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
U.S. Patent No. 8,478,903

DECLARATION OF DR. PRASHANT SHENOY

Mail Stop "PATENT BOARD"
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

I, Prashant Shenoy, Ph.D., do hereby declare as follows:

1. I am over the age of twenty-one and otherwise competent to make this Declaration. This Declaration is based on my personal knowledge as an expert in the relevant field. I could testify competently if asked to do so.

2. I have been retained by Sterne, Kessler, Goldstein & Fox, P.L.L.C. as a technical expert witness on behalf of Sandpiper CDN, LLC ("Patent Owner") to provide this Declaration concerning technical subject matter relevant to the *inter partes* review of U.S. Patent No. 8,478,903 ("'903 Patent") in IPR2026-00095.

3. I am being compensated for the time spent on this matter at my standard consulting rate of $650.00/hour. My compensation is not dependent on, and in no way affects, the substance of my statements in this Declaration. My findings are based on my education, experience, and background in the fields discussed below. I reside in Amherst, Massachusetts, USA.

4. In formulating my opinions, I have relied upon my experience in the relevant art, and I have considered each document cited herein, in light of the general knowledge in the art as of February 10, 1998, the earliest possible priority date of the '903 Patent. For example, I have considered the documents listed in Section II and any other documents cited in this declaration. To the best of my knowledge, these documents are true and accurate copies of what they purport to

be. An expert in the field would reasonably rely on them to formulate opinions such as those set forth in this declaration.

5. I will cite to the specification of a U.S. patent using the following format: (Patent, Col: Line Number(s)). For example, the citation to the '903 Patent, 1:1-10 points to the '903 Patent specification at column 1, lines 1-10.

6. I am familiar with the technology at issue and the state of the art at the time the application leading to the '903 Patent was filed, including content delivery networks, web protocols, network protocols, and proxy servers.

7. I have been asked to provide my technical review, analysis, insights, and opinions regarding the references. I reserve the right to supplement this Declaration in response to additional evidence that may come to light.

I. MY BACKGROUND AND QUALIFICATIONS

8. I earned a Bachelor of Technology degree from the Department of Computer Science and Engineering at the Indian Institute of Technology in Bombay, India. I later received my Master of Science (M.S.) and Doctor of Philosophy (Ph.D.) degrees in 1994 and 1998, respectively, in Computer Science from the University of Texas at Austin.

9. Over the years, I have led research and industry collaboration projects in areas such as storage and file systems, Web and Internet systems, content

distribution, streaming media, database systems, cloud computing, virtualization, RFID and sensor networks, and energy systems.

10. I am a tenured professor in the College of Information and Computer Sciences at the University of Massachusetts Amherst ("UMass"), where I currently hold the position of Distinguished Professor of Computer Science. I also served as Associate Dean in the College of Information and Computer Sciences from 2016 to 2025. I have been a faculty member at UMass since 1998.

11. I received my Bachelor of Technology (B.Tech.) degree in 1993 from the Department of Computer Science and Engineering at The Indian Institute of Technology ("IIT"), Mumbai, India. IIT awarded me the Institute Medal for being the top-ranking undergraduate student. In 1994, I received my Master of Science (M.S.) degree in Computer Sciences from The University of Texas at Austin. In 1998, I received my Doctor of Philosophy (Ph.D.) also from the Department of Computer Sciences at The University of Texas at Austin. My dissertation, which was in the area of file systems, was recognized by the department as the Best Doctoral Dissertation of 1998-1999.

12. I have been involved in the research and development of distributed computing systems for more than 30 years. At UMass, I perform research and teach in the areas of distributed systems, operating systems, computer networking, data management, and multimedia systems, among others.

13. A major focus of my research has been the design of large-scale distributed computing systems. Over the years, I have led research and industry collaboration projects in areas such as computer networking, distributed systems, storage and file systems, Web and Internet systems, content distribution, streaming media, database systems, cloud computing, virtualization, Radio Frequency Identification ("RFID") and sensor networks, and energy systems. During that time, I have consulted with both large companies and startups in a technical advisory role related to product architecture, intellectual property ("IP") issues, and applied research.

14. For example, starting in the early 1990s, my research focused on the development of storage and file systems. During that time, I developed multimedia file system and data management technologies that were designed for efficient storage and retrieval of large amounts of video and audio data. I collaborated with AT&T Bell Laboratories to design a new media server for video streaming. These efforts led to resulted in U.S. Patent No. 6,079,028 on fault-tolerant architectures for continuous media servers.

15. In the late 1990s, I led the design of a new operating system called QLinux that was based on the open-source Linux operating system. As part of this research, I developed resource management and scheduling techniques that provided isolation and service guarantees to applications. This technology

provided the underpinnings for fair resource management and isolation techniques used in operating system ("OS") virtualization. It was adopted by a startup company for providing resource isolation in its web hosting products.

16. In the early 2000s, I led several projects on web systems and cloud computing. For example, I led research projects on web technologies such as web caching and scalable web applications. I also led research projects on cloud computing platforms (which were known as web hosting platforms at the time) involving virtualization, dynamic resource management, replication, fault tolerance, and storage. My group at UMass collaborated with technology companies such as IBM, AT&T, Sprint, Intel, and many others as part of this research.

17. My research on virtualization has involved designing technologies for virtual machines and container resource management and deploying them in cloud computing systems and data centers. I am co-inventor on U.S. Patent No. 8,473,557, jointly assigned to AT&T and UMass, on virtual machine migration.

18. Over the years, I have authored or co-authored more than 300 articles in peer-reviewed journals and conference proceedings that have been cited over 27,000 times. Seventeen papers authored by myself and members of my research group on this project won awards for their technical contributions. Two papers

have won Test-of-Time awards, an award given to a publication that has had the most impact over a ten year period since its publication.

19. I have served as an Editor for six leading journals and have served as the Technical Program Chair for twenty conferences and workshops in my field. I founded the ACM Special Interest Group on Energy Systems and also served on a steering committee of conferences and journals.

20. I am also a Fellow of the ACM, IEEE, AAAS, and AAIA, four professional societies in computing and science.

21. I am a named inventor on three U.S. Patents and one published application (U.S. Application No. 2014/0040206 and U.S. Patent Nos. 8,510,596, 8,473,557, and 6,079,028). These documents are generally related to distributed and operating systems, virtualization, and cloud data centers.

22. As a faculty member at UMass, I have also taught undergraduate and graduate courses in operating systems, distributed systems, networking, and related topics for more than twenty years.

23. These courses teach concepts related to operating systems, virtual machines, containers, cloud computing, file systems, and distributed systems. They also involve hands-on course projects on these topics.

24. My *Curriculum Vitae* is attached as EX2006, which contains further details on my education, experience, publications, and other qualifications to render an expert option.

II. LIST OF MATERIALS CONSIDERED

Exhibit No.	Description
1001	U.S. Patent No. 8,478,903 to Farber et al. ("'903 Patent")
1002	Prosecution History of U.S. Application No. 11/065,412
1003	Declaration of Dr. Todd C. Mowry ("Mowry Dec.")
1004	Curriculum Vitae of Dr. Todd. C. Mowry
1005	International Publication No. WO1996041285 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 Publication 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")
1012	Fielding et al., Hypertext Transfer Protocol – HTTP/1.1, RFC 2068, Jan. 1997, https://www.rfc-editor.org/rfc/rfc2068 ("HTTP/1.1")

Exhibit No.	Description
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)
2001	Infringement Contentions, *Sandpiper CDN, LLC v. Microsoft Corp.*, Case No. 2:25-cv-00664 (E.D. Tex. Sept. 16, 2025)
2002	Sotera Stipulation, *Sandpiper CDN, LLC v. Microsoft Corp.*, Case No. 2:25-cv-00664 (E.D. Tex. Sept. 16, 2025)
2003	FAQs for Interim Processes for PTAB Workload Management, accessed at https://www.uspto.gov/patents/ptab/faqs/interim-processes-workload-management
2004	Interim Processes for PTAB Workload Management Memorandum (March 26, 2025)
2006	*Curriculum Vitae* of Dr. Prashant Shenoy
2007	Deposition Transcript of Dr. Todd C. Mowry in IPR2025-00969, February 10, 2026
2008	Complaint for Patent Infringement, *Sandpiper CDN, LLC v. Microsoft Corporation*, Case No. 2:25-cv-00664 (E.D. Tex.), June 26, 2025

III. BACKGROUND OF THE TECHNOLOGY

25. Content delivery networks ("CDNs") provide critical services that enable content providers to quickly deliver online content to millions of consumers simultaneously over the Internet. In the early 1990s, the World Wide Web saw increasing adoption and became a household staple. This mass adoption led to data congestion issues from the ever-growing number of users seeking to simultaneously access Internet content.

26. A typical computer server in the 1990s could only handle a limited number of simultaneous connections before becoming overloaded. Moreover, data take time to traverse routers and move through content-congested internet cables. Consumers living far from the physical server(s) hosting content experienced sluggish load times and high latency due to problems such as overloaded servers, congested network segments, and geographic separation.

27. The availability of CDN servers around the world allowed service providers to replicate appropriate content from customers' origin servers to appropriate CDN servers and transparently route end users requesting that content to the "best" CDN server to deliver that content, all while providing their customers with control over their content and the user experience. This service and its architecture was quickly adopted by many others in the industry.

