

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

EMC CORPORATION AND VMWARE, INC.
Petitioner

v.

PERSONALWEB TECHNOLOGIES LLC
Patent Owner

Case IPR2013-00082 (JYC)
U.S. Patent No. 5,978,791

Before KEVIN F. TURNER, JONI Y. CHANG, and
MICHAEL R. ZECHER, *Administrative Patent Judges*.

ZECHER, *Administrative Patent Judge*

DECISION
Institution of *Inter Partes* Review
37 C.F.R. § 42.108

I. INTRODUCTION

EMC Corporation and VMware, Inc. (“EMC”) filed a petition (“Pet.”) requesting *inter partes* review of claims 1-4, 29-33, and 41 of U.S. Patent 5,978,791 (“the ’791 patent”). Paper No. 8. Patent owner, PersonalWeb Technologies LLC (“PersonalWeb”), filed a preliminary response (“Prelim. Resp.”). Paper No. 15. We have jurisdiction under 35 U.S.C. § 314.

The standard for instituting an *inter partes* review is set forth in 35 U.S.C. § 314(a), which provides:

THRESHOLD --The Director may not authorize an *inter partes* review to be instituted unless the Director determines that the information presented in the petition filed under section 311 and any response filed under section 313 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.

Taking into account PersonalWeb’s preliminary response, we conclude that the information presented in the petition demonstrates that there is a reasonable likelihood that EMC will prevail in challenging claims 1-4, 29-33, and 41 as unpatentable under 35 U.S.C. §§ 102 and 103. Pursuant to 35 U.S.C. § 314, we hereby authorize an *inter partes* review to be instituted as to claims 1-4, 29-33, and 41 of the ’791 patent.

A. *Related Matters*

EMC indicates that the ’791 patent was asserted against it in *PersonalWeb Technologies LLC v. EMC Corporation and VMware, Inc.*, Case No. 6:11-cv-00660-LED, pending in the U.S. District Court for the Eastern District of Texas. Pet. 1. EMC also filed five other petitions

seeking *inter partes* review of the following patents: U.S. Patent No. 6,415,280 (IPR2013-00083), U.S. Patent No. 7,945,544 (IPR2013-00084), U.S. Patent No. 7,945,539 (IPR2013-00085), U.S. Patent No. 7,949,662 (IPR2013-00086), and U.S. Patent No. 8,001,096 (IPR2013-00087). *Id.*

B. The Invention of the '791 Patent (Ex. 1001)

The invention of the '791 patent relates to a data processing system that identifies data items using substantially unique identifiers, otherwise referred to as True Names, which depend on all the data in the data item and only on the data in the data item. Ex. 1001, Spec. 1:14-18, 3:29-32, and 6:6-10. According to the '791 patent, the identity of a data item depends only on the data and is independent of the data item's name, origin, location, address, or other information not directly derivable from the data associated therewith. Ex. 1001, Spec. 3:33-35. The invention of the '791 patent also examines the identities of a plurality of data items in order to determine whether a particular data item is present in the data processing system. Ex. 1001, Spec. 3:36-39.

C. Illustrative Claims

Independent claims 1, 30, and 33 are illustrative:

1. In a data processing system, an apparatus comprising:

identity means for *determining, for any of a plurality of data items present in the system, a substantially unique identifier, the identifier being determined using and depending on all the data in the data item and only the data in the data*

item, whereby two identical data items in the system will have the same identifier; and

existence means for determining whether a particular data item is present in the system, by examining the identifiers of the plurality of data items.

Ex. 1001, claims—Spec. 39:14-23 (emphasis added).

30. A method of identifying a data item present in a data processing system for subsequent access to the data item, the method comprising:

determining a substantial unique identifier for the data item, the identifier depending on and being determined using all of the data in the data item and only the data in the data item, whereby two identical data items in the system will have the same identifier; and

accessing a data item in the system using the identifier of the data item.

Ex. 1001, claims—Spec. 42:58-67 (emphasis added).

33. A method of duplicating a given data item present at a source location to a destination location in a data processing system, the method comprising:

determining a substantially unique identifier for the given data item, the identifier depending on and being determined using all of the data in the data item and only the data in the data item, whereby two identical data items in the system will have the same identifier;

determining, using the data identifier, whether the data item is present at the destination location; and

based on the determining whether the data item is present, providing the destination location with the data item only if the data item is not present at the destination.

Ex. 1001, claims—Spec. 43:11-23 (emphasis added).

D. Prior Art Relied Upon

EMC relies upon the following prior art references:

Woodhill US 5,649,196 July 15, 1997 Ex. 1005

Shirley Browne et al., “*Location-Independent Naming of Virtual Distributed Software Repositories*,” University of Tennessee Technical Report CS-95-278 (Feb. 1995)(Ex. 1002)(hereinafter “Browne”).

Albert Langer, “*Re: dl/describe (File Descriptions)*,” post to the “alt.sources” newsgroup on Aug. 7, 1991 (Ex. 1003)(hereinafter “Langer”).

Frederick W. Kantor, “*FWKCS™ Contents_Signature System Version 1.22*,” Zipfile FWKCS122.ZIP (Aug. 10, 1993)(Ex. 1004)(hereinafter “Kantor”).

E. Alleged Grounds of Unpatentability

EMC seeks to have claims 1-4, 29-33, and 41 of the ’791 patent cancelled based on the following alleged grounds of unpatentability:

1. Claim 1-4, 29-33, and 41 as anticipated under 35 U.S.C. § 102(a) by Browne. Pet. 26-35.
2. Claims 1-4, 29-33, and 41 as unpatentable under 35 U.S.C. § 103(a) over Browne. *Id.* at 35.
3. Claims 1-4, 29-33, and 41 as unpatentable under 35 U.S.C. § 103(a) over the combination of Browne and Langer. *Id.* at 35-36.
4. Claims 1-4 and 29 as unpatentable under 35 U.S.C. § 103(a) over the combination of Browne and Woodhill. *Id.* at 36-37.
5. Claims 1-4, 29-33, and 41 as anticipated under U.S.C. § 102(b) by Langer. *Id.* at 37-43.

