



(51) International Patent Classification:
G06F 11/30 (2006.01)

(21) International Application Number:
PCT/US2015/054200

(22) International Filing Date:
6 October 2015 (06.10.2015)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
2966/CHE/2015 12 June 2015 (12.06.2015) IN

(71) Applicant: **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).

(72) Inventors: **CHIRUMAMILLA, Narendra**; Sy.No.192, Whitefield Road, Mahadevapura Post, Bangalore 560048 (IN). **MAHESH, Keshetti**; Sy.No.192, Whitefield Road,

Mahadevapura Post, Bangalore 560048 (IN). **NAYAKA B, Govindaraja**; Sy.No.192, Whitefield Road, Mahadevapura Post, Bangalore 560048 (IN). **BASIREDDY, Ranjith Reddy**; Sy.No.192, Whitefield Road, Mahadevapura Post, Bangalore 560048 (IN). **MOHANTA, Taranisen**; Sy.No.192, Whitefield Road, Mahadevapura Post, Bangalore 560048 (IN).

(74) Agents: **ORTEGA, Arthur S.** et al.; Hewlett-Packard Enterprise, 3404 E. Harmony Road, Mail Stop 79, Fort Collins, Colorado 80528 (US).

(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, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, 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, SA, SC,

[Continued on next page]

(54) Title: DATA REPLICATION




Fig. 1

(57) Abstract: Some examples described herein relate to data replication. In an example, state of storage volumes belonging to data replication groups in a first storage device may be monitored. A data replication priority for each data replication group may be determined, in an example, a data replication priority for a data replication group may be determined based on state of a storage volume belonging to the data replication group. Data replication may be performed, from the first storage device to a second storage device, for data replication groups according to respective data replication priorities, beginning with the data replication group having highest data replication priority.



SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, 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, ST, 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, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

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

DATA REPLICATION

Background

[001] Increased adoption of technology by businesses has led to an explosion of data. Organizations may be required to store data for various reasons. These may include business reasons, legal and compliance requirements, auditing functions, investigative purposes, etc. Thus, considering the importance of data for an enterprise, organizations may safeguard their data by storing same data on multiple storage devices.

Brief Description of the Drawings

[002] For a better understanding of the solution, embodiments will now be described, purely by way of example, with reference to the accompanying drawings, in which:

[003] FIG. 1 is a block diagram of an example computing environment for data replication;

[004] FIG. 2 is a block diagram of an example computer system for data replication;

[005] FIG. 3 is a flowchart of an example method of data replication; and

[006] FIG. 4 is a block diagram of an example system for data replication.

Detailed Description

[007] Data replication may include copying and maintaining data on multiple storage devices, which in an instance may be present at multiple locations. Since data changes on a storage device may be applied on each of the other storage devices,

data replication provides users with an alternate data access option even if one storage device becomes unavailable. In other words, data replication may ensure high availability of data storage volumes.

[008] In an example, data replication may be performed via data replication groups (DR groups) wherein DR groups on a primary storage device may be replicated to another storage device without any priority. This may not be an ideal mechanism in many scenarios. For instance, such process may not expedite replication of a DR group whose volume(s) may be in a degraded or failure predicted state. In the event, if a DR group whose volume(s) is/are in a degraded state is not assigned a priority during data replication, there is a possibility that the DR group may enter a failed state before data replication for the DR group is complete. This may lead to an incomplete data replication.

[009] To prevent these issues, the present disclosure describes various examples for data replication. In an example, state of storage volumes belonging to data replication groups in a first storage device may be monitored. A data replication priority for each data replication group may be determined. In an instance, a data replication priority for a data replication group may be determined based on state of a storage volume belonging to the data replication group. Next, data replication may be performed, from the first storage device to a second storage device, for each data replication group according to respective data replication priorities, beginning with the data replication group having highest data replication priority.

[0010] FIG. 1 is a block diagram of an example computing environment 100 for data replication. Computing environment 100 may include storage devices 102 and 104, and a data replication system 106. Although only two storage devices 102 and 104 are shown in FIG. 1, other examples of this disclosure may include more than two storage devices.

