

(19) World Intellectual Property  
Organization  
International Bureau



(43) International Publication Date  
2 June 2005 (02.06.2005)

PCT

(10) International Publication Number  
**WO 2005/050489 A1**

(51) International Patent Classification<sup>7</sup>: **G06F 17/30**,  
H02H 3/05

(21) International Application Number:  
PCT/US2004/038195

(22) International Filing Date:  
15 November 2004 (15.11.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
60/520,452 13 November 2003 (13.11.2003) US

(71) Applicant (for all designated States except US): **COMM-  
VAULT SYSTEMS, INC.** [US/US]; Two Crescent Place,  
Oceanport, NJ 07757-0090 (US).

(72) Inventors; and


(75) Inventors/Applicants (for US only): **LU, Jun** [CN/US];  
2118 Aldrin Road, Apt. 1A, Ocean, NJ 0771 (US).  
**PRASAD, Arun** [IN/IN]; 2406 E Block, Jothy, Illam,,  
Sahakara, Nagar, Bellary Road, Bangalore 560092 (IN).

(74) Agent: **SHANAHAN, Michael**; Brown Raysman Mill-  
stein, Felder & Steiner LLP, 900 Third Avenue, New York,  
NY 10022 (US).

(81) Designated States (unless otherwise indicated, for every  
kind of national protection available): AE, AG, AL, AM,  
AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN,  
CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI,  
GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE,  
KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD,  
MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG,  
PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM,  
TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM,  
ZW.

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR STORED DATA ARCHIVE VERIFICATION



(57) Abstract: Methods and systems are described for verifying stored data by receiving first set of metadata associated with first set of stored data (200). Generating second set of metadata associated with a second set of stored data (210) and comparing first set of metadata and the second set of metadata (220).



(84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

**Published:**

— *with international search report*

**SYSTEM AND METHOD FOR STORED DATA ARCHIVE VERIFICATION****COPYRIGHT NOTICE**

[0001] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the  
5 facsimile reproduction by anyone of the patent document or the patent disclosures, as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all copyright rights whatsoever.

**PRIORITY APPLICATIONS**

This application claims the benefit of U.S. provisional application no.  
10 60/520,452 titled System and Method For Stored Data Archive Verification, filed November 13, 2003, which application is incorporated herein by reference in its entirety.

**RELATED APPLICATIONS**

[0002] This application is related to the following pending applications, each of which is hereby incorporated herein by reference in its entirety:

- 15 • U.S. Patent No. 6,418,478, titled PIPELINED HIGH SPEED DATA TRANSFER MECHANISM, issued July 9, 2002, attorney docket number 4982/6;
- U.S. Provisional Patent Application Serial No. 60/460,234, titled  
20 SYSTEM AND METHOD FOR PERFORMING STORAGE OPERATIONS IN A COMPUTER NETWORK, filed April 3, 2003, attorney docket number 4982/35PROV;
- Application Serial No. 10/818,749, titled SYSTEM AND METHOD FOR DYNAMICALLY PERFORMING STORAGE

OPERATIONS IN A COMPUTER NETWORK, filed April 5, 2004,  
attorney docket number 4982/35;

- Application Serial No. 10/260,209, titled SYSTEM AND  
METHOD FOR ARCHIVING OBJECTS IN AN INFORMATION  
5 STORE, filed September 30, 2002, attorney docket number 4982/19;  
and
- Application Serial No. 09/495,751, titled HIGH SPEED DATA  
TRANSFER MECHANISM, filed February 1, 2000, attorney docket  
number 4982/7.

10 BACKGROUND OF THE INVENTION

[0003] The invention disclosed herein relates generally to data storage systems in  
computer networks and, more particularly, to improvements in storing and verifying  
accurate archiving of electronic data.

[0004] Storage architectures used by individual computers or data stores to store  
15 electronic data typically include volatile storage media such as Random Access Memory  
“RAM”, and one or more nonvolatile storage devices such as hard drives, tape drives,  
optical disks, and other storage devices that form a part of or are directly associated with  
an individual computer. A network of computers such as a Local Area Network “LAN”  
or a Wide Area Network “WAN”, typically store electronic data via servers or stand-  
20 alone storage devices accessible via the network. Stand-alone storage devices are  
generally connected to one individual computer or a network of computers. Network  
storage devices commonly known in the art typically include physical drives in which  
tapes or other storage media are stored and a robotic arm which is used to place the tapes  
or storage media into the drives. Examples of network storage devices include

networkable tape drives, optical libraries, Redundant Arrays of Inexpensive Disks  
“RAID”, CD-ROM jukeboxes, and other devices.