28. End-users would connect to an optimal edge server—a server that was closer to the end user or that had available capacity. Distributing content across a network of servers alleviated data congestion issues and allowed consumers to connect to edge servers located near them to reduce latency. Systems and methods for propagating data from origin servers to edge servers and storing them at end servers for serving user requests are known as "caching."

29. I understand that by 1996, a team at Sandpiper Networks Inc. had conceived techniques for delivering static and dynamic content, including streaming resources, such as audio and video, using Sandpiper's CDN. EX2008, ¶23.

30. I also understand that by at least May 1998, Sandpiper Networks was caching content and delivering cached content to end users of content providers using its CDN. EX2008, ¶¶25-26 (Footnote 12). Sandpiper Networks' first paying customer was the L.A. Times, which paid Sandpiper Networks to host the report of Independent Counsel Ken Starr on his investigation of President Bill Clinton ("the Starr Report") beginning on September 11, 1998. EX2008, ¶¶25-26. Sandpiper Networks' CDN was capable of caching and was used to cache and deliver Internet resources including pictures, text files, dynamic resources, and streaming multimedia resources. EX2008, ¶¶25-26.

31. In the early-to-mid 2000s, demand for CDNs exploded. Companies like Microsoft built CDN platforms to provide public CDN services to customers and to meet their own enormous data needs. *See* EX2001. For example, Microsoft built its Azure CDN and Azure Front Door platforms. Additionally, services like Microsoft's Xbox gaming, Office 365, Windows updates, and web resources like Linkedin and Bing rely upon CDN technology to meet its enormous content provision needs.

IV. OVERVIEW OF THE '903 PATENT

32. The '903 Patent is titled "Shared Content Delivery Infrastructure" and was filed as U.S. Patent Application No. 09/021,506 on February 23, 2005. The '903 Patent claims priority benefit to U.S. Application No. 09/612,598 ("the '598 Application"), filed July 7, 2000, now abandoned. The '598 Application was a divisional of U.S. Patent Application No. 09/021,506, filed on February 10, 1998, now U.S. Patent No. 6,185,598.

33. One fundamental issue in CDN technology is the delivery of resources from more than one content provider. The inventors of the '903 Patent understood that website owners and other content providers want to have their own internet domain shown in end users' browsers even when the content providers are using a CDN to serve their content.

34. The ’903 Patent is directed to CDN technology and to solving issues related to delivering resources from more than one content provider. In some embodiments, this involves replicating content from a source associated with a client of a CDN network onto CDN servers. End-user requests are then directed to the CDN servers instead of to the client’s source servers (generally referred to as “origin” servers). Embodiments of the ’903 Patent address issues such as load balancing and reducing traffic to client origin servers. When an origin server receives and must reply to multiple requests, delivery of content from a content provider can be slow. Methods of using CDN servers to deliver content as described in the ’903 Patent help to solve this issue.

35. The inventors devised the concrete solutions recited in the ’903 Patent to address unique problems related to providing content from multiple content providers using CDN technology. For instance, the inventors conceived of approaches to content delivery in a network for delivering resources associated with more than one content provider using a shared CDN server and alias names to provide resources in response to requests.

36. For instance, the ’903 Patent describes using a first/second alias name associated with a repeater server and a first/second resource of a first/second content provider and a table that can be consulted to determine a content provider associated with a particular resource. These innovations enabled CDN customers to

use their own domain name while using a provider's CDN services to serve content to end users.

A. Architecture and operation of "repeater server"

37. FIG. 1 of the '903 Patent illustrates a network environment including **reflector 108** at **origin server 102** and **repeaters 104a, 104b, and 104c**. '903 Patent, 4:22-26.

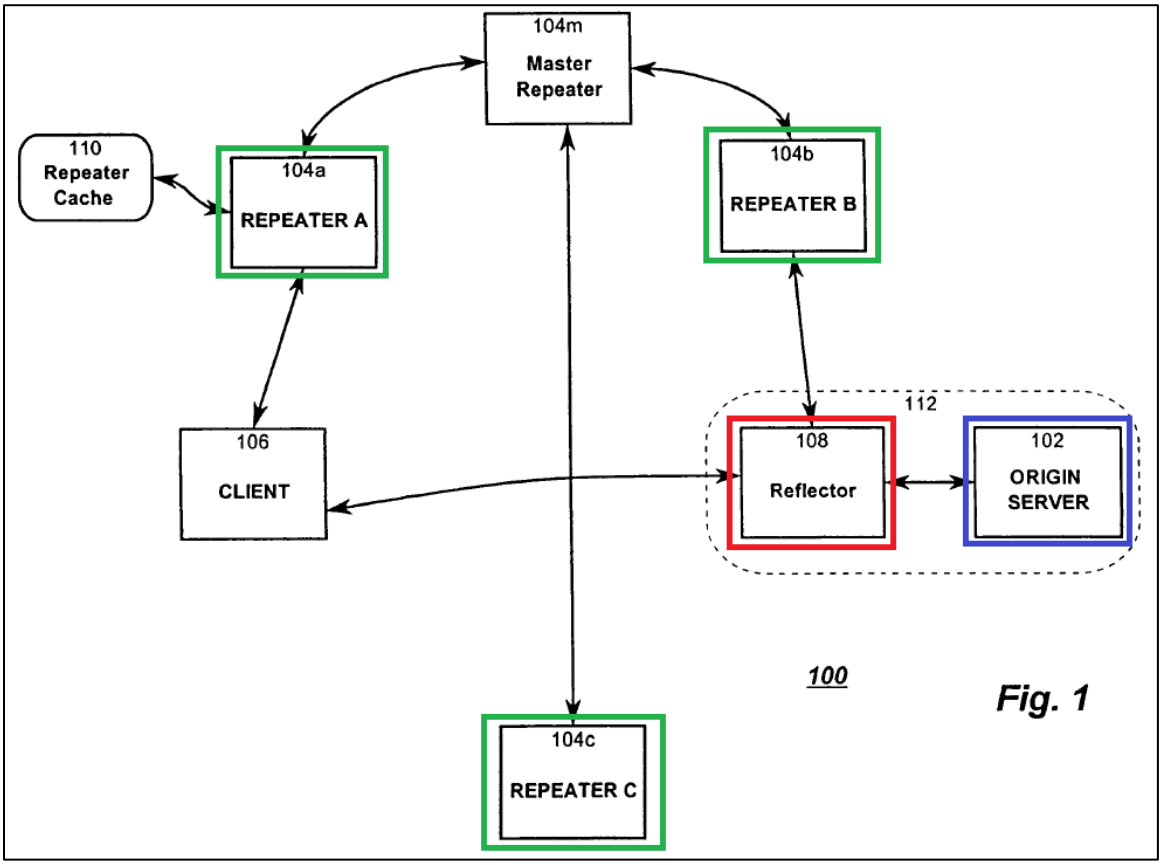

'903 Patent, FIG. 1 (annotated).

38. "Each **repeater 104a, 104b, and 104c** replicates some or all of the information on the **origin server 102** as well as information available on other origin servers in the network." *Id.*, 4:27-29; *see also id.*, 4:37-40 ("Thus, a repeater

can be considered as a dedicated proxy server that maintains a partial or sparse mirror of the **origin server 102**, by implementing a distributed coherent cache of the origin server.").

39. When a client requests resources from the origin server, **reflector 108** intercepts the request and either serves the request locally or reflects the request to a selected repeater. *Id.*, 5:9-22. **Reflector 108** is typically, though not necessarily, co-located with the origin server. *Id.*, 5:34-39. Essentially, **reflector 108** takes over **origin server 102's** IP address and port number so that when a client attempts to connect to **origin server 102**, the client connects to **reflector 108**. *Id.*, 7:40-49.

40. FIG. 3 of the '903 Patent illustrates the processing performed by a reflector. In **B1-B2**, the reflector analyzes a request to determine whether the requestor is a browser (i.e., received from a client) or a repeater. *Id.*, 7:59-60. A request from a repeater is served locally by the origin server. *Id.*, 7:62-63. This allows a repeater to build a cache of the content on the origin server because a repeater's requests to the origin server are always processed locally by the origin server.

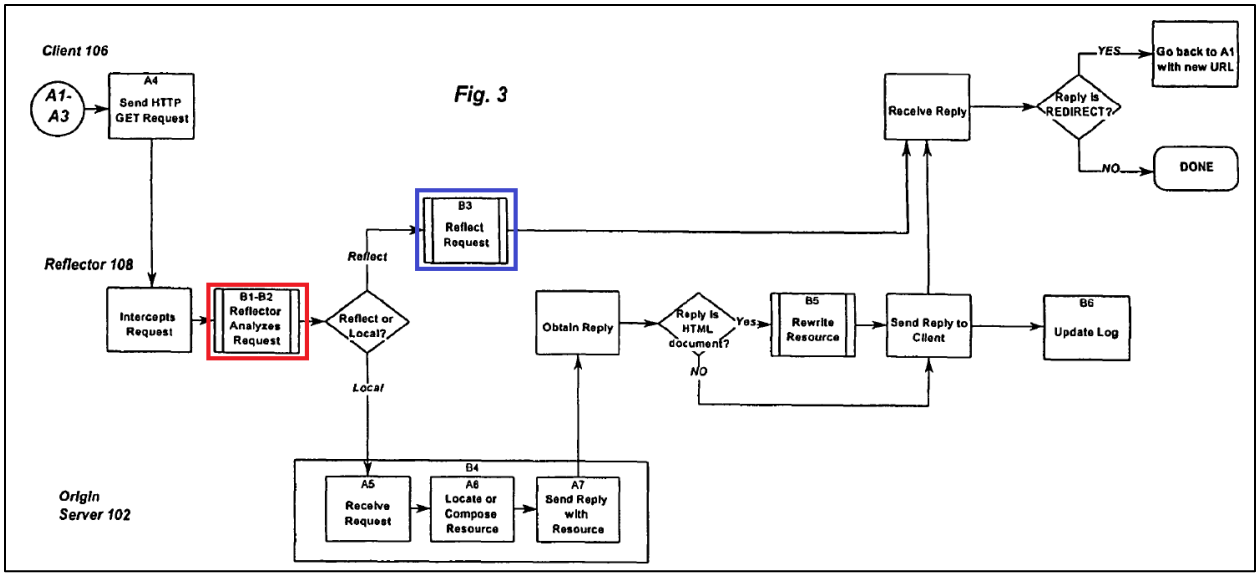

'903 Patent, FIG. 3 (annotated).