[0011] Components (i.e. 102, 104, and 106) of the computing environment 100 may be in communication with each other, for example, via a computer network. Computer network may be a wireless or wired network. Computer network may include, for example, a Local Area Network (LAN), a Wireless Local Area Network (WAN), a Metropolitan Area Network (MAN), a Storage Area Network (SAN), a Campus Area Network (CAN), or the like. Further, computer network may be a public network (for example, the Internet) or a private network (for example, an intranet). In an example, two or more of the components (i.e. 102, 104, and 106) of the computing environment 100 may be directly coupled to each other.

[0012] Storage devices 102 and 104 may each be, for example, a storage array, which may include one or more storage drives (for example, hard disk drives, solid state drives, etc.), a storage tape, a solid state drive, a USB drive, a Serial Advanced Technology Attachment (SATA) disk drive, a Fibre Channel (FC) disk drive, a Serial Attached SCSI (SAS) disk drive, a magnetic tape drive, an optical jukebox, and the like. In another example, storage devices 102 and 104 may each be a floppy disk, a hard disk, a CD-ROM, a DVD, a pen drive, a flash memory (e.g. USB flash drives or keys), a paper tape, an Iomega Zip drive, and the like. In another example, storage devices 102 and 104 may each be a Direct Attached Storage (DAS) device, a Network Attached Storage (NAS) device, a Redundant Array of Inexpensive Disks (RAID), a data archival storage system, or a combination of these devices.

[0013] In an example, physical storage space provided by storage device 102 and/or 104 may each be presented as a logical storage space to client computing devices (not shown). Such logical storage space (also referred as "logical volume", "virtual disk", or "storage volume") may be identified using a "Logical Unit Number" (LUN). In another instance, physical storage space provided by storage device 102 and/or 104 may each be presented as multiple logical volumes to a computing device. In such case, each of the logical storage spaces may be referred to by a separate LUN. Thus, if storage device is physical disk, a LUN may refer to the entire physical disk, or a

subset of the physical disk or disk volume. In another example, if a storage device is a storage array comprising multiple storage disk drives, physical storage space provided by the disk drives may be aggregated as a logical storage space. The aggregated logical storage space may be divided into multiple logical storage volumes, wherein each logical storage volume may be referred to by a separate LUN. LUNs, thus, may be used to identify individual or collections of physical disk devices for address by a protocol associated with, for example, a SCSI, iSCSI, or Fibre Channel (FC).

[0014] In an example, storage device 102 may act as primary storage device or source storage device, and storage device 104 may act as secondary storage device or target storage device. In another example, it may be vice versa. In a yet another example, in the event there are more than two storage devices in the computing environment, any of the storage device may act as primary storage device, and other storage devices may act as secondary storage devices. In an instance, data may be replicated from a primary storage device (or source storage device) to a secondary storage device (or target storage device).

[0015] In an example, storage volumes to be replicated, from a primary storage device to a secondary storage device, may be organized into one or more data replication (DR) groups. In an instance, the classification may be based on a host and/or an application's requirements. Thus, a DR group may include one or more storage volumes, and a storage device may include one or more data replication groups. The DR groups may be defined by a user (for example, a storage administrator) or system. In an instance, for every data replication group on a storage device, a log disk may be created that may be used to log I/O requests received on storage volumes belonging to a data replication group.

[0016] In an example, computing environment 100 may include one or more client computing devices (not shown). The client computing devices may each represent

any type of computing system capable of reading machine-executable instructions. Examples of client computing devices may include, without limitation, a server, a desktop computer, a notebook computer, a tablet computer, a thin client, a mobile device, a personal digital assistant (PDA), a phablet, and the like. In an example, client computing devices may store, modify, process, search, and/or delete data on storage device 102 and/or 104.

[0017] Data replication system 106 may include a computing device or a storage device.

In an example, data replication system 106 may include a monitoring module 110, a data replication priority module 112, and a data replication module 114. The term “module” may refer to a software component (machine readable instructions), a hardware component or a combination thereof. A module may include, by way of example, components, such as software components, processes, tasks, co-routines, functions, attributes, procedures, drivers, firmware, data, databases, data structures, Application Specific Integrated Circuits (ASIC) and other computing devices.

