

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

HEWLETT PACKARD ENTERPRISE CO.,
Petitioner,

v.

INTELLECTUAL VENTURES I LLC,
Patent Owner.

IPR2022-00290
Patent 6,779,082 B2

Before KEN B. BARRETT, NABEEL U. KHAN, and
STEPHEN E. BELISLE, *Administrative Patent Judges*.

KHAN, *Administrative Patent Judge*.

DECISION
Granting Institution of *Inter Partes* Review
35 U.S.C. § 314, 37 C.F.R. § 42.4

I. INTRODUCTION

On December 7, 2021, Hewlett Packard Enterprise Co. filed a Petition requesting *inter partes* review of claims 1–5 and 9–12 (the “challenged claims”) of U.S. Patent No. 6,779,082 B2 (Ex. 1001, “the ’082 patent”). Paper 2 (“Pet.”). Intellectual Ventures I LLC (“Patent Owner”) filed a Preliminary Response. Paper 7 (“Prelim. Resp.”).

We have authority to institute an *inter partes* review only if the information presented in the Petition shows “there is a reasonable likelihood that the petitioner would prevail with respect to at least 1 of the claims challenged in the petition.” 35 U.S.C. § 314(a) (2018). Upon consideration of the papers and evidence of record, we conclude that Petitioner has shown a reasonable likelihood of prevailing against at least one of the challenged claims. For the reasons explained below, we institute an *inter partes* review on the challenged claims of the ’082 patent.

II. BACKGROUND

A. *Related Proceedings*

The parties originally identified the following pending district court proceeding involving the ’082 patent: *Intellectual Ventures I et al. v. Hewlett Packard Enterprise Co.*, Case No. No. 6:21-CV-00596 (W.D. Tex.) (“parallel litigation”). Pet. 1; Paper 4, 1.

B. *Real Parties in Interest*

Petitioner identifies itself, “Hewlett Packard Enterprise Co.,” as the real party-in-interest. Pet. vii. Patent Owner identifies itself, “Intellectual Ventures I LLC,” as the real party-in-interest. Paper 4, 1.

C. *The ’082 Patent*

The ’082 patent is titled “Network-Based Disk Redundancy Storage System and Method.” Ex. 1001, code (54). The ’082 patent discloses “a

distributed and highly available data storage system” that “includes a network of data storage units that are controlled by an object management system.” *Id.* at 1:63–67. When data is written to one data storage unit, the object management system (OMS) makes a redundant copy of that data in another data storage. *Id.* at 1:67–2:3. In response to a file access request, the OMS then selects the distributed data storage unit according to the external inputs/outputs (I/O) with which the file access request is associated. *Id.* at 2:3–6.

Figure 1, reproduced below, “is a block diagram illustrating a data storage system according to an embodiment of the present invention.” Ex. 1001, 2:47–49.




Figure 1

As shown in Figure 1 above, the data storage system 100 includes a network switch 105 coupled to distributed data storage units 130a–130n and OMS managers 110a–110b. Ex. 1001, 3:22–27. Examples of the network switch may be a high-speed Ethernet switch or ATM network switch. *Id.* at

3:27–34. The application server 150 may be coupled to access the data storage system via the network switch 105. *Id.* at 3:34–36. Application programs such as voice message application programs may reside on the application server 150. *Id.* at 3:36–38. The distributed data storage units 130a–130n are the units of storage and disk redundancy. Each of the distributed data storage units has a plurality of external I/O lines for coupling to an external system, such as a public exchange (PBX) system. *Id.* at 3:39–44. The OMS of the data storage system resides on the distributed data storage units and two OMS managers 110a–110b. The OMS provides name translation, object location, and redundancy management for the system. *Id.* at 3:48–52. The OMS manager resides on two computer systems to provide high-availability and fault tolerance. *Id.* at 3:52–54.

Figure 2, reproduced below, is a block diagram illustrating the components of the distributed data storage unit 130a. Ex. 1001, 3:63–64.




Figure 2

As shown in Figure 2 above, data storage unit 130a includes a central processing unit (CPU) 202, a network interface 204 for coupling to network switch 105, a memory, a mass storage subsystem, an external I/O subsystem 210. Ex. 1001, 3:67–4:10. In one embodiment, the mass storage subsystem may include disks for storing voice mail messages and the external I/O may include one or more voice cards for communicating with a public service telephone network. *Id.* at 4:5–8. Programs and data structures stored in system memory 206 include an operating system 232, networking software 234, applications 236 related to the external I/O subsystem (e.g. an inbound voice message storage module for storing voice message in user voice mailboxes), OMS software such as a file naming service 242 and file copying service 244. *Id.* at 4:38–58.

Figure 3, reproduced below, illustrates the components of the distributed data storage unit 110a. Ex. 1001, 4:59–61.




Figure 3

As shown in Figure 3 above, OMS manager unit 110a includes a CPU 302, a network interface 304 for coupling to network switch 105, a memory 306 storage media), a mass-storage subsystem 308 (which may include a disk subsystem for storing voice mail messages), and one or more buses 312 for interconnecting the aforementioned elements of system 110a. Ex. 1001, 4:59–5:3. Programs and data structures stored in system memory 306 may include an OS 232, networking software 234, applications 236 related to the external I/O subsystem (e.g. an inbound voice message storage module for storing voice messages in user voice mailboxes), a file naming service, file copying service 244, OMS work queue 246, a unit selector module 248, an

OMS file mapping table 250, OMS file state table 252, and an OMS unit state table 254. *Id.* at 5:14–37.

The '082 patent explains that when an application such as voice message program running on application server 150 needs to create a new data file, the application sends a request to the OMS of the data storage system. Ex. 1001, 6:38–44. Preferably, the request for a new file has an association with an external 110 connection. *Id.* at 6:44–45. The request is preferably sent to the primary OMS manager unit 110a. Then, the file creation module 260 of the OMS 240 identifies and preferentially selects the distributed data storage unit that is associated with the external I/O connection, but if the data storage unit that is associated with the external I/O connection is unavailable, the OMS selects an available data storage unit. *Id.* at 6:47–52. For example, a voice message application detecting that a call is coming in, will preferentially create a new file for recording the voice stream of the call. *Id.* at 7:1–4. The request for the new file is sent to the distributed data storage unit associated with the incoming call. *Id.* at 7:4–6. In this way, the same data storage unit receiving the physical I/O stream will be used for recording the I/O stream. *Id.* at 7:6–8.

When an application retrieves a file, the application contacts OMS 240 with the name of the source file in the application. Ex. 1001, 7:66–8:2. The OMS queues the request in the OMS work queue 246. *Id.* at 8:2–3. When the OMS works through the OMS work queue 256 and finds the file retrieval request, the file retrieval module 280 of the OMS then looks up the name from the OMS file mapping table 250. *Id.* at 8:3–7. Assuming that multiple copies of the file are stored in the data storage system 100, the OMS will preferentially select a copy that is stored within the data storage unit with the most idle capacity. *Id.* at 8:7–9.

D. Illustrative Claim

Of the challenged claims, claims 1 and 9 are independent claims. Claims 2–5 depend directly from claim 1 and claims 10–12 depend directly from claim 9. Claim 1 recites:¹

[1pre] A data storage system comprising:

[1a] a network;

[1b] a plurality of distributed data storage units coupled to the network, the plurality of distributed data storage units having a plurality of external inputs and outputs; and

[1c] an object management system (OMS) manager unit coupled to the plurality of distributed data storage units via the network, the OMS manager unit and the plurality of distributed data storage unit implementing an object management system,

[1d] wherein the object management system preferentially selects a first one of the plurality of distributed data storage units for file access in response to a file access request provided that the file access request is associated with an external input/output of the first distributed data storage unit;

[1e] wherein the object management system is configured such that, in response to a file retrieval request that is associated with a data file and an external output of the first distributed data storage unit, the object management system preferentially returns a hostname and pathname of a copy of the data file that is stored within the first distributed data storage unit.

