



- (51) International Patent Classification:
G06F 3/06 (2006.01)
- (21) International Application Number:
PCT/US2016/054008
- (22) International Filing Date:
27 September 2016 (27.09.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/870,943 30 September 2015 (30.09.2015) US
- (71) Applicant: WESTERN DIGITAL TECHNOLOGIES, INC. [US/US]; 3355 Michelson Drive, Suite 100, Irvine, California 92612 (US).
- (72) Inventor: NEUMANN, Charles A.; 3355 Michelson Drive, Suite 100, Irvine, California 92612 (US).
- (74) Agents: TRUESDALE, Sabra-Anne et al.; Western Digital Technologies, Inc., 3355 Michelson Drive, Suite 100, Irvine, California 92612 (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, DJ, 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, KW, 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,

[Continued on next page]

(54) Title: REPLICATING DATA ACROSS DATA STORAGE DEVICES OF A LOGICAL VOLUME




FIG. 5

(57) Abstract: Systems and methods are disclosed for replicating data across data storage devices of a logical volume. A data storage system may create a logical volume by grouping the multiple data storage devices together. The data storage devices may be external standalone data storage devices. The data storage system may also replicate data across the logical volume.



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 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))

REPLICATING DATA ACROSS DATA STORAGE DEVICES OF A LOGICAL VOLUME

BACKGROUND

[0001] Distribution of data storage across one or more data storage devices can provide increased data security through data redundancy. One type of data storage device may be a direct-attached storage (DAS) device. DAS devices may provide one or more computing devices with direct access to data via a connection cable (e.g., via a direct or physical connection). Another type of data storage device may be a network-attached storage (NAS) device. NAS devices may provide access to data over computer networks (e.g., via a wired and/or wireless network).

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Various embodiments are depicted in the accompanying drawings for illustrative purposes, and should in no way be interpreted as limiting the scope of this disclosure. In addition, various features of different disclosed embodiments can be combined to form additional embodiments, which are part of this disclosure.

[0003] Figure 1 is a diagram of a data storage system, according to an embodiment.

[0004] Figure 2 is a diagram of a data storage system, according to an embodiment.

[0005] Figure 3 is a diagram of a data storage system, according to an embodiment.

[0006] Figure 4 is a diagram of an example file list, according to an embodiment.

[0007] Figure 5 is a flow diagram illustrating a process for replicating data across a logical volume, according to an embodiment.

[0008] Figure 6 is a flow diagram illustrating a process for adding a data storage device to a logical volume, according to an embodiment.

[0009] Figure 7 is a flow diagram illustrating a process for removing a data storage device from a logical volume, according to an embodiment.

[0010] Figure 8 is a flow diagram illustrating a process for temporarily removing a data storage device from a logical volume, according to an embodiment.

[0011] Figure 9 is a diagram of a computing device, according to an embodiment.

DETAILED DESCRIPTION

[0012] While certain embodiments are described, these embodiments are presented by way of example only, and are not intended to limit the scope of protection. Indeed, the novel methods and systems described herein may be embodied in a variety of other forms. Furthermore, various omissions, substitutions and changes in the form of the methods and systems described herein may be made without departing from the scope of protection.

[0013] The headings provided herein are for convenience only and do not necessarily affect the scope or meaning of the claimed invention. Disclosed herein are example implementations, configurations, and/or embodiments relating to replicating data (e.g., files) across a logical volume that may include multiple data storage devices (e.g., external standalone data storage devices).

Overview

[0014] Data storage devices/systems may provide one or more computing devices with file-level data storage. One example of a data storage device/system may be a DAS device. The DAS device may be coupled to a computing device (e.g., a laptop computer, a desktop computer, etc.) via a connection cable (e.g., a Thunderbolt cable, an external serial advanced technology attachment (eSATA) cable, a universal serial bus (USB) cable, etc.) coupled to a DAS interface (e.g., a communication interface such as USB 2.X, USB 3.X, Thunderbolt, eSATA, etc.) of the DAS device. A DAS device may provide a convenient mechanism for transferring data between computing devices, increasing the storage capacity (e.g., increase the storage space) of a computing device, and/or providing increased data security through data redundancy. Another example of a data storage device/system may be a NAS device. A NAS device may be coupled to a network via a NAS interface (e.g., a network interface or a communication interface such as Ethernet, 802.11 (Wi-Fi), etc.). A NAS device may provide file-level data storage over a network (e.g., a computer network), wherein access to the stored data is accessible to a group of clients. For example, a NAS device may include hardware, software, or a combination

of such elements, configured such that the NAS device operates as a file server. NAS devices/systems can provide a convenient mechanism for sharing data among multiple computers and/or remotely accessing data stored on the NAS devices/systems. As compared to traditional file servers, benefits of NAS devices/systems may include the ability to access data from various locations (e.g., remote locations), faster data access, easier administration, and/or simpler configuration.

[0015] Certain embodiments disclosed herein provide the ability for multiple data storage devices (e.g., external standalone data storage devices such as portable hard drives, NAS devices, etc.) to be grouped into a logical volume. The files stored on the multiple data storage devices may be replicated across the logical volume (e.g., copied onto other data storage devices in the logical volume). This may allow a user to leverage existing data storage devices to backup files and/or other data (to safeguard against data loss) more cheaply, more easily, more efficiently, and more quickly.

Data Storage Devices

[0016] Figure 1 is a diagram of a data storage system 100, according to an embodiment. In the data storage system 100, a data storage device 120 may be communicatively coupled to one or more client devices (e.g., computing devices) in order to provide file-based data storage services to the one or more client devices (e.g., one or more computing devices). Types of client devices (e.g., computing devices) that may have access to the data storage device 120 may include, but are not limited to, phones 137 (e.g., smart phones, cellular phones, etc.), cable set-top boxes 136, smart TV's 135, video game consoles 134, laptop computers 133, tablet computers 132, desktop computers 131, wearable computers and/or other types of computing devices. In one embodiment, the data storage device 120 may be an external standalone data storage device. For example, the data storage device 120 may be a data storage device that is not located within a computing device (e.g., not within a case or housing of a computing device). In another example, the data storage device 120 may be a data storage device that may provide access to data without being directly coupled to a computing device (e.g., may be a NAS device).

[0017] The data storage device 120 device may provide various client devices (e.g., phones 137, cable set-top boxes 136, smart TV's 135, video game

consoles 134, laptop computers 133, tablet computers 132, desktop computers 131) with access to various types of user data stored on the data storage device 120. The data storage device 120 may also allow users to store various types of user data on the data storage device 120. The data storage device 120 may comprise magnetic media, hard disk media, and/or solid-state media. While certain description herein refers to solid state memory generally, it is understood that solid state memory may comprise one or more of various types of solid state non-volatile memory devices such as flash integrated circuits, Chalcogenide RAM (C-RAM), Phase Change Memory (PC-RAM or PRAM), Programmable Metallization Cell RAM (PMC-RAM or PMCM), Ovonic Unified Memory (OUM), Resistance RAM (RRAM), NAND memory (e.g., single-level cell (SLC) memory, multi-level cell (MLC) memory, or any combination thereof), NOR memory, EEPROM, Ferroelectric Memory (FeRAM), Magnetoresistive RAM (MRAM), other discrete NVM (non-volatile memory) chips, or any combination thereof.

[0018] With further reference to Figure 1, the data storage device 120 (e.g., magnetic disk drive, hybrid hard drive, solid-state drive, etc.) may include a controller (not shown in Figure 1) configured to receive data commands and to execute such commands in one or more non-volatile memory components of the data storage device 120. Such commands may include data read/write commands, and the like. The controller may be configured to receive data commands from a communication interface (e.g., a NAS interface and/or a DAS interface) residing on a computing device. Data commands may specify a block address in the data storage device 120 and data may be accessed/transferred based on such commands. The data commands may also be referred to as data access requests.