41. If the request is from a client, the reflector looks up the requested resource in a rule base (e.g., a list of regular expressions and related attributes) to determine whether the resource requested is "repeatable" and either reflects the request (by sending to a repeater) or serves the request locally. *Id.*, 8:4-9; *see also id.*, 9:14-15 ("By using a rule base . . ., it is possible to selectively reflect resources.").

42. In **B3** (FIG. 4 illustrates the sub-steps of B3 in further detail), the reflector "determines the best repeater to reflect the request to, [and] creates a new resource identifier (URL) (using the requested URL and the best repeater) that identifies the same resource at the selected repeater." *Id.*, 8:22-28. In an embodiment, the reflector creates a single URL that contains both the "URL of the original resources" and "the identity of the selected repeater." *Id.*, 8:29-33. For

example, for an HTTP request, the single URL may be:

http://<repeater>/<server>/<path>. *Id.*, 8:34-41.

43. The reflector sends a "REDIRECT reply containing this new URL to the requesting client." *Id.*, 8:50-51. When the browser receives the REDIRECT, the browser reissues a request to the repeater for the resource using the new URL. *Id.*, 10:12-19.

B. Architecture and operation of "alias name"

44. To facilitate the operations that I describe above, the '903 patent describes the use of "alias names" that are "*associated with*" a shared repeater server and enables resources requests to be "*directed*" to the repeater server. This allows a repeater server that services requests for multiple origin servers to be identified using a common identification. EX1001, 9:44-47 (emphasis added). As I further explain below, the '903 patent also describes and represents the "alias name" as a host name that identifies a repeater server.

45. Having reviewed the specification of the '903 patent, I note that the terms "alias" and "alias name" appear solely in the passages from columns 9 and 10 of the '903 patent. '903 Patent, 9:44-10:7. Specifically, the specification provides examples of "alias names" for different repeaters:

> For instance, if www.example.com is the origin server, names for three repeaters might be created:

> wr1.example.com
>
> wr2.example.com
>
> wr3.example.com
>
> The name "wr1.example.com" would be an alias for repeater 1, which might also be known by other names such as "wr1.anotherExample.com" and "wr1.example.edu".
>
> If the repeater can determine by which name it was addressed, it can use this information (along with a table that associates repeater alias names with origin server names) to determine which origin server is being addressed. For instance, ***if repeater 1 is addressed as wr1.example.com, then the origin server is "www.example.com"; if it is addressed as "wr1.anotherExample.com", then the origin server is "www.anotherExample.com".***

'903 Patent, 9:44-63 (emphasis added).

46. In view of these alias name examples, a POSA would have recognized that an "alias name" would be a "host name." As the '903 patent explains, a "host name" identifies a particular computer such as a "server computer." '903 Patent, 6:42-57. For example, the '903 patent provides "www.uspto.gov" as an example of a "host name. *Id.* I understand that Petitioner's expert acknowledged that the alias names described in the '903 all also follow this structure. EX2007, 20:15-21:12 ("the text is consistent with a host name"; a host name is "a string of text with dots in it" and the alias name examples are

"consistent with the form of a host name"); *see also id.*, 22:12-22 ("a domain name might be a part of a host name").

47. Based on this description, the specification of the '903 patent supports the understanding that a POSA would understood an "alias name" to refer to a host name that is "*associated with*" a shared repeater server and that enables requests to be "*directed*" to the shared repeater server.

48. Additionally, I note that the plain language of the claims also supports this understanding. For example, independent claims 1, 28, 37, 40, and 43 recite *first* and *second alias names* being associated with **distinct** content providers but being associated with **the same** *at least one shared repeater server*. Similarly, claim 37 recites "*associating a* ***plurality of alias names*** *with the at least one shared repeater server*." In other words, the same *at least one shared repeater server* is addressable using multiple alias names (e.g., either or both of the *first alias name* and the *second alias name*). *See* '903 Patent, 9:59-63; *see also id.*, claim 2 ("wherein the [at] least one shared repeater server is addressable using the first alias name and the second alias name."). In this manner, the claims also reflect the passages and understanding conveyed in the specification of the '903 patent.

49. Based on the '903 patent specification and claims, a POSA would have understood the term "**alias name**" to refer to a host name that is "*associated*

with" a shared repeater server and that enables requests to be "*directed*" to the shared repeater server.

C. Architecture and operation of a reflected request

50. In FIG. 5, the '903 Patent illustrates how a repeater intelligently processes a reflected request (e.g., based on an alias name). '903 Patent, 10:16-64. In C1, the repeater "analyzes the request to determine the network address of the requested client and the path of the resource requested." *Id.*, 10:20-22, 10:22-23 ("Included in the path is an origin server name.").

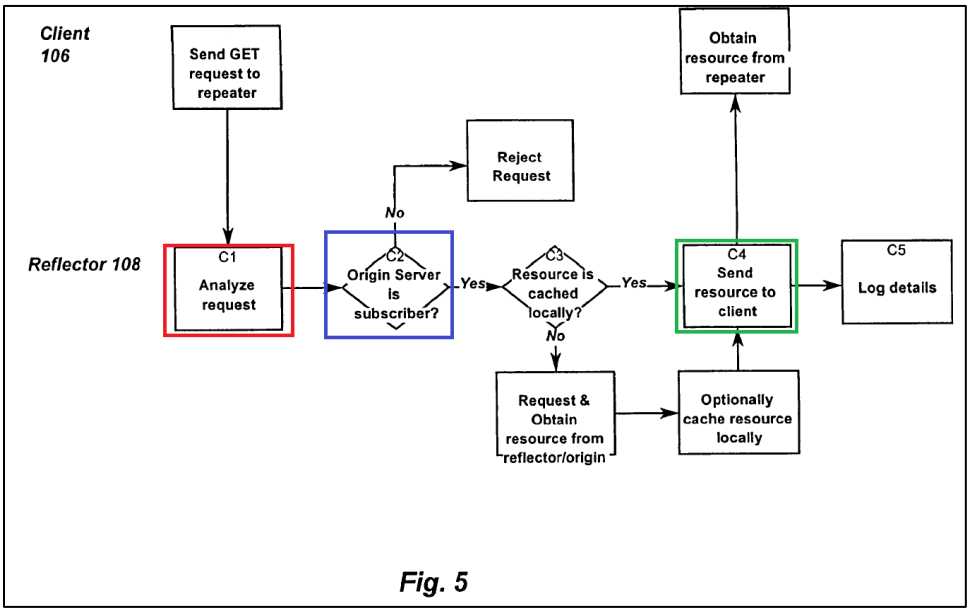

'903 Patent, FIG. 5 (annotated).

51. To determine the path of the resource, the repeater must know the name by which the client addressed the repeater. *Id.*, 9:56-59. A repeater determines the alias name used by inspecting the HOST tag in the request's HTTP header. *Id.*, 10:4-7.

52. For example, in **C2**, the repeater references a table that associates repeater aliases with origin server names to determine the requested origin server. *Id.*, 9:56-59. "For instance, if repeater 1 is addressed as wr1.example.com, then the origin server is 'www.example.com'; if it is addressed as 'wr1.anotherExample.com', then the origin server is 'www.anotherExample.com'." *Id.*, 9:59-63. The '903 Patent details this step:

> C2. The repeater uses an internal table to verify that the origin server belongs to a known "subscriber". A subscriber is an entity (e.g., a company) that publishes resources (e.g., files) via one or more origin servers. When the entity subscribes, it is permitted to utilize the repeater network. The subscriber tables described below include the information that is used to link reflectors to subscribers.

Id., 10:24-31. The '903 Patent provides additional detail about the subscriber table:

> When a new subscriber is added to the network, information about the subscriber is entered in a Subscriber Table at the master repeater and propagated to all repeaters in the network. This information includes the Committed Aggregate Information Rate (CAIR) for servers belonging to the subscriber, and **a list of the repeaters that may be used by servers belonging to the subscriber**.

Id., 17:51-57 (emphasis added).

53. Returning to FIG. 5, if the requested resource is not from a known subscriber, the repeater rejects the request. *Id.*, 10:32-33.