Ex. 1001, 10:25–47.

E. Evidence

The Petition relies on the following references:

¹ Reference letters in brackets mirror those provided by Petitioner. Pet. ix.

1. Coates, US 7,266,556 B1, issued Sep. 4, 2007. (“Coates”). Ex. 1004.
2. Moore et al., US 6,678,700 B1, issued Jan. 13, 2004. (“Moore”). Ex. 1005.
3. Moore, et al., “Collection-Based Long-Term Preservation,” San Diego Supercomputer Center (1999). (“Moore Publication”). Ex. 1006.

Petitioner also relies on the Declaration of Todd Mowry, Ph.D. Ex. 1003.

F. Prior Art and Asserted Grounds

Petitioner asserts that claims 1–5 and 9–12 are unpatentable on the following grounds:

Ground	Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
1	1–5, 9–12	102	Moore
2	1–5, 9–12	103	Moore ²

² Petitioner asserts grounds 2, 3, and 5 in view of the knowledge of a person of ordinary skill in the art. Pet. 2, 29, 50. We do not include the general knowledge of a person of ordinary skill in the art in listing the grounds themselves, recognizing that such knowledge is considered in every obviousness analysis. See 35 U.S.C. § 311(b) (*inter partes* review “only on the basis of prior art consisting of patents or printed publications”); *Koninklijke Philips N.V. v. Google LLC*, 948 F.3d 1330, 1337 (Fed. Cir. 2020) (“Although the prior art that can be considered in *inter partes* reviews is limited to patents and printed publications, it does not follow that we ignore the skilled artisan’s knowledge when determining whether it would have been obvious to modify the prior art. . . . Regardless of the tribunal, the inquiry into whether any ‘differences’ between the invention and the prior art would have rendered the invention obvious to a skilled artisan necessarily depends on such artisan’s knowledge.”); *Randall Mfg. v. Rea*, 733 F.3d 1355, 1362–63 (Fed. Cir. 2013) (“[T]he knowledge of [an ordinarily skilled] artisan is part of the store of public knowledge that must be consulted when considering whether a claimed invention would have been obvious.”); *Dow Jones & Co. v. Abblaise Ltd.*, 606 F.3d 1338, 1349 (Fed. Cir. 2010) (“[The obviousness] analysis requires an assessment of the . . . ‘background

Ground	Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
3	1–5, 9–12	103	Moore, Moore Publication,
4	1, 5, 9	102	Coates
5	1–5, 9–12	103	Coates

III. ANALYSIS OF THE ASSERTED GROUNDS

For the reasons below, we determine Petitioner has shown a reasonable likelihood of prevailing in its challenges to claims 1–5 and 9–12 of the '082 patent. Before we analyze the asserted grounds we first address two matters that underlie our analysis: the appropriate level of ordinary skill in the art and the constructions we will apply to claim terms.

A. Legal Standards

“In an [*inter partes* review], the petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable.” *Harmonic Inc. v. Avid Tech., Inc.*, 815 F.3d 1356, 1363 (Fed. Cir. 2016) (citing 35 U.S.C. § 312(a)(3) (requiring *inter partes* review petitions to identify “with particularity . . . the evidence that supports the grounds for the challenge to each claim”)); *see also* 37 C.F.R. § 42.104(b) (requiring a petition for *inter partes* review to identify how the challenged claim is to be construed and where each element of the claim is found in the prior art patents or printed publications relied upon).

A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference.” *Verdegaal Bros. Inc., v. Union Oil Co.*, 814 F.2d 628, 631

knowledge possessed by a person having ordinary skill in the art.” (citing *KSR, Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 401 (2007))).

(Fed. Cir. 1987); *see also Finisar Corp. v. DirecTV Group, Inc.*, 523 F.3d 1323, 1334 (Fed. Cir. 2008) (to anticipate a patent claim under 35 U.S.C. § 102, “a single prior art reference must expressly or inherently disclose each claim limitation”). Moreover, “[b]ecause the hallmark of anticipation is prior invention, the prior art reference—in order to anticipate under 35 U.S.C. § 102—must not only disclose all elements of the claim within the four corners of the document, but must also disclose those elements ‘arranged as in the claim.’” *Net MoneyIN, Inc. v. VeriSign, Inc.*, 545 F.3d 1359, 1369 (Fed. Cir. 2008) (quoting *Connell v. Sears, Roebuck & Co.*, 722 F.2d 1542, 1548 (Fed. Cir. 1983)). Whether a reference anticipates is assessed from the perspective of one of ordinary skill in the art. *See Dayco Prods., Inc. v. Total Containment, Inc.*, 329 F.3d 1358, 1368–69 (Fed. Cir. 2003) (“‘[T]he dispositive question regarding anticipation [i]s whether *one skilled in the art* would reasonably understand or infer from the [prior art reference’s] teaching’ that every claim element was disclosed in that single reference.” (second and third alterations in original) (quoting *In re Baxter Travenol Labs.*, 952 F.2d 388, 390 (Fed. Cir. 1991))).

A claim is unpatentable as obvious under 35 U.S.C. § 103(a) if the differences between the claimed subject matter and the prior art are such that the subject matter, as a whole, would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including: (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of skill in the art; and (4) where in evidence, so-called “secondary considerations,” including commercial success, long-felt but

unsolved needs, failure of others, and unexpected results. *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).

B. Level of Ordinary Skill in the Art

The level of ordinary skill in the pertinent art at the time of the invention is relevant to how we construe the patent claims. *See Phillips v. AWH Corp.*, 415 F.3d 1303, 1312–13 (Fed. Cir. 2005) (en banc). It is also one of the factual considerations relevant to obviousness. *See Graham*, 383 U.S. at 17–18.

To assess the level of ordinary skill, we construct a hypothetical “person of ordinary skill in the art,” from whose vantage point we assess obviousness and claim interpretation. *See In re Rouffet*, 149 F.3d 1350, 1357 (Fed. Cir. 1998). This legal construct “presumes that all prior art references in the field of the invention are available to this hypothetical skilled artisan.” *Id.* (citing *In re Carlson*, 983 F.2d 1032, 1038 (Fed. Cir. 1993)).

Petitioner contends a person of ordinary skill in the art (POSITA) “would have had at least a bachelor’s degree in electrical engineering, computer science, or an equivalent, plus three years of experience in the field of storage system design and architecture.” Pet. 16. Petitioner further contends that “[a]dditional education in the fields of electrical engineering or computer science such as a masters or doctorate degree may serve as a substitute for experience in the field . . . a person with less education but more relevant practical experience may also meet this standard. *Id.* (citing Ex. 1003 ¶ 26).

Patent Owner does not submit a proposed level of ordinary skill in the art in its Preliminary Response. *See generally* Prelim. Resp. For purposes of this Decision we adopt Petitioner’s proposed level of ordinary skill.

C. Claim Construction

In this *inter partes* review, the Board applies the same claim construction standard as that applied in federal courts. *See* 37 C.F.R. § 42.100(b). Under this standard, claim terms “are generally given their ordinary and customary meaning” as understood by a person of ordinary skill in the art in question at the time of the invention. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312–13 (Fed. Cir. 2005) (en banc) (citations omitted). The ordinary and customary meaning of a claim term “is its meaning to the ordinary artisan after reading the entire patent,” and “as of the effective filing date of the patent application.” *Phillips*, 415 F.3d at 1313, 1321. “In determining the meaning of the disputed claim limitation, we look principally to the intrinsic evidence of record, examining the claim language itself, the written description, and the prosecution history, if in evidence.” *DePuy Spine, Inc. v. Medtronic Sofamor Danek, Inc.*, 469 F.3d 1005, 1014 (Fed. Cir. 2006) (citing *Phillips*, 415 F.3d at 1312–17). Extrinsic evidence is “less significant than the intrinsic record in determining ‘the legally operative meaning of claim language.’” *Phillips*, 415 F.3d at 1317 (citation omitted).