6. Claims 1-4 and 29 as unpatentable under 35 U.S.C. § 103(a) over the combination of Langer and Woodhill. *Id.* at 43.

7. Claims 1-3, 29, and 33 as anticipated under U.S.C. § 102(b) by Kantor. *Id.* at 43-49.

8. Claims 4, 30-32, and 41 as unpatentable under 35 U.S.C. § 103(a) over Kantor. *Id.* at 49-50.

9. Claims 1-4 and 29 as unpatentable under 35 U.S.C. § 103(a) over the combination of Kantor and Langer. *Id.* at 50-51.

10. Claims 1-4, 29-33, and 41 as anticipated under U.S.C. § 102(e) by Woodhill. *Id.* at 51-59.

11. Claims 1-4 and 29 as unpatentable under 35 U.S.C. § 103(a) over Woodhill. *Id.* at 59.

12. Claims 1-4 and 29 as unpatentable under 35 U.S.C. § 103(a) over the combination of Woodhill and Kantor. *Id.* at 59-60.

II. FINDINGS OF FACT

The following findings of facts are supported by a preponderance of the evidence.

Woodhill

Woodhill generally relates to a system and method for distributed storage management on a networked computer system that includes a remote backup file server in communication with one or more local area networks. Ex. 1005, Spec. 1:11-17. Figure 1 of Woodhill illustrates the networked computer system. Ex. 1005, Spec. 2:56-58. Figure 1 of Woodhill is reproduced below.




FIG. 1

Figure 1 of Woodhill illustrates the networked computer system 10.

A remote backup file server 12 communicates with a wide area network 14 via data path 13, the wide area network 14 communications with a plurality of local area networks 16 via data paths 15, and each local area network 16 communications with multiple user workstations 18 and local computers 20 via data paths 17. Ex. 1005, Spec. 3:12-30. The storage space on each disk drive 19 on each local computer 20 is allocated according the hierarchy illustrated in Figure 2. Ex. 1005, Spec. 3:31-44.

Figure 2 of Woodhill illustrates a Distributed Storage Manager program that allocates storage space on each of the storage devices in the networked computer system. Ex. 1005, Spec. 2:59-62. Figure 2 of Woodhill is reproduced below.




FIG. 2

Figure 2 of Woodhill illustrates the Distributed Storage Manager program 24.

The Distributed Storage Manager program 24 builds and maintains the File Database 25 on the one or more disk drives 19 on each local computer 20 in the networked computer system 10. Ex. 1005, Spec. 3:45-49. The Distributed Storage Manager program 24 views a file as a collection of data streams. Ex. 1005, Spec. 4:13-15. Woodhill defines a data stream as a distinct collection of data within a file that may change independently from other distinct collections of data within the file. Ex. 1005, Spec. 4:15-18. Depending on the size of the data stream, the Distributed Storage Manager program 24 divides each data stream into one or more binary objects. Ex. 1005, Spec. 4:21-30.

Figure 3 of Woodhill illustrates the File Database used by the Distributed Storage Manager program. Ex. 1005, Spec. 2:63-64. Figure 3 of Woodhill is reproduced below.

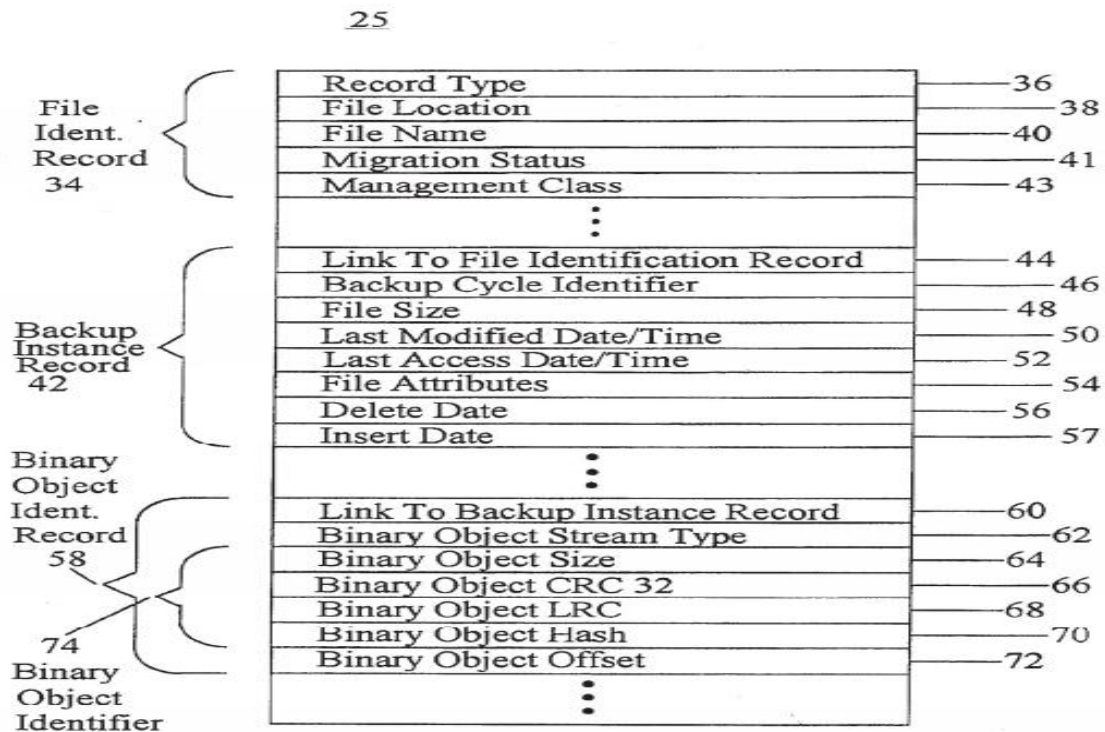


FIG. 3

Figure 3 of Woodhill illustrates the File Database 25.

