



- (51) **International Patent Classification:**
G06F 17/30 (2006.01)
- (21) **International Application Number:**
PCT/CN20 13/070 180
- (22) **International Filing Date:**
8 January 2013 (08.01.2013)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
13/453,252 23 April 2012 (23.04.2012) US
- (71) **Applicant: INTERNATIONAL BUSINESS MA-
CHINES CORPORATION [US/US];** New Orchard
Road, Armonk, New York 10504 (US).
- (71) **Applicant (for MG only): IBM (CHINA) CO., LIMITED**
[CN/CN]; 7F, Bldg 10, Zhangjiang Innovation Park, 399
Keyuan Road, Zhangjiang High-tech Campus, Pudong
New Area, Shanghai 201203 (CN).

- (72) **Inventors: FISKE, Rahul M.;** Rajiv Gandhi Infotech Pk
Phase 2, Plot No Pl-3 Midc, Hinjewadi, Village Limit Of
Marunji Pune, Mh 4 11057 (IN). **JONES, Carl Evan;** 9000
S Rita Rd Tucson, Arizona 85744-0002 (US). **ROY, Sub-
hojit;** Embassy Tech Zone, Plot No. 3, Rajeev Gandhi In-
fotech Park, Phase 2, Hinjewadi, Pune, Mh 4 11057 (IN).
- (74) **Agent: CCPIT PATENT AND TRADEMARK LAW
OFFICE;** 8th Floor, Vantone New World Plaza, 2
Fuchengmenwai Street, Xicheng District, Beijing 100037
(CN).
- (81) **Designated States (unless otherwise indicated, for every
kind of national protection available):** AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,
BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM,
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP,
KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD,
ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI,
NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU,
RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ,

[Continued on nextpage]

(54) **Title:** PRESERVING REDUNDANCY IN DATA DEDUPLICATION SYSTEMS BY INDICATOR

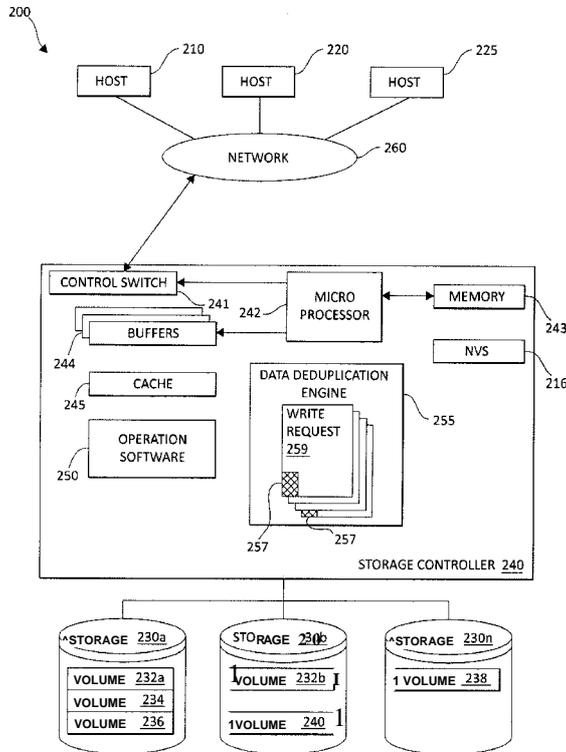


FIG. 2

(57) **Abstract:** Various embodiments for preserving data re-
dundancy in a data deduplication system in a computing en-
vironment are provided. In one embodiment, a method for
such preservation is disclosed. An indicator is configured.
The indicator is provided with a selected data segment to be
written through the data deduplication system to designate
that the selected data segment must not be subject to a dedu-
plication operation, such that repetitive data can be written
stored on physical locations despite being identical.

WO2013/159578 A1

TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA,
ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,

EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

PRESERVING REDUNDANCY IN DATA DEDUPLICATION SYSTEMS BY INDICATOR

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The present invention relates in general to computers, and more particularly to a method, system, and computer program product for preserving redundancy and other data security characteristics in computing environments in which data deduplication systems are incorporated.

Description of the Related Art

[0002] Computers and computer systems are found in a variety of settings in today's society. Computing environments and networks may be found at home, at work, at school, in government, and in other settings. Computing environments increasingly store data in one or more storage environments, which in many cases are remote from the local interface presented to a user.

[0003] These computing storage environments may use many storage devices such as disk drives, often working in concert, to store, retrieve, and update a large body of data, which may then be provided to a host computer requesting or sending the data. In some cases, a number of data storage subsystems are collectively managed as a single data storage system. These subsystems may be managed by host "sysplex" (system complex) configurations that combine several processing units or clusters of processing units. In this way, multi-tiered/multi-system computing environments, often including a variety of types of storage devices, may be used to organize and process large quantities of data.

SUMMARY OF THE INVENTION

[0004] Many multi-tiered/multi-system computing environments implement data deduplication technologies to improve storage performance by reducing the amount of duplicated storage across storage devices. Data deduplication systems are increasingly utilized because they help reduce the total amount of physical storage that is required to store data. This reduction is accomplished by ensuring that duplicate data is not stored multiple times. Instead, for example, if a chunk of incoming application WRITE data matches with an already stored chunk of data, a pointer to the original data is stored in the virtual storage map instead of allocating new physical storage space for the new chunk of data.

[0005] In certain situations, however, the behavior of deduplication i.e. single instancing of duplicate data, may go against the redundancy requirements of a hosted application, for example, or a storage policy, or other requirements. A need exists for a mechanism whereby data having the need to be stored multiple times is safeguarded, yet the benefits of deduplication systems are not diminished, by allowing deduplication to occur for remaining data not having such requirements.

[0006] In view of the foregoing, various embodiments for preserving data redundancy in data deduplication systems are disclosed. In one embodiment, by way of example only, a method for such preservation is disclosed. An indicator is configured. The indicator is provided with a selected data segment to be written through the data deduplication system to designate that the selected data segment must not be subject to a deduplication operation.

