



- (51) International Patent Classification:
G06F 3/06 (2006.01) G06F 17/30 (2006.01)
- (21) International Application Number:
PCT/US2018/033448
- (22) International Filing Date:
18 May 2018 (18.05.2018)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
15/695,686 05 September 2017 (05.09.2017) US
- (71) Applicant: WESTERN DIGITAL TECHNOLOGIES, INC. [US/US]; 5601 Great Oaks Parkway, San Jose, CA 95119 (US).
- (72) Inventors: VELUSAMY, Sivaraj; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). KANNAPPAN, Vithya; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). RAVIMOHAN, Narendhiran Chinnaanangur; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US). JAYARAMAN, Muralitharan; c/o Western Digital Technologies, Inc., 5601 Great Oaks Parkway, San Jose, CA 95119 (US).

(74) Agent: HETZ, Joseph F.; Brinks Gilson & Lione, P.O. Box 10087, Chicago, IL 60610 (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, JO, JP, KE, KG, KH, 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, 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).

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

(54) Title: STORAGE SYSTEM AND METHOD FOR EXECUTING FILE-BASED FIRMWARE COMMANDS AND COLLECTING RESPONSE DATA




FIG. 1B





FIG. 1C

(57) Abstract: A storage system and method for executing file-based firmware commands and collecting response data are provided, in one embodiment, a storage system is provided comprising a memory and a controller. The controller is configured to: receive a request from a host in communication with the storage system to write data in a file, wherein the file is identified by a file path name; determine whether the file path name matches a predetermined file path name; in response to determining that the file path name does not match the predetermined file path name, write the data in the file; and in response to determining that the file path name matches the predetermined file path name, execute a command represented by the data. Other embodiments are provided.



Storage System and Method for Executing File-Based Firmware Commands and Collecting Response Data

Background

[0001] Some storage systems are used by hosts that communicate with the storage system via an application layer, sending requests to read and write files in a particular file path in a file system directory. By communicating via the application layer, such hosts access the storage system via the file system and not via logical block addresses (LBAs). However, certain diagnostic and configuration functionality of the storage system may only be accessed via logical block addresses. As such, it may be necessary for a user to return the storage system to the manufacturer to perform that functionality.

Brief Description of the Drawings

[0002] Figure 1A is a block diagram of a non-volatile storage system of an embodiment.

[0003] Figure 1B is a block diagram illustrating a storage module of an embodiment.

[0004] Figure 1C is a block diagram illustrating a hierarchical storage system of an embodiment.

[0005] Figure 2A is a block diagram illustrating components of the controller of the non-volatile storage system illustrated in Figure 1A according to an embodiment.

[0006] Figure 2B is a block diagram illustrating components of the non-volatile memory storage system illustrated in Figure 1A according to an embodiment.

[0007] Figure 3 is a block diagram of a host and storage system of an embodiment.

[0008] Figure 4 is a flow chart of a method of an embodiment for executing a file-based firmware command.

[0009] Figure 5 is a flow chart of a method of an embodiment for collecting response data to an executed file-based firmware command.

Detailed Description

[0010] Overview

[0011] By way of introduction, the below embodiments relate to a storage system and method for executing file-based firmware commands and collecting response data. In one embodiment, a storage system is provided comprising a memory and a controller in communication with the memory. The controller is configured to: receive a request from a host in communication with the storage system to write data in a file, wherein the file is identified by a file path name; determine whether the file path name matches a predetermined file path name; in response to determining that the file path name does not match the predetermined file path name, write the data in the file; and in response to determining that the file path name matches the predetermined file path name, execute a command represented by the data.

[0012] In some embodiments, the controller is further configured to write, in the file, a response to the command after the command is executed.

[0013] In some embodiments, the controller is further configured to: receive a request from the host to read the file; and in response to receiving the request from the host to read the file, send the response to the command to the host.

[0014] In some embodiments, the command is encrypted, and wherein the controller is further configured to decrypt the command prior to executing the command.

[0015] In some embodiments, the controller is further configured to identify a logical block address associated with the file path name, and wherein the controller is configured to determine whether the file path name matches a predetermined file path name by determining whether the logical block address associated with the file path name matches a predetermined logical block address.

