

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

GOOGLE LLC,
Petitioner

v.

SANDPIPER CDN, LLC,
Patent Owner

Case IPR2025-00860
U.S. Patent No. 10,924,573

PATENT OWNER'S 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 TECHNOLOGY	4
III.	OVERVIEW OF THE '573 PATENT	9
IV.	PROSECUTION HISTORY	11
V.	PERSON OF ORDINARY SKILL IN THE ART (POSA)	13
VI.	CLAIM CONSTRUCTION	13
VII.	GROUND 1	13
VIII.	GROUND 2	18
A.	Overview of References	19
1.	Overview of Dilley.....	19
2.	Overview of Pai	20
3.	Overview of Wang.....	21
B.	A POSA would not be motivated to combine Dilley with Pai.	22
C.	The references fail to render obvious the claimed determining “a popularity designation” (elements 1(b)/11(b)/19(b)).....	26
D.	The references fail to teach or suggest “ <i>requesting the resource from a second server ... based at least on the popularity designation,</i> ” as recited in independent claim 19.	29
1.	Petitioner relies on conclusory expert testimony to fill gaps and conflates internal distribution with external delivery.....	30
2.	The Petition arguments for why Dilley and Pai are combinable for this limitation are based on impermissible hindsight.....	34
E.	The references fail to render obvious the claimed “redirect instruction” that is sent from the second server of the content delivery network to the first server (elements 1(d)/11(d)).....	36
1.	The Petition fails to properly address obviousness for element 1[d].	38
2.	Petitioner erroneously alleges that Dilley’s second server instructs the edge server to retrieve content from the	

	origin server, or that Dilley does not provide specifics on how that retrieval happens.	41
3.	Dilley’s edge server, not the second server, determines where the edge server should fetch content based on the metadata, without receiving any redirect instruction from the parent server.	43
4.	Petitioner’s combination of Dilley and Wang is improper.....	45
a.	Petitioner’s theory of motivation relies on improper hindsight reconstruction.....	46
b.	Petitioner fails to establish a motivation to combine because its purported benefit of reducing latency or improving server availability will not be achieved	47
c.	The proposed modification results in a superfluous and redundant architecture.....	49
F.	The references fail to teach or suggest that the second server is a peer server that sends a redirect instruction to the first server, as recited in dependent claims 5 and 15	51
G.	Petitioner’s inconsistent positions in district court preclude Petitioner from showing the “receiving an instruction” limitation of claim 3.	53
IX.	CONCLUSION.....	54

TABLE OF AUTHORITIES

Cases

<i>Abbott Labs. v. Andrx Pharm., Inc.</i> , 452 F.3d 1331 (Fed. Cir. 2006)	10
<i>Belden Inc. v. Berk-Tek LLC</i> , 805 F.3d 1064 (Fed. Cir. 2015)	12, 13
<i>Honeywell Int’l Inc. v. DSM IP Assets B.V.</i> , IPR2024-00493, Paper 7 (P.T.A.B. Aug. 21, 2024).....	12
<i>Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.</i> , 821 F.3d 1359 (Fed. Cir. 2016)	11, 12
<i>In re Johnson</i> , 558 F.2d 1008 (CCPA 1977)	17
<i>Koito Mfg. Co., Ltd. v. Turn-Key-Tech, LLC</i> , 381 F.3d 1142 (Fed. Cir. 2004)	12
<i>LizardTech, Inc. v. Earth Resource Mapping, Inc.</i> , 424 F.3d 1336 (Fed. Cir. 2005)	11, 12
<i>In re Magnum Oil Tools Int’l, Ltd.</i> , 829 F.3d 1364 (Fed. Cir. 2016)	10, 11
<i>MaxLite, Inc. v. Jiaxing Super Lighting Elec. Appliance Co., Ltd.</i> , IPR2020-00208, Paper 10 (P.T.A.B. June 24, 2020)	11
<i>NJOY, LLC v. JUUL Labs, Inc.</i> , IPR2024-00536, Paper 17 (P.T.A.B. Aug. 12, 2024).....	12
<i>Purdue Pharma L.P. v. Faulding Inc.</i> , 230 F.3d 1320 (Fed. Cir. 2000)	11
<i>Regents of the Univ. of Cal. v. Broad Inst., Inc.</i> , 903 F.3d 1286 (Fed. Cir. 2018)	12
<i>Sirona Dental Sys. GmbH v. Institut Straumann AG</i> , 892 F.3d 1349 (Fed. Cir. 2018)	11

Statutes

35 U.S.C. §103	8
----------------------	---

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
M.P.E.P. § 2173.05(i)	17

PATENT OWNER'S UPDATED 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>Intentionally Left Blank</i>
2004	U.S. Publication No. 2004/0221019 A1 to Swildens et al., published November 4, 2004
2005	<i>Intentionally Left Blank</i>
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	Curriculum Vitae of Dr. Prashant Shenoy
2016	Declaration of Prashant Shenoy in Support of Patent Owner's Response

Exhibit No.	Description
2017	Deposition Transcript of Dr. Todd Mowry (Feb. 24, 2026)
2018	“Level 3 Lands Netflix Streaming Business, Will Double Its Storage Capability,” TechCrunch, available at https://techcrunch.com/2010/11/11/level-3-lands-netflix-streaming-business-will-double-its-storage-capability/ (last accessed March 1, 2026)
2019	<i>Intentionally Left Blank</i>
2020	“Connect – Most Innovative New Products Award Winners Announced,” Connect NewsLetter, October 26, 1999, available at https://library.ucsd.edu/dc/object/bb60436752/_1.pdf (last accessed March 1, 2026)
2021	“Sandpiper Named ‘Best of the Best’ at Networld + Interop 99 Las Vegas,” PRNewswire, May 12, 1999, available at https://ubmchannelnews.mediaroom.com/index.php?s=29573&item=85841 (last accessed March 1, 2026)
2022	Entrepreneur of the Year, June 19, 2000
2023	“Sandpiper Networks’ Footprint Takes Aim at Internet Traffic,” Wall Street Journal Interactive Edition, June 17, 1999, available at https://url.us.m.mimecastprotect.com/s/wTVWCxklA1iDQQWjSYi2Syt1OK?domain=wsj.com (last accessed March 3, 2026)
2024	“Web-Page Distribution System Could Unclog Internet Traffic Jams,” The New York Times, April 26, 1999, available at https://archive.nytimes.com/www.nytimes.com/library/tech/99/04/biztech/articles/26sand.html (last accessed March 3, 2026)
2025	“Sandpiper Puts Its Footprint in Marketplace,” Los Angeles Times, September 28, 1998, available at https://www.latimes.com/archives/la-xpm-1998-sep-28-fi-27201-story.html

Exhibit No.	Description
2026	“CenturyLink named a Leader in IDC MarketScape: Worldwide Commercial CDN Vendors Report,” PR Newswire, Oct. 18, 2019, available at https://www.prnewswire.com/news-releases/centurylink-named-a-leader-in-idc-marketscape-worldwide-commercial-cdn-vendors-report-300933112.html (last accessed March 1, 2026)
2027	Google’s Preliminary Claim Construtions, (<i>Sandpiper CDN, LLC v. Google LLC</i> , No. 2:24-cv-03951)

I. INTRODUCTION

U.S. Patent No. 10,924,573 (the “’573 patent”) provides improved methods for handling content delivery in a content delivery network (“CDN”). The patent discloses a tiered CDN architecture in which a first-tier server determines a “popularity designation” for a resource and, based on that designation, the system determines how the content should be obtained. This process includes processing a redirect instruction that allows unpopular content to be delivered without caching at the edge server. The unique way in which the challenged claims leverage the popularity designation to direct traffic, reduce latency, and create efficiency in storing and delivering content is not found anywhere in the prior art.

First, the Board has already rightly found that Ground 1 of the Petition fails as a matter of law. The ’573 patent is entitled to its April 4, 2008, priority date because the original disclosure reasonably conveys possession of both caching and non-caching instructions, and therefore adequately describes the “instruction to not cache” limitation that formed the basis for the Petition’s priority challenge.

Petitioner’s expert was asked in his deposition to identify any error in the Institution Decision’s priority determination, and did not. Accordingly, as the Board observed, the primary reference (Newton-471) asserted in Ground 1 post-dates the ’573 patent’s priority date and is not prior art to the challenged claims.

Ground 2 also fails for multiple reasons. First, the Petition advances an internally inconsistent obviousness theory by relying on different subsets of references for different claim limitations. Specifically, the Petition proposes a combination of three references: Dilley, Pai, and Wang. But for each limitation, the Petition relies on different ways of combining those references, often discussing only a combination of two references or a modification to one. The Petition does not explain how or why a person of ordinary skill in the art would have been motivated to combine *all three* of the relied-upon references in the particular way that the Petition combines them, with a reasonable expectation of success.