[0005] Electronic data is typically archived according to a schedule, for example, data is designated to be copied and stored once a day. Generally, data is archived in the event that an original copy becomes unavailable, for example, the data is destroyed, lost or otherwise inaccessible. In general, the data is directed to a system component to be copied or backed up to storage media, such as an auxiliary copy, copy backup or quick recovery copy. Some systems check the copy or backup to ensure the copy or backup is accurate. Generally, the accuracy check includes steps such as analyzing each data item copied and comparing it to the original data, or other method. Such verification methods can be lengthy and time consuming for copies of large volumes of data, requiring significant use of system resources. Alternatively, some systems use cursory data checks, such as comparing filenames, which is less time consuming, however, it yields a less reliable data check.

## 15 BRIEF SUMMARY OF THE INVENTION

[0006] The present invention addresses the problems discussed above, and includes a method for verifying stored data. Methods and systems are described for verifying stored data. The system includes components such as a storage manager, media agent, storage device, jobs manager, archive check manager, index, copy manager, client computer, and data store. The storage system verifies data by receiving a first set of metadata associated with a first set of stored data, generating a second set of metadata associated with a second set of stored data, and comparing the first set of metadata and second set of metadata. Alternatively, the storage system generates a first set of metadata associated with a first set of stored data, generates a second set of stored data associated

with the first set of stored data, generates a second set of metadata associated with the second set of stored data, and compares the first set of metadata and the second set of metadata.

[0007] The method involves receiving a first set of metadata associated with a first set of stored data. The first set of metadata is typically generated based on the first set of stored data, or otherwise associated with the first set of stored data. The first set of metadata is received, for example, from a local system component, or external source. The first set of metadata may also be received in response to a retrieval storage operation initiated by a system component. The method generates a second set of metadata associated with a second set of stored data. Preferably, the second set of metadata is generated according to the same procedure used to generate or create the first set of metadata. Additionally, it is preferred that the second set of stored data be data that is associated with the first set of stored data, such as a backup, auxiliary copy, or other copy. The method compares the first set of metadata and the second set of metadata. In general, a comparison of the first set of metadata and the second set of metadata can reveal that the metadata sets are the same, or different. If the comparison reveals that the first and second sets of metadata are the same, then generally, the first set of stored data and the second set of stored data are also the same, and there was an accurate copy made in the second set of stored data. If the comparison of the first and second sets of metadata reveals differences, then the first and second sets of stored data are not the same, and there may have been an error in creating the second set of stored data. In the event that inconsistencies are found between the first and second sets of metadata, the method may initiate a storage operation, such as a copy or backup the first set of stored data.

[0008] In another embodiment of the invention, the method for verifying stored data includes performing a storage operation, for example, a copy or backup operation, or other storage operation as further described herein. In performing the storage operation, a first set of metadata is generated and is associated with a first set of stored data.

5 Preferably, the first set of metadata is generated based on the first set of stored data. The first set of stored data is preferably an original or primary copy, or other original set of data. Additionally, a second set of stored data is generated and is associated with the first set of stored data. Preferably, the second set of stored data is a backup, auxiliary copy, other copy, or otherwise based on the first set of stored data. The method generates a  
10 second set of metadata based on the second set of stored data. Preferably the second set of metadata is generated according to the same procedure as the first set of metadata. The first set of metadata and the second set of metadata are compared. The comparison of the first set of metadata and the second set of metadata provide an accuracy check for the first set of stored data and the second set of stored data.

15 BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The invention is illustrated in the figures of the accompanying drawings which are meant to be exemplary and not limiting, in which like references are intended to refer to like or corresponding parts, and in which:

[0010] Fig. 1 is block diagram showing a high-level view of the storage network  
20 architecture and components according to an embodiment of the invention;

[0011] Fig. 2 is a flow diagram for verifying stored data according an embodiment of the invention;

[0012] Fig. 3 is a flow diagram for verifying stored data according to an embodiment of the invention;

[0013] Fig. 4 is a block diagram showing a chunk, according to an embodiment of the invention;

5 [0014] Fig. 5 is a flow diagram for a storage operation according to an embodiment of the invention;

[0015] Fig. 6 is a flow diagram for verifying stored data, according to an embodiment of the invention;

[0016] Fig. 7 is a block diagram showing a stored data verification table  
10 according to an embodiment of the invention; and

[0017] Fig. 8 is a detailed flow diagram of a step of the flow diagram of Fig. 6 for verifying stored data, according to an embodiment of the invention.