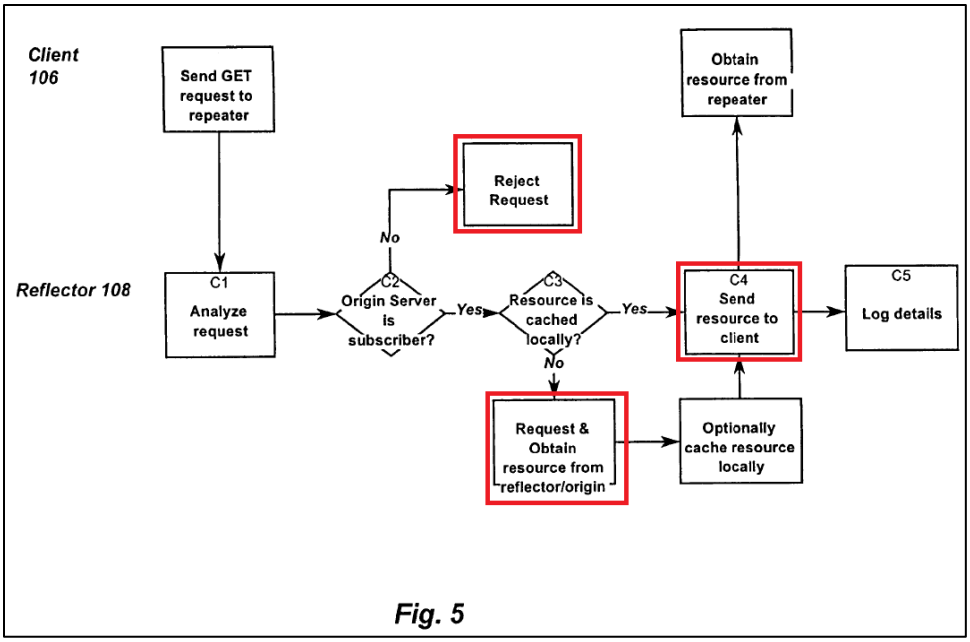

'903 Patent, FIG. 5 (annotated).

54. In C4, the repeater "constructs a reply including the requested resource (which was retrieved from the cache or from the origin server) and sends that reply to the requesting client." *Id.*, 10:54-57; *see also id.*, 10:42-44 ("Because this request to the originating reflector is from a repeater, the reflector always returns the requested resources rather than reflecting the request."). The client obtains the resource from the repeater, reducing the processing impact to the origin server. *Id.*, 2:62-65.

V. PROSECUTION HISTORY

55. The '903 Patent was examined as U.S. Application No. 11/065,412 ("'412 application"), which was filed on February 23, 2005. Applicant co-filed a First Preliminary Amendment that included one claim ("originally presented claim"). EX1002, 954, 1024.

56. On June 26, 2007, Applicant filed a Second Preliminary Amendment that cancelled the originally presented claim and introduced new claims 2-26 ("new claims"). *Id.*, 876-92. On July 12, 2007, January 25, 2008, and May 12, 2008, Applicant filed three separate Information Disclosure Statements. *Id.*, 788, 781, and 717. On August 5, 2008, Applicant filed a Third Preliminary Amendment that cancelled new claims 2-26 and introduced new claims 27-51 ("prosecution claims"). *Id.*, 687-96.

57. In a Non-Final Office Action mailed on September 18, 2008, Examiner Blair rejected the originally presented claim. *Id.*, 605-09. Recognizing that the Non-Final Office Action did not address the prosecution claims in the Third Preliminary Amendment, Applicant filed a Fourth Preliminary Amendment on December 10, 2008, that again cancelled all claims and re-advanced prosecution claims 27-51. *Id.*, 583-92, 599-600.

58. In a Non-Final Office Action mailed on January 9, 2009, Examiner Blair rejected prosecution claims 27-38, 43-45, 47, and 49-51 as anticipated by U.S. Patent No. 5,867,706 ("Martin") and claims 39-42, 46, and 48 as obvious over Martin in view of U.S. Patent Number 6,052,718 ("Gifford"). EX1002, 536-45.

59. On July 9, 2009, Applicant responded with amendments and remarks. *Id.*, 505-29. For example, Applicant amended claim 29:

29. (Currently amended) A ~~distributed hosting framework~~ content delivery system operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines, the ~~framework~~ system comprising:

at least one shared ~~content~~ 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, and to replicate at least some of the resources associated with a second content provider of said plurality of content providers; ~~and~~

~~at least one name server that provides domain name service (DNS) resolution,~~

wherein ~~resources~~ at least a first resource associated with the first content provider ~~are~~ is associated with a first alias name, said first alias name being associated with said at least one shared repeater server; and

wherein, ~~that causes the at least one name server to identify the at least one shared content server to serve the resources in response to~~ requests for the first resource ~~resources associated with the first content provider~~ made from said client machines are directed to the at least one shared repeater server, based at least in part on said first alias name; and

wherein ~~the resources~~ at least a second resource associated with the second content provider ~~are~~ is associated with a second alias name, said second alias name being associated with said at least one shared repeater server; and wherein

~~that causes the at least one name server to identify the at least one shared content server to serve the resources in response to~~ requests for the ~~resources~~ second resource ~~associated with the second content provider~~ 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.

EX1002, 506-07.

Applicant argued that Martin did not disclose "'at least one shared repeater server' that replicates at least some of the resources associated with a first content provider

and a second content provider" and that "Martin pertains to a server that serves resources only from and on behalf of a single content source (or content provider)." *Id.*, 520 (emphasis in original). Applicant also advanced new claims 53 and 54 that recited a "subscriber table", e.g., new claim 54 read:

> 54. (New) The content delivery system as in claim 45, further comprising:
>
> providing a subscriber table listing origin servers having content located thereon that 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; and
>
> wherein the at least one repeater server is further constructed and adapted to analyze, using the subscriber table, information about an alias name received with a client request for a resource to determine an origin server associated with the requested resource.

EX1002, 512-13.

60. On September 10, 2009, Examiner Blair issued a Notice of Non-Compliant Amendment that objected to additional text in the amendment between prosecution claims 51 and 52. *Id.*, 332-34. On September 11, 2009, Applicant filed a response to the notice of non-compliant amendment with a corrective amendment that removed the additional text. *Id.*, 319-28. The amendment was substantively similar to the claims filed on July 9, 2025, and again included new claims 53 and 54. *Id*.

61. On November 25, 2009, Examiner Blair issued a Restriction Requirement because prosecution claims 27 and 28 (Group I) were drawn to a

different invention than prosecution claims 29-42, 45-50, and 52-57 (Group II). EX1002, 303-08. On January 25, 2010, Applicant filed a response to the Restriction Requirement electing Group II. *Id.*, 297.

62. In a Final Office Action mailed on May 7, 2010, Examiner Blair rejected prosecution claims 29-38, 43-45, 47, and 55-57 as anticipated by U.S. Patent No. 6,085,193 ("Malkin") in view of Martin and claims 39-42, 46, and 48 as obvious over Malkin in view of Martin in view of Gifford. *Id.*, 241-50. Examiner Blair noted that claims 53 and 54 contained Allowable Subject Matter. *Id.*, 249.

63. On July 2, 2010, Applicant responded with amendments and remarks. *Id.*, 207-26. Applicant amended the independent claims to incorporate the subject matter deemed allowable by Examiner Blair in prosecution claims 53 and 54:

29. (Currently amended) 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:

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, and to replicate at least some of the resources associated with a second content provider of said plurality of content providers;

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,

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; ~~and~~

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; ~~and~~

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; ~~and~~

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

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.

EX1002, 210.

64. On October 9, 2012, the Examiner issued a First Notice of Allowance including an Examiner's Amendment that added "wherein the second content provider is distinct from the first content provider" to the independent claims:

> 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, 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;

EX1002, 82.

Examiner Blair noted that the closest prior art (Armbruster) "teaches [] a table listing aliases associated with content providers . . . [but] the cache servers in Armbruster do not perform the claimed analysis using the table." *Id.*, 94.

65. On October 17, 2012, Applicant paid the Issue Fee. *Id.*, 76. On October 31, 2012, Applicant filed a Petition to Withdraw an Application from Issue After Payment of the Issue Fee. *Id.*, 60-61. Applicant co-filed an Information Disclosure Statement listing a European Search Report. *Id.*, 62-64.

66. On April 29, 2013, Examiner Blair issued a Second Notice of Allowance. EX1002, 16. The Second Notice of Allowance also included the Examiner's Amendment recreated above. *Id.*, 23-34. On May 15, 2013, Applicant again filed the Issue Fee Payment form. *Id.*, 13-14. The '903 Patent issued on July 2, 2013. *Id.*, 2.

VI. OVERVIEW OF KENNER

67. Kenner discloses a "video clip storage and retrieval system" in which "video clips . . . can be requested and retrieved by a user" at a multimedia terminal. EX1005, Abstract; *see also id.*, 1:12-16, 6:20-22. Because "video files can be very large," Kenner recognizes a need "for a system capable of providing improved access to audio/video content on the Internet . . . ." *Id.*, 3:15-17, 4:5-7.

68. FIG. 1 of Kenner illustrates a video clip storage and retrieval system. *Id.*, 10:18-20. Kenner's system includes **user terminal 14**, **data sequencing interface ("DSI") 30**, **extended storage and retrieval units ("extended SRUs") 26**, and one or more index managers ("IM") including **primary index manager ("PIM") 22**.

