

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

MICROSOFT CORPORATION,
Petitioner

v.

SANDPIPER CDN, LLC,
Patent Owner

Case IPR2026-00174
U.S. Patent No. 9,762,692

**PATENT OWNER PRELIMINARY RESPONSE
UNDER 37 C.F.R. § 42.107(a)**

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

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	BACKGROUND OF THE TECHNOLOGY.....	3
III.	OVERVIEW OF THE '692 PATENT	8
IV.	PROSECUTION HISTORY	10
V.	OVERVIEW OF THE CITED PRIOR ART	11
A.	Seed	11
B.	Swildens	13
VI.	PERSON OF ORDINARY SKILL IN THE ART	15
VII.	CLAIM CONSTRUCTION	15
VIII.	LEGAL STANDARDS	16
IX.	ARGUMENT AGAINST INSTITUTION.....	17
A.	Independent Claims 1, 10, 15: Seed and Swildens, in any order and for both Grounds, do not teach the “first portion of the content provider’s library is distinct from the second portion,” as recited in claim 1 and similarly recited in claims 10 and 15.	19
B.	Independent Claim 1	22
1.	The Petitioner fails to show Seed-Swildens or Swildens-Seed renders obvious the claimed “ <i>if the resource is determined not to be popular, directing the client to a second server in a second tier of servers distinct from the first tier of servers.</i> ”	22
2.	The Petitioner fails to show either combination teaches the claimed “ <i>if the resource is determined to be popular, then the first server obtaining the resource and the first server serving the resource to the client.</i> ”	25
a.	Ground 1	26
b.	Ground 2	30
C.	Independent Claim 10	32
1.	Petitioner fails to show Seed-Swildens or Swildens-Seed renders obvious the claimed “ <i>provide the edge servers with a second set of information about which tier of servers should handle those requests for content.</i> ”	32

2.	Petitioner fails to show Seed-Swildens or Swildens-Seed renders obvious the claimed “ <i>determin[ing a tier of servers to handle a request] based, at least in part, on a measure of popularity of requested content.</i> ”.....	35
D.	Independent Claim 15: Seed and Swildens, in any order and for both Grounds, do not teach the claimed “ <i>redirecting being based, at least in part, on a current popularity value for the resource.</i> ”.....	36
E.	All Grounds: Petitioner improperly relies on declarant testimony to fill gaps in the prior art.....	38
X.	CHEVRON AND DEEPER FAVOR DENYING INSTITUTION.	40
XI.	CONCLUSION.....	41

PATENT OWNER'S UPDATED EXHIBIT LIST

Exhibit No.	Description
2001	First Amended Docket Control Order, <i>Sandpiper CDN, LLC v. Microsoft Corporation</i> , Case No. 2:25-cv-00664-JRG-RSP (E.D. Tex.), October 17, 2025
2002	Exhibit E101 to Microsoft's Invalidity Contentions (Subject Matter Ineligibility Contentions for U.S. Patent No. 9,762,692), <i>Sandpiper CDN, LLC v. Microsoft Corporation</i> , Case No. 2:25-cv-00664-JRG-RSP (E.D. Tex.), December 5, 2025
2003	"Enforcement and Non-Waiver of 37 C.F.R. § 42.104(b)(4) and Permissible Uses of General Knowledge in <i>Inter Partes</i> Reviews," United States Patent and Trademark Memorandum, July 31, 2025
2004	Success Rate of Motions to Stay Pending <i>Inter Partes</i> Review before Judge Rodney Gilstrap, Docket Navigator Report generated February 4, 2026
2005	Complaint for Patent Infringement, <i>Sandpiper CDN, LLC v. Microsoft Corporation</i> , Case No. 2:25-cv-00664-JRG-RSP (E.D. Tex.), June 26, 2025
2006	Microsoft's Invalidity Contentions, <i>Sandpiper CDN, LLC v. Microsoft Corporation</i> , Case No. 2:25-cv-00664-JRG-RSP (E.D. Tex.), December 5, 2025
2007	Declaration of Dr. Prashant Shenoy
2008	<i>Curriculum Vitae</i> of Dr. Prashant Shenoy
2009	Markman Order, <i>Sandpiper CDN, LLC v. Comcast Cable Comm 'ns, LLC</i> , Case No. 2:24-cv-00886-JRG-RSP (E.D. Tex.), February 18, 2026
2010	Infringement Contentions, <i>Sandpiper CDN, LLC v. Microsoft Corp.</i> , Case No. 2:25-cv-00664 (E.D. Tex. Sept. 16, 2025)

Exhibit No.	Description
2011	“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)
2012	“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)
2013	Entrepreneur of the Year, June 19, 2000
2014	“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)
2015	“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)
2016	“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
2017	“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)
2018	“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)

Exhibit No.	Description
2019	U.S. Patent Application Publication No. 2009/0252071 A1 ("Nurminen")
2020	U.S. Patent 7,937,477 B1 ("Day")

I. INTRODUCTION

The Director should deny institution of *inter partes* review because Petitioner has not met its burden to show a reasonable likelihood that any of the challenged claims 1, 2, 4-15, 17, and 18 (the “challenged claims”) of U.S. Patent No. 9,762,692 (the “’692 patent”) are unpatentable.

The Petition includes fatal legal and technical flaws stemming from Petitioner’s misapplication of *Seed* and *Swildens*, in any order, to the challenged claims. Petitioner significantly misconstrues the technical disclosures of these references and improperly relies on conclusory attorney and declarant statements to gap-fill the deficiencies of the applied art. Indeed, Petitioner fails to show *Seed-Swildens* (Ground 1) or *Swildens-Seed* (Ground 2) renders obvious numerous limitations of independent claims 1, 10, and 15, including, for example:

All independent claims:

- “*the second server comprises a first portion of the content provider’s library ... wherein at least one other server in the second tier comprises a second portion ... the first portion of the content provider’s library is distinct from the second portion*” (claim 1).
- “*distinct portions of the content provider’s library are logically partitioned across parent servers in the second tier*” (claim 10).
- “*the second tier server comprises a first portion of the content provider’s library ... wherein at least one other second tier server comprises a second portion of the content provider’s library distinct from the first portion*” (claim 15).