[0019] The data storage device 120 may be configured to store data in one or more magnetic recording disks and/or the solid state memory devices/arrays. In an embodiment, the data storage device 120 may comprise a cable box, a backup disk drive, a media storage unit, a streaming media device, a digital camera, or any other electronic device which may store data that may need to be accessed directly or wirelessly.

[0020] In certain embodiments, the data storage device 120 may store data received from a client device such that the data storage device 120 acts as data storage for the client device. To facilitate this function, the data storage device 120 may implement a logical interface. The logical interface can present to the client device memory as a set of logical addresses (e.g., sequential/contiguous addresses)

where data may be stored. Internally, the controller may map logical addresses to various physical memory addresses in the non-volatile memory of the data storage device 120. Mapping data indicating the mapping of logical addresses to physical memory addresses may be maintained in the data storage device 120.

[0021] In one embodiment, the data storage device 120 may be a DAS device. The DAS device may be directly coupled to client device (e.g., a desktop computer 131) via a connection cable 121. The connection cable 121 may be coupled to a communication interface (e.g., a USB 2.X interface, a USB 3.X interface, a Thunderbolt interface, etc.) of the DAS device (e.g., data storage device 120). In another embodiment, the data storage device 120 may also be a NAS device. The NAS device may also be coupled to the client devices (e.g., computing devices) 131-137 via a network 105. The NAS device may be coupled to the network 105 via a network interface (e.g., an Ethernet interface, an 802.11 (Wi-Fi) interface, etc.). Each of the client devices 131-137 may also be coupled to the network 105 via a network interface. In one embodiment, network 105 may include a public network (e.g., the Internet), a private network (e.g., a local area network (LAN)), a wide area network (WAN) such as the Internet, a wired network (e.g., Ethernet network), a wireless network (e.g., an 802.11 network or a Wi-Fi network), a cellular network (e.g., a Long Term Evolution (LTE) network), routers, hubs, switches, server computers, other types of computer networks, and/or a combination thereof.

[0022] In one embodiment, the data storage device 120 may receive data access requests from the client devices 131-137 via the communication interface (e.g., a DAS interface such as USB 2.X, USB 3.X, Thunderbolt) and/or via the network interface (e.g., Ethernet, 802.11, etc.). The data access requests may be messages, commands, and/or requests to access data on the data storage device. For example, a data access request may indicate that a computing device wants to read data from the data storage device 120. In another example, a data access request may indicate that a computing device wants to write data to the data storage device 120.

[0023] Although data storage devices/drives/systems may provide various benefits, as described above, in certain configurations, such systems may be associated with various issues and/or drawbacks. Users may purchase different data storage devices (e.g., different external standalone data storage devices such as NAS devices, portable external hard drives, etc.) which may have different storage capacities, may have different storage media (e.g., magnetic media, solid state media,

etc.), and may be manufactured by different companies/vendors. The user may store different data on each of the data storage devices (e.g., one portable external hard drive may store family photos, another portable external hard drive may store work-related documents, etc.). The user may not remember which files are stored on which data storage device. Even if the user does remember which files are stored on which data storage device, files may not be backed up unless the user remembers to perform a backup of the files. Although RAID systems may automatically backup files (and other data), RAID systems are often expensive which may discourage users from purchasing and/or using RAID systems. In addition, if one of the data storage devices in a RAID system fails or becomes inoperable, the RAID system may rebuild the failed data storage device which may be a time-consuming process (e.g., may take hours, or even days). Furthermore RAID systems may not be able to use data storage devices with different storage capacities (e.g., with different storage sizes) and/or different media.

[0024] It may be useful to provide a data storage system that is able to provide backup capabilities (e.g., is able to replicate or copy data, such as files) using multiple data storage devices (e.g., multiple external standalone data storage devices) of different storage capacities, different media, and/or manufactured by different vendors/companies. This may allow users to leverage their investment in existing data storage devices (e.g., existing external portable hard drives, NAS devices, etc.). The system may create a logical volume by grouping the multiple data storage devices together. The data storage system may also replicate data (e.g., automatically replicate/copy files) across the logical volume (as discussed in more detail below). The data storage system may provide RAID-like functionality without the costs associated with a RAID system. The data storage system may also allow additional storage space (e.g., add an additional external standalone data storage device such as a portable external hard drive) to be added more cheaply. The data storage system may also maintain the replication of the data when a data storage device is removed from the logical volume without rebuilding a drive (like a RAID system).

[0025] Figure 2 is a diagram of a data storage system 200 according to an embodiment. The data storage system 200 includes a computing device 211, a computing device 213, a connection cable 250, and a data storage device 220. Each of the computing devices 211 and 213 may be a laptop computer, a desktop computer, a server computer, a tablet computer, a smart phone, a set-top box, a smart TV, a

video game console, etc. The data storage device 220 may include a controller 230, a DAS interface 241 (e.g., USB 2.X, USB 3.X, Thunderbolt, eSATA, etc.), a NAS interface 243 (e.g., Ethernet, Wi-Fi, etc.), and non-volatile storage components 245. In one embodiment, the non-volatile storage components 245 may include non-volatile magnetic media, and/or solid-state memory, such as NAND flash. The controller 230 may provide overall control for the data storage device 220. In certain embodiments, the data storage device 220 may be a hard disk drive. The non-volatile storage components 245 may include one or more disks and the data storage device 220 may further include one or more heads (not shown) actuated radially over the disk(s), and a spindle motor for rotating the disk(s). Alternatively to, or in addition to, magnetic rotating media, solid-state memory and/or other non-volatile memory such as MRAM and/or phase change memory may be used. In certain embodiments, the data storage device 220 may be, for example, a hybrid hard drive including both magnetic media and solid-state media (e.g., the non-volatile storage components 245 may include both magnetic disks and solid state media/memory). In one embodiment, the non-volatile storage components 245 may be coupled to the controller via one or more serial bus connectors. Examples of serial bus connectors include, but are not limited to, serial ATA (SATA) connectors, peripheral component interconnect express (PCIe) connectors, and SATA Express connectors. In one embodiment, the data storage device 220 may be external standalone data storage devices (e.g., NAS devices/drives, portable external hard drives, etc.).

[0026] The controller 230 may receive data access requests (e.g., data and storage access commands) from a DAS interface 212 (e.g., a USB interface, a Thunderbolt interface) of the computing device 211. Data access requests communicated by the DAS interface 212 may include write and read commands issued by the computing device 211. The data access requests may specify a LBA, or range of LBAs, in the data storage device 220, and the controller 230 may execute the received data access requests in the non-volatile storage components 245. The controller 230 may also receive data access request from a NAS interface 214 (e.g., a communication interface such as an Ethernet interface, a Wi-Fi interface, etc.) of the computing device 213. The controller may determine an LBA and/or a range of LBAs in the data storage device 220 based on the data access requests and may execute the received data access requests in the non-volatile storage components 245. In a

hybrid hard drive, data may be stored in a magnetic media storage component as well as non-volatile solid-state memory.

[0027] The data storage device 220 may store data received from the computing devices 211 and 213, such that the data storage device 220 acts as memory for the computing devices 211 and 213. To facilitate this memory function, the controller 230 may implement a logical interface. The logical interface may present to the computing devices 211 and 213 the memory of the data storage device 220 as a set of logical addresses (e.g., contiguous address) where data can be stored. The controller 230 may map logical addresses to various physical memory addresses in the non-volatile storage components 245 and/or other memory module(s).

