

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC,
Petitioner

v.

SANDPIPER CDN, LLC,
Patent Owner

Case IPR2025-00969
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

I.	INTRODUCTION	1
II.	BACKGROUND OF THE TECHNOLOGY	3
III.	OVERVIEW OF THE '903 PATENT	6
IV.	PROSECUTION HISTORY	14
V.	OVERVIEW OF THE CITED PRIOR ART	18
VI.	PERSON OF ORDINARY SKILL IN THE ART	26
VII.	CLAIM CONSTRUCTION	26
VIII.	LEGAL STANDARDS	26
IX.	ARGUMENT AGAINST INSTITUTION	27
A.	Kenner does not teach or suggest <i>a first alias name and a second alias name</i> that are associated with <i>at least one shared repeater server</i> (all independent claims).	31
B.	Kenner does not teach or suggest <i>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</i> (all independent claims).	36
X.	CONCLUSION	39

PATENT OWNER'S EXHIBIT LIST

Exhibit No.	Description
2001	“Cloud CDN: content delivery network,” <i>Google Cloud</i> , accessed at https://cloud.google.com/cdn (last accessed July 2, 2025)
2002	<i>Intentionally Left Blank</i>
2003	<i>Sandpiper CDN, LLC v. Google LLC</i> , No. 2:24-cv-03951, ECF No. 1 (C.D. Cal. May 10, 2024)
2004	<i>Sandpiper CDN, LLC v. Google LLC</i> , No. 2:24-cv-03951, Invalidity Contentions (C.D. Cal. April 22, 2025)
2005	<i>Sandpiper CDN, LLC v. Google LLC</i> , No. 2:24-cv-03951, ECF No. 75 (C.D. Cal. May 23, 2025)
2006	<i>Sandpiper CDN, LLC v. Google LLC</i> , No. 2:24-cv-03951, ECF No. 82 (C.D. Cal. June 24, 2025)
2007	<i>Sandpiper CDN, LLC v. Google LLC</i> , No. 2:24-cv-03951, ECF No. 28 (C.D. Cal. Sept. 16, 2024)
2008	FAQs for Interim Processes for PTAB Workload Management, accessed at https://www.uspto.gov/patents/ptab/faqs/interim-processes-workload-management
2009	<i>Intentionally Left Blank</i>
2010	<i>Intentionally Left Blank</i>
2011	<i>Intentionally Left Blank</i>
2012	<i>Intentionally Left Blank</i>
2013	Interim Processes for PTAB Workload Management Memorandum (March 26, 2025)
2014	Declaration of Dr. Prashant Shenoy
2015	<i>Curriculum Vitae</i> of Dr. Prashant Shenoy

TABLE OF AUTHORITIES

	Page(s)
Cases	
<i>In-Depth Geophysical, Inc. v. ConocoPhillips Co.</i> , IPR2019-00850, Paper 56 (Sept. 3, 2020)	26
<i>In re Magnum Oil Tools Int’l, Ltd.</i> , 829 F.3d 1364 (Fed. Cir. 2016)	27
<i>Sirona Dental Sys. GmbH v. Institut Straumann AG</i> , 892 F.3d 1349 (Fed. Cir. 2018)	27
Statutes	
35 U.S.C. § 101	14
35 U.S.C. § 102	14
35 U.S.C. §103	14, 16
35 U.S.C. § 112	14
Other Authorities	
37 C.F.R. § 42.6(a)	1
37 C.F.R. § 42.6(e)	2
37 C.F.R. § 42.24(a)(1)	1
37 C.F.R. § 42.24(d)	1
37 C.F.R. § 42.104(b)(4)	30

I. INTRODUCTION

The Board should deny institution on the merits because the Petition fails to show a likelihood that any of the challenged claims is unpatentable. The Petition presents only a single, deeply flawed ground challenging the independent claims of U.S. Patent No. 9,478,903 (“the ’903 patent”). This single-reference ground of obviousness relies on International Publication No. WO1996041285 (“Kenner”). But the ’903 patent and Kenner disclose very different systems.