Independent claim 1:

- *“if the resource is determined to be popular, then the first server obtaining the resource and the first server serving the resource to the client.”*
- *“if the resource is determined not to be popular, directing the client to a second server in a second tier of servers distinct from the first tier of servers.”*

Independent claim 10:

- *“at least one popularity service constructed and adapted to obtain a first set of information from at least some of the edge servers about requests for content, and to provide the edge servers with a second set of information about which tier of servers should handle those requests for content.”*
- *“the popularity service making its determinations based, at least in part, on a measure of popularity of requested content.”*

Independent claim 15:

- *“selectively redirecting the request from the client to a second tier server, the second tier server being in a tier distinct from the tier containing the first server, the redirecting being based, at least in part, on a current popularity value for the resource.”*

Additionally, for all Grounds and throughout the Petition, Petitioner heavily relies on conclusory, unsupported statements from its declarant to satisfy claim terms not found in the applied references. Petitioner’s reliance on expert testimony to meet the claim terms, rather than the printed publications’ teachings asserted in its challenges, violates 35 U.S.C. § 311(b) and 37 C.F.R. § 42.104 (b)(4) (“Rule 104(b)(4)”). The Office has made clear that it will no longer waive the requirements of Rule 104(b)(4) that require Petitioner to “specify where each element of the claim is found in the prior art patents or printed publications relied upon,” which means “expert testimony ... may not be used to supply a missing

claim limitation.” EX2003, 1. Petitioner’s disregard for its statutory and regulatory obligations should not stand.

For at least these reasons, Petitioner has not met its burden to show a reasonable likelihood that the challenged claims are unpatentable. Accordingly, the Director should deny institution.

II. BACKGROUND OF THE TECHNOLOGY

The ’692 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. EX2007, ¶35.

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. EX2007, ¶36.

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. EX2007, ¶37.

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. EX2007, ¶38.

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. EX2007, ¶39.

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 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. EX2007, ¶40.

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. EX2011; EX2012. 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. EX2013. 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. EX2014; EX2015; EX2016.


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. EX2017. By 2019, CenturyLink offered CDN services with peak capacity exceeding 120 Tbps and served nine of the world’s ten largest media companies. EX2018. As the market expanded, other businesses entered the space—often by misappropriating Sandpiper’s innovations and directly competing with Sandpiper for market share. EX2007, ¶41.

One such company is Microsoft. Microsoft built a CDN, which uses technology described and claimed by Sandpiper’s patents. For example, Microsoft’s Azure CDN supports Xbox gaming, Office 365, Windows updates, and web resources like LinkedIn and Bing. *See* EX2010. Microsoft did not license this patented technology from Sandpiper, Level 3, and/or its predecessors. Accordingly, Sandpiper brought suit against Microsoft in district court for patent infringement. Microsoft filed IPR petitions against all of Sandpiper’s asserted patents in response, including by copying petitions previously filed by Google. EX2007, ¶42.

III. OVERVIEW OF THE '692 PATENT

The '692 patent presents a novel and innovative technique for serving content based on popularity levels. *See, e.g.*, EX1001, 6:41-7:10; EX2007, ¶43. As shown below, a CDN may include a plurality of tiers, such as a tier of edge servers, a tier of origin servers, and a tier of parent servers. EX1001, 2:52-66, 4:5-8, FIG. 1; EX2007, ¶43.



EX1001, FIG. 1.

When a client requests an object, that request is first directed to an edge server. EX1001, 5:45-48; EX2007, ¶44. The edge server checks if the requested object is present and, if so, serves the object. EX1001, 5:66-6:3; EX2007, ¶44. If

the requested object is not available, the CDN returns one of three possible responses to the edge server. EX1001, 6:41-50; *see also* EX1001, FIG. 5; EX2007, ¶44. First, if the object has not reached a first/minimal level of popularity, the popularity server instructs the edge server to redirect the request to an origin server. EX1001, 6:51-55; EX2007, ¶44. Second, if the object's popularity has exceeded the first/minimal level but has not yet exceeded a second, mid-tier threshold, the popularity server instructs the edge server to redirect the request to a parent server. EX1001, 6:56-61; EX2007, ¶44. Third, if the object's popularity has exceeded the mid-tier threshold, the popularity server instructs the edge server to serve the content itself. EX1001, 7:4-7; EX2007, ¶44.

Furthermore, when redirecting requests to parent servers, a determination of which particular parent server has the requested object is made. EX1001, 6:61-64; EX2007, ¶45. "Such a determination would be based on how the distinct subsets of a content provider's library are partitioned across parent servers in the same intermediate tier." EX1001, 6:64-67; EX2007, ¶45. For example, the '692 patent explains that when "a content provider's library comprises multiple content-types (e.g., movies, video games, music, software patches, etc.)," "it would be advantageous to store a proportionate amount of each content type on each participating intermediate tier server (e.g., 60% movies, 30% video games, 10% software downloads on each participating server in the tier)." EX1001, 12:29-36;

EX2007, ¶45. “In such a scenario, if the demand for a particular content-type increases during a given time period (and/or demand for a different content-type declines), then such fluctuations in demand will generally affect each participating server proportionately.” EX1001, 12:36-40; EX2007, ¶45.

IV. PROSECUTION HISTORY

During prosecution, the '692 patent received one non-final Office Action and one final Office Action. EX1002, 240-255, 2111-2127. In both Office Actions, the Examiner relied on Seed—also relied on by Petitioner in both Grounds 1 and 2 of this proceeding—for obviousness rejections. EX1002, 246-254, 2116-2125. The Examiner acknowledges in both Office Actions that Seed does not teach “at least one other second tier server comprises a second portion of the content provider’s library distinct from the first portion,” as recited in claim 17 (issued as independent claim 15) and similarly recited in issued independent claims 1 and 10. EX1002, 252, 2122-2123. The Applicant appealed the Examiner’s rejections in the final Office Action. EX1002, 2156.

