



- (51) **International Patent Classification:**
G06F 11/14 (2006.01) *G06F 12/16* (2006.01)
G06F 12/06 (2006.01)
- (21) **International Application Number:**
PCT/US2016/021799
- (22) **International Filing Date:**
10 March 2016 (10.03.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
6514/CHE/2015 4 December 2015 (04.12.2015) IN
- (71) **Applicant:** HEWLETT PACKARD ENTERPRISE DEVELOPMENT LP [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).
- (72) **Inventors:** MARISAMY, Venkatesh; Sy. No. 192, Whitefield Road, Mahadevapura Post, Bangalore, Karnataka 560048 (IN). VEDARAMAN, Kanthimathi; Sy. No. 192,

Whitefield Road, Mahadevapura Post, Bangalore, Karnataka 560048 (IN).

- (74) **Agents:** PATEL, Neel K. 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, 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,

[Continued on next page]

- (54) **Title:** CREATE SNAPSHOT OF SUBSET OF DATA

FIG. 4



- (57) **Abstract:** A subset of a data set is determined at a Logical Unit Number (LUN) to be replicated to a storage system. A physical address of the subset is determined. A snapshot of the subset corresponding to the determined address is created. The created snapshot includes less than an entirety of the LUN.



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).

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:

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

Declarations under Rule 4.17:

— *as to the identity of the inventor (Rule 4.17(i))*

CREATE SNAPSHOT OF SUBSET OF DATA

BACKGROUND

[0001] Today, the data centers are the hosting platforms for private and public cloud offerings. The data sets at these data centers is often protected and managed. Example data sets may include files as well as Virtual Machines (VMs) residing on a Network-Attached Storage (NAS) server or on an Array.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] The following detailed description references the drawings, wherein:

[0003] FIG. 1 is an example block diagram of a system to create a snapshot of a subset of data;

[0004] FIG. 2 is another example block diagram of a system to create a snapshot of a subset of data;

[0005] FIG. 3 is an example block diagram of a computing device including instructions for creating a snapshot of a subset of data; and

[0006] FIG. 4 is an example flowchart of a method for creating a snapshot of a subset of data.

DETAILED DESCRIPTION

[0007] Specific details are given in the following description to provide a thorough understanding of embodiments. However, it will be understood that embodiments may be practiced without these specific details. For example, systems may be shown in block diagrams in order not to obscure embodiments

in unnecessary detail. In other instances, well-known processes, structures and techniques may be shown without unnecessary detail in order to avoid obscuring embodiments.

[0008] In Big Data centers, the data management may be integral to continuous operations and for providing effective services. Generally, the data may be hosted on a Redundant Array of Independent Disks (RAID) based arrays and administrators may run array based snapshots to create a replicas for vertical and horizontal management of data. Snapshots may also be used as sources for data backups. These snapshots may be taken at a Logical Unit Number (LUN) level, which are not aware of the applications or databases since they are sitting on top of OS/Files system. Thus, no application based snapshots are possible. Further, a LUN snapshot meant for a particular application may include all types of data which is not relevant for the purpose it was created. This may result in overly expensive snapshots in terms of size and operational efficiency of the array.

[0009] Management and protection of the virtualized data may not be as straightforward as in the case of physical servers. To protect a VM, it may have to be in quiescent mode to maintain VM consistency first, and then a snapshot at array level may have to be taken. For example, if hundreds of VMs are sitting on single LUN, and there is a request is to create a snapshot of VMs belonging to a particular user (e.g. 10 VMs), the snapshot may be taken for the complete LUN which has all other 90 VMs which don't belong to that user. This does not address the actual requirement of data center management, also results in excessive use/wastage of processing cycles plus the space of the Array.

Further, current techniques do not provide a mechanism of application awareness at the block level to protect only blocks corresponding to the particular data set, which is chosen by the data/backup administrator to be protected.

[0010] Examples provide a method and/or system to take a snapshot of only the relevant data, such as a data set chosen by an administrator. In one example, a system may include a data protection unit, an address unit and a snapshot unit. The data protection unit may determine a subset of a data set at a Logical Unit Number (LUN) to be replicated to a storage system. The address unit may determine a physical address of the subset. The snapshot unit may create a snapshot of the subset corresponding to the determined address. The created snapshot may include less than an entirety of the LUN.