The File Database 25 includes three levels of records organized according to a predefined hierarchy: (1) the File Identification Record 34; (2) the Backup Instance Record 42; and (3) the Binary Object Identification Record 58. Ex. 1005, Spec. 3:54-4:47. The Binary Object Identification Record 58 includes, amongst other things, a Binary Object Identifier 74 that comprises a Binary Object Size 64, Binary Object CRC32 66, Binary Object LRC 68 and Binary Object Hash 70. Ex. 1005, Spec. 4:45-47, 7:64-8:1. The Binary Object Identifier 74 is a unique identifier for each binary object that is backed up. Ex. 1005, Spec. 4:45-47.

Because the Binary Object Identifier 74 uniquely identifies a particular binary object, Woodhill recognizes the importance of minimizing the possibility of assigning two different binary objects the same Binary Object Identifier 74. Ex. 1005, Spec. 8:33-36. While Woodhill discloses calculating the Binary Object Identifier 74 in various ways, *e.g.*, using a Binary Object size calculation, a Cyclical Redundancy Check calculation, a Longitudinal Redundancy Check calculation, and a Binary Hash algorithm (Ex.1005, Spec. 8:1-31), the key notion is that the Binary Object Identifier 74 is calculated from the content of the data instead of from an external or arbitrary source. Ex. 1005, Spec. 8:38-42. In other words, Woodhill recognizes that the critical feature in creating a Binary Object Identifier 74 is that the identifier should be based on the contents of the binary object such that the Binary Object Identifier 74 can change when the contents of the binary object change. Ex. 1005, Spec. 8:58-62. Therefore, duplicate binary objects, even if resident on different types of computers in a network, may be recognized by their identical Binary Object Identifiers 74. Ex. 1005, Spec. 8:62-65.

Woodhill discloses that the Distributed Storage Manager program 24 performs two backup operations concurrently. Ex. 1005, Spec. 9:30-31. First, the Distributed Storage Manager program 24 stores a compressed copy of each binary object it needs to restore the disk drives 19 on each local computer 20 somewhere on the local area network 16 other than on the local computer 20 where the binary object originally resided. Ex. 1005, Spec. 9:31-36. Second, the Distributed Storage Manager program 24 transmits

new or changed binary objects to the remote backup file server 12. Ex. 1005, Spec. 9:36-38.

Woodhill discloses that the Distributed Storage Manager program 24 performs auditing and reporting functions on a periodic basis in order to ensure that the binary objects, which already have been backed up, may be restored. Ex. 1005, Spec. 18:11-13. According to Woodhill, the Distributed Storage Manager program 24 initiates a restore of a randomly selected binary object identified by a Binary Object Identification Record 58 stored in the File Database 25. Ex. 1005, Spec. 18:16-19.

III. CLAIM CONSTRUCTION

During an *inter partes* review, the Board construes claims by applying the broadest reasonable interpretation in light of the specification. 37 C.F.R. § 42.100(b); *see also* Office Patent Trial Practice Guide, 77 *Fed. Reg.* 48756, 48766 (Aug. 14, 2012). Absent a special definition for a claim term being set forth in the specification, the definition that governs is the ordinary and customary meaning of the claim term as would be understood by one with ordinary skill in the art. *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312 (Fed. Cir. 2005) (en banc). In some cases, the ordinary and customary meaning of a claim term as would be understood by one with ordinary skill in the art may be apparent even to lay judges, and claim construction in such cases involves little more than the application of the widely accepted meaning of commonly understood words. *Id.* at 1314.

A. Claim Terms

EMC identifies five claim terms and its claim construction for each claim term. Pet. 4-6. Those claim terms are listed as follows: (1) “substantially unique identifier;” (2) “using the identifier;” (3) “data” and “data item;” (4) “location;” and (5) “True Name, data identity, and data identifier.” We will address each claim term identified by EMC in turn.

1. “Substantially unique identifier”

EMC construes the claim term “substantially unique identifier” as “an identity for a data item generated by processing *all* of the data in the data item, and *only* the data in the data item, through an algorithm.” Pet. 4 (emphasis in original). PersonalWeb agrees with EMC’s claim construction with respect to the claim term “substantially unique identifier.” Prelim. Resp. 4-5 (citing to Ex. 1001, Spec. 1:16-18 and 3:6-11).

While the Specification of the ’791 patent does not set forth an explicit or special definition for the claim term “substantially unique identifier,” we note that every challenged independent claim, *i.e.*, claims 1, 30, and 33, includes the following claim language:

a substantially unique identifier, the identifier being determined using and depending on all the data in the data item and only the data in the data item

In light of the Specification of the ’791 patent and the claim language of independent claims 1, 30, and 33, we construe the claim term “substantially unique identifier” as “an identity for a data item generated

being determined using and depending on all of the data in the data item, and only the data in the data item.”

2. *“Using the identifier”*

EMC construes the claim term “using the identifier” as “employing the unique identifier of the data item, with or without other information, to carry out the recited function.” Pet. 4. PersonalWeb agrees with EMC’s claim construction with respect to the claim term “using the identifier.”

Prelim. Resp. 6. Because the agreed upon claim construction is consistent with the Specification of the ’791 patent, we will adopt it as our own.

3. *“Data” and “data item”*

EMC construes that the claim terms “data” and “data item” as:

[a] sequence of bits. Thus a data item may be the contents of a file, a portion of a file, a page in memory, an object in an object-oriented program, a digital message, a digital scanned image, a part of a video or audio signal, or any other entity which can be represented by a sequence of bits.

Pet. 5 (citing to Ex. 1001, Spec. 1:54-60). EMC also indicates that the claim terms “data” and “data item” include the following:

data items (the data items being files, directories, records in the database, objects in objected-oriented programming, locations in memory or on a physical device or the like).

Pet. 5 (citing to Ex. 1001, Spec. 1:65-2:2). PersonalWeb only agrees with EMC that the claim terms “data” and “data item” may be construed as a “sequence of bits.” Prelim. Resp. 4 (citing to Ex. 1001, Spec. 1:54-55).

