

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

MICROSOFT CORPORATION,
Petitioner

v.

SANDPIPER CDN, LLC,
Patent Owner

Case IPR2026-00180
U.S. Patent No. 10,701,173

**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	4
III.	THE '173 PATENT	9
	A. Overview	9
	B. Prosecution history	10
IV.	PERSON OF ORDINARY SKILL IN THE ART (POSITA)	14
V.	CLAIM CONSTRUCTION	14
VI.	OVERVIEW OF PRIOR ART	15
	A. O'Rourke	15
	B. Middleton	18
	C. Blumofe	20
	D. MS-MSSO	22
VII.	ARGUMENTS AGAINST INSTITUTION.....	23
	A. The Petition fails to establish that O'Rourke renders obvious the claimed "when it is determined by the node that said content is cached locally, then, after determining that the content is cached locally, ... determining ... a current cache policy associated with the content" (Ground 1, all independent claims).	25
	1. Petitioner's reliance on Block 510 is erroneous	26
	2. Petitioner's alternative theory modifying block 510 fares no better.....	29
	B. The Petition fails to establish that Middleton renders obvious the claimed "determining, by the node, a current cache policy associated with the content, wherein determining the current cache policy comprises determining a time period for expiration of content in the cache" (Ground 2, all independent claims).....	32
	C. The Petition's improper reliance on expert testimony to fill the gaps in O'Rourke and Middleton warrants denial on the merits.	37
VIII.	CONCLUSION.....	39

TABLE OF AUTHORITIES

Cases

<i>Abbott Labs. v. Andrx Pharms., Inc.</i> , 452 F.3d 1331 (Fed. Cir. 2006)	11
<i>Chemours Company FC, LLC v. Daikin Industries, LTD.</i> , No. 20-1289 (Fed. Cir. July 22, 2021).....	19
<i>Henny Penny Corp. v. Frymaster LLC</i> , 938 F.3d 1324 (Fed. Cir. 2019)	29
<i>Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.</i> , 821 F.3d 1359 (Fed. Cir. 2016)	19, 21
<i>In re Magnum Oil Tools Int’l, Ltd.</i> , 829 F.3d 1364 (Fed. Cir. 2016)	18
<i>Regents of the Univ. of Cal. v. Broad Inst., Inc.</i> , 903 F.3d 1286 (Fed. Cir. 2018)	19
<i>Sirona Dental Sys. GmbH v. Institut Straumann AG</i> , 892 F.3d 1349 (Fed. Cir. 2018)	18
<i>Sisvel Int’l. S.A. v. Sierra Wireless, Inc.</i> , 81 F.4th 1231 (Fed. Cir. 2023)	11
<i>Virtek Vision Int’l ULC v. Assembly Guidance Sys., Inc.</i> , 97 F.4th 882 (Fed. Cir. 2024)	22
<i>Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.</i> , 200 F.3d 795 (Fed. Cir. 1999)	11
<i>Advanced Bionics, LLC v. MED-EL Elektromedizinische Geräte GmbH</i> , IPR2019-01469, Paper 6 (P.T.A.B. Feb. 13, 2020).....	8, 10, 13, 14
<i>In-Depth Geophysical, Inc. v. ConocoPhillips Co.</i> , IPR2019-00850, Paper 56 (PTAB Sept. 3, 2020).....	21

Statutes

35 U.S.C. § 311(b)	29
35 U.S.C. § 312(a)(3).....	29
35 U.S.C. § 314(a)	18
35 U.S.C. § 316(e)	18

Other Authorities

37 C.F.R. § 42.104(b)(4).....	29
Rule 104(b)(4).....	29, 30

PATENT OWNER'S 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	U.S. Patent Publication No. 2006/0095532 ("pub-O'Rourke")
2008	Declaration of Samuel H. Russ, Ph.D.
2009	<i>Curriculum Vitae</i> of Samuel H. Russ, Ph.D
2010	"CDN Stands For Content Delivery Network: History and Evolution," <i>BlazingCDN</i> , accessed at https://blog.blazingcdn.com/en-us/cdn-stands-for-content-delivery-network-history-and-evolution (last accessed August 1, 2025)
2011	"What is a CDN (Content Delivery Network)?" <i>Amazon Web Services</i> , accessed at https://aws.amazon.com/what-is/cdn/ (last accessed August 1, 2025)
2012	Basu, S., et al., "Adaptive TTL-Based Caching for Content Delivery," <i>IEEE/ACM Transactions on Networking</i> , 99:1-15 (April 2018)

Exhibit No.	Description
2013	“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)
2014	“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)
2015	Entrepreneur of the Year, June 19, 2000
2016	“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)
2017	“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)
2018	“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
2019	Infringement Contentions, <i>Sandpiper CDN, LLC v. Microsoft Corp.</i> , Case No. 2:25-cv-00664 (E.D. Tex. Sept. 16, 2025)

I. INTRODUCTION

U.S. Patent No. 10,701,173 (“the ’173 patent”) provides novel techniques for distributing content to end users via content delivery networks (CDNs).

Marking a departure from conventional cache management methods, the novel cache management techniques of the ’173 patent address how and when cache policies are checked to ensure that an up-to-date caching policy is used without requiring forced invalidations of cached content when a new policy is implemented. This revolutionary approach provides an iterative process for ensuring that caching and cache policies are always real-time current, without needing to perform disruptive global updates. Nothing in the cited prior art teaches or suggests the late-binding approach of the claimed cache policy, which turns conventional systems on their heads.