[0016] In some embodiments, the command comprises a diagnostic command.

[0017] In some embodiments, the command comprises a configuration command.

[0018] In some embodiments, the memory comprises a three-dimensional memory.

[0019] In some embodiments, the storage system is removably connected to the host.

[0020] In another embodiment, a method for executing a command in a storage system is provided. The method is performed in a storage system comprising a memory. The

method comprises: receiving a request from a host to write data in a file in the memory; based on a location of the file in a directory structure of the memory, determining that the data represents a command to be executed and not data to be written in the file; and executing the command.

[0021] In some embodiments, the method further comprises writing a response to the command after the command is executed.

[0022] In some embodiments, the method further comprises receiving a request from the host to read the file; and sending the response to the command to the host.

[0023] In some embodiments, the command is encrypted, and wherein the method further comprises decrypting the command prior to execution.

[0024] In some embodiments, the storage system identifies the location of the file in the directory structure based on an associated logical block address.

[0025] In some embodiments, the command comprises a diagnostic command.

[0026] In some embodiments, the command comprises a configuration command.

[0027] In some embodiments, the memory comprises a three-dimensional memory.

[0028] In some embodiments, the storage system is removably connected to the host.

[0029] In another embodiment, a storage system is provided comprising a memory; means for identifying a file-based request for executing a firmware command in the memory system; and means for executing the firmware command.

[0030] In some embodiments, the memory comprises a three-dimensional memory.

[0031] Other embodiments are possible, and each of the embodiments can be used alone or together in combination. Accordingly, various embodiments will now be described with reference to the attached drawings.

[0032] Embodiments

[0033] Storage systems suitable for use in implementing aspects of these embodiments are shown in Figures 1A-1C. Figure 1A is a block diagram illustrating a non-volatile storage system 100 according to an embodiment of the subject matter described herein. Referring to Figure 1A, non-volatile storage system 100 includes a controller 102 and non-volatile memory that may be made up of one or more non-volatile memory die 104. As used herein, the term die refers to the collection of non-volatile memory cells, and associated circuitry for managing the physical operation of those non-volatile memory

cells, that are formed on a single semiconductor substrate. Controller 102 interfaces with a host system and transmits command sequences for read, program, and erase operations to non-volatile memory die 104.

[0034] The controller 102 (which may be a non-volatile memory controller (e.g., a flash, Re-RAM, PCM, or MRAM controller)) can take the form of processing circuitry, a microprocessor or processor, and a computer-readable medium that stores computer-readable program code (e.g., firmware) executable by the (micro)processor, logic gates, switches, an application specific integrated circuit (ASIC), a programmable logic controller, and an embedded microcontroller, for example. The controller 102 can be configured with hardware and/or firmware to perform the various functions described below and shown in the flow diagrams. Also, some of the components shown as being internal to the controller can also be stored external to the controller, and other components can be used. Additionally, the phrase “operatively in communication with” could mean directly in communication with or indirectly (wired or wireless) in communication with through one or more components, which may or may not be shown or described herein.

[0035] As used herein, a non-volatile memory controller is a device that manages data stored on non-volatile memory and communicates with a host, such as a computer or electronic device. A non-volatile memory controller can have various functionality in addition to the specific functionality described herein. For example, the non-volatile memory controller can format the non-volatile memory to ensure the memory is operating properly, map out bad non-volatile memory cells, and allocate spare cells to be substituted for future failed cells. Some part of the spare cells can be used to hold firmware to operate the non-volatile memory controller and implement other features. In operation, when a host needs to read data from or write data to the non-volatile memory, it can communicate with the non-volatile memory controller. If the host provides a logical address to which data is to be read/written, the non-volatile memory controller can convert the logical address received from the host to a physical address in the non-volatile memory. (Alternatively, the host can provide the physical address.) The non-volatile memory controller can also perform various memory management functions, such as, but not limited to, wear leveling (distributing writes to avoid wearing out specific

blocks of memory that would otherwise be repeatedly written to) and garbage collection (after a block is full, moving only the valid pages of data to a new block, so the full block can be erased and reused).