[0011] Examples may reduce a number of times copy-on-write is performed, which is performed when any change is made in the blocks for which the snapshot is created. This is because examples may reduce the number of blocks in snapshot. Hence, examples may avoid copy-on-write for the blocks not related to the snapshot, so processing cycles and space usage may be improved or optimized only for required blocks.

[0012] Examples may provide continuous block level thin snapshots, which may remove the dependency on application level incremental backups. Examples may also store metadata based on the thin snapshots, which may then be presented instantaneously to choose the Files/VMs to be accessed for further data management, such as restoring and/or replica access.

[0013] Referring now to the drawings, FIG. 1 is an example block diagram of a system 100 to create a snapshot of a subset of data. The system 100 may include or be part of a microprocessor, a controller, a memory module or device, a notebook computer, a desktop computer, an all-in-one system, a server, a network device, a wireless device, a network and the like.

[0014] The system 100 is shown to include a data protection unit 110, an address unit 120 and a snapshot unit 130. The data protection, address and snapshot units 110, 120 and 130 may include, for example, a hardware device including electronic circuitry for implementing the functionality described below, such as control logic and/or memory. In addition or as an alternative, the data protection, address and snapshot units 110, 120 and 130 may be implemented as a series of instructions encoded on a machine-readable storage medium and executable by a processor.

[0015] The data protection unit to determine a subset 112 of a data set at a Logical Unit Number (LUN) to be replicated to a storage system. The data set may include a file, file system, virtual machine (VM), and the like. The storage system may include at least one of storage virtualization, a storage area network (SAN), and a Network attached storage (NAS).

[0016] The address unit 120 may determine a physical address 122 of the subset 112. The snapshot unit 130 may create a snapshot 132 of the subset 112 corresponding to the determined address 122. The created snapshot 132 may include less than an entirety of the LUN. The term snapshot may refer to a snapshot is the state of a system at a particular point in time and the term LUN

may refer to a virtual device that provides an area of usable storage capacity on one or more physical disk drive(s) in a computer system.

[0017] The may be located separately from the LUN or storage system. Similarly, the LUN may be located separately from the storage system. For example, the data protection, address and/or snapshot units 110, 120 and 130 may be located at a host while the LUN is located at a client. The storage system may be located at the host or at a site remote from both the host and the client. The system 100 is explained in greater detail below with respects to FIGS. 2-4.

[0018] FIG. 2 is another example block diagram of a system 200 to create a snapshot of a subset of data. The system 200 may include or be part of a microprocessor, a controller, a memory module or device, a notebook computer, a desktop computer, an all-in-one system, a server, a network device, a wireless device, a network and the like. Further, the system 200 of FIG. 2 may include at least the functionality and/or hardware of the system 100 of FIG. 1. For example, a data protection unit 210, an address unit 220 and a snapshot unit 230 of the system 200 of FIG. 2 may include at least the respective functionality and/or hardware of the data protection, address and/or snapshot units 110, 120 and 130 of the system 100 of FIG. 1. The system 200 of FIG. 2 is also shown to include the LUN 240 and the storage system 250. The storage system 250 may be, for example a NAS array.

[0019] As noted above, the data protection unit 210 may determine a subset 212 of a data set 241 at the LUN 240 to be replicated to the storage system 250. For example, an administrator, protection policy and/or a user may

determine the subset 212 to be replicated and then indicate the subset to the data protection unit 210, such as through a Graphical User Interface (GUI).

[0020] The address unit 220 may determine a physical address 222 of the subset 212, such as via plugin like Use tool like vmkfstools or hdparm. The snapshot unit 230 may create a snapshot of the subset 212 corresponding to the determined address 222. In one example, the address unit 220 may create a continuous filter string 224 indicating an address 222 of the subset 212. In this case, the snapshot unit 230 may use the filter string 224 to create the snapshot 232 in the storage system 250.

[0021] Here, the data set 241 at the LUN 240 is shown to include a set of 4 VMs 242-1 to 242-4. However, examples are not limited to VMs and may include other types of objects, such as file systems or files. Further, while FIG. 2 shows 4 VMs 242-1 to 242-4, examples of the data 241 set may include more or less than 4 VMs. The LUN 240 may be a thin-provisioned volume, which allows many VMs to be stored on a same data volume.

