Trials@uspto.gov
Tel: 571-272-7822

Paper 22
Entered: January 29, 2015

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

RACKSPACE HOSTING, INC.,
Petitioner,

v.

CLOUDING CORP.,
Patent Owner.

Case IPR2013-00519
Patent 7,032,089 B1

Before JAMESON LEE, JONI Y. CHANG, and DAVID C. MCKONE, *Administrative Patent Judges.*

MCKONE, *Administrative Patent Judge*.

FINAL WRITTEN DECISION
35 U.S.C. § 318(a) and 37 C.F.R. § 42.73

I. INTRODUCTION

A. Background

Rackspace Hosting, Inc. ("Petitioner") filed a Petition (Paper 2, "Pet.") to institute an *inter partes* review of claims 1–24 of U.S. Patent No. 7,032,089 B1 (Ex. 1001, "the '089 patent"). Clouding Corp.[1] ("Patent Owner") filed a Preliminary Response (Paper 5, "Prelim. Resp."). Pursuant to 35 U.S.C. § 314, in our Decision to Institute (Paper 6, "Dec."), we instituted this proceeding as to all of the challenged claims of the '089 patent.

During this trial, Patent Owner filed a Patent Owner Response (Paper 10, "PO Resp.") and Petitioner filed a Reply to the Patent Owner Response (Paper 12, "Reply"). An oral hearing was held on September 24, 2014 (Paper 21, "Tr.").

We have jurisdiction under 35 U.S.C. § 6(c). This decision is a final written decision under 35 U.S.C. § 318(a) as to the patentability of the challenged claims. Based on the record before us, Petitioner has demonstrated by a preponderance of the evidence that claims 1–24 are unpatentable.

[1] The Petition named Clouding IP, LLC, as the patent owner. Clouding Corp., however, represents that it obtained the '089 patent from Clouding IP, LLC. Paper 18, at 1.

B. Related Proceedings

Patent Owner has asserted the '089 patent against Petitioner in *Clouding IP, LLC v. Rackspace Hosting, Inc.*, Case No. 1:12-cv-00675 (D. Del.). Pet. 2–3; Paper 4, at 1–2.

Patent Owner has asserted the '089 patent in six additional proceedings, against other parties, in the United States District Court for the District of Delaware. Pet. 3; Paper 4, at 1–2.

C. References Relied Upon

Petitioner relies upon the following prior art references:

Gold (Ex. 1003)	WO 99/12098	Mar. 11, 1999
Pruett (Ex. 1005)	US 5,778,389	July 7, 1998

D. Grounds of Unpatentability

We instituted this proceeding based on the grounds of unpatentability set forth in the table below. Dec. 28.

References	Basis	Claims challenged
Gold	35 U.S.C. § 102(b)	1–24
Pruett	35 U.S.C. § 102(b)	1–3, 6, 7, 13, 15, 16

E. The '089 Patent

The '089 patent generally relates to synchronizing computer data replicated in different storage areas for backup and disaster recovery purposes. Ex. 1001, 1:15–18. For example, the '089 patent describes a system with a primary node (e.g., a computer) coupled by a network to a

secondary, or backup, node. *Id.* at Figs. 1, 3. One feature of this system is that the secondary node includes the same data as the primary node, such that, if the primary node fails, the secondary node, using the replicated data, can take over. *Id.* at 5:36–41, 6:56–62.

The '089 patent describes a "copy on read" replication technique that copies data from the primary node to the secondary node after the data is read from the primary node's storage area. *Id.* at 8:62–65. Figure 3, reproduced below, illustrates an example:

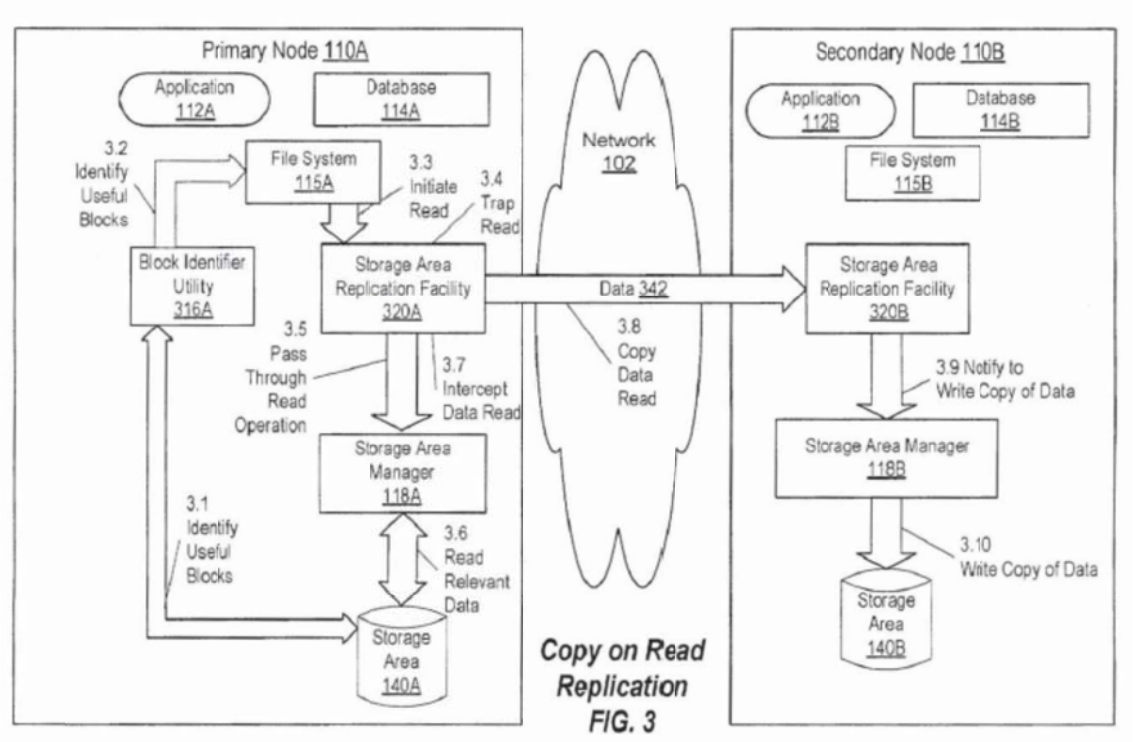

Figure 3 is a flow diagram of a copy-on-read operation used for replication. *Id.* at 4:17–19.

According to the '089 patent, not all data in the storage area of the primary node is "relevant" to an application managing the data on the secondary node. *Id.* at 8:12–17. To avoid replicating irrelevant data, the primary node of Figure 3 includes block identifier 316A at primary node

110A that identifies "useful" blocks of data ("storage objects") for copying to the secondary node *Id.* at 8:18–27.

Block identifier 316A identifies relevant blocks to read and sends a message to storage area manager 118A, which reads the relevant data from storage area 140A. *Id.* at 9:24–47. Storage area replication facility 320A intercepts this message and allows the message to pass through to storage area manager 118A, which retrieves the data from the primary node's storage area 140A. *Id.* at 9:48–55. After storage area manager 118A retrieves the relevant data, the data is intercepted by storage area replication facility 320A and sent to storage area replication facility 320B at secondary node 110B. *Id.* at 9:56–59.