[0028] The data storage device 220 may be configured to implement data redundancy, wherein user data stored in the non-volatile storage components 245 is maintained in one or more internal and/or external drives. For example, the controller 230 may include a data redundancy management module (not shown in Figure 2) configured to implement redundancy functionality. The data redundancy management module may implement redundant array of independent disks (RAID) technology, wherein the non-volatile storage components 245 includes a plurality of internal drives, disks, or other data storage partitions combined into a logical unit for the purposes of data redundancy and performance improvement. In addition, or alternatively, the data redundancy management module may be configured to implement RAID using one or more internal memory modules in combination with one or more external memory devices, as discussed in greater detail below.

[0029] For RAID purposes, the non-volatile storage components 245 may include an array of one or more storage devices, such as hard disks or other memory modules that are configured to store user data. In certain embodiments, such internal memory modules/disks may be each directly coupled to the controller 230 to provide a high bandwidth interface for reading and writing user data. The non-volatile storage components 245 may include one or more additional memory modules configured to store parity information.

[0030] The controller 230 may be configured to divide and replicate user data among the physical RAID memory modules, providing storage virtualization; the array may be accessed by the computing devices 211 and 213 as a single drive. Data may be distributed across the RAID memory modules/drives according to any desirable or practical RAID level, depending on the level of redundancy and/or

performance desired. For example, the data storage device 220 may be configured to implement RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, other RAID technology, or other erasure coding technology, depending on data reliability, availability, performance and/or capacity considerations or requirements.

[0031] The computing device 211 may be coupled to the data storage device 220 via connection cable 250. The connection cable 250 may directly connect the computing device 211 and the data storage device 220. The connection cable 250 may use one or more communication interfaces (such as a bus interface) and/or protocols that may allow the computing device 211 to communicate with the data storage device 220. The DAS interface 212 and 241 may be USB interfaces, Thunderbolt interfaces, serial attached SCSI (SAS), eSATA interface, etc.

[0032] In one embodiment, the connection cable 250 may include one or more data lines (e.g., one or more wires, pins, etc.) that allow the computing device 211 to communicate data with the data storage device 220. For example, the connection cable 250 may include data lines (not shown in Figure 2) that the computing device 211 may use to read data from and/or write data to the data storage device 220. The computing device 211 may communicate data to and from the data storage device using the DAS interface 212 (e.g., via the DAS interface 212). In another embodiment, the computing device 211 may provide an input voltage to the data storage device 220 and the data storage device 220 may use the input voltage to operate one or more components of the data storage device 220 (e.g., the controller 230, the non-volatile storage components 245, a motor, etc.). The connection cable 250 may include one or more voltage lines (e.g., wires, pins, etc.) that may receive the input voltage from the computing device 211 via the DAS interface 212. The one or more voltage lines may provide the input voltage (received from the computing device 211) to the data storage device 220 via the communication interface 240. In a further embodiment, the data storage device 220 may be coupled to a separate power source (e.g., may be coupled to a battery, to an AC adaptor, to a wall outlet, etc.).

[0033] In one embodiment, the connection cable 250 may include a bridge unit (not shown in Figure 2). For example, the connection cable 250 may include a USB bridge, a Thunderbolt bridge, or other type of bridge. The bridge unit may translate between two different types of communication interfaces and/or protocols. For example, if a connected storage device communicates in a first protocol, but not a

second protocol, the bridge unit may translate the second protocol to the first protocol, or vice versa.

[0034] The computing device 213 may be communicatively coupled to the data storage device 220 via a network 205 (e.g., one or more of a Wi-Fi network, a LAN, a cellular network, etc.). The computing device 213 may send data (e.g., files, images, movies, etc.) and/or data access requests to the data storage device 220 via the NAS interface 214 (e.g., a communication interface an Ethernet interface, a Wi-Fi interface, etc.) and the network 205. The data storage device 220 may receive the data and/or data access requests from the computing device 213 via the NAS interface 243.

[0035] In one embodiment, a data storage system may use multiple data storage devices (e.g., multiple data storage devices 220) to create a logical volume by grouping the multiple data storage devices together, as discussed in more detail below. The data storage system may also replicate data across the logical volume to provide RAID-like functionality without the costs associated with a RAID system, as discussed in more detail below.

Replicating Data Across a Logical Volume

[0036] Figure 3 is a diagram of a data storage system 300 according to an embodiment. The data storage system 300 includes a computing device 310, a network 305, a connector hub 306, and a logical volume 330. The network 305 may include a public network (e.g., the Internet), a private network (e.g., a local area network (LAN)), a wide area network (WAN) such as the Internet, a wired network (e.g., Ethernet network), a wireless network (e.g., an 802.11 network or a Wi-Fi network), a cellular network (e.g., a Long Term Evolution (LTE) network), routers, hubs, switches, server computers, other types of computer networks, and/or a combination thereof. The connector hub 306 may include multiple DAS interfaces (e.g., multiple connectors) that allow the connector hub 306 to be coupled to multiple data storage devices 320 and the computing device 310. For example, the connector hub 306 may be a USB hub that includes multiple USB ports/interfaces. Although the connector hub 306 is illustrated in Figure 3 as being separate from the computing device 310, it shall be understood that in other embodiments, the connector hub 306 may be part of the computing device 310. In one embodiment, the computing device 310 may be a NAS device.

[0037] The logical volume 330 includes multiple data storage devices 320. In one embodiment, the logical volume 330 may be a logical grouping or a set of data storage devices. The logical volume 330 may also be referred to as a set of data storage devices. The data storage devices 320 may be each be a NAS device and/or a DAS device. The data storage devices 320 may include different types of media. For example, the data storage devices 320 may include magnetic media, hard disk media, and/or solid-state media. In one embodiment, the data storage devices 320 may be external standalone data storage devices. For example, the data storage devices 320 may be external standalone NAS devices, or external DAS devices (e.g., portable external hard drives).

[0038] The computing device 310 includes a DAS interface 312 (e.g., a communication interface such as USB 2.X, USB 3.X, Thunderbolt, eSATA, etc.) and a NAS interface 314 (e.g., a network/communication interface such as Ethernet, 802.11 (Wi-Fi), etc.). The DAS interface 312 is coupled to the connector hub 306 which allows the computing device 310 to communicate (e.g., transmit and/or receive data) with the data storages devices 320 coupled to the connector hub 306. The NAS interface 314 is coupled to the network 305 which allows the computing device 310 to communicate (e.g., transmit and/or receive data) with the data storage device 320 coupled to the network 305. The computing device 310 also includes a storage module 311.

[0039] In one embodiment, the storage module 311 may create the logical volume 330. For example, the storage module 311 may receive input (e.g., user input received from a graphical user interface (GUI) presented/provided by the storage module 311, parameters from a configuration file, etc.) that identifies the data storage devices 320 that are to be included in the logical volume 330. The storage module 311 may group and/or aggregate the data storage devices 320 into a set of data storage devices (e.g., may associate the data storage devices 320 with the logical volume 330 or may include the data storage devices 320 in the logical volume 330). Files stored on the data storage devices 320 may be replicated (e.g., copied) onto one or more other data storage devices 320 in the logical volume 330, as discussed in more detail below.

[0040] In one embodiment, the storage module 311 may receive input (e.g., user input from a GUI presented/provided by the storage module 311, parameters from a configuration file, etc.) that indicate how much storage space on each of the data storage devices 320 should be allocated for replicating files. The storage module 311

may allocate the indicated amount of storage space on each of the data storage devices 320 for replicating files (as discussed in more detail below). For example, the storage module 311 may allocate one gigabyte of storage space on each data storage device 320 for replicating files.