On appeal, the Board reversed the Examiner’s rejection of independent claim 12 (issued as independent claim 10), independent claim 17 (issued as independent claim 15), and other dependent claims. EX1002, 2331-2334, 2339-2342. The Board found that the applied references, including Seed, do not teach that “the first portion of the content provider’s library is distinct from the second

portion,” as recited in dependent claim 2 and later incorporated into issued independent claim 1. EX1002, 2338-2339.

Upon return of the application to the Examiner, the Examiner conducted an interview with the Applicant and entered an Examiner’s amendment. EX1002, 2363-2367, 2377. The Examiner found the claims were allowable over the prior art and, accordingly, entered a Notice of Allowance. EX1002, 2361-2369.

V. OVERVIEW OF THE CITED PRIOR ART

A. Seed

Seed describes a system for “managed object replication and delivery over a network.” EX1005, ¶1; EX2007, ¶¶46-47. Seed explains that “[a] typical content delivery network (CDN) operator deploys one or more parent” and “one or more edge servers at the edge of the network to facilitate more cost-effective and efficient delivery of such objects to an end-user (client).” EX1005, ¶10; EX2007, ¶47. As shown below, “the network also optionally has access to one or more origin server sites 110” that “are typically owned and/or maintained by the network provider’s customers for storing and serving one or more objects.” EX1005, ¶20, FIG. 2; EX2007, ¶47.




Fig. 2

EX1005, FIG. 2

Seed describes that it “is typically not feasible to store all objects on the edge servers” because it can be cost-prohibitive and may not even be possible. EX1005, ¶11; EX2007, ¶48. One solution is to pre-populate edge servers with objects that are likely significant and in high demand. EX1005, ¶12; EX2007, ¶48. “However, it is difficult to predict popularity and difficult to manage pre-populating.” EX1005, ¶12; EX2007, ¶48. “Other solutions fetch objects on demand”—“when [the] requested object is not available on a handling edge server,” the handling edge server fetches the requested object from a parent server. EX1005, ¶13; EX2007, ¶48. Seed describes that doing so can use too much

bandwidth and create congestion. EX1005, ¶14; EX2007, ¶48.

In view of the issues, Seed's edge server redirects the client's request to a parent server when the requested object is not available. EX1005, ¶¶17-18, 22, 27; EX2007, ¶49. If the requested object is popular, it is replicated from a parent server to the edge server to be served in the future. EX1005, ¶¶17-18, 26-27; EX2007, ¶49.

B. Swildens

Swildens describes a system for “the automatic migration of data via a distributed computer network.” EX1006, ¶18; EX2007, ¶¶46, 50. Swildens' system includes edge servers, an origin server, and a staging site. EX1006, ¶¶42, 46, 53-54, FIG. 1; EX2007, ¶50. Swildens' system is depicted below in Figure 1.




Fig. 1

EX1006, FIG. 1.

Swildens’ “edge server receives a user’s request for content.” EX1006, ¶46; EX2007, ¶51. If the edge server does not have the content, it redirects the user to an origin server. EX1006, ¶46; EX2007, ¶51. If the origin server does not have the content, it redirects the user to a staging server. EX1006, ¶47; EX2007, ¶51. Swildens describes that servers “should typically have at least the combined storage and delivery capacity to store one copy of all of the content files.” EX1006, ¶58; EX2007, ¶51.

Swildens' system also includes an auto migration server that periodically (e.g., every hour) collects information from the edge servers to determine popularity levels of content. EX1006, ¶¶126-127; EX2007, ¶52. The auto migration server then selects the few most popular content items for auto migration to the edge servers. EX1006, ¶¶126-127; EX2007, ¶52. In addition, “[s]ingle edge serve[r]s can perform their own periodic evaluation of popular files.” EX1006, ¶128; EX2007, ¶52.

VI. PERSON OF ORDINARY SKILL IN THE ART

At this stage of the proceeding, solely for purposes of this Preliminary Response, Patent Owner does not dispute Petitioner’s level of skill for a person of ordinary skill in the art (POSITA). Pet., 9. Patent Owner reserves the right to contest Petitioner’s level of skill for a POSITA should trial be instituted. EX2007, ¶34.

VII. CLAIM CONSTRUCTION

In an *inter partes* review, claims are “construed using the same claim construction standard that would be used to construe the claim in a civil action under 35 U.S.C. 282(b).” 37 C.F.R. § 42.100(b). All claim terms must be given their ordinary and customary meaning as understood by a POSITA at the time of the invention, in light of the specification and prosecution history of the patent. *See Phillips v. AWH Corp.*, 415 F.3d 1303, 1312-13 (Fed. Cir. 2005) (en banc).

Petitioner asserts that no express constructions are necessary and it has “applied the claim terms’ ordinary meanings as understood by a POSITA.” Pet., 10. Petitioner further notes prior constructions made by the Board for independent claims 1 and 15 and that there were constructions made by a third party (Comcast) in related district court litigation. Pet., 10-11; *see also* EX2009, 7-21 (district court’s Markman order ruling on the disputed claim terms).

For purposes of institution, none of the arguments below relies on the construction of any specific claim term. Thus, no express construction is required for the Director to deny institution. *See Realtime Data, LLC v. Iancu*, 912 F.3d 1368, 1375 (Fed. Cir. 2019) (“The Board is required to construe ‘only those terms ... that are in controversy, and only to the extent necessary to resolve the controversy.’”) (quoting *Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999)). Patent Owner reserves the right to contest Petitioner’s implicit or explicit constructions should trial be instituted.

VIII. LEGAL STANDARDS

The Petition must identify, with particularity, the grounds and evidence that support finding the claims unpatentable. 35 U.S.C. § 312(a)(3). Petitioner must “specify where each element of the claim is found in the [relied upon] prior art patents.” 37 C.F.R. § 42.104(b)(4)). To comply with Rule 104(b)(4), the Petition may not rely on expert testimony, common sense, or other “general knowledge”

evidence to supply a missing claim limitation. EX2003.

