

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 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	2
III.	OVERVIEW OF THE '573 PATENT	6
IV.	PROSECUTION HISTORY	7
V.	PERSON OF ORDINARY SKILL IN THE ART (POSA)	10
VI.	CLAIM CONSTRUCTION	10
VII.	LEGAL STANDARDS	10
VIII.	REASONS TO DENY INSTITUTION.....	13
A.	Ground 1 fails because it relies on a baseless priority-date challenge.....	14
B.	Ground 2 (all independent claims): the combination of Dilley, Pai, and Wang fails to render obvious at least the feature of “determin[ing] a popularity designation associated with the requested resource.”	18
C.	Ground 2 (all independent claims) fails at least because there is no motivation to combine Dilley and Pai.	22
D.	Ground 2 (claims 1 and 11): the combination of Dilley, Pai, and Wang fails to render obvious at least the feature of “redirect[ing] instruction,” and even if taught, there is no motivation to combine Dilley and Wang.	24
E.	Ground 2 (claim 19): the combination of Dilley, Pai, and Wang fails to render obvious at least the feature of “requesting the resource from a second server of the content delivery network based at least on the popularity designation associated with the requested resource.”	28
IX.	CONCLUSION	29

TABLE OF AUTHORITIES

Cases

<i>Abbott Labs. v. Andrx Pharms., 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

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 two Grounds, each challenging all claims of the patent. Ground 1 presents a theory of anticipation, but relies on a reference that post-dates the priority date of U.S. Patent No. 10,924,573 (“the ’573 patent”). Because the earliest filed specification in the patent family provides adequate written-description support for the challenged claim language, the patent is entitled to its claimed earliest effective priority date and Ground 1 necessarily fails.

As a fallback position, Ground 2 cobbles together three prior-art references that predate the ’573 patent, but fail to disclose all of the limitations of the claims and are not obvious to combine. First, the proposed prior-art combination fails to show the key “popularity designation associated with the requested resource,” recited in all three independent claims, because that limitation is missing from Pai, the sole reference relied upon for that teaching.

Second, the proposed combination of Dilley, Pai, and Wang 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, “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, or “requesting the resource from a second server of the content delivery network based at least on the popularity designation associated with the requested resource,” as recited in independent claim 19.

Finally, Petitioner fails to present an adequate motivation to combine the three references as proposed—particularly combining Dilley and Pai to arrive at the “determin[ing] a popularity designation associated with the requested resource” limitation and combining Dilley and Wang to arrive at the “redirect instruction” limitation. Accordingly, Ground 2 suffers from at least two missing limitations from each of the three independent claims and a general lack of motivation to combine the references as proposed. The Petition therefore fails to show a likelihood of unpatentability for any claim. Institution should be denied.

II. BACKGROUND OF THE 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 was 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.

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 Patent Owner, Sandpiper CDN, LLC ("Sandpiper" or "PO"). Other businesses also entered the market 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

U.S. Patent No. 10,924,573 is entitled "Handling Long-Tail Content in a Content Delivery Network (CDN)," and it claims priority to U.S. Provisional Patent 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. If the first-tier server does not have the content, and the content is popular, 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 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 Pharms., 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 Preliminary Response, Patent Owner Sandpiper CDN, LLC agrees the *Phillips* standard applies. PO reserves the right to contest Petitioner’s implicit or explicit constructions during the proceeding, should the Board decide to institute review.

VII. LEGAL STANDARDS

In an IPR, “the burden of proving a proposition of unpatentability” is borne by (and remains at all times with) the petitioner. *In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1375 (Fed. Cir. 2016). To satisfy its burden, a petitioner must “articulate specific reasoning, based on evidence of record, to support the legal conclusion of obviousness”—“mere conclusory statements” will not suffice. *Id.* at 1380. Indeed, “[i]t is of the utmost importance that petitioners in the IPR proceedings adhere to the requirement that the initial petition identify ‘with particularity’ the ‘evidence that supports the grounds for the challenge to each

claim.”” *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1369 (Fed. Cir. 2016) (quoting 35 U.S.C. § 312(a)(3)).