#### DETAILED DESCRIPTION

[0018] Preferred embodiments of the invention are now described with reference  
15 to the drawings. An embodiment of the system of the present invention is shown in Fig. 1. The system and components of the system as presented in Fig. 1 are exemplary of a modular storage system such as the CommVault Galaxy™ backup and retrieval system and QiNetix™ storage management system, available from CommVault Systems, Inc. of Oceanport, NJ, and further described in U.S. Provisional Patent Application Serial No.  
20 60/460,234, titled SYSTEM AND METHOD FOR PERFORMING STORAGE OPERATIONS IN A COMPUTER NETWORK, filed April 3, 2003, which is incorporated herein by reference in its entirety.

[0019] As shown in Fig. 1, the system includes a data store 100, a client computer 110, a storage manager (or storage manager component) 120, an SM index 130, a jobs manager 140, an archive check manager 150, one or more media agents (or media management components) 160, an MA index 170, a copy manager 180, and one or more storage devices 190. The data store 100 contains data generated by a typical enterprise machine, such as client computer 110. For example, a data store 100 may contain Oracle data for a database, or mail server data. In the event that a data store 100 loses data, the data store 100 obtains the lost data from a backup or other data copy. The data store 100 is generally in communication with the client computer 110

10 [0020] The client computer 110 may be a single client computer 110 or a plurality of client computers 110. Client computer 110 generally refers to a computer with data, e.g. a primary data set, which may be backed up, such as a personal computer, a workstation, a server computer, a host computer, etc. The client computer 110 typically also includes software to facilitate data transfers from the data store 100 to the other system components. The client computer 110 is generally in communication with the storage manager 120. In some embodiments of the invention, the client computer 110 is also in communication with the media agent 160.

[0021] The storage manager 120 is generally a software module or modules capable of coordinating and controlling a storage system. In preferred embodiments, the storage manager 120 contains an SM index 130, a jobs manager 140, and an archive check manager 150. In other embodiments of the invention, the SM index 130, jobs manager 140, or archive check manager 150 may exist independently within the storage system, or reside within different components of the storage system, such as the media agent 160. The storage manager 120 communicates directly the client computer 110 and

media agent 160 to initiate, perform and manage system storage operations. Storage operations include, but are not limited to, creation, storage, retrieval, migration, deletion, and tracking of primary or production volume data, secondary volume data, primary copies, secondary copies, auxiliary copies, snapshot copies, backup copies, incremental  
5 copies, differential copies, HSM copies, archive copies, Information Lifecycle Management ("ILM") copies, and other types of copies and versions of electronic data. The storage manager 120 manages system storage operations, and together with the job manager 140 and archive check manager 150, manages and directs data verification, including, for example, updating the data verification information to SM index 130.

10 **[0022]** The storage manager 120 includes SM index 130, which is generally any storage index that is maintained to store index data for storage operations, such as data verification operations. Data stored to the SM index 130 in data verification operations includes for example, data copied or backed up, the metadata generated in association with data sets, the results of verification checks, and other information generated by the  
15 system during storage operations. The SM index 130 can also be a storage index known in the art, or other storage index, such as, the index cache of CommVault's Galaxy<sup>TM</sup> as further described in Application Number 10/260,209 which is incorporated herein by reference in its entirety. The SM index 130 stores data related to the storage system components, including information related to each data set, copies, backups, metadata, or  
20 other information. Index data is useful because, among other benefits, it provides the system with an efficient mechanism for locating information associated with performing storage operations.

**[0023]** The storage manager 120 also includes the jobs manager 140, which is a module (or modules) that schedules and monitors storage operations that are to be

performed, are being performed, or have been completed by the system. The jobs manager 140 can schedule storage operations automatically, for example, based on storage policies. Storage operations can also be scheduled manually, for example, a user may enter an input to the jobs manager 140, using, for example, a user interface to  
5 schedule a storage operation. The jobs manager 140 directs the archive check manager 150 to initiate data verification operations, as further described herein. The jobs manager 140 is preferably included in the storage manager 120, but may exist independently in the system, or within another system component, such as media agent 160.

**[0024]** A storage policy is generally a data structure or other information which  
10 includes a set of preferences and other storage criteria for performing a storage operation. The preferences and storage criteria may include, but are not limited to: a storage location, relationships between system components, network pathway to utilize, retention policies, data characteristics, compression or encryption requirements, preferred system components to utilize in a storage operation, and other criteria relating to a storage  
15 operation. A storage policy may be stored to a storage manager index, to archive media as metadata for use in restore operations or other storage operations, or to other locations or components of the system.