The determination regarding institution must be based on what the Petition actually presents and not what it could have reasonably presented had it been drafted differently. *In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1381 (Fed. Cir. 2016) (holding that the Board is not “free to adopt arguments on behalf of petitioners that could have been, but were not, raised by the petitioner during [the post-grant proceeding]”).

IX. ARGUMENT AGAINST INSTITUTION

The Petition has numerous flaws that each warrant denying institution. Both of the obviousness challenges contain key deficiencies in the obviousness analysis of independent claims 1, 10, and 15, and fail to establish that either proposed combination (Seed-Swildens or Swildens-Seed) teaches:

All claims: “*wherein the second server comprises a first portion of the content provider’s library ... wherein at least one other server in the second tier comprises a second portion of the content provider’s library, wherein the first portion of the content provider’s library is distinct from the second portion*” (element 1[d][ii], and similarly recited in elements 10[d] and 15[c]).

Claim 1: (a) *“if the resource is determined to be popular, then the first server obtaining the resource and the first server serving the resource to the client”* (element 1[c]) and (b) *“if the resource is determined not to be popular, directing the client to a second server in a second tier of servers distinct from the first tier of servers”* (element 1[d][i]).

Claim 10: (a) *“provid[ing] the edge servers with a second set of information about which tier of servers should handle those requests for content”* (element 10[c]) and (b) *“the popularity service making its determinations [of which tier of servers should handle those requests for content] based, at least in part, on a measure of popularity of requested content”* (element 10[c]).

Claim 15: *“selectively redirecting the request from the client to a second tier server ... based, at least in part, on a current popularity value for the resource”* (element 15[b]). EX2007, ¶¶53-55.

Further, in an attempt to remedy technical deficiencies with its references’ teachings and obviousness challenges, Petitioner turns to conclusory statements from its declarant, unsupported by the references themselves, which violates Rule 104(b) and is fatal to both of the Petition’s obviousness theories.

A. Independent Claims 1, 10, 15: Seed and Swildens, in any order and for both Grounds, do not teach the “first portion of the content provider’s library is distinct from the second portion,” as recited in claim 1 and similarly recited in claims 10 and 15.

Independent Claim 1 recites that “*the second server comprises a first portion of the content provider’s library ... wherein at least one other server in the second tier comprises a second portion ... the first portion of the content provider’s library is distinct from the second portion.*” Independent claim 10 recites that “*distinct portions of the content provider’s library are logically partitioned across parent servers in the second tier.*” Independent claim 15 recites that “*the second tier server comprises a first portion of the content provider’s library ... wherein at least one other second tier server comprises a second portion of the content provider’s library distinct from the first portion.*” For all three claims, Petitioner relies on its arguments for claim element 1[d][ii]. See Pet., 25-28, 37, 44-45, 58-59, 69-70, 77. However, Petitioner’s arguments regarding the claimed *distinct portions* are fundamentally flawed. EX2007, ¶56.

First, Petitioner argues that Seed discloses this feature. Pet., 25. This argument is defective for several reasons. EX2007, ¶57. To begin with, both the Examiner and the Board previously found that Seed does not disclose this feature. EX1002, 252, 2123, 2338-2342. Petitioner has not explained why the Examiner and Board were incorrect—nor can it because Seed simply does not disclose the claimed feature. EX2007, ¶57. Petitioner fails to identify any portion of Seed that

discloses partitioning of content amongst parent servers as claimed. Petitioner's theory solely relies on paragraph 29 of Seed. Pet., 25. But the cited portion of Seed discusses how an individual parent server handles a request when it does not have the requested object—it says nothing about how content is distributed. EX1005, ¶29; EX2007, ¶58. Indeed, Petitioner acknowledges that Seed “does not expressly disclose this limitation.” Pet., 25. Facing the absence of disclosure, Petitioner provides a conclusory sentence, relying improperly on expert testimony, that Seed *suggests* its servers have distinct content portions. Petitioner wrongly states that “Seed recognizes that not every parent-level server has the same objects.” Pet., 25. It does not. EX2007, ¶58. Mr. Lanning also acknowledges that “Seed does not expressly specify the extent of any overlap between parent servers.” EX1003, ¶124.

Faced with this gap, Petitioner argues that this limitation “would have been obvious in view of Swildens.” Pet., 25-26, 58-59. This argument hinges again on improper expert testimony and an assumption that Swildens' hash function would uniquely identify a single server. *See* Pet., 25-26, 58-59 (“[A] POSITA would have understood that a hash function would typically result in each server having distinct, non-overlapping portions of the library, because each content identifier would be assigned to one server based on its hash value.”). Mr. Lanning similarly argues that “running a content identifier (key) through a hash function would

typically produce only one output, which a POSITA would have understood would be the server to which it is assigned (e.g., a one-to-one key-to-slot mapping).”

EX1003, ¶134. However, Petitioner and Mr. Lanning misread Swildens because these assumptions contradict Swildens’ disclosure. EX2007, ¶59.

Swildens describes that “[s]ervers should typically have at least the combined storage and delivery capacity to store one copy of **all** of the content files.” EX1006, ¶58 (emphasis added). Swildens further describes that a group of LAN-connected servers can be configured to “access each other’s local storage devices.” EX1006, ¶¶88-89. In other words, the servers share common storage and all the servers have access to the *same* content. EX2007, ¶¶59-60. Swildens explains that using “a hash included in the URL (e.g., in the host name or first component of the path),” an origin server is able to “determine which **other of the LAN-connected origin servers in its group** should hold the content on its local storage device.” EX1006, ¶60 (emphasis added).