[0041] In one embodiment, the storage module 311 may identify a set of files (e.g., one or more files) stored on a data storage device 320 that are to be replicated across the logical volume 330 (e.g., that are to be copied onto one or more other data storage devices 320 of the logical volume). The storage module 311 may identify the set of files based on file types. For example, the storage module 311 may identify image files (e.g., files that include digital images), video files (e.g., files that include digital videos), word processing files/documents, etc. The storage module 311 may also identify the set of files based on file sizes. For example, the storage module 311 may identify files that are greater than a threshold size (e.g., greater than one megabyte), less than a threshold size (e.g., less than one hundred megabytes), and/or between a range of sizes (e.g., between five and fifty megabytes). The storage module 311 may also identify the set of files based on the number of times the files have been accessed. For example, the storage module 311 may identify files that have been accessed more than a threshold number of times or less than a threshold number of times. The storage module 311 may also identify the set of files based on the extensions (e.g., file extensions) of the set of files. For example, the storage module 311 may identify files that have a “.jpg” extension (e.g., Joint Photographic Experts Group (JPEG) files). In another example, the storage module 311 may identify files that have a “.mpg” extension (e.g., Motion Picture Experts Group (MPEG) files). The storage module 311 may also identify the set of files based on input received from a user. For example, a user may select one or more files via a GUI (presented/provided by the storage module 311) to identify the set of files. In another example, a user may select a folder via the GUI and the storage module 311 may identify the files in the folder and subfolders of the folder.

[0042] In one embodiment, the storage module 311 may determine a protection level for the set of files. The set of files may be replicated across the logical volume 330 based on the protection level, as discussed in more detail below. In one embodiment, the protection level may be data/information indicative of a priority or an importance of the set of files. For example, a higher protection level may indicate that the set of files has a higher importance (or higher priority) and a lower protection level

may indicate that the set of files has a lower importance (or lower priority). It shall be understood that in some embodiments, different numbers of protection levels may be used (e.g., two protection levels, five protection levels, etc.). The storage module 311 may create more copies of a set of files more when the protection level is higher and may create fewer copies of the set of files less when the protection level is lower. For example, the storage module 311 may copy the set of files onto more data storage devices 320 (e.g., four data storage devices 320) when the protection level is higher and may copy the set of files onto fewer data storage devices 320 (e.g., two data storage devices 320) when the protection level is lower. In another embodiment, the protection level may be data/information indicative of a how many copies of the set of files should exist across the logical volume 330. The protection level may indicate the number of data storage devices 320 that the set of files should be replicated on (e.g., copied on) and/or may indicate the number of copies of the set of files. For example, a protection level of “3” may indicate that the set of files should be copied onto three different data storage devices 320 (e.g., there should be three copies of the set of files across the logical volume 330). In one embodiment, the storage module 311 may maintain a minimum number of copies of a file exists across the logical volume based on the protection level. For example, the storage module 311 may maintain at least two copies or at least three copies of a file on different data storage devices 320 of the logical volume 330. In one embodiment, the protection level may be applied across the logical volume 330. For example, the protection level may indicate the minimum number of copies of a file (or set of files) that the storage module 311 should maintain across all of the data storage devices 320 of the logical volume 330.

[0043] In one embodiment, a user may update the protection level for a file or a set of files. For example, the user may provide user input (via a GUI provided/present by the storage module 311) to indicate that the protection level for a file should be increased or decreased. The storage module 311 may create additional copies of the file and/or remove copies of a file based on the updated protection level. The storage module 311 may update the file list to indicate that copies of the file have been created or removed. The storage module 311 may also update the file list to indicate that the protection level for a file or set of files has been updated.

[0044] In one embodiment, different sets of files may be associated with different protection levels. For example, a user may indicate that image files (e.g., family pictures) should have a high protection level and may indicate that music files

(e.g., digital music files such as MPEG-2 Audio Layer III (MP3) files) should have a lower protection level. The storage module 311 may receive data indicative of the protection levels via a GUI (which may be provided/presented by the storage module 311).

[0045] As discussed above, the storage module 311 may replicate one or more sets of files across the logical volume 330. For example, a first file stored on a first data storage device 320 (of the logical volume 330) may be replicated (e.g., copied) onto one or more additional data storage devices 320 (of the logical volume 330) based on the protection level. In one embodiment, the store module 311 may identify the one or more additional data storage devices 320 to replicate the set of files on, based on one or more storage parameters. Examples of storage parameters may include, but are not limited to, a type of the first file (e.g., image file, video file, word processing document/file, etc.), a size of the first file, one or more health statuses of the one or more additional data storage devices 320, and a number of times the first file has been accessed. In one embodiment, the set of files stored on a data storage device may be unique across the logical volume 330 prior to replicating the set of files. For example, a first file may be originally stored on only a first data storage device 320 (e.g., an original data storage device) prior to replicating the first file. After the first file is replicated/copied onto other data storage devices, the first file may no longer be unique across the logical volume because other data storage devices may have a copy of the first file.

[0046] In one embodiment, the storage module 311 may replicate certain types of files onto different data storage devices 320. For example, the storage module 311 may identify data storage devices that have magnetic media and solid state media (e.g., flash memory). Image files may be replicated (copied) onto data storage devices 320 with magnetic media and word processing documents may be copied onto data storage devices with solid state media (e.g., flash memory). In another embodiment, the storage module 311 may replicate files with different files sizes onto different data storage devices 320. For example, files with sizes greater than a threshold size may be copied onto a first data storage device 320 and files with sizes less than a threshold size may be copied onto a second data storage device 320.

[0047] In one embodiment, the storage module 311 may identify health statuses of the data storage devices 320 (of the logical volume 330). For example, the storage module 311 may determine obtain operating information about a data storage

device 320. The storage module 311 may obtain the operating information using various methods, such as Self-Monitoring, Analysis and Reporting Technology (SMART) diagnostics. The operating information may include data/information such as how long a data storage device 320 has been in operation, how many read operations and/or write operations have been performed on a data storage device 320, a number of errors (e.g., read/write errors) that have occurred on the data storage device 320, an estimated time to failure for the data storage device 320, etc. The storage module 311 may determine a health status for the data storage device 320 based on the operating information. For example, the storage module 311 may determine (e.g., generate/calculate) metric or value (or a set of metrics/values) to indicate the health of a data storage device 320. The storage module 311 may identify data storage devices 320 that have a health statuses above (e.g., greater than) or value or set of values.

[0048] In one embodiment, the storage module 311 may create and/or update a file list to indicate the one or more data storage devices 320 where a set of files has been copied. For example, a file may be replicated on (e.g., copied to) three data storage devices 320 (of the logical volume 330). The storage module 311 may update the file list to indicate that the file has been replicated on the three data storage devices. The file list is discussed in more detail below.

[0049] In one embodiment, the storage module 311 may monitor one or more files (e.g., a set of files across the logical volume). For example, the storage module 311 may monitor files stored on a data storage device and the copies of the files that are replicated on one or more additional data storage devices 320. The storage module 311 may determine that one or more copies of a first file has been modified based on the analysis of the one or more copies of the first file.