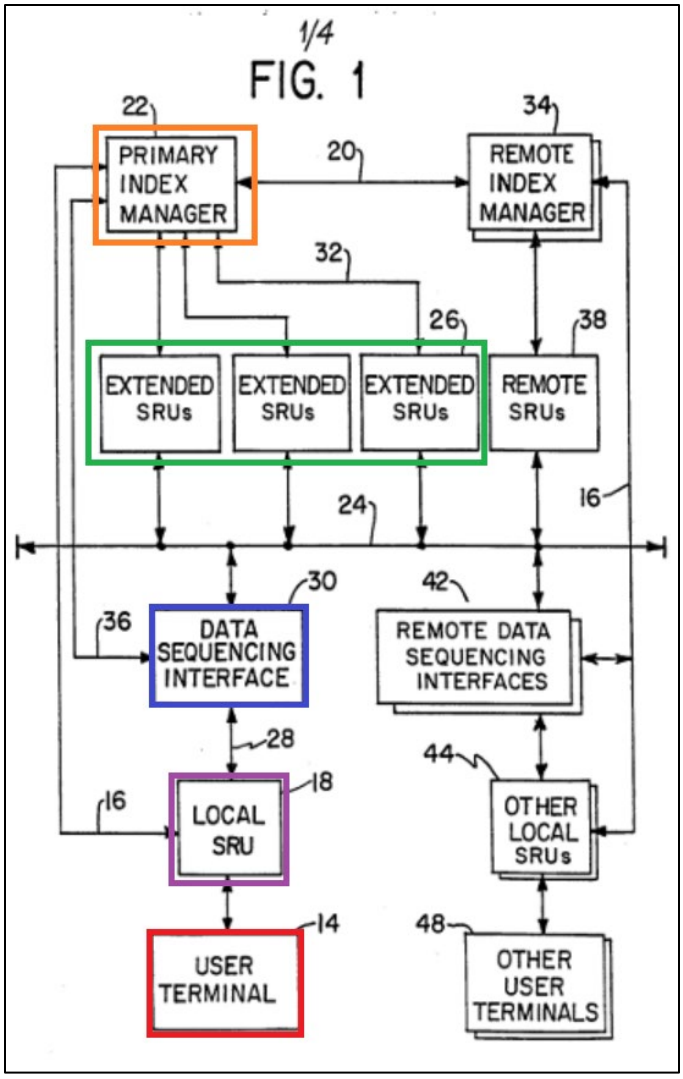

EX1005, FIG. 1 (annotated).

69. **User terminal 14** allows a "user to manipulate retrieved video clips"—e.g., play, stop, pause, fast forward, etc. *Id.*, 13:17-20. **User terminal 14** also "provides the user access to a database or index which can be interrogated for desired video clips and other information." *Id.*, 13:13-24.

70. **Local SRU 18** provides large capacity hard drives for temporarily storing video clips. *Id.*, 13:31-14:19. For example, **local SRU 18** may be a "file

server for a local area network, with one or more integral or connected storage devices." *Id.*, 14:3-5.

71. **PIM 22** is the "primary search engine and database management module of the invention." *Id.*, 16:2-3, 31:19. **PIM 22** "locates the requested video clips" and directs "the efficient download of the video clips to the user terminal 14." *Id.*, 10:27-28.

72. **Extended storage retrieval units ("SRUs") 26** provide the "principle storage facility for the system and [are] used to store audio-visual data in a plurality of audio-visual storage media." *Id.*, 17:30-32. **SRUs 26** transmit requested data to **user terminal 14**. *Id.*, 18:13-15.

73. **DSI 30** "manages the download of video clips and other information to **local SRU 18** from the various locations." *Id.*, 15:4-6. "A **DSI 30** is created and/or initialized by **PIM 22** whenever a user requests audiovisual information that is not stored within the **local SRU**." *Id.* 19:4-5; *see also id.* 31:26-28 ("[O]ne or more transient DSI processes 58 . . . may be created by the **PIM 64**, as required, for each user receiving audio-video content.").

74. FIG. 4 of Kenner illustrates a "preferred embodiment of an Internet-related video clip storage and retrieval system." *Id.*, 31:2-3. In this embodiment, "**terminal 50** is a personal computer running an HTML browser 82 with an audio-video decoding and playback 'browser extension' 84." *Id.*, 31:10-12.

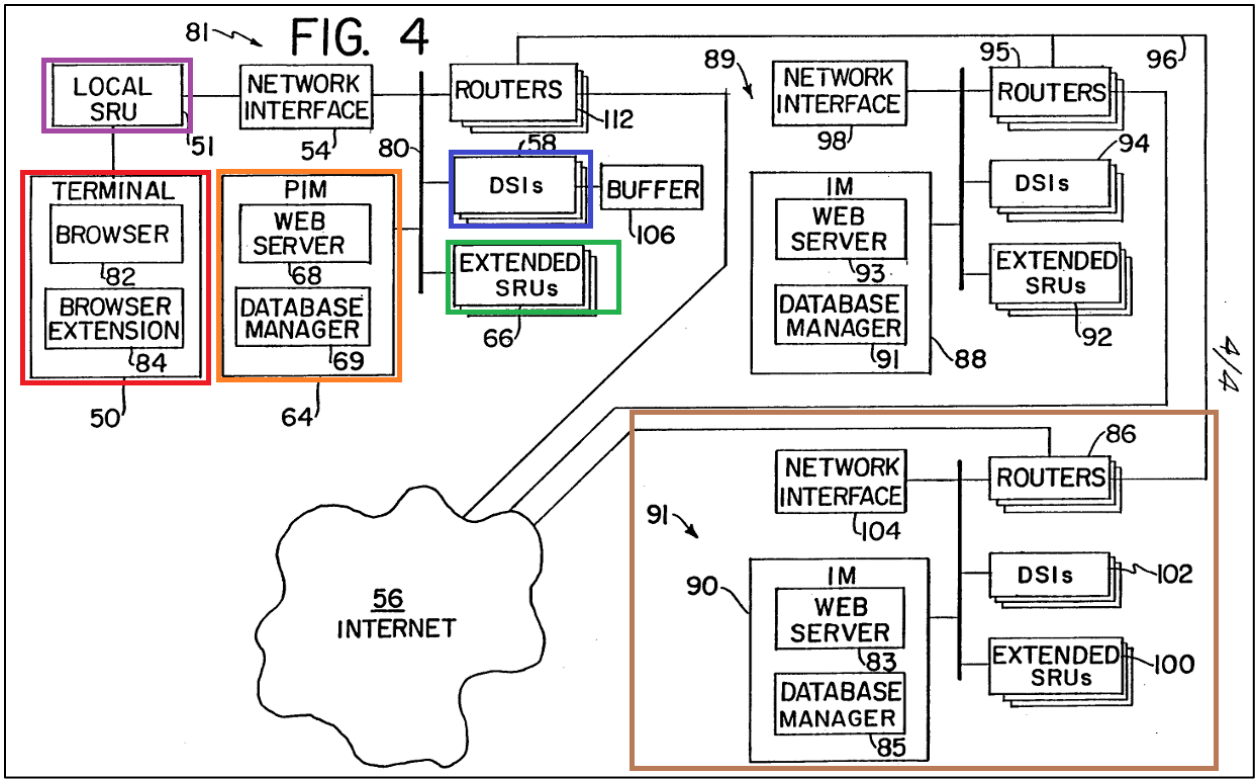

EX1005, FIG. 4 (annotated).

Content provider's region 91 includes an "IM 90 hosting a Web server 83 and a database manager 85, at least one extended SRU 100, one or more transient DSIs 102, and router 86." *Id.*, 32:14-16.

75. A content provider's server creates web pages for the user at **terminal 50** that inform the user of available content. *Id.*, 35:17-36:9. "A reference to a desired clip is embedded within the HTML of the Web page"—e.g., a video ID that is unique identifier for the selected clip referenced within an HTML EMBED tag. *Id.*, 36:15-18. The video ID includes:

> a text name of the file as defined by the content provider; the content provider's account number as provided by the organization running the

> subscription service; a category coordinate, possibly a representation of a hierarchical portion of a category tree, a geographic coordinate used to determine where the file is relevant (e.g. a region, state, or city); a time stamp, and a time period over which the file is relevant.

Id., 36:32-37:5; *see also id.*, 36:28-29 ("The video ID consists of a multidimensional set of content-characterization coordinates plus a unique file name."), 44:1-14 (discussing creation of the video ID when the content provider uploads a clip).

76. When the user selects a video clip, **PIM 64** receives a "virtual URL" that takes the form: "'http://' plus the Internet address of the **PIM 64**, plus the user's subscriber ID number, plus the video ID." *Id.*, 37:7-10; *see also id.*, 36:22-24 (noting that the browser invokes local SRU 51 to receive the video data). When **PIM 64** receives the virtual URL, "the virtual URL is decomposed into the video ID and subscriber ID components, which are then used to access the PIM's internal databases." *Id.*, 37:15-17.