[0018] Some of the example functionalities that may be performed by monitoring module 110, data replication priority module 112, and data replication module 114 are described in reference to FIG. 2 below.

[0019] FIG. 2 is a block diagram of an example computer system 200 for data replication. In an example, computer system 200 may be analogous to data replication system 106 of FIG. 1, in which like reference numerals correspond to the same or similar, though perhaps not identical, components. For the sake of brevity, components or reference numerals of FIG. 2 having a same or similarly described function in FIG. 1 are not being described in connection with FIG. 2. Said components or reference numerals may be considered alike.

[0020] Computer system 200 may generally represent any type of computing device capable of reading machine-executable instructions. Examples of computer system 200 may include, without limitation, a server, a desktop computer, a notebook computer, a tablet computer, a thin client, a mobile device, a personal digital assistant (PDA), a phablet, and the like. In an example, computer system 200 may include a storage system or storage device.

[0021] In the example of FIG. 2, computer system 200 may include a monitoring module 110, a data replication priority module 112, and a data replication module 114.

[0022] Monitoring module 110 may monitor state of storage volumes belonging to data replication groups in a first storage device. In an instance, storage volumes in a first storage device (also called as a primary storage device or source storage device) may be organized into one or more data replication (DR) groups. Monitoring module 110 may monitor state of the storage volumes belonging to various data replication groups. In an example, the monitoring may include determining one or more parameters indicating potential failure of a storage volume, degradation of a storage volumes, and degradation of a log disk corresponding to a data replication group.

[0023] Data replication priority module 112 may assign a data replication priority to data replication groups in a first storage device. A data replication priority for a data replication group may define when a data replication group may be replicated from a first storage device to a second storage device in comparison to one or more other data replication groups in the first storage device. In an example, data replication priority may include a non-numerical value (for instance, LOW, MEDIUM, HIGH, and IMMEDIATE). In another example, data replication priority may include a numerical value (for instance, 1, 2, 3, and 4).

[0024] In an instance, data replication priority module 112 may assign a default data replication priority (for example, LOW) to each newly created data replication group

in a first storage device. Data replication priority module 112 may maintain a list of data replication groups assigned to a particular priority. Thus, for instance, data replication module may maintain a list of data replication groups that are assigned a LOW priority. Likewise, data replication module may maintain separate lists of data replication groups based on respective priorities. For example, separate lists may be maintained for LOW, MEDIUM, HIGH, and IMMEDIATE priorities. In an instance, these priority lists may be stored in memory.

[0025] Data replication priority module 112 may dynamically determine a data replication priority for each data replication group in a first storage device. In an example, a dynamic determination of a data replication priority for a data replication group may comprise a determination based on state of a storage volume belonging to the data replication group. In this regard, in an instance, data replication priority module 112 may obtain or receive information related to state of a storage volume(s) belonging to a data replication group from monitoring module.

[0026] In an example, determining a data replication priority for a data replication group based on the state of a storage volume(s) belonging to the data replication group may comprise determining whether the storage volume belonging to the data replication group is in a pre-defined state. A pre-defined state for a storage volume(s) may be defined by a user or system. Some non-limiting examples of a pre-defined state for a storage volume(s) may include a volume degradation state and a volume failure predicted state. If data replication priority module 112 identifies a storage volume in a pre-defined state, data replication priority module 112 may determine a data replication priority for the storage volume. In an example, a data replication priority for the storage volume may be determined based on RAID level of the storage volume, number of disk failures in a RAID type, and/or failure prediction for a disk(s) in a RAID type. To provide an example, data replication priority for a storage volume may be determined as shown in the priority table below.