[0022] The data protection unit 210 may select the subset 212 of the set 241 of VMs 242. Here, it may be decided, such as by the administrator or user, to select the first and third VMs 242-1 and 242-3. Therefore, the data protection unit 210 may determine the subset 212 to correspond to only the subset 212 of the VMs 242, such the first and third VMs 242-1 and 242-3. The snapshot unit 230 may then create the snapshot 232 to correspond to the subset 212 of VMs 242. The created snapshot 232 may not include any of a remainder of the set 241 of VMs 242. Thus, the snapshot 232 created at the storage system 250

may only include the first and third VMs 242-1 and 242-3. Further, the created snapshot 232 may include less than an entirety of the LUN.

[0023] In another example, the data protection unit 210 may further determine the data set 241 to correspond to only an application (not shown) of the subset 212 of VMs 242. In this case, the snapshot unit 230 may create a snapshot 232 of only blocks of data corresponding to the application.

[0024] The data protection unit 210 may trigger the snapshot unit 230 to create a new snapshot 232 and/or modify the current snapshot 232, if at least one block of data of the subset 212 is changed. For example, a copy-on-write may be carried out for a change to any selected block of data corresponding to the subset 212. Conversely, a copy-on-write may be skipped for a change to a block of data not corresponding to the subset 212. In another example, the snapshot unit 230 may also store metadata 260, such as a mapping table, of the subset 212 separately for data management when creating the snapshot 232.

[0025] FIG. 3 is an example block diagram of a computing device 300 including instructions for creating a snapshot of a subset of data. In the embodiment of FIG. 3, the computing device 300 includes a processor 310 and a machine-readable storage medium 320. The machine-readable storage medium 320 further includes instructions 322, 324 and 326 for creating a snapshot of a subset of data.

[0026] The computing device 300 may be included in or part of, for example, a microprocessor, a controller, a memory module or device, a notebook computer, a desktop computer, an all-in-one system, a server, a network device, a

wireless device, or any other type of device capable of executing the instructions 322, 324 and 326. In certain examples, the computing device 300 may include or be connected to additional components such as memories, controllers, etc.

[0027] The processor 310 may be, at least one central processing unit (CPU), at least one semiconductor-based microprocessor, at least one graphics processing unit (GPU), a microcontroller, special purpose logic hardware controlled by microcode or other hardware devices suitable for retrieval and execution of instructions stored in the machine-readable storage medium 320, or combinations thereof. The processor 310 may fetch, decode, and execute instructions 322, 324 and 326 to implement creating the snapshot of the subset of data. As an alternative or in addition to retrieving and executing instructions, the processor 310 may include at least one integrated circuit (IC), other control logic, other electronic circuits, or combinations thereof that include a number of electronic components for performing the functionality of instructions 322, 324 and 326.

[0028] The machine-readable storage medium 320 may be any electronic, magnetic, optical, or other physical storage device that contains or stores executable instructions. Thus, the machine-readable storage medium 320 may be, for example, Random Access Memory (RAM), an Electrically Erasable Programmable Read-Only Memory (EEPROM), a storage drive, a Compact Disc Read Only Memory (CD-ROM), and the like. As such, the machine-readable storage medium 320 can be non-transitory. As described in detail below, machine-readable storage medium 320 may be encoded with a series of executable instructions for creating the snapshot of the subset of data.

[0029] Moreover, the instructions 322, 324 and 326, when executed by a processor (e.g., via one processing element or multiple processing elements of the processor) can cause the processor to perform processes, such as, the process of FIG. 4. For example, the receive instructions 322 may be executed by the processor 310 to receive a request to take a snapshot of a subset of a data set.

[0030] The determine instructions 324 may be executed by the processor 310 to determine a physical address of the subset. The create instructions 326 may be executed by the processor 310 to create a snapshot of the subset based on the determined address. The created snapshot may not include a remainder of the data set. Further, the snapshot may only be modified if a block of data of the subset changes.