According to the '089 patent, the block to be copied to the secondary node can be compared to a corresponding block already stored on the secondary node's storage area, for example, by comparing their checksums. *Id.* at 10:46–56. If the checksums match, there is no need to copy the data—it already is replicated. *Id.* at 10:59–62.

Claim 1, reproduced below, is illustrative of the claimed subject matter:

> 1. A method comprising:
>
> selecting a first storage object of a first storage area;
>
> determining whether the first storage object contains relevant data;
>
> if the first storage object contains relevant data, performing the following:
>
> reading the first storage object; and

when the reading is completed, copying contents of the first storage object to a second storage object of a second storage area; and

if the first storage object does not contain relevant data,

selecting a second storage object of the first storage area without reading the first storage object and without copying the contents of the first storage object to the second storage object.

II. ANALYSIS

A. Claim Construction

The Board interprets claims of an unexpired patent using the broadest reasonable construction in light of the specification of the patent in which they appear. *See* 37 C.F.R. § 42.100(b).[2] Claim terms generally are given

[2] Patent Owner contends that we should apply the claim construction framework used by district courts rather than the broadest reasonable interpretation. PO Resp. 4–5 n.1. Patent Owner argues that applying the standard set forth in *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc), would result in adoption of its proposed construction of "relevant data." *Id.* Patent Owner's recitation of the claim construction standard for *inter partes* review is incorrect. According to our rules, the Board interprets claims of an unexpired patent using the broadest reasonable construction in light of the specification of the patent in which they appear. *See* 37 C.F.R. § 42.100(b). Moreover, even if we applied the *Phillips* standard, Patent Owner's proposed construction or "relevant data" still would be inappropriate. As we explain below, the Specification of the '089 patent provides examples in which relevant data are data operated on by an application and, thus, relevant data are not restricted to data that an application uses to function. *See Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1584 (Fed. Cir. 1996) (A claim construction that excludes a preferred embodiment "is rarely, if ever, correct and would require highly persuasive evidentiary support.").

their ordinary and customary meaning, as would be understood by one of ordinary skill in the art in the context of the entire disclosure. *See In re Translogic Tech., Inc.*, 504 F.3d 1249, 1257 (Fed. Cir. 2007).

1. "relevant data" / "determin[e/ing] whether the first storage object contains relevant data"

Claim 1 recites "determining whether [a] first storage object contains *relevant data*" (emphasis added). Independent claims 7, 9, 13, 17, and 21 include similar recitations. In the Decision to Institute, we preliminarily construed "relevant data" to mean data relevant to the purpose for which storage objects are being copied. Dec. 7.

In its Response, Patent Owner proposes construing "relevant data" to mean "data relevant with respect to or in the context of an application that will make use of the data." PO Resp. 8. According to Patent Owner, several passages in the '089 patent's Specification demonstrate that relevance is measured in the context of an application that will make use of the data, rather than with respect to the purpose for which the data is being copied, as we preliminarily construed the term. *Id.* For example, Patent Owner points to the '089 patent's description that block identifier utility 316A can be implemented with a "'replication-aware' application enhanced to read *all relevant blocks necessary for operation of the application*, such as an enhanced version of application 112A, database 114A, or file system 115." Ex. 1001, 9:6–10 (emphasis added). The passage continues by describing that "if the application reading the data is a file system, a 'file dump' utility may exist that reads *all relevant file data* and provides the contents of those

blocks directly to storage area replication facility 120A." *Id.* at 9:14–17 (emphasis added).

Asked at the oral hearing to differentiate its proposed construction from our preliminary construction, Patent Owner argued that relevant data are data that an application relies upon for its operation rather than the data it operates upon. Tr. 33:23–34:15. Patent Owner made this distinction in addressing Gold, which, as explained below, describes a file backup utility. According to Patent Owner, under its proposed construction, relevant data are those data that the backup utility needs to operate while, under our preliminary construction, relevant data are limited to the files that are backed up by the utility. *Id.*

Patent Owner's characterization of our preliminary construction is incorrect. Data relevant to the purpose for which storage objects are being copied does not exclude either the data relevant to the operation of an application or the data operated on by the application. Moreover, Patent Owner's distinction is not captured in its proposed construction. "Data relevant with respect to or in the context of an application that will make use of the data," by its terms, covers both data used by an application to function and the data on which the application functions. Thus, we do not see a meaningful distinction between Patent Owner's proposed construction and our preliminary construction.

In any case, we disagree that the Specification limits the claims in the way argued by Patent Owner. For example, the Specification includes several other passages discussing "relevant data," including:

> "Storage objects containing data and information *relevant to managing the data by a particular application . . . .*" Ex. 1001, 4:65–66 (emphasis added);
>
> "Not all of the data in each block is necessarily *relevant for operation of the application managing the data* on a remote node." *Id.* at 8:13–15 (emphasis added);
>
> "In 'Identify Primary Storage Object with Useful Contents' step 510, an application or utility capable of identifying *storage objects with data or information used for managing the data* identifies a storage object with useful contents." *Id.* at 10:42–46 (emphasis added);
>
> "Establishing a 'copy on-read' mode in the storage area replication facility is the first step; however, a utility must be used that will *identify relevant data* blocks. For example, some file systems include utilities that walk through all the on-disk data structures, *but most of them will not necessarily read every useful data block of a file*. Some file systems include a 'dump' utility that traverses all the data structures as well as the data blocks, so such utilities can be good candidates for identifying relevant blocks." *Id.* at 11:27–35 (emphases added).

These passages describe various different examples of relevant data, including data relevant to an application's operation, data relevant to an application's management of data, and the data operated on by an application (traversed by a file dump utility). In another example, the '089 patent explains that, "[i]n 'Identify Storage Object with Useful Contents' step 410, an application or utility capable of identifying storage objects of data or information used for managing the data identifies a useful storage

object." *Id.* at 10:16–19. Here, useful data[3] are either data the application needs to function or data on which the application operates.

Thus, assuming Patent Owner's proposal does exclude the data operated on by an application, it is too narrow in light of the Specification's examples in which relevant data are the files traversed (operated on) by a utility rather than the data necessary for the utility itself to function. Ex. 1001, 11:27–35. Accordingly, we maintain our construction of "relevant data" as "data relevant to the purpose for which storage objects are being copied."

2. "storage area"

Claim 1 recites "selecting a first storage object of a first *storage area*" and "copying contents of the first storage object to a second storage object of a second *storage area*" (emphases added). Independent claims 7, 9, 13, 17, and 21 include similar recitations. In the Decision to Institute, we construed "storage area" as "a location in which data is stored," and made clear that "storage area" is not limited to "a storage volume," and that it is distinct

[3] The parties continue to dispute whether "relevant data" and "useful data" are synonymous. PO Resp. 9–11; Reply 3–6. We note that, in identifying structure corresponding to "determining means for determining whether the first storage object contains relevant data," Patent Owner identifies disclosure alternately describing a block identifier utility performing an algorithm to identify "relevant blocks," "useful blocks," and "relevant, useful blocks." PO Resp. 15 (citing Ex. 1001, 8:25–27, 9:11–13, 9:37–44). These are just some of the examples supporting our conclusion (Dec. 6) that the '089 patent uses "relevant" and "useful" synonymously, a conclusion we maintain. Neither party, however, adequately explains why this distinction, or lack thereof, has any bearing on the issues in this case.

from a file system. Dec. 7–8. Petitioner and Patent Owner do not dispute this construction. We see no reason to change this construction based on the full record and maintain it for purposes of this Decision.