[0036] Non-volatile memory die 104 may include any suitable non-volatile storage medium, including resistive random-access memory (ReRAM), magnetoresistive random-access memory (MRAM), phase-change memory (PCM), NAND flash memory cells and/or NOR flash memory cells. The memory cells can take the form of solid-state (e.g., flash) memory cells and can be one-time programmable, few-time programmable, or many-time programmable. The memory cells can also be single-level cells (SLC), multiple-level cells (MLC), triple-level cells (TLC), or use other memory cell level technologies, now known or later developed. Also, the memory cells can be fabricated in a two-dimensional or three-dimensional fashion.

[0037] The interface between controller 102 and non-volatile memory die 104 may be any suitable flash interface, such as Toggle Mode 200, 400, or 800. In one embodiment, storage system 100 may be a card based system, such as a secure digital (SD) or a micro secure digital (micro-SD) card. In an alternate embodiment, storage system 100 may be part of an embedded storage system.

[0038] Although, in the example illustrated in Figure 1A, non-volatile storage system 100 (sometimes referred to herein as a storage module) includes a single channel between controller 102 and non-volatile memory die 104, the subject matter described herein is not limited to having a single memory channel. For example, in some storage system architectures (such as the ones shown in Figures 1B and 1C), 2, 4, 8 or more memory channels may exist between the controller and the memory device, depending on controller capabilities. In any of the embodiments described herein, more than a single channel may exist between the controller and the memory die, even if a single channel is shown in the drawings.

[0039] Figure 1B illustrates a storage module 200 that includes plural non-volatile storage systems 100. As such, storage module 200 may include a storage controller 202 that interfaces with a host and with storage system 204, which includes a plurality of non-volatile storage systems 100. The interface between storage controller 202 and non-volatile storage systems 100 may be a bus interface, such as a serial advanced technology

attachment (SATA), peripheral component interface express (PCIe) interface, or dual-data-rate (DDR) interface. Storage module 200, in one embodiment, may be a solid state drive (SSD), or non-volatile dual in-line memory module (NVDIMM), such as found in server PC or portable computing devices, such as laptop computers, and tablet computers.

[0040] Figure 1C is a block diagram illustrating a hierarchical storage system. A hierarchical storage system 250 includes a plurality of storage controllers 202, each of which controls a respective storage system 204. Host systems 252 may access memories within the storage system via a bus interface. In one embodiment, the bus interface may be an NVMe or fiber channel over Ethernet (FCoE) interface. In one embodiment, the system illustrated in Figure 1C may be a rack mountable mass storage system that is accessible by multiple host computers, such as would be found in a data center or other location where mass storage is needed.

[0041] Figure 2A is a block diagram illustrating components of controller 102 in more detail. Controller 102 includes a front end module 108 that interfaces with a host, a back end module 110 that interfaces with the one or more non-volatile memory die 104, and various other modules that perform functions which will now be described in detail. A module may take the form of a packaged functional hardware unit designed for use with other components, a portion of a program code (e.g., software or firmware) executable by a (micro)processor or processing circuitry that usually performs a particular function of related functions, or a self-contained hardware or software component that interfaces with a larger system, for example. Modules of the controller 102 may include a file-system-to-LBA module 111, which is discussed in more detail below, and can be implemented in hardware or software/firmware. The file-system-to-LBA module 111 can be configured to perform the algorithms and methods discussed below and shown in the attached drawings.

[0042] Referring again to modules of the controller 102, a buffer manager/bus controller 114 manages buffers in random access memory (RAM) 116 and controls the internal bus arbitration of controller 102. A read only memory (ROM) 118 stores system boot code. Although illustrated in Figure 2A as located separately from the controller 102, in other embodiments one or both of the RAM 116 and ROM 118 may be located

within the controller. In yet other embodiments, portions of RAM and ROM may be located both within the controller 102 and outside the controller.

[0043] Front end module 108 includes a host interface 120 and a physical layer interface (PHY) 122 that provide the electrical interface with the host or next level storage controller. The choice of the type of host interface 120 can depend on the type of memory being used. Examples of host interfaces 120 include, but are not limited to, SATA, SATA Express, SAS, Fibre Channel, USB, PCIe, and NVMe. The host interface 120 typically facilitates transfer for data, control signals, and timing signals.

