

(51) International Patent Classification:
G06F 17/30 (2006.01)(21) International Application Number:
PCT/US2015/035576(22) International Filing Date:
12 June 2015 (12.06.2015)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant: **HEWLETT PACKARD ENTERPRISE DEVELOPMENT LP** [US/US]; 11445 Compaq Center Drive West, Houston, TX 77070 (US).(72) Inventors: **DONAGHY, Dave**; Longdown Avenue Stoke, Gifford, Bristol BS34 8QZ (GB). **WHELAN, Josh**; Longdown Avenue Stoke, Gifford, Bristol BS34 8QZ (GB). **BUTT, John**; Longdown Avenue Stoke, Gifford, Bristol BS34 8QZ (GB).(74) Agents: **GONCALVES, Brian** et al.; Hewlett Packard Enterprise, 3404 E. Harmony Road, Mail Stop 79, Fort Collins, CO 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, 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))

[Continued on next page]

(54) Title: DISK STORAGE ALLOCATION




FIG. 2

(57) Abstract: According to an example, to allocate disk storage space, a request to store a new data chunk in a container is received. For a list of available containers, an exposure count for the container is fetched in the event that a container comprises a data chunk from a same source as the new data chunk. An optimal container to store the new data chunk is determined, and in the event that an exposure count of the optimal container is less than an exposure count limit, the new data chunk in the optimal container is stored.



— *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))*

DISK STORAGE ALLOCATION

BACKGROUND

[0001] Computing systems, devices, and electronic components in general may utilize content in the form of digital files. Systems may create files, store files, or receive files, as examples. A user may retain or have access to a large number of files stored across a multitude of devices, systems, ecosystems, accounts, and/or platforms. In some systems, a file may be written and/or stored in separate parts or chunks. Failure of a single disk component at either the hardware or software level may have a significant impact on a user.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] FIG. 1 illustrates a disk container exposure count module and data storage containers including an exposure count, according to an example of the present disclosure;

[0003] FIG. 2 illustrates a flowchart of storing a new data chunk in a container, according to an example of the present disclosure; and

[0004] FIG. 3 illustrates a schematic representation of a computing device that may be used as a platform for storing a new data chunk in a container, according to an example of the present disclosure.

DETAILED DESCRIPTION

[0005] Various examples described below provide for allocating disk storage space on a drive or drives according to, for example, an allocation policy. An allocation policy may specify that a specific disk container or other storage

element on a disk may only store data related to a limited number of data items or data chunks, such that if the container were to fail, the “exposure” or impact on a user would be minimized.

[0006] Generally, computer systems include storage devices to store parts of a data item or entity such as a file, database table, memory page, communication stream, or other data for subsequent retrieval. In some examples, data items may be written and/or stored as separate parts in multiple storage locations, e.g., as data chunks. In some systems, the number of data items stored across a disk or disks may number in the billions, and may be stored across a complex configuration of hardware and software elements.

[0007] Data may also be de-duplicated or compressed to save disk storage space. For example, a single chunk of data that is replicated across multiple files may be stored only once, with the associated files or data items having a pointer or reference to the storage location of the common data chunk. In general, various techniques may be used on a disk system, including across multiple disks, to optimize data storage, read and write times, efficiency, and other metrics affecting performance, reliability, and other factors.

[0008] While such techniques improve the overall user and management experience and lower storage costs, they may increase the impact that failure of a disk, component, or disk sector could have on the overall system. For example, a disk may store a large number of data chunks in a single container on the disk. A container may be, in an example, an internal file used by a disk storage system to store data.

[0009] Data chunks in a container may relate to or “map” to a larger number of data items, particularly in the case where the data chunks are stored once and mapped to multiple data items in a one-to-many or even many-to-many relationship or dependencies. If, for example, the disk sector or disk area storing that container were to fail, e.g., such as a hardware failure or a software-level corruption, the large number of data items storing data chunks in that container could be damaged, unrepairable, or unrecoverable based on a failure of the single container.