Petitioner asserts that in the parallel litigation it “has proposed to construe ‘commodity disk server’ in Challenged Claim 5 to mean ‘non-RAID server.’” Pet. 16. Petitioner argues, however, that “the Board need not address this construction to resolve the issues in this Petition because, as discussed below, the prior art in each ground discloses non-RAID servers.” *Id.* at 16–17. Petitioner contends that it interprets the remaining claim terms in accordance with the ordinary and customary meaning as understood by one of ordinary skill in the art. *Id.* at 17.

Patent Owner disagrees with Petitioner's proposed construction of "commodity disk server" proffered in the parallel litigation but "agrees that construction of this term is not necessary at this stage in the proceeding." Prelim. Resp. 7. As for the remaining terms Patent Owner argues "no express construction is necessary to resolve any controversy. *Id.*

We determine no terms need to be construed to resolve the disputes between the parties at this stage of the proceeding.

D. Description of Prior Art References

Petitioner's challenges rely on Moore, the Moore Publication, Coates, and knowledge of a POSITA. *See Pet., passim.* We provide a brief overview of Moore, the Moore Publication, and Coates below.

1. Moore (Ex. 1005)

Moore is directed to "maintaining data objects in containers across a network of containers across a network of distributed heterogeneous resources in a manner which is transparent to a client." Ex. 1005, code (57). Moore explains that data objects, such as drawings, images, and e-mail, may be stored and distributed across heterogeneous network resources such as database management systems, archival storage systems, and file systems. *Id.* at 1:23–40. The management of data stored across heterogeneous systems may be complicated due to the diverse nature of these heterogeneous systems. *Id.* at 1:40–43. An objective of Moore is, therefore, to provide a unified framework for accessing the data objects from these numerous and diverse sources. *Id.* Other objectives are to provide a system for managing data objects distributed across heterogeneous resources in a way the reduces or eliminates the delay or latency characteristics of conventional systems and in a way in which the physical location of the stored data is transparent to the user. *Id.* at 3:37–45.

Moore discloses a Storage Resource Broker (SRB) Architecture in which its system may be deployed. Ex. 1005, 6:13–15. A simplified view of the SRB is depicted in Figure 1 reproduced below:




FIG. 1

Figure 1 above depicts a simplified implementation of the SRB architecture with an SRB client 1, one or more SRB servers 3, which are accessible by the client over network 2, a meta-data catalog (MCAT) 7, and distributed system resources such as database management systems 4, archival storage systems 5, and file systems 6 which are accessible by one or more SRB servers. Ex. 1005, 6:13–29. The MCAT stores meta-data associated with data sets, users, and resources managed by the SRB system. *Id.* at 6:30–31. Client applications are provided with a set of APIs for sending requests and receiving response to/from the SRB servers 3 which are responsible for carrying out the client requests. *Id.* at 6:36–39. These tasks include interacting with the MCAT 7, and performing input/output (I/O) on behalf of the clients. *Id.* at 6:39–41. A client uses the same APIs to

access every storage system resource managed by the SRB System. *Id.*

The complex tasks of interacting with various types of storage system resources, operating system (OS) and hardware architectures, are handled by the SRB servers 3. *Id.* at 6:42–45.

Moore describes an implementation example in which containers are used to store many small files in an archival storage system. Ex. 1005, 16:60–67. Moore explains that a “container” is used to aggregate data sets in much the same way as a folder can be used to aggregate paper documents. *Id.* at 10:27–30. Data objects stored in a container have meta-data kept in the MCAT for translating the in-container object to file path, offset and size. *Id.* at 17:1–5. An example of meta-data is illustrated in Figure 13, reproduced below.

RESOURCE	RNUM	PHY RSRC_NAME	RSRC TYPE	LOCN DESC	RSRC CLASS
test-cont	0	cont-hpss-caltech	hpss file system	johnny caltech	archival
	1	cont-unix-wisc	unix file system	vulture wisconsin- madison	cache
	2	cont-cache-maryland	unix file system	ape01 umiacs	cache
	3	cont-pri-hpss-sdsc	hpss file system	mds-18 sdsc	primary archival
	4	cont-pri-unix-sdsc	unix file system	ghidotah sdsc	primary cache

FIG. 13

Figure 13 above depicts the meta-data that describes the physical resources associated with the logical resource assigned to a container. Ex.

1005, 17:11–13. Item 101 is the name of the logical resource and items 102–106 describe each of the physical resources assigned to the logical resource. *Id.* at 17:13–16. For each physical resource, item 102 is a consecutive number identifying the resource; item 103 is the name of the resource; item 105 is the location of the resource (host address and the path for file system type resource); item 104 is the type of resource (UNIX file system, HPSS file system and DB large object, etc.) and item 106 is the class of the resource. *Id.* at 17:17–23.

Moore explains that when a client wishes to store objects in a container, a three-step process is performed. Ex. 1005, 17:55–57. In the first step, a container is created. *Id.* at 17:57–58. In the second step, a cache copy of the container is located, and the location in the cache copy where the object is to be stored is determined. In the third step, the object is written into the cache copy at the specified location. *Id.* at 17:58–61.

When a client wants to read objects from a container, a client can use the `srbbObjOpen()` call to open objects for I/O operations such as read and write. Ex. 1005, 20:34–38. The client does not need to be aware of whether an object is stored in a container because the information is already in MCAT. *Id.* at 20:38–40. The `srbbObjOpen()` call specifies the object name to be opened and returns an SRB file descriptor specifying the offset where the object is stored in a cache copy of the container. *Id.* at 20:40–42, 46–48. The client then issues an `srbbObjRead()` call specifying the SRB file descriptor which returns a copy of the object to the client. *Id.* at 20:42–44.

2. *Moore Publication (Exhibit 1006)*

“Collection-Based Long-Term Preservation” (the Moore Publication) is a technical report by Reagan Moore *et al.* The Moore Publication identifies the problem storage communities face when trying to store large

amounts of data for very long periods of time. Ex. 1006, 5³. The Moore Publication explains that the challenge faced is to maintain discoverability and access of data stored in supercomputer centers, digital libraries and archival storage communities while the technology used to manage the archive evolves or becomes obsolete. *Id.* In the Moore Publication, the authors present “a scalable architecture for managing media migration, and an information model for managing migration of the structure of the context” that permits the data to be interpreted. *Id.*

In order to manage heterogeneity of data resources, the Moore Publication implements the SRB architecture discussed in Moore and provides additional details regarding the implementation of this architecture. Ex. 1006, 12–13.

a) Prior Art Status of the Moore Publication

Petitioner argues the Moore Publication was published and publicly accessible by November 11, 2000. Pet. 22 (citing Ex. 1013 ¶¶ 78–83). Petitioner supports its argument with the declaration of Sylvia D. Hall-Ellis, Ph.D. (Exhibit 1013). Dr. Hall-Ellis explains that the Machine-Readable Cataloging (MARC) standard is an industry-wide standard method of storing and organizing library catalog information. Ex. 1013 ¶¶ 23–25. Dr. Hall-Ellis testifies that a MARC record may be viewed online using Online Computer Library Center’s (OCLC) bibliographic database. *Id.* ¶ 27. When an OCLC member institution acquires a work, it creates a MARC record for this work in its computer catalog system in the ordinary course of its business. *Id.* ¶ 30. This record is then made available to any other OCLC

³ Citations to the Moore Publication (Exhibit 1006) refer to the page numbers added by Petitioner to the document rather than the original page numbers.

member and is therefore made available to the public. *Id.* A MARC record comprises several fields including field 040 which identifies the library that created the MARC record and field 005, which identifies the date the MARC record was created in year, month, day format (YYYYMMDD). *Id.* ¶¶ 27, 31.