[0044] Back end module 110 includes an error correction controller (ECC) engine 124 that encodes the data bytes received from the host, and decodes and error corrects the data bytes read from the non-volatile memory. A command sequencer 126 generates command sequences, such as program and erase command sequences, to be transmitted to non-volatile memory die 104. A RAID (Redundant Array of Independent Drives) module 128 manages generation of RAID parity and recovery of failed data. The RAID parity may be used as an additional level of integrity protection for the data being written into the memory device 104. In some cases, the RAID module 128 may be a part of the ECC engine 124. A memory interface 130 provides the command sequences to non-volatile memory die 104 and receives status information from non-volatile memory die 104. In one embodiment, memory interface 130 may be a double data rate (DDR) interface, such as a Toggle Mode 200, 400, or 800 interface. A flash control layer 132 controls the overall operation of back end module 110.

[0045] The storage system 100 also includes other discrete components 140, such as external electrical interfaces, external RAM, resistors, capacitors, or other components that may interface with controller 102. In alternative embodiments, one or more of the physical layer interface 122, RAID module 128, media management layer 138 and buffer management/bus controller 114 are optional components that are not necessary in the controller 102.

[0046] Figure 2B is a block diagram illustrating components of non-volatile memory die 104 in more detail. Non-volatile memory die 104 includes peripheral circuitry 141 and non-volatile memory array 142. Non-volatile memory array 142 includes the non-volatile memory cells used to store data. The non-volatile memory cells may be any

suitable non-volatile memory cells, including ReRAM, MRAM, PCM, NAND flash memory cells and/or NOR flash memory cells in a two dimensional and/or three dimensional configuration. Non-volatile memory die 104 further includes a data cache 156 that caches data. Peripheral circuitry 141 includes a state machine 152 that provides status information to the controller 102.

[0047] Turning again to the drawings, Figure 3 is a block diagram of a host 50 and storage system 100 of an embodiment. For simplicity, this diagram only shows some of the possible components in the host 50 and storage system 100. Any suitable type of storage system and host can be used. For illustration purposes only, the host 50 can take the form of a mobile computing device (e.g., a mobile phone), and the storage system 100 can take the form of a memory card (e.g., a microSD card). It should be noted that these are merely examples and that other forms of a host and storage system can be used.

[0048] As shown in Figure 3, in this embodiment, the host 50 comprises a processor 54 that runs an application (“app”) 52. For example, computer-readable program code for the app 52 can be read from memory 56 and executed by the processor 54. In this embodiment, the app 52 communicates with the storage system 100 using read and write file requests. While the app 52 understands the file system structure of the memory 104 in the storage system 100, the app does not address logical block addresses (LBAs) in the memory 104. Instead, the storage system 100, which is not aware of files, uses the file-system-to-LBA module 111 in the controller 102 to convert the read/write file requests from the app 52 in the host 50 to logical block addresses, which can then be converted to physical addresses in the memory 104 by the controller 102.

[0049] As mentioned above, in situations, such as in Figure 3, where the host 50 communicates with the storage system 100 on an application layer and not via logical block addresses, the host 50 cannot access diagnostic and configuration functionality of the storage system 100, as such functionality is only accessible via logical block addresses. This may occur, for example, in certain situations when a microSD card is used with an Android host. As such, it may be necessary for a user to return the storage system 100 to the manufacturer to perform that functionality.

[0050] For example, once the storage system 100 is released to the market, there may be no way for the host 50 to perform a diagnostic function and collect debug information

from the field because the host 50 cannot execute firmware commands from the application layer. This is because LBA-level access is blocked by certain hosts. So, reporting firmware field failures and performing usage analysis may require the user to return the storage system 100 to the manufacturer, so the manufacturer can access the “raw” LBAs on the storage system 100 and run diagnostic commands.