Therefore, Swildens does not disclose, teach or suggest the claimed “first portion of the content provider’s library is distinct from the second portion,” but rather describes the content may be stored on multiple local storage devices in the LAN (not a single server as Petitioner alleges). EX2007, ¶61. Swildens describes that the origin server to which the content request was redirected may then use the hash function to determine which local storage devices the origin server itself can

access to serve the content (i.e., it does not redirect the request because *all* of the LAN servers have access to the *same* content). EX1006, ¶¶60, 88-89; EX2007, ¶61. Consistent with Swilden’s requirement, each origin server in the LAN has “the combined storage and delivery capacity” to serve “all of the content files and satisfy the peak delivery load.” EX1006, ¶58; EX2007, ¶61.

Accordingly, Petitioner fails to show that Seeds or Swildens, either separately or in combination, teach the claimed “first portion of the content provider’s library is distinct from the second portion,” recited in claims 1, 10, and 15. EX2007, ¶62.

B. Independent Claim 1

There are two additional failures of proof for claim 1. EX2007, ¶63.

- 1. The Petitioner fails to show Seed-Swildens or Swildens-Seed renders obvious the claimed “*if the resource is determined not to be popular, directing the client to a second server in a second tier of servers distinct from the first tier of servers.*”**

Independent claim 1 recites “*if the resource is determined not to be popular, directing the client to a second server in a second tier of servers distinct from the first tier of servers.*” Thus, when the resource is unpopular, the first server redirects the request to another server. Petitioner fails to show that either Ground teaches this element. EX2007, ¶64.

In **Ground 1**, Petitioner relies on Seed for this element. Pet., 23-24. Seed explicitly discusses that the edge server redirects the request when the requested object is not available. EX1005, ¶¶17-18, 27; EX2007, ¶65. Thus, contrary to Petitioner's argument, Seed's redirecting is based on availability, not popularity. Seed states that it does not include pre-populating, even with content likely to be popular. EX1005, ¶12; EX2007, ¶65. Furthermore, Seed explicitly explains that an unavailable object can be either popular or unpopular. For example, Seed explains that "[i]f the selected edge server does not have the requested object, a check is initiated (at 315) for the edge server to determine whether the requested object is popular and if so, to replicate the popular requested object to the edge server." EX1005, ¶26. The fact that Seed checks to determine whether the unavailable object is popular indicates that both popular and unpopular objects may be unavailable based on the specific disclosure in Seed. EX2007, ¶65. This point is further illustrated in Figure 3(b), reproduced below, which illustrates the steps employed (at 315) to determine whether the requested object is popular. EX1005, ¶26, FIG. 3B; EX2007, ¶65.




Fig. 3(b)

EX1005, FIG. 3B.

As shown in Figure 3B, check step 350 has a “Yes” outcome branch indicating that the requested object is unavailable and popular. EX2007, ¶66. Thus, contrary to the Petition’s assumption, when the requested object is unavailable, it does not need to be unpopular. EX2007, ¶66. Because the Petition’s entire theory is based on this flawed assumption, the obviousness challenge based on Seed-Swildens fails. EX2007, ¶66.

For **Ground 2**, Petitioner asserts that Swildens discloses this limitation for the same flawed assumption relied on in Ground 1. Pet., 54-55 (“[W]hen the edge server does not have the requested [object]—because the object either is not popular, or the edge server has not yet determined that the [object] is popular”) Thus, Petitioner again incorrectly assumes that unavailable content must be unpopular. Just like Seed, Swildens also explains that the edge server redirects the request when the requested content is unavailable and that the redirection is not based on whether the content is popular. EX1006, ¶68; EX2007, ¶67.

Because the Petition’s entire theory is based on this flawed assumption, Petitioner also fails to show that Swildens teaches directing the client to a second server based on a determination that the resource is not popular. EX2007, ¶67.

2. The Petitioner fails to show either combination teaches the claimed “*if the resource is determined to be popular, then the first server obtaining the resource and the first server serving the resource to the client.*”

Independent claim 1 recites “*if the resource is determined to be popular, then the first server obtaining the resource and the first server serving the resource to the client.*” Thus, when the resource is popular, the first server takes two actions: “*obtaining*” and “*serving.*” Petitioner fails to show that either Seed-Swildens or Swildens-Seed renders obvious this feature. EX2007, ¶68.

a. Ground 1

Petitioner makes several arguments in Ground 1 based on Seed supporting when the resource is popular, the first server takes two actions: “*obtaining*” and “*serving*.” Pet., 18-23. All fail. EX2007, ¶69.

First, Petitioner’s purported theory improperly separates the claim elements and relies on two separate actions: (1) in response to a first request for a popular but not-yet-replicated object, obtain the content, and (2) in response to a subsequent—and different—request after replication, serving the content. Pet., 19-20. However, these two actions are mutually exclusive because a request is either a first request or a subsequent request, but it cannot be both. EX2007, ¶70. Thus, under Petitioner’s theory, the edge server either performs “*obtaining*” in response to the first request or performs “*serving*” in response to the subsequent request, not both. EX2007, ¶70. Petitioner’s theory fails to explain how Seed’s edge server discloses both “*obtaining*” and “*serving*” if the resource is determined to be popular. EX2007, ¶70. In both of Petitioner’s scenarios Seed’s edge server performs only “*serving*” or “*obtaining*.” EX2007, ¶70. Therefore, Petitioner’s obviousness theory fails to explain how Seed discloses both “*obtaining the resource*” and “*serving the resource*” when “*the resource is determined to be popular*.” EX2007, ¶70.

Second, Petitioner alleges that Seed’s edge server is configured to fetch and serve the requested object based on popularity in response to the first request. Pet., 19-20 (“[i]f, however, the claim requires the edge server serve the object to the client after only a first request ...”). This is also incorrect. EX2007, ¶71. For the first request, Petitioner argues that “in response to a first request for a popular but not-yet replicated object, the client is redirected to another server that has the object, and the object is then replicated to the edge server for ‘future’ requests for the same object.” Pet., 19. But Seed explains that

[A] request for a streaming object will be served by a handling edge server if that handling edge server has a copy of that object.