Based on our review of the Specification of the '791 patent, we broadly, but reasonably construe the claim term “data item” as a “sequence of bits,” which includes one of the following: (1) the contents of a file; (2) a portion of a file; (3) a page in memory; (4) an object in an object-oriented program; (5) a digital message; (6) a digital scanned image; (7) a part of a video or audio signal; (8) a directory; (9) a record in a database; (10) a location in memory or on a physical device; and (11) any other entity which can be represented by a sequence of bits.

Further, we note that every challenged independent claim, *i.e.*, claims 1, 30, and 33, includes the claim language “all of the *data* in the *data item* and only the *data* in the *data item*.” Emphasis added. As such, independent claims 1, 30, and 33 treat the claims terms “data” and “data item” as separate and distinct elements. “In the absence of evidence to the contrary, we must presume that the use of different terms in the claims connotes different meanings.” *CAE Screenplates Inc. v. Heinrich Fiedler GmbH & Co. KG*, 224 F.3d 1308, 1317 (Fed. Cir. 2000). Therefore, we construe the claim term “data” as a subset of a “data item.”

4. “Location”

EMC construes the claim term “location” with respect to a data processing system as “any of a particular processor in the system, a memory of a particular processor, a storage device, a removable storage medium (such as a floppy disk or compact disk), or any other physical location in the system.” Pet. 5 (citing to Ex. 1001, Spec. 5:65-6:4). PersonalWeb agrees with EMC’s claim construction with respect to the claim term “location.”

Prelim. Resp. 4 (citing to Ex. 1001, Spec. 5:65-6:3). Because the agreed upon claim constructions is consistent with the Specification of the '791 patent, we will adopt it as our own.

5. *“True Name, data identity, and data identifier”*

EMC construes the claim terms “True Name, data identity, and data identifier” as a “substantially unique data identifier for a particular element.” Pet. 5-6 (citing to Ex. 1001, Spec. 6:7-10; *see also* Spec. 14:1-39).

PersonalWeb challenges EMC’s claim construction with respect to the claim term “True Name.” Prelim. Resp. 5. While PersonalWeb agrees with EMC that the claim term “True Name” amounts to a “substantially unique identifier,” PersonalWeb contends that the claim term “True Name” is narrower than a “data identifier” because it is further defined in the Specification of the '791 patent—specifically the “True Name” is calculated in accordance with the description at column 12, line 54 through column 13, line 9. *Id.* Upon reviewing the portion of the Specification of the '765 patent cited by PersonalWeb, we do not find an explicit or special definition for the claim term “True Name.” However, we note that the portion of the Specification of the '765 patent cited by EMC does provide an explicit or special definition for the claim term “True Name.” Pet. 5-6 (citing to Spec. 6:7-10). Therefore, we agree with EMC that the claim term “True Name” should be construed as a “substantially unique data identifier for a particular item.”

B. Means-Plus-Function Limitations

When construing a means-plus-function limitation under 35 U.S.C. § 112, ¶ 6,¹ we first must identify the claimed function, and then we look to the specification to identify the corresponding structure that actually performs the claimed function. *Med. Instrumentation & Diagnostics Corp. v. Elekta AB*, 344 F.3d 1205, 1210 (Fed. Cir. 2003); *Cardiac Pacemakers, Inc. v. St. Jude Med., Inc.*, 296 F.3d 1106, 1119 (Fed. Cir. 2002). The corresponding structure of a means-plus-function limitation, however, must be more than simply a general-purpose computer or microprocessor to avoid pure functional claiming. *Aristocrat Techs. Austl. Pty Ltd. v. Int'l Game Tech.*, 521 F.3d 1328, 1333 (Fed. Cir. 2008). That is, the specification must disclose “enough of an algorithm to provide the necessary structure under § 112, ¶ 6” or a disclosure that can be expressed in any understandable terms, *e.g.*, a mathematical formula, in prose, or as a flowchart. *Finisar Corp. v. The DirectTV Group*, 523 F.3d 1323, 1340 (Fed. Cir. 2008).

EMC identifies several claim limitations as means-plus-function limitations invoking 35 U.S.C. § 112, ¶ 6, and their corresponding structure for performing the claimed function. Pet. 6-8. At the outset, we agree that each limitation identified by EMC is a means-plus-function limitation because: (1) each limitation uses the term “means for”; (2) the term “means

¹ Section 4(c) of the AIA re-designated 35 U.S.C. § 112, ¶ 6, as 35 U.S.C. § 112(f). Because the '791 patent has a filing date before September 16, 2012 (effective date), we will refer to the pre-AIA version of 35 U.S.C. § 112.

for” is modified by functional language; and (3) the term “means for” is not modified by sufficient structure recited in the claim to perform the claimed function. We will address the claimed function and corresponding structure for each means-plus-function limitation identified by EMC in turn.

- 1. Identifying means for determining, for any of a plurality of data items present in the system, a substantially unique identifier, the identifier being determined using and depending on all of the data in the data item and only the data in the data item, whereby two identical data items in the system will have the same identifier (Claim 1)*

Both parties agree that the claimed function for this means-plus-function limitation is:

determining, for any of a plurality of data items present in the system, a substantially unique identifier, the identifier being determined using and depending on all of the data in the data item and only the data in the data item, whereby two identical data items in the system will have the same identifier.

Pet. 6; Prelim. Resp. 7. EMC contends that the corresponding structure for this means-plus-function limitation is the processor illustrated in Figure 1(b) programmed to execute the “Calculate True Name” mechanism depicted in Figures 10(a) and 10(b), where the message digest (“MD”) function is one of the MD4, MD5, and secure hash algorithm (“SHA”) functions. Pet. 6 (citing to Ex. 1001, Spec. 4:64-6:19, 12:54-14:39, 14:51-53, 31:32-50, and 32:54-64; Ex.1019). In response, PersonalWeb contends that the corresponding structure identified by EMC is incorrect because the portions of the specification cited by EMC do not indicate the structure necessary for

performing the claimed function. Prelim. Resp. 10. PersonalWeb argues that the specification clearly identifies the corresponding structure as at least one processor programmed in accordance with the Calculate True Name mechanism. *Id.* (citing to Ex. 1001, Spec. 7:62-63, 12:54-13:19, and 14:1-39).