77. **PIM 64** "maintains information on the audio-visual clips stored on its extended SRUs 66 in a clip database." *Id.*, 34:3-5. The database can contain the video ID and other information for each clip:

Item Name	Format	Description
Counter	numeric	Primary index for the records. Each record represents one video clip.
Video ID	text	The globally unique name of the video clip, as specified above.
Extended SRUs	IP array	The IP addresses of all the extended SRUs 66 which contain the file.
Copyright	boolean	A flag to indicate that the file is copyrighted and must be protected.
Charge Mechanism	numeric	A code representing the mechanism for charging for the file (pay per view, one-time fee, etc.).
Charge Parameters	numeric array	The amount charged per use under the specified charge mechanism.
Expiration Date	date	Date after which the file is to be removed from the system.
Size of File	numeric	Size of the file, in bytes.
Date	date	Date the file was made by the content provider.
Time	time	Time the file was made by the content provider.
Category	numeric	The subject category of the file, used for load projections as discussed below.
Usage Count	numeric array	The historical frequency of clip access across days and hours, used for load projections.
Segment Info	text/numeric	If the file is segmented, this is the array of segment descriptors and pointers into the file.
Link Info	text/numeric	If the file has been annotated with links to other files, this is the array of link names, URLs, and pointers into the file.

Id., 34:5-35:3.

78. The PIM "checks the user's subscription rights in its user database, and if authorized and necessary, initiates a DSI process 58 to download the desired clip to the user's terminal 50." *Id.*, 37:18-20; 9:5-8 (noting that a user may establish a subscription account with a content provider or ISP). The PIM "queries

its clip database to determine on which extended SRUs 66 the desired clip is stored." *Id.*, 38:32-39:1. The DSI then selects the least loaded SRU from the list and oversees the transfer of the video file from the selected SRU to the terminal 50. *Id.*, 39:6-9.

79. However, if PIM 64 determines that none of the extended SRUs 66 has the desired clips, PIM 64 will query the closest IMs to determine if any of them have the desired clip. *Id.*, 40:29-34. If none of the neighboring IMs has the desired clip, the PIM 64 will contact the source IM 90, where the content provider first uploaded the file. *Id.*, 41:13-16.

80. FIG. 3 of Kenner illustrates how "the invention would operate to search and download data." *Id.*, 25:7-8. A user "builds a data query at the user terminal 14 from the text database." *Id.*, 25:8-9; *see also id.*, 6:32.

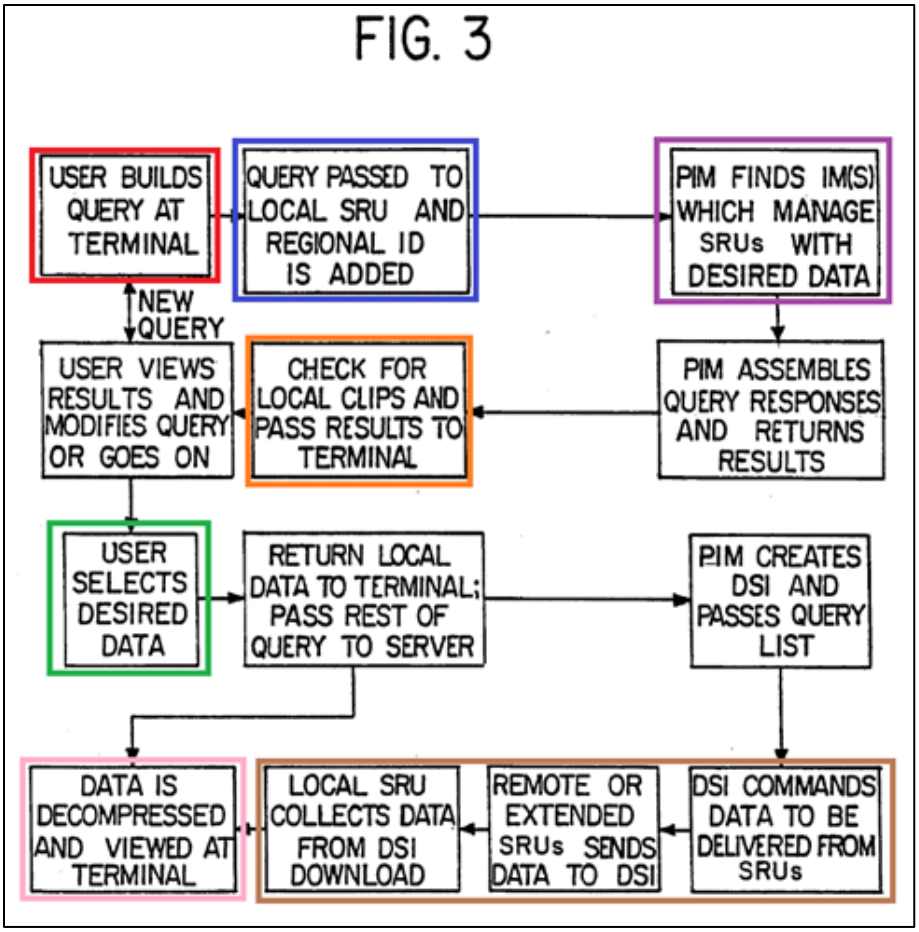

EX1005, FIG. 3 (annotated).

The local SRU then "(1) attaches a regional identifier to the query; and (2) searches its own database and flags each request that is stored at the local SRU . . . ." *Id.*, 25:11-14; *see also id.* 7:3-5 ("The local SRU provides temporary storage for the user's most requested video clips, and before the query is sent to the user's PIM, the local SRU is polled for requested video clips."). The PIM identifies locations of the audio-visual data stored at the extended and remote storage units and "downloads a list of all available video clips to the user's terminal 14." *Id.*, 25:14-22; *see also id.*, 6:32-7:1 ("The request is transmitted to the user's ("PIM")

via a local storage and retrieval unit (local SRU).") The user then "identifies and selects individual records or groups of records for further viewing and manipulation." *Id.*, 25:25-27. The data is retrieved from the storage units via the DSI and downloaded to the user terminal for viewing. *Id.*, 25:27-26:7; *see also id.*, 7:18-20 ("The DSI collects the requested video clips from the appropriate extended and remote SRUs and transmits this information to the local SRUs."); 7:21-22 ("The requested video clips . . . are then displayed at the user's terminal.").

VII. UNDERSTANDING OF RELEVANT LEGAL STANDARDS

81. In formulating my opinions set forth in this Declaration, I applied the following legal principles.

A. Person of Ordinary Skill in the Art ("POSA")

82. I understand that patent law analyses are performed from the viewpoint of a POSA. I understand that a POSA is a hypothetical person who is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity. For purposes of this declaration, I have applied Petitioner's articulation of a POSA. *See* Pet., 5; EX1003, ¶48 ("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"). It is my opinion that the definition of a POSA would not materially affect the limited analysis required for deciding institution and presented in this declaration. If called to give further opinions in this proceeding, I reserve the right to reconsider the definition of a POSA and arrive at my own conclusion.

B. Claim Construction

83. I understand that in an IPR, patent claim terms are given their ordinary and customary meaning as understood by a POSA at the time of the invention, in view of the patent's specification and its prosecution history—unless the patent explicitly defines the claim term. I also understand that when a patent explicitly defines a claim term in the specification, the patent's definition controls.

84. I understand that a dependent claim contains all limitations of the claim from which it depends.

C. Anticipation

85. I have been informed and I understand that a patent claim is anticipated when a single piece of prior art describes every element of the claimed invention, either expressly or inherently, arranged in the same way as in the claim. I have been told to consider elements to be inherently described only if they are necessarily present in the piece of prior art.

D. Obviousness

86. I understand that an obviousness analysis involves comparing a claim to the prior art to determine whether the claimed subject matter would have been obvious to a POSA in view of the prior art and general knowledge in the prior art at the time of the patent's priority date.

87. I understand that obviousness must be assessed from the viewpoint of a POSA at the time of the invention. I further understand that to establish obviousness, a party must perform the following factual inquiries: (a) determining the scope and content of the prior art; (b) ascertaining the differences between the claimed invention and the prior art; and (c) resolving the level of skill in the art. I understand that, in determining the scope and content of the prior art and ascertaining the differences between the claimed invention, a patent challenger must specify where each element of the claim is found in the prior art.

88. I understand that one way of establishing obviousness is by establishing that a POSA would have had both (i) a reason to modify or combine the teachings of the prior art to achieve the claimed invention and (ii) a reasonable expectation of success in doing so. I understand that the reason to combine prior art references can come from a variety of sources, not just the prior art itself or the specific problem the patentee was trying to solve. And I understand that the references themselves need not provide a specific hint or suggestion of the

alteration needed to arrive at the claimed invention; the analysis may include recourse to logic, judgment, and common sense available to a POSA that does not need to be explicit in any reference.

89. I understand that a "reasonable expectation of success" is assessed in view of the prior art and general knowledge in the art from the viewpoint of a POSA at the time of the patent's priority date, and it does not require an absolute certainty of success. Furthermore, I understand that determining whether the prior art would have suggested to a POSA to carry out the claimed invention with a reasonable expectation of success must be done after determining that the claimed elements are present in the prior art.

90. I also understand that before reaching a conclusion that the claimed invention would have been obvious one must consider any objective evidence of non-obviousness if it is available. The objective evidence of non-obviousness can include evidence of commercial success attributable to the claimed invention, evidence of industry praise for the claimed invention, evidence of a long-felt need that was solved by the claimed subject matter, evidence that others copied the claimed subject matter, or evidence that the claimed subject matter achieved an unexpected, superior result relative to the closest prior art. I understand that such evidence must have a nexus, or causal relationship, to the claimed subject matter

beyond what was available in the prior art and must be commensurate in scope with the patent claim(s) at issue.