Second, the cited references—in combination—fail to teach determining “a popularity designation” as required by all independent claims (claims 1, 11 and 19). The claimed “popularity designation” determines whether a resource is eligible for popularity-based handling. This eligibility-based step is not disclosed in any of the prior-art references. The Petition purports to rely on Pai for this limitation, but Pai merely teaches using numerical popularity values for storage (distribution/redistribution). Pai treats all content as inherently subject to redistribution based on popularity values, which is fundamentally different from the claimed designation step that precedes popularity-based logic. Neither Dilley nor Wang discloses any concept of eligibility or a prerequisite popularity

determination step. Accordingly, the combination does not cure the deficiency in Pai. Silence plus silence does not equal disclosure.

Third, the references fail to teach or suggest “requesting [a] resource from a second server ... based at least on the popularity designation,” as required by independent claim 19. Dilley selects retrieval paths based on object-specific metadata, not popularity designations. Dilley’s infrastructure is, therefore, already based on different logic, and a POSA would have no reason to modify that architecture. Pai addresses internal storage optimization rather than real-time content delivery to end users. Accordingly, even if a POSA was motivated to modify Dilley’s architecture to incorporate Pai, the incorporation of Pai’s popularity service would solve a different problem than the problem addressed in the ’573 patent. The Petition attempts to bridge this gap through conclusory expert testimony that improperly conflates popularity values with the claimed popularity designation and conflates internal content distribution with external content delivery. That testimony cannot properly fill the gap in the prior art.

Fourth, the references do not disclose the claimed “redirect instruction” sent from a second server to a first server as required by independent claims 1 and 11. Dilley already discloses complete and autonomous retrieval mechanisms driven by metadata evaluated at the edge server, without any redirect instruction issued by a second server. Dilley’s logic therefore does not have any shortcoming for a

secondary reference to overcome. Even if it did, Wang's redirect mechanism is incompatible with Dilley's architecture and would be redundant of functionality Dilley already performs through other means. There is no motivation to combine Dilley with Wang's retrieval mechanism, and no reasonable expectation of success in arriving at the claimed redirect instruction.

Finally, dependent claims 5 and 15 recite that the second server be a peer server that sends a redirect instruction. The Petition fails to identify any disclosure in the prior art of a peer server issuing such an instruction—only a parent server. This deficiency is fatal to the obviousness theories for claims 5 and 15 because the alleged latency benefits for modifying and combining the references do not apply to peer-to-peer configurations. They only apply to the parent-edge server hierarchy discussed in the references. There is therefore no reason that a POSA would modify this parent server functionality to apply to peer servers as claimed.

The Board should uphold the patentability of all challenged claims.

II. BACKGROUND OF TECHNOLOGY

The challenged patent arises from pioneering technology directed to delivering content across a network. Content delivery networks (CDNs) provide the 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. EX2014,

¶25.

In the early 1990s, the World Wide Web saw increasing adoption, becoming 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 content-congested internet cables. Consumers living far from the physical server(s) hosting content experienced sluggish load times and high latency due to problems such as overloaded servers, congested network segments, and geographic separation. 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 appropriate content from customers' origin servers to appropriate CDN servers, and transparently rendezvous end users requesting that content to the "best" CDN server to deliver that content, all while providing their customers with control over their content and

user experience. This service and its architecture were quickly imitated by many others in the industry. 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 them or that had available capacity. Distributing content across a network of servers alleviated data congestion issues and allowing consumers to connect to edge servers located near them reduced 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. EX2014, ¶28.

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

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 patent challenged in this Petition.

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 inter alia, pictures, text files, dynamic resources, and streaming multimedia resources.

The industry-wide acclaim for Sandpiper Networks further confirms the non-obvious nature of its innovations. Sandpiper's pioneering CDN technology was honored with the "Best of the Best" award at Networld+ Interop 99 Las Vegas, and the 1999 Most Innovative New Product Award from the University of California, San Diego (UCSD) CONNECT program. EX2020; EX2021. This technical leadership was mirrored by executive recognition, including the Ernst &

Young 2000 Entrepreneur of the Year Award bestowed upon Mr. Swart and Mr. Farber. EX2022. Beyond formal awards, Sandpiper's disruption of the CDN space was extensively documented by premier national publications, including The Wall Street Journal, The New York Times, and The Los Angeles Times, all of which highlighted the company's novel approach to content delivery. EX2023; EX2024; EX2025.

In the early-to-mid 2000s, demand for CDNs exploded. To meet the era's critical demand for increased CDN bandwidth, Sandpiper and its successors invented several patented technologies that enabled massive scaling for video streaming. These legacy assets, including the challenged patents, were eventually transferred to Patent Owner, Sandpiper CDN, LLC ("Sandpiper" or "PO"). Sandpiper Networks continued to commercialize its technology and compete in this expanding market, doing business as various entities including Digital Island, Inc., a company called "Level 3," and CenturyLink (later changed name to Lumen). In 2010, Netflix selected Level 3 as a primary CDN provider to support its entire streaming library, requiring Level 3 to double its storage and add nearly 3 terabits per second of capacity. EX2018. By 2019, CenturyLink offered CDN services with peak capacity exceeding 120Tbps and served nine of the world's ten largest media companies. EX2026. As the market expanded, other businesses

entered the space—often by misappropriating Sandpiper’s innovations and directly competing with Sandpiper for market share.

One such company is Google. Google built its own in-house 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, as well as for Google’s video search results (powered by YouTube). 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 ’573 PATENT

The ’573 patent is entitled “Handling Long-Tail Content in a Content Delivery Network (CDN),” and it claims priority to U.S. Provisional Application No. 61/042,412, filed on April 4, 2008. EX2014, ¶30. CDN technology generally includes multiple servers used to serve content, which can create problems relating to responding to requests for content. For example, technical issues arise in the context of content delivery using a network and multiple servers, relating to handling requests for resources and timely providing resources to requesting devices. EX2014, ¶31.

The inventors of the '573 patent described solutions to these problems in the '573 patent. The '573 patent describes a CDN including a tiered server system. In some embodiments, a first-tier server attempts to respond to a user request for content. The first-tier server determines whether the requested content is subject to popularity-based handling by accessing a popularity service to determine a popularity designation for the resource. Based on that designation, the system determines how the content should be obtained. Specifically, if the first-tier server does not have the content, and the content is designated, that server will request the content from the higher-tier server and subsequently cache the content for future delivery, for example. The disclosed invention addresses the reality of storage limitations by setting forth a framework that automatically caches only popular content, thus both speeding up content delivery and preserving memory space for popular content, in some cases. EX2014, ¶32.

The inventors devised solutions, as shown in the embodiments claimed by the '573 patent—for example, methods for using tiers of servers and specific processes for obtaining content for a requesting device. The solutions recited in the claims of the '573 patent include particular combinations that were not conventional at the time of the invention of the '573 patent, such as specific technical improvements to content delivery networks, and the '573 patent describes how the solutions are implemented. EX2014, ¶33.

IV. PROSECUTION HISTORY

The '573 patent was examined as U.S. Application No. 16/284,791 (“’791 application”), which was filed on February 25, 2019. EX1001, (21), (22); EX2014, ¶34. In a Non-Final Office Action mailed on July 24, 2020, Examiner Winder rejected prosecution claim 11 on the ground of non-statutory double patenting over claim 8 of U.S. Patent No. 8,930,538 (“Fullagar”), and rejected claims 1-2, 4-7, 11-12, and 14-18 on the ground of non-statutory double patenting over claims 1-7 and 9-14 of U.S. Patent No. 10,348,848 (“Newton”). EX1007, 60-71. Additionally, claims 1, 3-8, 10-11, and 14-20 were rejected under 35 U.S.C. § 103 as unpatentable over U.S. Patent No. 7,552,233 (“Raju”) in view of U.S. Publication No. 2004/0221019 (“Swildens”), and claims 9 and 12-13 were rejected under 35 U.S.C. § 103 as unpatentable over Raju in view of Swildens and further in view of U.S. Patent No. 7,024,466 (“Outten”). EX1007, 71-77; EX2014, ¶35.

On September 17, 2020, Applicant filed a terminal disclaimer to overcome the non-statutory double patenting rejections and responded to the § 103 rejections by filing amendments and remarks. EX1007, 40-50. Specifically, Applicant amended claim 1 to include subject matter disclosed in the originally filed claim 2. Independent claims 11 and 19 are amended similarly. For example, claims 1 and 2 are amended as follows:

1. (Currently Amended) A method of content delivery in a content delivery network comprising:

receiving, at a first server of a first tier of servers of the content delivery network, a request from a requesting device for a resource available from the content delivery network;

accessing a popularity service associated with the content delivery network to determine a popularity designation associated with the requested resource;

requesting the resource from a second server of the content delivery network;

processing, at the first server of the first tier of servers, a redirect instruction from the second server of the content delivery network to obtain the resource from a content server of the content delivery network;

receiving an instruction to not cache the portion of the resource at the first server of the first tier of servers when the portion of the resource is obtained from the content server of the content delivery network; and

providing the obtained resource to the requesting device.