[0051] As another example, by the host 50 not being able to execute firmware commands from the application layer, once the storage system 100 is released to the market, there may be no way for the host 50 to change firmware configurations. Configuring the storage system parameters for high performance may be important because host vendors can recommend different classes of products for different performance and endurance requirements of their applications (e.g., high performance or endurance partitioning). More specifically, different categories of storage systems may have different levels of endurance and performance. For a particular application that requires more performance or endurance, a more-expensive storage system 100 may be needed. For example, a 4k resolution video recording application with a 30MBps speed requirement would not be able to use a storage system that operates only at 4MBps or 10MBps. Similar constraints may apply for a low-endurance product category with respect to the application for which it can be used. Thus, it is desirable for the host 50 to be able to configure the storage system 100 to suit a specific application in the field

[0052] The following embodiments can be used to allow the host 50 to send application layer commands to the storage system 100, so that the storage system 100 can execute diagnostic and/or configuration commands from the field. In general, the storage system 100 has hardware and/or software to identify a file-based request from the host 50 for executing a firmware command in the memory system 100. In one embodiment, the storage system 100 receives a request from the host 50 to write data in a file identified by a file path name. The storage system 100 then determines whether the file path name matches a predetermined file path name. For example, the storage system 100 can use the file-system-to-LBA module 111 to identify a logical block address associated with the file path name and then determine whether that logical block address matches a predetermined logical block address. If the file path name does not match the predetermined file path name, the storage system 100 can treat the write request as a

normal write request and write the data in the file in memory 104 (e.g., at the physical address mapped to the determined logical block address). However, if the file path name matches the predetermined file path name, the storage system 100 recognize that the data actually represents a command to be executed and is not actually data to be written in the file. As such, the storage system 100 can execute the command.

[0053] Figure 4 is a flow chart illustrating the operation of this embodiment. Here, the app 52 running on the host 50 is responsible for interacting with the storage system 100. Of course, this is just one example, and other implementations are possible. As shown in Figure 4, in this example, the app 52 in the host 500 first creates a file with a specific name in the app-specific directory of the app (act 410). As noted above, in this embodiment, the app 52 does not have direct access the LBAs of the storage system 100, as it accesses the storage system using a file system. By creating a file in a specific path in the file system of the storage system 100, the storage system 100 can intelligently detect that the app 52 is sending the storage system 100 a particular firmware custom command and execute it. In one embodiment, the specific file path is hardcoded in the storage system 100 as the triggering file path for executing a special command (and not merely a location for storage of data). In this particular example, the triggering file is called “.debug” under the app specific directory “/Manufacturer/data/com.company.app/files/”. Of course, this is merely an example, and any desirable, suitable file path can be used. Next, the app 52 optionally encrypts the custom command (act 420). The app 52 then writes the command as data to the created file (act 430).

[0054] When the storage system 100 receives the write command, it checks to see if the file path matches the hardcoded, triggering file path (e.g., based on a location of the file in a directory structure of the memory 104) (act 440). This hardcoded path within the file system is where the storage system 100 anticipates where the diagnostic file writes will happen. For example, in one implementation, in the fresh state of the firmware, the storage system 100 monitors writes to LBA0. LBA0 either has the master boot record (MBR) or the partition boot record (PBR). LBA0 is typically written during creation of the file system image by the host 50. The MBR points to the PBR, and the PBR points to the root directory. This address is stored in the storage system 100. Writes to the root directory are monitored to see if the next leaf path is created, and the directory address of

that folder is acquired. For example, if the file path is `../Manufacturer/data/com.company.app`, the root directory is monitored for creation of the “Manufacturer” directory, and, using that, the address of that directory is acquired. Writes to this directory are monitored to determine if the next leaf path “data” directory is created or not and so on until the final destination directory’s address is acquired. Writes to this directory can be monitored to see if the special diagnostic file is created. If creation of this file is detected, the contents of this file are read. If the signature matches the diagnostic signature, then the card goes ahead and detects the opcode and translates it to specific functions like error log read, hot count data read, etc.