[0007] In addition to the foregoing exemplary embodiment, various system and computer program embodiments are provided and supply related advantages.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] In order that the advantages of the invention will be readily understood, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments that are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings, in which:

[0009] Fig. 1 is a block diagram illustrating an exemplary computing environment which may implement aspects of the present invention;

[0010] Fig. 2 is a block diagram showing a hardware structure of a data storage system, again in which aspects of the present invention may be implemented;

[0011] Fig. 3 is an exemplary write request, in which aspects of the present invention may be implemented;

[0012] Fig. 4 is a block diagram of an exemplary process of writing data through a data deduplication engine, again in which aspects of the present invention may be implemented;

[0013] Fig. 5 is an exemplary method for ensuring redundancy in data deduplication systems according to the present invention; and

[0014] Fig. 6 is an exemplary method for processing data in which aspects of the present invention may be implemented.

DETAILED DESCRIPTION OF THE DRAWINGS

[0015] Data deduplication in storage controllers typically works "behind the scene" of an application, and may sometimes operate contrary to the interests of the application when deduplication operations are performed against the needs of the application. This situation may arise if an application writes multiple copies of the same data, and intends to retain multiple physical copies, while the deduplication subsystem (deduplication engine) finds these matching copies and ends up deduplicating the copies while storing the data. This can be detrimental to the application, which expects to find multiple copies at various locations, and is made to believe that it has done so by the storage subsystem, but in reality only a single copy of the data has been written.

[0016] Consider the following example. File systems usually prefer to write multiple physical copies of the "Superblock," or a segment of metadata describing the file system on a block-based storage device, (or other metadata information) on a virtual disk to ensure redundancy. Since the contents of the Superblock are the same, data deduplication would result in retaining a single, physical copy of the Superblock and point multiple virtual addresses to the same physical block. This situation is highly inadvisable, because the loss of a single block on the physical copy of the Superblock may render the file system totally unusable, as there are no redundant copies of the Superblock. Conventionally, there are no existing methodologies that directly address this problem in data deduplication systems.

[0017] Various indirect methodologies may be employed to attempt to address this problem. In one example, the storage pool from which the data deduplication subsystem carves out physical storage can be mirrored (i.e., contains 2 or 3 copies of the same data). Hence multiple redundant copies can be created despite deduplication. However, this is inadequate protection for the application because of the following reasons. First, the application may wish to keep, for example, ten (10) copies of the same data. However, if

the storage pool is two-way mirrored, it may only retain a maximum of two (2) copies. Second, since data deduplication carves out physical storage pools that span across large amounts of storage and multiple file systems, it is likely that multiple applications and file systems share the same physical storage pool. Hence it is possible that some critical copies of data (like the Superblock) get physically placed on the same disk. Since deduplication would prevent multiple copies of the same data to be written to multiple physical locations, the number of copies of critical data reduces and they can get placed on the same physical disk for multiple file systems. This increases the risk of single failures becoming fatal.

[0018] The illustrated embodiments provide multiple mechanisms for addressing the issues discussed previously. One goal of these mechanisms is to ensure that the deduplication subsystem in the storage controller (or wherever it may be located) balances the benefits of reducing the number of copies of data against application requirements for physical allocating multiple copies of critical data. Each of the methodologies described in the following illustrated embodiments may be used in a variety of circumstances and may have attendant benefits specific to those circumstances.

[0019] In one such embodiment, an indicator may be provided by the application for a selected data segment to the target storage controller to designate that the particular data segment is not subject to deduplication operations. This action then forces the storage controller to allocate new, physical storage for the specified data blocks. One such indicator may include a bit, which is set high or low depending on whether the desired deduplication should be performed for the selected data segment / block.

[0020] In view of the described embodiment, by allowing the application to dictate whether a write must be deduplicated, the application is allowed flexibility to implement storage policy associated with the data it generates. This way, the application is in a better position than the deduplication system to determine whether selected data blocks, even

though identical, must still be located in separate physical locations. In addition, the storage controller (or other storage management device) continues to perform its role of data reduction by deduplication, and at the same time allowed enough control to the application to rule out deduplication when required.

[0021] By providing write commands with an indicator to indicate to the storage controller (or again, other storage management devices) whether the selected data must skip deduplication, very fine-grained control is thereby provided to the application, allowing for flexibility in implementation while still retaining advantages of deduplication functionality and retaining redundancy for key data.

[0022] In the following description, reference is made to the accompanying drawings which form a part hereof and which illustrate several embodiments of the present invention. It is understood that other embodiments may be utilized and structural and operational changes may be made without departing from the scope of the present invention.

[0023] Turning to Fig. 1, an example computer system 10 is depicted in which aspects of the present invention may be realized. Computer system 10 includes central processing unit (CPU) 12, which is connected to mass storage device(s) 14 and memory device 16. Mass storage devices may include hard disk drive (HDD) devices, which may be configured in a redundant array of independent disks (RAID). The data management operations, in which aspects of the present invention may be implemented as further described, may be executed on device(s) 14, located in system 10 or elsewhere. Memory device 16 may include such memory as electrically erasable programmable read only memory (EEPROM) or a host of related devices.

[0024] Memory 16 is shown including an application program 15, and an application program 17, in which a file system 19 is operational. Application 15 and application 17

may create, delete, or otherwise manage segments of data, such as data chunks or data blocks, which are physically stored in devices such as mass storage device 14. File system 19 provides a means to organize data expected to be retained after the application program 17 terminates by providing procedures to store, retrieve, and update data, as well as manage the available space on the device(s) that contain it. The file system 19 organizes data in an efficient manner, and is tuned to the specific characteristics of the device (such as computer 10 and/or memory 16). In one embodiment, application 17 may be an operating system (OS) 17, and file system 19 retains a tight coupling between the OS 17 and the file system 19. File system 19 may provide mechanisms to control access to the data and metadata, and may contain mechanisms to ensure data reliability such as those necessary to further certain aspects of the present invention, as one of ordinary skill in the art will appreciate. File system 19 may provide a means for multiple application programs 15, 17 to update data in the same file at nearly the same time.