Otherwise, the request is redirected, directly or indirectly, to a parent server for service of the requested streaming object to the client. If the requested streaming object is popular, the object is replicated from a parent server that has the requested streaming object to the handling edge server so that the handling edge server will serve the object from the edge of the network when the object is requested in the future.

EX1005, ¶18. In other words, in response to the first request, Seed’s edge server *redirects the request* to a parent server to serve the object and then replicates the object for future requests. EX1005, ¶18, EX2007, ¶71. Thus, Seed’s edge server only performs “*obtaining*,” not “*serving*,” because the request is redirected to and served by a parent server. EX2007, ¶71.

Petitioner then relies on Seed’s discussion of prior-art solutions that fetch objects on demand. Pet., 20. However, Seed does not disclose or suggest that fetching objects on demand is performed based on the popularity level. EX1005, ¶¶13-14; EX2007, ¶72. In fact, Seed does not mention popularity at all when describing these prior-art solutions. *See* EX1005, ¶¶13-14; EX2007, ¶72. Therefore, Petitioner fails to demonstrate that *the first server* performs “*obtaining*” and “*serving*” when the resource is popular, as claimed. EX2007, ¶72.

Attempting to cure this deficiency, Petitioner asserts that “[b]ased on these disclosures, a POSITA would have understood that Seed suggests having edge servers fetch and then serve popular objects.” Pet., 20. However, no explanation is provided to support this conclusory assertion; rather, Petitioner improperly relies on declarant testimony that itself provides no additional supporting evidence for this naked assertion. *See* Pet., 20; EX1003, ¶113; EX2007, ¶73.

Indeed, Patent Owner’s expert, Dr. Shenoy, demonstrates that the prior art approach described in Seed is fundamentally different from the claims and their advantageous solution. EX2007, ¶73. For example, Dr. Shenoy states

A solution to that issue, as recited in claim 1, is “obtaining” and “serving” based on the popularity level so that only popular objects are obtained. In this way, the edge servers would not be overloaded and resources for obtaining and storing the objects in the edge servers are efficiently used. Thus, I believe the on-demand fetching approach

described in Seed lacks any consideration of content popularity, rendering it critically deficient and materially different from the approach required by claim 1.

EX2007, ¶73.

Finally, Petitioner asserts that a POSITA would have been motivated to modify Seed to satisfy this element. Pet., 20-21. The Petition argues that there were only three typical ways to fulfill a request for a missing object: (1) obtain (fetch) the requested object from another server and then serve it to the client; (2) redirect the client to another server that has the object; or (3) forward the request to another server that has the object. Pet., 20-21. However, none of the three typical ways identified by the Petition satisfies the limitation. EX2007, ¶74. Thus, modifying Seed based on these alleged typical ways still does not satisfy the limitation.

EX2007, ¶74.

The first typical way involves fetching and serving, which ignores the claimed popular objects feature, as was shown was missing in Seed. EX2007, ¶75. Additionally, Seed explains that the fetching and serving approach is costly. EX1005, ¶¶13-15. This was one problem the '692 patent specifically addresses and overcomes by “obtaining” and “serving” popular objects. EX2007, ¶75; EX1001, 2:35-51, 6:48-7:10. Thus, the first typical way fails to meet the claim. EX2007, ¶75.

The second and third alleged typical ways are both related to redirecting the request and thus are not relevant to the claimed “*obtaining*” and “*serving*” the resource “*if the resource is determined to be popular.*” EX2007, ¶76.

Ground 1 fails to meet claim 1. EX2007, ¶77.

b. Ground 2

Petitioner makes several arguments for when the resource is popular, the first server takes two actions: “*obtaining*” **and** “*serving*” in Ground 2 based on Swildens. Pet., 51-54. All fail for similar reasons as the arguments based on Seed. EX2007, ¶78.

First, like Ground 1, Petitioner’s arguments are based on two separate actions: (1) in response to a first request for a popular but not-yet-replicated object, and (2) in response to subsequent requests after replication. Pet., 51-52. However, these two actions are mutually exclusive because a request is either a first request or a subsequent request, but it cannot be both. EX2007, ¶79. Thus, the edge server either performs “*obtaining*” in response to the first request or performs “*serving*” in response to the subsequent request, not both. EX2007, ¶79. Petitioner’s theory thus fails to explain how Seed’s edge server discloses both “*obtaining*” and “*serving*” if the resource is determined to be popular. EX2007, ¶79. Thus, in both of Petitioner’s scenarios Seed’s edge server performs only “*serving*” or “*obtaining.*” EX2007, ¶79. Petitioner’s obviousness theory fails to explain how

Seed discloses both “*obtaining the resource*” and “*serving the resource*” when “*the resource is determined to be popular*.” EX2007, ¶79.

Second, Petitioner argues that a POSITA would have been motivated to modify Swildens to fetch popular content rather than redirect clients based on Seed’s disclosure of fetching. Pet., 52-53. Petitioner again relies on Seed’s disclosure of fetching on demand in paragraph 13 to meet the claims. Pet., 52-53. As discussed in Ground 1, Seed describes prior-art solutions that fetch objects on demand but does not disclose that the fetching is performed based on the popularity level. EX1005, ¶¶13-14; EX2007, ¶80. In fact, Seed does not mention popularity at all when describing these prior-art solutions. *See* EX1005, ¶¶13-14. Therefore, Petitioner fails to demonstrate that *the first server* performs “*obtaining*” **and** “*serving*” when the resource is popular, as claimed. EX2007, ¶80.

Finally, Petitioner asserts that a POSITA would have been motivated to modify Swildens to satisfy this element. Pet., 53-54. Petitioner argues that there were only three typical ways for a server to ensure that requests for an object missing from its cache are fulfilled: (1) fetch-and-serve; (2) redirect the client; or (3) forward the request. Pet., 53. However, none of the three typical ways identified by the Petition satisfies the limitation of the claims. EX2007, ¶81.