2. (Currently Amended) The method as recited in claim 1 wherein the popularity designation associated with the requested resource indicates that the requested resource is not popular, ~~the method further comprising: receiving an instruction to not cache the resource at the first server of the first tier of servers when the resource is obtained from the content server of the content delivery network.~~

EX1007, 40-41; EX2014, ¶36.

On December 31, 2020, the Examiner issued a Notice of Allowance.

Examiner Winder noted that the closest art (Raju and Swildens) do not teach

“sending an instruction to not cache the content item,” as recited in the claims.

EX1007, 17; EX2014, ¶37.

V. PERSON OF ORDINARY SKILL IN THE ART (POSA)

A POSA is presumed to be aware of all pertinent art. The POSA possesses the “understandings and knowledge reflected in the prior art” and thinks along conventional wisdom. *See Abbott Labs. v. Andrx Pharm., Inc.*, 452 F.3d 1331, 1336 (Fed. Cir. 2006). At this stage, PO 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.

VI. CLAIM CONSTRUCTION

For purposes of this Response, Patent Owner Sandpiper CDN, LLC agrees the *Phillips* standard applies. Patent Owner reserves the right to contest Petitioner’s implicit or explicit constructions during the proceeding.

VII. GROUND 1

The Board has correctly agreed that there is support for “instruction not to cache” in U.S. Provisional Application No. 61/042,412 removing Ground 1. DI, 13-14. Specifically, the Institution Decision correctly noted that “even with the words ‘instruction to not cache’ [] not literally recited in the Provisional Application, caching and not caching are alternative options that indicate the instruction to the edge server that will be implemented for each option.” DI, 14 (citing EX2014, ¶56).

The Board went on to note that “these disclosures provide support for Dr. Shenoy’s testimony that the disclosures ‘reasonably convey to those skilled in the art that Patent Owner had possession of the edge server receiving an instruction to not cache the resource (and instead simply forward the redirect).’” DI, 14 (citing EX2014, ¶56). The Institution Decision also correctly explained that “there is sufficient evidence, in the context of the disclosures of written description support in the Provisional Application, providing support for Dr. Shenoy’s testimony that Patent Owner had possession of an edge server receiving an instruction to not cache the resource.” *Id.*, 15.

Accordingly, the Institution Decision rightly concluded that “the ’573 patent has a priority date of April 2, 2008, which predates the earliest possible date on which Newton-471 could be effective as prior art, September 30, 2014 (the date on which the earliest application in the priority chain of Newton-471 was filed).” *Id.* The Board preliminarily found “that Newton-471 is not prior art to the challenged claims,” finding, in turn, that “Petitioner has not demonstrated a reasonable likelihood that it would prevail in showing that claims 1–20 are anticipated by Newton-471.” *Id.* The Board should maintain its holding that Ground 1 fails for these reasons.

As Patent Owner explained, the claimed instruction to not cache is supported by the written description of the ’573 patent, the Provisional Application No.

61/042,412 (“’412 application”) to which the ’573 patent claims priority, as well as all the written descriptions of the patents in the priority chain of the ’573 patent.

EX2014, ¶51; EX1001, 7:25-30; EX1011, ¶44. For instance, the written descriptions of the ’573 patent and the ’412 application teach two alternate instructions that lead to two corresponding actions. EX1001, 7:25-30; EX1011, ¶44; EX2014, ¶52.

A redirect from the popularity service without the follow-me tag is an instruction to not cache and to forward the redirect. EX1001, 7:25-30; EX1011, ¶44; EX2014, ¶52. Alternatively, a redirect with the follow-me tag is an instruction to cache the resource and serve it to the client. EX1001, 7:25-30; EX1011, ¶44; EX2014, ¶52. The ’573 patent and the ’412 application specifications thus disclose both the function of the instruction to not cache (forwarding without caching) and the form the instruction takes (a redirect without the follow-me tag). EX2014, ¶53.

Additionally, the written description of the ’573 patent, the Provisional Application No. 61/042,412 to which the ’573 patent claims priority, as well as all the written descriptions of the patents in the priority chain of the ’573 patent, describe several technical reasons for instructing the edge server not to cache unpopular resources. *See, e.g.*, EX1001, 10:17-41; EX1011, ¶¶68-69; EX2014, ¶54. For instance, the written descriptions of the ’573 patent and the ’412 application disclose that caching unpopular resources at an edge server “increase[s]

bandwidth costs,” “adds latency,” and “risks forcing actually popular resources out of cache.” EX1001, 10:17-41; EX1011, ¶¶68-69; EX2014, ¶54. Therefore, the written descriptions of the ’573 patent and the ’412 application not only teach two alternate instructions that lead to two corresponding actions—(1) cache or (2) not cache at an edge server—but also teach not to cache unpopular resources at an edge server for various reasons. EX1001, 7:25-30, 10:17-41; EX1011, ¶¶44, 68-69; EX2014, ¶55.

The fact that the words “instruction to not cache” are not literally present in the ’573 patent and the ’412 application specifications does not matter. *In re Johnson*, 558 F.2d 1008, 1019 (CCPA 1977); M.P.E.P. § 2173.05(i). The ’573 patent and the ’412 application specifications present caching and not caching as alternative options and describe an embodiment that illustrates what instruction to the edge server will implement each option. EX1001, 7:25-30; EX1011, ¶44. The ’573 patent and the ’412 application specifications also describe technical reasons for not caching unpopular resources at an edge server. EX1001, 10:17-41; EX1011, ¶¶68-69; EX2014, ¶56. Those disclosures reasonably convey to those skilled in the art that Patent Owner had possession of the edge server receiving an instruction to not cache the resource (and instead simply forward the redirect) as of the earliest priority date. *Novartis Pharm. Corp. v. Accord Healthcare, Inc.*, 38 F.4th 1013, 1020 (Fed. Cir. 2022) (quoting *Ariad Pharm. Inc. v. Eli Lilly & Co.*,

598 F.3d 1336, 1351 (Fed. Cir. 2010)). That satisfies the written description requirement for priority.

Petitioner's expert does not appear to disagree with this aspect of the Institution Decision. Specifically, when asked if he had read the Board's Institution Decision finding, at least preliminarily, that the challenged claims were entitled to the April 4th, 2008 priority date, Dr. Mowry replied that he had. *See* EX2017, 14:13-17. When asked whether he thought the Board's priority decision was wrong, he did not say that it was and did not offer any basis to disagree. EX2017, 14:18-15:6. To the contrary, he stated "I accept the decision of the [B]oard," in that respect. *Id.* He also observed that he relied upon the Wang reference elsewhere in his report for expressly disclosing the "instruction not to cache" limitation, and characterized the relied-upon disclosure in Wang as "identical" to the written description of the instruction not to cache in the '573 patent (and the '412 application). *See* EX2017, 83:12-84:11.

Accordingly, the priority date of the '573 patent should be April 4, 2008—the filing date of the '412 application. Petitioner does not dispute that if the '573 patent is afforded the benefit of this priority date, then the art cited for Ground 1 does not qualify as prior art. Ground 1 therefore fails.

VIII. GROUND 2

Although Ground 2 relies on a proposed combination of three references—Dilley, Pai, and Wang—it never articulates how or why a POSA would have been motivated to combine all three, or would have had an expectation of success in doing so. Instead, the Petition presents different permutations of Dilley with Pai or Dilley with Wang for different limitations. That legal defect alone is fatal and cannot be cured in the Reply. Even those sub-combinations are flawed, in that Petitioner takes disparate teachings from the references, modifies them to resolve their incompatibilities, and then combines them to arrive at the claimed system with clear hindsight bias. That is another legal error.

Even in combination, the cited references do not teach or suggest all of the limitations of the claims. First, the combined references do not teach the claimed determination of a “popularity designation,” in the independent claims. Petitioner improperly equates Pai’s internal content-storage logic with the claims’ real-time delivery requirements. The Petition also fails to show how the references disclose the claimed redirect instruction sent from a second server to a first server. Once again, the Petition relies on hindsight-driven expert testimony to manufacture gaps that do not exist in the prior art and proposes combinations that are redundant, ignore the references’ express teachings, or defy their principles of operation.

Petitioner also fails to show the additional limitations of independent claim 19, as well as dependent claims 3, 5, and 15. Specifically, the references fail to render obvious the use of the popularity designation to determine where to request the resource (claim 19), the receiving an instruction to cache (claim 3), and the second server is a peer server that sends a redirect instruction to the first server (claims 5 and 15).