Petitioner presents six Grounds challenging the claims of the ’173 patent:

Ground	Art	Basis	Claims (Independent Bolded)
1A	O’Rourke (EX1005)	§ 103	1 , 2, 4-6, 8 , 9, 11-13
1B	O’Rourke, MS-MSSO (EX1008)	§ 103	3, 10
1C	O’Rourke, Blumofe (EX1007)	§ 103	7, 14
2A	Middleton (EX1006)	§ 103	1 , 2, 4-6, 8 , 9, 11-13
2B	Middleton, MS-MSSO	§ 103	3, 10
2C	Middleton, Blumofe	§ 103	7, 14

Pet., 2.

Ground 1 (i.e., Grounds 1A-1C, collectively) relies on O’Rourke alone to disclose all of the limitations for the independent claims, and most of the limitations of the dependent claims. This single-reference obviousness theory, however, admits that the novel features of the claims are not actually found in O’Rourke, and instead tries to use expert opinion and alleged common sense to fill these gaps. Petitioner’s argument ultimately fails, however, for many of the same reasons it failed during prosecution, when the Examiner considered the teachings of O’Rourke and nonetheless found the claims allowable. Petitioner adds MS-MSSO and Blumofe only to fill in a limitation in dependent claims 3 and 10, or 7 and 13, respectively. Neither is alleged to overcome the shortcomings of O’Rourke.

Specifically, for Ground 1A (including both independent claims), the Petition fails to establish that O’Rourke discloses or renders obvious “*when it is determined by the node that said content is cached locally, then, after determining that the content is cached locally[,]* ... determining ... *a current cache policy associated with the content,*” as recited by both independent claims. Regarding this claim feature, the ’173 patent explains that requests for cached content in CDNs employ “a late-binding of the caching policy, thereby assuring that the latest internal caching policy is used without having to force invalidations” (which are potentially disruptive and resource-intensive processes). EX1001, 3:66-4:3.

Although O’Rourke generally describes processes for checking the freshness of cached content, O’Rourke never describes the claimed checking a current cache policy and conditioning the check on a request for content and subsequent determination that the content is locally cached at a node. These claimed functions and their logic structures are completely missing from O’Rourke.

Ground 2 (i.e., Grounds 2A-2C, collectively) essentially repeats the flawed analysis of Ground 1A, replacing O’Rourke with a different but equally deficient reference, Middleton. Petitioner similarly relies on Middleton alone to teach all limitations of the independent claims and most of the dependent claim limitations, relying once again on MS-MSSO and Blumofe to fill in a limitation in some dependent claims. Neither MS-MSSO nor Blumofe is alleged to remedy any shortcoming or fill any gap in Middleton for the independent claims. While the Office did not expressly consider Middleton during prosecution, it is cumulative of O’Rourke, Petitioner uses it for all the same teachings, and it fails for largely the same reasons.

For instance, in Ground 2A (including both independent claims), the Petition fails to establish that Middleton teaches or suggests at least the claimed *“determining, by the node, a current cache policy associated with the content, wherein determining the current cache policy comprises determining a time period for expiration of content in the cache.”* While Middleton generally describes

processes for checking staleness of cached content, Middleton never describes the claimed checking for a “current” cache policy.

Accordingly, Petitioner fails to establish that either of the ’173 patent’s independent claims 1 and 8 are rendered obvious under any presented Ground. Because the Petition cannot show that any independent claim is obvious, it cannot show that any of the dependent claims depending therefrom are obvious. For at least these reasons, the Board should deny institution of this proceeding.

II. BACKGROUND

The challenged patent arises from pioneering technology directed to delivering content across a network. Content delivery networks (CDNs) provide critical services that enable content providers to quickly deliver online content to millions of consumers simultaneously over the Internet. EX2008, ¶26. 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 (“Patent Owner” or “Sandpiper”).

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.

EX2008, ¶26. 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. EX2008, ¶27. 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. EX2008, ¶27.

In the mid-1990s, Andrew Swart and David Farber were among the first individuals to develop services that allowed content providers to distribute their content over the Internet, while avoiding the common congestion and performance issues that plagued Internet transmission at that time. Sandpiper's solution was to deploy CDN servers around the world, replicate 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. EX2008, ¶28.

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. EX2008, ¶29. 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.

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 beginning on September 11, 1998. Sandpiper Networks' CDN was capable of caching, and was used to cache and deliver Internet resources including, *inter alia*, pictures, text files, dynamic resources, and streaming multimedia resources.

The industry-wide acclaim for Sandpiper Networks further confirms the non-obvious nature of its innovations. Sandpiper's pioneering CDN technology was honored with the "Best of the Best" award at Networld+ Interop 99 Las Vegas and the 1999 Most Innovative New Product Award from the University of California, San Diego (UCSD) CONNECT program. EX2013; EX2014. 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. EX2015. 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. EX2016; EX2017; EX2018.

In the early-to-mid 2000s, demand for CDNs exploded. EX2008, ¶30. 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 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.

One such company is Microsoft. Microsoft built its own in-house 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* EX2019. Microsoft did not license this patented technology from Sandpiper, Level 3,

and/or its predecessors. Accordingly, Sandpiper brought suit against Petitioner Microsoft Corporation (“Petitioner” or “Microsoft”) in district court for patent infringement. Microsoft filed IPR petitions against most of Sandpiper’s asserted patents in response.

III. THE ’173 PATENT

A. Overview

The ’173 patent is entitled “Caching in a Content Delivery Framework,” and was filed as U.S. Application No. 15/588,935 on May 8, 2017. It is a continuation of U.S. Patent No. 9,648,127, filed on December 15, 2014. The ’173 patent is directed to improvements in CDN technology, specifically addressing how CDNs manage caching policies, prevent stale content from being served, and handle subscriber-specific caching controls. CDN caches store copies of subscriber content so that end-user requests can be served by CDN nodes. EX2008, ¶31.