The first typical way involves fetching and serving, which ignores the claimed popular objects feature, as was shown to be missing in Swildens. EX2007,

¶82. Thus, the first typical way fails to meet the claim. EX2007, ¶82. The second and third alleged typical ways are both related to redirecting the request, and thus are not relevant to the claimed “*obtaining*” and “*serving*” the resource “*if the resource is determined to be popular.*” EX2007, ¶82. Because none of the three typical ways satisfy the claim limitation, modifying Swildens based on these alleged typical ways does not satisfy the limitation. EX2007, ¶82.

Accordingly, Petitioner fails to show the Swildens-Seed combination teaches “*obtaining*” and “*serving*” the resource “*if the resource is determined to be popular.*” EX2007, ¶83.

C. Independent Claim 10

There are two additional failures of proof for claim 10. EX2007, ¶84.

1. Petitioner fails to show Seed-Swildens or Swildens-Seed renders obvious the claimed “*provide the edge servers with a second set of information about which tier of servers should handle those requests for content.*”

Independent claim 10 recites “*at least one popularity service constructed and adapted ... to provide the edge servers with a second set of information about which tier of servers should handle those requests for content.*” Petitioner admits that Seed does not expressly disclose a separate “popularity service” that receives information from edge servers and provides information back to edge servers. Pet., 36. Instead, for both Grounds 1 and 2, Petitioner relies on Swildens for purportedly teaching this feature. Pet., 36-37, 69. Yet Petitioner then simply provides a single

conclusory sentence that wrongly asserts Swildens' auto migration is the same as the claimed "*provid[ing] the edge servers with a second set of information about which tier of servers should handle those requests for content.*" See Pet., 36-37, 69. EX2007, ¶85.

The silence in Petitioner's explanation is not accidental because Swildens auto migration and the claimed "*provid[ing] the edge servers with a second set of information about which tier of servers should handle those requests for content*" are two disparate operations. EX2007, ¶86. Swildens' auto migration server selects the content and performs the auto migration without providing any information to the edge servers. EX2007, ¶86. For example, Swildens describes that at a specified time interval, content that falls above a threshold value is sorted and the "top few" are selected for auto migration. EX1006, ¶126.

This is further evident by Swildens' characterization of the auto migration server. EX2007, ¶87. For example, Swildens explains that "[a] centralized server can periodically collect logs from all of the edge serve[r]s." EX1006, ¶127. To do that, "[a] single auto migration daemon can combine data from multiple edge servers located in a single edge location to determine the most popular content" and "the popularity is based on the aggregate data from all of the edge servers that are accessing that common storage." EX1006, ¶130. Thus, centralization is the fundamental concept of the auto migration process. Trying to use this to "*provide*

the edge servers with a second set of information about which tier of servers should handle those requests for content” would defeat the entire reason for Swildens’ invention. EX2007, ¶87.

Furthermore, even if Swildens’ auto migration server provides information to the edge server, it would not be the claimed “*second set of information about which tier of servers should handle those requests for content.*” EX2007, ¶88.

Swildens explains that the auto migration server selects a list of content, such as the top 20 popular content. EX1006, ¶126. Thus, any information provided by Swildens’ auto migration server would be that list of content, not “*a second set of information about which tier of servers should handle those requests for content,*” as recited in claim 10. EX2007, ¶88.

In sum, Petitioner fails to show “*at least one popularity service constructed and adapted ... to provide the edge servers with a second set of information about which tier of servers should handle those requests for content,*” and thus fails to establish Seed-Swildens or Swildens-Seed teaches this element of independent claim 10. EX2007, ¶89.

2. Petitioner fails to show Seed-Swildens or Swildens-Seed renders obvious the claimed “*determining a tier of servers to handle a request based, at least in part, on a measure of popularity of requested content.*”

Independent claim 10 recites “*the popularity service making its determinations,*” which includes “*which tier of servers should handle those requests for content,*” “*based, at least in part, on a measure of popularity of requested content.*” The Petition admits that Seed does not disclose this feature, but attempts to cure the deficiency using Swildens in both Grounds 1 and 2. Pet., 36, 69. However, Swildens does not teach the claimed “*which tier of servers should handle those requests for content,*” “*based, at least in part, on a measure of popularity of requested content.*” EX2007, ¶90.

Swildens uses popularity level to determine which content to auto-migrate, not the claimed “*which tier of servers should handle those requests for content.*” EX1006, ¶126; EX2007, ¶91. Rather than addressing this mismatch, Petitioner simply equates, without any explanation, Swildens’ tabulation and selection of content for auto-migration with the claimed determining a tier of server to handle the request. Pet., 36-37, 69. However, this conclusion is logically flawed on its face and contradicted by Swildens. EX2007, ¶91. Determining a list of content to auto-migrate is different from the claimed determining “*which tier of servers should handle those requests for content.*” EX2007, ¶91. Dr. Shenoy explains that “a list of content is digital data, but servers are physical devices.” EX2007, ¶91. In other

words, digital data is virtual and is different from servers that physically exist.

Thus, determining digital data, such as the list of content, is different from determining ‘*which tier of servers should handle those requests for content,*’ as recited in claim 10.” EX2007, ¶91.

Furthermore, Petitioner ignores Swildens’ explicit explanation of how it determines which server to handle a request. EX2007, ¶92. Swildens explains that when its edge server receives a request, it redirects the user to an origin server if the edge server does not have the requested content. EX1006, ¶46. If the origin server does not have the requested content, it redirects the user to a staging server. EX1006, ¶47. Thus, Swildens determines which server to handle a request based on availability (with respect to redirecting the user to a staging server), not the claimed “*the popularity service making its determinations.*” EX2007, ¶92.

Petitioner thus fails to establish either Ground 1 or 2 renders obvious “*the popularity service making its determinations based, at least in part, on a measure of popularity of requested content.*” EX2007, ¶93.

D. Independent Claim 15: Seed and Swildens, in any order and for both Grounds, do not teach the claimed “*redirecting being based, at least in part, on a current popularity value for the resource.*”

Independent claim 15 requires that “*the redirecting being based, at least in part, on a current popularity value for the resource.*” The Petition fails to show that Seed-Swildens or Swildens-Seed teaches this feature. EX2007, ¶94.