**[0025]** The archive check manager 150 is also included in the storage manager 120, and is a module (or modules) that manages data verification operations within the  
20 system, as further described herein. For example, the archive check manager 150 identifies data to be verified in the data verification operation and manages the data verification operation. The archive check manager 150 is preferably included in the storage manager 120, but may exist independently in the system, or within another system component, such as media agent 160.

[0026] The media agent 160 is generally a module or modules capable of managing storage operations, such as copying and backups, or other storage-related operation. The media agent 160 preferably includes MA index 170 and copy manager 180. The media agent 160 is in communication with the storage manager 120 and storage  
5 device 190, for example, via a local connection or a network. In some embodiments of the invention, the media agent 160 is also in communication with the client computer 110. The media agent 160 is capable of communicating instructions from a storage manager 120 to a storage device 190 to undertake a storage operation.

[0027] Additionally, the media agent 160 is capable of directing the copy manager  
10 180, which is responsible for initiating and performing a copy operation. The media agent 160 stores data related to storage operations and media, such as auxiliary copies, to MA index 170. Examples of data stored by the media agent 160 to the MA index 170 include metadata or other data associated with storage operations.

[0028] The MA index 170 may be an index, such as SM index 130 described  
15 herein, the media agent index described in U.S. Provisional Patent Application Serial No. 60/460,234, titled SYSTEM AND METHOD FOR PERFORMING STORAGE OPERATIONS IN A COMPUTER NETWORK, filed April 3, 2003, or other storage index known in the art. The copy manager 180 is a module or modules used to initiate, manage and perform copy operations or other copy operations. A copy operation is a  
20 storage operation that generally includes creating a duplicate copy of data for storage. Copies are typically made for disaster recovery purposes.

[0029] The copy manager 180 preferably resides within the media agent 160, but can also exist as an independent system component, or within another system component.

[0030] The storage device 190 is generally one or more devices as described herein. The storage device 190 can be a CommVault Galaxy™ storage device, or other storage device capable of performing a storage operation. Storage device 190 typically includes components such as media, slots, media drives and a robotic arm or arms.

5 [0031] Data generated by the system or otherwise obtained by the system, is typically received by data store 100, or client computer 110. A copy of the data, for example, an auxiliary copy or other copy, is made according to storage policies, as further described herein. When data is copied, the copy is generally transferred to the storage device 190 via the storage manager 120 or media agent 160. The data copy is checked to  
10 verify that the contents of the copy are identical to the original data.

[0032] Referring to Fig. 2, which depicts a flow chart of a method of an embodiment of the invention for verifying data, a first set of metadata associated with a first set of stored data is received, step 200. Generally, when a copy of the original or first set of data is created, a set of metadata related to the first set of data is created. In  
15 preferred embodiments, the metadata is generated by the copy manager, or other system component, and is received by a media agent, copy manager or other system component. The first set of metadata is associated with a first set of stored data, for example, the first set of metadata is generated based on the first set of stored data. In preferred  
embodiments, the first set of stored data is an original set of data, primary copy, other  
20 original data copy or verified copy. The first set of metadata is preferably associated with the first set of stored data according to a procedure known to the storage system. For example, the first set of metadata can be generated according to a routine based on the first set of stored data, obtained using variables based on the first set of stored data, generated as a chunk header, as further described herein, or other procedure for obtaining

metadata based on the first set of stored data, or other data. In preferred embodiments, the first set of metadata is stored to an index, such as the SM index or MA index.

[0033] A second set of metadata associated with a second set of stored data is generated, step 210. In preferred embodiments, the second set of stored data is associated with the first set of stored data. For example, the second set of stored data is typically a backup, auxiliary copy or other copy of the first set of stored data. The second set of metadata is generated based on the second set of stored data. Preferably, the procedure for generating the second set of metadata is the same procedure used to generate the first set of metadata. In one embodiment of the invention, the second set of metadata is generated by a copy manager, a media agent, or other system component. In some embodiments, the second set of metadata is generated in a partial copy operation, wherein, a set of data is read, metadata based on the set of data is generated, but no copy of the data file is made. In another embodiment of the invention, the second set of metadata is generated by the archive check manager. In preferred embodiments, the second set of metadata is stored to an index, such as the SM index or MA index .

