

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

COHESITY INC.,
Petitioner,

v.

COMMVAULT SYSTEMS, INC.,
Patent Owner.

IPR2021-00934
Patent 7,725,671 B2

Before KARL D. EASTHOM, JENNIFER MEYER CHAGNON, and
DAVID COTTA, *Administrative Patent Judges*.

EASTHOM, *Administrative Patent Judge*.

DECISION
Granting Institution of *Inter Partes* Review
35 U.S.C. § 314

Cohesity Inc. (“Petitioner”) filed a Petition (Paper 1, “Pet.”) requesting an *inter partes* review of claims 1, 4, 8, and 15 of U.S. Patent No. 7,725,671 B2 (Ex. 1001, “the ’671 patent”). Petitioner filed a Declaration of Erez Zadok, Ph.D. (Ex. 1003) with its Petition. Commvault Systems, Inc. (“Patent Owner”) filed a Preliminary Response (Paper 8, “Prelim. Resp.”). Patent Owner filed a Declaration of Mark T. Jones, Ph.D. (Ex. 2002) with its Preliminary Response.

The Board has authority to determine whether to institute an *inter partes* review (“IPR”). *See* 35 U.S.C. § 314(b) (2018). Under 35 U.S.C. § 314(a), the Board may not authorize an *inter partes* review unless the information in the Petition and the Preliminary Response “shows that there is a reasonable likelihood that the petitioner would prevail with respect to at least 1 of the claims challenged in the petition.” For the reasons that follow, we institute an *inter partes* review as to the challenged claims, claims 1, 4, 8, and 15 of the ’671 patent, on the sole ground of unpatentability presented.

I. BACKGROUND

A. *Real Parties-in-Interest*

Patent Owner and Petitioner each identifies itself as the real party-in-interest. Paper 4, 1; Pet. 65.

B. *Related Proceedings*

The parties indicate that Patent Owner asserted the ’671 patent in the following district court proceedings: *Commvault Systems, Inc. v. Rubrik Inc.*, No. 1:20-cv-00524-MN (D. Del. filed April 21, 2020); and *Commvault Systems, Inc. v. Cohesity Inc.*, No. 1:20-cv-00525-MN (D. Del. filed April 21, 2020) (“District Court” or “District Court proceeding”). *See* Pet. 65; Paper 4, 1.

The '671 patent was the subject of a separate *inter partes* review in *Rubrik, Inc. v. Commvault Systems, Inc.*, IPR2021-00589 (PTAB, filed Feb. 26, 2021), but the Board terminated the proceeding prior to institution based on the parties' settlement. *See* Pet. 65; Paper 4, 1; IPR2021-00589, Paper 11 (decision granting joint motion to terminate) (PTAB Sept. 1, 2021).

C. *The '671 patent*

The '671 patent generally describes “data classification” methods and systems “to facilitate and improve data management.” Ex. 1001, code (57). More specifically, the '671 patent describes copying data and metadata from a first storage location to another by using metadata stored at the first storage location associated with the data there. *See id.* at 5:36–6:26, Fig. 14. Using the metadata “significantly reduces the amount of time required to obtain data object information by substantially eliminating the need to obtain information from the source data, and furthermore minimizes the involvement of network resources in this process, substantially reducing the processing burden on the host system.” *Id.* at 6:12–17.

The '671 patent describes “metadata” as “data or information about data, and may include, for example, data relating to storage operations or storage management, such as data locations, storage management components associated with data, storage devices used in performing storage operations, index data, data application type, or other data.” Ex. 1001, 6:2–7. Other examples of metadata include “information about data objects or characteristics associated with data objects,” including


the data owner, . . . the last modified time, . . . the data size, . . . information about the content, . . . to/from information for email, . . . creation date, . . . file type, . . . application type, . . . location/network, . . . frequency of change, . . . and aging information (e.g., *a schedule*, which may include a time period,

in which the data object is migrated to secondary or long term storage), etc.

Id. at 7:21–42 (emphasis added). Another example of metadata includes “storage management components associated with data.” *Id.* at 6:5–6.

The '671 patent describes a “metabase” as a “collection of metadata and/or other information regarding the network data.” Ex. 1001, 5:67–6:2. With a collection of metadata in a metabase, “if it is desired to obtain information regarding network data, a system administrator or system process may simply consult the metabase for such information rather than iteratively access and analyze each data item in the network.” *Id.* at 6:8–12.

Figure 13 of the '671 patent follows:



In Figure 13 above, the arrows indicate transfer operations of the metadata between metabases 1350 and 1360, and of data between data stores 1330 and 1340, of associated computing devices 1310 and 1320. Ex. 1001,

22:33–48. A “[d]ata store . . . generally represents application data such as production volume data.” *Id.* at 8:28–29.

D. Illustrative Independent Claims 1 and 15

Of the four challenged claims, the two independent claims, claims 1 and 15, are illustrative and reproduced below with added bracketing:

1. A method of distributing metadata throughout a network, comprising:

[a.] receiving an indication to perform a copy operation on one or more data objects;

[b.] analyzing a first metabase to identify data objects in a first data store associated with the first metabase that satisfy a selection criteria;

[c.] copying the identified data objects from the first data store to a second data store; and

[d.] copying metadata associated with the identified data objects from the first metabase to a second metabase, wherein the second metabase is associated with the second data store.

15. A computer-readable medium encoded with instructions, which when executed by a computer system, perform a method of moving data objects from a primary data store to a secondary data store, the method comprising:

[a.] analyzing a primary metabase to determine data objects stored within a primary data store satisfying a data object query;

[b.] identifying a secondary data store based on metadata stored within a secondary metabase to receive the determined data objects;

[c.] copying the data objects from the primary data store to the secondary data store; and

[d.] updating the secondary metabase with information describing the copied data objects.

E. The Asserted Ground

Petitioner challenges claims 1, 4, 8, and 15 of the '671 patent on the following ground (Pet. 7):

Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
1, 5, 8, 15	103 ¹	Chefalas, ² Tivoli Redbook ³

II. ANALYSIS

A. Legal Standards

Under 35 U.S.C. § 103(a), a claim is unpatentable 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). Tribunals generally resolve obviousness based on underlying factual determinations in the record, 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) so-called secondary considerations. *See Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966). Prior art references must be “considered together with the knowledge of one of ordinary skill in the

¹ The Leahy-Smith America Invents Act (“AIA”), Pub. L. No. 112-29, 125 Stat. 284, 287–88 (2011), amended 35 U.S.C. § 103. The '671 patent's effective filing date precedes March 16, 2013 (the effective date of the relevant amendment), so the pre-AIA version of § 103 applies.

² Chefalas et al., US 6,976,039 B2, issued Dec. 13, 2005. Ex. 1004.

³ International Business Machines, “IBM Tivoli Storage Management Concepts,” (4th ed. 2003). Ex. 1005.

pertinent art.” *In re Paulsen*, 30 F.3d 1475, 1480 (Fed. Cir. 1994) (citing *In re Samour*, 571 F.2d 559, 562 (CCPA 1978)).

B. Level of Ordinary Skill in the Art

Relying on the testimony of Dr. Zadok, Petitioner contends that

[a] person having ordinary skill in the art (PHOSITA) at the time of the '671 Patent would have been a person having at least a Bachelor's Degree in an Engineering discipline such as Electrical, Computer, or Software Engineering, or a Bachelor's Degree in Computer Science, Information Systems, or equivalent degree, and at least two years of relevant experience in the research, design, development and/or administration of data storage systems, with additional education substituting for less experience and *vice versa*.

Pet. 6 (citing Ex. 1003 ¶¶ 89–90). The Preliminary Response does not address Petitioner's proposed level of skill.

The present record, including the '671 patent specification and prior art, supports Petitioner's proposed level of ordinary skill in the art, and we adopt it, at least for purposes of institution. *See Okajima v. Bourdeau*, 261 F.3d 1350, 1355 (Fed. Cir. 2001) (explaining that specific findings regarding ordinary skill level are not required “where the prior art itself reflects an appropriate level and a need for testimony is not shown” (quoting *Litton Indus. Prods., Inc. v. Solid State Sys. Corp.*, 755 F.2d 158, 163 (Fed. Cir. 1985))).