A. Overview of References

The three cited prior-art references present three different approaches to the issue of storing and retrieving content online. They differ from each other and from the system described and claimed in the '573 patent in both structure and function. EX2016, ¶57.

1. Overview of Dilley

Dilley discloses a tiered CDN using a cache hierarchy feature. EX1003, Abstract, 5:31-36. Dilley's tiered CDN includes edge servers, distribution hubs (or parent servers), and origin servers. EX1003, 4:64-5:6. The CDN system routes an object request from a client to an edge server. EX1003, 8:3-11. If the edge server does not have the requested object, Dilley describes two "request routing mechanism[s]" for the edge server to serve the request object: (1) a default mechanism where the edge server contacts the origin server directly, or (2) serve the requested object using a parent server in the hierarchy. EX1003, 6:10-15, 8:34-

42. Dilley selects the routing mechanism based on “object-specific metadata.”

EX2016, ¶58; EX1003, 4:23-26, 6:14-15, 8:23-48.

More specifically, Dilley’s system uses “object-specific metadata” to handle a request for an object. EX1003, 4:23-26. Dilley’s metadata includes a “use-hierarchy” tag. EX1003, 6:25-61. Dilley explains that the default setting of “OA” for the tag indicates that the edge server should “go direct to the (O)rigin-server specified in the (A)RL” without using the cache hierarchy feature. *Id.* Setting the tag to be “coremap” turns on the cache hierarchy feature and instructs the edge server to do a DNS lookup for a parent server. *Id.*; EX2016, ¶59.

According to Dilley, popular content or large files are serviced using the cache hierarchy so that the origin servers are not overwhelmed. EX1003, 5:16-21, 6:4-7. Additionally, the cache hierarchy “reduces average latency” and “improves connectivity” for “poorly-connected edge regions.” EX1003, 6:7-10; EX2016, ¶60. Dilley does not disclose a popularity service or popularity designation. Because Dilley’s cache hierarchy is fully functional using object metadata, there is no role that a popularity designation or popularity service would play in Dilley’s system.

2. Overview of Pai

Pai discloses a storage optimization system that stores (distributes or redistributes) media content within an on-demand media system based on the popularity value of the media content. EX1004, ¶¶13, 20, FIGs. 6, 7. Recognizing

that media content is “very expensive to store,” Pai proposes to store more popular content closer to the client. EX1004, ¶¶2, 13. Pai also dynamically redistributes media content to media sources of different storage levels based on updated popularity values. EX1004, ¶¶39-51, FIGs. 6, 7; EX2016, ¶61. Pai’s popularity logic is based on storage and solving storage issues. It is not related to retrieval of content. As such, there is no determination of a popularity designation as part of the retrieval process in Pai.

3. Overview of Wang

Wang discloses a CDN with a hierarchical service network that uses light weight directory access protocol (LDAP) to exchange directory information between service managers of the network. EX1005, Abstract, ¶¶4, 8, 22. For example, when a user requests content from the CDN, an integrated service switch directs the request to a caching service engine. If it does not have the content, then the network sends LDAP search requests up the hierarchy (Level 1 to Level 2, etc.). EX1005, ¶123. If Wang’s level 3 service manager returns the origin server IP address with an indication of whether the content is cacheable or not, then the caching service engine will either “http-redirect” the client to the origin server, or “initiate” a new http session on behalf of the client to the origin server based on the indication. EX1005, ¶¶123-124; EX2016, ¶62. This logic is based on whether the content is determined to be cacheable, and not whether the content has a particular

popularity designation. In other words, everything that can be cached will be cached in Wang. There is no concept or infrastructure in Wang for choosing not to cache something that is cacheable based on popularity (or any other attribute).

B. A POSA would not be motivated to combine Dilley with Pai.

Petitioner's motivation to combine Dilley and Pai rests on a fundamental mischaracterization of both references. First, Petitioner erroneously alleges that Dilley determines where to "cache content" as a function of metadata associated with the requested object. Pet., 37 (citing EX1003, 2:42-56). Next, Petitioner conflates Pai's popularity value with the claimed *popularity designation*. Section VIII.C. Compounding those errors, Petitioner contends that Dilley's "object specific metadata" can be used to facilitate control over a "tiered distribution of content" using Pai's popularity-based method. Pet., 37.

While the Board's Institution Decision correctly acknowledged that Dilley's metadata is used to determine where to fetch content, DI, 24, the Board suggested that Pai's popularity service would be used for "content placement" in the tiered cache hierarchy of Dilley. *Id.* However, this finding fails to bridge the gap between the references' teaching of content *placement* (internal storage) based on popularity value of Pai and content *delivery* (external routing) based on the popularity designation as recited by the claim.

Dilley’s “object specific metadata” is used exclusively to select which server to **fetch** content, not to determine which server will **cache** content as claimed. *See* Section VIII.A.1; EX1003, 4:23-26, 6:14-15, 8:23-48. More importantly, Petitioner improperly conflates Pai’s popularity value with the claimed popularity designation, and Pai’s internal **storage/distribution** with Dilley’s **external delivery/routing**. Therefore, Petitioner fails to provide a coherent motivation to combine the references or satisfy the claim requirement of “*requesting the resource ...based at least on the popularity designation.*” EX2016, ¶81.

Petitioner’s inconsistent treatment of the “popularity service” limitation is further highlighted by Petitioner’s proposed claim construction in district court. In the parallel litigation, Petitioner proposed to construe the term “popularity service” as “a component or functionality that determines the popularity of content and provides information used to determine where the content will be served from.” EX 2027, 3. Petitioner then relies on Pai to provide the inner workings of the popularity service in the proposed Dilley-Pai-Wang combination. But Pai does not provide information used to determine where the content will be served from. Pai’s popularity functionality only determines where content will be stored. Petitioner thus relies upon a much broader construction of “popularity service” in this proceeding—relying on a reference that would not satisfy Petitioner’s own claim construction in district court.

The Office has made clear that allowing petitioners to propose different constructions at the Board and the district court, such as seeking narrower constructions in litigation to avoid infringement liability, is contrary to the goal of “providing greater predictability and certainty in the patent system.” *Revvo Techs., Inc. v. Cerebrum Sensor Techs., Inc.*, IPR2025-00632, Paper 20 at 4 (P.T.A.B. Nov. 3, 2025) (Squires, Director) (precedential) (internal citations omitted). Rules 104(b)(3) and 104(b)(4) require a petitioner to explain “how the challenged claim is to be construed” and “how the construed claim is unpatentable.” 37 C.F.R. §§ 42.104(b)(3)-(4). Petitioner’s treatment of this limitation illustrates precisely why the Board disfavors that approach.

Furthermore, a POSA would not have been motivated to combine Dilley and Pai because their teachings are incompatible with each other, rendering one or both inoperable for their intended purpose or changing their principle of operation. Petitioner intentionally conflates Dilley’s edge servers with its distribution hubs or parent servers to manufacture a rationale for combination that a POSA would have found technically incompatible. Specifically, Petitioner incorrectly alleges that “[b]ecause the edge servers are used to prevent the origin servers from becoming ‘overwhelmed with requests for popular content,’ unpopular content would be served from parent servers while ‘popular content would be served from the edge.’”

Pet., 83-84 (citing EX1003, 5:10-21). This is a misreading of the prior art.

EX2016, ¶¶83-84.

In reality, the very passage Petitioner cites clarifies that Dilley’s “**distribution hubs**” (parent servers)—which are distinct from edge servers—act as buffers to guarantee that the origin site is not “overwhelmed with requests for **popular content** or large files.” EX1003, 5:10-21. Dilley explicitly distinguishes its “edge servers 302a-n” from “distribution hubs 304a-n,” teaching that its tiered infrastructure uses “a set of edge servers and a hierarchical set of ‘parent’ or hub regions that service the edge servers.” EX1003, 4:63-5:7, FIG. 3; EX2016, ¶85.

On the other hand, Pai’s level one media sources store the most popular content while level two media sources (the second server) store less popular content. EX1004, ¶¶38-45. Therefore, Pai’s system stores *less* popular content in a second server while Dilley serves *more* popular content from a second server (e.g., hubs or parent servers). EX2016, ¶86. This difference is not merely a design choice that a POSA would adapt in combining the two references. The reasons that Dilley and Pai take opposite approaches to storing popular content is because they propose fundamentally different solutions to the problem of managing network traffic. This technical incompatibility would render one or both references inoperable for their intended purpose and/or change their principle of operation, which demonstrates the lack of obviousness to combine the references. This error


infects every argument of Ground 2, all of which rely on the same flawed reading of the references to combine Dilley with Pai.