[0034] The first set of metadata and the second set of metadata are compared, step 220. The first set of metadata and the second set of metadata are compared to determine whether the sets of metadata are similar, equivalent, or otherwise indicate that the underlying data which generated the metadata are substantially similar. If the comparison of the first set of metadata and the second set of metadata indicates that the sets of metadata are similar, the underlying data used to generate the first and second sets of metadata, in this case, the first and second sets of stored data, are substantially similar. Thus, the second set of stored data is determined to be an accurate or true copy of the first set of stored data. When the first set of metadata and the second set of metadata are

determined to be different or inconsistent, or otherwise indicate that the underlying data used to generate the first and second sets of metadata are different, the first and second sets of stored data may not be equivalent, and therefore, the second set of stored data is likely an inaccurate or otherwise defective copy of the first set of stored data.

5   **[0035]**        In preferred embodiments, a copy manager compares the first set of metadata and the second set of metadata. In another embodiment, an archive check manager compares the first set of metadata and the second set of metadata. The first set of metadata and the second set of metadata are compared to determine whether the sets of metadata are the same. In some embodiments, one or both of the sets of metadata are  
10   obtained from an index, such as the MA index, SM index, or are generated as needed, for example, by performing a partial storage operation, further described herein, or other method for obtaining metadata.

**[0036]**        Referring to Fig. 3, which depicts a flowchart of another embodiment of the invention, a storage operation is initiated, step 250. The storage operation is normally  
15   initiated by a storage manager, for example, by a jobs manager, which, as described herein, schedules storage operations, for example, based on storage policies. In other embodiments, other system components initiate system storage operations. The jobs manager, or other system component, schedules storage operations and directs an archive check manager, or other system component, to initiate a storage operation. The storage  
20   operation initiated at step 250 is preferably an operation similar to a copy operation. For example, a data file in a client data store is copied. A copy manager, or other system component, generates a first set of metadata associated with a first set of stored data, step 260. For example, the copy manager initiates a copy operation, and obtains the first set of stored data. The copy manager reads the first set of stored data and generates a first set of

metadata. As described herein, the first set of metadata is generated according to a procedure, known to the system, such as according to a routine based on the first set of stored data, or obtained using a variable based on the first set of stored data, generated as a chunk header, or other procedure for obtaining metadata. The first set of stored data is  
5 generally an original set of data, such as a primary copy, other original copy or verified copy.

[0037] The copy manager or other system component generates a second set of stored data associated with first set of stored data, step 270. Preferably, the second set of stored data is a backup, auxiliary copy, or other copy of the first set of stored data. The  
10 second set of stored data is generally obtained by directing the copy manager to read the first set of stored data and write a copy, such as the second set of stored data, but may also be generated using other copy methods and system components. The copy manager, or other system component, generates a second set of metadata associated with the second set of stored data, step 280. Specifically, the second set of metadata is generated based on  
15 the second set of stored data. The second set of metadata is preferably obtained using the same procedure used to generate the first set of metadata in step 260, or other procedure for generating metadata, for example, generating metadata based on data from a primary copy. The copy manager or other system component, compares the first set of metadata and the second set of metadata, step 290. The first set of metadata and the second set of  
20 metadata are compared to determine whether the first set of metadata and the second set of metadata are equivalent. As described herein, the comparison of the first and second sets of metadata may be used as an indicator of whether the second set of stored data is an accurate copy of the first set of stored data.

[0038] When data is copied, for example, when the media agent or copy manager copies a first set of stored data and creates a second set of stored data, the data is generally copied and transferred between system components using chunks, such as the system and methods described in U.S. Patent No. 6,418,478, titled PIPELINED HIGH SPEED DATA TRANSFER MECHANISM, issued July 9, 2002. Chunks are useful because they can be used to assemble data into packages that are transmitted across system networks, for example according to system resource availability.

[0039] Referring to Fig. 4, which depicts a chunk according to an embodiment of the invention, chunks are generally a fixed file size, such as 512 MB, and include a chunk header 300 and a chunk payload 310. The chunk header 300 includes data or information such as an archive file ID 320, a chunk sequence number 330, a chunk volume 340, an offset 350, a chunk volume file mark number 360, and other information 370. A chunk payload 310 includes data files or a subset of data files that were copied, e.g. an initial or first set of data. Each of the data items or information included in the chunk header 300 are related to the file data 380-389 in the chunk payload 310. The chunk header 300 data items include metadata generally derived from the stored data or files in the chunk payload 310. For example, a copy manager or other system component which creates, packages, or otherwise handles chunks, reads the data contained in the chunk payload 310 to generate metadata or other information in the chunk header 300.