[0025] In the illustrated embodiment, memory device 16 and mass storage device 14 are connected to CPU 12 via a signal-bearing medium. In addition, CPU 12 is connected through communication port 18 to a communication network 20, having an attached plurality of additional computer systems 22 and 24. The computer system 10 may include one or more processor devices (e.g., CPU 12) and additional memory devices 16 for each individual component of the computer system 10 to execute and perform each operation described herein to accomplish the purposes of the present invention.

[0026] Fig. 2 is an exemplary block diagram 200 showing a hardware structure of a data storage system in a computer system according to the present invention. Host computers 210, 220, 225, are shown, each acting as a central processing unit for performing data processing as part of a data storage system 200. The hosts (physical or virtual devices), 210, 220, and 225 may be one or more new physical devices or logical devices to accomplish the purposes of the present invention in the data storage system 200.

In one embodiment, by way of example only, a data storage system 200 may be implemented as IBM® System Storage™ DS8000™. A Network connection 260 may be a fibre channel fabric, a fibre channel point to point link, a fibre channel over ethernet fabric or point to point link, a FICON or ESCON I/O interface, any other I/O interface type, a wireless network, a wired network, a LAN, a WAN, heterogeneous, homogeneous, public (i.e. the Internet), private, or any combination thereof. The hosts, 210, 220, and 225 may be local or distributed among one or more locations and may be equipped with any type of fabric (or fabric channel) (not shown in Figure 2) or network adapter 260 to the storage controller 240, such as Fibre channel, FICON, ESCON, Ethernet, fiber optic, wireless, or coaxial adapters. Data storage system 200 is accordingly equipped with a suitable fabric (not shown in Fig.2) or network adapter 260 to communicate. Data storage system 200 is depicted in Fig.2 comprising storage controller 240 and storage 230.

[0027] To facilitate a clearer understanding of the methods described herein, storage controller 240 is shown in Fig. 2 as a single processing unit, including a microprocessor 242, system memory 243 and nonvolatile storage ("NVS") 216, which will be described in more detail below. It is noted that in some embodiments, storage controller 240 is comprised of multiple processing units, each with their own processor complex and system memory, and interconnected by a dedicated network within data storage system 200. Storage 230 may be comprised of one or more storage devices, such as storage arrays, which are connected to storage controller 240 by a storage network.

[0028] In some embodiments, the devices included in storage 230 may be connected in a loop architecture. Storage controller 240 manages storage 230 and facilitates the processing of write and read requests intended for storage 230. The system memory 243 of storage controller 240 stores program instructions and data that the processor 242 may access for executing functions associated with managing storage 230. In one embodiment, system memory 243 includes, is associated, or is in communication with the operation

software 250, and configured in part for accomplishing functionality of the present invention. As shown in Figure 2, system memory 243 may also include or be in communication with a cache 245 for storage 230, also referred to herein as a "cache memory", for buffering "write data" and "read data", which respectively refer to write/read requests and their associated data. In one embodiment, cache 245 is allocated in a device external to system memory 243, yet remains accessible by microprocessor 242 and may serve to provide additional security against data loss, in addition to carrying out the operations as described in herein.

[0029] In some embodiments, cache 245 is implemented with a volatile memory and non-volatile memory and coupled to microprocessor 242 via a local bus (not shown in Figure 2) for enhanced performance of data storage system 200. The NVS 216 included in data storage controller is accessible by microprocessor 242 and serves to provide additional support for operations and execution of the present invention as described in other figures. The NVS 216, may also referred to as a "persistent" cache, or "cache memory" and is implemented with nonvolatile memory that may or may not utilize external power to retain data stored therein. The NVS may be stored in and with the cache 245 for any purposes suited to accomplish the objectives of the present invention. In some embodiments, a backup power source (not shown in Figure 2), such as a battery, supplies NVS 216 with sufficient power to retain the data stored therein in case of power loss to data storage system 200. In certain embodiments, the capacity of NVS 216 is less than or equal to the total capacity of cache 245.

[0030] Storage 230 may be physically comprised of one or more storage devices, such as storage arrays. A storage array is a logical grouping of individual storage devices, such as a hard disk. In certain embodiments, storage 230 is comprised of a JBOD (Just a Bunch of Disks) array or a RAID (Redundant Array of Independent Disks) array. A collection of physical storage arrays may be further combined to form a rank, which

dissociates the physical storage from the logical configuration. The storage space in a rank may be allocated into logical volumes, which define the storage location specified in a write/read request.

[0031] In one embodiment, by way of example only, the storage system as shown in Figure 2 may include a logical volume, or simply "volume," may have different kinds of allocations. Storage 230a, 230b and 230n are shown as ranks in data storage system 200, and are referred to herein as rank 230a, 230b and 230n. Ranks may be local to data storage system 200, or may be located at a physically remote location. In other words, a local storage controller may connect with a remote storage controller and manage storage at the remote location. Rank 230a is shown configured with two entire volumes, 234 and 236, as well as one partial volume 232a. Rank 230b is shown with another partial volume 232b. Thus volume 232 is allocated across ranks 230a and 230b. Rank 230n is shown as being fully allocated to volume 238—that is, rank 230n refers to the entire physical storage for volume 238. From the above examples, it will be appreciated that a rank may be configured to include one or more partial and/or entire volumes. Volumes and ranks may further be divided into so-called "tracks," which represent a fixed block of storage. A track is therefore associated with a given volume and may be given a given rank.

[0032] The storage controller 240 includes a control switch 241 for controlling the fiber channel protocol to the host computers 210, 220, 225, a microprocessor 242 for controlling all the storage controller 240, a nonvolatile control memory 243 for storing a microprogram (operation software) 250 for controlling the operation of storage controller 240, data for control and each table described later, cache 245 for temporarily storing (buffering) data, and buffers 244 for assisting the cache 245 to read and write data, a control switch 241 for controlling a protocol to control data transfer to or from the storage devices 230, and compression operation module 255 and compression operation list module 257 in which information may be set. Multiple buffers 244 may be implemented

with the present invention to assist with the operations as described herein.