Claimed embodiments of the ’173 patent address technical issues with handling requests and timely providing content. The ’173 patent describes enhancing this process by dynamically determining the appropriate caching policy at the time a request for content is received, enabling the CDN to honor up-to-date rules while improving efficiency and flexibility. EX2008, ¶32. The inventors pioneered the described unconventional technical

solutions for maintaining content freshness and enforcing subscriber-specific caching directives in a scalable, low-overhead manner to address problems associated with stale content, inconsistent policy application, and the operational burden of widespread invalidations.

B. Prosecution history

The '173 patent was given a thorough examination and included consideration of the primary reference used in Ground 1 in this Petition (O'Rourke). Via Information Disclosure Statements, Patent Owner cited O'Rourke, the published application of O'Rourke ("pub-O'Rourke"), and an International Preliminary Report on Patentability (IPRP) that applied pub-O'Rourke to Patent Owner's then-pending claims (on anticipatory grounds). EX1002, 278-284, 304. Along with issuance of the first Office Action, the Examiner indicated they considered O'Rourke, pub-O'Rourke, and the IPRP. EX1002, 272, 277; EX2008, ¶33.

The Examiner, after being presented with the IPRP containing anticipatory arguments based on pub-O'Rourke, chose to issue a different anticipation rejection using a different reference (U.S. Patent No. 9,984,166), which was found to be closer prior art than O'Rourke. EX1002, 262-270. A nonstatutory double patenting (NDP) rejection with respect to

the issued parent was also issued. *Id.* Patent Owner responded with rebuttal arguments. EX1002, 250-258; EX2008, ¶34.

Thereafter, a series of Office Actions and responses ensued, involving an obviousness rejection (relying on a combination of U.S. Patent No. 9,984,166 and U.S. Patent No. 2015/0347248), an Examiner interview, Request for Continued Examination, Patent Owner's submission of a terminal disclaimer, and amendments to the claims. EX1002, 104, 128-136, 139, 142-155, 160-168, 171-182, 186-192, 208-220, 232, 238-247; EX2008, ¶35. Amendments included the addition of "after determining that the content is cached locally" to step (B) and "wherein determining the current cache policy comprises determining a time period for expiration of content in the cache" to step (B)(1) of the independent claims. EX1002, 112-114. These amendments successfully distinguished the closest prior art, and they distinguish both primary references cited in the present Petition. After these amendments, the Examiner entered a Notice of Allowance. EX1002, 18-25. Patent Owner then paid the issue fee and the patent issued. EX1002, 8-17; EX2008, ¶36.

Petitioner's challenges to the '173 patent rely on *the same art and analysis* that the Office previously considered. *See Advanced Bionics, LLC v. MED-EL Elektromedizinische Geräte GmbH*, IPR2019-01469, Paper 6 at 7-8 (P.T.A.B. Feb.

13, 2020) (precedential). Patent Owner even agrees that “the Examiner was aware of (and considered) the Written Opinion of the International Searching Authority from the ’173Pat.’s PCT counterpart.” Pet., 7 (citing EX1002, 277-284).

Petitioner’s application of O’Rourke substantially parallels the IPRP that was considered by the Office. EX1002, 283. The chart below illustrates the overlap in Petitioner’s O’Rourke arguments and the IPRP considered by the Examiner, specifically for elements of challenged claims 1 and 8.

’173 Patent—Claim 1/8 Elements	O’Rourke Portions Cited by Petitioner	Pub-O’Rourke Portions Cited in IPRP
(A) in response to a request for content, determining, by a node in the CDN, if the content is cached locally at the node,	Block 502, block 505; EX1005, 6:1-12, 6:14-15; Pet., 15-18.	EX2007, ¶[0039] (covers block 502 and EX1005, 6:1-15); EX1002, 283.
(B) when it is determined by the node that said content is cached locally, then, after determining that the content is cached locally:	Block 504; EX1005, 6:28-31, 6:31-34; Pet., 18-20.	EX2007, ¶[0041] (covers EX1005, 6:28-54, which implicate policies at block 504); EX1002, 283.
(B)(l) determining, by the node, a current cache policy associated with the content, wherein determining the current cache policy comprises determining a time period for expiration of content in the cache; and	Block 504; EX1005, 6:1-10, 6:12-15; Pet., 26-28.	EX2007, ¶[0039] (covers block 504 and EX1005, 6:1-15); EX1002, 283.

'173 Patent—Claim 1/8 Elements	O'Rourke Portions Cited by Petitioner	Pub-O'Rourke Portions Cited in IPRP
then (B)(2) determining, by the node and based on said current cache policy associated with the content, whether it is acceptable to serve the content that is cached locally	EX1005, 6:31-34, 6:37-40; Pet., 29.	EX2007, ¶[0041] (covers EX1005, 6:28-54); EX1002, 283.

Furthermore, Petitioner's mappings for the elements of claims 1 and 8 designated (B)(3), (B)(3)(i), and (B)(3)(ii) identify O'Rourke's blocks 510 and 514 and associated text. This mapping duplicates the IPRP's reliance on pub-O'Rourke at paragraph [0041], describing blocks 510 and 514. Pet., 29-35; EX1002, 283. The Petition's use of the "same or substantially the same prior art or arguments previously ... presented to the Office," should fail now for the same reasons it failed in *Advanced Bionics*. *Advanced Bionics*, IPR2019-01469, Paper 6 at 7-8.

Indeed, Petitioner has not "demonstrated that the Office erred in a manner material to the patentability of [the] challenged claims." *Id.* at 8. Petitioner does not articulate a "material error" by the Examiner (e.g., a misread claim term, an overlooked passage in O'Rourke, or a factual mistake). Instead, Petitioner emphasizes that O'Rourke and the IPRP were in the file yet "never applied," and presents a slightly modified interpretation of O'Rourke's flow to fit the "after determining that ... a current cache policy" sequence in the claims. Pet., 7, 12, 22-

26. That is, Petitioner implies the Office got it wrong by omission, but does not identify a concrete misapprehension of law or fact. There is none.

