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

PATENT OWNER'S PRELIMINARY RESPONSE

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

PATENT OWNER'S EXHIBIT LIST

Exhibit No.	Description
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)
2005	Declaration of Dr. Prashant Shenoy
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

TABLE OF AUTHORITIES

Cases

Statutes

I. INTRODUCTION

The instant Petition is a copy-cat joinder Petition copying IPR2025-00969 challenging U.S. Patent No. 9,478,903 ("the '903 patent").

Although the Board instituted on the '969 IPR, this POPR presents new arguments not previous addressed by the Board that demonstrate the Board should deny institution on the merits here because both Petitions fail to show a likelihood that any of the challenged claims are unpatentable. In the end, all Grounds are premised on a flawed analogy between the International Publication No. WO1996041285 ("Kenner") video retrieval system and the '903 patent because Kenner discloses a disparate technical architecture and operation for at least two reasons.

The '903 patent claims a network of shared multi-tenant repeater servers that optimize content transmission to clients from multiple content providers using a novel "alias name" architecture and operation. Requests for content are "reflected" by a reflector to an appropriate repeater based on an "alias name." The claims require an "*alias name*" that is "*associated with*" a shared repeater server and enables resource requests to be "*directed*" to the repeater server. The claims also require a system that uses a plurality of alias names (e.g., first or second alias names) associated with the at least one shared repeater server.

In contrast, Kenner discloses a video retrieval system. A user terminal accesses a primary index manager (a search engine) that accesses information about available videos, retrieves videos from storage units (local or remote), and provides videos to the terminal for viewing and manipulating. First, Kenner does not teach or suggest an "alias name" that is associated with *at least one shared repeater server*. The Petition asserts that Kenner's video ID teaches the claimed *alias name* and that Kenner's primary index manager ("PIM") teaches the *at least one shared repeater server*. Pet., 19-23. But Kenner's video ID is not a host name that identifies the PIM. Because Kenner's video ID is the unique identifier for a video used consistently throughout Kenner's system, Kenner's video ID is not an alias name as understood in the context of the '903 patent.

Second, Kenner does not teach or suggest directing a resource request received from a client machine to the at least one shared repeater server *based at least in part on an alias name*. In Kenner, all client terminal requests are directed to a corresponding PIM. Kenner does not describe using the video ID or inspecting the video ID to identify the PIM for receiving the request. Because requests from a client machine in Kenner are not directed using the video ID, Kenner does not describe directing a request based at least in part on an alias name.

Institution should be denied.

II. BACKGROUND OF THE TECHNOLOGY

The ’903 patent arises from pioneering technology directed to delivering content across a network using a content delivery network (“CDN”). CDNs provide critical services that enable platform and service providers like Microsoft to provide CDN platforms and cloud services to millions of consumers simultaneously over the Internet. The technology underlying these CDNs was made possible through a series of inventions developed by a California start-up called Sandpiper Software Consulting, LLC, a predecessor of Patent Owner Sandpiper CDN, LLC. *See* EX2005, ¶29; EX2008, ¶¶23-25.

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. A typical computer server in the 1990s could only handle a limited number of simultaneous connections before becoming overloaded. Moreover, signals take time to traverse routers and move through 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. *See* EX2005, ¶¶25-27.

In the mid-1990s, Andrew Swart and David Farber were among the first individuals to develop services that allowed content providers to distribute their

content over the Internet while avoiding the common congestion and performance issues that plagued Internet transmission at that time. Sandpiper's solution was to deploy CDN servers around the world, replicate content from customers' origin servers to appropriate CDN servers, and transparently route end users requesting that content to the "best" CDN server, all while providing customers with control over their content and the user experience. This service and its architecture was quickly imitated by many others in the industry. *See id.*, ¶¶29-30; EX2008, ¶¶23-26.

Using solutions developed by Mr. Swart and Mr. Farber, end users would connect to an optimal edge server—one that that was closer to the user or that had available capacity. Distributing content across a network of servers alleviated data congestion issues and allowed consumers to connect to nearby edge servers to reduce latency. Mr. Swart and Mr. Farber developed and built systems and methods for propagating data from origin servers to edge servers (known as "caching") based on network demand and for seamlessly routing users to the optimal edge server with the correct content. Using caching to perform the replication process within the CDN framework flew in the face of conventional wisdom. *See* EX2005, ¶28; EX2008, ¶¶23-26.

In 1996, Mr. Swart and Mr. Farber founded Sandpiper Networks Inc. to further develop and commercialize their novel concept for a CDN. Sandpiper

Networks was based in Thousand Oaks, California. There, the inventors designed and built a CDN ultimately referred to as "Footprint." By at least May 24, 1996, the Sandpiper team had conceived techniques for delivering static and dynamic content, including streaming resources, such as audio and video, using Sandpiper's CDN. *See* EX2008, ¶¶23-26.

Sandpiper Networks labored not only to build and implement its CDN but also to protect their groundbreaking innovation through the patent system. Recognizing that its invention could revolutionize content delivery worldwide, Sandpiper Networks (and its successor companies, including Patent Owner) filed numerous patent applications directed to its foundational CDN technology, including the '903 patent.

By at least May 1998, Sandpiper Networks was caching content and delivering cached content to end users of content providers using its CDN. 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. 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. *See* EX2005, ¶30; EX2008, ¶¶23-26.

In the early-to-mid 2000s, demand for CDNs exploded. Sandpiper Networks continued to commercialize its technology and compete in this expanding market, doing business as various entities including Digital Island, Inc. and a company called "Level 3." It also created Sandpiper CDN, LLC (Patent Owner). Other businesses also entered the market by misappropriating Sandpiper's innovations and directly competing with Sandpiper for market share. *See* EX2005, ¶31; EX2008, ¶¶27-32.