[0033] In one embodiment, the host computers or one or more physical or virtual devices, 210, 220, 225 and the storage controller 240 are connected through a network adaptor (this could be a fibre channel)260 as an interface i.e., via at least one switch called "fabric." In one embodiment, the operation of the system shown in Fig. 2 will be described. The microprocessor 242 may control the memory 243 to store command information from the host device (physical or virtual)210 and information for identifying the host device (physical or virtual)210. The control switch 241, the buffers 244, the cache 245, the operating software 250, the microprocessor 242, memory 243, NVS 216, compression operation module 255 and compression operation list module 257 are in communication with each other and may be separate or one individual component(s). Also, several, if not all of the components, such as the operation software 250 may be included with the memory 243. Each of the components within the devices shown may be linked together and may be in communication with each other for purposes suited to the present invention.

[0034] The storage controller 240 is shown including a data deduplication engine 255, in which a number of write requests 259 are processed. The data deduplication engine 255 may be structurally one complete module or may be associated and/or incorporated within other individual modules. Data deduplication engine 255 is configured for performing, in conjunction with other components of storage controller 240 such as microprocessor 242, data deduplication operations on write data passed through storage controller 240 to storage 230.

[0035] As previously indicated, storage controller 240 includes cache 245 (or what may be termed holistically as cache system) 245 accepts write data from hosts 210, 220, and 225, or similar devices, that is then placed in cache memory 245. Data deduplication

engine 255 then tests the write data for duplication in the cache memory 245. The write requests 259 that the application 17 (Fig. 1) desires to protect from deduplication operations are designated with indicators 257 as shown. Please note that the indicator is "set" in the write command, and not in the data to be written (i.e., the write command payload). Those of the write requests 259 having indicator 257 that are passed through the data deduplication engine 255 designate to the data deduplication engine 255 that they are not subject to a deduplication operation.

[0036] In one embodiment, the application 17 / file system 19 first determines whether a given data block must be stored multiple times on physical storage. This may be Superblock metadata associated with the file system 19 itself. The application 17 then may initiate a write request 259 to the underlying disk (device) driver, with the indicator 257 set (such as a high bit), designating that this data block must not be deduplicated, even if the block is determined to be identical to a previously written data block.

[0037] In a following step, the disk driver in turn prepares a command, such as a Small Computer Systems Interface (SCSI) command with the set bit to indicate "skip deduplication" to the storage controller 240. On receiving the designated data block, the storage controller 240, via data deduplication engine 255, skips the deduplication steps of fingerprint generation, matching, etc., as one of ordinary skill in the art will appreciate, and directly writes the selected data block onto a new physical disk location, even if there may have been earlier instances of identical data being stored on the storage controller 240.

[0038] In one embodiment, for applications executing at the user level, the writesystem call may be provided with an extended attribute indicating that a data block must be written to physical disk even if its contents match data written to disk earlier. If this is the case, additional steps are performed in similar fashion to that previously described depending on whether the write was to a block or a file interface, for example.

[0039] Turning now to Fig. 3, an embodiment of an input/output request, specifically a time-based write command 300 that may be sent from the host 210 (Fig. 2) to the storage controller 240. In one embodiment, the write request 300 is a SCSI command. Fig. 3 identifies each bit included in the 32 bytes of a command block. For example, in byte 0, an operation code identifies the type of command to be performed, i.e., a time-based write command. Bytes 2-9 are for the Logical Block Address (LBA) that identifies a first unit of data. Bytes 10-13 are for the transfer length, which indicates the number of blocks that are being written beginning with the first unit of data, and so on. One of ordinary skill in the art will appreciate that the various fields and/or various applicability of these fields in the write request 300 may vary according to a particular implementation.

[0040] In the illustrated embodiment shown in Fig. 3, the indicator 257 is shown making up bit 7 of byte 1 of the write command 300. In this manner, the indicator 257 may be a low (i.e., 0) or high (i.e., 1) bit to indicate that the designated data to be written is not subject to a data deduplication operation performed by data deduplication engine 255 (again, Fig. 2). Here again, those of ordinary skill in the art will appreciate that other similar indicators may be applied to the write request, or provided in conjunction with the write request and/or write data, to designate particular portions of data to be exempt from deduplication.

[0041] Turning now to Fig. 4, a block diagram of various functional aspects of the present invention are depicted as an exemplary flow. Specifically, cache system 245 is shown to accept write data 402 to be processed through data deduplication engine 255 as a write operation 404 to cache 245. As the write data 402 is processed through the data deduplication engine 255, and as previously described, the write data 402 is normally tested for duplication with existing storage. However, since the write request 259 contains indicator 257, no deduplication operation is performed, and the write data passes through

the deduplication engine 255 through the cache system 245 and ultimately to physically allocated storage 230.

[0042] In an alternative embodiment, which is not shown for purposes of illustrative convenience, the cache system 245 may be placed in front of the deduplication engine 255. In such cases, if the indicator 257 is set in the write command 259, the corresponding data segment is written through the cache 245. Alternatively, the cache 245 stores the bit indicator 257 such that when the segment is later flushed, the indicator 257 for a given set of segments is available to the deduplication engine 255. One of ordinary skill in the art will appreciate that other modifications to the functional aspects depicted in Fig. 4 such as operational flow and placement as previously indicated may be performed to suit a particular application.

[0043] Fig. 5, following, is an exemplary method 500 for ensuring data redundancy in storage subsystems having data deduplication systems, in which aspects of the illustrated embodiments may be implemented. Method 500 begins (step 502) with the configuration of an indicator to be provided with a selected data segment (i.e., such as incorporated into a SCSI write request) to be written through the data deduplication system to designate that the selected data segment must not be subject to a deduplication operation (step 504). The method 500 then ends (step 506).