Figures 1(a) and 1(b) illustrate the data processing system that implements the invention of the '791 patent. Ex. 1001, Spec. 4:44-46. Figure 1(b) is reproduced below.

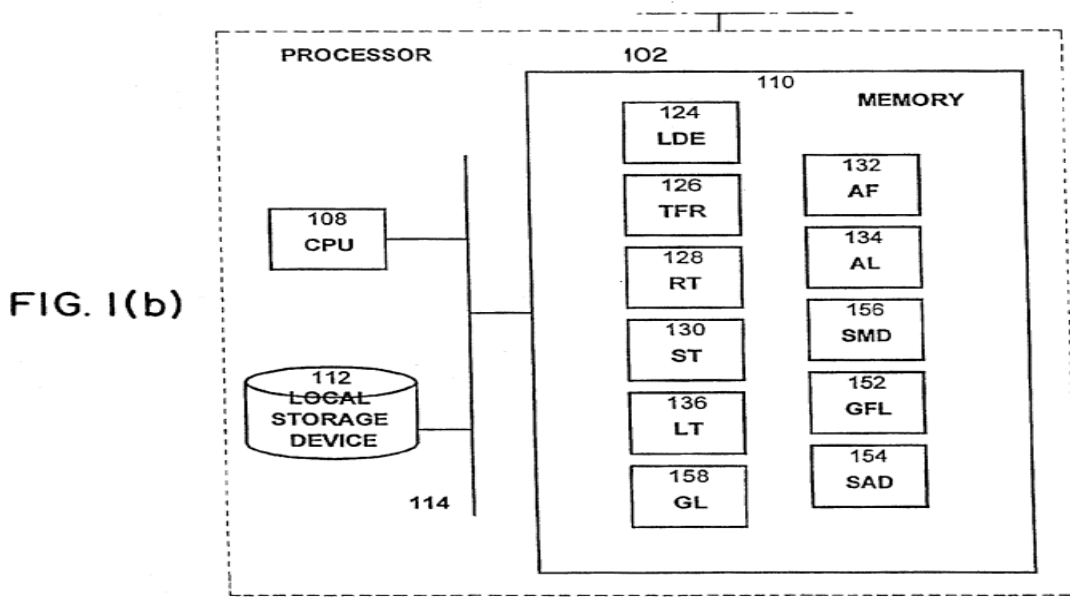


Figure 1(b) illustrates a typical data processor.

The Specification of the '791 patent discloses that each processor 102 includes a central processing unit 108, memory 110, and one or more local storage devices 112 connected via an internal bus 114. Ex. 1001, Spec. 4:64-67. The memory 110 in each processor 102 stores data structures that

are either local to the processor itself or shared amongst multiple processors in the data processing system. Ex. 1001, Spec. 7:61-8:18. According to the Specification of the '791 patent, “the [aforementioned] *data structures, stored in memory 110 of one or more processors 102 are used to implement the mechanisms* described herein.” Ex. 1001, Spec. 7:61-63 (emphasis added). Further, the Specification of the '791 patent discloses:

In the presently preferred embodiments, *either MD5 or SHA is employed as the basis for the computation of True Names.* Whichever of these two message digest functions is employed, the same function must be employed on a system-wide basis.

Ex. 1001, Spec. 13:15-19 (emphasis added).

We identify the corresponding structure for performing the recited function—namely

determining, for any of a plurality of data items present in the system, a substantially unique identifier, the identifier being determined using and depending on all of the data in the data item and only the data in the data item, whereby two identical data items in the system will have the same identifier

—to be a data processor programmed to perform a hash function, *e.g.*, MD5 or SHA.

2. *Existence means for determining whether a particular item is present in the system, by examining the identifiers of the plurality of data items (Claim 1)*

Both parties agree that the claimed function for this means-plus-function limitation is “determining whether a particular data item is present in the system, by examining the identifiers of the plurality of data items.”

Pet. 7; Prelim. Resp. 7. EMC contends that the corresponding structure for this means-plus-function limitation is the processor illustrated in Figure 1(b), which stores the True File Registry and is programmed to execute the “Locate Remote File” mechanism illustrated in Figures 16(a) and 16(b). Pet. 7 (citing to Ex. 1001, Spec. 4:64-6:19, 9:36-10:10, 16:38-17:9, 17:41-43, 25:10-13, and 35:51-55). In response, PersonalWeb contends that the corresponding structure identified by EMC is incorrect because the portions of the specification cited by EMC do not indicate the structure that is necessary for performing the claimed function. Prelim. Resp. 11. PersonalWeb also argues that certain citations provided by EMC are misplaced because they do not describe the claimed function. *Id.*

As discussed above, the Specification of the '791 patent uses the data structures stored in the memory of a data processor to implement the claimed functions. Ex. 1001, Figure 1(b), Spec. 7:61-63. The Specification of the '791 patent further discloses:

A mechanism for assimilating a data item (scratch file or segment) into a file system, given the scratch file [identification] ID of the data item, is [] described with reference to FIG. 11. The purpose of this mechanism is to add a given data item to the True File registry **126**. If the data item already exists in the True File registry **126**, this will be discovered and used during this process, and the duplicate will be eliminated. . . . *Next, look for an entry for the True Name in the True File registry **126** (Step S232) and determine whether a True Name entry, record **140**, exists in the True File registry.*

Ex. 1001, Spec. 14:41-56 (emphasis added). The Specification of the '791 patent also discloses:

The mechanism to link a path to a True Name is [] described with reference to FIG. 14. *First, if desired, confirm that the True Name exists locally by searching for it in the True Name registry or local directory extensions table 135 (Step S260).*

Ex. 1001, Spec. 15:52-56 (emphasis added).