One such company is Microsoft. Microsoft built a CDN, which uses technology described and claimed by Sandpiper's patents. For example, Microsoft's Azure CDN supports Xbox gaming, Office 365, Windows updates, and web resources like LinkedIn and Bing. *See* EX2001. Microsoft did not license this patented technology from Sandpiper, Level 3, and/or its predecessors. Accordingly, Sandpiper brought suit against Microsoft in district court for patent infringement. Copying Petitions previously filed by Google, Microsoft filed IPR Petitions against all of Sandpiper's asserted patents in response.

III. OVERVIEW OF THE '903 PATENT

The '903 patent is titled "Shared Content Delivery Infrastructure." EX1001, Title. 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 content providers want to have their own internet domain shown in end users'

browsers even when the content providers leverage a CDN to serve their content. *See* EX2005, ¶¶32-33.

Thus, the '903 patent solves issues in delivering resources from more than one content provider. This includes replicating content from a source associated with a client of the CDN network onto CDN servers—i.e., repeater servers. End user requests can be directed to the repeater servers instead of to the client's source servers (generally referred to as "origin" servers). This can reduce traffic on client origin servers. *See id.*, ¶34. These specific innovations enabled CDN customers to use their own domain name while using CDN services to serve content to end users. *See id.*, ¶¶35-36.

For instance, the inventors conceived of and claimed approaches to content delivery that involve a claimed **shared repeater server** and a claimed **alias name** that is a host name "*associated with*" a shared repeater server and enables requests to be "*directed*" to the repeater server (Section III.B). The claims also recite a first/second **alias name** associated with **a repeater server** and distinct content providers and a table that can be consulted by a repeater server to determine a content provider associated with a particular request. *See* Appendix A (comparing the independent claims).

A. Architecture and operation of "repeater server"

Figure 1 of the '903 patent illustrates a network environment including **reflector 108** at **origin server 102** and **repeaters 104a, 104b, and 104c**. EX1001, 4:22-26; *see also* EX2005, ¶¶37-39.

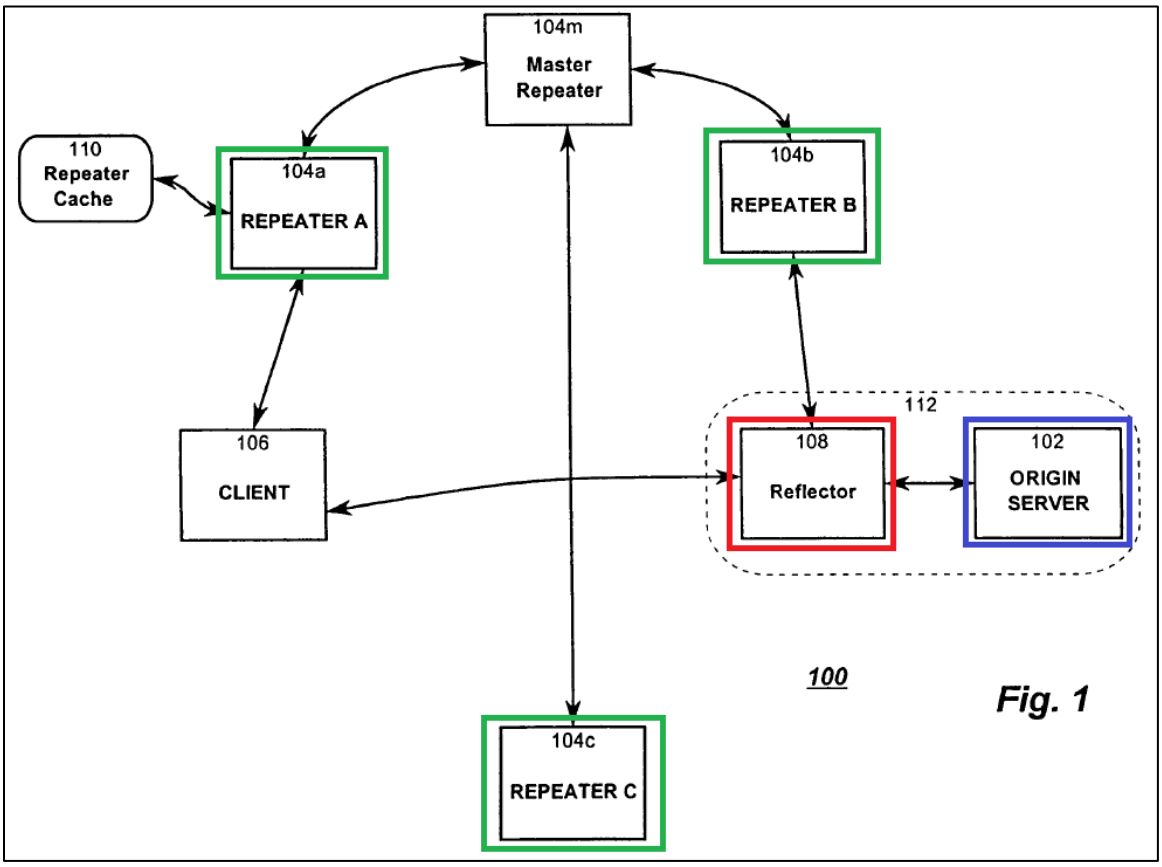

EX1001, FIG. 1 (annotated).

"Each **repeater 104a, 104b, and 104c** replicates some or all of the information available on the **origin server 102** as well as information available on other origin servers in the network." EX1001, 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.").

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.

Figure 3 of the '903 patent illustrates a reflector's processing of a request. In B1-B2, the reflector analyzes the request to determine whether the requestor is a browser (i.e., received from a client) or a repeater. *Id.*, 7:59-60; *see also* EX2005, ¶¶40-43. A request from a repeater is served locally by the origin server. EX1001, 7:62-63. This allows a repeater to build a cache of the content on the origin server because the repeater's requests to the origin server are always processed locally.

EX2005, ¶40.