VIII. DETAILED DISCUSSION OF OPINIONS

91. Upon considering the materials listed above and discussed herein, based on my extensive experience and education in the relevant art, in view of my understanding of the governing legal principles and the arguments presented in the Petition, I have formed the following opinions.

A. Kenner's "video ID" does not teach the claimed "alias name" (All Grounds; All Challenged Claims).

1. Kenner's video ID is not an alias name because (a) it is not associated with a repeater server and (b) it is not a host name.

92. While Petitioner argues that Kenner's video ID teaches the claimed *alias name*, I note that the video ID is not associated with any repeater server and also additionally is not a host name. Pet., 17-18. Based on this understanding, Kenner's video ID does not teach the claimed "alias name."

a. Kenner's video ID is not associated with a repeater server

93. As I explained in Section IV.B, a POSA would have understood the claimed "**alias name**" to be a host name "*associated with*" a shared repeater server and that enables resource requests to be "*directed*" to the repeater server.

94. With reference to Claim 1, the Petition argues that Kenner's video ID teaches the claimed *first/second alias name*. Pet., 19-23. Kenner's video ID is the

"globally unique name of [a] video clip." EX1005, 34:8-9. Kenner, however, refers to the video ID as a "unique reference identifier for each video clip" that "corresponds to an identifying field within [a] text database." EX1005, 22:1-2, 43:27-29. The video ID includes:

> a text name of the file as defined by the content provider; the content provider's account number as provided by the organization running the subscription service; a category coordinate, possibly a representation of a hierarchical portion of a category tree, a geographic coordinate used to determine where the file is relevant (e.g. a region, state, or city); a time stamp, and a time period over which the file is relevant.

Id., 36:32-37:5.

95. As seen from this passage, however, there is no teaching in Kenner's video ID that associates the video ID with a repeater server. Further, nothing in the video ID identifies the specific PIM (PIM 64) that Petitioner has identified as the claimed *at least one shared repeater server*. I also further explain in Section VII. A.2, that Kenner contemplates another component performing this work instead: a "virtual URL." The "virtual URL" contains the "Internet address of the PIM 64" that a request is directed to. EX1005, 37:7-17. Kenner's video ID therefore is not associated with a repeater server and does not teach an "alias name" as recited in the claims. Section IV.B.

96. I note that Petitioner argues that when an IM stores a particular clip, that IM becomes associated with the video ID. Pet., 20-21. I disagree with this conclusion. Even when storing a particular clip, the video ID is not associated with that particular IM because the video ID categorizes clip content. EX1005, 36:29-31, 41:1-6. This is because the video ID includes content coordinate data that generally identifies the types of *files* found in IMs rather than IMs. *Id.*, 36:29-31. For example, a video ID identifying "'news:sports:baseball' clips" simply identifies content. *Id.*, 38:3-8. In this manner, there is no association between the video ID and the IM that stores a particular clip. I understand that Petitioner's expert even acknowledged that Kenner's video ID's "content coordinate data," which is used to identify IMs *likely* to have a clip, might result in the return of "one, or more than one, or ***even possibly zero other IMs***." EX2007, 126:5-12. With this understanding, simply storing a clip does not create an association between the video ID and an IM.

b. Kenner's video ID is not a host name

97. I also note that Kenner's video ID is not a "host name" as the '903 patent describes as a characteristic of an alias name. I understand that Petitioner's expert confirmed this understanding during a deposition and stated that "I don't think one would normally characterize [a video ID] as being a host name." EX2007, 129:16-22. I agree with this assessment. As I previously explained, the

claimed "alias name" is a host name, such as "wr1.example.com" or "wr1.anotherExample.com." EX1001, 9:44-63. Such a host name corresponds to a repeater server. EX1001, 6:45-50. ***First***, Kenner does not refer to its video ID as a host name. Rather, Kenner's video ID is a "globally unique name of [a] video clip." EX1005, 34:8-9. There is therefore no suggestion in Kenner that a video ID is a host name. ***Second***, Kenner's video ID also does not follow a host name structure. EX1005, 36:28-37:5. Rather, Kenner's "video ID consists of a multidimensional set of content-characterization coordinates plus a unique file name" along with a "text name of the file" and several other parameters. EX1005, 36:28-37:6. This differs from the host name structure described in the '903 patent and by Petitioner's own expert. EX1001, 6:42-57; EX2007, 20:5-21:1. Based on this understanding, Kenner's video ID is not a host name and also does not use the same structure as a host name.

98. Thus, because Kenner's video ID is not associated with a repeater server and is not a host name, Kenner's video ID does not teach the claimed "alias name."

2. Kenner's virtual URL additionally demonstrates why Kenner's video ID is not an alias name

99. Kenner's video ID does not disclose the claimed *alias name* in view of Kenner's virtual URL. Because Kenner's virtual URL explicitly identifies a particular IM using an "Internet address" and directs requests to that IM based on

the "Internet address," a POSA would not have viewed the video ID as providing these functionalities. EX1005, 37:7-10. In this manner, the video ID does not teach the claimed *alias name* because the video ID does not provide the same *alias name* functionality recited in the claims (e.g., identification of *at least one shared repeater server* or directing a request to the *at least one shared repeater server* based on the alias name).

100. I note that Kenner's "'virtual URL' is constructed in the form 'http://' ***plus the Internet address of the PIM 64***, plus the user's subscriber ID number, plus the video ID." EX1005, 37:7-10 (emphasis added). Regardless of whether a local SRU 51 has previously stored a video clip, the local SRU 51 will pass the video ID to PIM 64 using the virtual URL. EX1005, 36:21-27, 37:7-10. Specifically, the Internet address of the PIM 64 is used to identify PIM 64 and pass along the video ID to PIM 64. EX1005, 36:21-27, 37:15-17; EX2007, 119:13-17 (Dr. Mowry describing the Internet address as a way for local SRU 51 to talk to PIM 64). "Upon receipt by the PIM 64, the virtual URL is decomposed into the video ID and subscriber ID components, which are then used to access the PIM's internal database." EX1005, 37:15-17.

101. In view of this description, Kenner's virtual URL and its included "Internet address" are the mechanisms used to identify a PIM and direct resource

requests to the PIM.[1] I further explain this issue in Section VIII.B. Kenner's video ID does not perform the function of directing a resource request from the client machine to the PIM. Based on the language in Claim 1 of the '903 patent, this means that Kenner's video ID functionality does not teach the claimed *said alias name being associated with said at least one shared repeater server* or that *requests* for a particular resource (e.g., video clip) are *directed to the at least one shared repeater server based at least in part on said first alias name*. EX1001, Claim 1; *see also id.*, claims 28, 37, 40, and 43. Specifically, Kenner's "Internet address" is associated with a particular IM (e.g., PIM 64)—not Kenner's video ID. Additionally, requests from a client machine are directed to the particular IM using the IM's Internet address—not Kenner's video ID. Kenner's video ID therefore does not perform the same functionality as the *alias name* recited in the claims. EX1001, Claim 1. Thus, Kenner's video ID does not teach the claimed *alias name*.

B. Kenner does not direct resource requests to "the at least one shared repeater server, based at least in part on" an alias name

[1] I note that I am not providing an opinion on whether the virtual URL is an alias name. Rather, I note that the virtual URL is the mechanism used to identify a PIM rather than that the video ID.

(All Grounds; Independent Claims 1, 28, and 37 and corresponding dependent claims).

102. Kenner does not describe directing resource requests to "**the** at least one shared repeater server, ***based* at least in part on***" an alias name. Because (1) Kenner directs all content requests from terminal 50 to PIM 64 regardless of the content of a video ID and (2) Kenner's description of a PIM 64 querying other IMs does not address directing a resource request received from a client machine, Kenner does not describe directing client machine resource requests to "***the*** at least one shared repeater server, ***based* at least in part on***" an alias name.

103. Based on my review, the claims require directing requests from client machines to a particular shared repeater servers "based at least in part on" an alias name:

- 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***;
- 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***;

EX1001, Claim 1.

104. Similarly, claims 28 and 37 recite directing these requests directly to the at least one shared repeater server based on the alias name so that the shared repeater server receiving the request can return the requested resource:

- 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***
- 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***,

EX1001, Claim 28; *see also id.*, Claim 37.

105. In this manner, the claims recite a requirement that the alias name impacts where to initially direct a resource request from a client machine. That is, a received alias name must be examined to identify ***the*** *at least one shared repeater server* (and not any other repeater server) that receives the original request from the client machine. Claims 28 and 37 underscore this understanding by stating that ***the*** *at least one repeater server* is the same server that (1) receives the request from the client based on the alias name and (2) returns the requested resource to the client. Based on the claims' reference "the at least one shared repeater server" and

"said at least one repeater server," a POSA would have recognized the claims as indicating that the particular "at least one repeater server" that is identified must be the same one that performs the claimed functionality and is the same one that is identified from the alias name. Kenner does not describe this functionality.

106. ***First***, Kenner does not teach directing requests to the PIM (i.e., Petitioner's mapping for *at least one shared repeater server*) based on an alias name because Kenner's video ID is not an alias name. Section VIII.A.1. As I explained in Section IV.B, the '903 patent describes directing a request to the repeater server based on an alias name associated with the repeater server. EX1001, 9:44-63. Kenner's video ID includes no such association with a repeater server. Section VIII.A.1. Because Kenner does not describe an alias name being associated with the claimed shared repeater server, Kenner does not direct any client requests ***based on*** an alias name.