3. "copy-on-read mode"

The term "copy-on-read mode" appears in dependent claims 10, 14, 18, and 22. We preliminarily construed "copy-on-read mode" to mean "a mode where data are read from one storage area and subsequently copied to another storage area." Petitioner and Patent Owner do not dispute this construction. We see no reason to change this construction based on the full record and maintain it for purposes of this Decision.

4. Means-Plus-Function Terms

Claims 9–12 recite claim elements in "means-plus-function" format. "An element in a claim for a combination may be expressed as a means . . . for performing a specified function without the recital of structure, material, or acts in support thereof, and such claim shall be construed to cover the corresponding structure, material, or acts described in the specification and equivalents thereof." 35 U.S.C. § 112, ¶ 6.[4] The United States Court of Appeals for the Federal Circuit stated: "Section 112, ¶ 6 recites a mandatory procedure for interpreting the meaning of a means- or step-plus-function claim element. These claim limitations 'shall be construed to cover the corresponding structure, material, or acts described in the specification and equivalents thereof.'" *Al-Site Corp. v. VSI Int'l, Inc.*, 174 F.3d 1308, 1320

[4] Now recodified as 35 U.S.C. § 112(f).

(Fed. Cir. 1999); *see also In re Donaldson Co., Inc.*, 16 F.3d 1189, 1193–94 (Fed. Cir. 1994) ("[P]aragraph six applies regardless of the context in which the interpretation of means-plus-function language arises, i.e., whether as part of a patentability determination in the PTO or as part of a validity or infringement determination in a court."). We construe such a limitation by determining what the claimed function is and identifying the structure or materials disclosed in the specification that correspond to the means for performing that function. *See Kemco Sales, Inc. v. Control Papers Co., Inc.*, 208 F.3d 1352, 1360 (Fed. Cir. 2000).

For computer-implemented inventions, this corresponding structure must be more than a general purpose computer or microprocessor. *See Aristocrat Techs. Australia Pty Ltd. v. Int'l Game Tech.*, 521 F.3d 1328, 1333 (Fed. Cir. 2008). Rather, "[w]hen dealing with a 'special purpose computer-implemented means-plus-function limitation,' [the Federal Circuit] require[s] the specification to disclose the algorithm for performing the function." *Function Media, L.L.C. v. Google, Inc.*, 708 F.3d 1310, 1318 (Fed. Cir. 2013); *accord Blackboard, Inc. v. Desire2Learn, Inc.*, 574 F.3d 1371, 1384 (Fed. Cir. 2009) ("[W]hen a computer is referenced as support for a function in a means-plus-function claim, there must be some explanation of how the computer performs the claimed function.").

Petitioner argues that the Specification of the '089 patent does not identify particular structure corresponding to the functions recited in the means-plus-function limitations of claims 9–12. *Id.* Petitioner then identifies other claims as providing the structure corresponding to these limitations. Pet. 12–15 (citing to claims 13, 17, and 21 for disclosure of structure corresponding to "selecting means," "determining means,"

"reading means," and "copying means"; citing to claims 15, 19, and 23 for disclosure of structure corresponding to "sending means"; and citing to claims 16, 20, and 24 for disclosure of structure corresponding to "second copying means").

We invited Patent Owner to address the issue in its PO Response. Dec. 10–11. Patent Owner responded with a more detailed identification of structure for these terms and supported it with the testimony of Dr. Prasant Mohapatra, a Professor in the Department of Computer Science at the University of California, Davis (Ex. 2002, "Mohapatra Decl."). PO Resp. 11–21. We address below Patent Owner's evidence for each of the means-plus-function terms.

In reply, Petitioner argues that Patent Owner's proposed constructions of means-plus-function terms "are improper because they do not point to any algorithm for making a relevancy determination." Reply 6. Petitioner, however, does not address any means-plus-function term with specificity or offer any evidence rebutting Patent Owner's proposals. At the oral hearing, Petitioner admitted that it did not introduce any evidence on whether a person of ordinary skill in the art would have considered the corresponding structure identified by Patent Owner to be sufficient. Tr. 24:6–21.

a. Two instances of "selecting means"

There are two instances of "selecting means" in claim 9. The function of the first "selecting means" is "selecting a first storage object of a first storage area." We proceeded on the basis that the first instance of "selecting means" is computer software executing on a general purpose computer implementing a block identifier utility and equivalents thereof. Dec. 11.

Patent Owner identifies block identifier 316A as corresponding to the selecting means. PO Resp. 12. According to Patent Owner and its Declarant, block identifier 316A receives, from an application program, database, or file system, a command for selecting a storage object and, in response, identifies a relevant block of data. *Id.* (citing Ex. 2002 ¶ 18). Patent Owner argues that a described "file dump" utility and a replication–aware database are specific examples using block identifiers that systematically read relevant blocks of data from storage areas. PO Resp. 12–13 (citing Ex. 2002 ¶ 18). Patent Owner further points to action 3.2 ("Identify Useful Blocks) of Figure 3, during which "block identifier utility 316A notifies file system 115A of the useful blocks." Ex. 1001, 9:45–47. We note that action 3.2 follows action 3.1, in which "block identifier 316A identifies useful blocks stored in storage area 140A." *Id.* at 9:24–25. According to the '089 patent:

> While block identifier utility 316A is shown as directly accessing storage area 140A to perform this identification, one of skill in the art will understand that several intermediate steps may be performed to provide this functionality. For example, typically a program reading a data block will call an interface to a storage area manager, such as storage area manager 118A, which deals with directly accessing the physical device.

Id. at 9:25–32. As Patent Owner points out, block identifier 316A is described as a module associated with a computer system that has a processor executing software instructions. PO Resp. 13 (citing Ex. 1001, 6:29–32). Accordingly, we agree with Patent Owner and Dr. Mohapatra that the block identifier utility 316A is described in terms of specific algorithms implemented as instructions executed on a computer system for selecting a storage object from a storage area.

The function of the second instance of "selecting means" is "selecting a second storage object of the first storage area without reading the first storage object and without copying the contents of the first storage object to the second storage object if the first storage object does not contain relevant data." We proceeded on the basis that the second instance of "selecting means" is computer software executing on a general purpose computer implementing a block identifier utility and equivalents thereof. Dec. 12.

Patent Owner identifies block identifier 316A as corresponding to the second selecting means for substantially the same reasons as it advanced for the first selecting means. PO Resp. 14. For the same reasons, we agree.

Accordingly, we maintain our constructions of the first instance of "selecting means" and the second instance of "selecting means."

b. "determining means"

The function of the "determining means" is "determining whether the first storage object contains relevant data." We proceeded on the basis that "determining means" is computer software executing on a general purpose computer implementing a block identifier utility and equivalents thereof. Dec. 13.