C. Claim Construction

In an *inter partes* review, the Board generally construes each claim “in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent.” 37 C.F.R. § 42.100(b) (2020). “[T]he ordinary and customary meaning of a claim term is the meaning that the term would

have to a person of ordinary skill in the art in question at the time of the invention.” *Phillips v. AWH Corp.*, 415 F.3d 1303, 1313 (Fed. Cir. 2005) (en banc). “There are only two exceptions to this general rule: 1) when a patentee sets out a definition and acts as his own lexicographer, or 2) when the patentee disavows the full scope of a claim term either in the specification or during prosecution.” *Thorner v. Sony Comput. Entm’t Am. LLC*, 669 F.3d 1362, 1365 (Fed. Cir. 2012).

Petitioner asserts that it adopts “the plain meaning of the terms in the Challenged Claims for the purposes of these IPR proceedings only.” Pet. 8. Patent Owner contends that the Petition adopts inconsistent claim construction positions relative to its positions in the related District Court proceeding. *See* Prelim. Resp. 4; *supra* § I.B. Specifically, Patent Owner contends that in the District Court proceeding, Petitioner construes “metabase” as a “metadata database,” and “primary” as “production or live.” Prelim. Resp. 4 (citing Pet. 8; Ex. 2003, 1–2).⁴ Patent Owner contends that because Petitioner takes inconsistent claim construction positions, “the

⁴ Subsequent to the filing of the Preliminary Response, the parties entered a “Joint Claim Construction Statement” in the District Court. Ex. 3001. The Joint Claim Construction Statement states as follows:

1. “The parties agree that ‘primary data store’ means ‘a data store for production or live data’”
2. “The parties agree that ‘primary metabase’ means ‘a metabase for production or live data.’” Ex. 3001, 1.

Also, the Joint Claim Construction Statement shows that Patent Owner proposes to construe a “metabase” as a “database that contains metadata” and Petitioner proposes to construe it as “metadata database.” Ex. 3001, 4. According to PACER, the District Court has not yet ruled on any claim construction terms.

Petition fails to demonstrate a reasonable likelihood of success as to any of the challenged claims and institution should be denied.” *See id.* at 13.

Petitioner contends that construction beyond the plain meaning is not necessary in this proceeding “[g]iven the nature of the [asserted] prior art and the proximity to the asserted claims.” Pet. 8. The record does not support Patent Owner’s argument that Petitioner’s claim construction here, which relies on the plain meaning of terms, is inconsistent with Petitioner’s proposed District Court construction, simply because they are not exactly the same. *See* Prelim. Resp. 4–13; *see also* Pet. 8 n.1 (“Petitioner reserves the right to assert claim constructions in other proceedings, such as in district court, *that are consistent with the plain and ordinary meaning of the claim terms* in view of the patent and file history.” (emphasis added)). Any difference between the parties in the District Court or with respect to the plain meaning as advanced in the Petition does not raise a material distinction on this preliminary record for purposes of institution.⁵

Based on the foregoing discussion, on this preliminary record, and for purposes of institution, the record does not support Patent Owner’s argument that the plain meaning of any term is inconsistent with Petitioner’s position

⁵ Patent Owner contends that “Petitioner points to Chefalas’ disclosure that this ‘data store’ stores ‘meta data’ (Pet., 27) but does not explain how that makes it a ‘metadata database’ that only has metadata per its District Court construction, much less a ‘production or live’ one at that.” Prelim. Resp. 6. Contrary to this argument, Petitioner does not contend in the District Court that a “metadata database” “only has metadata,” and Patent Owner’s proposed District Court construction indicates the same. *See* Ex. 3001, 4. As to the term “primary metabase” or “primary data store,” on this preliminary record, the phrase “for production or live data” appears to be a statement of intended use so that this agreed upon construction merely requires a metabase and data store to be able to store production or live data.

in the District Court in a material manner. *See supra* note 5. As to Patent Owner’s related contention that it is not “self-evident that Petitioner’s district court constructions are met” in the asserted ground (Prelim Resp. 5), as just noted, no material inconsistencies with the plain meaning of any terms at issue exist on this preliminary record, and the Petition addresses the plain meaning. *See supra* note 5; *infra* § II.D.

Patent Owner also argues that limitation 1d, as supported by the specification, requires “‘copying metadata associated with the *identified* data objects from the first metabase to a second metabase,’ not copying *all* metadata.” Prelim. Resp. 32 (emphasis to claim by Patent Owner).

The Petition, as addressed below, shows agreement by the parties that the claims at least read on copying only the metadata associated with a subset of data objects under Patent Owner’s narrow claim construction. The analysis below addresses Petitioner’s showing under Patent Owner’s narrow claim construction arguments. *See infra* § II.D.4.d.

However, contrary to Patent Owner’s proposed construction of step 1d, claim 1 does not recite “copying [*only*] metadata associated with the identified data objects from the first metabase to a second metabase.” Moreover, the specification describes copying all of the metadata at least in some situations, for example, “classification agent 312b may periodically copy or backup metabase 314 pursuant to a storage policy . . . such that the information in metabase 314 may be quickly restored if lost, deleted or is otherwise unavailable.” Ex. 1001, 9:36–40.

Furthermore, on this preliminary record, copying only a select group of metadata is not a “key aspect” of the invention, contrary to Patent Owner’s arguments. *See* Prelim. Resp. 14. Rather, the specification

indicates that regardless of whether the system copies unrelated metadata, the purpose of the metadata is “to quickly locate information”:

[s]toring metadata relating to data content in metabase 314 provides the capability to perform content searches for data in the metabase 314, instead of searching entries in mass storage 318. This allows the system to quickly locate information satisfying a content query that may be retrieved from mass storage 318, if necessary.

Ex. 1001, 9:62–67.

Similarly, with respect to Figure 14, which Patent Owner relies upon to support its narrow construction (Prelim. Resp. 29–32), the specification states that “at step 1410, *data objects and/or associated metadata may be identified for movement* from one computing device to another.” Ex. 1001, 22:51–53 (emphasis added). Then, it states that

[t]his may be accomplished by forming a query for certain data, such as a search for data that may be to be moved or copied pursuant to a data management or storage policy, or in response to a request to move data relating to certain processes or applications from one computing device to another, for any other suitable purpose such as disaster recovery, resource reallocation or reorganization, load balancing, etc.

Id. at 22:53–60. Accordingly, on this preliminary record, the specification is permissive—it contemplates 1) copying only some of the data and/or metadata (because copying “may be” selective), or 2) copying all of the data and/or metadata “for any . . . suitable purpose.” *See id.* at 22:51–60.

The ’671 specification also advises that “[t]he above detailed description of embodiments of the invention is not intended to be exhaustive or to limit the invention to the precise form disclosed above,” further indicating that claim 1 does not limit copying to copying only some of the

metadata, even if an embodiment describes that feature. Ex. 1001, 39:13–15.

For the reasons discussed above, on this preliminary record, claim 1 encompasses both copying all metadata and only selected metadata. During trial, the parties will have the opportunity to address any claim construction issues.

On this preliminary record and for purposes of institution, no claim terms require explicit construction. *See, e.g., Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017) (“[W]e need only construe terms ‘that are in controversy, and only to the extent necessary to resolve the controversy’ . . .” (quoting *Vivid Techs., Inc. v. Am. Sci. & Eng’g, Inc.*, 200 F.3d 795, 803 (Fed. Cir. 1999))).

D. Obviousness, Chefalas and Tivoli Redbook, Claims 1, 4, 8, 15

Petitioner contends that claims 1, 4, 8, and 15 would have been obvious over the combination of Chefalas and Tivoli Redbook. Pet. 15–54. Patent Owner disputes Petitioner’s contentions. Prelim. Resp. 13–50.

1. Chefalas (Ex. 1004)

Chefalas discloses methods for identifying files to copy/backup using relational metadata. Ex. 1004, 2:60–64. Specifically, Chefalas discloses a backup/copy method in which “a data store is queried containing meta data regarding files associated with the application in which the data store includes metadata describing the files accessed by the application” and the result of this query “is used as input to the backup program.” *Id.* at 2:64–3:4.