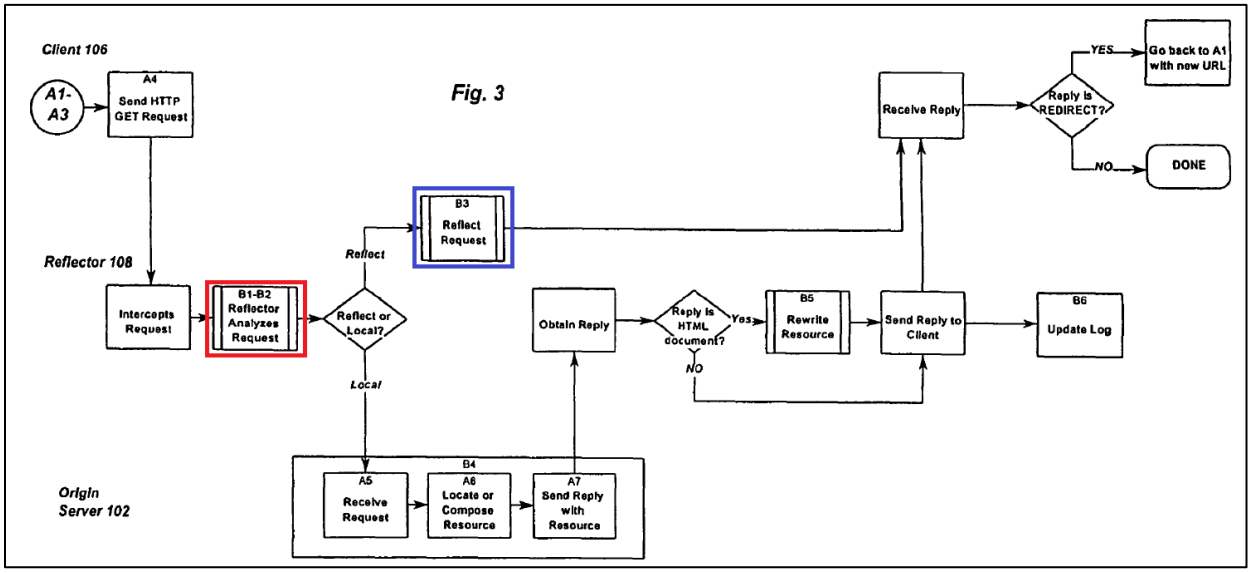

EX1001, FIG. 3 (annotated).

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. EX1001, 8:4-9; *see also id.*, 9:14-15 ("By using a rule base …, it is possible to selectively reflect resources.").

In **B3**, 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 resource" 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.

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"

To facilitate the operations just described, the '903 patent describes a novel use of the claimed "**alias name**" that is "*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; EX2005, ¶¶44-49. As further explained below, the '903 patent describes the "alias name" as a host name that identifies a repeater server. The '903 patent can employ a "family" of names for an origin server, with "each name identifying one of the repeaters used by that server." EX1001, 9:44-47. "For instance, if www.example.com is the origin server, names for three repeaters might be created: wr1.example.com[,] wr2.example.com[, and] wr3.example.com." *Id.*, 9:48-52. "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.'" *Id.*, 9:53-55 (emphasis added).

Notably, the terms "alias" and "alias name" appear solely in the passages from columns 9 and 10 of the '903 patent. *See id.*, 9:44-10:7. Specifically, the specification provides the meaning 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".***

Id., 9:44-63 (emphasis added).

In view of these alias name examples, a POSA would have recognized that an "alias name" would be a "host name." EX2005, ¶46. As the '903 patent explains, a "host name" identifies a particular computer such as a "server

computer." EX1001, 6:42-57. For example, the '903 patent provides "www.uspto.gov" as an example of a "host name." *Id.* As Petitioner's own expert acknowledged in a deposition concerning the '903 patent, 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").

In this manner, the specification of the '903 patent supports the understanding that a POSA would have understood "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. EX2005, ¶47.

Moreover, the plain language of the claims also supports this understanding as well. *Id.*, ¶48. 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* EX1001, 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. EX2005, ¶48.

Thus, in view of 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. *Id.*, ¶49.

C. Architecture and operation of a reflected request

In Figure 5, the '903 patent illustrates how a repeater intelligently processes a reflected request (e.g., based on an alias name). EX1001, 10:16-64; *see also* EX2005, ¶¶50-54. In **C1**, the repeater "analyzes the request to determine the network address of the requesting client and the path of the resource requested." EX1001, 10:20-22, 10:22-23 ("Included in the path is an origin server name.").