C. The references fail to render obvious the claimed determining “a popularity designation” (elements 1(b)/11(b)/19(b)).

The Petition fails to demonstrate that the prior art teaches “*accessing a popularity service associated with the content delivery network to determine a popularity designation associated with the requested resource,*” as recited in independent claims 1, 11, and 19. EX2016, ¶¶63-68. Petitioner relies on *Pai alone* for this teaching but fails to point to any disclosure in *Pai* that satisfies the limitation. Pet., 44-46. Nothing in *Dilley* remedies this deficiency and the Petition does not allege otherwise. Wang is not mentioned for this feature.

The '573 patent describes a specific operation. The system “determine[s] a popularity designation” for a requested resource—i.e., whether the specific resource is even eligible for popularity-based routing. EX1001, 6:27-31 (“whether this object has been designated so that the object’s popularity will be used to determine where it will be served from”). After determining a resource is so “designated” the system performs popularity-based routing: comparing the resource’s popularity (e.g., a “popularity count”) against specific thresholds to determine which server to fulfill the request. EX1001, 6:53-7:24, FIG. 5; EX2016, ¶64.

As the '573 patent explains, “not necessarily all” content uses the popularity service; rather, a “popularity check” is performed only for “certain designated content.” EX1001, 5:47-50, 5:62-64. Thus, this claimed “*determining a popularity designation*” (Step 506 of Figure 5) occurs *before* the “popularity checking” (Steps 512 and 516 of Figure 5) is ever initiated. EX1001, 6:27-41, FIG. 5; EX2016, ¶65.



EX1001, FIG. 5 (annotated).

In contrast, Pai discloses a system that distributes or redistributes *all* media content based on “current popularity values.” EX1004, ¶39. While Pai evaluates these numeric popularity values to manage storage of *all* media content across media sources, Pai lacks any teaching of **determining** a “*popularity designation*” for its media content. EX1004, ¶¶39-51. Indeed, Pai has no popularity designation, Pai has no reason to check for one, and there is no step or action to perform based on the popularity designation. Petitioner effectively admits this deficiency by pointing only to Pai’s determination of a “popularity value” associated with media content. Pet., 45; EX1004, ¶39. A POSA would instantly distinguish Pai’s “popularity value” from the claimed “*popularity designation*.” EX2016, ¶¶66-67.

Indeed, Petitioner’s own expert, Dr. Mowry, confirmed this confusion during his deposition. When asked to identify the component that determines a popularity designation, Dr. Mowry could only point to Pai’s teaching of “dynamically determin[ing] popularity values.” EX2017, 92:11-18 (“Q: In your proposed combination for Dilley, Pai, and Wang, what is the thing that determines a popularity designation? A: ... Pai teaches how to dynamically determine popularity values to determine popularity.”). Dr. Mowry further admitted that Pai’s determination of popularity value is “the same concept” as using the threshold levels described in the ’573 patent. EX2017, 92:18-93:1. However, as shown in the annotated version of Figure 5 above, the ’573 patent’s popularity checking steps

(steps 512 and 516) using thresholds occur *after* the determination of popularity designation (step 508). EX1003, FIG. 5.

Because Pai treats *all* media content as inherently subject to its distribution/redistribution logic, it never performs the claimed determination of a “popularity designation.” Accordingly, Pai does not “*determine a popularity designation*,” and Dilley or Wang are not alleged to, nor do they, cure the deficiencies of Pai. EX2016, ¶68. Any combination of Pai with Dilley and Wang would therefore treat all media content as subject to Pai’s distribution logic. The potential caching instruction dichotomy from Wang would be irrelevant to Pai’s popularity service, and do nothing to create the step of determining a popularity designation where none exists in Pai. Ground 2 therefore fails.

D. The references fail to teach or suggest “*requesting the resource from a second server ... based at least on the popularity designation*,” as recited in independent claim 19.

The Petition fails to demonstrate obviousness for at least “*requesting the resource from a second server ... based at least on the popularity designation associated with the requested resource*,” as recited in claim 19. EX2016, ¶¶69-86. No combination of the cited prior art teaches a structure or operation that uses the determined popularity designation to request a resource from a particular server. Petitioner relies exclusively on the testimony of Dr. Mowry to fill this gap. Pet.,

83-85 (citing EX1006, ¶¶352-355). As discussed in detail below, neither Dilley nor Pai discloses this feature. EX2016, ¶70.

Second, Petitioner mischaracterizes Dilley in order to advance its flawed motivation to combine Dilley and Pai. Instead, as discussed below, a POSA would not have been motivated to combine Dilley and Pai because their teachings are incompatible with each other. EX2016, ¶71.

1. Petitioner relies on conclusory expert testimony to fill gaps and conflates internal distribution with external delivery.

Petitioner relies exclusively on the conclusory testimony of Dr. Mowry to allege that the combination of Dilley and Pai, again with no mention of Wang, teaches requesting content from the second server “*based ... on the popularity designation.*” However, Petitioner fails to meet its burden because neither reference, alone or in combination, discloses this claim feature.

The statutory basis for *inter partes* review limits invalidity grounds to prior art “consisting of patents or printed publications.” 35 U.S.C. § 311(b). The corresponding regulation requires a petition to “specify where each element of the claim is found in the prior art patents or printed publications relied upon.” 37 C.F.R. § 42.104(b)(4); *see also* 35 U.S.C. § 312(a)(3); EX2003 (Office guidance on the proper application of Rule 104(b)(4); *Henny Penny Corp. v. Frymaster LLC*, 938 F.3d 1324, 1331-32 (Fed. Cir. 2019). The Petition here fails to do this. Instead, the expert testimony on which the Petition relies attempts to fill gaps in the prior

art with conclusory opinion about what a POSA would allegedly understand that the disclosed systems *could* do, rather than what they are described as doing.

First, Dilley determines *a second server* (a parent server) based on object-specific metadata, not based on a “*popularity designation*.” Section VIII.A.1. Dilley’s object-specific metadata includes a “use-hierarchy” tag, where specific values identify different destination, e.g., “OA” indicates the second server being an “(O)rigin-server specified in the (A)RL” and “coremap” indicates a hostname or IP address of a parent server (i.e., a second server). EX1003, 6:30-61, 7:14-17. Thus, by evaluating the metadata associated with a requested object, Dilley’s edge server “determines the name of the server region through which the request is funneled to fetch or validate the object being requested.” EX1003, 7:14-17. In fact, the Board previously agreed that Dilley’s metadata is used to determine where to fetch content from. DI, 24.

Petitioner relies only on Dr. Mowry’s unsubstantiated testimony to assert that “based on an unpopular designation, Dilley would retrieve content from the second server.” Pet., 84 (citing EX1006, ¶¶352-353). Even Dr., Mowry’s use of the hypothetical “would” betrays his reliance on unsupported extrapolation from the disclosure of the reference. Nothing in Dilley discloses a circumstance under which Dilley retrieves content from a second server based on an unpopular designation. Dilley’s object-specific metadata is fundamentally distinct from the

recited “*popularity designation*,” and Dilley never teaches determining a popularity designation. Thus, Dilley cannot teach “*requesting the resource from a second server... based at least on the popularity designation*.” EX2016, ¶74. Dr. Mowry’s hypothesis is based on what he alleges a POSA would expect Dilley to do if the system in Dilley performed steps that it does not perform.

Pai is equally deficient. Pai discloses a mechanism for distributing or storing media content within a CDN based on popularity values. EX1004, ¶¶13, 20, 39, FIGs. 6, 7. In other words, Pai’s distribution of media content is merely storing content **within the CDN**. However, the claimed feature requires requesting *the resource* from a specific *second server* for delivery to end users **outside the CDN**. Petitioner simply ignores this distinction and assumes that a POSA would modify the structure and operation of Pai using improper hindsight.

Additionally, as established above, Pai does not teach the claimed popularity designation. *See* Section VIII.C. Petitioner again exclusively relies on the conclusory testimony of Dr. Mowry to allege that “when Pai’s system receives a request for content that does not exceed the highest popularity threshold, the request would need to be forwarded to a second server.” Pet., 84 (citing EX1006, ¶354). Once again, Dr. Mowry opines on what Pai “would need” to do in a hypothetical situation that is not disclosed in Pai and that Pai’s system would not encounter as it is currently configured. In doing so, Dr. Mowry improperly

conflates Pai’s popularity value with the claimed popularity *designation*—two attributes that are different in kind and serve different purposes. Moreover, Pai does not teach or suggest forwarding a request to a second server based on a popularity threshold, much less based on a popularity designation, as claimed. Pai’s internal storage optimization does not involve selecting a delivery server *based at least on the popularity designation*. EX2016, ¶76.