Chefalas's Figure 4 follows:




Figure 4 illustrates a block diagram for a copy system including source 400 containing “operating system” 402, “file” 406, “application” 408, “backup application” 414, and “database” 412. Ex. 1004, 6:20–24; 6:61–7:6. Application 408 and backup application 414 communicate with operating system 402. *Id.* at Fig. 4. Operating system 402 includes “file functions” 404 and “device driver” 410 connected to database 412. *Id.* “[B]ackup application 414 may be any custom-written backup application or any commercially available backup applications such as Tivoli Storage Manager.” *Id.* at 7:58–61. “Storage system 416 may contain one or more storage devices, such as, for example, a hard disk drive, a tape drive, or a rewriteable compact disk drive.” *Id.* at 7:67–8:3.

Chefalas’s metadata “is stored in a data store and takes the form of a database.” Ex. 1004, 8:27–28. Specifically, Chefalas’s “relational meta data stored in database 412 is updated to reflect the association with multiple applications.” *Id.* at 8:8–10. Chefalas identifies “examples of meta data,”

including records 500, 502, and 504, “which may be stored in a database, such as database 412 in [Figure] 4.” *Id.* at 8:6–9:1.

During copying as initiated by a user or administrator, remotely via server 418 or locally via backup application 414 on source 400, files are “backed up by reading the relational meta data for each file that is related to the relevant application. The user can then accept or refuse for one or more files to be copied or backed up to the new system.” Ex. 1004, 8:31–41. Or “the backup through backup server 418 or backup application 414 may be employed to identify files without requiring a user to select files for backup.” *Id.* at 8:41–44.

Chefalas’s Figure 15 follows:




Figure 15 illustrates a backup/copy sequential process as follows: step 1500, receiving a request for backup data; step 1502, requesting relational metadata; step 1504, receiving relational metadata; and step 1506, identifying files associated with the relational metadata. Then, steps 1508, 1510, and 1512 involve requesting, receiving and transferring identified files to a backup device. *Id.* at 11:18–46, Fig. 15.

2. Tivoli Redbook (Ex. 1005)

Tivoli Redbook describes Tivoli Storage Manager software. Ex. 1005, 47. As noted above, Chefalas refers to Tivoli Storage Manager software as a “commercially available backup application” in reference to Chefalas’s backup application 414. Ex. 1004, 7:58–61.

Tivoli Redbook states that “[t]he specially designed IBM Tivoli Storage Manager database retains information about all client system and user files, business policies, disaster recovery, and the scheduling of client and administrative tasks. This database retains information called metadata, which means data that describes data.” Ex. 1005, 47.

Tivoli Redbook illustrates the components of the IBM Tivoli Storage Manager in Figure 4-1, reproduced below.




Figure 4-1 illustrates a centrally located “IBM Tivoli Storage Manager Server.” Ex. 1005, 58. To the right of the Server, Figure 4-1 illustrates a pyramid labeled “Data Storage Hierarchy” connected to the Server with a line labeled “Storage pools Policy Placement and Recovery.” *Id.* Below the Data Storage Hierarchy pyramid, Figure 4-1 illustrates “LAN-SAN Clients” workstations connected to the Server. *Id.* To the left of the Server, Figure 4-1 illustrates a green “Recovery Log” storage with an arrow pointing to an orange data storage. *Id.* Both green and orange data storages connect to the Server by separate lines labeled “Relational Database.” *Id.*

3. Independent claim 15⁶

Claim 15’s preamble recites “[a] computer-readable medium encoded with instructions, which when executed by a computer system, perform a

⁶ This Institution Decision tracks the Petition’s format by addressing independent claim 15 prior to independent claim 1. *See* Pet. 16 n.2 (“Petitioner addresses independent claim 15 first given that it includes many

method of moving data objects from a primary data store to a secondary data store, the method comprising.” “To the extent [the preamble] is limiting,” Petitioner relies on Chefalas’s computer implemented instructions for identifying files to copy using relational metadata. Pet. 16–17 (citing Ex. 1004, 2:60–3:7). Petitioner contends that Chefalas’s process of “copying data ‘from source 400 to storage system 416’” teaches the claimed method of moving data objects from a primary data store to a secondary data store. *Id.* at 17–18 (citing Ex. 1004, 8:45–51).


Patent Owner does not specifically address Petitioner’s showing regarding the preamble. Based on the foregoing summary and a review of the record, even if the preamble is limiting, Petitioner’s showing is sufficient for purposes of institution.

- a. “analyzing a primary metabase to determine data objects stored within a primary data store satisfying a data object query”*

Petitioner reads limitation 15a on Chefalas’s disclosure of “quer[ying] the relational meta data in database 412 and build[ing] a list of the files to transfer from source 400 to storage system 416.” Pet. 20 (citing Ex. 1004, 8:45–57). Petitioner contends that “database 412 is an example of a ‘metabase’ which stores metadata,” and quotes Chefalas as stating that “the meta data is stored in a data store and takes the form of a database. A data store is some data structure or system used to store meta data.” *Id.* at 19–20 (quoting Ex. 1004, 8:27–30).

of the elements and limitations of independent claim 1 and is otherwise more limiting in most aspects.”).

Petitioner annotates Chefalas's Figure 5 as reproduced below:



Annotated Figure 5 represents searchable metadata 506–516 in stored data block metadata records 500, 502, and 504. Ex. 1004, 8:63–9:15. As Petitioner annotates in Figure 5 above, the searchable metadata includes the following in sequential order: date of the last file update 506; time the file was last accessed 508; name of the file 510; location of the file 512; application making the call for the file 514; and user 516. Pet. 19–20 (citing Ex. 1004, 8:63–9:8, Fig. 5).


Petitioner contends that “given that Chefalas teaches backing up data from (1) a source computer system that may be utilized by a user (2) to a backup server and storage system, a PHOSITA would understand that the source system would contain primary data and would be a primary data store.” Pet. 21–22 (citing Ex. 1003 ¶¶ 104–112). Petitioner further contends that Chefalas’s metabase is a primary metabase because “it is located on the source computer system and includes metadata that describes the primary data in the source system.” *Id.* at 22 (citing Ex. 1003 ¶¶ 104–112).

Patent Owner does not specifically address Petitioner’s showing regarding limitation 15a. Based on the foregoing summary and a review of the record, Petitioner’s showing as to limitation 15a is sufficient for purposes of institution.

- b. “identifying a secondary data store based on metadata stored within a secondary metabase to receive the determined data objects”

Petitioner relies on the collective teachings of Chefalas and Tivoli Redbook in its analysis of limitation 15b. Pet. 22–30.

Petitioner annotates and reproduces Chefalas’s Figure 4 as follows:



In Chefalas’s Figure 4, as annotated by Petitioner, shaded backup server 418 and storage system 416 on the left constitute a “secondary data store” as recited in claim 1. See Pet. 23 (citing Ex. 1004, Fig. 4).

Petitioner contends that Chefalas teaches that backup server 418 and backup application 414 act as an agent or intermediary to copy and store data from source database 412 onto storage system 416. Pet. 23 (citing Ex. 1004, 8:45–51). Petitioner contends that “[a] PHOSITA would understand that the storage related to the Backup Server would be a secondary data store given that it stores the secondary copies of data copied from the source computer (*i.e.*, the primary data store).” *Id.* at 23–24 (citing Ex. 1003 ¶¶ 113–115).

Petitioner acknowledges that, although “*Chefalas* discloses multiple secondary data stores for backing up secondary copies of data, *Chefalas* does not expressly describe the backup operation on the server in detail, nor does it expressly disclose a secondary metabase used in identifying the secondary data store based on metadata.” Pet. 24. However, Petitioner contends that Tivoli Redbook suggests these features in its description of Tivoli Storage Manager, a known commercially available backup system that *Chefalas* specifically describes for use as backup application 414. See *id.* (citing Ex. 1004, 7:58–61; Ex. 1003 ¶¶ 113–115).