107. ***Second***, Kenner does not direct a request to "*the at least one shared repeater server*" (i.e., Kenner's PIM 64) ***based at least in part on*** an alias name because ***all*** of a particular client's requests are sent to that client's assigned PIM (i.e., Petitioner's alleged "*at least one shared repeaters server*"). Pet., 13; EX2007, 113:1-5 (Dr. Mowry acknowledging that all requests from terminal 50 are sent to PIM 64). In view of the claim language, a POSA would have understood that for such claim language to be satisfied, this would require an identification of the PIM

64 specifically from among other IMs and directing requests specifically to PIM 64 because of the video ID. Because Kenner does not describe a video ID dictating that a particular client request (e.g., from terminal 50) should be directed to PIM 64, Kenner does not describe the claimed functionality.

108. As seen from Figure 4, which I have reproduced below, Kenner describes all requests from terminal 50 being directed to PIM 64:

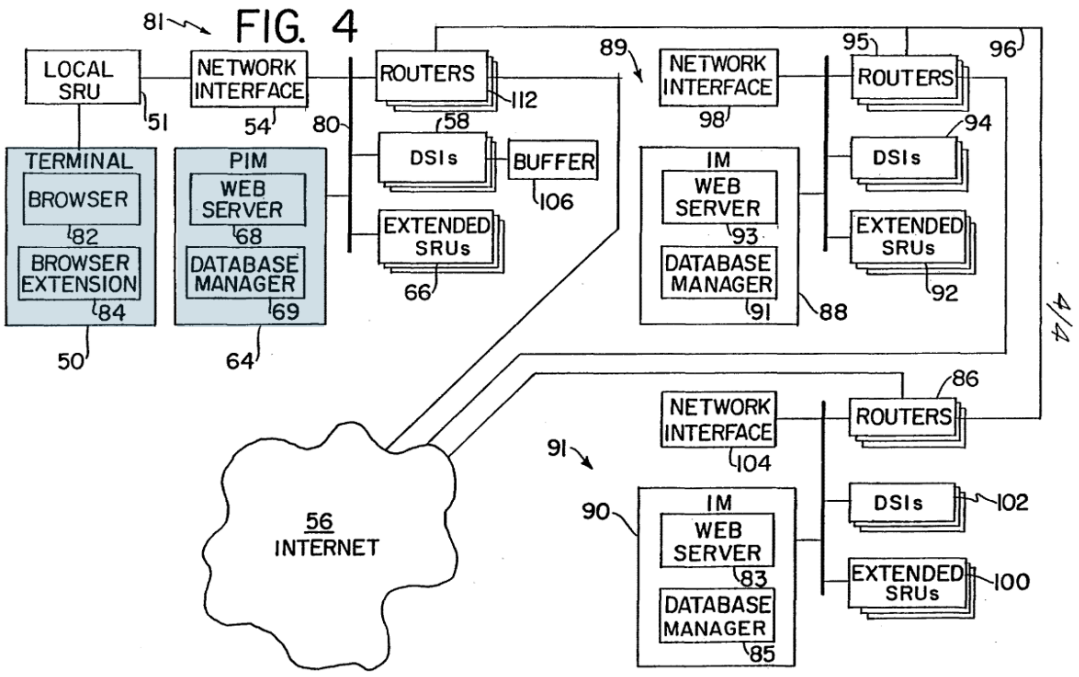

EX1005, FIG. 4 (annotated).

109. For example, when a terminal 50 requests a desired video clip, "the local SRU 51 belonging to the terminal 50 and the browser extension 84 are invoked to receive the data." EX1005, 36:21-22. Kenner continues and explains:

> First, the local SRU 51 intercepts a video ID, a unique identifier specifying the selected clip, which is stored within the EMBED field in

> the Web page. ***The local SRU 51 first determines if the desired clip is already stored locally. If not, the local SRU 51 passes the video ID to the PIM 64 associated with the user's terminal 50***. The local SRU 51 then awaits authorization from the PIM 64 to proceed with a data transfer.

EX1005, 36:22-27 (emphasis added).

110. Even if the desired clip is already located on local SRU 51, Kenner still passes the video ID to PIM 64 via a "virtual URL." EX1005, 37:7-17. "If the desired clip was located on the local SRU 51," local SRU 51 passes the virtual URL to PIM 64 so that PIM 64 can decompose the virtual URL into "video ID and subscriber ID components, which are then used to access the PIM's internal databases." EX1005, 37:7-17.

111. In this manner, Kenner describes directing all requests from terminal 50 to PIM 64 regardless of the content of the video ID. There is no examination of the content of the video ID to determine that PIM 64 (from among other IMs) should receive a particular client request from terminal 50. Further, Kenner does not describe *using* the video ID to determine that PIM 64 should receive the request from terminal 50 either. Based on this description, Kenner's video ID does not impact the decision of which server should receive the request at all. Rather, PIM 64—i.e., the server that Petitioner has identified as the *shared repeater server*—always receives client requests from terminal 50. EX1005, 36:22-27, 37:7-

17. Thus, because Kenner's video ID is not used to determine that a client request should be directed to PIM 64, Kenner does not describe directing resource requests to "the at least one shared repeater server, ***based at least in part on***" an alias name.

112. Appearing to recognize this deficiency, I note that the Petition presents an alternative theory where a request is directed to a primary IM (PIM) from a "neighboring IM." Pet., 21. Petitioner argues that if a given IM does not identify a desired clip, it will query the closest IMs to determine if a remote IM has the desired clip. Pet., 21 (citing EX1005, 40:29-33). Petitioner argues that the IM will use content coordinate data included in the video ID to identify neighboring IMs, like the primary IM. *Id.*

113. But Petitioner's alternative theory is also inconsistent and conflicts with the claim language. For example, as recited in Claim 1, requests for the first or second resource "***from said client machines***" are directed to ***the*** at least one shared repeater server ***based at least in part on*** the first or second alias name. EX1001, Claim 1. In this manner, the claims refer to handling the requests that are directly received from the client machines. Upon receiving those requests from the client machines, "*the at least one shared repeater server*" is identified to handle the request based on the particular "*alias name*" corresponding to the request.

114. Kenner's description of communication occurring between neighboring IMs does not address the specific functionality recited in the claims.

Specifically, a PIM querying another IM does not address the original handling of a request received from a client machine (e.g., Kenner's terminal 50). EX1005, 40:29-33, 41:1-5. While a PIM may query other IMs, a client terminal's corresponding PIM still handles all requests received from the client terminal. EX1005, 36:22-27, 37:7-17 (describing the passing of the video ID to PIM 64 regardless of whether a local copy of a desired file is already stored or not). For example, PIM 64 still handles all requests sent by terminal 50. In this manner, even if PIM 64 queries a neighboring IM—or even if a neighboring IM queries PIM 64—this functionality still does not teach the specific claim language requiring requests ***from said client machines*** are directed to ***the*** at least one shared repeater server ***based at least in part on said first alias name***.

115. Because Kenner does not describe (1) using the video ID to determine where to direct the original requests received from terminal 50 or (2) using the video ID to identify PIM 64 as "the at least one shared repeater" meant to receive the request, Kenner does not teach the claimed directing of requests from ***said client machines*** to the at least one shared repeater server ***based at least in part on*** an alias name.

IX. CONCLUSION

116. In signing this declaration, I recognize that the declaration will be filed as evidence in an *inter partes* review before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If cross-examination is required, I will appear for cross-examination within the United States during the time allotted.

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

Executed on this 13th day of February 2026, in Los Angeles, CA.

Prashant Shenoy

Dr. Prashant Shenoy

X. APPENDIX A: COMPARISON OF INDEPENDENT CLAIMS

Independent Claims Comparison

first/second alias name alias -> repeater association direction based on alias name *table* use of table

Claim 1	Claim 28	Claim 37	Claim 40	Claim 43
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: 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, 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; 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, wherein at least a first resource associated with the first content provider is associated	28. A method, in a content delivery system operative in a computer network for delivering content to client machines, the computer network comprising a plurality of origin servers, each of said origin servers having resources associated therewith, and 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: 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; 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	37. A method, in a content delivery system operative in a computer network for delivering content to client machines and comprising at least one shared repeater server operable to replicate resources stored on a plurality of origin servers, the method comprising: 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 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, 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;	40. A server operative in a computer network for delivering resources associated with a plurality of content providers to multiple client machines, 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: cache storage; a table listing content providers of said plurality of content providers having content that is authorized for delivery to client machines via the server, wherein the server is associated with *a first alias name* corresponding to said first content provider, and 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 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	43. A server operative in a computer network for delivering resources to multiple client machines, the server comprising: cache storage; a table listing a plurality of origin servers having content that is authorized for delivery to client machines via the server, wherein the server is associated with *a first alias name* corresponding to a first origin server of the plurality of origin servers, and 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 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 wherein, in response to a request for a resource, the server uses at least the table to analyze an

Claim 1	Claim 28	Claim 37	Claim 40	Claim 43
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; 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*; 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; 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 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.	resource from said at least one repeater server, wherein the second origin server is distinct from the first origin server; 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 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.	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 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 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.	second content provider, and 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.	alias name associated with the request to determine an origin server associated with the particular resource.