[0044] In Fig. 6, following, an exemplary method 600 for data processing is illustrated, here again in which various aspects of the present invention may be implemented. Method 600 begins (step 602), with the configuration of the indicator as previously described (step 604). In the illustrated embodiment, the indicator is provided in the context of a write request as a set bit in an SCSI write command (step 606). The command is then processed (step 608) by the storage controller. If the SCSI command includes a set bit for particular selected data (step 610), then a deduplication operation is withheld from being performed

on the associated data segment (step 612), and correspondingly, the associated data segment is directly written to a newly allocated physical storage location (step 614).

[0045] Returning to step 610, if the SCSI command does not include the set bit for the particular selected data segment, then the associated data segment is processed through various deduplication algorithms in the deduplication engine (such as being checked for deduplication in other versions of the data segment that were previously stored, and performing other deduplication functionality that would be apparent to those of ordinary skill in the art. The method 600 then ends (step 618).

[0046] The mechanisms of the illustrated embodiments may be applicable to write requests themselves as previously described, or, in other embodiments, be implemented in the context of the selected data itself. For example, in one embodiment, the Superblock metadata itself may be flagged with the indicator to designate the information as not subject to data deduplication operations. Other techniques for providing an indication in conjunction with selected data may be apparent to those of ordinary skill in the art as appropriate for a specific application.

[0047] As will be appreciated by one of ordinary skill in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may generally be referred to herein as a "circuit," "module," "process" or "system." Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

[0048] Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

[0049] Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wired, optical fiber cable, RF, etc., or any suitable combination of the foregoing. Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, or entirely on the remote computer or server. In the last scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

[0050] Aspects of the present invention are described below with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0051] These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks. The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0052] The flowchart and block diagrams in the above figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods

and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

[0053] While one or more embodiments of the present invention have been illustrated in detail, one of ordinary skill in the art will appreciate that modifications and adaptations to those embodiments may be made without departing from the scope of the present invention as set forth in the following claims.

Claims

What is claimed is:

1. A method for preserving data redundancy in a data deduplication system in a computing environment by a processor, comprising:
configuring an indicator provided with a selected data segment to be written through the data deduplication system to designate that the selected data segment must not be subject to a deduplication operation.
2. The method of claim 1, further including processing the indicator such that the deduplication operation is withheld from being performed on the selected data segment.
3. The method of claim 2, wherein processing the indicator includes preparing a command, incorporating the indicator, to be provided to a storage controller in the computing environment.
4. The method of claim 3, wherein processing the indicator includes:
receiving the command by the storage controller,
bypassing at least one deduplication operation, and
directly writing the selected data segment to a new physical storage location.
5. The method of claim 1, wherein configuring an indicator further includes, for an application operational in the computing environment, providing a system call with an extended attribute indicating that the selected data segment must be written to physical memory notwithstanding the selected data segment matching data written to the physical memory earlier in time.
6. The method of claim 1, wherein the selected data segment is at least one of a

Superblock and file system metadata required to be stored repetitively on physical storage, and configuring an indicator includes configuring at least one of a bit and a flag..

7. A system for preserving data redundancy in a data deduplication system in a computing environment, comprising:

a file system module; and

a storage controller in operable communication with the file system module,

wherein:

the file system module is adapted for configuring an indicator provided with a selected data segment to be written through the data deduplication system to designate that the selected data segment must not be subject to a deduplication operation.

8. The system of claim 7, wherein the storage controller is adapted for processing the indicator such that the deduplication operation is withheld from being performed on the selected data segment.

9. The system of claim 8, further including a storage driver in operable communication between the file system module and the storage controller, wherein the storage driver is further adapted for, pursuant to processing the indicator, preparing a command, incorporating the indicator, to be provided to the storage controller.

10. The system of claim 9, wherein the storage controller is further adapted for, pursuant to processing the indicator:

receiving the command,

bypassing at least one deduplication operation, and

directly writing the selected data segment to a new physical storage location.

11. The system of claim 7, further including an application, operationally

communicative with the file system module, wherein the application is adapted for providing a system call with an extended attribute indicating that the selected data segment must be written to physical memory notwithstanding the selected data segment matching data written to the physical memory earlier in time.

12. The system of claim 7, wherein the selected data segment is at least one of a Superblock and file system metadata required to be stored repetitively on physical storage.

13. The system of claim 7, wherein the indicator includes at least one of a bit and a flag.

14. The system of claim 13, wherein the at least one of the bit and the flag are associated with the selected data segment using an adapted small computer system interface (SCSI) write command.

15. A computer program product for preserving data redundancy in a data deduplication system in a computing environment by a processor, the computer program product comprising a computer-readable storage medium having computer-readable program code portions stored therein, the computer-readable program code portions comprising:

a first executable portion for configuring an indicator provided with a selected data segment to be written through the data deduplication system to designate that the selected data segment must not be subject to a deduplication operation.

16. The computer program product of claim 15, further including a second executable portion for processing the indicator such that the deduplication operation is withheld from being performed on the selected data segment.

17. The computer program product of claim 16, further including a third executable

portion for, pursuant to processing the indicator, preparing a command, incorporating the indicator, to be provided to a storage controller in the computing environment.

18. The computer program product of claim 17, further including a fourth executable portion for, pursuant to processing the indicator:

receiving the command by the storage controller,
bypassing at least one deduplication operation, and
directly writing the selected data segment to a new physical storage location.

19. The computer program product of claim 15, further including a second executable portion for, pursuant to configuring an indicator, for an application operational in the computing environment, providing a system call with an extended attribute indicating that the selected data segment must be written to physical memory notwithstanding the selected data segment matching data written to the physical memory earlier in time.

20. The computer program product of claim 15, wherein the selected data segment is at least one of a Superblock and file system metadata required to be stored repetitively in physical storage, and further including a second executable portion for, pursuant to configuring an indicator, configuring at least one of a bit and a flag.