[0040] The archive file ID 320 is an alphanumeric identifier for the file being archived, for example, in the payload 310. The chunk sequence number 330 provides the sequence information for the chunk. The chunk volume 340 indicates the volume of the chunk utilized, for example, the volume of the data or files on the chunk. Offset 350 indicates the distance the data item is offset, for example, which is information related to,

and used in reassembling the chunk files. Chunk volume file mark number 360 refers to information related to the volume of the file. The chunk header 300 also contains other information 370, which may be data, variables, information or other metadata related to files 380-389 in the payload 310. Files 380-389 of the payload represent that a subset of a data file, a set of data, a file or a block, or several sets of data, files or blocks may be stored in a chunk payload, for example, file 1 to file n. The information contained in the chunk header 300 is preferably stored to an index, such as an MA index or SM index. Storing the information in the chunk header 300 to the index is useful so that the information can be readily accessed by a media agent, storage manager, copy manager or other system component.

**[0041]** Referring to Fig. 5, which depicts a flow diagram of a storage operation, a storage operation is initiated, step 400. Preferably the storage operation is initiated by a storage manager, or storage manager components, such as a jobs manager, an archive check manager, a media agent, or media agent components, such as a copy manager or other system component. The storage operations initiated at step 400 include, for example, system backups, media movement, imports, exports, migrations, restores, data verification and other storage-related operations. Files are identified for the storage operation, step 410. The files are identified by the media agent, storage manager, archive check manager, copy manager, or other system component. In some embodiments, the files may be identified according to storage policies associated with a particular file. As further described herein, information related to files and storage policies are stored in, and may be obtained from an index, such as an SM index or MA index, or other locations within the system.

**[0042]** In general, the storage operation is assigned to a particular media agent to manage the operation. The media agent is typically associated with the file, storage device, or other system component utilized in the storage operation. Additionally, the media agent may be designated to manage the operation by storage policies. In some  
5 embodiments, the media agent may be dynamically reassigned during a storage operation according to storage policies, system resource availability or other variable. The copy manager used in storage operations is typically the copy manager associated with the media agent assigned to manage the operation. It is understood that the copy manager may be dynamically reassigned during the storage operation, for example, when its media  
10 agent is reassigned, or according to system resource availability or other variables.

**[0043]** A chunk header is generated for the files identified in step 410, and the files are packaged into chunks, step 420. The chunk header is generated according to a procedure known to the system, such as the procedure described herein in connection with generating metadata. For example, the chunk header is generated according to  
15 specific terms or variables based on the files, such as the data or information items included in the chunk header. Preferably, the copy manager creates the chunk header. The files are packaged into chunks, for example, into a payload. The chunk header information is updated to an index, such as an SM index or MA index.

**[0044]** The chunk is transferred, step 430, for example, via the networks using, or  
20 as directed by, a media agent, or other system component. In general, the chunk is transferred through the networks according to system resource availability. The chunk transfer destination is preferably a storage device assigned, for example, according to storage policies, or as directed by media agent, storage manager, or other system component.

[0045] The chunk is received, step 440 at the storage device, mentioned herein. The storage device preferably communicates data indicating that the chunk has been received to the media agent, storage manager or other system component. The chunk receipt data is updated to an index, such as the SM index or MA index. Upon receipt of  
5 the chunk, the chunk is processed, step 450. The chunk is generally processed by sequence number, or other information that is contained in the chunk header to reassemble the files. In preferred embodiments, the copy manager reads the chunk header and stores the data contained in the chunk header to an index, such as the SM index or MA index. In some embodiments, the data or files contained in the chunk payload are  
10 extracted from the chunk according to the chunk header information. The copy manager copies the chunk file information, for example, to storage media, such as disks, tapes, CD-ROMs or other media for storage. The data extracted from the chunk is stored, step 460. The data is generally stored to a storage device.

[0046] Referring to Fig. 6, which depicts a flow chart for verifying stored data  
15 according to one possible embodiment of the invention, a data verification operation is initiated, step 500. In preferred embodiments the data verification operation is scheduled by the jobs manager according to storage policies, and the job manager directs the archive check manager to initiate a data verification operation. In other embodiments, the data verification operation is initiated by the storage manager, media agent, or other system  
20 component. The data verification operation, as described herein, determines and verifies whether a copy of data is an accurate copy of an original data set.