In addition, 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. *Magnum Oil Tools*, 829 F.3d at 1381 (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).

Petitioner has the burden to establish that a patent lacks priority to an earlier application from which priority was claimed, and must identify any and all elements of the claims that the petitioner believes lack written description support in any and all earlier applications. *See MaxLite, Inc. v. Jiaxing Super Lighting Elec. Appliance Co., Ltd.*, IPR2020-00208, Paper 10 (P.T.A.B. June 24, 2020). Adequate written description does not require *in haec verba* support for the claimed subject matter at issue. *Purdue Pharma L.P. v. Faulding Inc.*, 230 F.3d 1320, 1323 (Fed. Cir. 2000). Because a patent specification is written for a POSA, and the POSA comes with the knowledge of what has come before, it is unnecessary to spell out every detail of the invention in the specification. *See LizardTech, Inc. v. Earth*

Resource Mapping, Inc., 424 F.3d 1336, 1345 (Fed. Cir. 2005). Only enough must be included to convince a POSA that the inventor possessed the invention and to enable such a person to make and use the invention without undue experimentation. *Id.* Furthermore, an expert’s conclusory testimony that ignores a disclosed embodiment is insufficient to show failure of written description. *See Koito Mfg. Co., Ltd. v. Turn-Key-Tech, LLC*, 381 F.3d 1142, 1155 (Fed. Cir. 2004).

“An obviousness determination requires finding that a [POSA] would have been motivated to combine or modify the teachings in the prior art *and* would have had a reasonable expectation of success in doing so.” *Regents of the Univ. of Cal. v. Broad Inst., Inc.*, 903 F.3d 1286, 1291 (Fed. Cir. 2018); *Intelligent Bio-Sys.*, 821 F.3d at 1367-68; *see also, e.g., NJOY, LLC v. JUUL Labs, Inc.*, IPR2024-00536, Paper 17 (P.T.A.B. Aug. 12, 2024) (denying institution where petitioner did not show a reasonable expectation of success in combining the asserted references); *Honeywell Int’l Inc. v. DSM IP Assets B.V.*, IPR2024-00493, Paper 7 (P.T.A.B. Aug. 21, 2024) (same).

“[O]bviousness concerns whether a skilled artisan not only could have made but would have been motivated to make the combinations or modifications of prior art to arrive at the claimed invention.” *Belden Inc. v. Berk-Tek LLC*, 805 F.3d 1064, 1073 (Fed. Cir. 2015). The fact that the combination might have been successful does not mean that a POSA would have been motivated to make the

combination. *See id.* It is not sufficient “for obviousness that a variation of the prior art would predictably work, but requires consideration of whether, in light of factors such as ‘design incentives and other market forces,’ the hypothetical skilled artisan would recognize the potential benefits and pursue the variation.” *Id.* at 1075 (citing *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 417 (2007)).

VIII. REASONS TO DENY INSTITUTION

The Petition fails to show a likelihood that any of the challenged claims is unpatentable. *First*, Ground 1 presents a theory of anticipation, but relies on a reference that postdates the priority date of the ’573 patent. Because the earliest filed specification in the patent family provides adequate written-description support for the challenged claim language, the patent is entitled to its claimed earliest effective priority date and Ground 1 fails. *Second*, while Ground 2 does not rely on a priority challenge, it suffers from multiple substantive deficiencies. Ground 2 cobbles together three prior-art references to present a theory of obviousness, but fails to show the key “popularity designation associated with the requested resource” recited in all three independent claims. Ground 2 also fails to identify the “redirect instruction” limitations of independent claims 1 and 11 or the “requesting a resource” limitation of independent claim 19 in the cited prior art. Finally, Petitioner fails to establish a motivation to combine the references as proposed. The Board should deny institution on the merits.