[0031] FIG. 4 is an example flowchart of a method 400 for creating a snapshot of a subset of data. Although execution of the method 400 is described below with reference to the system 100, other suitable components for execution of the method 400 can be utilized, such as the system 200. Additionally, the components for executing the method 400 may be spread among multiple devices (e.g., a processing device in communication with input and output devices). In certain scenarios, multiple devices acting in coordination can be considered a single device to perform the method 400. The method 400 may be implemented in the form of executable instructions stored on a machine-readable storage medium, such as storage medium 320, and/or in the form of electronic circuitry.

[0032] At block 410, the system 100 selects a subset 112 of a data set stored on a single logical unit number (LUN) of a storage system. At block 420, the system 100 determines a location of the subset 112. At block 430, the

system 100 creates a snapshot 132 of the blocks of data corresponding to the subset 112. The snapshot 132 is triggered if any the blocks of data corresponding to the subset 112 change. However, the snapshot 132 may not be triggered in response to a block of data changing that corresponds a remainder of the data set.

CLAIMSWe claim:

1. A system, comprising:
a data protection unit to determine a subset of a data set at a Logical Unit Number (LUN) to be replicated to a storage system;
an address unit to determine a physical address of the subset; and
a snapshot unit to create a snapshot of the subset corresponding to the determined address, wherein
the created snapshot includes less than an entirety of the LUN.
2. The system of claim 1, wherein,
the data set includes at least one of a file, file system and a virtual machine (VM), and
the storage system includes at least one of storage virtualization, a storage area network (SAN), and a Network attached storage (NAS).
3. The system of claim 2, wherein,
the LUN includes a set of VMs, and
a subset of the set of VMs is selected.
4. The system of claim 3, wherein,
the data protection unit is to determine the subset to correspond to only the subset of VMs, and
the snapshot unit is to create the snapshot to correspond to the subset of

VMs.

5. The system of claim 4, wherein,
the data protection unit is to further determine the data set to correspond to only an application of the subset of VMs, and
the snapshot unit is to create a snapshot of only blocks of data corresponding to the application.

6. The system of claim 4, wherein the created snapshot does not include any of a remainder of the set of VMs.

7. The system of claim 1, wherein at least one of an administrator, a user, and a protection policy are to determine the subset to be replicated.

8. The system of claim 1, wherein,
the address unit is to create a continuous filter string indicating an address of the subset, and
the snapshot unit is to use the filter string to create the snapshot in the storage system.

9. The system of claim 1, wherein the data protection unit is to trigger the snapshot unit to at least one of create a new snapshot and modify the current snapshot, if at least one block of data of the subset is changed.

10. The system of claim 10, wherein,
a copy-on-write is carried out for a change to any selected block of data corresponding to the subset, and
a copy-on-write is skipped for a change to a block of data not corresponding to the subset.

11. The system of claim 1, wherein the snapshot unit is to store metadata of the subset separately for data management when creating the snapshot.

12. A method, comprising:
selecting subset of a data set stored on a single logical unit number (LUN) of a storage system;
determining a location of the subset; and
creating a snapshot of the blocks of data corresponding to the subset,
wherein
the snapshot is triggered, if any of the blocks of data corresponding to the subset change.

13. The method of claim 12, wherein a snapshot is not triggered in response to a block of data changing that corresponds a remainder of the data set.

14. A non-transitory computer-readable storage medium storing instructions that, if executed by a processor of a device, cause the processor to:

- receive a request to take a snapshot of a subset of a data set;
- determine a physical address of the subset; and
- create a snapshot of the subset based on the determined address,

wherein

the created snapshot does not include a remainder of the data set.

15. The non-transitory computer-readable storage medium of claim 14, wherein the snapshot is modified only if a block of data of the subset changes.

FIG. 1




FIG. 2




FIG. 3





FIG. 4



INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/021799**A. CLASSIFICATION OF SUBJECT MATTER****G06F 11/14(2006.01)i, G06F 12/06(2006.01)i, G06F 12/16(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 11/14; G06F 12/16; G06F 3/06; G06F 17/30; G06F 12/00; G06F 12/06