Patent Owner and its declarant contend that the act of determining is the same as the act of identifying and that block identifier utility 316A corresponds to a determining means. PO Resp. 15 (citing Ex. 2002 ¶ 20). According to Patent Owner and its declarant, block identifier utility 316A identifies relevant blocks by directly accessing a physical storage device or calling an interface to a storage area manager, which access the physical

storage device. PO Resp. 15–16 (citing Ex. 1001, 9:11–13, 9:25–32); Ex. 2002 ¶ 21.

The '089 patent specifically identifies block identifier utility 316A as corresponding to a determining means: "Block identifier utility 316A is representative of an identifier module, means or instructions for identifying relevant, useful blocks to a module reading the data." Ex. 1001, 8:25–27. As explained above, in action 3.1 of Figure 3, "block identifier utility 316A identifies useful blocks stored in storage area 140A." Ex. 1001, 9:24–25. The Specification further explains that

> The identification or useful blocks is typically performed in response to a user command. A user command may be issued by a person or by an application providing a user interface. For example, a user interface may be provided to block identifier utility 316A and/or storage area replication facility 320A. A user command may start the resynchronization process, without necessarily requiring the user to be aware of the underlying implementation details.

Ex. 1001, 9:37–44. As explained above, action 3.1 is followed by action 3.2, in which "block identifier utility 316A notifies file system 115A of the useful blocks," after which a file system initiates a read operation. *Id.* at 9:44–46. This is further described in conjunction with Figure 4. Specifically, "[i]n 'Identify Storage Object with Useful Contents' step 410, an application or utility capable of identifying storage objects of data or information used for managing the data identifies a useful storage object." *Id.* at 10:16–19. As Patent Owner points out, block identifier 316A is described as a module associated with a computer system that has a processor executing software instructions. PO Resp. 16 (citing Ex. 1001, 6:29–32). Thus, we agree with Patent Owner and Dr. Mohapatra that the

block identifier utility 316A is described in terms of specific algorithms implemented as instructions executed on a computer system for determining whether a storage object contains relevant data.

Accordingly, we maintain our construction of "determining means."

c. "reading means"

The function of the "reading means" is "reading the first storage object if the first storage object contains relevant data." We proceeded on the basis that "reading means" is computer software executing on a general purpose computer implementing an application, database, file system, and/or storage area manager, and equivalents thereof. Dec. 13–14.

Patent Owner identifies application 112A, database 114A, file system 115A, and/or storage manager 118A as corresponding to a reading means. PO Resp. 16–17 (citing Ex. 2002 ¶ 22). The '089 patent specifically identifies these components as corresponding to a reading means:

> application 112A, database 114A, and/or file system 115A can read the data (in conjunction with storage area manager 118A), either directly or via one of the other components. For example, application 112A may use both file system 115A and storage area manager 118A to read the data. Therefore, application 112A, database 114A, file system 115A and/or storage area manager 118A can also be considered to form a reading module, means, or instructions.

Ex. 1001, 8:28–36.

As shown in Figure 3, at action 3.2, block identifier utility 316A notifies file system 115A that relevant blocks in storage area 140A have been identified. Ex. 1001, 9:45–46. At action 3.3, file system 115A initiates a read operation on the blocks of memory identified by block identifier

utility 316A. *Id.* at 9:46–47. In action 3.6, storage area manager 118A reads the identified relevant blocks from storage area 140A. *Id.* at 53–55. As Patent Owner points out, each of these components is described as a module associated with a computer system that has a processor executing software instructions. PO Resp. 17–18 (citing Ex. 1001, 6:29–32). Thus, we agree with Patent Owner and Dr. Mohapatra that application 112A, database 114A, file system 115A, and/or storage manager 118A are described in terms of specific algorithms implemented as instructions executed on a computer system for reading a first storage object if the first storage object contains relevant data. Accordingly, we maintain our construction of "reading means."

d. "copying means"

The function of the "copying means" is "copying contents of the first storage object to a second storage object of a second storage area if the first storage object contains relevant data." We proceeded on the basis that "copying means" is computer software executing on a general purpose computer implementing a storage area replication facility and equivalents thereof. Dec. 14.

Patent Owner argues that storage area replication facilities 320A and/or 320B correspond to the "copying means." PO Resp. 18. The '089 patent specifically identifies these components as corresponding to the copying means: "Either or both or storage area replication facility 320A and 320B are representative of a copying module, means or instructions used to replicate data to a secondary node." Ex. 1001, 8:51–53. We agree with Patent Owner and its declarant (PO Resp. 18–19 (citing Ex. 2002 ¶ 23)),

that, in connection with actions 3.4–3.9 of Figure 3, the '089 patent describes storage area replication facilities 320A and 320B in terms of an algorithm for copying contents of a first storage object to a second storage object:

> [I]n action 3.4, storage area replication facility 320A intercepts, or traps, the read operation because storage area replication facility 320A is operating in "copy on read" mode. In action 3.5, storage area replication facility 320A allows the read operation to pass through to storage area manager 118A. In action 3.6, storage area manager [118A] reads the data from the identified relevant blocks from storage area 140A.
>
> In action 3.7, the data read (data 342) are intercepted by storage area replication facility 320A. In action 3.8, storage area replication facility 320A provides data 342 to storage area replication facility 320B on secondary node 110B. In action 3.9, storage area replication facility 320B on secondary node 110B notifies storage area manager 118B on secondary node 110B to write the copy of the data read to storage area 140B on secondary node 110B.

Id. at 9:48–63. Accordingly, we maintain our construction of "copying means."

e. "sending means"

Claim 11 depends from claim 9 and recites "sending means." The function of the sending means is "sending the contents over a network connection from the first storage area to the second storage area." We proceeded on the basis that "sending means" is computer software executing on a general purpose computer implementing a storage area replication facility and equivalents thereof. Dec. 15.

Patent Owner identifies storage area replication facility 320A as corresponding to the sending means. PO Resp. 20. As explained above, storage area replication facility 320A is described in terms of an algorithm in conjunction with actions 3.4–3.9, depicted in Figure 3. Specifically, in action 3.8, "storage area replication facility 320A provides data 342 to storage area replication facility 320B on secondary node 110B." Ex. 1001, 9:57–59. According to Dr. Mohapatra, Figure 3 depicts storage area replication facility 320A sending the data over a network. Ex. 2002 ¶ 24.

We agree with Patent Owner and Dr. Mohapatra that a skilled artisan would understand storage area replication facility 320A to be described in terms of specific algorithms implemented as instructions executed on a computer system for sending the contents of a storage object over a network connection from a first storage area to a second storage area. Accordingly, we maintain our construction of "sending means."

f. "second copying means"

Claim 12 depends from claim 9 and recites "second copying means." The function of the second copying means is "copying second contents of the first storage object to the second storage object if the second contents are relevant data." We proceeded on the basis that "second copying means" is computer software executing on a general purpose computer implementing a storage area replication facility and equivalents thereof. Dec. 15–16.