The ’903 patent claims a network of shared repeater servers that optimize content transmission to clients from multiple content providers. Requests for content are “reflected” by a reflector to an appropriate repeater, and the repeater provides cached content to clients via redirect.

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.

With respect to the claims of the ’903 patent, Kenner does not teach or suggest at least two critical claim elements.

First, Kenner does not teach or suggest *a first alias name and a second alias name* that are associated with *at least one shared repeater server*. The Petition asserts that Kenner’s video ID teaches the claimed *first/second alias name* and that

Kenner's primary index manager ("PIM") teaches the *at least one shared repeater server*. Pet., 20-24. But Kenner's video ID is not associated with the PIM—Kenner's video ID does not identify the PIM or anything analogous to a repeater server. Indeed, Kenner's video ID is not an alias at all—it is the unique identifier for a video used consistently throughout Kenner's system.

Second, Kenner does not teach or suggest *the at least one shared repeater server is further constructed and adapted to analyze, using the table, an alias name received with a client request for a particular resource to determine a content provider associated with the particular resource*. Kenner does not teach or suggest any analysis of the video ID with respect to a *table*. In Kenner, a PIM need not analyze the video ID received in any request *using the table* because the provider's account number is in the video ID. Any such analysis of the video ID would be redundant. In contrast, in the claimed invention, a repeater analyzes the *alias name* received in the header of a request with reference to the *table* to determine an authorized content provider.

Importantly, these two features were recognized by the Examiner as novel during prosecution. *See* Section IV.

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 content providers like Google to quickly deliver online content 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* EX2014, ¶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* EX2014, ¶26.

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* EX2014, ¶27.

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* EX2014, ¶28.

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* EX2014, ¶29.

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* EX2014, ¶30.

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* EX2014, ¶31.

One such company is Google. Google built a CDN, which uses technology described and claimed by Sandpiper’s patents. Additionally, YouTube, which Google purchased in 2006, relies upon CDN technology to meet its enormous data streaming needs and to provide Google’s video search results (powered by YouTube). *See* EX2001. Google did not license this patented technology from Sandpiper, Level 3, and/or its predecessors. Accordingly, Sandpiper brought suit against Google in district court for patent infringement. Google 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* EX2014, ¶¶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* EX2014, ¶34.

For instance, the inventors conceived of approaches to content delivery which involve a shared repeater server and the use of alias names. As one example, the '903 patent claims 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). These specific innovations enabled CDN customers to use their own domain name while using CDN services to serve content to end users. *See* EX2014, ¶¶35-36.

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* EX2014, ¶¶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. EX1001, 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. EX1001, 7:59-60; *see also* EX2014, ¶¶40-43. 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 the repeater's requests to the origin server are always processed locally. EX2014, ¶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.” EX1001, 8:22-28. 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. Alternatively, 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.” *Id.*, 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 `repeater1`, which might also be known by other names such as ‘`wr1.anotherExample.com`’ and ‘`wr1.example.edu`.’” *Id.*, 9:53-55.

The reflector sends a “REDIRECT reply containing this new URL to the requesting client.” EX1001, 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.

In Figure 5, the '903 patent illustrates how a repeater intelligently processes a reflected request. EX1001, 10:16-64; *see also* EX2014, ¶¶44-48. In **C1**, the repeater “analyzes the request to determine the network address of the requesting client and the path of the resource requested.” *Id.*, 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.

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.

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).” *Id.*, 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. *Id.*, 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.” *Id.*, 94.


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

V. OVERVIEW OF THE CITED PRIOR ART

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; EX2014, ¶¶61-74. 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 **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 exact 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.

VII. CLAIM CONSTRUCTION

For purposes of this Preliminary Response, Patent Owner agrees the *Phillips* standard applies. Patent Owner reserves the right to contest Petitioner’s implicit or explicit constructions during the proceeding should the Board decide to institute review.