[0050] In one embodiment, the storage module 311 may using a hash function/algorithm (such as MD5) to determine whether any copies of the first file have been modified. For example, the storage module 311 may create hashes of the one or more copies of the first file (e.g., apply the hash function/algorithm to the one or more copies of the first file). The storage module 311 may compare the hashes of the one or more copies to determine if any copies of the first file have been modified (e.g., have been updated). For example, if a first hash of one copy of the first file does not match a second hash of another copy of the first file, the storage module 311 may determine that one or more copies of the first file have been modified. In another

embodiment, the storage module 311 may compare timestamps and/or file sizes of copies of the first file to determine whether any copies of the first file have been modified. For example, the storage module 311 may determine that a first timestamp and/or a first file size of one copy of the first file is different from a second timestamp and/or a second file size of another copy of the first file.

[0051] In one embodiment, the storage module 311 may update one or more files that have been modified. The storage module 311 may update the one or more files (that have been modified) across the logical volume 330 (e.g., across all data storage devices 320 that include the one or more files). For example, the storage module 311 may determine that one or more copies of a first file has been modified (based on hashes, timestamps, and/or file sizes, as discussed above). The storage module 311 may update the first file across the logical volume 330. For example, the storage module 311 may identify the latest copy of the first file and may copy the latest copy of the first file to all the data storage devices 320 that include the first file (e.g., may copy the latest copy of the first file to all data storage devices 320 on which the first file is replicated or stored).

[0052] In one embodiment, the storage module 311 may receive a request to add an additional data storage device (e.g., a standalone DAS device such as a portable external hard drive, a standalone NAS device, etc.) to the logical volume 330. For example, a user may connect or couple the additional data storage device to the connector hub 306 and may provide user input via a GUI (which may be provided and/or presented by the storage module 311) and/or a configuration file to indicate that the additional data storage device is to be added to the logical volume 330. In one embodiment, the storage module 311 may also receive user input (e.g., via a GUI or configuration file) indicating an amount of storage space on the additional data storage device that may be used to replicate files. For example, the storage module 311 may receive user input indicating that two gigabytes of storage space on the additional data storage device may be used to replicate files. In another example, the storage module 311 may receive user input indicate that fifty percent of the total storage space/capacity of the additional data storage device may be used to replicate files.

[0053] The storage module 311 may add the additional data storage device to the logical volume 330 based on the request. For example, the storage module 311 may update the file list to indicate that the additional storage device has been added to the logical volume 330. In one embodiment, the storage module 311 may identify

files on the additional data storage device that are unique across the logical volume 330 after the additional data storage device has been added to the logical volume 330. For example, the storage module 311 may identify files that have not been replicated on (e.g., copied to) other data storage devices 320 in the logical volume 330. In another example, a user may provide user input identifying one or more files that have not be replicated to other data storage devices 320. The storage module 311 may replicate one or more of the files (that have not been replicated) based on a protection level (as discussed above). In another embodiment, the storage module 311 may replicate (e.g., copy) files from other data storage device 320 onto the additional data storage device after the additional data storage device has been added to the logical volume 330.

[0054] In one embodiment, the storage module 311 may receive a request to remove a first data storage device 320 from the logical volume 330. For example, a user may disconnect/uncouple the first data storage device from the connector hub 306 and may provide user input via a GUI (which may be provided and/or presented by the storage module 311) and/or a configuration file to indicate that the first data storage device 320 is to be removed from the logical volume 330. The storage module 311 may identify a set of files (e.g., one or more files) that are replicated on the first data storage device 320. For example, the storage module 311 may identify a set of files that were copied onto the first data storage device 320 (e.g., replicated on the first data storage device 320) from other data storage devices 320 in the logical volume 330. The storage module 311 may copy the set of files (that are replicated on the first data storage device 320) to one or more other data storage devices 320 in the logical volume 330. For example, the storage module 311 may identify one or more other data storage devices 320 based on storage parameters (as discussed above) and may copy the set of files to the one or more other data storage devices 320. This may allow the storage module 311 to replicate the set of files across the logical volume 330 based on the one or more protection levels for the set of files. For example, the storage module 311 may copy a first file from the set of files to another data storage device 320 to maintain a minimum number of copies of the first file on different data storage devices 320, based on the protection level for the first file. The storage module 311 may remove the first data storage device 320 from the logical volume 330 based on the request. For example, the storage module 311 may disassociate the first data storage device 320 from the logical volume 330 and/or may update a file list to indicate

that the first data storage device is no longer part of the logical volume 330. The storage module 311 may also remove the set of files from the first data storage device 320 (e.g., may delete the set of files that are replicated on the first data storage device 320).

[0055] In another embodiment, the storage module 311 may determine that a first data storage device 320 from logical volume 330 is inoperable (e.g., has failed, is no longer operating, etc.). For example, the storage module 311 may periodically communicate (e.g., send/receive messages, frames, data, etc.) with the first data storage device 320. The storage module 311 may determine that the first data storage device is inoperable when the storage module 311 is no longer able to periodically communicate with the first data storage device 320. In another example, the storage module 311 may receive data from the first data storage device 320 (e.g., error messages) indicating that the first data storage device 320 is inoperable. The storage module 311 may identify a set of files that were copied onto the first data storage device 320 (e.g., replicated on the first data storage device 320) from other data storage devices 320 in the logical volume 330. The storage module 311 may replicate (e.g., copy) the set of files onto other data storage devices 320 to maintain a minimum number of copies of the set of files across the logical volume 330.

[0056] In one embodiment, the storage module 311 may determine that a first data storage device 320 of the logical volume 330 is to be temporarily removed from the logical volume 330 (e.g., temporarily unplugged or uncoupled from the connector hub 306 or the network 305). For example, a user may provide user input via a GUI (which may be provided and/or presented by the storage module 311) and/or a configuration file to indicate that the first data storage device 320 is to temporarily be removed from the logical volume 330 (e.g., indicate that the first data storage device 320 will be removed from the logical volume 330 and will be re-inserted into the logical volume 330 at a later time). The storage module may identify a set of files (e.g., one or more files) that are replicated on the first data storage device 320 and may copy the set of files (that are replicated on the first data storage device 320) to one or more other data storage devices 320 in the logical volume 330, as discussed above. This may allow the storage module 311 to replicate the set of files across the logical volume 330 based on the one or more protection levels for the set of files (e.g., to maintain a minimum number of copies of the set of files), as discussed above.

[0057] In another embodiment, the storage module 311 may determine that the first data storage device 320 (that was temporarily removed from the logical volume 330) has been re-inserted into the logical volume 330. For example, a user may re-connect or re-couple the data storage device 320 to the network 305 or the connector hub 306 or the user input via a GUI (which may be provided and/or presented by the storage module 311) and/or a configuration file to indicate that the first data storage device 320 has been re-inserted into the logical volume. The storage module 311 may remove the set of files that were replicated on the first data storage device 320 because the set of files was replicated (e.g. copied) onto other data storage devices 320 of the logical volume 330 when the first data storage device 320 was temporarily removed from the logical volume 330.

[0058] In one embodiment, the storage module 311 may cause files that are stored and/or replicated on the data storage devices 320 of the logical volume 330 to be encrypted. For example, the storage module 311 may include an encryption module (e.g., software, hardware, firmware or a combination thereof) that may encrypt the files. In another example, the data storage devices 320 may include encryption modules and the storage module 311 may send messages, instructions, data, etc., to the encryption modules indicating the files that are stored and/or replicated on the data storage devices 320 are to be encrypted.

[0059] In one embodiment, the storage module 311 may monitor the health statuses of the data storage devices 320. The storage module 311 may identify one or more data storage devices 320 with health statuses that are below or less than a set of values. For example, the storage module 311 may identify data storage devices 320 with an estimated life that is below a certain value. In another example, the storage module 311 may identify a data storage device 320 where errors that have occurred on the data storage device 320 (e.g., errors have occurred while reading/writing data on the data storage device 320). The storage module 311 may copy the files on the one or more data storage devices 320 (e.g., files that were originally stored on the one or more data storage devices 320 and files that were replicated on the one or more data storage devices 320) to other data storage devices 320 in the logical volume when the health statuses of the one or more drives are below the set of values.