Petitioner annotates and reproduces Tivoli Redbook’s Figure 4-1 as follows:




Figure 4-1 Divisions of the IBM Tivoli Storage Manager database

Petitioner annotates Tivoli Redbook’s Figure 4-1 as reproduced above to show a “Secondary Metabase” to the left of the IBM Tivoli Storage Manager Server. See Pet. 25 (citing Ex. 1005, 58, Figure 4-1). Petitioner contends that “*Tivoli Redbook* discloses a secondary database of metadata that is used to select where to store back-up (*i.e.*, secondary) copies of

selected primary data.” *Id.* at 26 (citing Ex. 1005, Figure 4-1). Petitioner contends that the secondary database is a metabase because Tivoli Redbook teaches that the “database retains information about all client system and user files, business policies, disaster recovery, and the scheduling of client and administrative tasks. *This database retains information called metadata, which means data that describes data.*” *Id.* (emphasis added to original by Petitioner) (quoting Ex. 1005, 47). Petitioner contends that Tivoli Redbook discloses that “[t]he stored metadata includes ‘*all policies for handling your data,*’ . . . ‘*information about each node or client’s files for which it manages data.*’ . . . ‘*file locations for each client,*’ and ‘*storage durations for each file.*’” *Id.* (emphasis in original by Petitioner) (quoting Ex. 1005, 57–58). Finally, Petitioner contends that “the IBM Tivoli Storage Manager *database looks at the rules section again to determine where to put the backup copy of the data,*” and then moves the data from the server to the storage pool previously defined for that client node.” *Id.* at 26–27 (citing Ex. 1005, 59).

Petitioner argues that a PHOSITA would have understood that a metabase stored on the backup system is a “secondary metabase” as recited in claim 15. Pet. 27 (citing Ex. 1003 ¶¶ 116–123). Petitioner contends that

the combination of *Chefalas* with *Tivoli Redbook* discloses a secondary metabase (*i.e.*, the server’s “relational database”) used in identifying the secondary data store (*i.e.*, backup storage pools managed by the server) based on metadata (*i.e.*, the metadata policies and rules stored within the database for determining backup data placement) as recited in claim 15(b).

Id. (citing Ex. 1003 ¶¶ 116–123).

As to reasons to combine the references, Petitioner first argues that *Chefalas* expressly teaches utilizing Tivoli Storage Manager as backup

application 414. Pet. 27–28 (citing Ex. 1004, 7:58–65; Ex. 1003 ¶¶ 124–125). Second, Petitioner argues that both Chefalas and Tivoli Redbook relate to IBM technologies, which in combination would provide a complete data management and backup system. *Id.* at 28 (citing Ex. 1004, code (73); Ex. 1005, i–ii; Ex. 1003 ¶ 126). Third, Petitioner argues that Tivoli Redbook “discloses a similar primary metabase associated with the primary storage that is used for incremental backup and maintains similar metadata, including modification times as disclosed in” Chefalas, so a person of ordinary skill in the art would have used Tivoli Redbook’s teachings to implement Chefalas’s backup system. *Id.* at 28–29 (citing Ex. 1003 ¶¶ 127–128). Finally, Petitioner argues that applying Tivoli Redbook’s secondary metabase to Chefalas’s backup functionality would maintain the metadata associated with the files that are backed up from source 400, resulting in more efficient usage of available storage and retrieval of stored data. *Id.* at 29 (citing Ex. 1003 ¶ 129).

Patent Owner argues that “Petitioner asserts that Tivoli Redbook’s database retains metadata and then concludes that the policies in the rules section of the database are metadata. Petitioner, however, offers no explanation or proof as to why these policies are metadata.” Prelim. Resp. 44 (citing Ex. 2002 ¶ 73). Patent Owner argues “[t]hat Tivoli Redbook’s database includes metadata is not proof that the policies also stored on the database are metadata.” *Id.* at 47. Patent Owner similarly argues that “Tivoli Redbook . . . does not teach that its policies are metadata.” *Id.* at 49 (citing Ex. 2002 ¶¶ 80–81; Ex. 1005, 51).

Notwithstanding Patent Owner’s arguments, on this preliminary record, the ’671 patent does not present a clear line of demarcation between

all policies and metadata. The '671 patent specification states that “[g]enerally, metadata refers to data or information about data, and may include, for example, data relating to storage operations or storage management, such as data locations, storage management components associated with data, storage devices used in performing storage operations, index data, data application type, or other data.” Ex. 1001, 6:2–7 (emphasis added). As Patent Owner also notes, one form of metadata includes “aging information (e.g., a schedule, which may include a time period, in which the data object is migrated to secondary or long term storage), etc.” Prelim. Resp. 48–50 (quoting Ex. 1001, 7:40–42; citing Ex. 2002 ¶ 79) (emphasis added). This “aging information” or “schedule” appears to be similar to a rule or policy that includes “storage durations for each file,” as described in Tivoli Redbook, because it applies a rule (i.e., a schedule) for migrating a data object. See Ex. 1005, 57–58. In addition, according to the '671 patent, “metadata refers to . . . information about data,” which may include “storage management components associated with data.” See Ex. 1001, 6:2–7 (emphasis added). This also suggests that “metadata” encompasses information in the form of policy or rule components for managing data storage, as disclosed in Tivoli Redbook. During trial, the parties will have the opportunity to clarify the record in this regard.

In addition, although Petitioner relies upon rule or policy information as metadata, Petitioner also relies on Tivoli Redbook’s stored “metadata, which means data that describes data.” See Pet. 26 (quoting Ex. 1005, 47). Petitioner also contends that Tivoli teaches that “[t]he stored metadata includes ‘all policies for handling your data, as well as the information about each node or client’s files for which it manages data,’ and that the

“database also keeps track of *file locations for each client and storage durations for each file.*” *Id.* (emphasis by Petitioner) (quoting Ex. 1005, 57–58). Tivoli also teaches that, “[c]lient data attributes *such as storage destination*, number of versions, and retention period can be assigned at the individual file level and *stored in the database.*” Ex. 1005, 47 (emphasis added).

To support its showing, Petitioner contends that “IBM Tivoli Storage Manager has defined some logical entities to classify the available storage resources. Most important is the logical entity called a storage pool.” Pet. 25 (quoting Ex. 1005, 40). Petitioner also notes that Tivoli Redbook’s rules rely on stored registration and other data about the client and that the rules implement “schedules”:

IBM Tivoli Storage Manager server must first check in with the database to see whether the client node is registered and authorized to work with this server. The client nodes are listed or registered in this Rules section, which contains instructions, policies, and schedules that the server must follow in order to complete the request it received from the client node.


Id. at 26–27 n.3 (quoting Ex. 1005, 58–59).

As indicated above, Tivoli Redbook’s server stores data to “keep[] track of data storage locations,” which constitutes metadata (data about data). Ex. 1005, 57. In other words, Petitioner sufficiently shows for institution purposes that even if the server’s policies and/or rules in Tivoli Redbook are not metadata, the policies and rules generally use metadata “to complete the request [the server] received from the client node.” *See id.* at 58–59. Petitioner also sufficiently shows that Tivoli Redbook’s metadata and policies at least suggest the same type of metadata contemplated by the ’671 patent specification. Based on the foregoing summary and a review of

the record, Petitioner's showing as to limitation 15b is sufficient for purposes of institution.

- c. *"copying the data objects from the primary data store to the secondary data store"*

Petitioner contends that Chefalas alone or in combination with Tivoli Redbook teaches this claim limitation. Pet. 30–32. Petitioner annotates and reproduces Chefalas's Figure 15 as follows:



Chefalas's Figure 15, as annotated by Petitioner, includes blue-shaded steps in the middle as follows: step 1506, "identify files associated with application using relational meta data"; step 1508, "request identified files"; step 1510, "receive identified files"; and step 1512, "transfer files to backup device." Pet. 31 (citing Ex. 1004, Fig. 15). Petitioner contends that Chefalas teaches that in step 1512, the identified files (data objects) result

from querying database 412. *Id.* (citing Ex. 1004, 11:30–38; Ex. 1003 ¶¶ 131–133).