Citing to its Declarant, Patent Owner contends the second copying means corresponds to the same structure as the copying means recited in claim 9. PO Resp. 21 (citing Ex. 2002 ¶ 25). We agree that both copying

means correspond to the same structure recited in the Specification and, accordingly, maintain our construction of "second copying means."

5. "Modules" and "Instructions" Terms

Claim 13 recites a "selecting module," "determining module," "reading module," "copying module," and second instance of "selecting module." Claim 15 recites a "sending module." Claims 17 and 21 each recite "selecting instructions," "determining instructions," "reading instructions," "copying instructions," and second instances of "selecting instructions." Claims 19 and 23 each recite "sending instructions."

We preliminarily construed each "module" of claims 13–16 to be a computer software component, implemented on a computer, configured to carry out the function recited in the claim element. For example, "a selecting module configured to select a first storage object of a first storage area" is a computer software component, implemented on a computer, configured to select a first storage object of a first storage area. Similarly, we preliminarily construed each set of "instructions" of claims 17–24 to be software code configured to carry out the function recited in the claim element. For example, the "selecting instructions configured to select a first storage object of a first storage area" of claims 17 and 21 are software code configured to select a first storage object of a first storage area.

Petitioner and Patent Owner do not dispute these constructions. We see no reason to change these constructions based on the full record and maintain them for purposes of this Decision.

B. Asserted Grounds of Unpatentability

1. Anticipation by Gold

Petitioner asserts that Gold anticipates each of claims 1–24. Pet. 15–16. Petitioner supports its Petition with the testimony of Dr. Jon B. Weissman, a Computer Science professor at the University of Minnesota (Ex. 1009, "Weissman Decl."). For the reasons below, Petitioner has shown, by a preponderance of the evidence, that Gold discloses each limitation of claims 1–24.

a. Overview of Gold

Gold describes a technique for backing up computer files stored on a client computer to tape media connected to the client by a network. Ex. 1003, p. 1, Abstract. Figure 3, reproduced below, illustrates an example:

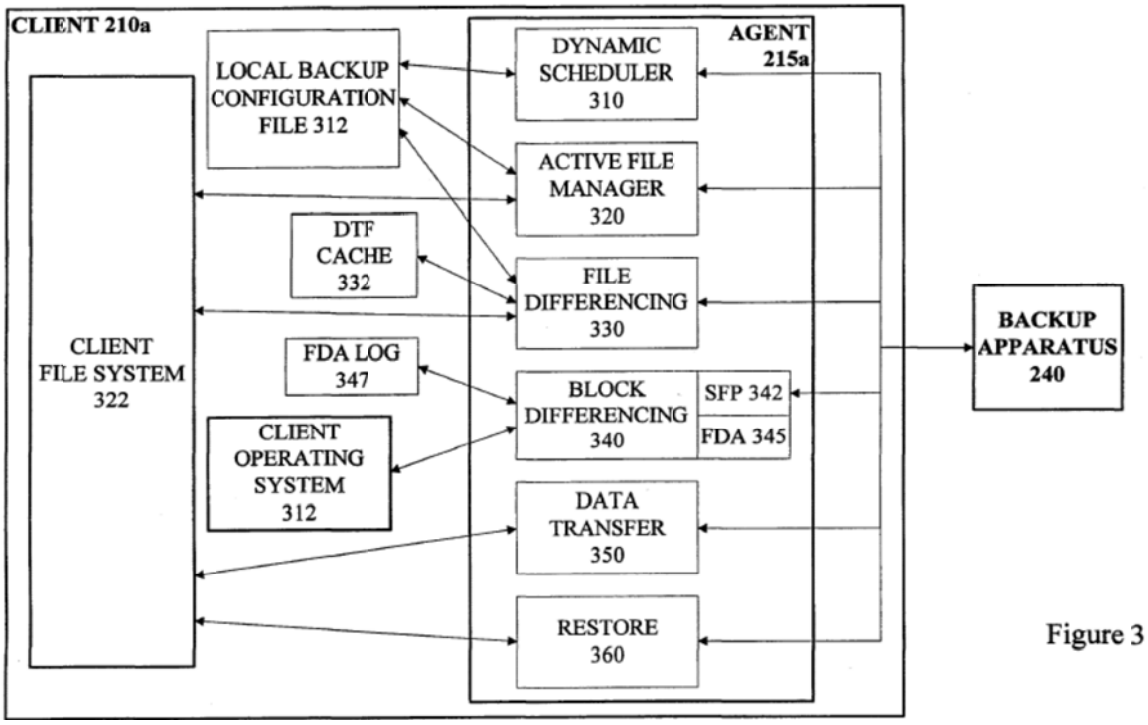

Figure 3

Figure 3 is a block diagram of a client. *Id.* at 3:15–16.

As shown in Figure 3, client 210a includes dynamic scheduler module 310 that assesses the client's resources at a scheduled target time to determine if a client backup can run without impacting seriously the system performance. *Id.* at 1:24–31. If a backup can be performed, file differencing module ("FDM") 330 selects the files on the client that will be backed up by determining which files have changed or have been added since a previous backup. *Id.* at 7:9–10. FDM 330 identifies files by reading a current directory tree of local client file system 332, which includes entries identifying each file's time and date of modification. *Id.* at 7:11–13. Those times and dates are compared to cached times and dates stored in a directory tree file. *Id.* FDM 330 also can filter out files, such as temporary files, Internet cache files, and swap files, that do not need to be backed up. *Id.* at 7:16–19.

FDM 330 also can communicate with tape backup apparatus 240 to determine whether a file already is backed up on the tape media (and, thus, backup would be redundant). *Id.* at 7:20–23. FDM 330 sends tape backup apparatus 240 a list of files to be backed up, along with calculated cyclic redundancy check ("CRC") codes (corresponding to file name, date and time stamp, and file size information). *Id.* at 7:23–25. The tape backup apparatus returns a list of the files it contains that match those on FDM 330's list, along with CRC checksums (corresponding to the actual data in the files). *Id.* at 7:26–28. For each file on tape backup apparatus 240's list that might be redundant, FDM 330 calculates a CRC checksum. *Id.* at 7:28–30. If the checksums match, the same version of the same file is already backed up on the tape media, and backing it up again would be unnecessary. *Id.*

b. Anticipation of Claims 1–24

We find that Gold's hard disk drive, on which local file system (client file system 322) resides, is a "first storage area" and that Gold's description of FDM 330 selecting files to be backed up by determining which files have changed or have been added, constitutes a disclosure of "selecting a first storage object of a first storage area," as recited in claims 1 and 7. Ex. 1003, 5:9–10, 7:9–11. Similarly, we find that Gold discloses a "selecting means for selecting a first storage object of a first storage area," recited in claim 9; a "selecting module configured to select a first storage object of a first storage area," recited in claim 13; and "selecting instructions configured to select a first storage object of a first storage area," recited in claims 17 and 21.