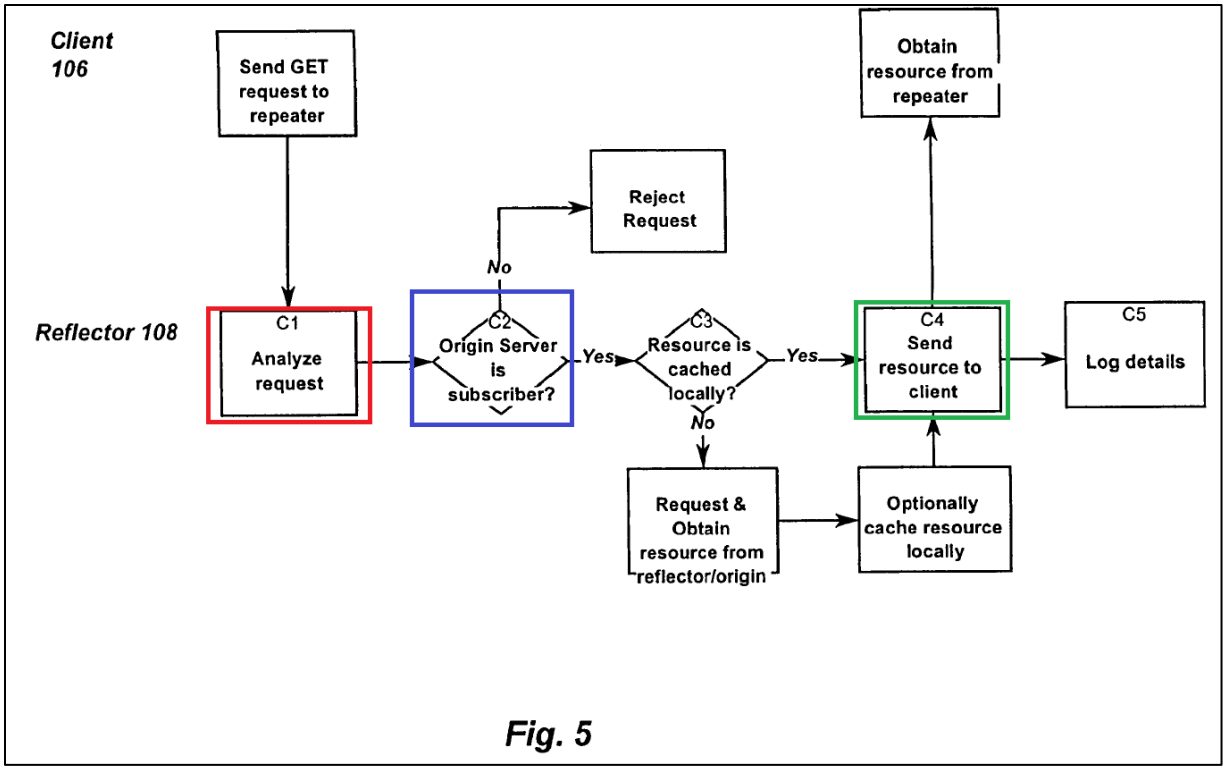

EX1001, FIG. 5 (annotated).

To determine the path of the resource, the repeater must know the name by which the client addressed the repeater. EX1001, 9:56-59. A repeater determines the alias used by inspecting the HOST tag in the request's HTTP header. *Id.* 10:4-8.

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).

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

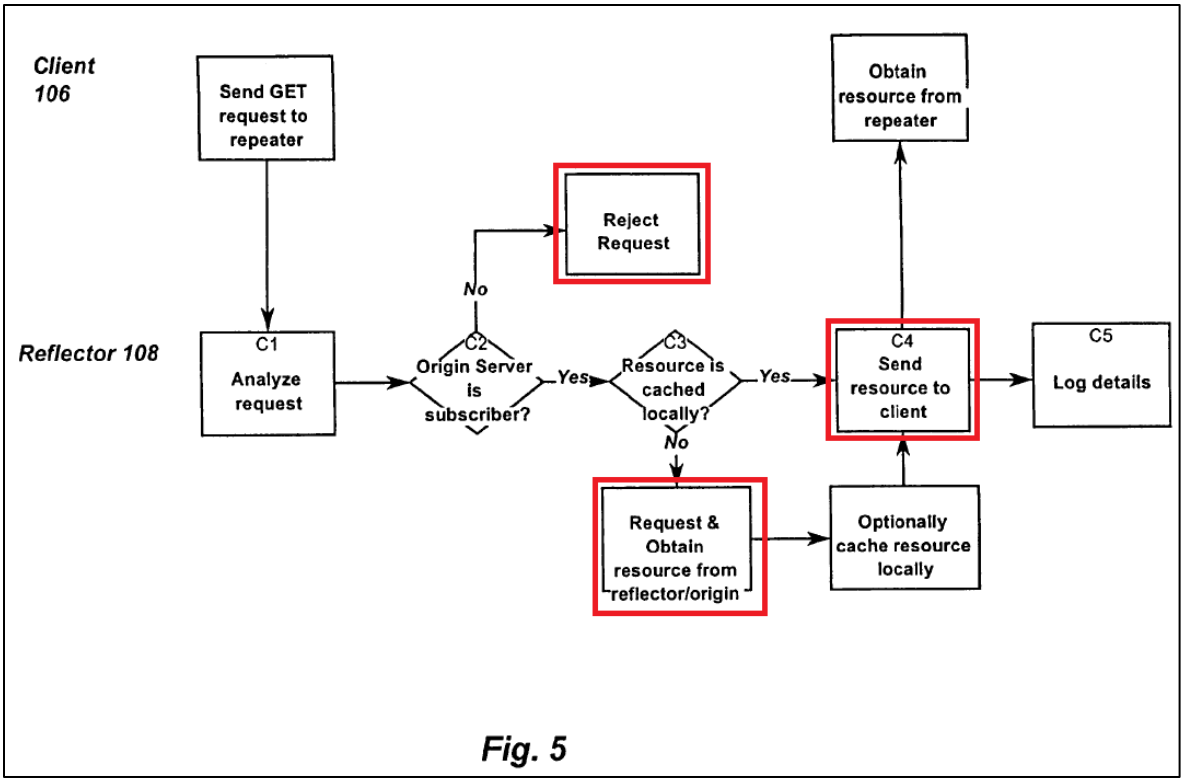

EX1001, FIG. 5 (annotated).

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." EX1001, 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 load on the origin server. *Id.*, 2:62-65.

IV. PROSECUTION HISTORY

The '903 patent was examined as U.S. Application No. 11/065,412 ("'412 Application"), which was filed on February 23, 2005. Prosecution focused on the core claim features of repeater servers and alias names. EX2005, ¶¶55-66.

To begin prosecution, Applicant filed four Preliminary Amendments that, in total, cancelled claims 1-26 and introduced new claims 27-51 ("prosecution claims"). EX1002, 1024, 954, 876-92, 788, 781, 717, and 697-96.

In a second Non-Final Office Action mailed on January 9, 2009, Examiner Blair rejected prosecution claims 27-38, 43-45, 47, and 49-51 under 35 U.S.C. § 102 as being anticipated by U.S. Patent No. 5,867,706 ("Martin") and claims 39-42, 46, and 48 under 35 U.S.C. § 103 as unpatentable over Martin in view of U.S. Patent Number 6,052,718 ("Gifford"). *Id.*, 536-45. Examiner Blair also rejected prosecution claims 27-44 under 35 U.S.C. § 101 and prosecution claims 27-51 under 35 U.S.C. § 112. *Id.*, 538-39.