[0060] Figure 4 is a diagram of an example file list 400, according to an embodiment. Although the file list 400 illustrated in Figure 4 is shown in the form of a

table, it shall be understood that in other embodiments, the data/information illustrated in the file list 400 may be represented using various other formats. The file list 400 includes multiple rows. In one embodiment, each row may represent a file that is stored and/or replicated on one or more data storage devices of a logical volume.

[0061] In one embodiment, the file list 400 includes a “File Name” column, a “Metadata” column, an “Original Drive” column, and a “Copy Drives” column. The File Name column may indicate the path (which may be optional) and the name of a file represented by a row. For example, the file list 400 indicates that the file represented by the first row has a path/file name of “(path1)/file1.” The Metadata column may indicate metadata associated with a file identified in the File Name column. For example, the file list 400 indicates that the file “file1” may have a timestamp, a file size, a hash (e.g., hash value generated using a hash function/algorithm), etc. It shall be understood that the metadata illustrated in Figure 4 are merely examples and that other metadata (e.g., an owner/author of a file, security attributes such as read only, etc.) may be included in the Metadata column.

[0062] The Original Drive column may indicate where a file was originally stored before the file was replicated across the logical volume 330 (e.g., before the file was copied onto one or more other data storage devices in the logical volume 330). For example, the file list 400 indicates that the file “file1” was originally stored on a data storage device identified as “Drive 2.” The Original Drive column may include a name, a serial number, a model number, and/or other data that may be used to identify data storage devices. The Copy Drive column may indicate which data storage devices a file has been replicated on. For example, the file list 400 indicates that the file “file1” has been replicated onto (e.g., copied onto) data storage devices identified as “Drive 1” and “Drive 3.” In one embodiment, the file list 400 may also include a list of all of the data storage devices that are in the logical volume. For example, the file list 400 may include names, serial numbers, model numbers, and/or other data that may be used to identify the data storage devices that are in the logical volume.

[0063] In one embodiment, the file list 400 may be created and/or updated by a storage module, as discussed above. For example, the storage module may update the file list 400 when data storage devices are added to and/or removed from the logical volume (as discussed above). The storage module may also update the file table when one or more protection levels for one or more files are updated (as discussed above).

[0064] Figure 5 is a flow diagram illustrating a process 500 for replicating data across a logical volume, according to an embodiment. The process 500 may be performed by a controller and/or a storage module, as illustrated and discussed above in conjunction with Figures 1-3. The controller and the storage module may each be processing logic that includes hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions run on a processor to perform hardware simulation), firmware, or a combination thereof.

[0065] The process 500 begins at block 505 where the process 500 creates a logical volume. For example, the process 500 may receive input identifying data storage devices that are to be included in the logical volume and may group/aggregate the data storage devices 320 (e.g., may associate the data storage devices 320 with the logical volume 330 or may include the data storage devices 320 in the logical volume 330), as discussed above. At block 510, the process 500 may identify a set of files to be replicated across the logical volume. For example, the process 500 may receive user input (via a GUI) identifying one or more files that are to be replicate across the logical volume, as discussed above.

[0066] At block 515, the process 500 determines the protection level for the set of files. For example, the process 500 may receive user input indicating one or more protection levels for the set of files, as discussed above. At block 520, the process 500 may replicate the set of files across the logical volume. For example, the process 500 may copy the set of files onto one or more additional data storage devices of the logical volume based on the one or more protection levels. At block 525, the process 500 may monitor the set of files. For example, the process 500 may analyze hashes of the set of files and/or timestamps of the set of files. The process 500 may determine that one or more copies of a file (or multiple files) has been modified. For example, the process 500 may determine that hashes of the copies of file are different or the timestamps of the copies of the file are different, as discussed above. The process 500 may update the file across the logical volume at block 535. For example, the process 500 may replicate a latest copy of the file onto other data storage devices that include a copy of the file, as discussed above.

[0067] Figure 6 is a flow diagram illustrating a process 600 for adding a data storage device to a logical volume, according to an embodiment. The process 600 may be performed by a controller and/or a storage module, as illustrated and discussed above in conjunction with Figures 1-3. The controller and the storage

module may each be processing logic that includes hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions run on a processor to perform hardware simulation), firmware, or a combination thereof.

[0068] The process 600 begins at block 605 where the process 600 receive a request to add a data storage device to the logical volume. For example, the process 600 may receive user input via a GUI and/or may detect that a data storage device has been coupled to a connector hub or network, as discussed above. At block 601, the process 600 may add the data storage device to the logical volume. For example, the process 600 may update a file table to indicate that the data storage device is part of the logical volume, as discussed above.

[0069] Figure 7 is a flow diagram illustrating a process 700 for removing a data storage device from a logical volume, according to an embodiment. The process 700 may be performed by a controller and/or a storage module, as illustrated and discussed above in conjunction with Figures 1-3. The controller and the storage module may each be processing logic that includes hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions run on a processor to perform hardware simulation), firmware, or a combination thereof.

[0070] The process 700 begins at block 705 where the process 700 determines that a data storage device is no longer operable or where the process 700 receives a request to remove the data storage device from the logical volume. For example, the process 700 may receive one or more error messages from the data storage device, as discussed above. In another example the process 700 may receive user input indicating that a user wants to remove the data storage device from the logical volume. At block 710, the process 700 may identify a set of files that are replicated on the data storage device, as discussed above. At block 715, the process 700 may copy the set of files to one or more other data storage devices in the logical volume (e.g., to maintain a minimum number of copies of the set of files across the logical volume) based on the protection level. For example, the process 700 may identify one or more other data storage devices based on health statuses of the one or more data storage devices and may copy the set of files the one or more other data storage devices, as discussed above. At block 720, the process 700 may remove the data storage device from the logical volume. For example, the process 700 may update a file list to indicate that the data storage device is no longer part of the logical volume, as discussed above.

[0071] Figure 8 is a flow diagram illustrating a process 800 for temporarily removing a data storage device from a logical volume, according to an embodiment. The process 800 may be performed by a controller and/or a storage module, as illustrated and discussed above in conjunction with Figures 1-3. The controller and the storage module may each be processing logic that includes hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions run on a processor to perform hardware simulation), firmware, or a combination thereof.

[0072] The process 800 begins at block 805 where the process 800 determines that a data storage device is to be temporarily removed from the logical volume. For example, the process 800 may receive user input (via a GUI) identifying the data storage device that is to be temporarily removed from the logical volume. At block 810, the process 800 may identify a set of files that are replicated on the data storage device, as discussed above. At block 815, the process 800 may copy the set of files to one or more other data storage devices in the logical volume (e.g., to maintain a minimum number of copies of the set of files across the logical volume) based on the protection level. For example, the process 800 may identify one or more other data storage devices based on health statuses of the one or more data storage devices and may copy the set of files to the one or more other data storage devices, as discussed above. At block 820, the process 800 may determine that the data storage device has been re-inserted into the logical volume. For example, the process 800 determine that the data storage devices has been reconnected to a connector hub or network, as discussed above. At block 825, the process 800 may remove the set of files (e.g., delete the set of files) that were replicated on the data storage device. For example, the process 800 may remove the set of files because the set of files were copied onto one or more other data storage devices at block 815.