Regarding the means-plus-function limitations of claims 9–12, Patent Owner confirmed at the oral hearing that it has not introduced evidence or argument disputing that Gold discloses the particular algorithms corresponding to those limitations, other than its arguments and evidence directed to Gold failing to disclose "relevant data," as discussed below. Tr. 52:24–54:21. We find that the structure and algorithms in Gold identified by Petitioner as disclosing the method steps of claims 1–3 and the structural elements of claims 17–20 (detailed below) also are the same as the structure and algorithms corresponding to the analogous means-plus-function limitations of claims 9–12. *See* Pet. 25–29.

The parties dispute whether Gold discloses "determining whether the first storage object contains relevant data," as recited in claims 1 and 7, and recited similarly in the "determining means" of claim 9, the "determining module" of claim 13, and the "determining instructions" of claims 17 and

21. Petitioner contends that Gold's description of filtering a list of files for excluded files (e.g., temporary files and Internet cache files) discloses this limitation. Pet. 16. Specifically, Dr. Weissman testifies that filtering out any such files that do not need to be backed up constitutes filtering out irrelevant files. Ex. 1009 ¶ 34.

Patent Owner makes two arguments regarding "relevant data." First, Patent Owner contends that Gold does not disclose a system in which a storage object is read only after it is identified as relevant. PO Resp. 22–26. Second, Patent Owner contends that Gold's description of filtering out files that will not be backed up does not constitute determining whether storage objects contain relevant data. PO Resp. 22–26.

Regarding the second argument, Patent Owner argues that Gold's filter does not take into account the context of an application that will make use of the files. *Id.* at 27. According to Patent Owner, determining whether to back up a file based on whether it has changed since a previous back up does not take into account the proper context. *Id.* at 28. At the oral hearing, Patent Owner clarified its distinction:

> what is relevant to the backup application is the data that it relies upon for its operation rather than the data it operates upon. The data that it operates upon is merely . . . what is being backed up. Whereas the data that is relevant for it would be its operating parameters, for example.

Tr. 33:25–34:5. However, as explained in Section II.A.1, "relevant data," or "data relevant to the purpose for which storage objects are being copied," includes data used by an application to function and data upon which an application operates. Thus, we are not persuaded by Patent Owner's second argument.

Regarding Patent Owner's first argument, Patent Owner contends that Gold describes reading a file to determine whether it should be backed up. Thus, assuming that determining whether a filed should be backed up is a relevancy determination about a storage object, such a determination is not made prior to reading the storage object. PO Resp. 23–25. At the oral hearing, Patent Owner argued that Petitioner's proposed application of Gold to claim 1 "would run afoul . . . of the prohibition on reading recited later in the claim if relevance is not determined." Tr. 37:9–11. Presumably, Patent Owner refers to the claim language "if the first storage object does not contain relevant data, selecting a second storage object of the first storage area *without reading the first storage object*," recited in claim 1 (emphasis added), and similar language recited in claims 7, 9, 13, 17, and 21.

Petitioner disputes that Gold discloses reading a file to determine whether to back it up. Instead, Petitioner argues, Gold discloses reading file attributes in a directory, similarly to the way described in the '089 patent. Reply 7–9 (citing Ex. 1003, 7:9–13; Ex. 1001, 9:24–36). We agree. In determining whether to back up a file, Gold "achieves this by reading the current directory tree of the local file system 322 and checking each file's modified time/date against the entries in a cached Directory Tree File (DTF) 332 generated from the last backup." Ex. 1003, 7:11–13. The '089 patent determines relevancy of a file in the same manner: "block identifier utility 316A is shown as directly accessing storage area 140A to perform this identification." Ex. 1001, 9:25–27.

In any case, Petitioner argues, the claims of the '089 patent are open-ended "comprising" claims that do not preclude reading a storage object to determine its relevance. Reply 9–10. According to Petitioner, if the

determination to back up a file includes a read of the file, the file is read a second time when a data transfer module transfers the backup data to a tape backup apparatus. *Id.*; Pet. 18 (citing Ex. 1003, 9:13–14; Ex. 1009 ¶ 37).

We agree with Petitioner. "[T]he word 'comprising' in the transitional phrase of a patent claim creates a presumption that the body of the claim is open." *Crystal Semiconductor Corp. v. TriTech Microelectronics Int'l, Inc.*, 246 F.3d 1336, 1350 (Fed. Cir. 2001). The Federal Circuit has cautioned, however, that "'[c]omprising' is not a weasel word with which to abrogate claim limitations." *Spectrum Int'l, Inc. v. Sterilite Corp.*, 164 F.3d 1372, 1380 (Fed. Cir. 1998). In this case, the claim language "if the first storage object does not contain relevant data, selecting a second storage object of the first storage area without reading the first storage object" pertains to the behavior of the claimed method after a relevancy determination has been made. Patent Owner does not explain persuasively how this language addresses the behavior of the method prior to a relevancy determination. We conclude that nothing in claim 1 precludes reading before or during a relevancy determination. After considering Patent Owner's arguments, we nevertheless are persuaded that Gold discloses "determining whether the first storage object contains relevant data," as recited in claims 1 and 7 and the similar limitations in claims 9, 13, 17, and 21. *See* Ex. 1003, 7:16–17; Ex. 1009 ¶ 34.

We also find that Gold discloses, after this determination, "if the first storage object contains relevant data . . . reading the first storage object," as recited in claim 1. *See* Ex. 1003, 9:13–14, 18:23–26; Ex. 1009 ¶ 36. Specifically, we agree with Dr. Weissman that Gold's description of transferring a file to be backed up from the client to the tape backup

apparatus necessarily requires reading that file, and, thus, Gold discloses "reading the first storage object," as recited in claim 1. Ex. 1009 ¶ 36. For the same reasons, Gold discloses the "reading means," "reading module," and "reading instructions," as recited in claims 9, 13, 17, and 21. *See* Pet. 26, 29–30, 34, 38.

We further find that Gold's description of sending files to be backed up to a data transfer module to be compressed and sent to the tape backup apparatus discloses "copying contents of the first storage object to a second storage object of a second storage area," as recited in claim 1. Ex. 1003, 9:5–14; Pet. 18–19. For the same reasons, Gold discloses the "copying means," "copying module," and "copying instructions" of claims 9, 13, 17, and 21. *See* Pet. 26–27, 30–31, 34–35, 38–39.

Because the files excluded by Gold's FDM, which do not contain relevant data, are filtered rather than read and copied, Gold also discloses "if the first storage object does not contain relevant data, selecting a second storage object of the first storage area without reading the first storage object and without copying the contents of the first storage object to the second storage object," as recited in claims 1 and 7. Ex. 1003, 7:16–19; Pet. 19. For the same reasons, Gold discloses the second "selecting means," second "selecting module," and second "selecting instructions" of claims 9, 13, 17, and 21. *See* Pet. 27, 31, 35, 39.

Petitioner points to Gold's description of FDM 330 consulting with the tape backup apparatus to determine whether a file to be backed up is already on the tape as constituting a disclosure of the following limitations of claim 7:

> "identifying a second storage object of a second storage area corresponding to the first storage object of the first storage area";
>
> "wherein contents of the second storage object were previously copied from contents of the first storage object";
>
> "comparing the contents of the first storage object to the contents of the second storage object"; and
>
> "when the contents of the first storage object and the contents of the second storage object do not match, copying the contents of the first storage object to the second storage object."