The Board’s Institution Decision suggested that Patent Owner argued Pai fails merely because “it does not disclose ‘requesting the resource from a second server of the content delivery network.’” DI, 28 (citing POPR, 29). To be clear, Pai’s failure is more specific and more serious than that. Pai fails because it does not disclose requesting the resource from a second server *based at least on the popularity designation*. Pai’s system does not and cannot take any action based on a popularity designation that it does not have. To the extent Pai uses popularity at all, it is used only to determine where to store content within the CDN. Popularity plays no role in determining where to fetch content in Pai’s CDN. Pai’s logic and infrastructure does not provide for the use of a popularity designation in requesting the resource from a second server.

Therefore, neither Dilley nor Pai teaches “*requesting the resource from a second server ... based at least on the popularity designation associated with the requested resource*,” as recited in claim 19, alone *or in combination*. Wang is not

mentioned in the Petition for this feature, and also does not cure the deficiency of Dilley and Pai. The Board should find Ground 2, and its support of claim 19, fails.

2. The Petition arguments for why Dilley and Pai are combinable for this limitation are based on impermissible hindsight.

Independent from a POSA's lack of motivation to combine Dilley and Pai generally, discussed above, Petitioner's proposed combination of Dilley with Pai for the "requesting a resource" limitation of claim 19 is specifically flawed. Petitioner's proposed combination of Dilley and Pai is a classic product of impermissible hindsight by improperly conflating **internal distribution** (a background storage logic that stores content closer to end-users) with **external delivery** (real-time requesting and serving content to an end user). This fundamental mismatch is exacerbated by Petitioner's and Dr. Mowry's failure to distinguish between Pai's numerical "popularity value" with the claimed *popularity designation*. See Section VIII.C. There is simply no suggestion in either reference that a background storage metric (i.e., Pai's popularity value) could, or should be, used as a real-time routing trigger. Petitioner's leap from Pai's storing based on "popularity values" to the claimed "*requesting ... based ... on [] popularity designation*" is a narrative driven by what is claimed, but is not supported by the prior art without distorting its teachings to meet the claims. EX2016, ¶78.

The Board’s Institution Decision appears to suggest Patent Owner addressed these references “separately instead of in combination.” DI, 28. But that is incorrect. Patent Owner’s position is that the proposed three-way combination of prior art is based *not* on reading the teachings together, but on significantly modifying both Dilley and Pai in order to achieve a result that neither reference contemplates based entirely on what a POSA “would” understand according to Dr. Mowry’s say so. Indeed, Petitioner admits—and the Board appears to agree—that the proposed combination teaches only “distributing” or **storing** popular content at Dilley’s edge “closer to the requesting end-users.” Pet., 84; DI, 28. The storage-centric logic remains structurally and functionally distinct from the claim requirement of *requesting the resource from a second server* where that specific second server is selected *based at least on the popularity designation*. EX2016, ¶79. In other words, the proposed combination of references teaches a system for storing popular content closer to edge servers. It does not teach selecting servers from which to fetch content based on a popularity designation. Petitioner only arrives at that additional element by modifying the prior art to perform a step that it does not perform even in combination, in order to achieve a hindsight-driven result that neither reference suggests. Thus, Ground 2 remains unproven.

E. The references fail to render obvious the claimed “redirect instruction” that is sent from the second server of the content delivery network to the first server (elements 1(d)/11(d)).

Petitioner relies on a proposed Dilley-Pai combination to disclose the claim elements 1[a]-[c] and 11[a]-[c] , but then switches to a proposed Dilley-Wang combination for limitation 1[d]/11[d]. As a result, rather than presenting one coherent combination of references that allegedly discloses all of the limitations of claims 1 and 11, Petitioner treats each limitation as a separate task that can be found obvious by a separate combination. That is improper as a matter of law.. *See* Pet., 34-40, 50-52.

Furthermore, the proposed combination of Dilley and Wang fails to render obvious the challenged independent claims 1 and 11 at least because the combination fails to teach or suggest “*processing, at the first server of the first tier of servers, a redirect instruction from the second server of the content delivery network to obtain the resource from a content server of the content delivery network,*” as recited in independent claim 1, or “*transmitting a redirect instruction intended for the requesting device to the edge server to obtain the resource from a content server,*” as recited in independent claim 11. EX2016, ¶¶87-112.

Petitioner and the Board rely on the flawed testimony of Dr. Mowry, who improperly assumes that Dilley’s second server (“a peer server or a parent server”) would “instruct the first server to request the content from server even higher on

the hierarchy, such as a parent or origin server.” Pet., 50; EX1006, ¶251; DI, 25-

26. There is no teaching in the reference itself that supports this assumption. There is no suggestion that Dilley’s second server sends redirect instructions at all.

Petitioner alleges that “[w]hile Dilley discloses that the content can be retrieved from the origin server ... it does not provide specifics on how that retrieval happens.” EX1006, ¶252. But that silence is not only an invitation to speculate on how Dilley will retrieve the content, but to propose that it does so by performing a step that Dilley is not even capable of performing. A POSA would not understand Dilley’s parent server to send a “redirect instruction” to an edge server to obtain the resource from the origin server because Dilley does not manage traffic through redirection instructions. A POSA would have no reason to change Dilley’s operation other than Petitioner’s hindsight-driven desire to implement such a redirect instruction from Wang. Pet., 50 (citing EX1006, ¶253).

Petitioner also argues that combining Dilley[-Pai] and Wang would “decrease latency” by removing the parent server as a “middleman,” supposedly allowing “the edge server to communicate with and receive content directly from the origin server.” Pet., 51 (citing EX1006, ¶¶255-257). However, that motivation is irrelevant because Dr. Mowry’s understanding of Dilley is misplaced. Dilley’s parent server is not a middleman to begin with, and modifying Dilley’s parent server to send a redirect instruction that it does not currently send gives the parent

server more of a middleman role, not less. Dilley's established retrieval path for requested object already functions as intended and the proposed combination offers no actual latency benefit. Dr. Mowry is proposing to solve a problem in Dilley that does not exist. A POSA would not have been motivated to combine Dilley and Wang to implement the claimed "*redirect instruction*" because there is no benefit to Dilley's system in doing so. EX2016, ¶89.

1. The Petition fails to properly address obviousness for element 1[d].

In addition to a POSA's lack of motivation to combine Dilley and Pai generally, discussed above, Petitioner further fails to demonstrate how Dilley, *Pai*, and Wang are combinable for limitation 1[d] specifically. In other words, the proposed obviousness combination is a three-way combination of Dilley, Pai, and Wang. Element 1[d], however, presents only a combination of Dilley and Wang. The combination does not include Pai, and does not explain how the limitation would be satisfied by a three-way combination of all references. This failure is significant because Pai's disclosure is incompatible with Dilley in this context. Pai's popularity logic is based on storage and solving storage issues. It is not related to retrieval of content. As such, there is no determination of a popularity designation as part of the retrieval process in Pai. The Petition therefore cannot rely on Pai to provide a popularity service, and then fail to explain how that popularity service would work in performing a function it cannot perform.

Rather than resolve that incompatibility, Petitioner simply ignores Pai in discussing this limitation. Indeed, the only arguments in the Petition for combining all three references are directed to a completely different feature. Pet., 39-40 (discussing the “instructions to not cache” feature when discussing how, why, and expectation of success of the three-reference combination, not the “redirect instruction” feature).

The Petition’s strategy of presenting a different combination of references for each limitation of claim 1 is part of a larger pattern. For elements 1[pre]-1[a], Petitioner relies on Dilley; element 1[b] relies on Dilley in view of Pai; element 1[c] goes back to relying on Dilley alone; element 1[d] relies on Dilley in view of Wang, ignoring Pai or how Wang would be combinable with Dilley and Pai; and element 1[f] relies once again on Dilley alone. While Petitioner asserts Ground 2 based on all three references, its analysis of element 1[d] relies exclusively on the combination of Dilley and Wang, omitting any mention of Pai or explaining the role of Pai in the combination for this limitation. *See* Pet., 34-40, 50-52.

To be clear, there is nothing problematic about finding a particular teaching in a particular reference. But where the Petition relies on a combination of references to teach the claim, the Petition must explain how the *combination* teaches each limitation, and not how isolated portions of the references or sub-combinations might teach the limitation. The test for obviousness is “what the

combined teachings of the references would have suggested to one of ordinary skill in the art,” and *all* teachings in the prior art must be considered. *In re Young*, 927 F.2d 588, 591 (Fed. Cir. 1991) (emphasis added).

This deficiency in the Petition is particularly problematic where, as here, the references are not combinable or compatible in certain respects. The law requires that once two references are required for a claim feature, adding a third reference for a subsequent claim feature requires articulating all *Graham* factors: what RefA+RefB teach, what RefA+RefB lack, what RefC teaches, and why it would be obvious to combine RefA, RefB and RefC. *See Graham v. John Deere Co.*, 383 U.S. 1, 17-18 (1966). This analysis is missing for at least element 1[d]. Crucially, this omission creates a fundamental gap in the record that Petitioner cannot remedy now that the trial is underway.