IV. PERSON OF ORDINARY SKILL IN THE ART (POSITA)

A POSITA is presumed to be aware of all pertinent art. A POSITA possesses the “understandings and knowledge reflected in the prior art,” and thinks along conventional wisdom. *See Abbott Lab’ys. v. Andrx Pharms., Inc.*, 452 F.3d 1331, 1336 (Fed. Cir. 2006). At this stage, Patent Owner submits that the Board does not need to resolve the appropriate level of skill for a POSITA for purposes of deciding institution because the Petition fails under any reasonable articulation of a POSITA.

V. CLAIM CONSTRUCTION

Absent lexicography or disavowal, claim terms in an IPR are given their plain and ordinary meaning in light of the specification and prosecution history. *Sisvel Int’l. S.A. v. Sierra Wireless, Inc.*, 81 F.4th 1231, 1236 (Fed. Cir. 2023). At this preliminary stage, no term requires express construction beyond its plain and ordinary meaning, as would have been understood by a POSITA at the time of the invention in light of the specification and the patent’s prosecution history. *See Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999).


Patent Owner notes, however, that the term “cache policy” as used in the challenged claims refers to a very different structure and logic than the “cache policy” disclosed in the applied prior art. This difference does not necessarily imply a different definition for the words “cache” or “policy” themselves, but merely reflects that in this field, the plain meaning of the term “cache policy” is highly contextual, and can denote different functionalities. As discussed in detail below, a POSITA reading this term in the context of the ’173 patent would understand it to be referring to a different function, performed by a different program, than the “cache policy” described in the O’Rourke and Middleton references.

VI. OVERVIEW OF PRIOR ART

Petitioner relies on different combinations of four references across the various Grounds. The independent claims only rely on O’Rourke and Middleton, singularly, to allege obviousness. EX2008, ¶37.

A. O’Rourke

O’Rourke discloses a method and apparatus for selecting cache and proxy policy in a cache server used to distribute media content in a content delivery network. EX1005, Abstract, 1:13-15; EX2008, ¶38. Figure 5 illustrates a procedure for handling a client request for pre-recorded content.



EX1005, FIG. 5.

In O'Rourke's procedure 500, after a client requests content (block 502), the cache server first "identifies the policies related to requests for pre-recorded content" at block 504. EX1005, 6:12-14; EX2008, ¶39. Block 504 is keyed to the type of request and identifies which pre-defined cache

policies may apply to handling that request. EX2008, ¶39. O'Rourke does not describe block 504 as being conditioned on whether the requested content is cached locally, nor as determining any policy specific to a cached object as claimed. *Id.*


At block 505, the cache server “identifies whether the requested content is stored on the cache server.” EX1005, 6:14-15; EX2008, ¶40. This is a cache-presence check that determines only “whether the requested content is available on the cache server or whether the requested content is stored on a different server.” EX1005, 6:14-19; EX2008, ¶40. O'Rourke does not describe any cache-policy determination as being conditioned on the outcome of block 505 as claimed. *See* EX1005, 6:20-37; EX2008, ¶40.

If the content is stored on the cache server, O'Rourke proceeds to block 510, where the cache server determines “whether the cached version of the requested content is an acceptable version.” EX1005, 6:31-37; EX2008, ¶41. O'Rourke explains that this determination concerns freshness or version matching, such as whether the cached version matches the origin server's version or falls within a permitted age window. EX1005, 6:31-37; EX2008, ¶41. Block 510 is described as a version-acceptability or freshness check, not as a determination of a cache policy associated with the content as claimed. EX1005, 6:31-37, 7:64-67; EX2008, ¶41.

Thus, O'Rourke discloses separate steps for (i) identification of policies related to a request (block 504), (ii) determination of whether content is cached (block 505), and (iii) evaluation of whether a cached copy is acceptable or fresh (block 510) as distinct steps, without describing block 510 as a determination of a current cache policy made after confirming that the content is cached locally. EX1005, 6:12-37; EX2008, ¶42.

B. Middleton

Middleton discloses a cache server architecture for streaming content in a content delivery network, including mechanisms for caching content, streaming content to clients, and revalidating cached content. EX1006, ¶¶[0033]-[0034]; EX2008, ¶43. In Middleton, “[g]enerally, each content item has an associated ... [TTL] value.” EX1006, ¶[0052]; EX2008, ¶43. Middleton describes a cache server that manages cached content by comparing content age to an associated time-to-live (TTL) value and reacting when cached content becomes stale. EX1006, ¶[0080], FIG. 5; EX2008, ¶43.



EX1006, FIG. 5.

If the content has not exceeded its TTL, the cached content is streamed to the requester (operation 508). EX1006, ¶[0082]; EX2008, ¶44. When the cached content is stale, Middleton describes sending a HEAD request to a media access server to attempt revalidation (operation 510).

EX1006, ¶¶[0081]-[0082]; EX2008, ¶44. If the media access server indicates that the cached content remains valid, the cache server updates the TTL associated with the content and continues to treat the cached content as usable (operation 514). EX1006, ¶[0082]; EX2008, ¶44. If revalidation fails, the stale content is deleted and the cache server retrieves a fresh copy of the content (operations 516, 518). EX1006, ¶[0083]; EX2008, ¶44.