Petitioner further contends that Tivoli Redbook teaches a “journal-based backup” in which “data objects are identified for backup by accessing a local database (*i.e.*, a primary metabase) that tracks files that have been updated over time.” Pet. 32 (citing Ex. 1005, 112). Petitioner quotes Tivoli Redbook’s disclosure that “‘backup archive client obtains the backup candidate list by contacting the Journal Based Backup Daemon,’ which ‘manages and maintains a journal database of change activity.’” *Id.* (quoting Ex. 1005, 112). Petitioner contends that a “PHOSITA would understand that obtaining the candidate list of files to backup is identifying said list.” *Id.*

Patent Owner does not specifically address Petitioner’s showing regarding limitation 15c. Based on the foregoing summary and a review of the record, Petitioner’s showing as to limitation 15c is sufficient for purposes of institution.

d. “updating the secondary metabase with information describing the copied data objects.”

Petitioner relies on the collective teachings of Chefalas and Tivoli Redbook in its analysis of limitation 15d. Pet. 32–37. Petitioner acknowledges that Chefalas “does not expressly disclose updating metadata in a secondary metabase.” *Id.* at 33. Petitioner contends that Tivoli Redbook teaches that this is a feature of Tivoli Storage Manager, which Chefalas lists as a backup application. *Id.* at 33–34. Specifically, Petitioner contends that the Tivoli Storage Manager moves data from the server to the storage pool and also “looks into the database *to store the metadata from this client node*. Each file or object that it receives from this client node is noted

in the files section with all of the relevant attributes about it.” *Id.* (emphasis by Petitioner) (quoting Ex. 1005, 59). Petitioner further quotes Tivoli Redbook as teaching that for “journal-based (*i.e.*, incremental) backup, “[t]he backup-archive client, as always, still sends data (files) to the IBM Tivoli Storage Manager server and, as has always been the case, the IBM Tivoli Storage Manager server populates the IBM Tivoli Storage Manager database with the file details and location.”” *Id.* at 34 (emphasis by Petitioner) (quoting Ex. 1005, 113).

As motivation to update the metadata, Petitioner first argues that a person of ordinary skill in the art would have understood the necessity of updating the tracking locations for the secondary copies in the secondary data store to locate the data as needed for restoration. Pet. 34–35 (citing Ex. 1003 ¶¶ 141–142). Second, Petitioner argues that a person of ordinary skill would have recognized that during restoration, “the list of files available to restore would have to be queried on the secondary metabase, because (1) it has the most up-to-date files that can be restored and (2) the primary storage system (and its local metabase) may have been destroyed or damaged.” *Id.* at 35–36 (citing Ex. 1003 ¶¶ 143–144). Third, Petitioner argues that both Chefalas and Tivoli Redbook teach continuous or real-time updating of the metadata upon file changes. *Id.* at 36 (citing Ex. 1004, 7:7–12; Ex. 1005, 112). Finally, Petitioner contends that “[a] PHOSITA would understand that based on” the cited teachings in Tivoli Redbook and Chefalas of “updating the metadata in the primary metabase . . . it would be likewise beneficial to update the metadata (*i.e.*, in a secondary metabase) in the backup system, so that both the data and metadata would be restored consistently.” *Id.* (citing Ex. 1003 ¶¶ 145–146).

Patent Owner does not specifically address Petitioner’s showing regarding limitation 15d. Based on the foregoing summary and a review of the record, Petitioner’s showing as to step 15d is sufficient for purposes of institution.

e. Summary of claim 15

Considering the arguments and evidence presented by Patent Owner as summarized above, Petitioner’s showing is sufficient for purposes of institution. Based on the foregoing discussion and a review of the preliminary record, Petitioner demonstrates a reasonable likelihood that it would prevail with respect to showing the unpatentability of claim 15.

4. Independent Claim 1

Claim 1 is similar to claim 15, but claim 1 recites a “method” and other claim features addressed below. *See supra* § I.D. (listing claims 1 and 15). To address the similar limitations of claim 1, Petitioner largely relies on or repeats its analysis with respect to claim 15. *See* Pet. 37–48. For example, Petitioner contends that “[t]o the extent the preamble is limiting, *Chefalas* in combination with *Tivoli Redbook* discloses a ‘method of distributing metadata throughout a network’ for the same reasons detailed above for claim 15(pre) . . . and [] 15(d).” *Id.* at 37 (citing Ex. 1003 ¶ 148).

Patent Owner does not specifically address Petitioner’s showing regarding the preamble. Based on the foregoing summary and a review of the record, even if the preamble is limiting, Petitioner’s showing is sufficient for purposes of institution.

a. “receiving an indication to perform a copy operation on one or more data objects”

Petitioner contends that *Chefalas* teaches that a user desiring to backup data for an application may select a backup function through backup

application 414 or initiate a remote backup through backup server 418. Pet. 38 (citing Ex. 1004, 8:31–36, Fig. 4). Petitioner argues that Chefalas illustrates this process in Figure 15, beginning with step 1500, “receiving a request from a requestor to backup data for an application The request may be made by a user or by a process.” *Id.* at 39 (quoting Ex. 1004, 11:18–23). Petitioner argues that Chefalas teaches transferring the identified files to a backup device, so that copied files are those identified by the query of database 412. *Id.* at 39–40 (citing Ex. 1004, Fig. 15, steps 1506, 1512).


Patent Owner does not specifically address Petitioner’s showing regarding step 1a. Based on the foregoing summary and a review of the record, Petitioner’s showing as to step 1a is sufficient for purposes of institution.

- b. *“analyzing a first metabase to identify data objects in a first data store associated with the first metabase that satisfy a selection criteria”*

Petitioner contends that “*Chefalas* discloses this claim element for similar reasons as discussed above in claim 15(a),” where the “first data store” and “first metabase” of claim 1 correspond to the “primary data store” and “primary metabase” of claim 15, respectively. Pet. 40. Petitioner contends that “a PHOSITA would understand that the first data store disclosed in *Chefalas* is associated with database 412 as discussed above in claim 15(a) at least given that database 412 is located at and utilized to identify the data for copying from source 400.” *Id.* at 40–41 (citing Ex. 1003 ¶¶ 153–154). As to identifying data objects, Petitioner contends that Chefalas “discloses querying or searching the metabase, for example, which, ‘*specifies an application name as well as criteria, which may be used*

to search for records, such as records 500, 502, and 504 in FIG. 5 within database 412 in FIG. 4.” *Id.* at 41 (quoting Ex. 1004, 9:9–17).

Petitioner contends that Chefalas’s Figure 15, as annotated and reproduced below, illustrates a process that identifies a list of files for backup:



Petitioner annotates Chefalas’s Figure 15 by shading step 1502, “request for relational meta data for the application from database”; step 1504, “receive relational meta data”; and step 1506, “identify fields associated with application using relational meta data.” Pet. 42 (citing Ex. 1004, Fig. 15).

Patent Owner does not specifically address Petitioner’s showing regarding step 1b. Based on the foregoing summary and a review of the

record, Petitioner's showing as to step 1b is sufficient for purposes of institution.

- c. *"copying the data objects from the primary data store to the secondary data store"*

Petitioner contends Chefalas alone or in combination with Tivoli Redbook teaches step 1c. Pet. 42–43. Petitioner contends that "given that the secondary data stores in *Chefalas* and *Tivoli Redbook* are backup systems that receive data from the primary or first data store, these data stores read on the claimed second data store for the same reasons as discussed in claim 15(c)." *Id.* at 43 (citing Ex. 1003 ¶ 158).

Patent Owner does not specifically address Petitioner's showing regarding step 1c. Based on the foregoing summary and a review of the record, Petitioner's showing as to step 1c is sufficient for purposes of institution.