Dr. Hall-Ellis testifies that the MARC record for the Moore Publication was created by the National Science Library/Bibliothèque Scientifique Nationale (National Science Library) and that the “Entered” field shows that a cataloger at the National Science Library created the OCLC record on November 11, 2000. Ex. 1013 ¶ 80, Attachment 6. Dr. Hall-Ellis testifies that based on the information in the MARC record the Moore Publication was publicly available on or before November 11, 2000, through any library with access to the OCLC bibliographic database or through an individual library. *Id.* ¶ 82.

Patent Owner argues that Petitioner has not established that the Moore Publication is prior art. Prelim. Resp. 25–28. Patent Owner argues that the Moore Publication itself does not have any information regarding when it was publicly accessible. Prelim. Resp. 25. Primarily though, Patent Owner takes issue with statements made in the Hall-Ellis declaration. For example, Patent Owner argues that the “MARC record presented in Attachment 6 of the Hall-Ellis declaration does not show a clear label of ‘entered’ or ‘20001111.’” *Id.* at 26. To illustrate its point, Patent Owner submits an annotated figure from Attachment 6 of the Hall-Ellis declaration in its Preliminary Response which we have reproduced below:

OCLC 45324187 No holdings in XXX - 2 other holdings

Books		Re En 200 Re 2021062	
c tere 011 place 5202213			
Stat d 11 d .6			
Type a	ELvl 7	Srcce d	Audn
BLvl m	Form	Conf 0	Biog
Cont t	GPub f	LitF 0	MRec
Desc	Ills	Fest 0	Indx
		DtSt a	Dates 1999
			Lang eng
			Ctry

EX1013, 120 (Attachment 6) (annotated).

The figure above from Patent Owner’s Preliminary Response shows a portion of Attachment 6 from the Hall-Ellis declaration which Patent Owner has annotated to show that the labels “entered” and “20001111” are not clear. *Id.* Next, Patent Owner argues that that MARC record at the official website of the National Science Library/Bibliothèque Scientifique Nationale (the “CAI Library”) does not show the alleged publicly accessible date of November 11, 2000 for the Moore Publication” and that “[t]here is no “Entered” field (“20001111”) in the MARC record provided by the official website of the CAI Library.” *Id.* at 27 (citing Ex. 2001). Finally, Patent Owner argues Petitioner does not demonstrate that the Moore Publication downloaded from the Defense Technical Information Center (DTIC) website and presented as Exhibit 1006 is identical to the cataloged report at the CAI library. *Id.* at 27.

At this stage of the proceeding we are persuaded by Petitioner’s arguments and evidence. We determine that Petitioner has identified sufficient evidence establishing a reasonable likelihood that the Moore Publication was publicly accessible before the critical date of the challenged patent. Specifically, Petitioner provides credible un rebutted testimony from Dr. Hall-Ellis that a MARC record was entered on November 11, 2000 signifying that the Moore Publication was acquired by the National Science Library and made publicly available. Ex. 1013 ¶¶ 78–83. We disagree with

Patent Owner that a screen shot of the MARC record in Attachment 6 of the Hall-Ellis declaration is unclear regarding its “entered” and “20001111” label. Attachment 6, depicted above, clearly shows the words “Entered” and the number “20001111” next to it, albeit broken up over three lines instead of one. Furthermore, the MARC record provided as Attachment 6 to Exhibit 1006 describes the Moore Publication as a 138-page document titled “Collection-Based Long-Term Preservation” and with a description that is consistent with the Exhibit 1006 relied upon by Petitioner, indicating that the MARC record does, in fact, correspond with the copy of the Moore Publication downloaded from the DTIC website. *See* Ex. 1006 (a 138-page document titled “Collection-Based Long-Term Preservation”).

3. *Coates (Exhibit 1004)*

Coates is directed to a distributed storage system that include multiple intelligent storage nodes arranged in one or more storage centers. Ex. 1006, 3:3–4. The intelligent storage nodes, which store the files for the distributed storage system, are combined with one or more distributed object storage managers (DOSMs). *Id.* at 3:4–7. A network couples the DOSMs to the intelligent storage nodes. *Id.* at 3:7–8. The DOSMs manage requests from clients, such as content servers and end-user computers, to download files from a storage center. *Id.* at 3:8–11. The distributed storage system is arranged to support uninterrupted delivery of files in the event of a failure. *Id.* at 3:12–14. The failover architecture includes storing, for each file, a duplicate file in a different intelligent storage node. *Id.* at 3:14–16. In the event the distributed storage system enters a failover condition, the DOSM determines, for files stored in the failed intelligent storage node, a location (i.e., intelligent storage node) for the duplicate files. *Id.* at 3:16–20. The intelligent storage nodes, which store duplicate files, are located in different

storage centers. *Id.* at 3:30–32. The storage centers are located in different geographic areas (e.g., west coast and east coast). *Id.* at 3:32–34. The storage centers are mirrored so that the files stored in one intelligent storage node in the first storage center are the same as the files stored in one intelligent storage node in the second storage center. *Id.* at 3:34–38.

Figure 3, reproduced below, illustrates an embodiment of a storage cluster.




Figure 3

As shown in Figure 3 above, the storage cluster consists of DOSMs 320 and intelligent storage nodes 340. Ex. 1004, 7:3–5. The storage cluster 300 receives upload, download, and delete operations that include the storage resource location (SRL). *Id.* at 6:67–7:2. The SRL is then used to uniquely identify a client file. *Id.* at 7:2–3. The SRL has several fields, including the storage cluster field which includes the name or IP address of a storage center and an arbitrary-customer-uri field which may include the

filename and extension of the file being downloaded. *Id.* at 7:50–8:3. The load balancing fabric 310 distributes upload and download requests to one of a plurality of DOSMs based on DOSM availability. *Id.* at 7:18–20. The DOSM handles hundreds of simultaneous download transactions. *Id.* at 7:23–24. The DOSMs 320 communicate with the intelligent storage nodes 340 via an interconnect fabric 330. *Id.* at 7:34–37. The interconnect fabric 330 consists of a high-speed, high bandwidth fabric to ensure that all the DOSMs 320 communicate with every intelligent storage node at all times. *Id.* at 7:34–37.

Coates explains that when the storage cluster receives a download request, including the unique file identifier in the SRL, the load balancing fabric will select an available DOSM which will determine whether the requested object resides in its cache. Ex. 1004, 9:16–33. If the object is cached, then the object is transmitted to the recipient, but if not, then the DOSM attempts to identify the location of the object in one of the intelligent storage nodes using a lookup table. *Id.* at 9:36–45. Table 610 of figure 6, reproduced below, illustrates a DOSM file lookup table.

DOSM File Look-up Table
610

File Id	IP Addr	Disk Id
file1.MD5	10.3.100.1	3
file2.MD5	10.3.098.1	1
file3.MD5	10.3.050.1	6
file4.MD5	10.3.100.1	2
file5.MD5	10.3.098.1	1
file6.MD5	10.3.050.1	8

Each entry in table 610 above, includes a file identification of the object file, an IP address of the intelligent storage node that stores the object file, and a disk ID which stores an integer value that identifies the specific

disk drive on the intelligent storage node that stores the object file. Ex. 1004, 11:12–27.

E. Anticipation over Moore (Ground 1), Obviousness over Moore (Ground 2) and Obviousness over Moore and the Moore Publication (Ground 3)

Petitioner argues that claims 1–5 and 9–12 of the '082 patent are anticipated by Moore (Ground 1), would have been obvious over Moore (Ground 2), and would have been obvious over Moore and the Moore Publication (Ground 3). Pet. 2, 29–49. Based on our review of the Petition and Preliminary Response, we determine that Petitioner has not demonstrated a reasonable likelihood that Moore anticipates the challenged claims. Specifically, as explained further below in our analysis, we determine that Petitioner has not sufficiently shown Moore discloses limitation 1c which requires an object management system (OMS) manager unit to be coupled to the plurality of distributed data storage units via a network. Petitioner also does not sufficiently explain how this limitation would have been obvious over Moore in light of the knowledge of a POSITA. We do, however, determine that Petitioner has sufficiently shown limitation 1c is taught by the combination of Moore and the Moore Publication. We, therefore, determine Petitioner has demonstrated a reasonable likelihood of prevailing in showing that claims 1–5 and 9–12 are unpatentable as obvious over Moore and the Moore Publication.