The Petition must identify “with particularity” the evidence supporting the grounds for challenge, and a Petitioner may not raise “an entirely new rationale” in a reply to fill gaps in its initial theory. *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1369-70 (Fed. Cir. 2016). By failing to integrate Pai into the analysis of element 1[d], Petitioner fails to explain how a POSA would have been motivated to combine the teachings of *all* three references to achieve the claimed invention. Petitioner’s burden is not merely to show that the elements of

1[d] existed in Wang, but how, why, and the expectation of success of adding Wang to Dilley-Pai to meet the missing feature. The Petition fails in that respect.

2. Petitioner erroneously alleges that Dilley’s second server instructs the edge server to retrieve content from the origin server, or that Dilley does not provide specifics on how that retrieval happens.

Even in combination, the prior art does not disclose limitation 1(d)/11(d) because Dilley’s second server does not perform the operations that Petitioner alleges. Petitioner and the Board improperly rely on Dr. Mowry’s testimony, which erroneously alleges that Dilley’s second server “instruct[s] the first server” to request content from an origin server if the content is not available at the first and second servers. EX1006, ¶251 (citing EX1003, 7:31-40, 8:53-58). To the contrary, Dilley’s disclosure explicitly teaches that it is the second server that requests the content from the origin server—it does not instruct the first server to do so. Specifically, Dilley states that the origin server is contacted “when a server in a parent region needs to fetch content” and that the “parent serves the request... by making a request to the origin server” *Id.*, 7:31-40, 8:53-58. There is no disclosure or suggestion in Dilley that a second server provides a redirect instruction to the first server to request content from the origin server directly. *Id.*; EX2016, ¶91.

Petitioner’s argument rests on the false premise that Dilley “does not provide specifics” on how content is retrieved from the origin server, and that Petitioner is therefore entitled to hypothesize whatever system it wants in order to fill that

silence. EX1006, ¶252. It is not proper do to so, especially where Dilley provides a clear, specific disclosure of how to retrieve content from an origin server. In other words, Dilley is not silent, and therefore Petitioner is not entitled to fill that silence with a POSA's alleged creativity (even if such an analysis were appropriate). EX2016, ¶92.

Specifically, Dilley's edge server evaluates object-specific metadata to determine whether to retrieve content from an origin server directly, or through a parent server, which fetches the object from the origin server. By ignoring this established mechanism, Petitioner attempts to create a "gap" in the prior art where none exists in order to justify the introduction of Wang's redirect instruction. EX1006, ¶252. Because Dilley already teaches a complete retrieval path, a POSA would have seen no reason to implement a redirect instruction to facilitate a process that Dilley's existing hierarchy already performs. EX2016, ¶92.

The Board's preliminary finding alleges that "Petitioner relies upon portions of Dilley that disclose that the edge server may contact the origin server for objects directly." DI, 25 (citing EX1003, 7:15-30). But the fact that the edge server may contact the origin server directly does not imply that it does so pursuant to a redirect instruction. Indeed, the disclosure of Dilley suggests that it is not. Dilley discloses an **autonomous default routing mechanism** allowing the edge server to retrieve the requested object from the origin server directly; bypassing parent

servers entirely. Section VIII.A.1; EX1003, 6:10-15, 8:34-42. If the parent server is bypassed, then it did not send a redirect instruction. Indeed, a POSA would have understood that a parent server *cannot* provide the claimed “redirect instruction” for such a process from which it is structurally excluded, as there is no communication loop between the edge server and the parent server in this scenario.

Thus, Petitioner’s conclusion that “a POSITA would have been motivated to implement this redirect instruction from Wang” is built upon mischaracterization of Dilley and must be rejected. EX2016, ¶93.

3. Dilley’s edge server, not the second server, determines where the edge server should fetch content based on the metadata, without receiving any redirect instruction from the parent server.

Contrary to Petitioner’s allegations, Dilley provides clear, specific disclosure of how to retrieve content from the origin server by using metadata associated with the content, and Dilley does not disclose or even suggest that the second server would instruct the first server to request content from the origin server because Dilley’s system uses metadata associated with the content for routing decisions. EX1003, 7:14-18; EX2016, ¶¶94-98.

Dilley discloses a tiered content distribution system using a cache hierarchy feature. EX1003, 5:31-36. Dilley explains, only “a relatively small, specific subset of objects” are serviced using the hierarchy, while most objects are fetched by the edge server contacting the origin server. EX1003, 6:10-14. More specifically,

Dilley's system uses "object-specific metadata" to handle a request for an object.

Id., 4:23-26. Dilley's metadata includes a "use-hierarchy" tag. EX1003, 6:30-61.

Dilley explains that the default setting of "OA" for the tag indicates that the edge server should "go direct to the (O)origin-server specified in the (A)RL" without using the cache hierarchy feature. *Id.* And setting the tag to be "coremap" turns on the cache hierarchy feature and instructs the edge server to do a DNS lookup for a parent server. *Id.*; EX2016, ¶95.

Additionally, Dilley discloses a "request routing mechanism" that is implemented as a software "running in each edge server" and includes "a guide process," and it is the guide process instructs the edge server where to get the content. EX1003, 7:46-8:2, 8:22-42. According to Dilley, the edge server ("*the first server*") determines whether it should fetch an object directly from the origin server or from a parent server by "quer[ying] its associated guide process."

EX1003, 8:22-38, FIG. 6. The "guide process" of the edge server "examine[s] object metadata for a given tag" and makes the determination to request the object from a parent server (if the object is a candidate to use the cache hierarchy feature), or to request the object from the origin server (if the object is not a candidate for distribution through the hierarchy). *Id.*, *see also* EX1003, 7:14-18 ("[A]pplication of metadata determines the name of the server region through which the request is

funneled to fetch or validate the object being requested. This is the origin server by default, but it can also be a cache hierarchy parent ...”). EX2016, ¶96.

Dilley also discloses that if the guide process instructs the edge server to request the object from a parent server, and the parent server does not have it, a guide process of the parent server instructs the parent server to “make[] a request to the origin server.” EX1003, 8:22-38, 8:53-58, FIG. 6. In other words, the parent server fetches the object from the origin server, instead of instructing the edge server to do so. EX2016, ¶97.

Accordingly, Dilley provides a clear, specific disclosure of how to retrieve content from the origin server based on metadata associated with the content, e.g., contacting the origin server directly if the requested object is not a candidate for using the cache hierarchy feature, otherwise, contacting a parent server in the cache hierarchy and the parent server fetches the content from the origin server if it does not have it. Dilley does not even suggest or imply that the parent server would instruct the edge server to retrieve the content directly from the origin server. Therefore, a POSA would not have been looking to other references, such as Wang, for a teaching of the claimed “*redirect instruction*.” EX2016, ¶98.

4. Petitioner’s combination of Dilley and Wang is improper.

The Petitioner’s proposed combination of Dilley and Wang is a result of hindsight-driven reconstruction that ignores the technical realities and express

objectives of the prior art. Because the combination results in a redundant architecture with no motivation to combine, it would not have been obvious to a POSA at the time of the invention. EX2016, ¶99.

a. Petitioner’s theory of motivation relies on improper hindsight reconstruction

The Petitioner’s assertion—that a POSA would find it “obvious” to bypass a “middleman” parent server to reduce latency—is a classic example of hindsight bias that “uses the patent in suit as a guide through the maze of prior art.” *See In re NTP, Inc.*, 654 F.3d 1279, 1299 (Fed. Cir. 2011); MPEP § 2143.01. Under the standards set forth in *KSR*, the PTAB must guard against “the distortion caused by hindsight bias” and must not rely on “conclusory statements” that a modification was a matter of simple common sense. *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398 (2007).

Petitioner’s theory of “fewer nodes equals less latency” is a simplistic generalization that fails to account for the specific technical context of the prior art. Pet., 51. Dilley teaches that its “middleman” parent server in the cache hierarchy “maintains a persistent connection to the origin server, which “reduces average latency and improves connectivity” for a “poorly-connected” edge server. EX1003, 5:10-12, 6:7-11; EX2016, ¶100.

Only by looking at the ’573 Patent’s teachings can Petitioner justify the counter-intuitive step of removing a persistent connection bridge. *See Mintz v.*



Dietz Watson, Inc., 679 F.3d 1372, 1379 (Fed. Cir. 2012). Because Petitioner offers no contemporaneous evidence that a POSA would have disregarded Dilley's express connectivity benefits in favor of Wang's redirect, the proposed combination is an improper hindsight reconstruction and must be rejected.

b. Petitioner fails to establish a motivation to combine because its purported benefit of reducing latency or improving server availability will not be achieved

The Petitioner's proposed combination of Dilley and Wang fails because it relies on "benefits" that are technically illusory and contrary to the express teachings of the prior art. EX2016, ¶102.