Throughout this process, Middleton treats TTL values as pre-existing attributes associated with content and uses them as part of a freshness or staleness check. EX1006, ¶[0052]; EX2008, ¶45. Middleton does not describe determining or selecting a cache policy at request time as claimed. EX2008, ¶45. Instead, it explains that each content item has an associated TTL value that is read and compared against content age to assess staleness, with the cache server reacting based on the pre-associated TTL. EX1006, ¶[0080]; EX2008, ¶45. TTL refresh occurs only after a staleness event and in response to a revalidation exchange, functioning as a reactive update to expiration information content rather than a determination of which cache policy governs the request. EX1006, ¶¶[0081]-[0082]; EX2008, ¶45.

C. Blumofe

Blumofe discloses extending a content delivery network (CDN) into mobile or wireline networks by deploying an intermediate caching

component referred to as an “Extender.” EX1007, ¶¶[0025]-[0027]. “The Extender operates as a passive web caching proxy beyond the edge of the CDN.” EX1007, ¶[0027]. The Extender “interacts with the CDN by forwarding requests to the CDN edge and caching/serving content based on response headers received from the edge.” *Id.* When the Extender receives a request for CDN-managed content, it forwards the request to a CDN edge server and applies cache-handling directives returned by the CDN, such as TTL expiration times, no-store instructions, or purge indications. EX1007, ¶¶[0031], [0037]. The decision logic governing cache suitability and duration resides at the CDN edge server. EX1007, ¶[0054].

Petitioner relies on Blumofe only for narrow teachings regarding CDN-supplied directives for cache handling, invoking Blumofe’s Extender to supply externally dictated cache behavior rather than any request-responsive cache-policy determination by the node. Pet., 46, 53, 89, 90. That reliance does not purport to address the deficiencies of the primary references identified in Grounds 1A and 2A, which concern whether a node determines a current cache policy after determining that content is cached locally. Nor does Petitioner contend in Grounds 1C or 2C that Blumofe supplies the missing claim elements from the primary references in Grounds 1A and 2A. Instead, Blumofe is cited only as a supplemental reference, and

its disclosure of cache behavior controlled remotely by a CDN edge server does not remedy the gaps identified in the primary art. *Id.*

D. MS-MSSO

MS-MSSO explains that it “provides an overview of the Media Streaming Server System Overview Protocol Family” and “describes the intended functionality of the Media Streaming Server System and how it interacts with systems or applications that need to stream media,” framing the document as an architectural overview rather than a disclosure of application-level request handling logic. EX1008, 2. With respect to caching, MS-MSSO states that “[i]f the server is acting as a caching server the proxy server will request a stream from the origin server and allow multiple clients to stream the content,” and that, for on-demand content, “a reverse proxy server can reduce the load on the origin server by caching the content from the origin server and streaming it to clients from its cache.” EX1008, 16.

MS-MSSO similarly explains that “[a]s a Cache/Proxy server, [the media server] receives content from the origin server and caches the content for further distribution,” without describing how cache behavior is selected, controlled, or determined in response to a request. EX1008, 51. MS-MSSO further clarifies that the supported protocols “do not share any data or state

between the protocols supported by the system,” underscoring that the document does not define shared decision logic, policy rules, or data structures governing cache behavior. EX1008, 38.

Petitioner relies on MS-MSSO only for narrow, background teachings regarding the existence of streaming servers within a media streaming architecture. Pet., 40, 85. That limited reliance does not address the deficiencies of the primary references identified in Grounds 1A and 2A. Nor does Petitioner contend in Grounds 1B or 2B that MS-MSSO supplies the missing claim elements from the primary references in Grounds 1A and 2A. Instead, MS-MSSO is cited only as a secondary reference invoked through expert testimony, and Petitioner does not identify any express disclosure in MS-MSSO that remedies the deficiencies in the primary references.

VII. ARGUMENTS AGAINST INSTITUTION

The Petition presents two main grounds of obviousness, each of which relies on a single reference to disclose the majority of claims and limitations. For Ground 1, Petitioner relies on O’Rourke alone to teach or suggest all limitations of the independent claims, and most of the limitations of the dependent claims. Pet., 2. Ground 2 takes the same tact with Middleton. *Id.* Both grounds fail to render obvious the independent claims, because neither O’Rourke nor Middleton teaches at least the determining operations or cache

policy of the independent claims. Because the Petition fails to show all of the limitations of the independent claims in the prior art, it necessarily fails to show that any of the challenged claims are likely obvious.

In an IPR, “the burden of proving a proposition of unpatentability” is borne by (and remains at all times with) the petitioner. 35 U.S.C. §§ 314(a), 316(e); *In re Magnum Oil Tools Int’l, Ltd.*, 829 F.3d 1364, 1375 (Fed. Cir. 2016). Defects in establishing a ground of unpatentability in a petition are not curable during trial since the petition cannot be corrected or enhanced during trial. *See Sirona Dental Sys. GmbH v. Institut Straumann AG*, 892 F.3d 1349, 1356 (Fed. Cir. 2018). 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*, 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*, 892 F.3d at 1356.

An obviousness determination based on a single prior-art reference requires the Petitioner to show a motivation modify the teachings in the prior art and that a POSITA would have had a reasonable expectation of success

in doing so. *See 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.

Arguments that a POSITA would be motivated and able to modify the prior art to bridge the gap between the challenged claims and what is expressly or inherently disclosed must be supported by substantial evidence. *See*

Chemours Company FC, LLC v. Daikin Industries, Ltd., No. 20-1289 (Fed.

Cir. July 22, 2021). For the reasons discussed below, Petitioner fails to meet that burden here.

A. The Petition fails to establish that O’Rourke renders obvious the claimed “when it is determined by the node that said content is cached locally, then, after determining that the content is cached locally, ... determining ... a current cache policy associated with the content” (Ground 1, all independent claims).