1. Claim 1

a) [1pre] A data storage system comprising: [1a] a network [1b] a plurality of distributed data storage units coupled to the network, the plurality of distributed data storage units having a plurality of external inputs and outputs; and

Petitioner presents evidence that Moore discloses the preamble and limitations 1a and 1b. Pet. 29–33. Patent Owner does not separately dispute

Petitioner's contentions regarding the preamble and limitations 1a and 1b. *See* Prelim. Resp. 13–24. At this stage of the proceeding we are persuaded by Petitioner's contentions which we find are supported by the cited evidence.

For example, Petitioner identifies Moore's Storage Resource Broker (SRB) system as the recited data storage system. Pet. 29 (citing Ex. 1003 ¶ 82). Petitioner argues that "[t]he SRB system 'provides distributed clients with uniform access to diverse storage resources in a heterogeneous computing environment.'" *Id.* (citing Ex 1005, 6:13–17, 6:24–29). We note that because we are persuaded that Coates discloses the preamble we need not decide whether the preamble limits the claim.

Petitioner further argues that Moore's SRB system comprises a network as shown in Figure 14C of Moore, which depicts SRB clients connected via a network to SRB servers. Pet. 30 (citing Ex. 1005, 1:15–20, 6:18–24, 6:36–45, Figs. 1, 14C).

Petitioner identifies Moore's "'physical resources,' including known off-the-shelf database management systems, archival storage systems, and file systems" as the recited distributed data storage units. Pet. 31 (citing Ex. 1005, 6:24–29, Figs. 4, 14c; 1003 ¶¶ 85–87). Petitioner argues that each distributed data storage unit has physical external input/output to a driver interface to receive and respond to requests from SRB clients. *Id.* (citing Ex. 1005, 8:58–63). Finally, Petitioner argues the physical resources are coupled to the network via SRB servers. *Id.* at 32 (citing Ex. 1005, code (57), 10:28–30, 16:20–24, 16:60–67, 9:65–67; Ex. 1003 ¶ 87).

For the aforementioned reasons, at this stage of the proceeding we are persuaded that Moore discloses the preamble and limitations 1a–1b.

- b) *[1c] an object management system (OMS) manager unit coupled to the plurality of distributed data storage units via the network, the OMS manager unit and the plurality of distributed data storage unit implementing an object management system*

Petitioner identifies Moore's SRB agent, which includes dispatcher 10, high-level request handler 11, and low-level request handler 12, and the MCAT as the recited OMS. Pet. 33 (citing Ex. 1003 ¶ 89). Petitioner argues that the SRB agent handles each connection and receives and services all client requests, including returning results to the client. Pet. 34 (citing Ex. 1005, 7:47–59, 8:8–19, 8:33–40, 14:37–49, Figs. 15a–15b).

Petitioner argues that Moore's SRB agent is coupled to the distributed data storage units (Moore's physical resources) via a network. Pet. 35–36. For example, Petitioner argues that communications from the SRB to the distributed data storage units are translated to internet protocol communications, such as HTTP which shows that the SRB agent is coupled to the distributed data storage units via a network. Pet. 35–36 (citing 8:48–9:9, 9:65–67). Petitioner further argues that Moore discloses that the distributed resources are accessible by SRB servers over File Transfer Protocol (FTP), which, according to Petitioner, further demonstrates that the storage units are connected to the SRB server via a network. Pet. 35–36 (citing Ex. 1005, 6:24–29).

Petitioner argues, in the alternative that if Moore does not disclose that the OMS manager is coupled to the distributed data storage unit via a network, that it would have been obvious over the combination of Moore, the Moore Publication and knowledge of a POSITA. Pet. 36–37. Specifically, Petitioner argues that the Moore Publication describes an exemplary distributed data storage unit—the IBM High Performance Storage System (HPSS)—that is accessed over high-speed networks and that

is connected to the SRB servers via a network. Pet. 36–37 (citing Ex. 1006, 15, 29, Fig. 4–1; Ex. 1003 ¶ 92). Petitioner argues that it would have been obvious for a person of ordinary skill in the art “to combine Moore with the Moore Publication because the Moore Publication describes additional details of the exact same system (the SRB system), working predictably.” Pet. 37 (Ex. 1003 ¶ 92).

Patent Owner disputes Petitioner’s contention that Moore’s SRB agent is coupled to the distributed resources via a network. Prelim. Resp. 21–25. Specifically, Patent Owner argues that Moore’s SRB server only communicates with a local storage unit and that if a requested storage system is not local, the low-level handler dispatches a call to the appropriate SRB agent which then communicates with its local storage unit. Because only local communication is disclosed by Moore, Patent Owner argues that the SRB agent is not connected to the storage units via a network. *Id.* at 22–23 (citing Ex. 1005, 7:1–9, 7:25–42, 7:44–52, 8:58–9:3, Fig. 3). Indeed, according to Patent Owner even Moore’s disclosure of FTP and HTTP drivers are directed to local I/O operations. *Id.* at 24.

At this stage of proceeding we are unpersuaded by Petitioner’s arguments that Moore discloses that its SRB agent is connected to its storage units via a network. The evidence cited by Petitioner does not affirmatively disclose such a network connecting the SRB servers and the storage units. References to certain protocols, such as HTTP and FTP, are insufficient to overcome Patent Owner’s argument that these disclosures are instead directed to local input operations.

We determine, however, that Petitioner has presented sufficient evidence that the Moore Publication teaches a storage unit, the HPSS, that is connected to the SRB servers via a network. Pet. 36–37. We also determine

Petitioner has provided a rational reason for combining Moore with the Moore Publication because the Moore Publication describes additional details of the HPSS already disclosed in Moore. At this stage, Patent Owner does not dispute Petitioner's contentions regarding the Moore Publication other than to contend that the Moore Publication is not prior art. *See* Prelim. Resp. 25–28. In this regard, as we explained above, we determine Petitioner has presented sufficient evidence at this stage of the proceeding demonstrating that the Moore Publication was publicly available before the critical date of the '082 patent. *See supra* § III.D.2. Accordingly, at this stage of the proceeding we are persuaded that the combination of Moore and the Moore Publication teaches limitation 1c.

- c) *[1d] wherein the object management system preferentially selects a first one of the plurality of distributed data storage units for file access in response to a file access request provided that the file access request is associated with an external input/output of the first distributed data storage unit*

Petitioner identifies Moore's "srbObjOpen() call issued by a client, which opens an object (file) in a container" as the recited file access request. Pet. 37–38 (citing Ex. 1005, 20:34–21:30, Figs. 9–10; Ex. 1003 ¶ 95). Petitioner argues the srbObjOpen() call is associated with an external input/output of a specific distributed data storage unit because Moore sets up an association between a file access request and a preferred external input/output of a distributed data storage unit when it creates containers for objects stored in the SRB system. Pet. 38 (citing Ex. 1003 ¶ 96). According to Petitioner, meta-data in Figure 13 of Moore, such as fields 103, 104 and 106, associates a container with a specific physical resource and that the type of resource invokes the appropriate low-level storage driver to handle the request to the named resource and thus identifies the resource's external

input/output connected to that driver. Pet. 38 (citing Ex. 1005, 7:30–33, 1003 ¶ 97). Petitioner contends that field 106, in particular, specifies whether a physical resource is a primary cache for that container and that the primary cache for a particular client is determined by physical proximity. Pet. 39 (citing Ex. 1005, 12:17–32).