First, Petitioner contends that bypassing Dilley's parent server to request objects directly from the origin would necessarily "decrease latency," Pet. 50; EX1006, ¶¶255-257, this assertion is technically unsound.

As illustrated in the diagrams below, the proposed combination of Dilley and Wang requires **four discrete steps** to retrieve content from the origin server, which is **identical** to Dilley's standalone system.



EX2016, ¶104.

A POSA would have understood that removing Dilley's parent server as a "middleman" does not inherently decrease latency. In fact, as demonstrated below, integrating Wang's redirect procedure into Dilley's cache hierarchy can actually **increase** latency. *Id.*; EX2016, ¶105.

Dilley expressly teaches that its parent server is not a bottleneck, but a functional necessity that "reduces average latency" by providing a persistent connection for "poorly-connected" edge servers. Section VIII.E.4.a; EX1003, 6:7-11. A POSA would have understood that removing the parent server—and thus the "middleman" providing that persistent connection—would inevitably increase latency by forcing a direct, unstable connection between the poorly-connected edge server and the origin server. EX2016, ¶106.

Second, Petitioner contends that another benefit of combining Wang and Dilley would be improving parent server availability. Pet., 51; EX1006, ¶257. However, as established by Dr. Shenoy, a POSA would have known that a parent server at the time of the invention was inherently capable of servicing multiple concurrent requests. EX2016, ¶107. Because concurrent processing was already a baseline technical reality, the “benefit” alleged by Petitioner provides no actual improvement or motivation to modify the prior art. *Id.*

Furthermore, as previously established, Dilley’s parent server provides essential service for an edge server in a poorly connected region. A POSA would have understood that terminating such service, merely to accommodate other requests, would run counter to the express teachings and intent of Dilley. Thus, a POSA would have had no motivation to implement Wang’s redirect instruction for the sake of server availability. EX2016, ¶108.

Because the proposed combination would likely result in a degradation of system performance rather than an improvement, a POSA would not have expected success. EX2016, ¶109.

c. The proposed modification results in a superfluous and redundant architecture.

Dilley discloses two “request routing mechanism[s]”: (1) a default mechanism that retrieves the requested object from the origin server directly, and (2) a second mechanism that retrieves the requested object using the cache

hierarchy. Section VIII.A.1; EX1003, 6:10-15, 8:34-42. For the second mechanism, as already established in Section VIII.E.4.b, there is no reasonable expectation of success for the combination to improve performance. For the default routing mechanism, a POSA would have had no motivation to implement a modification that is entirely redundant of existing functionality in the prior art. *See In re Urbanski*, 809 F.3d 1237, 1243 (Fed. Cir. 2016) (noting that a POSA is not motivated to undertake modifications that increase complexity without a corresponding functional benefit). EX2016, ¶110.

Dilley already discloses a “guide process” that evaluates metadata to determine if an object request should bypass the parent server and the object be retrieved directly from the origin server. Section VIII.E.3; EX1003, 7:46-8:2, 9:7-13. Petitioner’s proposal to implement Wang’s redirection procedure in Dilley’s parent server creates a “superfluous” mechanism to achieve a bypass result that Dilley already performs at the edge server level. EX2016, ¶111.

A POSA would not have been motivated to add unnecessary logic to the parent server to handle requests that, by definition, would have already bypassed that parent server under Dilley’s existing “guide process.” EX2016, ¶112. Because Petitioner’s combination results in a “redundant and unnecessary implementation” that provides no actual improvement, Petitioner has failed to prove a rational reason to combine these specific references. *See KSR*, 550 U.S. at 418.

F. The references fail to teach or suggest that the second server is a peer server that sends a redirect instruction to the first server, as recited in dependent claims 5 and 15

Claim 5 depends on claim 1 and further recites that the “*second server*” (which provides the redirect instruction) is a “*peer server*” of the first tier. Petitioner only relies on Dilley in its arguments. Thus, the reference must teach or suggest a peer-to-peer redirect that instructs the first server *to obtain the resource from a content server of the content delivery network*. Claim 15 depends on claim 11 and recites similar limitations. EX2016, ¶¶113-117.

First, Petitioner relies on the same flawed arguments presented for claim limitation 1(d) to allege that a peer server sends the redirect instruction to the first server when the request resource is unavailable. Pet., 69. The combination of Dilley and Wang fails to teach or suggest that any “*second server*” – regardless a peer server or a parent server - provides a redirect instruction to the first server. *Supra* VIII.E; EX2016, ¶114.

Second, Petitioner’s reliance on Dilley is misplaced. Petitioner cites portions of Dilley to suggest that the second server can be a peer server or a parent server, and “[t]his second server will either serve the content from its cache or instruct the first server to request the content from a server higher on the hierarchy, such as an origin server.” Pet., 50 (citing EX1003, 7:31-40, 8:53-58). However, these citations exclusively describe actions performed by a **parent server**, not a peer server.

Indeed, Petitioner’s own argument admits that “[a] POSITA would have understood that to implement the described process, **when the second server is a parent server**, Dilley would send a redirect instruction to the edge server, directing the edge server to obtain the resource from the origin server instead.” Pet., 50. Dilley actually teaches that if “no peer copy of the object is available,” the edge server queries its **own** guide process to identify where to go get the object – it does not receive a redirect from a peer server. EX1003, 8:18-23, FIG. 6; EX2016, ¶115.

Third, Petitioner’s purported rationale for combining Wang and Dilley is logically inapplicable to peer servers. Petitioner argues that the combination removes **the parent server** as a “middleman” to decrease latency and increase availability of **the parent server**. Pet., 51 (citing EX1006, ¶¶255-257). These purported benefits are irrelevant to a peer server, which by definition, is not a “middleman” in Dilley’s hierarchy. EX2016, ¶116.

In sum, the prior art fails to teach “*the second server is a peer server*” that provides “*a redirect instruction*” to the first sever, and Petitioner fails to prove by a preponderance of evidence that claims 5 and 15 are unpatentable. EX2016, ¶117.

G. Petitioner’s inconsistent positions in district court preclude Petitioner from showing the “receiving an instruction” limitation of claim 3.

Petitioner has taken the position in the parallel district court litigation that the limitation “receiving an instruction to cache the resource at the first server of the first tier of servers when the resource is obtained from the content server of the content delivery network” is indefinite. EX2027, 3-4. Petitioner’s expert then renders his obviousness opinions for this limitation by “assuming” that the claim is not indefinite. EX1006 ¶¶288-298. The Office has found the pursuit of such inconsistent positions across forums to warrant denying a Petition. It should do so here as well.

“The Board’s claim construction rules are designed to ensure that the Board correctly construes claim terms and to minimize inconsistency in claim construction between forums.” *Revvo*, IPR2025-00632, Paper 20 at 4. A petitioner that advances different claim construction positions in different fora needs to explain why those different positions are warranted. *Id.* at 3-5. That a petitioner “cannot raise indefiniteness challenges in an *inter partes* review is not a sufficient explanation” because such a statement amounts to an assertion that a petitioner should be permitted to raise inconsistent invalidity challenges. *Tesla, Inc. v. Intellectual Ventures II LLC*, IPR2025-00340, Paper 18 at 3-4 (P.T.A.B. Nov. 5, 2025) (informative).

The impact on the Board when a petitioner takes inconsistent positions cannot be understated. As one panel explained the issue, “[i]f, according to Petitioner many limitations in [the claims] are indefinite, then how does Petitioner propose that the Board construe such limitations to assess Petitioner’s invalidity challenges?” *Am. Airlines, Inc. v. Intellectual Ventures I LLC*, IPR2025-01055, Paper 11 at 13 (P.T.A.B. Nov. 21, 2025).

Here, rather than providing an explicit construction or explain how the claim is to be interpreted Petitioner simply “assumes” that the claim is not indefinite and proceeds with the obviousness analysis. EX1006 ¶¶288-298. That approach is improper and should be rejected.

IX. CONCLUSION

Grounds 1 and 2 fail to show that any claim is unpatentable.

Respectfully submitted,
STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason Eisenberg, Reg. No. 43,447
Daniel Block, Reg. No. 68,395
Kristina Caggiano Kelly, Reg. No. 58,164
Xiaowei Jin, Reg. No. 81,041
Counsel for Patent Owner Sandpiper CDN, LLC

Date: March 3, 2026

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

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 11,637 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 Sandpiper CDN, LLC

Date: March 3, 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 RESPONSE** and associated Exhibits 2016-2027 were served in their entireties on March 3, 2026, upon the following parties via electronic mail:

Erika H. Arner (Lead Counsel) erika.arner@finnegan.com
Daniel C. Tucker (Back-Up Counsel) daniel.tucker@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
Anthony J. Berlenbach (Back-Up Counsel) anthony.berlenbach@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 Sandpiper CDN, LLC

Date: March 3, 2026

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