Volume RAID Type	No. of Disk Failures	Failure Predicted	Data Replication Priority
RAID 0	0	FALSE	LOW
	0	TRUE	IMMEDIATE
RAID 1	0	FALSE	LOW
	0	TRUE	MEDIUM
	1	FALSE	HIGH
	1	TRUE	IMMEDIATE
RAID 5	0	FALSE	LOW
	0	TRUE	MEDIUM
	1	FALSE	HIGH
	1	TRUE	IMMEDIATE
RAID 6	0	FALSE	LOW
	0	TRUE	LOW
	1	FALSE	LOW
	1	TRUE	MEDIUM
	2	FALSE	HIGH
	2	TRUE	IMMEDIATE

[0027] For instance, a data replication priority for the storage volume based on RAID 0 may be determined as LOW if number of disk failures is 0 (zero) and failure prediction for the disk(s) in RAID 0 is FALSE. However, for the same storage volume, data replication priority may be determined as IMMEDIATE if number of disk failures is 0 (zero) but failure prediction for the disk(s) is TRUE (i.e. one or more disks in the RAID is expected to fail).

[0028] To provide another example, a data replication priority for the storage volume based on RAID 1 may be determined as LOW if number of disk failures is 0 (zero) and failure prediction for the disk(s) in RAID 1 is FALSE. However, for the same storage

volume, data replication priority may be determined as MEDIUM if number of disk failures is 0 (zero) but failure prediction for the disk(s) is TRUE (i.e. one or more disks in the RAID is expected to fail). Further, for the same storage volume, data replication priority may be determined as HIGH if number of failed disks is 1 but no further failure is predicted for a disk(s). Further still, for the same storage volume, data replication priority may be determined as IMMEDIATE if further failure prediction for the disk(s) is TRUE (i.e. one or more disks in the RAID is expected to fail).

[0029] If a storage volume belonging to a data replication group is not in a pre-defined state, the existing data replication priority for the storage volume may be considered as the data replication priority for the storage volume.

[0030] Once a data replication priority for a storage volume is determined, the data replication priority for the storage volume may be compared with an existing data replication priority for the data replication group. In an example, an existing data replication priority for a data replication group may include the default priority for the data replication group, which is defined by the data replication priority module when the data replication group is created. In another example, an existing data replication priority for a data replication group may include an earlier priority determined by the data replication priority module.

[0031] Upon comparison, if the existing data replication priority for the data replication group is lower than the data replication priority for the storage volume, data replication priority module may assign a new data replication priority to the data replication group and consider the new data replication priority as the data replication priority for the data replication group. In an instance, the new data replication priority for the data replication group may correspond to the data replication priority for the storage volume. In other words, if the existing data replication priority for the data replication group is lower than the data replication

priority for the storage volume, the data replication group may be assigned a new higher data replication priority that matches with the data replication priority for the storage volume.

[0032] On the other hand, if the existing data replication priority for the data replication group is not lower than the data replication priority for the storage volume, data replication priority module 112 may consider the existing data replication priority as the data replication priority for the data replication group. In other words, in such case, no change may be made in the data replication priority for the data replication group. Data replication priority module 112 may thus determine a new data replication priority for each data replication group in a first storage device.

[0033] In an instance, each data replication group in a first storage device may be classified into one of the priority lists maintained by data replication priority module 112 according to respective data replication priorities. For example, if data replication priority for a DR group is determined as HIGH, the DR group may be placed in a priority list that includes other DR groups having HIGH priority. In the event, a new priority is determined for a DR group, the DR group may be classified into a priority list corresponding to the new priority. For example, if new data replication priority for a DR group is determined as HIGH (as against an existing LOW priority), the DR group may be place in a HIGH priority list.

[0034] In an instance, data replication groups within a priority list may be further organized based on a parameter. The parameter may include, for example, number of I/O requests received on storage volumes belonging to respective data replication groups and number of storage volumes belonging to respective data replication groups.

[0035] Data replication module 114 may perform data replication, from the first storage device to a second storage device, for data replication groups according to their

respective data replication priorities. Data replication may begin with the data replication group having highest data replication priority, and continue with other DR groups in a decreasing order of priority.