VIII. LEGAL STANDARDS

It is Petitioner’s burden to demonstrate unpatentability by a preponderance of the evidence. Petitioner “must ‘specify where each element of the claim is found in the [relied upon] prior art patents.’” *In-Depth Geophysical, Inc. v.*

ConocoPhillips Co., IPR2019-00850, Paper 56, *27 (Sept. 3, 2020) (quoting 37 C.F.R. § 42.104(b)(4)). The Board makes its determination regarding institution based on what the Petition actually presents and not what it could have reasonably contained had it been drafted differently. *In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1381 (Fed. Cir. 2016) (holding that the Board is not “free to adopt arguments on behalf of petitioners that could have been, but were not, raised by the petitioner during [a post-grant proceeding]”). The Board cannot “deviate from the grounds in the petition and raise its own” theories of invalidity. *Sirona Dental Sys. GmbH v. Institut Straumann AG*, 892 F.3d 1349, 1356 (Fed. Cir. 2018).

IX. ARGUMENT AGAINST INSTITUTION

The Petition is deeply flawed. Even if the Board were inclined to decipher the Petition on the Petitioner’s behalf, the deficiencies of Kenner—the lone reference relied upon to invalidate the independent claims—are unavoidable.


The ’903 patent claims are directed to a CDN that serves multiple content providers. The claimed innovations allow a repeater server to receive a reflected request, identify authorized origin servers, and serve content from authorized content providers to clients. *See* EX2014, ¶86.

Kenner does not disclose such a CDN. Instead, Kenner discloses delivering content to many simultaneous end users using distributed index managers. Pet., 8-12 *citing* EX1005, 4:32-5:5 (describing general advantages of parallel systems),

5:21-26 (describing capability of supporting thousands of simultaneous users), 28:3-6 (describing remote storage units), 33:4-11 (describing subscription service for users), 35:11-16 (describing user accessing videos from a subscription service), 40:24-28 (describing multiple index managers to server a large number of concurrent users), 41:13-19 (noting that a PIM can contact a source index manager), 54:12-16 (describing a large number of geographically dispersed users).

In this regard, the “subscribers” in Kenner and the ’903 patent are not the same. *See, e.g.,* Pet., 19. In the ’903 patent, a “subscriber” is “an entity (e.g., a company) that publishes resources (e.g., files) via one or more origin servers.” EX1001, 10:24-31 (“When the entity subscribes, it is permitted to utilize the repeater network.”). In Kenner, a “subscriber” is the end user. *See* EX1005, 52:7-12. The Petition erroneously conflates the two. *See* Pet., 19; *see also* EX2014, ¶¶87-88.

Unable to find a disclosure of serving multiple content providers in Kenner, the Petition fabricates disclosure using annotation to modify Figure 4 of Kenner to include “an additional source IM” where none exists in the actual disclosure.



Pet., 11 (annotating a modified FIG. 4).

Similarly, the Petition falsely equates Kenner's PIM to the claimed *at least one shared repeater server*. See Pet., 14-15; EX2014, ¶89. Kenner's PIM is more akin to a web server with access to local and remote storage than a repeater server—e.g., Kenner does not describe a proxy, a cache, etc. See EX1001, 4:37-40 (“[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.”).

The Petition tries to bridge these significant gaps between the '903 patent and the prior art by stretching and reworking the disclosure of Kenner as far as it can, and then purporting to fill the remaining gaps in its own retelling of Kenner with gap filling through its expert declaration in violation of § 42.104(b)(4). Yet somewhat paradoxically, the testimony of Petitioner's expert, Dr. Mowry, offers little substantive detail or insights to actually address the deficiencies of Kenner.

Specifically, Kenner fails to teach or suggest: (1) a *first alias name* and a *second alias name* that are associated with *at least one shared repeater server* and/or (2) *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*. For ease of reference, the discussion below refers to the language of claim 1, but each of the independent claims recites substantially similar features. *See* Appendix A (comparing the independent claims); EX2014, ¶90.

Notably, these are the same features that the Examiner recognized as novel during prosecution. *See* Section IV. Each failing is discussed in more detail below.

A. Kenner does not teach or suggest *a first alias name and a second alias name* that are associated with *at least one shared repeater server* (all independent claims).