[0073] Figure 9 is a diagram of a computing device 900, according to an embodiment. The computing device 900 may execute instructions that may cause the computing device 900 to perform any one or more of the methodologies (e.g., operations, methods, functions, etc.) discussed herein, may be executed. The computing device 900 may be a mobile phone, a smart phone, a netbook computer, a rackmount server, a router computer, a server computer, a personal computer, a mainframe computer, a laptop computer, a tablet computer, a desktop computer etc., within which a set of instructions, for causing the machine to perform any one or more

of the methodologies discussed herein, may be executed. In alternative embodiments, the machine may be connected (e.g., networked) to other machines in a LAN, an intranet, an extranet, or the Internet. The machine may operate in the capacity of a server machine in client-server network environment. The machine may be a personal computer (PC), a set-top box (STB), a server, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while only a single machine is illustrated, the term “machine” shall also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

[0074] The example computing device 900 includes a processing device (e.g., a processor, a controller, a central processing unit (CPU), etc.) 902, a main memory 904 (e.g., read-only memory (ROM), flash memory, dynamic random access memory (DRAM) such as synchronous DRAM (SDRAM)), a network-access interface 908, a direct-access interface 909, an output device, 910, an input device 912, and a data storage device 918, which communicate with each other via a bus 930.

[0075] Processing device 902 represents one or more general-purpose processing devices such as a microprocessor, central processing unit, or the like. More particularly, the processing device 902 may be a complex instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or a processor implementing other instruction sets or processors implementing a combination of instruction sets. The processing device 902 may also be one or more special-purpose processing devices such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor (DSP), network processor, or the like. The processing device 902 is configured to execute storage module 311 for performing the operations and steps discussed herein.

[0076] The computing device 900 may include a network-access interface 908 (e.g., a network interface card, a Wi-Fi interface, etc.) which may communicate with a network (e.g., network 105 illustrated in Figure 1). The computing device may also include a direct access interface 909 (e.g., a USB interface, an eSATA interface, a Thunderbolt interface, etc.). The computing device 900 also may include an output device 910 (e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT)), and an input device 912 (e.g., a mouse, a keyboard, etc.). In one embodiment, the output

device 910 and the input device 912 may be combined into a single component or device (e.g., an LCD touch screen).

[0077] The data storage device 918 may include a computer-readable storage medium 928 on which is stored one or more sets of instructions (e.g., storage module 311) embodying any one or more of the methodologies or functions described herein. The storage module 311 may also reside, completely or at least partially, within the main memory 904 and/or within the processing device 902 during execution thereof by the computing device 900. The main memory 904 and the processing device 902 may also constitute computer-readable media. The instructions may further be transmitted or received over via the network-access interface 908 and/or direct-access interface 909.

[0078] While the computer-readable storage medium 928 is shown in an example embodiment to be a single medium, the term “computer-readable storage medium” should be taken to include a single medium or multiple media (e.g., a centralized or distributed database and/or associated caches and servers) that store the one or more sets of instructions. The term “computer-readable storage medium” shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by the machine and that cause the machine to perform any one or more of the methodologies of the present disclosure. The term “computer-readable storage medium” shall accordingly be taken to include, but not be limited to, solid-state memories, optical media and magnetic media.

Additional Embodiments

[0079] Those skilled in the art will appreciate that in some embodiments, other types of distributed data storage systems can be implemented while remaining within the scope of the present disclosure. In addition, the actual steps taken in the processes discussed herein may differ from those described or shown in the figures. Depending on the embodiment, certain of the steps described above may be removed, others may be added.

[0080] While certain embodiments have been described, these embodiments have been presented by way of example only, and are not intended to limit the scope of protection. Indeed, the novel methods and systems described herein may be embodied in a variety of other forms. Furthermore, various omissions,

substitutions and changes in the form of the methods and systems described herein may be made. The accompanying claims and their equivalents are intended to cover such forms or modifications as would fall within the scope and spirit of the protection. For example, the various components illustrated in the figures may be implemented as software and/or firmware on a processor, ASIC/FPGA, or dedicated hardware. Also, the features and attributes of the specific embodiments disclosed above may be combined in different ways to form additional embodiments, all of which fall within the scope of the present disclosure. Although the present disclosure provides certain preferred embodiments and applications, other embodiments that are apparent to those of ordinary skill in the art, including embodiments which do not provide all of the features and advantages set forth herein, are also within the scope of this disclosure. Accordingly, the scope of the present disclosure is intended to be defined only by reference to the appended claims.

[0081] The words “example” or “exemplary” are used herein to mean serving as an example, instance, or illustration. Any aspect or design described herein as “example” or “exemplary” is not necessarily to be construed as preferred or advantageous over other aspects or designs. Rather, use of the words “example” or “exemplary” is intended to present concepts in a concrete fashion. As used in this disclosure, the term “or” is intended to mean an inclusive “or” rather than an exclusive “or”. That is, unless specified otherwise, or clear from context, “X includes A or B” is intended to mean any of the natural inclusive permutations. That is, if X includes A; X includes B; or X includes both A and B, then “X includes A or B” is satisfied under any of the foregoing instances. In addition, the articles “a” and “an” as used in this disclosure and the appended claims should generally be construed to mean “one or more” unless specified otherwise or clear from context to be directed to a singular form. Moreover, use of the term “an embodiment” or “one embodiment” or “an implementation” or “one implementation” throughout is not intended to mean the same embodiment or implementation unless described as such. Furthermore, the terms “first,” “second,” “third,” “fourth,” etc., as used herein are meant as labels to distinguish among different elements and may not necessarily have an ordinal meaning according to their numerical designation.

[0082] All of the processes described above may be embodied in, and fully automated via, software code modules executed by one or more general purpose or special purpose computers or processors. The code modules may be stored on any

type of computer-readable medium or other computer storage device or collection of storage devices. Some or all of the methods may alternatively be embodied in specialized computer hardware.

WHAT IS CLAIMED IS:

1. An apparatus, comprising:
 - a communication interface coupled to a set of external standalone data storage devices;
 - a controller coupled to the communication interface, the controller configured to:
 - create a logical volume comprising the set of external standalone data storage devices;
 - identify a set of files stored to be replicated across the logical volume;
 - determine a protection level for the set of files, wherein the protection level is applied across the logical volume; and
 - replicate the set of files across the logical volume based on the protection level.
2. The apparatus of claim 1, wherein the controller is further configured to:
 - monitor the set of files across the logical volume;
 - determine that one or more copies of a first file from the set of files has been modified; and
 - update the first file across the logical volume.
3. The apparatus of claim 2, wherein the controller is configured to determine that one or more copies of a first file from the set of files has been modified by:
 - comparing one or more hashes of the one or more copies of the first file.
4. The apparatus of claim 2, wherein the controller is configured to determine that one or more copies of a first file from the set of files has been modified by comparing one or more of:
 - timestamps associated with the one or more copies of the first file, or
 - file sizes associated with the one or more copies of first.
5. The apparatus of claim 1, wherein replicating the set of files across the logical volume comprises:

identifying a first file from the set of files, wherein the first file is located on a first external standalone data storage device of the set of external standalone data storage devices;

identifying one or more additional data storage devices from the set of external standalone data storage devices based on a set of storage parameters; and

copying the first file to the one or more additional data storage devices based on the protection level.