In Ground 1A, Petitioner’s theory rests on the flawed recasting of O’Rourke’s block 510—an “acceptable version” or freshness check—as the claimed determination of a current cache policy, based solely on its placement later in the flow. Pet., 26-29; EX2008, ¶47. Petitioner asserts that because O’Rourke’s block 510 is reached only after determining that content is cached locally, the “acceptable version” decision performed at block 510 constitutes the claimed determination of a current cache policy. Pet., 26; EX2008, ¶47. Under Petitioner’s theory, any freshness or staleness evaluation of block 510 is treated as a cache-policy determination,

disregarding the fact that O’Rourke explicitly determines cache policies *prior* to locally cached content determinations and describes block 510 *only* as a version acceptability check. EX1005, 6:12-14, 6:31-34; EX2008, ¶47. It fails for at least this reason.

Specifically, Petitioner’s theory is wrong because it conflates policy identification with cache-freshness evaluation and ignores O’Rourke’s explicit teachings for each of its blocks/steps. EX2008, ¶48. Policy identification involves determining the cache policy governing how a CDN operates. *Id.* Cache-freshness evaluation is a separate process that may or may not apply a cache policy, and evaluates cached content based on that policy. *Id.*

1. Petitioner’s reliance on Block 510 is erroneous

Petitioner’s reliance on O’Rourke’s block 510 is erroneous in two respects. EX2008, ¶49. *First*, Petitioner’s reading of the claimed “*when it is determined by the node that said content is cached locally, then, after determining that the content is cached locally: (B)(1) determining, by the node, a current cache policy associated with the content,*” is overbroad. *Id.* Petitioner effectively reads out the order of steps and cause-and-effect recited through the words “when” and “then, after.” *Second*, even with

Petitioner's overbroad reading of the claims, O'Rourke's block 510 still does not teach the cache policy. *Id.*

Block 510 determines whether the cached version of the requested content is an "acceptable version" (e.g., whether it matches the origin server's version or falls within a freshness window). EX1005, 6:34-37; EX2008, ¶50. Block 510 is a version/freshness assessment, not a determination of a content policy. EX2008, ¶50. That distinction matters because O'Rourke's "cache policies" refer to a broader set of operational behaviors (e.g., "distributing media content, handling cache misses, and prefetching media content) that govern how a CDN operates. EX1005, 2:61-62; EX2008, ¶50.

Even if the claimed "determining a current cache policy associated with requested content" is interpreted broadly to include any identification or determination of a policy relevant to the request, O'Rourke places that concept squarely at block 504, not block 510. EX2008, ¶51. That error alone is fatal to this Ground. *In-Depth Geophysical, Inc. v. ConocoPhillips Co.*, IPR2019-00850, Paper 56 at 27 (PTAB Sept. 3, 2020) (Petitioner "must 'specify where each element of the claim is found in the [relied upon] prior art patents.'" (quoting 37 C.F.R. § 42.104(b)(4))); *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1369 (Fed. Cir. 2016).

O'Rourke is clear about what blocks 504 and 510 perform.

O'Rourke's "cache policies" that include "distributing media content, handling cache misses, and prefetching media content. EX1005, 2:61-62; EX2008, ¶52. Such cache policies are identified at block 504. EX1005, 6:12-14; EX2008, ¶52. Block 510, by contrast, answers a narrower question: whether cached content remains fresh enough to serve (not which cache policy governs the request). EX2008, ¶52. O'Rourke instead teaches identification of relevant policies at block 504 ("identifies the policies related to requests for pre-recorded content"). EX1005, 6:12-14; EX2008, ¶52.

Petitioner's argument tries to impute the policy inquiry of block 504 into the freshness check of block 510. EX2008, ¶53. That move trades on a superficial similarity in terminology. *Id.* While the '173 patent's cache policies govern how a *cached* object is treated at request time, O'Rourke's block 510 merely asks "is this cached content fresh enough" to serve. *Id.* This type of evaluation already existed in conventional HTTP, which the '173 patent contemplates with mention of evaluations like TTL. EX1001, 5:55-6:3; EX2008, ¶53. But this teaching in O'Rourke does not teach or suggest the specific cache policy of the challenged claims. EX2008, ¶53.

Thus, Petitioner’s recasting of block 510 as “determining a current cache policy” is inconsistent with O’Rourke’s description of the functions at blocks 504 and 510. EX2008, ¶54. O’Rourke does not teach or suggest that its cache policy identification step (block 504) occurs “after determining that the content is cached locally.” *Id.* And block 510 does not perform the policy determination Petitioner imputes to it. *Id.*

2. Petitioner’s alternative theory modifying block 510 fares no better.

Petitioner attempts to avoid the clear discrepancy between the challenged claims and O’Rourke using alternative theories that redefine which of O’Rourke’s blocks supposedly perform the claimed determination. Pet., 12, 22-23; EX2008, ¶55. But Petitioner’s stated rationale for rearranging the logic and flow of O’Rourke—that neither function depends on the other is (1) incorrect, and (2) would not provide a reason to substitute one for the other even if it were true. EX2008, ¶55; *see Virtek Vision Int’l ULC v. Assembly Guidance Sys., Inc.*, 97 F.4th 882, 887 (Fed. Cir. 2024) (“[t]he mere fact that these possible arrangements existed in the prior art does not provide a reason that a skilled artisan would have substituted” one arrangement for the other).

For instance, Petitioner’s first theory conflating and switching the functionalities of blocks 504 and 510 fails as discussed above:

[O'Rourke's] procedure 500 transitions at line 501A from prior identification of policies related to requests for pre-recorded content (504) to application of cache policies for the requested content (e.g., block 510). ...Thus, after line 501A (which occurs after determining that the content is cached locally as recited in 1[b]/8[b]) *O'Rourke determin[es]...a current cache policy (to determine whether a cached version of the content is an acceptable version at block 510).*

Pet., 28 (emphasis added); EX2008, ¶56.