Pet. 17, 22–23 (citing Ex. 1003, 7:23–30, 8:30–37, 9:5–23, 10:52–59). Petitioner supports its contentions with the testimony of Dr. Weissman. Ex. 1009 ¶ 36. Patent Owner does not present argument or evidence rebutting Petitioner's evidence as to these limitations. We are persuaded by Petitioner's evidence that Gold discloses the additional limitations of claim 7.

In sum, we have considered the evidence and argument presented by Petitioner and Patent Owner. We are persuaded that Petitioner's evidence shows that Gold discloses each limitation of claims 1, 7, 9, 13, 17, and 21, arranged as in those claims. We are not persuaded by Patent Owner's arguments that Gold fails to disclose claim limitations regarding "relevant data." Accordingly, we determine that Petitioner has proved by a preponderance of the evidence that claims 1, 7, 9, 13, 17, and 21 are anticipated by Gold.

Claims 2–6 depend from claim 1; claim 8 depends from claim 7; claims 10–12 depend from claim 9; claims 14–16 depend from claim 13; claims 18–20 depend from claim 17; and claims 22–24 depend from claim 21. Patent Owner does not present any evidence or argument regarding

these dependent claims specifically. Having reviewed Petitioner's evidence of unpatentability for these dependent claims (Pet. 19–22, 24–25, 27–29, 31–32, 35–37, 39–41), we conclude that Petitioner also has proved by a preponderance of the evidence that claims 2–6, 8, 10–12, 14–16, 18–20, and 22–24 are anticipated by Gold.

2. Anticipation by Pruett

Petitioner contends that Pruett anticipates claims 1–3, 6, 7, 13, 15, and 16. Pet. 41. Petitioner supports its Petition with the testimony of Dr. Weissman. For the reasons below, Petitioner has shown, by a preponderance of the evidence, that Pruett discloses each limitation of claims 1–3, 6, 7, 13, 15, and 16.

a. Overview of Pruett

Pruett describes a technique for automatically synchronizing files and directories between two computers connected by a network. Ex. 1005, Abstract. Figure 1, reproduced below, illustrates an example:

FIG. 1

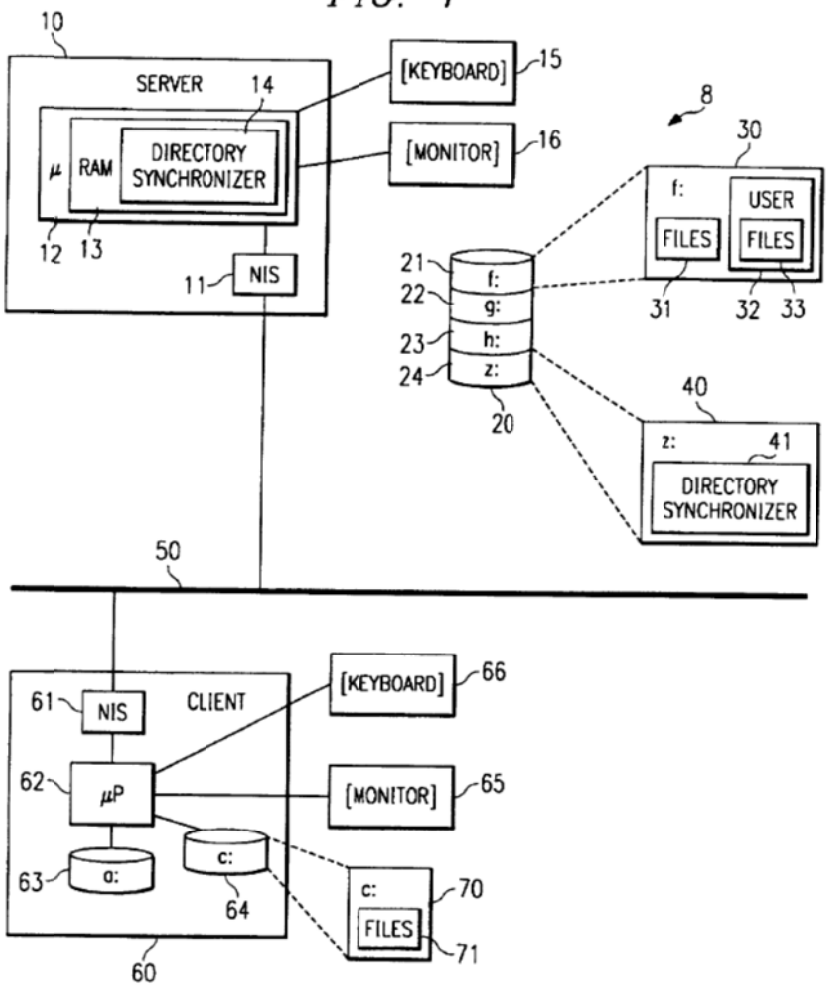

Figure 1 is a block diagram of a local area network. *Id.* at 2:35–36.

File directory 30 stored in mass storage device 20 on file server 10 can be synchronized with a directory stored in hard disk drive 64 on client computer 60. *Id.* at 2:43–67. A user initiates a synchronization by entering, on keyboard 66 of client computer 60, a command invoking synchronizing program 41 and identifying the source and target directories along with any optional "qualifiers" that affect the synchronization. *Id.* at 3:1–39. One of those qualifiers, the command sequence "/H," specifies that hidden files or

directories in the source directory are not synchronized with the target directory. *Id.* at 8:29–40.

Synchronizing program 41 proceeds by identifying whether a file in the source directory is a subdirectory, and, if not, determining whether a file with that name exists in the target directory. *Id.* 6:18–30. If the file does not exist in the target directory, it is copied from the source directory to the target directory. *Id.* at 6:31–33. If the file exists in the target directory, the synchronizing program determines whether the files in the source and target directories have identical contents (e.g., by comparing date stamps of the files), and, if so, does not copy the contents of the source file into the target file. *Id.* at 6:42–64. Synchronizing program 41 continues this process sequentially through the files in the source directory until there are no more files to synchronize. *Id.* at 8:22–28. If the "/H" qualifier is included, and a file is hidden, the synchronizing program skips determining whether the hidden file is a subdirectory, determining whether it is already in the target directory, or copying the file to the target directory. *Id.* at 8:29–40; Fig. 2B (path 200h).

The user also can specify, using a "/I" qualifier, an "action" file that specifies an action to be taken when the synchronizing program reaches various source or target subdirectories. *Id.* at 3:55–65. One of the actions specified in the action file can be to ignore a source directory or subdirectory. *Id.* at 5:32–54. If a directory is to be ignored, the synchronizing program skips determining whether the files in that directory are subdirectories, determining whether they are already in the target directory, or copying the files to the target directory. *Id.* at 5:50–56; Fig. 2A (path 130i).

b. Anticipation of Claims 1–3, 6, 7, 13, 15, and 16

We find that the source directory 32 of the server 10's mass storage device 20 in Pruett is a first storage area and a subdirectory or file in that source directory is a first storage object, as recited in claims 1, 7, and 13. Pruett's description of determining whether a file in the source directory is a file or subdirectory and subsequently selecting the file or subdirectory for processing is a disclosure of "selecting a first storage object of a first storage area," as recited in claims 1 and 7, and "a selecting module" as recited in claim 13. Ex. 1005, 5:21–31, 6:18–26, 8:22–28; Pet. 43–44, 49, 52–53.