A. Ground 1 fails because it relies on a baseless priority-date challenge.

Petitioner first alleges that the effective priority date for the '573 patent is February 25, 2019, the filing date of the patent application that lead to the '573 patent. Petitioner argues that the patent is not entitled to its claimed priority date because new matter, i.e., the “instruction to not cache” limitation, was introduced in February 2019 and “has no support in any application to which the '573 Patent claims priority.” Pet., 8; EX2014, ¶50. The record clearly establishes the opposite. The limitation has written-description support throughout the priority chain. EX2014, ¶51.

The 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. For example, the '573 patent and the '412 application both disclose:

If the edge server 108 receives a redirect from the popularity service 102 without the “follow me” flag set (cases 1 and 2 above), it simply forwards the redirect to the client 104 (S4a, 522, 524). If the edge server 108 receives a “follow me” redirect, it obtains and caches the resource (at 526) and serves it to the client (at 528).

EX1001, 7:25-30; EX1011, ¶44. Accordingly, 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. For example, the '573 patent and the '412 application both disclose:

There are several reasons why a publisher/subscriber may not want the so-called “long-tail” content served from a caching CDN service, for example:

If cache fills are done from the subscriber's origin server and the subscriber pays for cache fill bandwidth, unnecessary fills for unpopular resources increase bandwidth costs without providing any value.

Serving an unpopular resource through a CDN cache adds latency (due to the cache fill), and risks forcing actually popular resources out of cache, potentially causing thrashing. The result is lower efficiency and the risk of degraded service.

For related or similar reasons, a CDN provider generally also does not want to serve long-tail content from an edge cache:

If cache fills are done from origin storage, unnecessary cache fills consume bandwidth both from the storage and cache systems, increasing (doubling) the internal cost to serve. This lowers efficiency, requiring relatively more bandwidth internally to serve the same content to the outside.

The second argument above also applies from the CDN's perspective: a CDN operator wants to minimize latency as well as the risk of thrashing in order to satisfy all of its subscribers.

EX1001, 10:17-41; EX1011, ¶¶68-69; EX2014, ¶54. Accordingly, 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) (“If alternative elements are positively recited in the specification, they may be explicitly excluded in the claims.”). 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). That satisfies the written description requirement for priority.

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.

B. Ground 2 (all independent claims): the combination of Dilley, Pai, and Wang fails to render obvious at least the feature of “determin[ing] a popularity designation associated with the requested resource.”

The proposed combination of Dilley, Pai, and Wang fails to render obvious any of the challenged independent claims at least because the combination fails to teach or suggest “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. Petitioner relies on Pai alone for this teaching but fails to point to any disclosure in Pai that satisfies the limitation. EX2014, ¶57.


Specifically, Petitioner alleges that Pai discloses “accessing a popularity service to determine the popularity associated with the requested resource.” Pet., 45 (citing EX1004, ¶¶39, 47). Petitioner mischaracterizes Pai. Pai only discloses determining popularity values “associated with *currently maintained media content*.” EX1004, ¶¶39, 47 (emphasis added). The popularity-value determination is not determined in association with “the requested resource,” as claimed, because the claimed requested resource is not Pai’s “currently maintained media content.” EX2014, ¶58.

In other words, the claim requires determining a “popularity designation,” which can then be used to “determine where [the requested resource] will be served from.” EX1001, 6:27-31; *see also, e.g.*, EX1001, 13:55-64. Particularly, the

popularity of the requested resource (e.g., the resource's "popularity count") is compared with different thresholds to determine where to serve the requested resource. EX1001, 6:53-7:24, FIG. 5. Pai, on the other hand, uses a currently known popular value to distribute or redistribute currently maintained media content within an on-demand media system, *before* any of the content is requested by a user. EX1004, ¶13; EX2014, ¶59.