We identify the corresponding structure for performing the recited function—namely “determining whether a particular data item is present in the system, by examining the identifiers of the plurality of data items”—to be a data processor programmed according to step S232 illustrated in Figure 11 or step S260 illustrated in Figure 14.

3. Local existence means for determining whether an instance of a particular data item is present at a particular location in the system, based on the identifier of the data item (Claims 2 and 3)

Both EMC and PersonalWeb agree that the claimed function for this means-plus-function limitation is “determining whether an instance of a particular data item is present at a particular location in the system, based on the identifier of the data item.” Pet. 7-8; Prelim. Resp. 8. EMC contends that the corresponding structure for this means-plus-function limitation is the processor illustrated in Figure 1(b), which stores the True File Registry and is programmed to execute the “Locate True File” mechanism illustrated in Figure 28. Pet. 7-8 (citing to Ex. 1001, Spec. 4:64-6:19, 9:36-10:10, 16:48-51, 23:52-24:28, 32:42-45, 35:51-55, and 36:65-66; Ex. 1019). In response, PersonalWeb contends that the corresponding structure identified by EMC is incorrect because the portions of the specification cited by EMC do not

indicate the structure necessary for performing the claimed function. Prelim. Resp. 11.

As discussed above, the Specification of the '791 patent uses the data structures stored in the memory of a data processor to implement the claimed functions. Ex. 1001, Figure 1(b), Spec. 7:61-63. In addition, Figure 14 of the '791 patent illustrates Step S260, which confirms that the True Name exists locally by searching for it in the True Name registry or local directory extensions table. Ex. 1001, Spec. 15:54-56.

We identify the corresponding structure for performing the recited function—namely “determining whether an instance of a particular data item is present at a particular location in the system, based on the identifier of the data item”—to be a data processor programmed according to step S260 illustrated in Figure 14.

4. *Data associating means for making and maintaining, for a data item in the system, an association between the data item and the identifier of the data item (Claim 4)*

Both EMC and PersonalWeb agree that the claimed function for this means-plus-function limitation is “making and maintaining, for a data time in the system, an association between the data item and the identifier of the data item.” Pet. 8; Prelim. Resp. 8. EMC contends that the corresponding structure for this means-plus-function limitation is the processor illustrated in Figure 1(b), which stores the True File Registry and is programmed to execute the “Assimilate Data Item” mechanism illustrated in Figure 11. Pet. 8 (citing to Ex. 1001, Spec. 4:64-6:19, 9:36-10:10, 14:40-15:4, 15:41-44,

16:29-31, 18:34-36 and 43-45, 19:30-37, 24:34-35 and 51-52, 28:30-33, 30:55-57, 32:54-33:9, and 33:33-39; Ex. 1019). In response, PersonalWeb contends that the corresponding structure identified by EMC is incorrect because the portions of the specification cited by EMC do not indicate the structure necessary for performing the claimed function. Prelim. Resp. 11-12.

As discussed above, the Specification of the '791 patent uses the data structures stored in the memory of a data processor to implement the claimed functions. Ex. 1001, Figure 1(b), Spec. 7:61-63. In addition, Figure 11 of the '791 patent illustrates a mechanism for assimilating a data item into the True File registry using the scratch file ID of the data item. Ex. 1001, Spec. 14:41-43. According to the Specification of the '791 patent, the purpose of that mechanism is to add a given data item to the True File registry, or if the data item already exists in the True File registry, discover and use the pre-existing data item before eliminating the duplicate. Ex. 1001, Spec. 14:43-47. The relevant steps set forth in Figure 11 of the '791 patent are listed as follows:

First, determine the True Name of the data item corresponding to the given scratch File ID using the Calculate True Name primitive mechanism (Step **S230**). Next, look for an entry for the True Name in the True File registry **126** (Step **S232**) and determine whether a True Name entry, record **140**, exists in the True File registry **126**. If the entry record includes a corresponding True File ID or compressed File ID (Step **S237**), delete the file with the scratch File ID (Step **S238**). Otherwise store the give True File ID in the entry record (step **S239**).

Ex. 1001, Spec. 14:51-60.

We identify the corresponding structure for performing the recited function—namely “making and maintaining, for a data time in the system, an association between the data item and the identifier of the data item”—to be a data processor programmed according to the steps S230, S232, and S237-S239 illustrated in Figure 11.

5. Access means for accessing a particular data item using the identifier of the data item (Claim 4)

Both EMC and PersonalWeb agree that the claimed function for this means-plus-function limitation is “accessing a particular data item using the identifier of the data item.” Pet. 8; Prelim. Resp. 8-9. EMC contends that the corresponding structure for this means-plus-function limitation is the processor illustrated in Figure 1(b), which stores the True File Registry and is programmed to execute the “Make True File Local” mechanism illustrated in Figures 17(a) and 17(b). Pet. 8 (citing to Ex. 1001, Spec. 4:64-6:19, 9:36-10:10, 17:10-45, 18:4-8 and 53-55, 20:65-21:5, 24:54-55, 33:51-59, and 36:52-55 and 61-64). In response, PersonalWeb contends that the corresponding structure identified by EMC is incorrect because the portions of the specification cited by EMC do not indicate the structure necessary for performing the claimed function. Prelim. Resp. 11-12.

As discussed above, the Specification of the '791 patent uses the data structures stored in the memory of a data processor to implement the claimed functions. Ex. 1001, Figure 1(b), Spec. 7:61-63. In addition, the Specification of the '791 patent discloses:

Make True File Local[—] *This mechanism is used when a True Name is known and a locally accessible copy of the corresponding file or data item is required. . . .* The mechanism takes a True Name and returns when there is a local, accessible copy of the True File in the True File registry **126**. This mechanism is described here with reference to the flow chart of FIGS. **17(a)** and **17(b)**.