- d. *"copying metadata associated with the identified data objects from the first metabase to a second metabase, wherein the second metabase is associated with the second data store"*

Petitioner relies on the combined teachings of Chefalas and Tivoli Redbook to teach or suggest limitation 1d. Pet. 43–48. As discussed above with respect to steps 15b ("identifying a secondary data store") and 15d ("updating the secondary metabase"), Petitioner contends that Chefalas discloses a first metabase, database 412, and Tivoli Redbook discloses a second metabase, the IBM Tivoli Storage manager database. *See supra* §§ II.D.3.b., II.D.3.d.; Pet. 43 (citing Ex. 1003 ¶ 159). Petitioner contends that "[a] PHOSITA would . . . understand that in updating the metadata at the second metabase, one would need to copy the metadata from the first metabase (*i.e.*, database 412 disclosed in *Chefalas*) to the second metabase

(i.e., the IBM Tivoli Storage Manager database disclosed in *Tivoli Redbook*) for four reasons.” *Id.* at 43–44.⁷

As its first reason, Petitioner contends that “metadata must be preserved at the source computer system prior to backup or migration to the second system in order for the metadata to be an accurate representation of the source system.” Pet. 44. Petitioner contends that

[a] PHOSITA would recognize that the metadata . . . must be preserved at the source because if the metadata is derived at a later time, such as after the data object is copied to the secondary data store), the metadata may change (e.g., the modification/creation time) and will not accurately describe the data object.

Id. at 45 (citing Ex. 1003 ¶ 162).

As a second reason, Petitioner contends that *Tivoli Redbook* and *Chefalas* describe storing similar metadata. *See* Pet. 45. For example, Petitioner contends *Tivoli Redbook* teaches that stored data in its second metabase includes “date and time, file size, file creation date and time, file modification date and time, file location structure on the client node, and other client-relevant information.” *Id.* (quoting Ex. 1005, 59). Petitioner contends that *Chefalas* depicts storing “much of this same metadata” in first metabase 412 as depicted in Figure 5, including “‘date of the last file update,’ ‘time the file was last accessed,’ ‘name of the file,’ ‘location of the file,’ ‘application making the call for the function,’ and the ‘user.’” *Id.* at 45 (quoting Ex. 1004, 8:63–9:8, Fig. 5). Petitioner further contends that *Tivoli Redbook* suggests that similar metadata exists in the journal database (first

⁷ As noted above, Patent Owner does not address Petitioner’s showing with respect to updating. *See supra* § II.D.3.d.

metabase) on the source computer prior to backup. *Id.* at 45–46 (citing Ex. 1005, 112; Ex. 1003 ¶¶ 163–164).

To illustrate some of the metadata similarities, Dr. Zadok produces the following table (Ex. 1003 ¶ 164):

<u>Chefalas First Metabase</u> (<i>Chefalas</i> at Fig. 5; 8:63-9:8)	<u>Tivoli Redbook First Metabase</u> (<i>Tivoli Redbook</i> at 112)	<u>Tivoli Redbook Second Metabase</u> (<i>Tivoli Redbook</i> at 59)
“time the file was last accessed”	“Access time/date changes”	“date and time”
-	“Size changes”	“file size”
-	“Objects created, deleted, or renamed on the file system”	“file creation date and time”
“date of the last file update”	“Modification time/date changes”	“file modification date and time”
“location of the file”	-	“file location structure on the client node”

From the first column on the left to the third column on the right, the table above, as produced by Dr. Zadok, lists similar metadata in Chefalas’s first metabase, Tivoli Redbook’s first metabase, and Tivoli Redbook’s second metabase, respectively.

As a third reason, Petitioner contends that because the source computer preserves metadata in a first metabase “the most natural and efficient method of updating the metadata would be to copy the metadata from the first to the second metabase as suggested by the combination of *Chefalas* and *Tivoli Redbook*.” Pet. 46. Petitioner contends that

[a] PHOSITA would understand that accessing and copying the metadata from the metabase (*i.e.*, through standard APIs) would be more efficient and less resource intensive than re-creating the metadata at the source computer system (which would require

analyzing and deriving the information from the data objects themselves) for transmitting to the secondary data store.

Id. at 46–47 (citing Ex. 1003 ¶ 165).

As its final reason, Petitioner contends that “copying metadata from a first to a second computer system was a well-known and conventional method for migrating, copying, or transferring data or metadata and thus would have been within the knowledge of a PHOSITA.” Pet. 47. To support this argument, Petitioner cites to Osaki⁸ and Takata.⁹ *Id.* at 47–48 (citing Ex. 1006 ¶ 26, Fig. 3; Ex. 1008 ¶ 173; Ex. 1003 ¶¶ 166–168).

Relying on the testimony of Dr. Jones, Patent Owner argues that Chefalas and Tivoli Redbook do not teach or suggest copying metadata from a first metabase to a second metabase. Specifically, Patent Owner argues that neither Chefalas nor Tivoli Redbook individually contains two metabases. Prelim. Resp. 17–18. Patent Owner explains that although Chefalas identifies files by relying “on metadata stored in a first alleged metabase, database 412 on source 400,” “backup server 418 and storage system 416 do not have a corresponding database.” *Id.* at 18–19 (quoting Ex. 2002 ¶ 36). Addressing Petitioner’s argument that the IBM Tivoli Storage Manager’s journal database is a first metabase similar to Chefalas’s first metabase, Patent Owner quotes the testimony of Dr. Jones, who testifies that “Petitioner . . . does not explain why this journal database is a metabase” and that “Tivoli Redbook does not state that these journal entries contain metadata.” *Id.* at 19–20 (quoting Ex. 2002 ¶¶ 36–38). Patent Owner also argues “Petitioner does not argue copying metadata from a first *metabase* to

⁸ Osaki et al., US 2006/0218365 A1, published Sept. 28, 2006. Ex. 1006.

⁹ Takata et al., US 2004/0215878 A1, published Oct. 28, 2004. Ex. 1008.

a second *metabase* was well-known or conventional,” in view of Osaki and Takata. *Id.* at 26 (citing Ex. 2002 ¶¶ 49–50).

Patent Owner’s arguments do not undermine Petitioner’s showing here for institution purposes. As Patent Owner notes, Petitioner relies on Tivoli Redbook’s journal database as a first metabase that is similar to Chefalas’s first metabase at the source (e.g., client) of data for copying, to support copying to Tivoli Redbook’s second metabase (e.g., server). *See* Pet. 45–46 (*comparing* Ex. 1004, Fig. 5, 8:63–9:8 (first metabase), *with* Ex. 1005, 59 (second metabase), 112 (first/journal metabase); citing Ex. 1003 ¶¶ 163–164 (similar comparison—*see* metadata table reproduced above)). Dr. Jones’s testimony that the journal database does not contain metadata does not account for Petitioner’s showing related to “[t]he type of change activity that generates journal entries,” including “[o]bjects created, deleted, or renamed on the file system,” “[s]ize changes,” “[m]odification time/date changes,” “[a]ttribute changes,” and “[s]ecurity (ACL) changes.” Ex. 1005, 112 (describing the “journal database” on a client “backup candidate[]”); *see* Pet. 45–56 (describing the journal database (first metabase) (quoting Ex. 1005, 112; citing Ex. 1003 ¶¶ 163–164)); Ex. 1003 ¶ 164 (creating table as reproduced above showing metadata in respective databases, including the journal database). Also, Tivoli Redbook states that “the backup-archive client does not carry out *the initial metadata conversation*,” indicating the system transfers metadata from the client at some point. *See* Ex. 1005, 113 (emphasis added). Therefore, on this preliminary record, Petitioner shows sufficiently that at least some of Tivoli Redbook’s clients create metadata in the form of “[j]ournal database entries generated by real-time file system change activity,” in order to “specify objects to back up or expire.” *See id.*

at 112; Ex. 1003 ¶ 164 (“Tivoli Redbook . . . suggests . . . metadata is stored in its ‘journal database’: ‘Journal database entries are generated by real-time file system change activity, and they specify objects to back up or expire.’” (quoting Ex. 1005, 112) (emphasis omitted)).

To support the obviousness of copying metadata, Petitioner also refers to its similar showing with respect to limitation 1d that recites updating, which Patent Owner does not address. *See* Pet. 43. Discussing limitation 1d, the Petition quotes Tivoli Redbook, which states that

the server also looks into the database to store the metadata from this client node. Each file or object that it receives from this client node is noted in the files section with all of the relevant attributes about it. The IBM Tivoli Storage Manager database stores the file by client node, date and time, file size, file creation date and time, file modification date and time, file location structure on the client node, and other client-relevant information.