[0010] A user or administrator of a system may desire to minimize the number of data items that can be lost should a single container fail, corrupt, or otherwise be unreadable or unrecoverable by limiting the number of data chunks and/or data items that may be written to a single container. Such an exposure count limit may be expressed as an absolute number, e.g., only one item may be lost from a single failure, or may expressed as a proportion or percentage, e.g., only 10% of data items stored on a disk may be lost from a single failure.

[0011] In another example, a combination of criteria may be used to determine an exposure count limit, such as that only 10% of items, but in no case no more than 5 items, may be lost due to a single failure. In other examples, the limit may be adaptive in nature based on, for example, whether data is stored on a personal or enterprise system, whether data is classified as mission critical, whether data is stored on one disk or across many disks, or whether data is stored on a system enabled with a redundant array of inexpensive disks or RAID.

[0012] In an example of allocating storage on a disk or disks to minimize the impact of data storage failure, when writing a new data chunk to a disk, a count of

items stored in a container, e.g., an exposure count, is determined for each container that is available or accessible to store the new data chunk. An optimal container, e.g., a container with the lowest exposure count, is then determined such that new data may be stored in a container that would minimize impact on a user should the container fail.

[0013] In such an example, in the event that the exposure count of the optimal container is less than an exposure count limit set in the allocation policy, the new data chunk may be stored to the container with the lowest exposure count, e.g., the optimal container. If the exposure count is equal to or greater than the exposure count limit in the allocation policy, a new container may be created and the new data chunk may be stored in the new container.

[0014] FIG. 1 illustrates a disk container exposure count module and data storage containers including an exposure count, according to an example of the present disclosure.

[0015] A disk storage system (hereinafter simply a “disk”) may comprise a container or containers, e.g., containers 108 and 110, which may be internal files stored on the disk for purposes of housing a data item or items, or a data chunk or data chunks. For example, a disk may store data items or files 102, 104, and 106, which may be comprised of data chunks 102a-d, 104a-d, and 106a-d, respectively.

[0016] Disk container exposure count module 112 may contain instructions to store the data chunks 102a-d, 104a-d, and 106a-d across containers on the disk, including on, but not limited to, containers 108 and 110. The instructions stored in disk container exposure count module 112 are discussed below in more detail with respect to FIGS. 2 and 3.

[0017] In the example of FIG. 1, an allocation policy for the disk, or more generally for a system or systems housing or accessing the disk of FIG. 1, may indicate that the maximum number of data items exposed to a container is two. More specifically, the allocation policy controlling the disk of FIG. 1 may indicate that once chunks from two data items, e.g., two sources, are stored in a single container, attempts to store data chunks from a third data item will result in an instruction to store the data chunks of the third item in a different container that has not exceeded the exposure count limit of the allocation policy.

[0018] For example, in FIG. 1, data chunks 102a-d and 104a-d may be stored in container 1, resulting in an exposure count of 2 for container 1. However, the allocation policy of the disk shown in FIG. 1 would then store data chunks 106a-d in a different container, e.g., container 2, or a new container if no containers that satisfied the allocation policy exposure count limit were available.

[0019] FIG. 2 illustrates a flowchart of storing a new data chunk in a container, according to an example of the present disclosure.

[0020] In block 202, a request to store a new data chunk (or chunks) in a container is received. The request may originate from an operating system, disk system, disk controller, software program or application, network communication, user input, or other system, application, or protocol capable of instructing a disk to write data.

[0021] In block 204, a loop through containers available for storage on the disk may be commenced. For each container, in block 206, a determination may be made as to whether the container already stores a data chunk from the same data item or source as the new data chunk to be written to the disk.

[0022] In block 210, if the container already stores a data chunk from the same data item or source as the new data chunk to be written to the disk, the existing exposure count for the container is fetched. In such cases, the exposure count of the container would not be increased by storing the new data chunk to the container, since a data chunk from the data item is already stored in the container and thus the data item is already exposed.

[0023] In block 208, if the container does not already store a data chunk from the same data item or source as the new data chunk to be written to the disk, the existing exposure count for the container is fetched and incremented by one. In such an example, storing the new data chunk to the container would increase the exposure count of the container, since no data from that data item has been stored in the container.