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)." EX1002, 520. Applicant also advanced new claims 53 and 54:

> 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.

In a Final Office Action mailed on May 7, 2010, Examiner Blair rejected prosecution claims 29-38, 43-45, 47, and 55-57 under 35 U.S.C. § 103 as being anticipated by U.S. Patent No. 6,085,193 ("Malkin") in view of Martin and claims 39-42, 46, and 48 under 35 U.S.C. § 103 as unpatentable over Malkin in view of Martin in view of Gifford. EX1002, 241-50. Examiner Blair noted that claims 53 and 54 (reproduced above) contained Allowable Subject Matter. *Id.*, 249.

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, 210-211.

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 (U.S. Patent No. 6,243,760 to 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." EX1002, 94.

The '903 patent issued on July 2, 2013. *Id.*, 2.

V. OVERVIEW OF KENNER

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; EX2005, ¶¶67-80. 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 …." EX1005, 3:15-17, 4:5-7.

Figure 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 a primary index manager **PIM 22**.

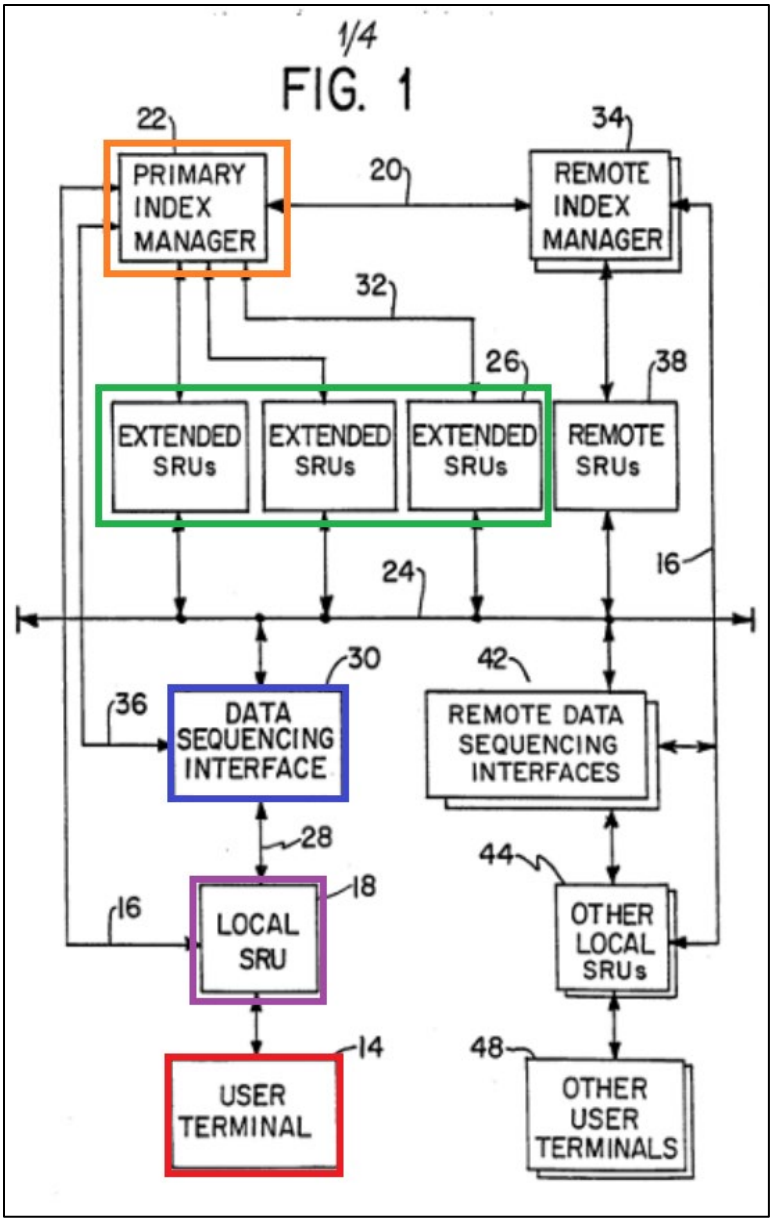

EX1005, FIG. 1 (annotated).

User terminal 14 allows a "user to manipulate retrieved video clips"—e.g., play, stop, pause, fast forward, etc. EX1005, 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.

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.

PIM 22 is the "primary search engine and database management module of the invention." *Id.*, 16:2-3. **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.

Extended storage retrieval units ("SRUs") 26 provide the "principle storage facility for the system and is 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.

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.").

Figure 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.