The parties dispute whether Pruett discloses "determining whether the first storage object contains relevant data," as recited in claims 1 and 7, or the "determining module configured to determine whether the first storage object contains relevant data," as recited in claim 13. Petitioner contends that Pruett's description of skipping hidden files and only analyzing and copying non-hidden files discloses these limitations. Pet. 44–45, 49–50, 53.

Patent Owner contends that Petitioner offers insufficient evidence to equate Pruett's hidden files with files that are not relevant. PO Resp. 29. According to Patent Owner, "[s]imply because a user does not wish to back up a particular file or directory (hidden or otherwise) does not mean that the application that will make use of that file does not consider the file to be relevant." *Id.* at 30. Patent Owner argues that Pruett's determination as to whether to copy a file is made by the user according to whether the user wishes to back up the file, rather than as to the role the file would play in an application. *Id.*

In reply, Petitioner argues that nothing in the claim language precludes a relevancy determination that is made pursuant to a user

configuration. Reply 12. We agree with Petitioner. As explained in Section II.A.1, "relevant data" are data relevant to the purpose for which storage objects are being copied, and include data used by an application to function and data upon which an application operates. If a user, through configuration flags, indicates to a backup application that certain files are not to be backed up, the user is indicating that data not flagged for exclusion are relevant to the purpose for which the data are being copied. Patent Owner does not point to anything in the claim language or Specification of the '089 patent precluding a user from making this contextual consideration.

At the oral hearing, Patent Owner clarified its argument, stating that:

> [w]e are not taking the position that a user selection is somehow precluded by the claims. Our argument with respect to Pruett really is that the Petitioner has bootstrapped the notion of copying to provide relevance. In other words, because the file is copied, somehow it is relevant. Whereas the claim requires the opposite. Relevance has to be determined prior to copying.

Tr. 47:2–9. We are not persuaded by this argument. As can be seen in Figure 2B, Pruett's method makes a determination as to whether a file is hidden (and thus is not subject to backup) before copying the file. *See also* Ex. 1005, 8:29–40:

> Returning to step 200 in FIG. 2A, if a "/H" qualifier was specified upon execution, indicating a desire to synchronize only those files which are not "hidden," then the method proceeds to decisional step 205, as illustrated by path 200h. At decisional step 205, it is determined whether the "hidden" tag for the currently selected source file has been triggered, indicating that the currently selected file is a hidden file. If the file is not hidden, the method proceeds to decisional step 210. If the file is hidden, the method proceeds to decisional step 270. Thus, if the user has included the "/H" qualifier, hidden files in

> the source directory are not synchronized with the target directory.

Copying (at steps 230 and 242 of Figure 2B) only takes place if the method proceeds to step 210 rather than 205. Thus, relevance is determined before copying. We find that Pruett discloses "determining whether the first storage object contains relevant data," as recited in claims 1 and 7, and the similar limitation in claim 13.

We also find that Pruett discloses, after this determination, "if the first storage object contains relevant data, . . . reading the first storage object," and "when the reading is completed, copying contents of the first storage object to a second storage object of a second storage area" as recited in claim 1. *See* Ex. 1005, 6:27–33; 6:64–67. We agree with Dr. Weissman (Ex. 1009 ¶ 56) that a file must be read before it can be copied. *See also* Ex. 1005, 3:16–19 ("To completely synchronize the target and source directories using program 41, the user should have . . . read permission for the source directory."). For the same reasons, Pruett discloses the "reading module" and "copying module" of claim 13. *See* Pet. 53–54.

We further find that Pruett's description of sequentially processing files and sub-directories, and skipping those that are hidden, constitutes a disclosure of "if the first storage object does not contain relevant data, selecting a second storage object of the first storage area without reading the first storage object and without copying the contents of the first storage object to the second storage object," as recited in claims 1 and 7. Ex. 1005, 5:43–44, 8:37–40, 9:31–37; Pet. 46, 52. For the same reasons, Pruett discloses the second "selecting module" recited in claim 13. *See* Pet. 54–55.

Petitioner points to Pruett's description of identifying files with the same name in both the source and target directories, comparing the time stamps of those files, and copying the contents of the source files to the target files when the time stamps do not match, as constituting a disclosure of the following limitations of claim 7:

> "identifying a second storage object of a second storage area corresponding to the first storage object of the first storage area";
>
> "wherein contents of the second storage object were previously copied from contents of the first storage object";
>
> "comparing the contents of the first storage object to the contents of the second storage object"; and
>
> "when the contents of the first storage object and the contents of the second storage object do not match, copying the contents of the first storage object to the second storage object."

Pet. 50–52 (citing Ex. 1005, 5:45–49, 5:58–60, 6:2730, 6:42–67, 7:38–41). Patent Owner does not present argument or evidence rebutting Petitioner's evidence as to these limitations. We are persuaded by Petitioner's evidence that Pruett discloses the additional limitations of claim 7.

In sum, we have considered the evidence and argument presented by Petitioner and Patent Owner. We are persuaded that Petitioner's evidence shows that Pruett discloses each limitation of claims 1, 7, and 13, arranged as in those claims. We are not persuaded by Patent Owner's arguments that Pruett fails to disclose claim limitations regarding "relevant data." Accordingly, we determine that Petitioner has proved by a preponderance of the evidence that claims 1, 7, and 13 are anticipated by Pruett.

Claims 2, 3, and 6 depend from claim 1; claims 15 and 16 depend from claim 13. Patent Owner does not present any evidence or argument

regarding these dependent claims specifically. Having reviewed Petitioner's evidence of unpatentability for these dependent claims (Pet. 46–49, 55–56), we conclude that Petitioner also has proved by a preponderance of the evidence that claims 2, 3, 6, 15, and 16 are anticipated by Pruett.

III. CONCLUSION

Petitioner has demonstrated by a preponderance of the evidence that claims 1–24 are unpatentable based on the following grounds of unpatentability:

(1) Claims 1–24 as anticipated under 35 U.S.C. § 102(b) by Gold; and

(2) Claims 1–3, 6, 7, 13, 15, and 16 as anticipated under 35 U.S.C. § 102(b) by Pruett.

IV. ORDER

For the reasons given, it is

ORDERED, based on a preponderance of the evidence, that claims 1–24 of U.S. Patent No. 7,032,089 B1 are held unpatentable; and

FURTHER ORDERED, because this is a final written decision, the parties to this proceeding seeking judicial review of our Decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

For Petitioner:

Erika Harmon Arner
Jeffrey C. Totten
FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER, LLP
erika.arner@finnegan.com
jeffrey.totten@finnegan.com

For Patent Owner:

Tarek N. Fahmi
Megan M. Gallant
ASCENDA LAW GROUP, PC
tarek.fahmi@ascendalaw.com
megan.gallant@ascendalaw.com