Importantly, Pai's popular values of *currently maintained media content* are not used to determine the claimed where *the requested content* will be served from. EX1004, ¶13; EX2014, ¶59. Indeed, Petitioner agrees that Pai's system retrieves a user request resource "based on the content's current storage location," and *not* based on the content's popularity value. Pet., 46. Even though the content's current storage location may be based on the current popularity of the content, identifying the current storage location for retrieving the content is not a determining of the popularity designation. Accordingly, Pai fails to disclose "determin[ing] a popularity designation associated with the requested resource." EX2014, ¶60.

With regards to at least this distinguishing feature, the '573 patent discloses an exemplary embodiment with a mechanism of delivering content in a multi-tier content delivery network (CDN) with a popularity service. *See* EX1001, 5:27-43, FIG. 3. Figure 5 of the '573 patent shows an exemplary process of "*determin[ing] a popularity designation associated with the requested resource.*" EX2014, ¶61.



EX1001, FIG. 5 (annotated).

Therefore, contrary to Petitioner's argument, Pai does not disclose "determin[ing] a popularity designation associated with the requested resource." Rather, Pai discloses a mechanism of dynamically distributing/redistributing media content among multiple media sources within an on-demand media system based on the currently known popularity of the media content. EX1004, ¶¶13, 20, FIGs. 6, 7. For distribution or redistribution of media content, Pai discloses determining "current popularity values associated with currently maintained media content" and distributing or redistributing the media content to different media sources based on its currently known popularity value. EX1004, ¶¶39-51, FIGs. 6, 7. Pai discloses determining "current popularity values" associated with "currently maintained media content." *Id.* In other words, Pai checks popularity values of currently maintained media contents within an on-demand media system ("currently maintained media content") to determine which server to store each of the content. EX2014, ¶62.

Accordingly, Pai does not "determine a popularity designation associated with the requested resource." The proposed prior-art combination therefore fails to render obvious "accessing a popularity service associated with the content delivery network to determine a popularity designation associated with the requested resource." Ground 2 therefore fails.

C. Ground 2 (all independent claims) fails at least because there is no motivation to combine Dilley and Pai.

The Petition fails to establish a motivation to combine Dilley and Pai to arrive at the challenged claims. EX2014, ¶78. Petitioner alleges that a POSA would have found it obvious to combine Dilley with Pai. Pet., 37. Specifically, Petitioner argues that “Dilley determines where to cache content ‘as [a] function of metadata associated with the given object request.’” *Id.* (citing EX1003, 2:42-56). Petitioner then argues that “[a] POSITA would have understood that using ‘object specific metadata’ to facilitate control over a tiered distribution of content in a CDN could be done using content popularity,” and “a POSITA would have been motivated to implement Pai’s more specific teachings by providing dynamic popularity-based media distribution in Dilley.” Pet., 37-38. Again, Petitioner mischaracterizes Dilley and Pai, and a POSA would not have been motivated to combine Dilley with Pai. EX2014, ¶79.

First, contrary to Petitioner’s argument, Dilley discloses that “application of metadata 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 (emphasis added). In other words, Dilley’s metadata is used to determine where to “fetch” the request object, not to determine where to cache content. *Id.* Dilley does not “determine[] where to cache content ... [based on] metadata associated with the given object request,” as alleged by Petitioner. Pet., 37; EX2014, ¶80.

More specifically, Dilley discloses a “tiered distribution” architecture including “a set of edge servers and a hierarchical set of ‘parent’ or hub regions that service the edge servers.” EX1003, 4:63-5:1. Dilley explains that “only a relatively small, specific subset of objects are serviced using [the tiered distribution architecture]” and the subset of objects are selected using “object ‘metadata.’” *Id.*, 6:10-15. Dilley’s object metadata, or “metadata tag” includes a “specified domain name”—“a specified hostname or IP address.” *Id.*, 6:43-63. Pai, on the other hand, discloses a mechanism of dynamically distributing/redistributing media content among multiple media sources within an on-demand media system based on the “popularity of the media content.” EX1004, ¶13. Pai’s “popularity of the media content” is “determined [] by monitoring a number of received requests for the media content over time.” *Id.*; *see also id.*, ¶29 (discussing updating popularity values associated with the media content); EX2014, ¶81.