Ex. 1001, 17:10-18. The relevant steps set forth in Figure 17(a) are listed as follows:

First, look to the True File registry **126** for a True File entry record **140** for a corresponding True Name (Step **S292**). . . . If there is already a True File ID for the entry (Step **S294**), this mechanism's task is complete.

Ex. 1001, Spec. 17:10-23.

We identify the corresponding structure for performing the recited function—namely “accessing a particular data item using the identifier of the data item”—to be a processor programmed according to steps **S292** and **S294** illustrated in Figure 17(a).

IV. ANALYSIS

A. 35 U.S.C. § 102(e) Ground of Unpatentability—Woodhill

Claims 1-4, 29-33, and 41

EMC contends that claims 1-4, 29-33, and 41 are anticipated under 35 U.S.C. § 102(e) by Woodhill. Pet. 51-59. In particular, EMC argues that Woodhill describes the claimed subject matter recited in independent claims 1, 30, and 33. *Id.* at 56-59. EMC also relies upon the Declaration of Douglas W. Clark (Ex. 1009) to support its positions and an attached claim

chart (Ex. 1041) to explain where Woodhill describes the claimed subject matter recited in claims 1-4, 29-33, and 41. *Id.* In response, PersonalWeb contends that Woodhill does not describe the claimed subject matter recited in independent claims 1, 30, and 33, and in dependent claim 41. Prelim. Resp. 35-47.

We begin our analysis by addressing PersonalWeb's arguments with respect to each disputed claim in turn. PersonalWeb contends that the "identity means" recited in independent claim 1 requires a mechanism for determining if a data item is a simple or compound data item. Prelim. Resp. 37. PersonalWeb argues that when the "identity means" determines that the data item is a compound data item, a hash of hashes is required. *Id.* With that context in mind, PersonalWeb alleges that, while Woodhill discloses applying a hash function to a binary object for use in the hash field 70 of the Binary Object Identifier 74, Woodhill does not describe a hash of hashes when it determines that the data item is a compound data item. *Id.* at 39.

PersonalWeb's argument is not commensurate in scope with the "identity means" recited in independent claim 1. Based on our claim construction above, the claimed function for the "identity means" is:

determining, for any of a plurality of data items present in the system, a substantially unique identifier, the identifier being determined using and depending on all of the data in the data item and only the data in the data item, whereby two identical data items in the system will have the same identifier.

The corresponding structure for the aforementioned function is a data processor programmed to perform a hash function, *e.g.*, MD5 or SHA.

Neither the Specification of the '791 patent, nor the claim itself, indicates that the “identity means” requires a hash of hashes, as asserted by PersonalWeb. Consequently, we are persuaded that EMC has demonstrated that Woodhill describes the “identity means” recited in independent claim 1.

Next, PersonalWeb contends that Woodhill does not describe the “existence means” recited in independent 1. Prelim. Resp. 39-40. In particular, PersonalWeb argues that Woodhill does not disclose determining whether a particular binary object is present in the system by examining a plurality of Binary Object Identifiers 74. *Id.* at 40. PersonalWeb alleges that if Woodhill creates a Binary Object Identifier 74, it necessarily means that the corresponding binary object already exists in the system and, therefore, there is no need or reason to determine whether the binary object is already present in the system. *Id.*

Woodhill recognizes duplicate binary objects residing on different types of computers in a network by their identical Binary Object Identifiers 74. Ex. 1005, Spec. 8:62-65. Therefore, consistent with our claim construction above, Woodhill determines whether a particular binary object is present in the networked computer system 10 by examining the Binary Object Identifiers 74 of the plurality of binary objects located within the system. As such, we are persuaded that EMC has demonstrated that Woodhill describes the “existence means” recited in independent claim 1.

With respect to independent claim 30, PersonalWeb contends that Woodhill does not describe the claimed “accessing” method step. Prelim. Resp. 41. In particular, PersonalWeb argues that Woodhill does not disclose

using the Binary Object Identifier 74 to access a binary object, but instead merely uses the Binary Object Identifier 74 to determine whether a particular binary object has been modified and, as a result, should be backed up. *Id.* PersonalWeb also alleges that Woodhill’s auditing technique does not use the Binary Object Identifier 74 to access a binary object. *Id.* at 41-42.

Woodhill discloses that the Distributed Storage Manager program 24 performs auditing and reporting functions on a periodic basis in order to ensure that the binary objects which already have been backed up may be restored. Ex. 1005, Spec. 18:11-13. According to Woodhill, the Distributed Storage Manager program 24 initiates a restore of a randomly selected binary object identified by a Binary Object Identification Record 58 stored in the File Database 25. Ex. 1005, Spec. 18:16-19. As previously disclosed in Woodhill, the Binary Object Identification Record 58 includes, amongst other things, a Binary Object Identifier 74, which is a unique identifier for each binary object. Ex. 1005, Spec. 4:35-47, 7:64-8:1. In light of the cited disclosures, EMC’s position that Woodhill relies upon a unique identifier—the Binary Object Identifier 74 of the Binary Object Identification Record 58—to access a binary object is reasonable. Therefore, we are persuaded that EMC has demonstrated that Woodhill describes the “accessing” method step recited in independent claim 30.

With respect to independent claim 33, PersonalWeb contends that Woodhill fails to describe the following claim limitations:

determining, using the data identifier, whether the data item is present at the destination location; and

based on the determining whether the data item is present [at the destination location], providing the destination location with the data item only if the data item is not present at the destination.

Prelim. Resp. 45. In particular, based on the assumption that Woodhill's remote backup server 12 is the claimed "destination location," PersonalWeb alleges that Woodhill does not determine whether the binary object is present at the remote backup file server. *Id.*

As discussed above, Woodhill determines whether a particular binary object is present in the networked computer system 10 by examining the Binary Object Identifiers 74 of a plurality of binary objects located within the system. Ex. 1005, Spec. 8:62-65. Woodhill further discloses that the Distributed Storage Manager program 24 backups each binary object by storing a compressed copy of the binary object in two locations—(1) on a disk drive 19 associated with a local computer 20; and (2) on the remote backup file server 12. Ex. 1005, Spec. 9: 31-38. If the binary object was ever lost or destroyed at the disk drive 19 on the local computer 21, Woodhill indicates that the binary object stored at the remote backup file server 12 may be copied to the disk drive 19 on the local computer 21. Ex. 1005, Abstract; Spec. 2:25-27.

