8000

Sarah Fowler (Reg. No. 54,854)
stahnke-ptab@perkinscoie.com
PERKINS COIE LLP
3150 Porter Drive
Palo Alto, CA 94304
Telephone: (650) 838-4300

Petitioner consents to electronic service at the following email address: Microsoft-Sandpiper-IPR@perkinscoie.com.

X. Grounds for Standing

Petitioners certify the ’903 patent is available for *inter partes review* and that Petitioners are not barred or estopped.

XI. Conclusion

Petitioner requests institution of *inter partes* review and cancellation of the challenged claims.

Respectfully submitted,
/ *Jessica Kaiser* /

Dated: November 3, 2025

Attorney for Petitioner

CLAIM APPENDIX

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

[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,

[1.a.ii] and 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;

[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;

[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;

[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;

[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;

[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; and

[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.

[2] 2. The content delivery system as recited in claim 1, wherein the least one shared repeater server is addressable using the first alias name and the second alias name.

[22] 22. The content delivery system as recited in claim 1, wherein at least one of the first resource and the second resource is an embedded object.

[23] 23. The content delivery system as recited in claim 1, wherein the requests comprise host tag names that identify specific origin servers from which the resources originate.

[24] 24. A content delivery system as recited in claim 1, 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.

[26] 26. A method as recited in claim 1, 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).

[27] 27. The content delivery system as recited in claim 23, wherein the host tag names are included in HTTP headers.

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

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

[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, the method comprising:

[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;

[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;

[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; and

[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.

[29] 29. The method of claim 28, 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.

[30] 30. The method of claim 28, 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.

[31] 31. The method of claim 28, 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.

[32] 32. The method as recited in claim 28, wherein the resources

comprise an HTTP-based resource.

[33] 33. The method as recited in claim 28, wherein the resources comprise a File Transfer Protocol (FTP)-based resource.

[34.a] 34. A method as recited in claim 28, 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,

[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.

[35] 35. A method as recited in claim 28, 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).

[36] 36. A method as recited in claim 28, further comprising:

rejecting the client request for the particular resource if the alias name received

with the client request is not included the table.

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

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

[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, and

[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,

[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;

[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; and

[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; and

[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.

[38] 38. The method as in claim 37, wherein a copy of the at least one table is located on each of a plurality of repeater servers.

[39] 39. A method as recited in claim 37, further comprising:

rejecting the client request for the resource if the alias name received with the client request is not included the table.

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

[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, the server comprising:

[40.a] cache storage;

[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,

[40.c] wherein the server is associated with a first alias name corresponding to said first content provider, and

[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, and

[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, and

[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.

[41] 41. A server as recited in claim 40, 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).

[42] 42. A server as recited in claim 40, 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 the table.

[43.pre] 43. A server operative in a computer network for delivering resources to multiple client machines, the server comprising:

[43.a] cache storage;

[43.b] a table listing a plurality of origin servers having content that is authorized for delivery to client machines via the server,

[43.c] wherein the server is associated with a first alias name corresponding to a first origin server of the plurality of origin servers, and

[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, and

[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, and

[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.

[44] 44. A server as recited in claim 43, wherein the first origin server and the second origin server are associated with distinct content providers located at distinct physical locations.

[45] 45. A server as recited in claim 43, 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).

[46] 46. A server as recited in claim 43, 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 the table.

CERTIFICATE OF COMPLIANCE

The undersigned hereby certifies that the foregoing **Petition for *Inter Partes* Review** contains 12,901 words, excluding those portions identified in 37 C.F.R. §42.24(a), as measured by the word-processing system used to prepare this paper.

Respectfully submitted,

/ *Jessica Kaiser* /

Jessica Kaiser
Reg. No. 58,937
Attorney for Petitioner

Dated: November 3, 2025

CERTIFICATE OF SERVICE

The undersigned certifies that the foregoing Petition for *Inter Partes* Review of U.S. Patent No. 8,478,903, the associated Power of Attorney, and Exhibits EX-1001 through EX-1017 were served on November 3, 2025 by FedEx Priority Overnight® on the correspondence address of record indicated in the Patent Office's Patent Center system for U.S. Patent No. 8,478,903.

STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.
1101 K Street, NW
10th Floor
WASHINGTON, DC 20005

A courtesy copy was also sent via electronic mail to Patent Owner's litigation counsel listed below:

Andrea Leigh Fair - andrea@millerfairhenry.com
Carolyn E Drell - cdrell@bsfllp.com
Garrett C. Parish - garrett@millerfairhenry.com
Mark D. Schafer - mschafer@bsfllp.com
Philip Eckert - peckert@bsfllp.com
Ryan D. Dykal - rdykal@bsfllp.com

Respectfully submitted,

/ *Jessica Kaiser* /

Jessica Kaiser
Reg. No. 58,937
Attorney for Petitioner

Dated: November 3, 2025