Accordingly, Dilley’s “object specific metadata” are domain names or IP addresses used to identify names of the servers/server region to “fetch” requested objects. EX1003, 6:43-63, 7:14-17. Dilley’s “object specific metadata” is completely different from Pai’s “popularity of the media content” (or “popularity values associated with the media content”), which is a counting number. EX1004, ¶¶13, 29. A POSA would not have been motivated to use Pai’s content popularity

as Dilley’s “object specific metadata” to “facilitate control over a tiered distribution of content,” as alleged by Petitioner. Pet., 37; EX2014, ¶82.

Additionally, Petitioner conflates Dilley’s “tiered distribution” of requested content with Pai’s popularity-based media content distribution/redistribution. Dilley’s “tiered distribution” is “content delivery to end users” from a CDN with tiered distribution hubs. EX1003, 5:1-5. Pai’s popularity-based media content distribution/redistribution distributes/redistributes media content based on popularity “within an on-demand media system,” e.g., *within* a CDN. EX1004, ¶¶13, 29, FIGs. 6, 7; EX2014, ¶83. A POSA would not have been motivated to “implement Pai’s more specific teachings by providing dynamic popularity-based media distribution in Dilley,” Pet., 38, because Pai’s popularity-based media distribution is different from Dilley’s tiered distribution. EX2014, ¶84.

D. Ground 2 (claims 1 and 11): the combination of Dilley, Pai, and Wang fails to render obvious at least the feature of “redirect[ing] instruction,” and even if taught, there is no motivation to combine Dilley and Wang.

The proposed combination of Dilley, Pai, 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. EX2014, ¶63.

Petitioner relies on the combination of Dilley and Wang for this teaching. Specifically, Petitioner alleges 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. Petitioner goes on to allege that “a POSITA would have been motivated to implement this redirect instruction from Wang.” *Id.*; EX2014, ¶64. However, Petitioner’s understanding of Dilley is misplaced and a POSA would not have been motivated to combine Dilley and Wang for disclosing the “redirect instruction,” as recited in independent claims 1 and 11. EX2014, ¶65.

First, contrary to Petitioner’s argument, Dilley does not disclose that the second server would “send a redirect instruction to the edge server, directing the edge server to obtain the resource from the origin server instead.” Pet., 50 (citing EX1003, 7:15-30); EX2014, ¶66. Dilley discloses that if “a given object request [to an edge server] cannot be serviced in the given edge region,” the request is provided to a parent region “as a function of metadata associated with the given object request.” EX1003, Abstract, 2:46-52. Dilley discloses that the metadata “determines” a server “through which the request is funneled to fetch or validate

the object being requested.” *Id.*, 7:15-18; EX2014, ¶67. In other words, the edge server (“*the first server*”) uses the object metadata to determine where to obtain the requested object (e.g., from a parent server). EX1003, 7:15-18; *see also id.*, 8:22-28 (“[T]he edge server queries its associated guide process to identify where to go to get the object. . . . [T]his determination is made by having the guide process examine object metadata for a given tag.”). Dilley further discloses that the parent server (“*the second server*”) serves the request either from its cache if it has it, “or by making a request to the origin server.” *Id.*, 8:53-55. In other words, Dilley’s parent server (“*the second server*”) does **not** send a redirect instruction to the edge server (“*the first server*”). Instead, Dilley’s parent server obtains the requested object from the origin server if it does not have it and serves the requested object to the edge server. *Id.*; EX2014, ¶68.

Second, a POSA would not have been motivated to combine Dilley and Wang. As discussed earlier, Dilley’s edge server (“*the first server*”) uses object metadata to identify a parent server (“*the second server*”), and the parent server obtains the object from the origin server if it does not have it, then serves the requested object to the edge server. EX1003, 8:53-55. Dilley’s parent server does **not** redirect the edge server to the origin server. Wang’s caching engine one (“*the first server*”), on the other hand, receives a redirect instruction to obtain the content from the origin server. EX1005, ¶123. Therefore, Dilley and Wang use two

different methods to achieve a same goal, e.g., to obtain requested content.