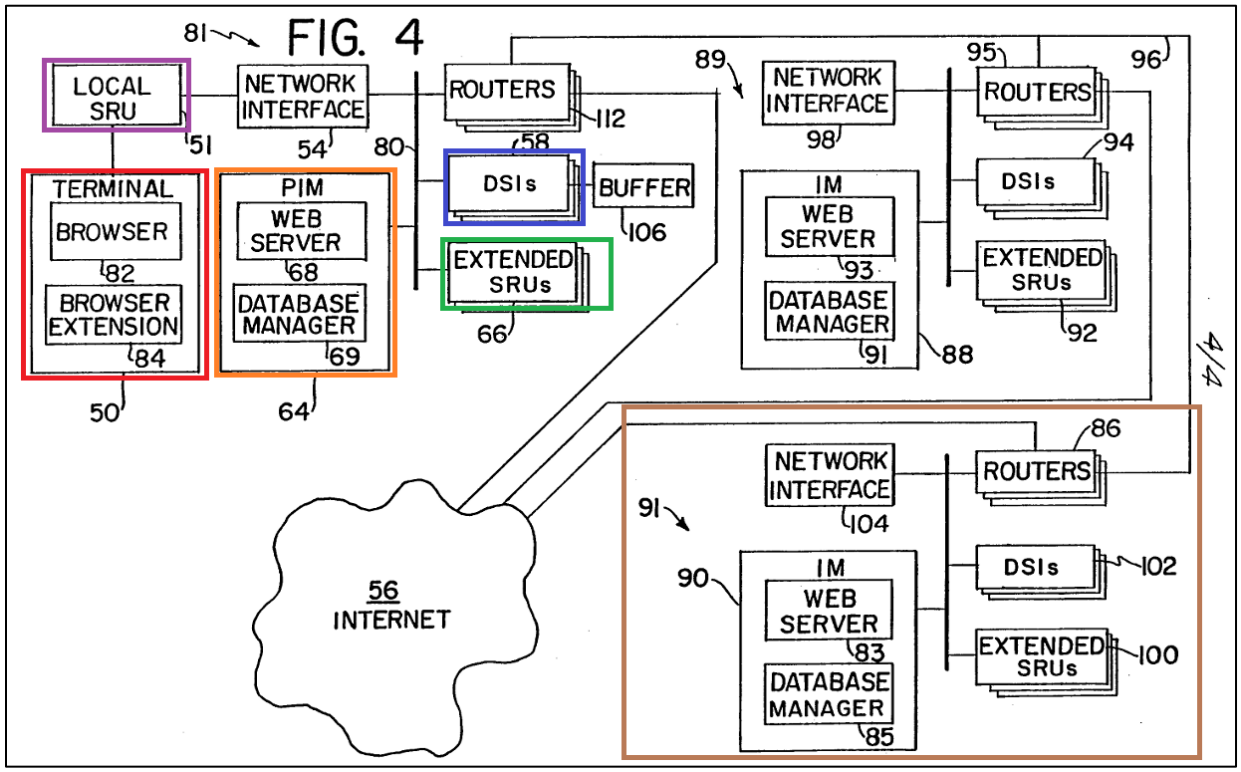

EX1005, FIG. 4 (annotated).

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 a router 86." EX1005, 32:14-16.

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).

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). "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 databases." *Id.*, 37:15-17.

The PIM "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 following 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.

EX1005, 34:5-35:3.

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." EX1005, 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.

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.*, 40:29-33. 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.

Figure 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.

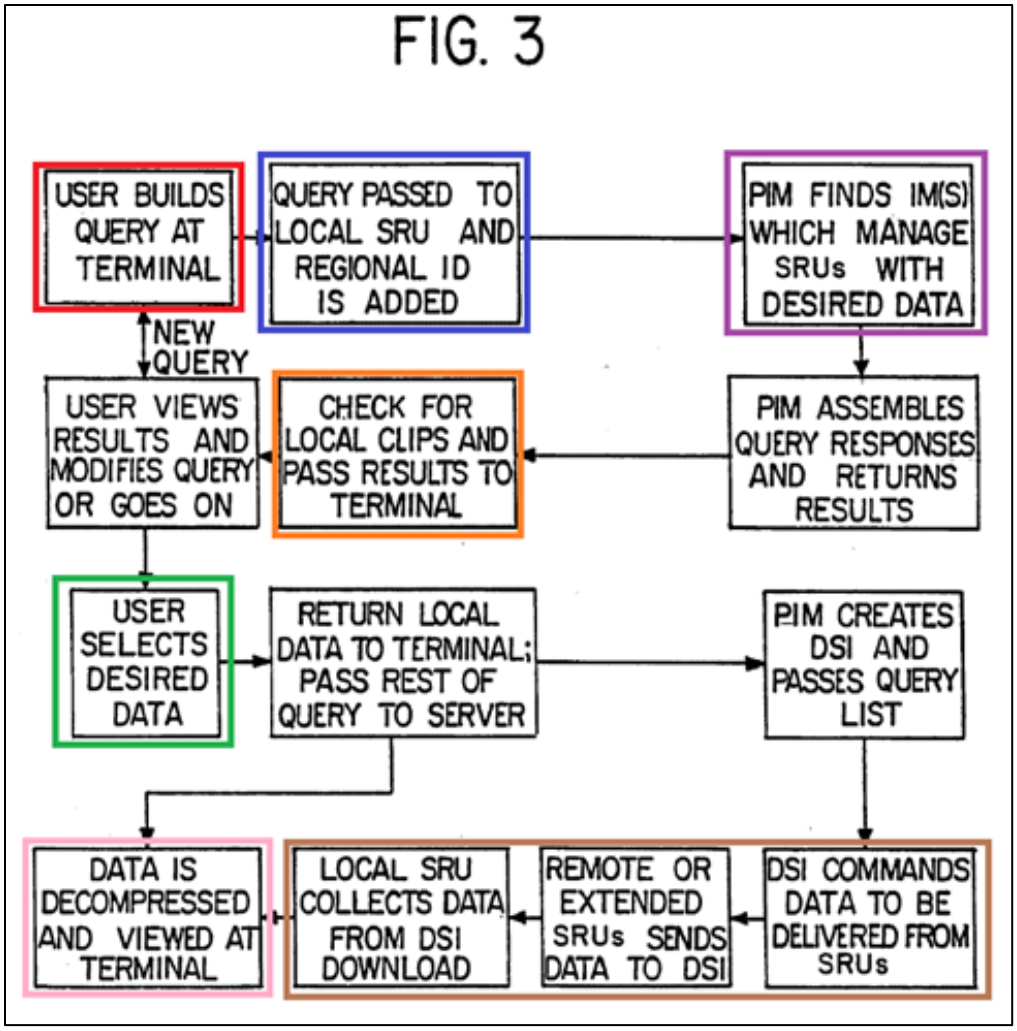

EX1005, FIG. 3 (annotated).

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 …." EX1005, 25:11-14; *see also id.* 7:3-4 ("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.").

VI. PERSON OF ORDINARY SKILL IN THE ART

At this stage, Patent Owner submits that the Board does not need to resolve the appropriate level of skill for a POSA for purposes of deciding institution because the Petition fails under any articulation of a POSA. EX2005, ¶82.