Each of the independent claims recites a *first alias name* and a *second alias name*. See Appendix A (comparing the independent claims); EX2014, ¶91. In claim 1, the *first alias name* and the *second alias name* are claimed to be associated with both: (1) *at least one shared repeater server*; and (2) a *first/second resource associated with the first/second content provider*. See EX2014, ¶96. Each of the remaining independent claims (28, 37, 40, and 43) recites substantially similar features. See Appendix A (comparing the independent claims); EX2014, ¶¶97-98.

In the '903 patent, the claimed *alias name* is an alternative address for a repeater server that also contains the URL of a resource. EX1001, 8:29-33, 9:53-55. For example, an *alias name* identifies both a repeater (“wr1”) and an origin server (“example.com”, “anotherExample.com”, or “example.edu”). See EX2014, ¶92.

Repeater Name	Example Alias Names
repeater1	“wr1.example.com” “wr1.anotherExample.com” “wr1.example.edu”

EX1001, 9:44-55.

The claimed *first alias name* and the *second alias name* are associated with **distinct** content providers. But both the *first alias name* and the *second alias name*

are claimed to be associated with **the same** *at least one shared repeater server*. In other words, the same *at least one shared repeater server* is addressable using either/both 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.”). This is reflected in the examples provided by the ’903 patent because “repeater1” can be addressed using “wr1.example.com,” “wr1.anotherExample.com,” or “wr1.example.edu.” Each of the example alias names is also associated with a distinct content provider—i.e., “example.com,” “anotherExample.com,” and “example.edu.” EX2014, ¶99.

The Petition asserts that Kenner’s video ID teaches the claimed *first/second alias name*. Pet., 20-24. Kenner’s video ID is a “unique reference identifier for each video clip” that “corresponds to an identifying field within the text database.” EX1005, 22:1-2, 43:27-29 (“The software tool will also communicate with a master database of video ID numbers to ensure that each clip uploaded, no matter from where, will have a unique video ID.”).

But Kenner’s video ID is not associated with both a *content provider* and *at least one shared repeater server*, as required by the claims, because Kenner’s video ID does not identify a repeater server or anything analogous to a repeater server. See EX2014, ¶¶92-106. The Petition asserts that Kenner’s PIM is the *at*

least one shared repeater server. See Pet., 15-16. But Kenner's video ID does not identify a PIM or otherwise associate the video ID with a PIM. See EX2014, ¶100.

Kenner's 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. But nothing in the video ID indicates the PIM (*at least one shared repeater server*). *See EX2014, ¶101.*

The Petition asserts that Kenner's clip database provides the association between the *first/second alias name* and the *at least one shared repeater server*. *Pet.*, 20-21. Kenner's clip database includes the video ID and the retrieval units that store the referenced file.

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.

EX1005, 34:7-11.

Petitioner essentially asserts that by storing a table that includes video IDs on a PIM, Kenner teaches an association between a video ID and the PIM. But this interpretation is radically divorced from the claimed context. *See* EX2014, ¶102.

Namely, the claims require requests for a resource to be directed to the repeater *based at least in part on said [first/second] alias name*. In Kenner, a user terminal always accesses the same PIM (which the Petition asserts is the claimed *repeater server*). Nothing is *directed to* the PIM in any manner that is *based on* the video ID. Again, Kenner’s video ID identifies with particularity the video file. *See* EX2014, ¶103. The video ID does not identify a PIM and the user terminal directs all requests to the same PIM. *See, e.g.,* EX1005, 37:7-38:16.

This is to be contrasted with the ’903 patent, where the claims require that the *alias name* be associated with both the *resource* and the *at least one shared repeater server*. A reflector then directs the request to a repeater in the content delivery network using the *alias name*, and the repeater can determine the content provider and verify that the content provider is an authorized subscriber upon receiving the request. *See* EX1001, 9:59-63; EX2014, ¶104.