[0036] FIG. 3 is a flowchart of an example method 300 for data replication. The method 300, which is described below, may at least partially be executed on data replication system 100 of FIG. 1 or computer system 200 of FIG. 2. However, other computing devices may be used as well. At block 302, state of storage volumes belonging to data replication groups in a first storage device may be monitored. At block 304, a data replication priority for each data replication group may be determined. In an example, a data replication priority for a data replication group may be determined based on state of a storage volume belonging to the data replication group. At block 306, data replication may be performed, from the first storage device to a second storage device, for data replication groups according to respective data replication priorities, beginning with the data replication group having highest data replication priority.

[0037] FIG. 4 is a block diagram of an example system 400 for data replication. System 400 includes a processor 402 and a machine-readable storage medium 404 communicatively coupled through a system bus. In an example, system 400 may be analogous to data replication system 100 of FIG. 1 or computer system 200 of FIG. 2. Processor 402 may be any type of Central Processing Unit (CPU), microprocessor, or processing logic that interprets and executes machine-readable instructions stored in machine-readable storage medium 404. Machine-readable storage medium 404 may be a random access memory (RAM) or another type of dynamic storage device that may store information and machine-readable instructions that may be executed by processor 402. For example, machine-readable storage medium 404 may be Synchronous DRAM (SDRAM), Double Data Rate (DDR), Rambus DRAM (RDRAM), Rambus RAM, etc. or a storage memory media such as a floppy disk, a hard disk, a CD-ROM, a DVD, a pen drive, and the like. In an example, machine-readable storage

medium 404 may be a non-transitory machine-readable medium. Machine-readable storage medium 404 may store instructions 406, 408, and 410. In an example, instructions 406 may be executed by processor 402 to monitor state of storage volumes belonging to data replication groups in a source storage device. Instructions 408 may be executed by processor 402 to determine a data replication priority for each data replication group, wherein a data replication priority for a data replication group is determined based on state of a storage volume belonging to the data replication group. Instructions 410 may be executed by processor 402 to perform data replication, from the source storage device to a target storage device, for data replication groups according to respective data replication priorities, beginning with the data replication group having highest data replication priority.

[0038] For the purpose of simplicity of explanation, the example method of FIG. 3 is shown as executing serially, however it is to be understood and appreciated that the present and other examples are not limited by the illustrated order. The example systems of FIGS. 1, 2 and 4, and method of FIG. 3 may be implemented in the form of a computer program product including computer-executable instructions, such as program code, which may be run on any suitable computing device in conjunction with a suitable operating system (for example, Microsoft Windows, Linux, UNIX, and the like). Embodiments within the scope of the present solution may also include program products comprising non-transitory computer-readable media for carrying or having computer-executable instructions or data structures stored thereon. Such computer-readable media can be any available media that can be accessed by a general purpose or special purpose computer. By way of example, such computer-readable media can comprise RAM, ROM, EPROM, EEPROM, CD-ROM, magnetic disk storage or other storage devices, or any other medium which can be used to carry or store desired program code in the form of computer-executable instructions and which can be accessed by a general purpose or special purpose computer. The computer readable instructions can also be accessed from memory and executed by a processor.

[0039] It may be noted that the above-described examples of the present solution is for the purpose of illustration only. Although the solution has been described in conjunction with a specific embodiment thereof, numerous modifications may be possible without materially departing from the teachings and advantages of the subject matter described herein. Other substitutions, modifications and changes may be made without departing from the spirit of the present solution. All of the features disclosed in this specification (including any accompanying claims, abstract and drawings), and/or all of the steps of any method or process so disclosed, may be combined in any combination, except combinations where at least some of such features and/or steps are mutually exclusive.

Claims:

1. A method of data replication, comprising:

monitoring state of storage volumes belonging to data replication groups in a first storage device;

determining a data replication priority for each data replication group, wherein a data replication priority for a data replication group is determined based on state of a storage volume belonging to the data replication group; and

performing data replication, from the first storage device to a second storage device, for data replication groups according to respective data replication priorities, beginning with data replication group having highest data replication priority.