VII. ARGUMENT AGAINST INSTITUTION

Kenner fails to teach all elements of independent claims 1, 28, 37, 40, and 43 because Kenner does not teach at least two critical claim features: (1) "*alias name*" and (2) directing resource requests to "*the at least one shared repeater server, based at least in part on*" an alias name.

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.

Kenner's video ID (i.e., Petitioner's alleged *alias name*) is not associated with any repeater server and also is not a host name. Pet., 17-18. Accordingly, Kenner's video ID does not teach the claimed "alias name." EX2005, ¶92.

a. Kenner's video ID lacks association with a repeater server.

As explained in Section III.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. EX2005, ¶¶93-94.

With reference to Claim 1, the Petition asserts 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. But Kenner refers to the video ID as a "unique reference identifier for each video clip" that "corresponds to an identifying field within [a] text database." *Id.*, 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.

As seen from this description, nothing in Kenner's video ID associates the video ID with a repeater server. EX2005, ¶95. Moreover, nothing in the video ID identifies the specific PIM that Petitioner relies on to allegedly teach the claimed *at least one shared repeater server*. *Id.* In fact, Kenner explicitly contemplates other components performing this work: the "virtual URL" contains the "Internet address of the PIM 64" that a request is directed to. EX1005, 37:7-17; *see also infra* Section VII.A.2. In this manner, Kenner's video ID is not associated with a repeater server and therefore is not an "alias name" as claimed. Section III.B.

Petitioner argues that when an IM stores a particular clip, that IM becomes associated with the video ID. Pet., 20-21. But this theory is also flawed. 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. EX2005, ¶96. Indeed, Petitioner's expert even acknowledges that the 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. Thus, 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"

Additionally, Kenner's video ID is not a "host name" as the '903 patent describes as a characteristic of an alias name. EX2005, ¶97. Petitioner's expert confirmed this understanding stating that "I don't think one would normally characterize [a video ID] as being a host name." EX2007, 129:16-22. As 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. EX2005, ¶97. 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. EX2005, ¶97. ***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. In this manner, Kenner's video ID is not a host name and also does not use the same structure as a host name. EX2005, ¶97.

***

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." EX2005, ¶98.

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

Kenner's video ID does not disclose the claimed *alias name* in view of Kenner's virtual URL. EX2005, ¶¶99-100. 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). EX2005, ¶¶99-100.

At the outset, 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. *Id.*, 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. *Id.*, 36:21-27, 37:15-17; EX2007, 119:13-17 (Petitioner's expert 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.

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. EX2005, ¶101; *see also infra* Section VII.B. Kenner's video ID does not perform this functionality. In view of 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; EX2005, ¶101. 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; EX2005, ¶101. 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 (All Grounds; Independent Claims 1, 28, and 37 and corresponding dependent claims).

Kenner does not describe directing resource requests to "**the** at least one shared repeater server, ***based at least in part on***" an alias name. EX2005, ¶102. 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. *Id.*

At the outset, the claims require directing requests from client machines to 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. 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.

In this manner, the claims recite a requirement that the alias name impacts where to initially direct a resource request from a client machine. EX2005, ¶¶103-105. 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. *Id.* 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 antecedent basis reflected in the claims, 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.

First, Kenner does not teach directing requests to the PIM (i.e., alleged *at least one shared repeater server*) based on an alias name because Kenner's video ID is not an alias name. Section VII.A.1. Again, as explained in Section III.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 VII.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. EX2005, ¶106.

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 (Petitioner's expert acknowledging that all requests from terminal 50 are sent to PIM 64). In view of the antecedent basis in the claims, to actually meet the claim language would require an identification of the PIM 64 specifically from among other IMs and a directing of the requests specifically to PIM 64 because of the video ID. But 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. EX2005, ¶107.

With reference to Figure 4, Kenner describes *all* requests from terminal 50 being directed to PIM 64:

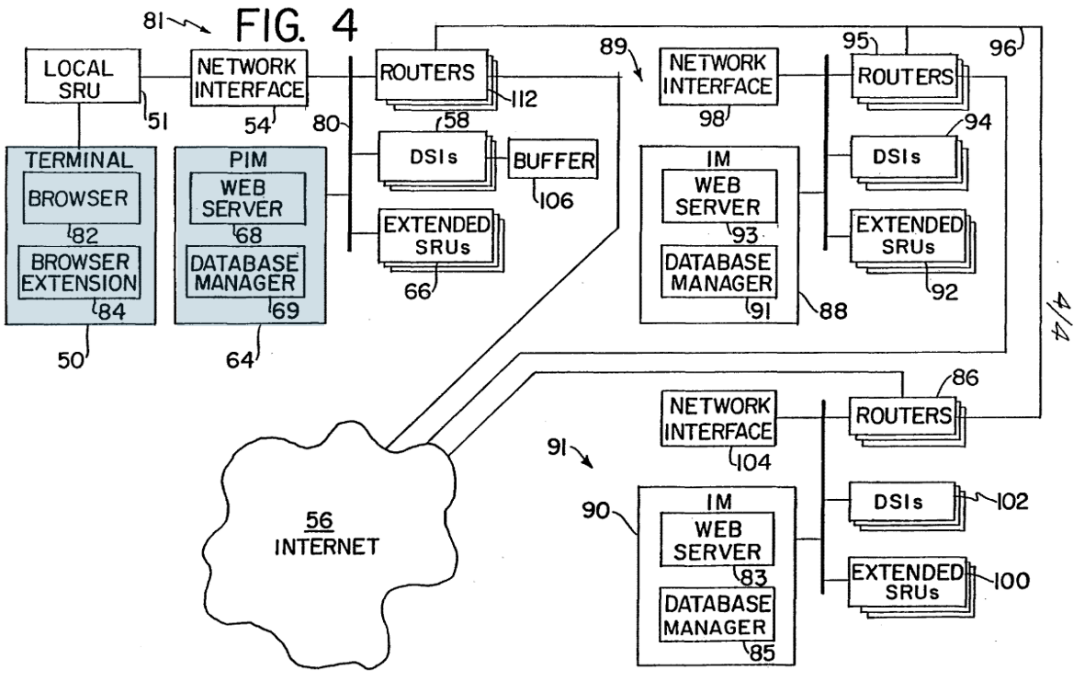

EX1005, FIG. 4 (annotated).