[0047] Data to be verified in the data verification operation is identified, step 510. In preferred embodiments, the data to be verified is identified, for example, by referencing a stored data verification table, such as the stored data verification table 600

depicted in Fig. 7. In other embodiments, the data to be verified is identified according to information contained in an index, such as an SM index, an MA index, storage policy, or other system component, or as directed by a user. Preferably, the archive check manager identifies the data to be verified by referring to the stored data verification table 600, the  
5 index, such as the SM index or MA index, or other data source. In other embodiments of the invention, the storage manager, jobs manager, media agent, copy manager or other system component identifies the data to be verified. In some embodiments, the stored data verification table 600 data may be stored, for example to the SM index or MA index.

[0048] Referring to Fig. 7, the data verification table 600 includes data such as the  
10 data shown in a column titled chunk ID 610, a column titled chunk verified 620, a column titled result 630, or other information. Chunk ID 610 is an alphanumeric identifier for the file, block, chunk or data set stored, such as the files, blocks, chunks, or data sets stored to chunk payload. Chunk verified 620 provides a yes or no flag indicating whether the chunk or data has been verified, for example, in a data verification operation, as described  
15 herein. Result 630 indicates the result of the data verification operation, such as good copy, unsuccessful, bad file, bad incremental block, or other result. In addition to verifying whether the data is accurate, the data verification operation identifies with particularity which part of the file or copied data is not a good copy, such as the result in column result 630 which indicates that file 5 is bad.

20 [0049] Referring again to Fig. 6, as mentioned herein, in general, the jobs manager schedules data verification operations and other storage operations, and directs the archive check manager to identify data to be verified. In some embodiments, the archive check manager refers to an index, such as an MA index to identify data to be verified. In other embodiments, the archive check manager identifies the data to be

verified in step 510 by referring to the stored data verification table. In some embodiments, the archive check manager refers to the stored data verification table by chunk identifier to determine whether the chunk has been verified, for example, according to the data contained in a chunk verified column. In other embodiments, the archive  
5 check manager refers to the stored data verification table by data contained in the chunk verified column to determine which chunks have not been verified, and obtaining the chunk identifier for unverified chunks.

[0050] The archive check manager obtains the list of data for verifying, as described herein, and sends it to the copy manager, or other system component, step 520.  
10 The list of data for verifying can be sent, for example as data stored to an index, such as an SM index or MA index, or in a communication to the copy manager. Generally, upon receipt of the list of data to be verified, the copy manager verifies the data, step 530.

[0051] An embodiment of step 530 is depicted in the detailed flowchart of Fig. 8. The copy manager, or other system component, initiates a copy operation, step 700. The  
15 copy manager receives a first set of stored data, or a primary copy, or other original copy of data, and a first set of metadata, step 710. The first set of stored data is typically stored to a storage device and information associated with the first set of stored data, such as the first set of metadata, is stored to an index, or other system component. The index, or other system component, provides the source of the first set of stored data and the first set  
20 of metadata received by the copy manager. In another embodiment, the copy manager, or other system component, generates the first set of metadata, as described herein, for the first set of stored data. In preferred embodiments, the copy manager performs a copy read operation, wherein, the first set of stored data is read, and a set of metadata is generated. The first set of metadata is stored to an index, or other system component.

[0052] The copy manager, or other system component, generates a second set of metadata, step 720. In preferred embodiments, the copy manager reads a copy of the first set of stored data, such as an auxiliary copy, or other copy and performs a partial storage operation. The partial storage operation is generally a copy operation, such as a copy operation used to create a copy of the first set of stored data, auxiliary copy or other copy. The partial storage operation includes reading the data file, creating metadata associated with the data file, but does not complete the copy operation. Thus, only metadata is generated in the partial storage operation. In preferred embodiments, the procedure for generating the metadata in the partial storage operation is the same as the procedure used to generate the first set of metadata, and the first and second sets of metadata are preferably chunk headers, such as chunk header 300. A copy of the second set of metadata can be stored for later reference to an index, or used immediately in a comparison of the first set of metadata and second set of metadata, step 730.

[0053] Referring back to Fig. 6, the result of the data comparison, for example, the comparison of the first set of metadata is updated to the stored data verification table, step 540. If the first and second sets of metadata are the same, the copy manager updates the stored data verification table to reflect the result to indicate the second set of stored data is a good copy. Alternatively, if the first and second sets of metadata are different, the copy manager can determine which of the files, for example, which of the files in the chunk, are inconsistent, and update the results data accordingly, for example, the data included in the column results 630 of Fig. 7, such as file 5 is bad, bad incremental block, the verification was unsuccessful, or other result. In preferred embodiments, when the data verification result is something other than good copy, the copy manager initiates another copy or backup operation to copy or back up the first set of stored data.