2. The method of claim 1, wherein determining the data replication priority for the data replication group based on the state of the storage volume belonging to the data replication group, comprises:

determining whether the storage volume belonging to the data replication group is in a pre-defined state;

if the storage volume is in the pre-defined state, determining a data replication priority for the storage volume;

comparing the data replication priority for the storage volume with an existing data replication priority for the data replication group;

if the existing data replication priority for the data replication group is lower than the data replication priority for the storage volume, assigning a new data replication priority to the data replication group and considering the new data replication priority as the data replication priority for the data replication group; and

if the existing data replication priority for the data replication group is not lower than the data replication priority for the storage volume, considering the existing data replication priority as the data replication priority for the data replication group.

3. The method of claim 2, wherein the new data replication priority corresponds to the data replication priority for the storage volume.
4. The method of claim 2, wherein the existing data replication priority includes a default priority for the data replication group.
5. The method of claim 1, further comprising classifying each data replication group into one of priority lists according to respective data replication priorities.
6. The method of claim 5, further comprising classifying data replication groups within a priority list according to a parameter.
7. A system for data replication, comprising:
 - a monitoring module to monitor state of storage volumes belonging to data replication groups in a primary storage device;
 - a data replication priority module to determine a data replication priority for each data replication group, wherein a data replication priority for a data replication group is determined based on state of a storage volume belonging to the data replication group; and
 - a data replication module to perform data replication, from the primary storage device to a secondary storage device, for data replication groups according to respective data replication priorities, beginning with data replication group having highest data replication priority.
8. The system of claim 7, wherein to determine the data replication priority for the data replication group based on the state of the storage volume belonging to the data replication group, the data replication module to:
 - determine whether the storage volume belonging to the data replication group is in a pre-defined state;

if the storage volume is in the pre-defined state, determine a data replication priority for the storage volume;

compare the data replication priority for the storage volume with an existing data replication priority for the data replication group;

if the existing data replication priority for the data replication group is lower than the data replication priority for the storage volume, modify the existing data replication priority for the data replication group to correspond to the data replication priority for the storage volume and consider the modified data replication priority as the data replication priority for the data replication group; and

if the existing data replication priority for the data replication group is not lower than the data replication priority for the storage volume, consider the existing data replication priority as the data replication priority for the data replication group.

9. The system of claim 7, wherein to monitor includes determining one of a parameter indicating potential failure of the storage volumes, degradation of the storage volumes, and degradation of a log disk corresponding to the data replication group.

10. A non-transitory machine-readable storage medium comprising instructions for data replication, the instructions executable by a processor to:

monitor state of storage volumes belonging to data replication groups in a source storage device;

determine a data replication priority for each data replication group, wherein a data replication priority for a data replication group is determined based on state of a storage volume belonging to the data replication group; and

perform data replication, from the source storage device to a target storage device, for data replication groups according to respective data replication priorities, beginning with data replication group having highest data replication priority.

11. The storage medium of claim 10, wherein instructions to determine the data replication priority for the data replication group based on the state of the storage volume belonging to the data replication group, comprises instructions to:

determine whether the storage volume belonging to the data replication group is in a pre-defined state;

if the storage volume is in the pre-defined state, determine a data replication priority for the storage volume;

compare the data replication priority for the storage volume with an existing data replication priority for the data replication group;

if the existing data replication priority for the data replication group is lower than the data replication priority for the storage volume, assign a new data replication priority to the data replication group and consider the new data replication priority as the data replication priority for the data replication group; and

if the existing data replication priority for the data replication group is not lower than the data replication priority for the storage volume, consider the existing data replication priority as the data replication priority for the data replication group.

12. The storage medium of claim 11, further comprising instructions to classify each data replication group into one of priority lists according to respective data replication priorities, wherein data replication groups within a priority list are organized according to a parameter.

13. The storage medium of claim 12, wherein the parameter includes number of I/O requests received on storage volumes belonging to respective data replication groups.

14. The storage medium of claim 12, wherein the parameter includes number of storage volumes belonging to respective data replication groups.

15. The storage medium of claim 11, wherein the pre-defined state includes one of a volume degradation state or a volume failure predicted state.