Petitioner further argues that Moore’s OMS preferentially selects the primary distributed data storage unit for file access in response to receiving the file access request associated with that distributed data storage unit. Pet. 39 (citing 1003 ¶ 99). For example, Petitioner argues that in response to the `srbObjOpen()` call, the SRB system preferentially selects the resource designated as the primary cache in the container to handle the request but that if the primary cache is not available, other resources assigned to the same container can be used to fulfill the request. *Id.* at 39–40 (citing Ex. 1005, 19:16–18; 20:58–67).

Patent Owner disputes Petitioner’s contention that Moore teaches a file access request that is associated with a first distributed data storage unit. Prelim. Resp. 14–17. Specifically, Patent Owner argues that Moore’s `srbObjOpen()` call only specifies the object name to be opened and is not associated with a data storage unit because the client does not need to be aware of whether an object is stored in a container. Prelim. Resp. 14–15 (citing Ex. 1005, 20:35–47, Fig. 9). Patent Owner adds that Moore’s client is not even aware of the association between the file descriptor and the storage node. *Id.* at 15 (citing Ex. 1005, 10:30–39, Fig. 9). According to Patent Owner, the association of the container and its meta-data is established before the open object request is issued. *Id.* at 15 (citing Ex. 1005, 14:11–16, 10:30–44). Patent Owner further argues that Moore’s container creation request is also not associated with a data storage unit

because the association of Moore's container and its physical resources is not created according to the request of the client and the container creation request initiated by the client is not associated with any physical resources. Prelim. Resp. 16 (citing Ex. 1005, 14:37–49, 14:54–59, 17:62–65, Figs. 5, 16A).

Patent Owner also disputes that Moore preferentially selects a data storage unit in response to a file access request. Prelim. Resp. 17–21. Patent Owner argues that Moore reads from the cache by default without preferentially selecting from other resources assigned to the same container in response to the file access request. Prelim. Resp. 17–18 (citing Ex. 1005, 14:37–49, 18:29–36, Figs. 5, 16A). Patent Owner argues that if a copy of the container is not available in cache then Moore stages an archived copy of the container in cache then retrieves data from the cached copy and through this process always reads from a cached copy of the container. *Id.* at 18 (citing Ex. 1005, 10:30–44, 14:16–22, Fig. 16D). This is true, according to Patent Owner, even in the situation raised by Petitioner regarding Moore determining the primary cache based on physical proximity. *Id.* at 21 (citing Ex. 1005, 12:16–38).

At this stage of the proceeding, we are persuaded by Petitioner's arguments that Moore discloses limitation 1d. Specifically, we are persuaded that the `srbObjOpen()` request of Moore is associated with an external input/output of Moore's distributed data storage unit because the call identifies the requested data object which is then used to query the MCAT for the meta-data associated with the object which then identifies the physical resources in which the object is stored. Pet. 38 (citing Ex. 1005, 7:30–33, 1003 ¶ 97). We are also persuaded that Moore preferentially selects the primary cache from amongst resources and that these other

resources may be used if the primary resource is unavailable. Ex. 1005, 20:58–67. Petitioner’s arguments are supported by credible, un rebutted testimony from Dr. Mowry. Ex. 1003 ¶¶ 95–99.

We disagree with Patent Owner’s arguments regarding the client being unaware of the association between the file descriptor and the storage node because the claim does not require the client to be aware of the association between the request and the external input/output of the distributed data storage unit. We also disagree with Patent Owner’s argument that Moore does not preferentially select a storage unit because it stages files in a cache. As explained above, Moore discloses that it stores data objects in multiple physical resources and that it may access these resources under certain conditions. Ex. 1005, 20:58–67.

d) [1e] wherein the object management system is configured such that, in response to a file retrieval request that is associated with a data file and an external output of the first distributed data storage unit, the object management system preferentially returns a hostname and pathname of a copy of the data file that is stored within the first distributed data storage unit

Petitioner argues that the `srbObjOpen()` and `srbObjRead()` calls together disclose the claimed file retrieval request. Pet. 40 (citing Ex. 1005, 20:35–48, 21:7–15, Figs. 9–10; Ex. 1003 ¶ 101). Petitioner argues that the `srbObjOpen()` and `srbObjRead()` calls are associated with a data file and an external output of the first distributed data storage unit. *Id.* (citing Ex. 1003 ¶ 102). Petitioner argues that in response to `srbObjOpen()` request, the OMS preferentially returns a hostname and pathname of the data file in the form of a descriptor specifying the location from which to read the file. *Id.* at 41 (citing Ex. 1005, 20:58–67, Fig. 9).

Patent Owner does not dispute Petitioner's arguments regarding limitation 1e. For the aforementioned reasons, at this stage of the proceeding we are persuaded that Moore discloses limitation 1e.

e) Conclusion as to Claim 1

For the reasons explained above and based on the arguments and evidence presented in the Petition, we determine that Petitioner has not shown claim 1 is anticipated by Moore or would have been obvious over Moore. We determine that Petitioner has shown sufficiently that the combination of Moore and the Moore Publication would have taught the limitations of claim 1 and that a POSITA would have been motivated to combine the Moore Publication with the teachings of Moore. We, therefore, determine that, at this stage of the proceedings, Petitioner has demonstrated a reasonable likelihood that claim 1 would have been obvious over the combination of Moore and the Moore Publication.

2. Claims 2–5 and 9–12

Petitioner contends claims 2–5 and 9–12 of the '082 patent are unpatentable as anticipated over Moore (Ground 1), as obvious over Moore (Ground 2), and as obvious over Moore and the Moore Publication (Ground 3). Petitioner provides detailed analysis and evidence that Moore discloses the limitations of each of these claims. Pet. 42–49. Patent Owner relies on the same or similar arguments here as the ones discussed above regarding claim 1.

For the reasons explained above with regard to claim 1, we determine Petitioner does not present sufficient evidence to establish that Moore anticipates claims 2–5 and 9–12 (Ground 1) or that the same claims would have been obvious over Moore (Ground 2). We do, however, determine that Petitioner presents sufficient evidence to establish a reasonable likelihood

that Moore and the Moore Publication teach the limitations of claims 2–5 and 9–12 and that a POSITA would have been motivated to combine the Moore Publication with the teachings of Moore. We, therefore, determine that, at this stage of the proceedings, Petitioner has demonstrated a reasonable likelihood that claims 2–5 and 9–12 would have been obvious over the combination of Moore and the Moore Publication (Ground 3).

F. Anticipation by Coates (Ground 4) and Obviousness over Coates (Ground 5)

Petitioner argues that claims 1, 5, and 9 of the '082 patent are anticipated by Coates (Ground 4) and claims 1–5 and 9–12 would have been obvious over Coates (Ground 5). Pet. 2, 50–74. Based on our review of the Petition and Preliminary Response, we determine that Petitioner has demonstrated a reasonable likelihood that Coates anticipates claims 1, 5, and 9 and that claims 1–5 and 9–12 would have been obvious over Coates.

1. Claim 1

a) [1pre] A data storage system comprising: [1a] a network [1b] a plurality of distributed data storage units coupled to the network, the plurality of distributed data storage units having a plurality of external inputs and outputs; and [1c] an object management system (OMS) manager unit coupled to the plurality of distributed data storage units via the network, the OMS manager unit and the plurality of distributed data storage unit implementing an object management system

Petitioner presents evidence that Coates discloses the preamble and limitations 1a–1c. Pet. 50–57. Patent Owner does not separately dispute Petitioner's contentions regarding the preamble and limitations 1a–1c. *See* Prelim. Resp. 28–45. At this stage of the proceeding we are persuaded by Petitioner's contentions which we find are supported by the cited evidence.

For example, Petitioner argues Coates discloses the preamble by disclosing a data storage system of intelligent storage nodes that are

arranged in storage clusters across a network. Pet. 50 (citing Ex. 1004, 3:3–14, 3:31–39, 19:5–21:39; Ex. 1003 ¶ 124). Petitioner contends the storage clusters are geographically replicated to provide load balancing and disaster protection. Pet. 50 (citing Ex. 1004, 19:13–22, 19:27–42). We note that because we are persuaded that Coates discloses the preamble we need not decide whether the preamble limits the claim.