Petitioner then turns to reordering O'Rourke's disclosed steps. In the second theory, Petitioner concedes that O'Rourke's blocks 504 and 505 “do not depend on” one another, asserting that:

[i]t would have been obvious to, in response to a request for pre-recorded content, identify whether the requested content is stored on the cache server (505) and then identify the policies related to requests for pre-recorded content (504), *because neither function depends on the other function.*

Pet., 11 (emphasis added); EX2008, ¶57.

That concession is dispositive. EX2008, ¶58. It confirms that that O'Rourke does not satisfy the claimed dependency requiring that, “when it is determined [] that said content is cached locally, then, after determining that the content is cached locally,” the node determines a current cache policy. In O'Rourke, block 504 is keyed to requests (“cache server identifies the policies related to request for pre-recorded content”) and block 505 is

keyed to whether cached content is present (“cache server identifies whether the requested content is stored on the cache server”). EX1005, 6:12-14; EX2008, ¶58. Petitioner’s own premise establishes that O’Rourke’s block 504 is not conditioned on a prior determination that the content is cached locally—it undercuts the conditional dependencies and sequencing recited in the independent claims. EX2008, ¶58.

And even if O’Rourke’s blocks 504 and 505 could be interchanged (which Patent Owner does not concede), the claims still require that the subsequent determination of the expiration period be tied to the determination that the content is cached locally, not merely temporally sequenced after it. EX2008, ¶59; *see* claims 1 and 8 (“wherein determining the current cache policy comprises determining a time period for *expiration of content in the cache*”). Petitioner’s reordering may supply a temporal “after,” but it does not create the required causal link—namely a determination of an expiration period made because the content was found cached locally. EX2008, ¶59; *see* claims 1 and 8 (“when it is *determined* by the node that said *content is cached locally*, then, *after determining that the content is cached locally*”).

O’Rourke’s block 504 is never described as determining an expiration period based on content having been found in the cache. EX2008, ¶60. That

is, merely performing block 505 first does not transform block 504 into a determination that uses the “cached locally” result to determine an expiration for content in “the cache.” By conceding that blocks 504 and 505 “do not depend on each other,” Petitioner forecloses any conditional linkage. *Id.*

In summary, Petitioner’s fallback theory defeats itself. EX2008, ¶61. By asserting that O’Rourke’s blocks 504 and 505 “do not depend on each other,” Petitioner confirms that O’Rourke does not disclose the conditional linkage the claims require—namely, that after the node determines content is cached locally, the node then determines a current cache policy based on that determination. *Id.* The conceded independence between blocks 504 and 505 forecloses any suggestion that O’Rourke conditions cache-policy determination on cache presence, and no reordering of those steps can supply the dependencies the claims require. *Id.*

B. The Petition fails to establish that Middleton renders obvious the claimed “determining, by the node, a current cache policy associated with the content, wherein determining the current cache policy comprises determining a time period for expiration of content in the cache” (Ground 2, all independent claims).

Ground 2A fails at least because Middleton does not disclose or suggest the claimed “determining [] a current cache policy.” EX2008, ¶62. Middleton fails to teach or suggest this limitation for a related, but even

more straightforward reason than O’Rourke. *Id.* Rather than identifying any determination of a cache policy in Middleton, Petitioner collapses Middleton’s time-to-live (TTL) staleness check into a policy determination. *Id.* But a TTL staleness check and a policy determination are two different procedures that serve two different functions and follow two different logical structures. Petitioner misapprehends the technology. *Id.*

Petitioner’s entire theory is that Middleton’s “content’s TTL value *is a cache policy* associated with the content” and that comparing “the age of the locally cached content to a TTL value associated with the content” satisfies the determining claimed step. Pet., 67; EX1006, ¶[0080]; EX2008, ¶63. A TTL value is not a cache policy—not even as Middleton describes it. *Id.* Petitioner is simply wrong to equate these things. *Id.*

Critically, Middleton does not describe determining which cache policy is “current” (the claims require a determination of a “*current* cache policy”) because Middleton has no mechanism and no need for making that determination. EX2008, ¶64. Middleton operates in a fundamentally different way: by relying on a pre-existing TTL associated with cached content. EX1006, [0052] (“Generally, each content item has an associated time-to-live (TTL) value.”); EX2008, ¶64. Middleton uses the pre-existing TTL as an input to a staleness check. EX1006, [0052] (“check the age of the

content to determine if the content is stale”); EX2008, ¶64. That staleness check applies a pre-associated expiration value that is already bound to the cached content; Middleton does not describe determining, at request time, which expiration directive or rule is currently operative for the content. EX2008, ¶64.

The ’173 patent identifies a different architectural problem that arises when expiration behavior can change after content is cached—namely, “invalidat[ing] cached content every time a caching policy for that content is set,” an approach that is potentially expensive. EX1001, 3:64-4:3; EX2008, ¶65. The claimed determining addresses the problem by determining the current cache policy at request time, ensuring that the latest cache policy applies without requiring forced invalidations when an update is made to the cache policy. EX2008, ¶65. Middleton does not disclose any request-time mechanism for determining whether updated caching rules govern cached content and instead describes applying TTL values that are already associated with the cached content. EX1006, ¶[0052]; EX2008, ¶65.