For instance, Woodhill may determine whether a binary object is present at a disk drive 19 on a local computer 21, *i.e.*, destination location, by examining the Binary Object Identifiers 74 located on the disk drive 19. If the binary object is not present on the disk drive 19, the remote backup file server 12 could copy the binary object to the disk drive 19.

Alternatively, Woodhill may determine whether a binary object is present at the remote backup file server 12, *i.e.*, destination location, by examining the Binary Object Identifiers 74 located on the remote backup file server 12. If the binary object is not present on the remote backup file server 12, the disk driver 19 on the local computer 21 could copy the binary object to the remote backup file server 12. Therefore, we are persuaded that EMC has presented sufficient evidence demonstrating that Woodhill describes the disputed subject matter recited in independent claim 33.

With respect to dependent claim 41, PersonalWeb contends that EMC does not provide a detailed analysis pertaining to how Woodhill describes that claimed subject matter. Prelim. Resp. 43. Contrary to PersonalWeb's argument, EMC provides a brief discussion of how Woodhill describes the claimed subject matter recited in dependent claim 41. Pet. 58. In addition, EMC provides a detailed claim chart that indicates the textual portions of Woodhill relied upon to describe the method steps recited in dependent claim 41. Ex. 1041, p. 21. Because the explanations provided by EMC as to how Woodhill describes dependent claim 41 have merit and are otherwise un rebutted, we are persuaded that EMC has demonstrated that Woodhill describes the claimed subject matter recited in dependent claim 41.

For the foregoing reasons, we conclude that there is a reasonable likelihood that EMC will prevail on its assertion that claims 1, 30, 33, and 41 of the '791 patent are anticipated by Woodhill. We authorize an *inter partes* review on this ground of unpatentability. PersonalWeb does not present separate and distinct arguments with respect to claims 2-4, 29, 31, and 32.

The explanations provided by EMC as to how Woodhill describes the claimed subject matter recited in those claims have merit and are otherwise un rebutted. Therefore, we conclude that there is a reasonable likelihood that EMC will prevail on its assertion that claims 2-4, 29, 31, and 32 of the '791 patent are anticipated by Woodhill. We also authorize an *inter partes* review on this ground of unpatentability.

*B. 35 U.S.C. § 103(a) Grounds of Unpatentability—Woodhill
Claim 1-4 and 29*

EMC contends that claims 1-4 and 29 are unpatentable under 35 U.S.C. § 103(a) over Woodhill. Pet. 59. In particular, EMC argues that one with ordinary skill in the art would have found it obvious to calculate Woodhill's Binary Object Identifier using an MD5 hash function. *Id.* (citing to Ex. 1005, Spec. 8:52-58; Ex. 1009, ¶¶ 97 and 98). EMC also provides an articulated reason with a rational underpinning to justify the legal conclusion of obviousness. *Id.*

To rebut the ground of unpatentability asserted by EMC against claims 1-4 and 29, PersonalWeb relies upon essentially the same arguments it presented in rebutting the grounds of unpatentability asserted by EMC against claims 1-4, 29-33, and 41. Prelim. Resp. 35-47. We already have addressed those arguments in our previous discussion of claims 1-4, 29-33, and 41.

For the same reasons set forth above in our discussion of claims 1-4, 29-33, and 41, we conclude that there is a reasonable likelihood that EMC will prevail on its assertion that claims 1-4 and 29 of the '791 patent are

unpatentable over Woodhill. We authorize an *inter partes* review on this ground of unpatentability as well.

C. Remaining Grounds of Unpatentability

Claims 1-4, 29-33, and 41

EMC contends that claims 1-4, 29-33, and 41 are unpatentable under 35 U.S.C. §§ 102 and 103, based in whole or in part on Browne, Langer, Kantor, or Woodhill. Pet. 26-51, 59-60. Those grounds of unpatentability are redundant to the grounds of unpatentability on which we initiate an *inter partes* review. Accordingly, we do not authorize an *inter partes* review on the remaining grounds of unpatentability asserted by EMC against claims 1-4, 29-33, and 41 of the '791 patent. *See* 37 C.F.R. § 42.108(a).

V. ORDER

It is **ORDERED** that pursuant to 35 U.S.C. § 314(a), an *inter partes* review is hereby instituted as to claims 1-4, 29-33, and 41 of the '791 patent for the following grounds of unpatentability:

A. Claims 1-4, 29-33, and 41 as anticipated under U.S.C. § 102(e) by Woodhill.

B. Claims 1-4 and 29 as unpatentable under 35 U.S.C. § 103(a) over Woodhill.

It is **FURTHERED ORDERED** that 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. The trial will commence on the entry date of this decision.

Case IPR2013-00082
U.S. Patent No. 5,978,791

It is **FURTHER ORDERED** that an initial conference call with the Board is scheduled for 2PM on June 3, 2013. The parties are directed to the Office Trial Practice Guide, 77 *Fed. Reg.* 48756, 48765-66 (Aug. 14, 2012) for guidance in preparing for the initial conference call, and should come prepared to discuss any proposed changes to the Scheduling Order entered herewith and any motions the parties anticipate filing during the trial.

Case IPR2013-00082
U.S. Patent No. 5,978,791

For PETITIONER:

Peter Dichiara
David L. Cavanaugh
WILMER CUTLER PICKERING HALE AND DORR LLP
Peter.Dichiara@wilmerhale.com
David.Cavanaugh@wilmerhale.com

For PATENT OWNER

Joseph A. Rhoa
Updeep S. Gill
NIXON & VANDERHYE P.C.
jar@nixonvan.com
usg@nixonvan.com

cu