Furthermore, Petitioner argues Coates discloses the recited “network” because it repeatedly describes its system as a network storage system. Pet. 51 (citing Ex. 1004, code (57), 5:11–37). According to Petitioner, the geographically separate storage clusters within Coates’ system are coupled to a wide area network. *Id.* at 51 (citing Ex. 1004, 19:33–35). Moreover, a storage cluster also has a network connection within the cluster. Pet. 51 (citing Ex. 1004, 3:4–8, 7:33–44, 10:16–22, 12:7–10; Ex. 1003 ¶ 127).

Petitioner argues that the storage nodes 340 disclose the recited “distributed data storage units and that these storage nodes are coupled to a network, specifically interconnect fabric 330, which can be a public access network such as the Internet.” Pet. 53 (citing Ex. 1004, 3:4–8, 7:33–44, 10:16–22, 12:7–10; Ex. 1003 ¶ 129). Additionally, Petitioner contends that Coates’s storage nodes 340 have a plurality of external inputs and outputs in the form of a physical network interface used to receive inputs from and send outputs to the rest of the storage system. *Id.* (citing Ex. 1004, 11:63–12:6, 12:7–10; Ex. 1003 ¶ 130).

Petitioner identifies the virtual file server (VFS), the traffic management service, and the DOSMs together as the recited OMS manager. Pet. 54 (citing Ex. 1003 ¶¶ 131–136). Petitioner argues that the VFS, traffic management service, and DOSMs are coupled to the distributed data storage units via a network. *Id.* at 56–57 (citing Ex. 1003 ¶ 135). Specifically,

Petitioner relies on Figure 16 of Coates and argues the VFS and traffic management service are coupled to the DOSMs and the DOSMs are coupled to the storage nodes via the interconnect fabric 330. *Id.* (citing Ex. 1004, 3:4–8, 5:49–53, 7:33–44, 10:16–22, 12:7–10; Ex 1003 ¶ 135). Finally, Petitioner argues that the OMS manager unit (VFS, traffic management service, DOSMs) and the distributed data storage units implement an object management system because the VFS, traffic management service, and DOSMs manage client requests and control access to the storage nodes and that storage nodes store files for the distributed storage system. *Id.* at 57 (citing Ex. 1004, 3:3–11, Ex. 1003 ¶ 136).

For the aforementioned reasons, at this stage of the proceeding we are persuaded that Coates discloses the preamble and limitations 1a–1c.

b) *[Id] wherein the object management system preferentially selects a first one of the plurality of distributed data storage units for file access in response to a file access request provided that the file access request is associated with an external input/output of the first distributed data storage unit*

Petitioner identifies “a file download request including a storage resource locator (‘SRL’)” as the recited “file access request.” Pet. 57 (citing Ex 1004, 7:50–53, 6:67–7:3, 26:55–67; Ex 1003 ¶ 138). Petitioner argues that the SRL includes a unique file identifier field identifying a particular file and a storage-cluster field identifying a name or IP address of a storage cluster. *Id.* at 57–58 (citing Ex. 1004, 7:2–3, 7:50–8:65, 9:40–45, 11:8–27, 20:51–57; Ex 1003 ¶ 138).

Petitioner argues the SRL in the file access request is associated with an external input/output of the first distributed data storage unit because the storage cluster field in the SRL identifies the cluster preferred by the client and the unique file identifier field in the SRL identifies a network interface

for a specific node within the cluster containing that file. Pet. 58 (citing Ex. 1004, 5:46–48, 9:19–21, 9:40–45, 11:8–27, 20:51–57; Ex. 1003 ¶ 139).

Petitioner contends that the file identifier field in the SRL identifies a specific storage node in the preferred cluster based on a lookup table storing a predetermined mapping of file IDs to IP addresses of storage nodes. *Id.* at 59 (citing Ex. 1004, 11:8–10, 20:51–57, 29:47–49, Fig. 29; Ex 1003 ¶ 141). Thus, according to Petitioner, when a client requests to download (access) a specific file in the client’s preferred cluster, the request’s SRL associates the request with the network interface of a specific storage node in that cluster containing that file. *Id.* at 60 (citing Ex. 1004, 7:2–3, 7:50–8:65, 11:8–27).

Petitioner argues that the initial selection of a storage node based on the SRL is a preference that is subject to change by the OMS manager unit. *Id.* at 61 (citing Ex. 1003 ¶ 143). According to Petitioner each file is replicated in multiple storage nodes across multiple clusters and upon receiving a request for a specific file in a local cluster, if the client’s preferred cluster becomes overloaded with requests, new requests are automatically diverted to the next closest storage center. *Id.* at 61 (citing Ex. 1004, 19:19–20, 19:35–37; Ex. 1003 ¶ 144). Additionally, Petitioner argues that after a cluster is selected, the DOSM in that cluster can pick a different node than the one specified by the SRL’s unique file identifier, as might be the case if the local storage node containing the requested file is overloaded. *Id.* at 61–62 (citing Ex. 1004, 3:12–39, 11:50–60, 19:35–39, 20:59–21:18). Petitioner argues this is an additional reason that the initial selection of a local storage node in response to an SRL is preferential. *Id.* at 62 (citing Ex. 1003 ¶ 145).

Patent Owner argues Coates does not disclose a file access request that is associated with a data storage unit. Prelim. Resp. 29–34.

Specifically, Patent Owner argues that although the SRL identifies the storage cluster and file, none of the fields in the SRL identify a storage node inside the storage cluster. *Id.* at 29–30 (citing Ex. 1004, 7:50–8:7). Instead, according to Patent Owner, Coates’ DOSM maintains a file lookup table that maps the object file to the storage node such that the download SRL does not need to be associated with a storage node. *Id.* at 30–31 (citing Ex. 1004, 11:8–27, Fig. 6). Additionally, Patent Owner argues that the DOSM mapping between file identifiers and storage nodes is not predetermined and instead is dynamically generated. *Id.* at 32 (citing Ex. 1004, 3:40–50, 7:26–27, 9:48–67, 30:47–31:11, Fig. 31).

Patent Owner further argues that Coates does not preferentially select a data storage unit in response to the file access request. Prelim. Resp. 34–41. Instead, according to Patent Owner, Coates preferentially selects a DOSM and preferentially downloads the data object from the selected DOSM’s data cache. *Id.* at 34 (citing Ex. 1004, 9:21–39, Fig. 4). If the object is not cached then the DOSM attempts to identify the location of the object in one of the intelligent storage nodes by first using its look-up table to find the node with the object and if the object is not located in the look-up table then attempting to find the object through broadcasting. *Id.* at 38–39 (citing Ex. 1004, 9:40–10:28, 30:51–31:11, Fig. 4). Thus, Patent Owner argues, once the data object is found in one of the storage nodes the DOSM connects with that storage node to open the file with the requested data object. *Id.* at 38–39 (citing Ex. 1004, 9:48–10:18, Fig. 4). Moreover, addressing Petitioner’s contention that Coates geographically load balances between different storage centers, Patent Owner argues that even if Coates selects a remote storage center because of load balancing, Coates does not

preferentially select a storage node in the remote storage center. *Id.* at 40 (citing Ex. 1004, 19:17–19, 5:38–48, 6:62–7:49, Figs. 1, 3, 16).