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: Logical Unit Number(LUN), subset, snapshot, storage, VMs, 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.
X	US 7,243,198 B2 (BRIAN THOMAS BERKOWITZ et al.) 10 July 2007 See column 2, lines 21-23; column 3, lines 63-66; column 4, lines 56-67; column 6, lines 7-14; column 7, lines 23-26; column 10, lines 49-51; column 11, line 57 - column 12, line 30; column 12, line 53 - column 13, line 39; column 27, lines 16-19; column 28, lines 10-13, 32-36; and figures 1 and 3-5.	1-8, 11, 14
Y		9-10, 12-13, 15
Y	US 2003-0229651 A1 (YOICHI MIZUNO et al.) 11 December 2003 See paragraphs [0038] and [0044]; and figure 1.	9-10, 12-13, 15
A	US 2010-0077160 A1 (PETER CHI-HSIUNG LIU et al.) 25 March 2010 See paragraphs [0180]-[0184]; claim 1; and figure 3.	1-15
A	US 2015-0212893 A1 (COMMVault SYSTEMS, INC.) 30 July 2015 See paragraphs [0188]-[0189] and figure 1A.	1-15
A	CN 104407935 A (HUAWEI DIGITAL TECHNOLOGY (CHENGDU) CO., LTD.) 11 March 2015 See paragraphs [0039]-[0046] and figure 2.	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

08 August 2016 (08.08.2016)

Date of mailing of the international search report

08 August 2016 (08.08.2016)

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. +82-42-481-8578

Authorized officer

LEE, EUN KYU

Telephone No. +82-42-481-3580



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/021799

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 7243198 B2	10/07/2007	US 2003-0149736 A1	07/08/2003
		US 2005-0004979 A1	06/01/2005
		US 2007-0283111 A1	06/12/2007
		US 6826666 B2	30/11/2004
		US 7721057 B2	18/05/2010
US 2003-0229651 A1	11/12/2003	JP 2003-280964 A	03/10/2003
		US 7120768 B2	10/10/2006
US 2010-0077160 A1	25/03/2010	AT 497210 T	15/02/2011
		AT 555438 T	15/05/2012
		AU 2006-262045 A1	04/01/2007
		AU 2006-262045 B2	03/11/2011
		AU 2006-262046 A1	04/01/2007
		AU 2006-262046 B2	12/05/2011
		CA 2613359 A1	04/01/2007
		CA 2613419 A1	04/01/2007
		CA 2613419 C	02/12/2014
		CN 101243413 A	13/08/2008
		CN 101243413 B	14/08/2013
		CN 101495970 A	29/07/2009
		CN 101495970 B	17/07/2013
		CN 102117233 A	06/07/2011
		CN 102117233 B	17/07/2013
		EP 1907935 A2	09/04/2008
		EP 1907935 A4	22/04/2009
		EP 1907935 B1	26/01/2011
		EP 1907972 A2	09/04/2008
		EP 1907972 A4	20/04/2011
		EP 1907972 B1	25/04/2012
		EP 2477114 A2	18/07/2012
		EP 2477114 A3	24/10/2012
		EP 2477114 B1	15/01/2014
		EP 2713271 A2	02/04/2014
		EP 2713271 A3	16/04/2014
		EP 2713271 B1	12/08/2015
		HK 1119800 A1	14/03/2014
		HK 1136048 A1	25/04/2014
		HK 1159285 A1	11/04/2014
		HK 1173805 A1	15/08/2014
		JP 2008-547123 A	25/12/2008
		JP 2009-506399 A	12/02/2009
		JP 2012-198918 A	18/10/2012
		JP 2013-254513 A	19/12/2013
		JP 5420242 B2	19/02/2014
		JP 5430711 B2	05/03/2014
		US 2009-0222496 A1	03/09/2009
		US 2011-0218968 A1	08/09/2011
		US 2012-0297246 A1	22/11/2012

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/021799

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 2013-0086347 A1	04/04/2013
		US 2014-0325267 A1	30/10/2014
		US 7937547 B2	03/05/2011
		US 8255651 B2	28/08/2012
		US 8346726 B2	01/01/2013
		US 8706992 B2	22/04/2014
		US 9116847 B2	25/08/2015
		WO 2007-002397 A2	04/01/2007
		WO 2007-002397 A3	16/04/2009
		WO 2007-002398 A2	04/01/2007
		WO 2007-002398 A3	21/06/2007
US 2015-0212893 A1	30/07/2015	None	
CN 104407935 A	11/03/2015	None	