Recognizing the unsteadiness of its position, the Petition alternatively points to a “request [that] may be directed to the [PIM] from a neighboring IM” as teaching or suggesting *directed to the at least one shared repeater server based at least in part on said first alias name*. Pet., 22 citing EX1005, 40:29-33. The

Petition asserts that the “content coordinate data included in the video ID” can be used to “determine which IMs are likely to have the desired clip.” Pet., 22. But Kenner’s content coordinate data specifies “which categories of content are to be made available to the subscribers in its region.” EX1005, 44:27-30. In this fashion, a PIM can identify regional IMs that have subscribed to the category of the video ID. EX1005, 36:20-30 (“The content coordinates enable a match with the regional IMs that have subscribed to that type of file”). But this behavior does not teach the video ID (the purported *first/second alias name*) itself being associated with a particular PIM (the purported *at least one repeater server*) as required by the claims. If anything it just teaches the PIM (the purported *at least one repeater server*) being associated with other IMs. *See* EX2014, ¶105.

Indeed, Kenner’s video ID is not used as an “alias” for anything, rather, it is the only unique identifier of a video. EX2014, ¶94. For example, Kenner’s PIM, SRUs, and DSI each reference the same unique identifier of the video ID in a database. *See* EX1005, 21:1-6, 22:1-5, 22:30-23:9, 23:30-24:5.

’903 Patent <i>alias names</i>	Kenner’s Video ID
<i>Alternative identifiers</i> for a single repeater server	A <i>unique</i> identifier for a video .
Examples: “wrl.example.com” “wrl.anotherExample.com” “wrl.example.edu”	Constructed from: “a text name of the file,” the “content provider’s account number,” “a category coordinate,” a “geographic coordinate”, a “time stamp,” and a “time period” of relevance.

Compare EX1001, 9:44-55 *with* EX1005, 36:32-37:6.

To summarize, Kenner's video ID is not an alias name at all, let alone an alias name that is associated with a *resource* and *at least one shared repeater* and used to direct requests from a client machine to a repeater server.

B. Kenner does not teach or suggest *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* (all independent claims).

Claim 1 requires *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*. Each of the remaining independent claims (28, 37, 40, and 43) recites substantially similar features. *See* Appendix A (comparing the independent claims); EX2014, ¶¶107-10.

In the '903 patent, the repeater server receives a reflected request that includes an alias name (e.g., the repeater can examine the host tag of the HTTP header). EX1001, 10:4-7. Using the alias name, the repeater determines the origin server being addressed by referencing a table that "associates repeater alias names with origin server names." *Id.*, 9:56-59. The table indicates whether the origin server is a valid subscriber and whether that subscriber may use that particular repeater. *See id.*, 17:51-57. In this fashion, the repeater can identify the origin

server for the resource and assess whether a corresponding content provider is authorized to use that particular repeater. *Id.*, 9:56-59, 17:54-57; EX2014, ¶93.

Kenner does not teach or suggest this feature. *See* EX2014, ¶¶111-112. The Petition points to a “clip database” maintained on Kenner’s PIM as teaching the claimed *table*. Pet., 18. Kenner’s clip database includes the video ID and other information about the videos stored in the extended SRUs. EX1005, 34:3-6.

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.

EX1005, 34:7-11.

But the claims require that the *table* list content from content providers that are *authorized to have resources delivered to client machines via the at least one shared repeater server*. EX2014, ¶113. Kenner does not teach or suggest any such authorization across different content providers. The Petition points to the fact that Kenner’s video ID includes a content provider’s account number. *See* Pet., 19 *citing* EX1005, 36:32-37:6. The mere inclusion of the provider’s account number in the video ID does not indicate that a content provider *is authorized to have resources delivered to client machines via the at least one shared repeater server*. Kenner does not disclose authorization of content providers—presumably all

content providers are authorized to use the PIM in Kenner. The only permissions checked in Kenner are the access permissions of end users. *See id.*, 33:4-11.