At this stage of the proceeding we are persuaded by Petitioner’s arguments that Coates discloses limitation 1d. Specifically, we are persuaded that the SRL in Coates’s file access request is associated with an external input/output of the first distributed data storage unit because the SRL identifies the cluster and file requested and because the file identifier in turn identifies the storage node in which the file is stored. Pet. 58–59 (citing Ex. 1004, 3:12–39, 5:46–48, 7:2–3, 7:50–8:65, 9:19–21, 9:40–45, 11:8–27, 11:50–60, 19:19–20, 19:35–39, 20:51–57, 29:47–49, Fig. 29). We are also persuaded that because Coates replicates files across multiple storage nodes in different clusters and storage centers and that Coates can load balance between storage centers by diverting new requests to the next closest storage center, Coates preferentially selects the storage node from which to access the file. Pet. 61–62; Ex. 1004, 19:24–45. Petitioner’s arguments are supported by the evidence of record cited above and by credible un rebutted testimony from Dr. Mowry. Ex. 1003 ¶¶ 138–145.

We disagree with Patent Owner’s argument that the file request is not associated with the storage node because the SRL of the file request does not identify the storage node. Petitioner’s arguments show that the SRL identifies a file and that the file identifier is used to identify the storage node through a mapping stored in Coates’s lookup table. Pet. at 59 (citing Ex. 1004, 11:8–10, 20:51–57, 29:47–49, Fig. 29; Ex 1003 ¶ 141). At this stage, without further evidence or reasoning that the claim precludes the aforementioned association, we determine Petitioner has established the recited association between the request and the storage unit.

We also disagree with Patent Owner's argument that Coates does not preferentially select a storage node because it only selects a storage center when load balancing rather than selecting a storage node within a storage cluster. Prelim. Resp. 39–41. As Patent Owner acknowledges (*id.* at 40), a storage center includes storage nodes and therefore, selecting a different storage center due to load balancing would also necessarily be a selection of a different storage node from which to access the file.

For the aforementioned reasons, at this stage of the proceeding we are persuaded that Coates discloses limitation 1d.

- c) *[1e] wherein the object management system is configured such that, in response to a file retrieval request that is associated with a data file and an external output of the first distributed data storage unit, the object management system preferentially returns a hostname and pathname of a copy of the data file that is stored within the first distributed data storage unit*

Petitioner argues that Coates discloses this limitation because when a client issues a request with a download SRL and this request is received by a DOSM, the DOSM will look up the specific storage node containing the file. Pet. 62 (citing Ex. 1003 ¶¶ 146–153). This lookup operation returns an IP address for the node, which Petitioner argues is the recited hostname. *Id.* (citing Ex. 1004, 6:67-7:3, 7:50-59, 9:19-21, 20:51-57, Ex. 1003 ¶¶ 147–148). Petitioner argues each file identifier also identifies a pathname which associates a file identifier with a folder ID and parent folder ID. *Id.* at 63 (citing Ex. 1004, 11:8–27, Fig. 12; Ex. 1003 ¶ 149). Petitioner argues that the hostname and file identifier are then returned from the lookup operation and used to retrieve the file from the storage node at that unique IP address. *Id.* (citing Ex. 1004, 12:36–43; Ex. 1003 ¶ 149).

Petitioner argues in the alternative that if “Patent Owner argues that this limitation requires a hostname and pathname to be returned to a *client* this limitation is also disclosed when the VFS initially provides an SRL to a client.” Pet. 64 (citing Ex. 1003 ¶ 150). In the scenario where the client requests to download a file for which it does not yet have an SRL, the VFS returns a name, file identification, download SRL and parent folder identification to the client. *Id.* at 64–65 (citing Ex. 1004, 25:57–67, 26:29–32, 26:43–45; Ex. 1003 ¶ 150). Petitioner argues that it would have been obvious to preferentially route the request to a local cluster in view of Coates’ express teaching of automatic geographic load balancing. Pet. 64 (citing Ex. 1004, 19:14–6). According to Petitioner, the returned download SRL will identify a cluster, which corresponds to a hostname, and the file identifier corresponds to a pathname. *Id.* 65 (citing Ex. 1004, 15:27–30; Ex. 1003 ¶ 152).

Patent Owner argues that Coates does not disclose that a hostname and pathname of the data object stored in the storage node is returned to the client. Prelim. Resp. 41–42 (citing Ex. 1004, 9:16–9:45, 9:48–10:22, Fig. 4). Instead, according to Patent Owner, after the object is located, either in the cache of a DOSM or through the look-up table, the object itself, not the hostname and pathname of the object, is transmitted to the recipient (the client). *Id.* 41 (citing Ex. 1004, 9:33–36, 9:40–45, 10:16–22, Fig. 4).

Responding to Petitioner’s alternative contentions, Patent Owner argues that the VFS only returns XML encoded documents with the list of files and or directories, not a hostname and pathname. Prelim. Resp. 44–45 (citing Ex. 1004, 5:54–6:4, 16:51–58, 26:29–54).

At this stage of the proceeding we are persuaded by Petitioner’s arguments that Coates discloses limitation 1e. Specifically, we are

persuaded that the lookup operation returns an IP address for the node and a file identifier. We agree with Petitioner that the IP address and file identifier disclose the recited hostname and pathname respectively. Patent Owner's argument that the client receives the object rather than the hostname and pathname, implicitly limits the claim as requiring that the object management system return the hostname and pathname to the *client*. On its face, the claim does not specify the entity to which the hostname and pathname are returned. At this stage of the proceeding, Patent Owner does not provide sufficient evidence or reasoning why the claim should be limited in this manner.

For the aforementioned reasons, at this stage of the proceeding we are persuaded that Coates discloses limitation 1e.

d) Conclusion as to Claim 1

For the reasons explained above and based on the arguments and evidence presented in the Petition, we determine that Petitioner has shown Coates discloses the limitations of claim 1. We therefore determine Petitioner has established a reasonable likelihood that Coates anticipates claim 1. We also determine that Petitioner has established a reasonable likelihood that claim 1 would have been obvious over Coates.

2. Claims 2–5 and 9–12

Petitioner contends claims 5 and 9 of the '082 patent are unpatentable as anticipated (Ground 4) by Coates. Pet. 70–73. Petitioner provides sufficient evidence and reasoning that Coates discloses the limitations of each of these claims. Pet. 70–73. Patent Owner relies on the same or similar arguments here as the ones discussed above regarding claim 1. We determine Petitioner has presented sufficient evidence to establish that Coates discloses the limitations of claims 5 and 9.

Petitioner also contends claims 2–5 and 9–12 would have been obvious over Coates (Ground 5). At this stage of the proceeding we determine Petitioner provides sufficient evidence to demonstrate that claims 2–5 and 9–12 would have been obvious over Coates.

IV. CONCLUSION

For the foregoing reasons and on the present record, we determine that the Petition demonstrates a reasonable likelihood that Petitioner would prevail in establishing the unpatentability of at least one of the challenged claims of the '082 patent. At this stage of the proceeding, we have not made a final determination with respect to the patentability of any of the challenged claims.

V. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that, pursuant to 35 U.S.C. § 314(a), an *inter partes* review of the challenged claims of the '082 patent is instituted with respect to all grounds of unpatentability set forth in the Petition; and

FURTHER ORDERED that, pursuant to 35 U.S.C. § 314(c) and 37 C.F.R. § 42.4(b), *inter partes* review of the '082 patent shall commence on the entry date of this Order, and notice is hereby given of the institution of a trial.

IPR2022-00290
Patent 6,779,082 B2

For PETITIONER:

Eliot Williams
Morgan Mayne
Jacob McDonald
BAKER BOTTS L.L.P.
Eliot.williams@bakerbotts.com
Morgan.grissum@bakerbotts.com
Jacob.mcdonald@bakerbotts.com

For PATENT OWNER:

Daniel Block
Byron Pickard
James Hietala
Steven Pappas
Todd Thurheimer
Russell Rigby
STERNE, KESSLER, GOLDSTEIN & FOX P.L.L.C.
Dblock-ptab@sternekessler.com
Bpickard-ptab@sternekessler.com
Jhietala-ptab@sternekessler.com
Spappas-ptab@sternekessler.com
Tthurheimer-ptab@sternekessler.com
rrigby@intven.com