[0055] As discussed above, the structure of the file path in this example is that the MBR has the PBR address, the PBR has the root directory’s address, the root directory has the next child directory’s address, and so on. In case of a closed-ended write command (e.g., block count input by the host 50 in advance), the implementation can have overhead only if the command contains the already identified MBR, PBR, directory, etc. The front-end firmware can monitor the start LBA and the block count and if any of the special structures already learned during the format time are part of the command, then the RAM location of the incoming data to those structures can be parsed by the firmware to understand if the address of the child structures are modified. If the addresses of the child structures are modified, they can be updated in the memory 104 (e.g., in a special firmware structure). In case of open-ended commands, the front-end firmware may not be able to predict if the addresses of the special structures it has already remembered are going to be part of that write or not. A hardware mechanism called auto transfer can handle the transfer of data from the host interface module to the flash interface module with minimal firmware involvement. In this case, there can be several ways to minimize the overhead. For example, the firmware can input the addresses of the special structures to a hardware engine, and an interrupt mechanism can be used to wake up a component to parse the data to these special addresses to see if the next child file system structure’s address is updated or not. As another example, low-level firmware can perform these functions.

[0056] Returning to Figure 4, if the file path does not match the hardcoded, triggering file path, the storage system 100 writes the data to the file in a normal write operation (act

450). However, if the file path does match the hardcoded, triggering file path, the storage system 100 recognizes that the data is actually a custom command. If the custom command was encrypted, the storage system 100 can decrypt it (act 470). The storage system 100 then determines if the custom command exists in the storage system 100 (act 470). If the custom command does not exist, the storage system 100 writes the data in the file in a normal write operation (act 450). However, if the custom command does exist, the storage system 100 executes the command and writes the response to the command (e.g., as required per the custom command spec) in the file (act 490). For example, if the command was a diagnostic command, the storage system 100 can write the debug response to the file.

[0057] Turning again to the drawings, Figure 5 is a flow chart of a method of an embodiment for collecting response data to an executed file-based firmware command. As shown in Figure 5, after the storage system 100 acknowledges that the command was successfully executed, the app 52 can send a request to the storage system 100 to read the file (act 510). The storage system 100 responds to this request by performing a normal read operation of the file and returning the data stored therein, which, in this example, is the result of a diagnostic operation (e.g., an error log read and a hot count data read) (act 520). The app 52 then decodes the returned data and can send it to a server or another external component (e.g., after getting permission from the user) (act 530).

[0058] It should be noted that the command sent by the app 52 can take any suitable form and can be for any suitable operation. For example, through an option in the app 52, a user can create permanent or temporary high-performance partitions or folders by permanently or temporarily sacrificing capacity. In one implementation, the user can create a folder and select the high performance option and then select an amount of data to give high performance for. This can be at the cost of capacity loss based on the high performance amount of data, which can require approved from the user. If the partition is a temporary high performance partition, after its use is completed, the user can select an option to clean up the partition during idle time, thus gaining back the lost capacity. Also, an already-created partition or folder can be converted to high performance. There can be several options to create a high performance partition/folder without

compromising capacity but with endurance compromise. The endurance percentage loss can be announced to the user.

[0059] As another example, through an option in the app 52, a user can create a permanent or temporary high-endurance partition or folder. In one implementation, copied data will have more endurance, and the endurance benefit will be announced to the user. There can be several options (e.g., a high-endurance partition with capacity reduction and better performance, and a high-endurance partition without capacity reduction but compromising performance). The capacity or performance reduction percentage can be announced to the user. Also, an already-created partition or folder can be converted to high endurance, so that it will be applicable to files to be copied after that.

[0060] There are several advantages associated with these embodiments. For example, these embodiment can enable low-cost storage systems to be used for high performance/endurance applications. These embodiments also allow the host 50 to collect usage data and failure analysis in real-time without the need to return the storage system 100 to the manufacturer. Such data can be used to design better storage system in the future.

[0061] Finally, as mentioned above, any suitable type of memory can be used. Semiconductor memory devices include volatile memory devices, such as dynamic random access memory ("DRAM") or static random access memory ("SRAM") devices, non-volatile memory devices, such as resistive random access memory ("ReRAM"), electrically erasable programmable read only memory ("EEPROM"), flash memory (which can also be considered a subset of EEPROM), ferroelectric random access memory ("FRAM"), and magnetoresistive random access memory ("MRAM"), and other semiconductor elements capable of storing information. Each type of memory device may have different configurations. For example, flash memory devices may be configured in a NAND or a NOR configuration.