**[0054]** Systems and methods disclosed herein are used generally to check or verify all data in a storage system. In some embodiment, the systems and methods can be applied selectively to portions of data. For example, in some embodiments, data within the system can be arbitrarily grouped into backup sets or in CommVault system backup sets, generally denoting arbitrary groups of data wherein data in one backup set can also be in another backup set, such that it is not exclusively in one backup set. At the sub-client level, data in one sub-client group is not in another sub-client group and is exclusive to a sub-client. As described herein, backup sets may have several sub-clients in a set, at the application level, or as specified in a storage policy for a group of data.

10 **[0055]** While the invention has been described and illustrated in connection with preferred embodiments, many variations and modifications as will be evident to those skilled in this art may be made without departing from the spirit and scope of the invention, and the invention is thus not to be limited to the precise details of methodology or construction set forth above as such variations and modification are intended to be included within the scope of the invention.

15

**Claims:**

What is claimed is:

1. A method for verifying stored data in a storage management system, the method comprising:
  - 5 performing a first storage operation on a first set of stored data,  
wherein the first storage operation generates a first set of metadata associated with the first set of stored data and a second set of stored data associated with the first set of stored data;  
generating a second set of metadata associated with the second set of  
10 stored data; and  
comparing the first set of metadata and the second set of metadata.
2. The method of claim 1 wherein generating the second set of metadata comprises performing a second storage operation on the second set of stored data.
3. The method of claim 2 wherein the first storage operation and the second storage  
15 operation are the same type of storage operation.
4. The method of claim 2 wherein the second storage operation is a partial storage operation.
5. The method of claim 1, wherein the first set of metadata and the second set of metadata are chunk headers.
- 20 6. The method of claim 1, further comprising:  
initiating a backup of the first set of stored data when the comparing of the  
first set of metadata and the second set of metadata indicates that the first set of  
metadata and the second set of metadata are not equivalent.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6

600

STORED DATA VERIFICATION TABLE

| 610<br>CHUNK ID | 620<br>CHUNK VERIFIED | 630<br>RESULT            |
|-----------------|-----------------------|--------------------------|
| 1               | YES/NO                | GOOD COPY                |
| 2               | YES/NO                | UNSUCCESSFUL             |
| 3               | YES/NO                | GOOD COPY                |
| 4               | YES/NO                | FILE 5 BAD               |
| 5               | YES/NO                | BAD INCREMENTAL<br>BLOCK |
| 6               | YES/NO                | GOOD COPY                |
| 7               | YES/NO                | GOOD COPY                |
| 8               | YES/NO                | GOOD COPY                |
| 9               | YES/NO                | FILE 10 BAD              |
| n               | YES/NO                | BAD INCREMENTAL<br>BLOCK |

FIG. 7




FIG. 8

# INTERNATIONAL SEARCH REPORT

International application No.

PCT/US04/38195

## A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 17/30; H02H 3/05  
US CL : 707/204, 104.1; 714/7, 18

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)  
U.S. : 707/204, 104.1

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category * | Citation of document, with indication, where appropriate, of the relevant passages                                                           | Relevant to claim No. |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| X          | US 6,502,205 B1 (YANAI et al.) 31 December 2002 (31.12.2002) Abstract column 2, lines 32-60, column 7, lines 45-67 and column 8, lines 1-50. | 1-4 and 6             |
| ---        |                                                                                                                                              | -----                 |
| Y          |                                                                                                                                              | 5                     |
| Y          | US 5,933,104 (KIMURA) 3 August 1999 (3.08.1999) column 5, lines 20-41 and see fig. 8                                                         | 5                     |
| A          | US 6,073,220 (GUNDERSON) 6 June 2000 (06.06.2000)                                                                                            | 1-6                   |
| A          | US 5,813,009 (JOHNSON et al.) 22 September 1998 (22.09.1998)                                                                                 | 1-6                   |

☐ Further documents are listed in the continuation of Box C.

☐ See patent family annex.

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

14 January 2005 (14.01.2005)

Date of mailing of the international search report

15 FEB 2005

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US  
Commissioner for Patents  
P.O. Box 1450  
Alexandria, Virginia 22313-1450

Facsimile No. (703) 305-3230

Authorized officer

JEAN CORRIELUS

Telephone No. 571 272-4032