Id. at 33–34 (emphasis by Petitioner) (quoting Ex. 1005, 59). This passage and others noted below appear to describe transmitting metadata generated at the client to store at the server database.¹⁰ The Petition also explains that a person of ordinary skill in the art “would be motivated to update the metadata” in order to locate the data objects during a “disaster recovery,” otherwise “*the backup system would be unable to locate the proper data in order to perform the restoration, rendering the secondary copies useless.*”

¹⁰ This passage also appears to disclose copying “metadata from this client node.” Similarly, Tivoli Redbook describes “two sessions” between “[t]he IBM Tivoli Storage Manager server and client node . . . : *one to move metadata and the other to move data.*” Ex. 1005, 59 (emphasis added). During trial, the parties will have the opportunity to address any difference between moving or storing metadata from the client, as described by Tivoli Redbook, and copying metadata, as claimed.

Id. at 34–35 (emphasis added) (citing Ex. 1003 ¶¶ 141–142). Petitioner applies this updating motivation to copying, and the parties agree that copying includes updating. *See* Pet. 43–44; Ex. 2002 ¶ 41 (“‘Updating,’ in some circumstances, can be performed by ‘copying’ but is not limited to being performed by ‘copying.’”); Prelim. Resp. 21 (quoting Ex. 2002 ¶ 41).

Similarly, Tivoli Redbook states that during journal-based backup “[t]he backup-archive client, as always, still sends data (files) to the IBM Tivoli Storage Manager server and, as has always been the case, *the IBM Tivoli Storage Manager server populates the IBM Tivoli Storage Manager database with the file details and location.*” Ex. 1005, 113 (emphasis added); *see* Pet. 34 (quoting same; citing Ex. 1003 ¶¶ 138–140). Such populating of “file details and location” further suggests for institution purposes copying metadata associated with the file details at the second metabase, as Petitioner’s comparison to Chefalas’s Figure 5 and other metadata also suggests. *See* Pet. 45–46. And “[i]n all cases, ‘[w]hen the file or object has been committed to a storage pool device, the IBM Tivoli Storage Manager database is updated to keep track of its location in the database Tracking section.’” *Id.* at 34 (quoting Ex. 1005, 59) (emphasis omitted). On this preliminary record, Tivoli Redbook discloses updating metadata, and at least suggests copying metadata, as Petitioner argues. *See* Pet. 34–36 (updating), 43–48 (copying).

Moreover, Patent Owner’s arguments largely address the individual teachings of the references, whereas Petitioner relies on the combined teachings and knowledge of the ordinary artisan to allege the obviousness of copying metadata. As noted above, Petitioner contends that copying metadata from a first metabase to a second metabase would have been

obvious as an efficient use of existing source metadata as opposed to a re-creation of the metadata at the second database. *See* Pet. 46–47 (citing Ex. 1003 ¶ 165). Patent Owner notes that Petitioner cites Osaki as “teach[ing] a system to copy data more efficiently and that Osaki teaches reading metadata then copying to another storage system.” Prelim. Resp. 26 (citing Pet. 47). Patent Owner’s argument that “Petitioner does not argue copying metadata from a first *metabase* to a second *metabase* was well-known or conventional” mischaracterizes Petitioner’s reliance on the two references.

For example, Petitioner specifically quotes Osaki as “cop[ying the metadata from one storage system] to the secondary storage system 113 as shown by step 302.” Pet. 47 (quoting Ex. 1006 ¶ 26) (emphasis omitted). Osaki teaches primary 101 and secondary 113 storage systems including metabase 105 and metabase 705, respectively. *See* Ex. 1006, Fig. 6. Similar remarks apply to Takata. *See* Pet. 47 (contending that Takata’s system copies metadata from one logical storage unit to another (citing Ex. 1008, code (57), ¶ 173)). Patent Owner’s argument that Takata does “not apply” because “the resulting combination has different metabases on each side” (Prelim. Resp. 27) does not undermine Petitioner’s showing, because claim 1 does not preclude different metabases on each side.

On this preliminary record, Takata and Osaki reflect the knowledge of an artisan of ordinary skill and support Petitioner’s reasons for combining Chefalas and Tivoli Redbook to suggest copying. *See Paulsen*, 30 F.3d at 1480 (reasoning that prior art references must be “considered together with the knowledge of one of ordinary skill in the pertinent art.” (citing *Samour*, 571 F.2d at 562)); *KSR*, 550 U.S. at 418 (“Often, it will be necessary for a

court to look to interrelated teachings of multiple patents; the effects of demands known to the design community or present in the marketplace; and the background knowledge possessed by a person having ordinary skill in the art, all in order to determine whether there was an apparent reason to combine the known elements in the fashion claimed by the patent at issue.”).

Patent Owner also distinguishes “updating” metadata in Tivoli Redbook’s system from “copying” metadata as claimed. Prelim. Resp. 21, 33. In support, Patent Owner quotes the testimony of Dr. Jones, who testifies that updating “is not limited to being performed by ‘copying.’ If, for instance, you were to modify a document, the associated metadata might be ‘updated’ to reflect the new ‘last modified’ time. That would not suggest or entail ‘copying’ the metadata, however, in order to update it.” *Id.* at 21–22 (quoting Ex. 2002 ¶ 41). Patent Owner also characterizes as “unsubstantiated” and “errant” “Petitioner’s premise that the metadata in Chefalas’s database 412 is ‘substantially the same’ as the metadata found in Tivoli’s database.” *Id.* at 24 (citing Ex. 2002 ¶ 45; quoting Pet. 46). Patent Owner contends that “the differences between the metadata in Chefalas’s alleged [first] metabase and the metadata in Tivoli Redbook’s alleged [second] metabase are enough to offset any efficiencies hypothetically gained.” *Id.*

These arguments do not undermine Petitioner’s obviousness showing, as outlined above, which relies on the combined teachings to suggest copying the source metadata for efficiency reasons and addresses differences between updating and copying metadata. *See* Pet. 44–48. Dr. Jones’s table (reproduced above) indicates that at least some of the metadata is similar. Ex. 1003 ¶ 164. Moreover, claim 1 encompasses different first and second

metabases each with different metadata, so it is not clear how the alleged difference in some of the metadata between the Tivoli Redbook and Chefalas systems undercuts Petitioner's obviousness showing. As also indicated above, Petitioner relies on the combined teachings of the two references and the knowledge of an artisan of ordinary skill. "A person of ordinary skill is also a person of ordinary creativity, not an automaton" who would have been unable to retain efficiencies via copying metadata, a well-known practice, even if some differences in metadata teachings exist in the two references. *See KSR*, 550 U.S. at 421.

With further regard to alleged differences in metadata, Patent Owner's contention that "substantial metadata is stored in Tivoli that is not present in Chefalas's database 412" does not explain why it would not have been obvious to copy metadata from Chefalas's source database as Petitioner proposes. *See* Prelim. Resp. 24 (citing Pet. 46–47). The argument is analogous to arguing that one cannot copy a sentence from a paper into a book because the book contains more sentences than a paper.

Setting aside the analogy, Petitioner relies on Chefalas for its "disclos[ure of] querying or searching the metabase, for example, which, 'specifies an application name as well as criteria, which may be used to search for records, such as records 500, 502, and 504 in FIG. 5 within database 412'" Pet. 41 (quoting Ex. 1004, 9:9–17) (emphasis omitted). Petitioner proposes copying the associated metadata in metabase 412, as disclosed in Chefalas, "to the second metabase (*i.e.*, the IBM Tivoli Storage Manager database disclosed in *Tivoli Redbook*) for four reasons," as discussed above. *See id.* at 43–44. Contrary to the reasoning of Dr. Jones, on this preliminary record, an artisan of ordinary skill would not have

contemplated a “need to traverse Chefalas’s file system as [the Tivoli Redbook system] normally does to retrieve the data objects themselves and to retrieve the additional metadata Tivoli requires” merely to copy the metadata already existing in Chefalas’s system. *See* Prelim. Resp. 25 (quoting Ex. 2002 ¶¶ 46–48). Nothing on this preliminary record shows that Tivoli Redbook’s alleged requirement for “another process to transfer metadata” (*see id.*) to track data copying from multiple clients detracts from Petitioner’s preliminary showing of efficiency of copying existing metadata from a single client or source. *See* Ex. 2002 ¶ 46 (testifying that “Tivoli Redbook describes a commercial product which supports ‘[m]ore than 39 different operating platforms’ and has ‘[f]ull compatibility with hundreds of storage devices.’” (quoting Ex. 1005, 3)); Prelim. Resp. 24 (quoting same).