[0062] The memory devices can be formed from passive and/or active elements, in any combinations. By way of non-limiting example, passive semiconductor memory elements include ReRAM device elements, which in some embodiments include a resistivity switching storage element, such as an anti-fuse, phase change material, etc.,

and optionally a steering element, such as a diode, etc. Further by way of non-limiting example, active semiconductor memory elements include EEPROM and flash memory device elements, which in some embodiments include elements containing a charge storage region, such as a floating gate, conductive nanoparticles, or a charge storage dielectric material.

[0063] Multiple memory elements may be configured so that they are connected in series or so that each element is individually accessible. By way of non-limiting example, flash memory devices in a NAND configuration (NAND memory) typically contain memory elements connected in series. A NAND memory array may be configured so that the array is composed of multiple strings of memory in which a string is composed of multiple memory elements sharing a single bit line and accessed as a group. Alternatively, memory elements may be configured so that each element is individually accessible, e.g., a NOR memory array. NAND and NOR memory configurations are examples, and memory elements may be otherwise configured.

[0064] The semiconductor memory elements located within and/or over a substrate may be arranged in two or three dimensions, such as a two dimensional memory structure or a three dimensional memory structure.

[0065] In a two dimensional memory structure, the semiconductor memory elements are arranged in a single plane or a single memory device level. Typically, in a two dimensional memory structure, memory elements are arranged in a plane (e.g., in an x-z direction plane) which extends substantially parallel to a major surface of a substrate that supports the memory elements. The substrate may be a wafer over or in which the layer of the memory elements are formed or it may be a carrier substrate which is attached to the memory elements after they are formed. As a non-limiting example, the substrate may include a semiconductor such as silicon.

[0066] The memory elements may be arranged in the single memory device level in an ordered array, such as in a plurality of rows and/or columns. However, the memory elements may be arrayed in non-regular or non-orthogonal configurations. The memory elements may each have two or more electrodes or contact lines, such as bit lines and wordlines.

[0067] A three dimensional memory array is arranged so that memory elements occupy multiple planes or multiple memory device levels, thereby forming a structure in three dimensions (i.e., in the x, y and z directions, where the y direction is substantially perpendicular and the x and z directions are substantially parallel to the major surface of the substrate).

[0068] As a non-limiting example, a three dimensional memory structure may be vertically arranged as a stack of multiple two dimensional memory device levels. As another non-limiting example, a three dimensional memory array may be arranged as multiple vertical columns (e.g., columns extending substantially perpendicular to the major surface of the substrate, i.e., in the y direction) with each column having multiple memory elements in each column. The columns may be arranged in a two dimensional configuration, e.g., in an x-z plane, resulting in a three dimensional arrangement of memory elements with elements on multiple vertically stacked memory planes. Other configurations of memory elements in three dimensions can also constitute a three dimensional memory array.

[0069] By way of non-limiting example, in a three dimensional NAND memory array, the memory elements may be coupled together to form a NAND string within a single horizontal (e.g., x-z) memory device levels. Alternatively, the memory elements may be coupled together to form a vertical NAND string that traverses across multiple horizontal memory device levels. Other three dimensional configurations can be envisioned wherein some NAND strings contain memory elements in a single memory level while other strings contain memory elements which span through multiple memory levels. Three dimensional memory arrays may also be designed in a NOR configuration and in a ReRAM configuration.

[0070] Typically, in a monolithic three dimensional memory array, one or more memory device levels are formed above a single substrate. Optionally, the monolithic three dimensional memory array may also have one or more memory layers at least partially within the single substrate. As a non-limiting example, the substrate may include a semiconductor such as silicon. In a monolithic three dimensional array, the layers constituting each memory device level of the array are typically formed on the layers of the underlying memory device levels of the array. However, layers of adjacent

memory device levels of a monolithic three dimensional memory array may be shared or have intervening layers between memory device levels.