EX2014, ¶69.

A POSA would not have been motivated to replace Dilley’s method with Wang’s method. Petitioner alleges that combining Dilley and Wang “would remove [Dilley’s] parent server as a ‘middleman,’ and would allow [Dilley’s] edge server to communicate with and receive content directly from the origin server.” Pet., 51. However, this is squarely against the fundamental design philosophy of Dilley. EX2014, ¶70. Indeed, Dilley discloses that “the tiered distribution hubs act as accelerators for the end users, ensuring that the users obtain their desired content quickly and reliably, and the hubs **acts as buffers** for the origin site’s internal infrastructure, guaranteeing that **the site is not overwhelmed with requests** for popular content or large files.” EX1003, 5:16-21. Accordingly, Dilley’s tiered distribution hubs (e.g., the parent server) act exactly as a “middleman” (e.g., “buffers”) and Dilley does not want edge servers to overwhelm the origin server by limiting direct communication with the origin server. EX2014, ¶71. Therefore, a POSA would not have been motivated to combine Dilley with Wang. EX2014, ¶72.

E. Ground 2 (claim 19): the combination of Dilley, Pai, and Wang fails to render obvious at least the feature of “requesting the resource from a second server of the content delivery network based at least on the popularity designation associated with the requested resource.”

The proposed combination of Dilley, Pai, and Wang fails to render obvious independent claim 19 at least because the combination fails to teach or suggest “requesting the resource from a second server of the content delivery network based at least on the popularity designation associated with the requested resource,” as recited in independent claim 19. EX2014, ¶73. Petitioner relies on Dilley and Pai for this teaching but fails to point to any disclosure in Dilley or Pai that satisfies the limitation. EX2014, ¶74.

Specifically, Petitioner alleges that a POSA would have understood that “based on an unpopular designation, Dilley would retrieve content from the second server.” Pet., 84. Petitioner argues that in Dilley, “the servers in the tiered distribution hubs [(i.e., the parent or origin servers)] will provide the edge servers with the requested content’ when the ‘edge servers do not have the appropriate file when it is requested.’” Pet., 83 (citing EX1003, 5:6-10). Indeed, when Dilley’s edge server does not have the requested object, it contacts “either a core region or a dedicated parent region ... as a function of given object service metadata.” EX1003, 7:31-35. Dilley explains the metadata associated with the requested object is used to “determine[] the name of the server region through which the

request is funneled to fetch or validate the object being requested.” *Id.*, 7:14-17.

Dilley’s object-specific metadata is not “the popularity designation associated with the requested resource.” Accordingly, Dilley fails to disclose “requesting the resource from a second server of the content delivery network based at least on the popularity designation associated with the requested resource.” EX2014, ¶75.

Pai also fails to disclose the same feature. Pai only discloses a mechanism of dynamically distributing/redistributing media content among multiple media sources within an on-demand media system based on “current popularity values associated with currently maintained media content.” EX1004, ¶¶13, 20, 39, FIGs. 6, 7. In other words, Pai’s distribution and redistribution of “currently maintained media content” are within the CDN. Pai does not disclose “requesting the resource from a second server of the content delivery network,” let alone a request “based at least on the popularity designation associated with the requested resource.” EX2014, ¶76. Accordingly, the proposed prior-art combination fails to render obvious “requesting the resource from a second server of the content delivery network based at least on the popularity designation associated with the requested resource.” EX2014, ¶77.

IX. CONCLUSION

The Board should deny institution of *inter partes* review of U.S. Patent No. 10,924,573 on the merits.

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

Date: September 10, 2025

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 6,081 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 10, 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 10, 2025, 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

Date: September 10, 2025

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