As a threshold matter, claim 15 does not have multiple meanings. EX2007, ¶95. Petitioner misconstrues the relationship between terms in claims 15 and 18 and inappropriately implies there is an indefiniteness issue¹ with claim 18 when none exists. *Contra* Pet., 40-44. Petitioner wrongly assumes that the “*second tier server*” is a server in the “*second tier*.” Pet., 41. Contrary to the Petition’s assumption, element 15[b] recites “*the second tier server being in a tier distinct from the tier containing the first server*” (i.e., it does not recite the second tier server is in the second tier, but rather that it is a server in a tier distinct from the tier of the first server). EX2007, ¶95. Claim 18 further recites that “*the first server is a parent server in the second tier and the second [tier] server is a server in the origin tier*.” Thus, claim 18 also places the “*first server*” and the “*second tier server*” in different tiers, which does not contradict element 15[b]. EX2007, ¶95.

Regardless, Petitioner’s obviousness theory for element 15[b] relies on Seed to teach “*redirecting would be based, at least in part, on a current popularity value for the resource*.” Pet., 44 (citing EX1005, ¶26). But, as discussed in Section IX.B.1, this is not what Seed discloses. Rather, Seed discloses that redirecting is done because “the selected edge server does not have the requested object,” which

¹ 35 U.S.C. § 311(b) sets forth that a petitioner may request to cancel claims *only* under Sections 102 or 103.

applies equally to both popular and unpopular content according to Seed. EX1005, ¶26; EX2007, ¶96. Thus, Petitioner fails to show element 15[b] is obvious over Swildens-Seed or Seed-Swildens. EX2007, ¶96.

E. All Grounds: Petitioner improperly relies on declarant testimony to fill gaps in the prior art.

The Petition often presents conclusory statements with little to no explanation on how the cited evidence supports the assertion. For example, in the analysis of limitation 1[c], Petitioner relies on the assertion that a “POSITA would have understood that Seed suggests having edge servers fetch and then serve popular objects to the client in response to a request,” only citing a paragraph of declarant testimony, which itself does not cite any evidence supporting the assertion. Pet., 20-23, 52-54; *see also* Pet., 40-41, 73 (relying on a similar conclusory assertion for 15[c]). As another example, the Petition makes a number of conclusory assertions in its analysis of element 1[d.ii] that attempt to fill gaps in Swildens by relying on declarant testimony to support the erroneous allegation (for reasons discussed above) that “a POSITA would have understood that a hash function would typically result in each server having distinct, non-overlapping portions of the library.” Pet., 26, 58-59; *see also* Pet., 45, 77-78 (relying on similar declarant testimony for element 15[d.ii]). These are just a few examples; there are many others.

The Director should reject these arguments because they are insufficient to satisfy Petitioner's burden of persuasion. *See Parus Holdings, Inc. v. Google LLC*, 70 F.4th 1365, 1371 (Fed. Cir. 2023) (noting that a party's burden "cannot be met simply by throwing mountains of evidence at the Board without explanation or identification of the relevant portions of that evidence"). 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).

Such cursory and conclusory analysis falls well short of Petitioner's obligation to perform reasoned analysis based on the references asserted in its challenges. Because this deficiency is replete throughout the entire Petition, the Director should deny institution.

X. *CHEVRON AND DEEPER FAVOR DENYING INSTITUTION.*

The merits consideration phase of an IPR includes determining whether the petition presents a sufficient number of grounds/challenges that meet the reasonable likelihood standard. When “a petition meets the standards for institution in relation to fewer than all the challenges presented, or where deficiencies in the petition ... affect only a portion of the challenges presented, the [Board] will evaluate all the challenges and determine whether, in the interests of efficient administration of the Office and integrity of the patent system, the entire petition should be denied.” Trial Practice Guide, § II.D.3 (citing 35 U.S.C. §§ 316(b), 326(b)); *accord Chevron Oronite Co., LLC v. Infineum USA L.P.*, IPR2018-00923, Paper 9 (PTAB Nov. 7, 2018) (informative); *Deeper, UAB v. Vexilar, Inc.*, IPR2018-01310, Paper 7 (PTAB Jan. 24, 2019) (informative).

As discussed throughout this Preliminary Response, the Petition suffers from numerous fatal deficiencies. Given these defects, institution of trial should be denied even if some of the claims or Grounds meet the standard for institution (which they do not).

It would not be an efficient use of the Board's valuable time and resources to institute review when resolution of this proceeding would be costly and time-consuming to address challenges that are flawed from the outset. *See* Interim Process Memo for PTAB Workload Management², 3; 35 U.S.C. § 316(b).

XI. CONCLUSION

The Petition fails to show a reasonable likelihood that any of the challenged claims are unpatentable. Therefore, the Director should deny institution.

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg (Reg. No. 43,447)
Daniel S. Block (Reg. No. 68,395)
Kristina Caggiano Kelly (Reg. No. 58,164)
Wenhao Xiong (Reg. No. 78,365)
Counsel for Patent Owner Sandpiper CDN, LLC

Date: March 19, 2026

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

² Available at

<https://www.uspto.gov/sites/default/files/documents/InterimProcesses-PTABWorkloadMgmt-20250326.pdf>.

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 8,320 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 in 14-point Times New Roman font.

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg
Registration No. 43,447
Attorney for Patent Owner Sandpiper CDN, LLC

Date: March 19, 2026

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

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

I certify that the above-captioned **PATENT OWNER PRELIMINARY RESPONSE UNDER 37 C.F.R. § 42.107(a)** and associated Exhibits 2007-2020 were served in their entirety on March 19, 2026, upon the following parties via electronic mail:

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

Respectfully submitted,

STERNE, KESSLER, GOLDSTEIN & FOX PLLC

/Jason D. Eisenberg/

Jason D. Eisenberg
Registration No. 43,447
Attorney for Patent Owner Sandpiper CDN, LLC

Date: March 19, 2026

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