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.

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

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." *Id.*, 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." *Id.*, 37:7-17.

In this manner, Kenner describes directing all requests from terminal 50 to PIM 64 regardless of the content of the video ID. EX2005, ¶¶108-111. 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. Moreover, Kenner does not describe *using* the video ID to determine that PIM 64 should receive the request from terminal 50 either. *Id.* Indeed, the video ID does not impact the decision of which server should receive the request at all. *Id.* Rather, PIM 64—Petitioner's alleged *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. EX2005, ¶111.

Seeming to recognize this deficiency, 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. *Id.* (citing EX1005, 40:29-33). Petitioner alleges that the IM will use content coordinate data included in the video ID to identify neighboring IMs, like the primary IM. Pet., 21.

But Petitioner's alternative theory is also fundamentally flawed in view of 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. EX2005, ¶¶112-113. 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. *Id.*

Kenner's description of communication occurring between neighboring IMs does not address the specific functionality recited in the claims. *Id.*, ¶114. 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. *Id.*, 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***. EX2005, ¶114.

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. *Id.*, ¶115.

C. The Board should deny Microsoft's Petition on the merits because Petitioner has failed to provide any indication that it can produce Dr. Mowry for deposition, thereby rendering his declaration unusable.

As explained in Patent Owner's Opposition to Petitioner's Motion for Joinder (Paper 6), Microsoft has provided no evidence of actually speaking to or retaining Dr. Mowry. Without taking at least these steps, there is no guarantee that Microsoft would be able to produce Dr. Mowry for cross-examination. While

Microsoft has provided an email (EX1019) indicating that Google does not oppose Microsoft retaining Dr. Mowry and that Dr. Mowry has no conflicts, Microsoft still provides no indication that it has actually contacted Dr. Mowry or received consent from Dr. Mowry to submit his declaration in this IPR proceeding.

Indeed, during a deposition of Dr. Mowry related to the '903 patent in the Google proceeding (IPR2025-00969), Dr. Mowry had no recollection of speaking to any of Microsoft's representatives at all:

> Q: It's also my understanding that you were approached by Microsoft as well; is that accurate?
>
> A: I don't recall, to be honest.
>
> Q: Okay. Are you aware that Microsoft filed a copy of your declaration in another IPR proceeding?
>
> A: I don't remember that off the top of my head.

EX2007, 10:5-13.

> Q: Okay. Do you recall talking to any attorneys from any other law firms about the '903 patent?
>
> A: Not -- not specifically, not off the top of my head.

EX2007, 11:4-8.

These admissions demonstrate that Microsoft has failed to diligently retain Dr. Mowry—an error that the Board has admonished and applied as reason for denying institution. *See OpenSky Industries, LLC v. VLSI Tech. LLC*, IPR2021-01056, Paper 18, 7-8 (P.T.A.B. Dec. 23, 2021) (the Board denying institution and

noting that "we have no evidence from Dr. Singh that he would be willing to testify here" and "[t]here is no indication that Petitioner ever spoke to Dr. Singh or attempted to retain him for this proceeding or secure his availability for cross examination before filing his declaration").

In view of these admissions and Microsoft's improper reliance on an expert declaration where it has not retained said expert for cross examination purposes, Microsoft cannot rely on Dr. Mowry's declaration to support its arguments. *Id.*, 9. Accordingly, the Board should also deny Microsoft's petition on the merits because the theories presented are not properly supported by expert testimony. *Id.*, 9-10.

VIII. CONCLUSION

The Board should deny *inter partes* review of the copy-cat joinder Petition against the '903 patent.

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg, Reg. No. 43,447
Daniel S. Block, Reg. No. 68,395
Kristina Caggiano Kelly, Reg. No. 58,164
Timothy L. Tang, Reg. No. 75,187
Todd Thurheimer, Reg. No. 76,231
Counsel for Patent Owner

Date: February 13, 2026

1101 K Street, NW
10th Floor
Washington, DC 20005
(202) 371-2600

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.

CERTIFICATE OF WORD COUNT (37 C.F.R. § 42.24(d))

1. This Patent Owner Preliminary Response complies with the type-volume limitation of 14,000 words, comprising 7,970 words, excluding the parts exempted by 37 C.F.R. § 42.24(a)(1).

2. This Patent Owner Preliminary Response complies with the general format requirements of 37 C.F.R. § 42.6(a) and has been prepared using Microsoft® Word 2016 in 14-point Times New Roman font.

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg
Registration No. 43,447
Counsel for Patent Owner

Date: February 13, 2026

1101 K Street, NW, 10th Floor
Washington, DC 20005
(202) 371-2600

CERTIFICATE OF SERVICE (37 C.F.R. § 42.6(e))

I certify that the above-captioned **PATENT OWNER PRELIMINARY RESPONSE** and associated Exhibits 2005-2008 were served in their entireties on February 13, 2026, upon the following parties via electronic mail:

Jessica Kaiser (Lead Counsel)	kaiser-ptab@perkinscoie.com
Christopher Marando (Back-Up Counsel)	marando-ptab@perkinscoie.com
Matthew Lembo (Back-Up Counsel)	lembo-ptab@perkinscoie.com
Steven Beigelmacher (Back-Up Counsel)	beigelmacher-ptab@perkinscoie.com
Sarah Fowler (Back-Up Counsel)	stahnke-ptab@perkinscoie.com

Microsoft-Sandpiper-IPR@perkinscoie.com
PERKINS COIE LLP

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg
Registration No. 43,447
Counsel for Patent Owner

Date: February 13, 2026

1101 K Street, NW, 10th Floor
Washington, DC 20005
(202) 371-2600