6. The apparatus of claim 5, wherein the controller is further configured to:
update a file list to indicate that the set of files is copied on the one or more additional data storage devices.
7. The apparatus of claim 5, wherein the set of storage parameters comprises one or more of:
a type of the first file;
a size of the first file;
one or more health statuses of the one or more additional data storage devices;
one or more types of the one or more additional data storage devices; or
a number of times the first file has been accessed.
8. The apparatus of claim 1, wherein creating the logical volume comprises:
aggregating set of external standalone drives to form the logical volume.
9. The apparatus of claim 1, wherein creating the logical volume comprises:
allocating an amount of space for replicating files on each external standalone data storage device of the set of external standalone data storage devices.
10. The apparatus of claim 1, wherein the set of files is identified based on one or more of:
file types of the set of files;
file sizes of the set of files;
numbers of times the set of files have been accessed; or
extensions of the set of files.

11. The apparatus of claim 1, wherein the set of files is unique across the logical volume prior to replicating the set of files.
12. The apparatus of claim 1, wherein the controller is further configured to:
 - receive a request to add an additional external standalone data storage device to the logical volume; and
 - add the additional external standalone data storage device to the logical volume based on the request.
13. The apparatus of claim 1, wherein the controller is further configured to:
 - receive a request to remove a first external standalone data storage device from set of external standalone data storage devices;
 - identify a second set of files that are replicated on the first external standalone data storage device;
 - copy the second set of files to other external standalone data storage devices in the set of external standalone data storage devices; and
 - remove the first external standalone data storage device from the set of external standalone data storage devices.
14. The apparatus of claim 1, wherein the controller is further configured to:
 - determine that a first external standalone data storage device from set of external standalone data storage devices is inoperable;
 - identify a second set of files that are replicated on the first external standalone data storage device;
 - copy the second set of files to other external standalone data storage devices in the set of external standalone data storage devices; and
 - remove the first external standalone data storage device from the set of external standalone data storage devices.
15. The apparatus of claim 1, wherein the controller is further configured to:
 - cause the set of files to be encrypted.
16. The apparatus of claim 1, wherein the controller is further configured to:

determine that a first external standalone data storage device from set of external standalone data storage devices is to be temporarily removed from the logical volume;

identify a second set of files that are replicated on the first external standalone data storage device; and

copy the second set of files to other external standalone data storage devices in the set of external standalone data storage devices.

17. The apparatus of claim 16, wherein the controller is further configured to:

determine that the first external standalone data storage device has been re-inserted into the logical volume; and

remove the second set of files from the first external standalone data storage device.

18. The apparatus of claim 1, wherein the set of external standalone data storage devices comprises:

one or more of a direct-access storage device or a network-access storage device.

19. A method, comprising:

connecting to set of external standalone data storage devices;

creating a logical volume comprising the set of external standalone data storage devices;

identifying a set of files stored to be replicated across the logical volume;

determining a protection level for the set of files, wherein the protection level is applied across the logical volume; and

replicating the set of files across the logical volume based on the protection level.

20. A non-transitory computer-readable medium having stored thereon instructions that, when executed by a processor, cause the processor to perform operations comprising:

creating a logical volume comprising a set of external standalone data storage devices;

identifying a set of files stored to be replicated across the logical volume;

determining a protection level for the set of files, wherein the protection level is applied across the logical volume; and

replicating the set of files across the logical volume based on the protection level.

1 / 8




FIG. 1




FIG. 2




FIG. 3


400



File Name	Metadata	Original Drive	Copy Drives
(path1)/file1	timestamp, file size, hash	Drive 2	Drive 1, Drive 3
(path2)/file2	timestamp, file size, hash	Drive 3	Drive 2
(path3)/file3	timestamp, file size, hash	Drive 1	Drive 5
(path4)/file4	timestamp, file size, hash	Drive 2	Drive 4, Drive 5
(path5)/file5	timestamp, file size, hash	Drive 3	Drive 1, Drive 4
(path6)/file6	timestamp, file size, hash	Drive 5	Drive 2, Drive 3, Drive 4
...			...
(pathN)/fileN	timestamp, file size, hash	Drive 1	Drive 5

FIG. 4

5 / 8

**FIG. 5**




FIG. 7






FIG. 6

7 / 8

**FIG. 8**

8 / 8

**FIG. 9**

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/054008**A. CLASSIFICATION OF SUBJECT MATTER****G06F 3/06(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 3/06; H04L 9/08; G06F 12/16; G06F 17/00; G06F 7/00; H04L 9/00; G06F 13/28; G06F 13/00

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: storage, distribution, protection, level, replicate, copy

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 8706755 B2 (SUJAL M. PATEL et al.) 22 April 2014 See column 1, lines 63-66; column 3, lines 27-30; column 4, lines 35-64; column 5, lines 52-60; column 12, lines 21-23; column 16, lines 9-30; column 20, lines 32-35; claim 1, 5-6, 11; and figure 1.	1-12, 18-20
Y		13-17
Y	US 7844775 B2 (WILLIAM E. SNAMAN, JR. et al.) 30 November 2010 See column 1, lines 63-67; claim 11; and figure 1.	13-14, 16-17
Y	US 9020151 B1 (TREND MICRO INCORPORATED) 28 April 2015 See column 1, lines 27-41; claim 1; and figure 1.	15
A	US 8484356 B1 (FREDERICK DOUGLIS et al.) 09 July 2013 See column 9, lines 14-23; and figure 2B.	1-20
A	US 2012-0079424 A1 (RICHARD ROGERS et al.) 29 March 2012 See paragraph [0067]; and figure 9.	1-20



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

11 January 2017 (11.01.2017)

Date of mailing of the international search report

12 January 2017 (12.01.2017)

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

BYUN, Sung Cheal

Telephone No. +82-42-481-8262



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/054008

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 8706755 B2	22/04/2014	AU 2003-284888 A1	25/05/2004
		AU 2003-284888 A8	25/05/2004
		CN 101515273 A	26/08/2009
		CN 101515273 B	11/05/2011
		CN 1547714 A	17/11/2004
		CN 1547714 C	26/11/2008
		EP 1421520 A1	26/05/2004
		EP 1421520 A4	12/11/2008
		EP 2330518 A2	08/06/2011
		EP 2330518 A3	31/10/2012
		JP 04504677 B2	14/07/2010
		JP 2004-537813 A	16/12/2004
		US 2003-0033308 A1	13/02/2003
		US 2003-0135514 A1	17/07/2003
		US 2006-0277432 A1	07/12/2006
		US 2008-0021907 A1	24/01/2008
		US 2008-0243773 A1	02/10/2008
		US 2010-0235413 A1	16/09/2010
		US 2010-0257219 A1	07/10/2010
		US 7146524 B2	05/12/2006
		US 7509524 B2	24/03/2009
		US 7685126 B2	23/03/2010
		US 7743033 B2	22/06/2010
		US 7962779 B2	14/06/2011
		US 8112395 B2	07/02/2012
		WO 03-012699 A1	13/02/2003
		WO 03-012699 A8	14/08/2003
		WO 2004-040412 A2	13/05/2004
		WO 2004-040412 A3	23/09/2004
US 7844775 B2	30/11/2010	CA 2560625 A1	23/03/2007
		EP 1770508 A2	04/04/2007
		EP 1770508 A3	19/09/2007
		EP 1770508 B1	14/07/2010
		JP 04562196 B2	13/10/2010
		JP 2007-122698 A	17/05/2007
		US 2007-0073990 A1	29/03/2007
		US 2007-0083723 A1	12/04/2007
US 9020151 B1	28/04/2015	None	
US 8484356 B1	09/07/2013	US 9436292 B1	06/09/2016
US 2012-0079424 A1	29/03/2012	US 2014-181446 A1	26/06/2014
		US 2016-004760 A1	07/01/2016
		US 8700571 B2	15/04/2014
		US 9146689 B2	29/09/2015