1/4




Fig. 1

2/4




Fig. 2

3/4


300

Fig. 3

4/4

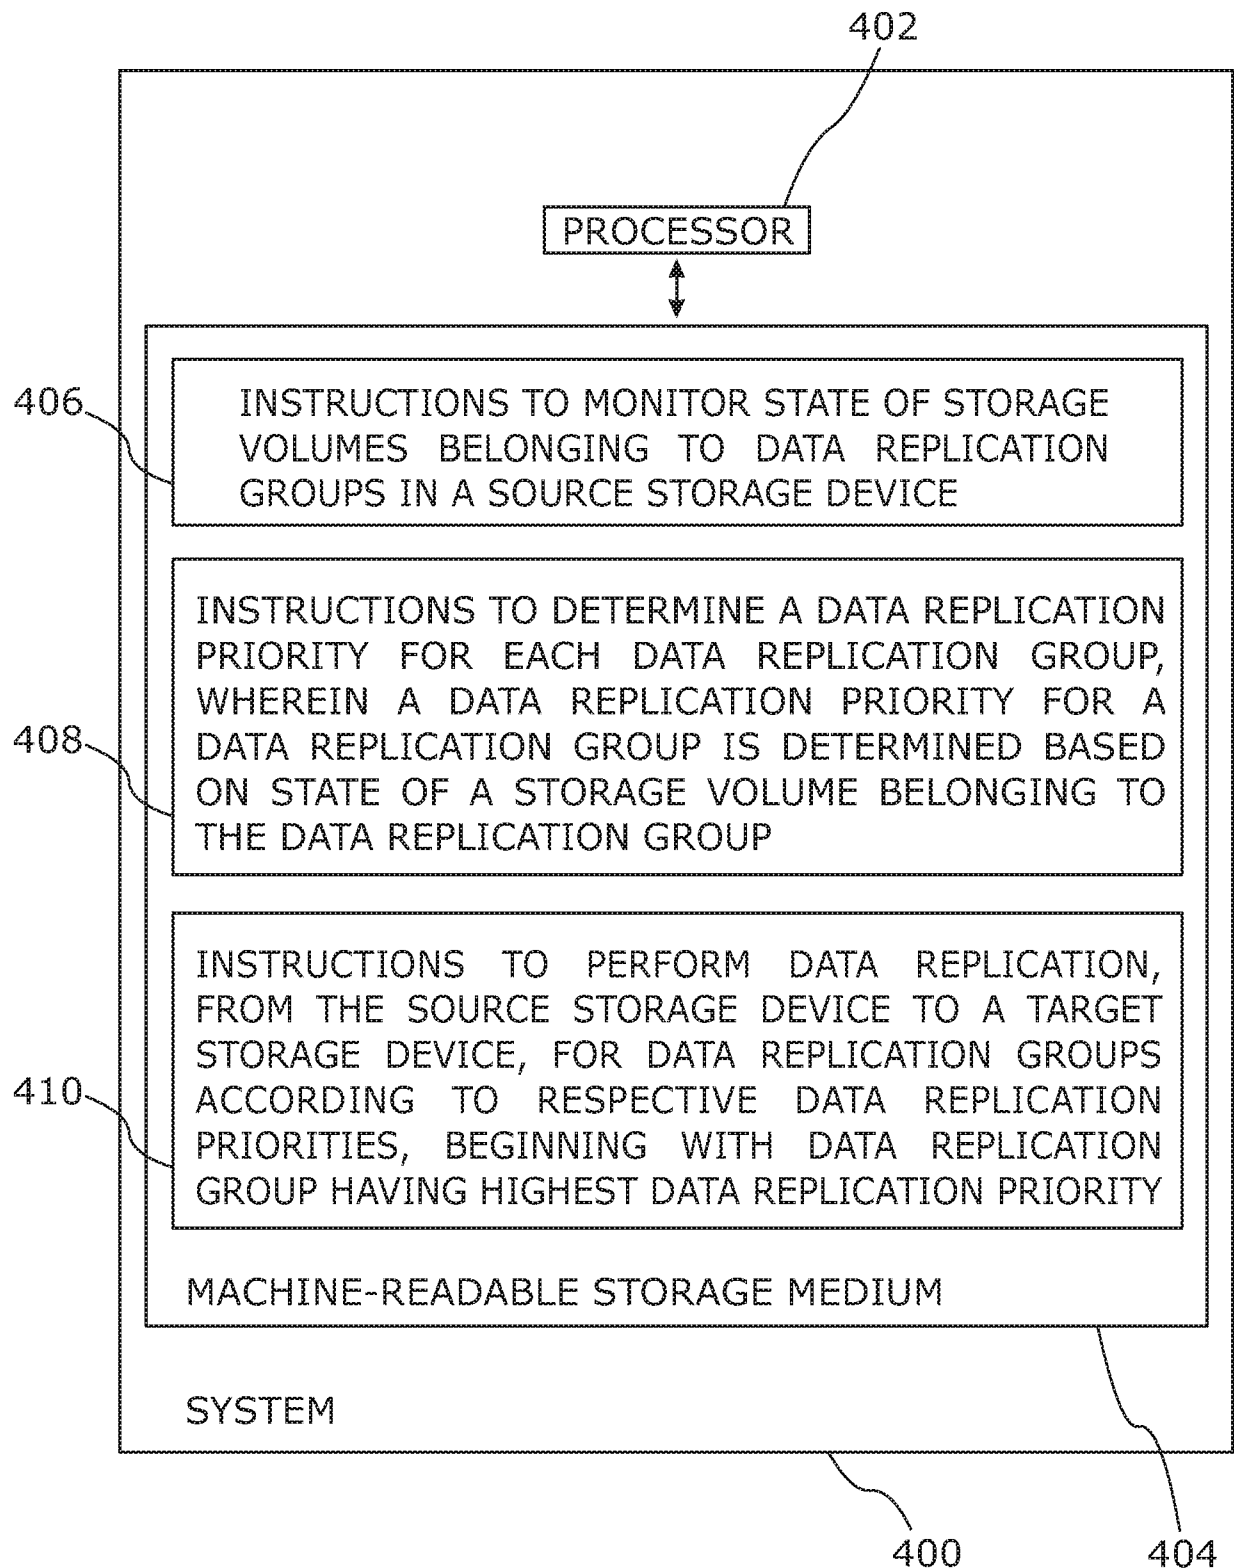


Fig. 4

A. CLASSIFICATION OF SUBJECT MATTER**G06F 11/30(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)

G06F 11/30; G06F 15/16; G06F 3/06; G06F 17/30; G06F 7/00

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: data replication, priority, logical unit, volume, degradation, data group, and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2014-0324774 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 30 October 2014 See paragraphs [0030]–[0041] and figure 3.	1–15
A	US 2010-0299447 A1 (NILESH ANANT SALVI et al.) 25 November 2010 See paragraphs [0016]–[0022] and figure 2.	1–15
A	US 2012-0185433 A1 (THEODORE T. HARRIS, JR. et al.) 19 July 2012 See paragraphs [0024]–[0027] and figure 1.	1–15
A	US 8,341,121 B1 (CHRISTOPHER HERCULES CLAUDATOS et al.) 25 December 2012 See column 4, line 60 – column 5, line 19; and figure 2.	1–15
A	EP 2098950 A2 (HITACHI LTD.) 09 September 2009 See paragraphs [0262]–[0264] and figure 38.	1–15



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 but 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

29 February 2016 (29.02.2016)

Date of mailing of the international search report

29 February 2016 (29.02.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/054200

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014-0324774 A1	30/10/2014	CN 104123198 A	29/10/2014
US 2010-0299447 A1	25/11/2010	None	
US 2012-0185433 A1	19/07/2012	None	
US 8341121 B1	25/12/2012	None	
EP 2098950 A2	09/09/2009	CN 101526884 A	09/09/2009
		CN 101526884 B	27/03/2013
		EP 2098950 A3	28/09/2011
		JP 2009-217379 A	24/09/2009
		JP 5410025 B2	05/02/2014
		US 2009-0228676 A1	10/09/2009
		US 2011-0258407 A1	20/10/2011
		US 2012-0191941 A1	26/07/2012
		US 8151080 B2	03/04/2012
		US 8271761 B2	18/09/2012
		US 8423746 B2	16/04/2013