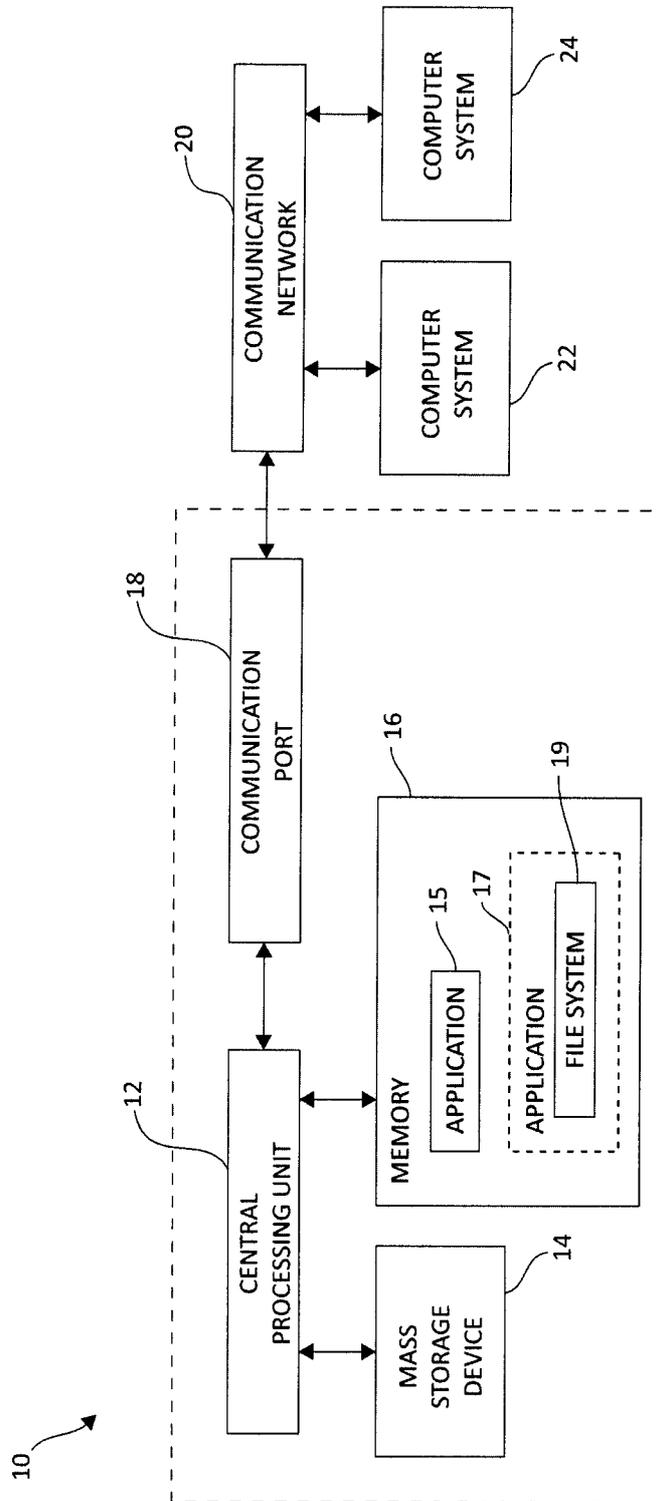


FIG. 1

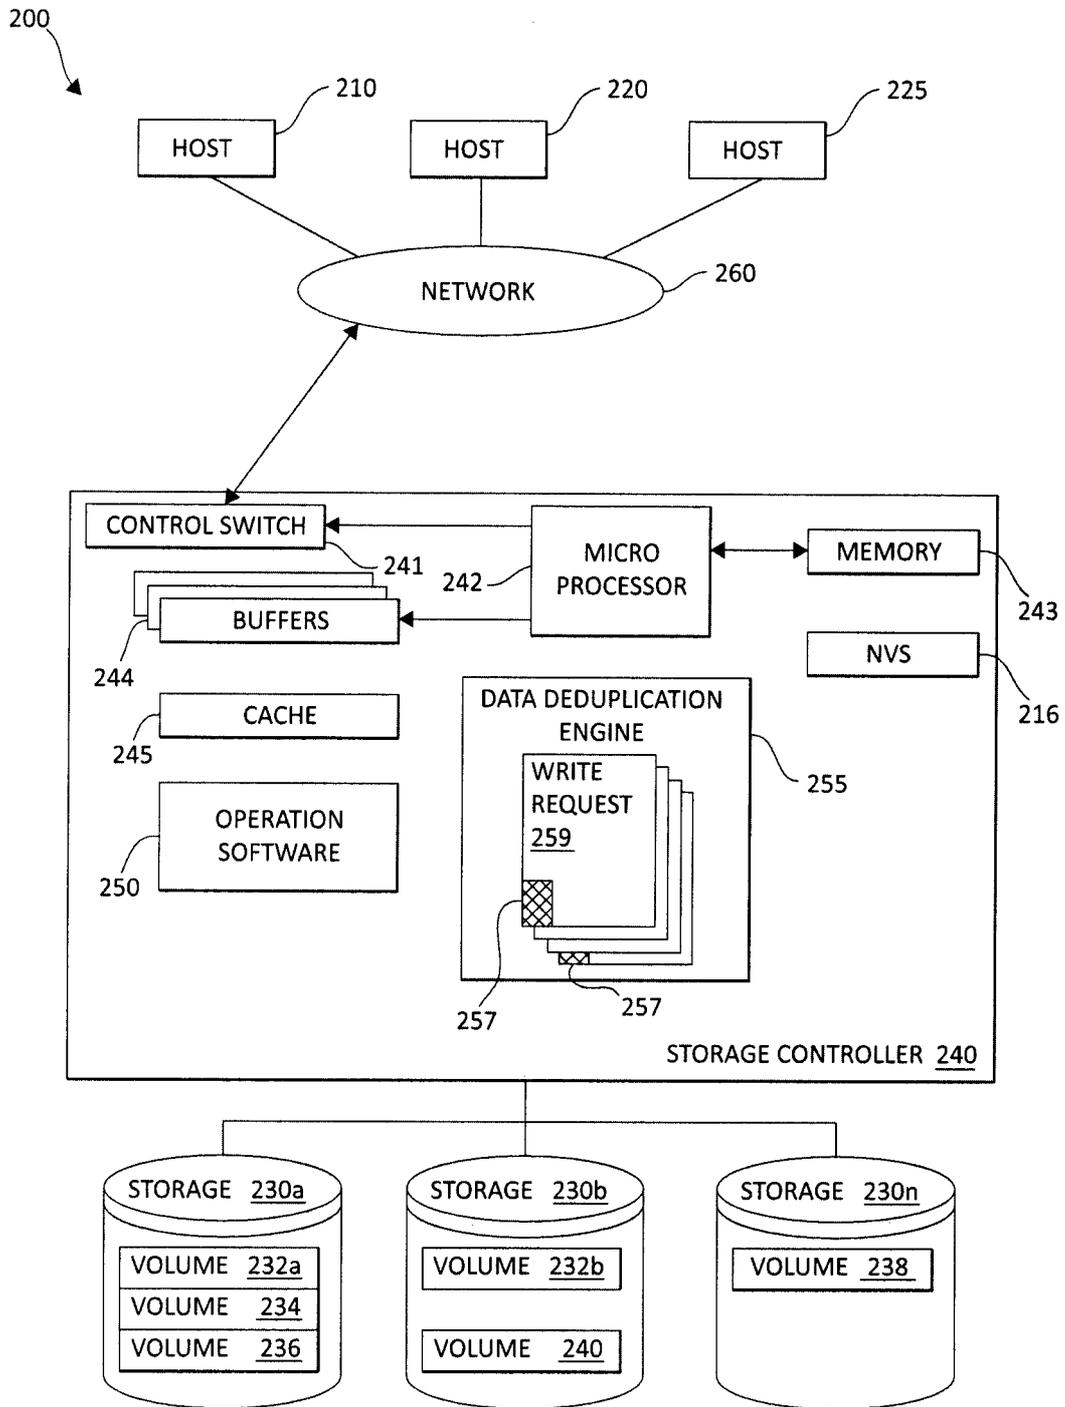


FIG. 2

300

Byte/Bit	7	6	5	4	3	2	1	0
0	Operation Code (COh)							
1	Reserved	DPO	Reserved	EBP	ByteChk	RelAddr		
2	Logical Block Address							
9								
10	Transfer Length							
13								
14	Reserved							
15								
16	RelativeChk							
17	Time							
24								
25	Reserved							
30								
31	Control							