Even if a TTL value were treated as a “policy” or part of one (which Patent Owner does not concede), Middleton still never discloses determining that policy’s *current* status. EX2008, ¶66. It does not even contemplate the kind of policy changes that would make such a status something to be

determined. *Id.* Middleton functions entirely by reading whatever TTL is already associated with the content and checking whether the content has expired. *Id.* But such a comparison against an assumed, static expiration value is not a determination of a “current cache policy” in a dynamic system. *Id.*

Petitioner’s theory again trades on a superficial similarity in terminology rather than substantive disclosure. EX2008, ¶67. In Middleton, a TTL is simply expiration metadata associated with content and used to assess staleness. EX1006, ¶¶[0052], [0080]; EX2008, ¶67. It is not identified, selected, or determined by the node as a governing cache policy at request time. EX2008, ¶67. By contrast, the ’173 patent’s claimed cache policies are rules that govern how cached content is handled when a request is received, including determining an expiration period as part of that policy. *See* EX1001, 5:55-6:3; EX2008, ¶67.

Treating Middleton’s TTL comparison as a policy determination collapses, and thereby erases, a legally meaningful distinction that the claim recites. EX2008, ¶68. It reads out an expressly recited order of steps, and an expressly recited cause-and-effect relationship between two operations. *Id.* In doing so, Petitioner improperly substitutes Middleton’s routine freshness

check for the claimed node's affirmative determination of a current cache policy. *Id.*

Put another way, Middleton's process contains no inquiry into whether the TTL in use reflects the policy currently in effect at the node, nor any selection among policies. EX2008, ¶69. Middleton's comparison succeeds regardless, without ever ascertaining whether the policy is current. *Id.* Nor does Middleton's revalidation or TTL refresh cure this deficiency. *Id.* Middleton updates TTL values only after a staleness event and only reactively, in response to "a HEAD request [] sent to a media access server (MAS)." EX1006, ¶[0082]; EX2008, ¶69. That reactive update mechanism does not involve determining what cache policy is currently in effect at the node; it simply replaces one expiration value with another after the system has already determined that content is stale. EX2008, ¶69.

In summary, Middleton describes a staleness-checking and reactive update mechanism, not the claimed determination of a current cache policy. EX2008, ¶70. The node never determines which caching policy governs the content before making its serve decision; it merely compares content age to a pre-existing expiration value and proceeds accordingly. *Id.* That is insufficient under any reasonable construction of "determining, by the node, a current cache policy." *Id.*

C. The Petition’s improper reliance on expert testimony to fill the gaps in O’Rourke and Middleton warrants denial on the merits.

The Petition improperly uses alleged knowledge of a POSITA and state of the art to modify the prior art and fill missing limitations in the Grounds. This is improper, even for a single-reference obviousness theory. The statutory basis for *inter partes* review limits invalidity grounds to prior art “consisting of patents or printed publications.” 35 U.S.C. § 311(b). The corresponding regulation requires a petition to “specify where each element of the claim is found in the prior art patents or printed publications relied upon.” 37 C.F.R. § 42.104(b)(4); *see also* 35 U.S.C. § 312(a)(3); EX2003 (Office guidance on the proper application of Rule 104(b)(4); *Henny Penny Corp. v. Frymaster LLC*, 938 F.3d 1324, 1331-32 (Fed. Cir. 2019). The Petition here fails to do this. Compounding this error, the expert testimony on which the Petition relies to fill these gaps is conclusory, and provides no analysis, evidence, or explanation.

For example, in its introduction of Ground 1A, Petitioner does not identify any disclosure in O’Rourke that performs the claimed determination in the required sequence. Instead, Petitioner’s expert rewrites O’Rourke’s disclosed order of blocks by invoking what a POSITA allegedly would have understood, relying solely on expert testimony to assert that the “information identifying functions” above line 501A could be performed in a different order. *See* Pet., 12-13 (citing EX1003, ¶¶79-80) (asserting that it would have been obvious to perform block 505

before block 504 because “neither function depends on the other”). The POSITA-driven reordering supplies the missing sequence not from O’Rourke’s disclosure, but from hindsight reconstruction.

In its analysis of limitation 1[b][i], Petitioner similarly relies not on the disclosure of O’Rourke, but on the alleged understanding of a POSITA and a purported obviousness modification to O’Rourke that is supported only by conclusory expert testimony. *See* Pet., 24 (relying only on expert testimony to support the allegation that a POSITA would have understood that O’Rourke suggests “to determine whether content is too stale, the cache policy must be able to determine when the change from acceptably stale ... to too stale ... occurs” and that it “would have been obvious to a POSITA that this change, from acceptably stale to too stale, is a time period for expiration.”).

Petitioner’s attempts to satisfy the database limitation of claims 5 and 12 similarly relies on asserted POSITA understanding rather than any express disclosure in O’Rourke. *See* Pet., 39 (asserting that “a POSITA would have understood” that storing cache policies on a different computer system “includes being stored on a database” (citing EX1003, ¶141). Although O’Rourke states only that cache plug-in modules can run on a different computer system, Petitioner contends that a POSITA would have understood that storing cache policies on such a system necessarily includes storage in a database. *Id.* Other than expert testimony

supplying that inference, Petitioner does not identify any disclosure in O'Rourke that teaches or suggests use of a database for storing cache policies.

The Petition's attempts to supply missing limitations with the *ipse dixit* of its expert falls well short of Petitioner's obligation to perform reasoned analysis—based on the references asserted in its challenges—and runs afoul of the Office guidance that compliance with Rule 104(b)(4) precludes reliance on expert testimony for claim limitations. EX2003. Accordingly, this is an additional reason the Board should deny institution on the merits.

VIII. CONCLUSION

The Petition fails to show a reasonable likelihood that any of the challenged claims is unpatentable in view of the cited art. The Board 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)
Christopher Y. Kim (Reg. No. 79,393)
Counsel for Patent Owner Sandpiper CDN, LLC

Date: April 6, 2026

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

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

1. This Patent Owner Preliminary Response complies with the type-volume limitation of 14,000 words, comprising 7,714 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
Attorney for Patent Owner Sandpiper CDN, LLC

Date: April 6, 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

2008-2019 were served in their entirety on April 6 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: April 6, 2026

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