Similar remarks apply Dr. Jones’s other quoted testimony cited in support of Patent Owner’s argument that different amounts of metadata in the two systems would “offset any efficiencies hypothetically gained” in copying the metadata from Chefalas’s metabase. For example, Dr. Jones quotes Tivoli Redbook as stating that “IBM Tivoli Storage Manager backups first inspect the client system to determine *what data exists on the node*, then transmit *that metadata* to the IBM Tivoli Storage manager server and store it in the database.” Ex. 2002 ¶ 48 (emphasis by Dr. Jones) (quoting Ex. 1005, 58); Prelim. Resp. 26 (quoting same). According further to Dr. Jones, Tivoli Redbook’s process shows that “during a subsequent selective backup, the metadata associated with the identified data objects is already stored in Tivoli Redbook’s alleged secondary metabase.” Ex. 2002 ¶ 48; Prelim. Resp. 26 (quoting same).

It is not clear how this testimony undermines Petitioner's argument about the obviousness of copying metadata from Chefalas's first metabase to a second metabase as suggested by the combined teachings of the references. Rather, Patent Owner appears here to concede (through the testimony of Dr. Jones) that Tivoli Redbook's client system transmits metadata for storage into the IBM Tivoli Storage manager server database. On this preliminary record, this supports Petitioner's showing related to the obviousness of copying metadata from one storage system (e.g., client) to another (e.g., server).

Further addressing alleged differences in metadata in Chefalas and Tivoli Redbook, Dr. Jones testifies that "[t]here is no need for Chefalas to store information on the size of the file, the access control lists, or extended attributes as these attributes are not used to determine which files to back up." Ex. 2002 ¶ 45; Prelim. Resp. 23–24 (quoting same). This testimony about specific metadata in Tivoli Redbook's journal database does not address Petitioner's reliance on the combined teachings of the references, including the storing of metadata in Chefalas's system so that a user can select files to copy based on different metadata about the files. *See* Pet. 39–40 (discussing Chefalas's flow diagram of Figure 15 as teaching how to identify files based on querying database 412 using relational metadata), 41 (arguing that "Chefalas discloses querying or searching the metabase, for example, which, 'specifies an application name as well as criteria, which may be used to search for records, such as records 500, 502, and 504 in FIG. 5 within database 412 in FIG. 4'" (quoting Ex. 1004, 9:9–17) (emphasis omitted)), 45 (discussing Chefalas's Figure 5, which shows different types of metadata for locating files to be copied). Similar to Chefalas's metadata

for specifying and locating desired files, Tivoli Redbook's journal database's metadata that represents "change activity" "is configurable by the user." Ex. 1005, 112.

Patent Owner also argues that Petitioner fails to show "copying metadata associated with the identified data objects." Prelim. Resp. 29–32. Generally, Patent Owner argues that "in Tivoli Redbook, the alleged second metabase contains the metadata for all the files on the client node and, consequently, Petitioner's proposed combination does not render obvious 'copying *metadata* associated with the *identified* data objects' 'that satisfy a selection criteria.'" *Id.* at 29.

Again, this argument does not address Petitioner's reliance on Chefalas's system of using metadata to select specific files, as described in the previous paragraph, and copying that metadata. *See* Pet. 44–48. Nevertheless, to support this line of argument, Patent Owner proposes a claim construction for step 1d. Specifically, Patent Owner argues that the step of copying metadata associated with identified data objects that satisfy certain selection criteria excludes "copying *all* metadata." Prelim. Resp. 32. In a related argument, Patent Owner argues that Petitioner never addresses "the proper construction of limitation 1[d] which requires 'copying metadata *with the identified data objects*' 'that satisfy a selection criteria.'" *Id.* at 39.

Patent Owner argues that in contrast to its proposed claim construction, Tivoli Redbook "teaches that its alleged metabase contains metadata for all data on a client node," and that "Petitioner's argument thus depends on a claim construction wherein this limitation may be satisfied where 'all' metadata in an entire volume is copied without regard to whether it is 'associated with the identified data objects' 'that satisfy a selection

criteria.” Prelim. Resp. 33. Patent Owner similarly argues that Tivoli Redbook teaches retaining information about all client system and user files, not information relating to a specific subset of copied data. *Id.* at 39 (citing Ex. 2002 ¶¶ 65–66). Patent Owner also argues that Chefalas does not teach copying metadata, and the combined references do not teach step 1d. *Id.* at 40–42 (citing Ex. 2002 ¶¶ 68–69).

This preliminary record does not support these arguments. Even if the systems retain large quantities of metadata, this does not address Petitioner’s showing of copying specific metadata. As noted above, Petitioner relies on using different types of metadata to locate specific files and copying that metadata, where Chefalas and Tivoli Redbook describe flexible systems for describing metadata. *See, e.g.*, Ex. 1004, 8:67–9:1 (describing “examples of meta data, which may be stored in a data base”); Ex. 1005, 112 (“The type of change activity that generates journal entries is configurable by the user”). Alternatively, this preliminary record does not support Patent Owner’s claim construction argument that claim 1 precludes copying all metadata. *See supra* § II.C. Based on the foregoing summary and a review of the record, Petitioner’s showing as to step 1d is sufficient for purposes of institution.

e. Summary of claim 1

Considering the arguments and evidence presented by Patent Owner as summarized above, Petitioner’s showing with respect to claim 1 is sufficient for purposes of institution. Based on the foregoing discussion and a review of the preliminary record, Petitioner demonstrates a reasonable likelihood that it would prevail with respect to showing the unpatentability of claim 1.

5. *Dependent Claims 4 and 8*

Petitioner contends that the combined teachings of Chefalas and Tivoli Redbook would have rendered claims 4 and 8 obvious. Pet. 49–54. Patent Owner does not dispute these assertions or raise arguments separately from those discussed above with respect to claim 1.

Each of these claims depend from independent claim 1. Petitioner builds on its showing with respect to claim 1, providing citations to the record to support its allegations of obviousness with respect to claims 4 and 8. Pet. 49–54. For reasons similar to those discussed above in connection with claim 1, and based on the Petition’s obviousness showing with respect to claims 4 and 8, the Petition establishes a reasonable likelihood of prevailing in showing that the combined teachings of Chefalas and Tivoli Redbook would have rendered claims 4 and 8 obvious.

III. CONCLUSION

After considering the evidence and arguments presented in the Petition and the Preliminary Response, we determine that Petitioner demonstrates a reasonable likelihood that it would prevail with respect to its unpatentability challenges. Accordingly, we institute an *inter partes* review on the challenged claims on the sole ground presented in the Petition. At this stage of the proceeding, we have not made a final determination as to the patentability of the challenged claims.

IV. ORDER

Accordingly, it is

ORDERED that pursuant to 35 U.S.C. § 314, *inter partes* review is instituted as to the challenged claims of the ’671 patent with respect to the sole ground of unpatentability presented in the Petition; and

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

For PETITIONER:

David Cochran
Matthew Johnson
Marlee Hartenstein
JONES DAY
dcochran@jonesday.com
mwjohnson@jonesday.com
mhartenstein@jonesday.com

For PATENT OWNER:

Kenneth Weatherwax
Bridget Smith
Flavio Rose
Parham Hendifar
Patrick Maloney
Jason Linger
LOWENSTEIN & WEATHERWAX LLP
weatherwax@lowensteinweatherwax.com
smith@lowensteinweatherwax.com
rose@lowensteinweatherwax.com
hendifar@lowensteinweatherwax.com
maloney@lowensteinweatherwax.com
linger@lowensteinweatherwax.com