257

FIG. 3

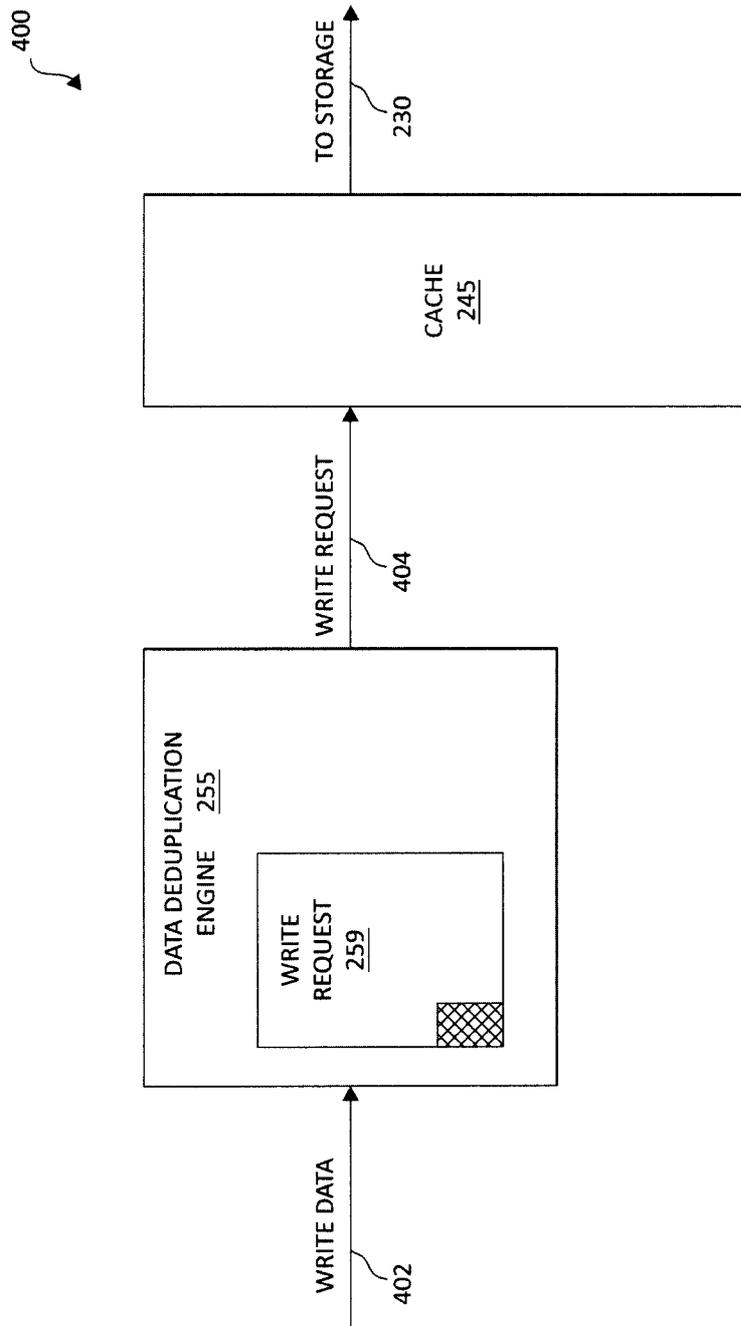


FIG. 4

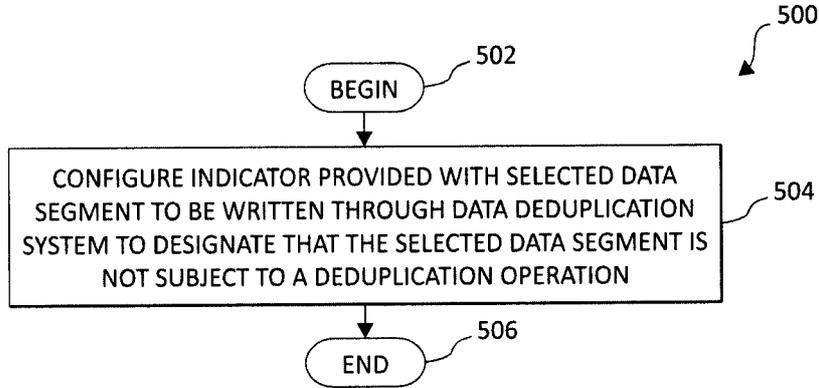


FIG. 5

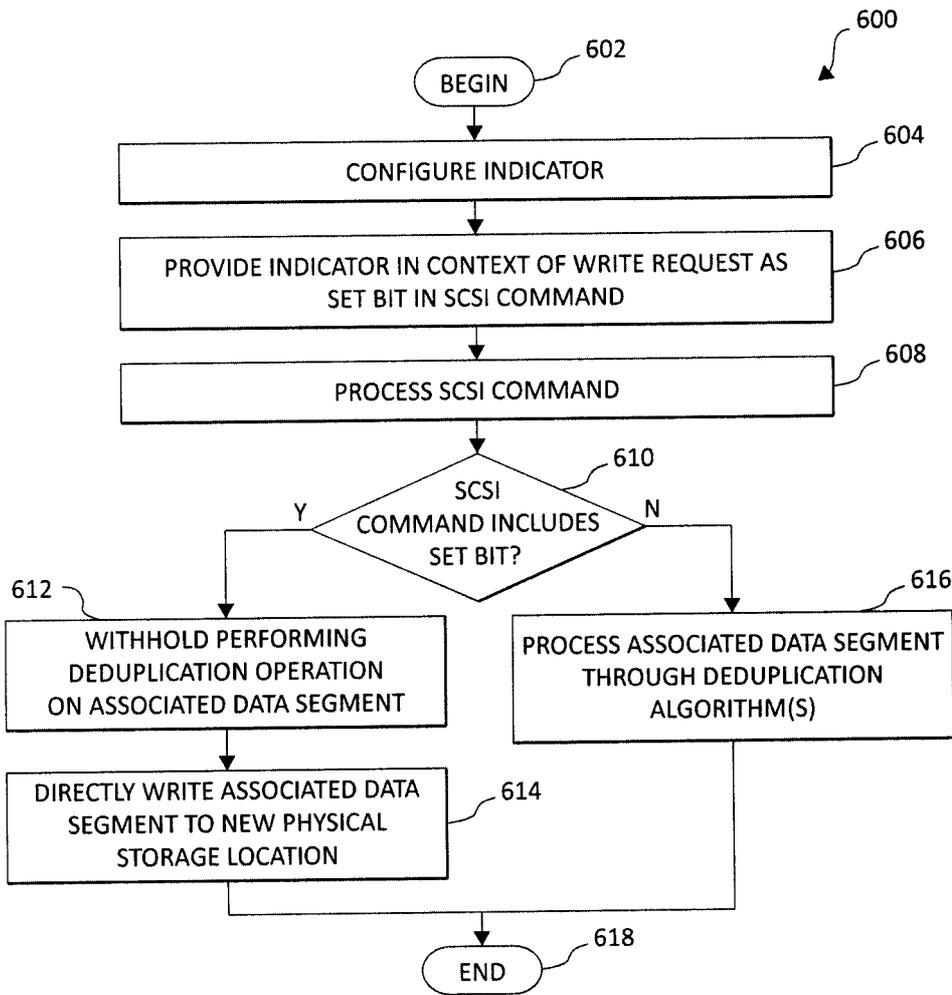


FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/CN2013/070180

A. CLASSIFICATION OF SUBJECT MATTER

G06F 17/30 (2006.01) i

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)

IPC: H04L, G06F

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)

DWPI, CNABS, USTXT, WOTXT, EPTXT, CATXT, CNTXT, CNKI, GOOLGE: deduplicat+, de-duplicat+, indicator+, delet+, duplica+, indicator, not

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category ^{1*}	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012/0089574 A1 (QUANTUM CORPORATION) 12 Apr. 2012 (12.04.2012) paragraphs 30-33 in the description, figures 1-2	1-20
A	US 2011/0225214 A1 (SYMANTEC CORPORATION) 15 Sep. 2011 (15.09.2011) the whole document	1-20
A	US 2010/0299311 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 25 Nov. 2010 (25.11.2010) the whole document	1-20

1-1 Further documents are listed in the continuation of Box C. See patent family annex.

<p>* Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"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)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p>	<p>"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</p> <p>"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</p> <p>"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</p> <p>"&" document member of the same patent family</p>
--	---

Date of the actual completion of the international search 08 Mar. 2013 (08.03.2013)	Date of mailing of the international search report 28 Mar. 2013 (28.03.2013)
--	--

Name and mailing address of the ISA/CN
The State Intellectual Property Office, the P.R.China
6 Xitucheng Rd., Jimen Bridge, Haidian District, Beijing, China
100088
facsimile No. 86-10-62019451

Authorized officer
LI, Ping
Telephone No. (86-10) 624 13326

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/CN2013/070180

Patent Documents referred in the Report	Publication Date	Patent Family	Publication Date
US 2012/0089574A1	12.04.2012	NONE	
US 20 11/022521 4 A1	15.09.201 1	WO 201 1112230 A1 US 8108447 B2 AU 2010347733 A1 EP 2545435 A1 CN 10291 8487 A	15.09.201 1 31.01 .2012 04. 10.2012 16.01 .201 3 06.02.201 3
US20 10/02993 11 A 1	25. 11.2010	US 8280859 B2 US 2009234892 A1 US 7814074 B2	02. 10.2012 17.09.2009 12. 10.2010