[0024] In some examples, the flow of blocks 204 through 210 may be executed each time a data chunk is to be stored, or at a particular interval. In some examples, the exposure counts may be stored on a disk system, while in other examples the exposure counts may be stored in an indexed database.

[0025] In block 212, an optimal container to store the new data chunk is determined. In some examples, the optimal container may be the container on the disk with lowest exposure count. In examples where multiple containers have equally low exposure counts, other factors may influence or determine which of those containers are selected, such as disk storage optimization techniques.

[0026] In block 214, a determination is made as to whether the optimal container has an exposure count that is less than the exposure count limit set in the allocation policy. In the event that the exposure count of the container is less

than the exposure count limit, the new data chunk may be stored in the optimal container. In the event that the exposure count is equal to or greater than the exposure count limit, a new container may be created, and the new data chunk may be stored in the new container.

[0027] FIG. 3 illustrates a schematic representation of a computing device that may be used as a platform for storing a new data chunk in a container, according to an example of the present disclosure.

[0028] In an example, device 300 comprises a storage device or storage controller 302, memory (or machine-readable medium) 304, processing resource or CPU 306, and power source 308, all of which may be coupled by a bus or other interconnect. In some examples, device 300 may also comprise a computer readable medium, which may comprise an operating system, network applications, and/or disk storage allocation module.

[0029] Some or all of the operations set forth in the figures may be contained as machine-readable instructions as a utility, program, or subprogram in any desired computer readable storage medium, or embedded on hardware. In addition, the operations may be embodied by machine-readable instructions. For example, they may exist as machine-readable instructions in source code, object code, executable code, or other formats. The computer readable medium may also store other machine-readable instructions, including instructions downloaded from a network or the internet.

[0030] The computer-readable medium may also store machine-readable instructions that may perform basic tasks such as recognizing input from input devices, such as a keyboard or a keypad; sending output to a display; keeping track

of files and directories on a computer readable medium; and managing traffic on a bus. The network applications may include various components for establishing and maintaining network connections, such as machine readable instructions for implementing communication protocols including but not limited to TCP/IP, HTTP, HTTPS, Ethernet, USB, and FireWire.

[0031] In an example where the device of FIG. 3 is a device for allocating disk storage, device 300 may include instructions, e.g., instructions 312, to carry out the steps of FIG. 2.

[0032] In another example, instructions 312 may fetch a disk storage allocation policy, which as discussed above may include an exposure count limit for the disk. Instructions 312 may receive a request to store a new data item, such as from an operating system, disk system, disk controller, software program or application, network communication, user input, or other system, application, or protocol capable of instructing a disk to write data.

[0033] Instructions 312 may then determine a list of containers available to store the new data item and an exposure count for each of the containers. An optimal container to store the new data item may then be determined.

[0034] A new container may be created on the disk, e.g., new storage container 310, in the event that the optimal container exposure count is equal to or greater than an exposure count limit. The new data item or data chunks may then be stored in the new container.

[0035] The above discussion is meant to be illustrative of the principles and various embodiments of the present disclosure. Numerous variations and modifications will become apparent to those skilled in the art once the above

disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

CLAIMS

What is claimed is:

1. A method of allocating data storage on a computing device, comprising:

receiving, with a processor, a request to store a new data chunk in a container;

looping, with the processor, through a list of available containers;

in the event that a container comprises a data chunk from a same source as the new data chunk, fetching an exposure count for the container;

determining, with the processor, an optimal container to store the new data chunk, and

in the event that an exposure count of the optimal container is less than an exposure count limit, storing the new data chunk in the optimal container.
2. The method of claim 1, wherein the exposure count is determined by a number of files dependent on a container.
3. The method of claim 1, wherein the exposure count limit is determined by a disk drive storage allocation policy.
4. The method of claim 1, wherein determining an optimal container to store the new data chunk comprises determining a container with a lowest exposure count.