Moreover, in the '903 patent, the claims require that the authorization be particular to the repeater server—i.e., *via the at least one shared repeater server*. *See* EX1001, 10:29-31 (“The subscriber tables described below include the information that is used to link reflectors to subscribers.”) That is, a particular content provider has permission or lacks permission to use a particular repeater server in the CDN. *Id.*, 17:56-57 (describing the subscriber table having a list of the repeaters that may be used by a subscriber’s servers). Kenner’s clip database does not include any information related to authorization, let alone authorization for a content provider to have resources delivered via a specific PIM (the purported *repeater server*). *See* EX1005, 34:5-35:3; EX2014, ¶114.

The repeater uses the table to *analyze an alias name received with a client request*—another feature not disclosed, taught, or suggested by Kenner. This deficiency highlights the Petition’s fundamental error in equating Kenner’s video ID to an *alias name*. Pet., 25 (identifying a PIM locating content at a source IM as teaching [1.g]). In Kenner, a PIM could identify the source IM with the video ID because the video ID itself contains the provider’s account number. *See* EX1005, 36:32-37:5. Kenner’s PIM need not analyze the video ID received in any request *using the table*. In Kenner, any additional analysis of the video ID to determine the

provider with respect to the clip database would be redundant. In contrast, in the claimed invention, a repeater analyzes the *alias name* received in the header of a request with reference to the *table* to determine an origin server/content provider. See EX2014, ¶115.

Thus, entirely absent from Kenner is any teaching of analyzing the purported *alias name* (“video ID”) with respect to a *table listing a plurality of alias names* (the PIM’s clip database). See EX2014, ¶116. Dr. Mowry provides no further insights to cover this gap in Kenner. Compare EX1003, ¶96, with Pet., 48-49.

X. CONCLUSION

The Board should deny *inter partes* review of 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: September 16, 2025

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

Appendix A: Comparison of Independent Claims

Independent Claims Comparison

<i>first/second alias name</i>	<i>alias -> repeater association</i>	<i>direction based on alias name</i>	<i>table</i>	<i>use of table</i>
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 <i>a first alias name</i>, wherein requests for a first resource located on a first origin server are directed, based at least in part on <i>said first alias name</i>, 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 <i>a second alias name</i>, wherein requests for a second resource located on a second origin server are directed, based at least in part on <i>said second alias name</i>, 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 <i>a first alias name</i> of said plurality of alias names is associated with a first origin server, and at least <i>a second alias name</i> 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 <i>a first alias name</i> corresponding to said first content provider, and wherein the server is associated with <i>a second alias name</i> corresponding to said second content provider, <i>said second alias name</i> being distinct from <i>the first alias name</i>, 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 <i>a first alias name</i> corresponding to a first origin server of the plurality of origin servers, and wherein the server is associated with <i>a second alias name</i> corresponding to a second origin server of the plurality of origin servers, said second origin server being distinct from the first origin server and <i>said second alias name</i> being distinct from <i>the first alias name</i>, 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 <i>first alias name</i> of the plurality of alias names, <i>said first alias name</i> 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 <i>said first alias name</i>; wherein at least a second resource associated with the second content provider is associated with a <i>second alias name</i> of the plurality of alias names, <i>said second alias name</i> 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 <i>said second alias name</i>; 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 <i>said first alias name</i>, 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 <i>said second alias name</i>, 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 6,673 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: September 16, 2025

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 2014-2015 were served in their entirety on September 16, 2025, upon the following parties via electronic mail:

Erika H. Arner (Lead Counsel)	erika.arner@finnegan.com
Kara A. Specht (Back-Up Counsel)	kara.specht@finnegan.com
Cara E. Regan (Back-Up Counsel)	cara.regan@finnegan.com
Sydney R. Kestle (Back-Up Counsel)	sydney.kestle@finnegan.com
Daniel C. Tucker (Back-Up Counsel)	daniel.tucker@finnegan.com
Christopher B. Anderson (Back-Up Counsel)	christopher.anderson@finnegan.com
Google-Sandpiper-IPRs@finnegan.com	

FINNEGAN, HENDERSON, FARABOW, GARRETT & DUNNER, LLP

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

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

Date: September 16, 2025

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