5. The method of claim 1, wherein the list of available containers is stored in an indexed database.

6. A computing device comprising:

a processor;

a memory;

a storage controller; and

a disk container exposure count module,

wherein the disk container exposure count module comprises instructions executable by the processor to allocate disk storage of a file; fetch an exposure count for a plurality of containers from a database; and determine an optimal container from the plurality of containers to store a new data chunk based on a plurality of exposure counts of the plurality of containers, and

wherein the storage controller is to store the new data chunk in the optimal container.

7. The system of claim 6, wherein the optimal container is an existing container with an exposure count less than an exposure count maximum value.

8. The system of claim 7, wherein the exposure count maximum value is equal to a maximum number of items that are permitted to depend on a single container.

9. The system of claim 7, wherein the exposure count maximum value is determined by a percentage of items on a disk that are permitted to depend on a single container.

10. The system of claim 6, wherein the optimal container is a new container.
11. A non-transitory computer readable storage medium comprising a set of instructions executable by a processor to:
 - fetch a disk storage allocation policy for a disk storage system;
 - receive a request to store a new data item in a container on the disk storage system;
 - determine a list of containers available to store the new data item and an exposure count for each of the containers;
 - determine an optimal container to store the new data item;
 - create a new container in the event that the optimal container exposure count is equal to or greater than an exposure count limit, and
 - store the new data item in the new container.
12. The non-transitory computer readable storage medium of claim 11, wherein the exposure count limit is stored in the disk storage allocation policy.
13. The non-transitory computer readable storage medium of claim 11, wherein the optimal container is the container with a lowest exposure count.
14. The non-transitory computer readable storage medium of claim 11, wherein the new data item comprises a plurality of data chunks.

15. The non-transitory computer readable storage medium of claim 14, wherein the plurality of data chunks are deduplicated.




FIG. 1




FIG. 2




FIG. 3

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2015/035576**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)

G06F 17/30; G06F 12/02; G06F 12/16; G06F 3/06; G05B 11/01; G06F 12/08

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: allocate, data chunk, container, exposure count, limit, optimal, storage, 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 2014-0122795 A1 (INTERNATIONAL BUSINESS MACHINES CORP.) 01 May 2014 See paragraphs [0004]–[0006], [0017], [0085], and [0101]; claims 1 and 17; and figures 4–5.	1–15
A	US 2012-0173838 A1 (DIETMAR NOLL et al.) 05 July 2012 See paragraphs [0005]–[0015], [0025], and [0032]–[0036]; and figure 2.	1–15
A	WO 2012-178032 A1 (COHORTFS, LLC) 27 December 2012 See paragraphs [0011]–[0012] and [0075]–[0087]; and figures 5–6.	1–15
A	US 2006-0248273 A1 (RICHARD P. JERNIGAN IV et al.) 02 November 2006 See paragraphs [0006]–[0010]; and figure 15B.	1–15
A	US 2002-0111699 A1 (BRUNO P. MELLI et al.) 15 August 2002 See paragraphs [0005] and [0028]–[0032]; and figure 4.	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 Metropolitan City, 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/US2015/035576

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014-0122795 A1	01/05/2014	CN 103793336 A US 2015-161011 A1 US 9009424 B2	14/05/2014 11/06/2015 14/04/2015
US 2012-0173838 A1	05/07/2012	US 8862833 B2 WO 2011-069705 A1	14/10/2014 16/06/2011
WO 2012-178032 A1	27/12/2012	EP 2724226 A1 US 2012-331249 A1 US 9098201 B2	30/04/2014 27/12/2012 04/08/2015
US 2006-0248273 A1	02/11/2006	EP 1875354 A1 EP 1875354 A4 EP 1875354 B1 JP 2008-541207 A JP 5291456 B2 US 7617370 B2 WO 2006-119021 A1	09/01/2008 08/04/2009 14/08/2013 20/11/2008 18/09/2013 10/11/2009 09/11/2006
US 2002-0111699 A1	15/08/2002	FR 2824160 A1 US 6738783 B2	31/10/2002 18/05/2004