[0071] Then again, two dimensional arrays may be formed separately and then packaged together to form a non-monolithic memory device having multiple layers of memory. For example, non-monolithic stacked memories can be constructed by forming memory levels on separate substrates and then stacking the memory levels atop each other. The substrates may be thinned or removed from the memory device levels before stacking, but as the memory device levels are initially formed over separate substrates, the resulting memory arrays are not monolithic three dimensional memory arrays. Further, multiple two dimensional memory arrays or three dimensional memory arrays (monolithic or non-monolithic) may be formed on separate chips and then packaged together to form a stacked-chip memory device.

[0072] Associated circuitry is typically required for operation of the memory elements and for communication with the memory elements. As non-limiting examples, memory devices may have circuitry used for controlling and driving memory elements to accomplish functions such as programming and reading. This associated circuitry may be on the same substrate as the memory elements and/or on a separate substrate. For example, a controller for memory read-write operations may be located on a separate controller chip and/or on the same substrate as the memory elements.

[0073] One of skill in the art will recognize that this invention is not limited to the two dimensional and three dimensional structures described but cover all relevant memory structures within the spirit and scope of the invention as described herein and as understood by one of skill in the art.

[0074] It is intended that the foregoing detailed description be understood as an illustration of selected forms that the invention can take and not as a definition of the invention. It is only the following claims, including all equivalents, that are intended to define the scope of the claimed invention. Finally, it should be noted that any aspect of any of the embodiments described herein can be used alone or in combination with one another.

What is claimed is:

1. A storage system comprising:
 - a memory; and
 - a controller in communication with the memory, wherein the controller is configured to:
 - receive a request from a host in communication with the storage system to write data in a file, wherein the file is identified by a file path name;
 - determine whether the file path name matches a predetermined file path name;
 - in response to determining that the file path name does not match the predetermined file path name, write the data in the file; and
 - in response to determining that the file path name matches the predetermined file path name, execute a command represented by the data.
2. The storage system of Claim 1, wherein the controller is further configured to write, in the file, a response to the command after the command is executed.
3. The storage system of Claim 2, wherein the controller is further configured to:
 - receive a request from the host to read the file; and
 - in response to receiving the request from the host to read the file, send the response to the command to the host.

4. The storage system of Claim 1, wherein the command is encrypted, and wherein the controller is further configured to decrypt the command prior to executing the command.
5. The storage system of Claim 1, wherein the controller is further configured to identify a logical block address associated with the file path name, and wherein the controller is configured to determine whether the file path name matches a predetermined file path name by determining whether the logical block address associated with the file path name matches a predetermined logical block address.
6. The storage system of Claim 1, wherein the command comprises a diagnostic command.
7. The storage system of Claim 1, wherein the command comprises a configuration command.
8. The storage system of Claim 1, wherein the memory comprises a three-dimensional memory.
9. The storage system of Claim 1, wherein the storage system is removably connected to the host.
10. A method for executing a command in a storage system, the method comprising:

performing the following in a storage system comprising a memory:

receiving a request from a host to write data in a file in the memory;

based on a location of the file in a directory structure of the memory,

determining that the data represents a command to be executed and not data to be written in the file; and

executing the command.

11. The method of Claim 10 further comprising writing a response to the command after the command is executed.

12. The method of Claim 11 further comprising:

receiving a request from the host to read the file; and

sending the response to the command to the host.

13. The method of Claim 10, wherein the command is encrypted, and wherein the method further comprises decrypting the command prior to execution.

14. The method of Claim 10, wherein the storage system identifies the location of the file in the directory structure based on an associated logical block address.

15. The method of Claim 10, wherein the command comprises a diagnostic command.

16. The method of Claim 10, wherein the command comprises a configuration command.
17. The method of Claim 10, wherein the memory comprises a three-dimensional memory.
18. The method of Claim 10, wherein the storage system is removably connected to the host.
19. A storage system comprising:
 - a memory;
 - means for identifying a file-based request for executing a firmware command in the memory system; and
 - means for executing the firmware command.
20. The storage system of Claim 19, wherein the memory comprises a three-dimensional memory.




FIG. 1A




FIG. 1B




FIG. 1C




FIG. 2A

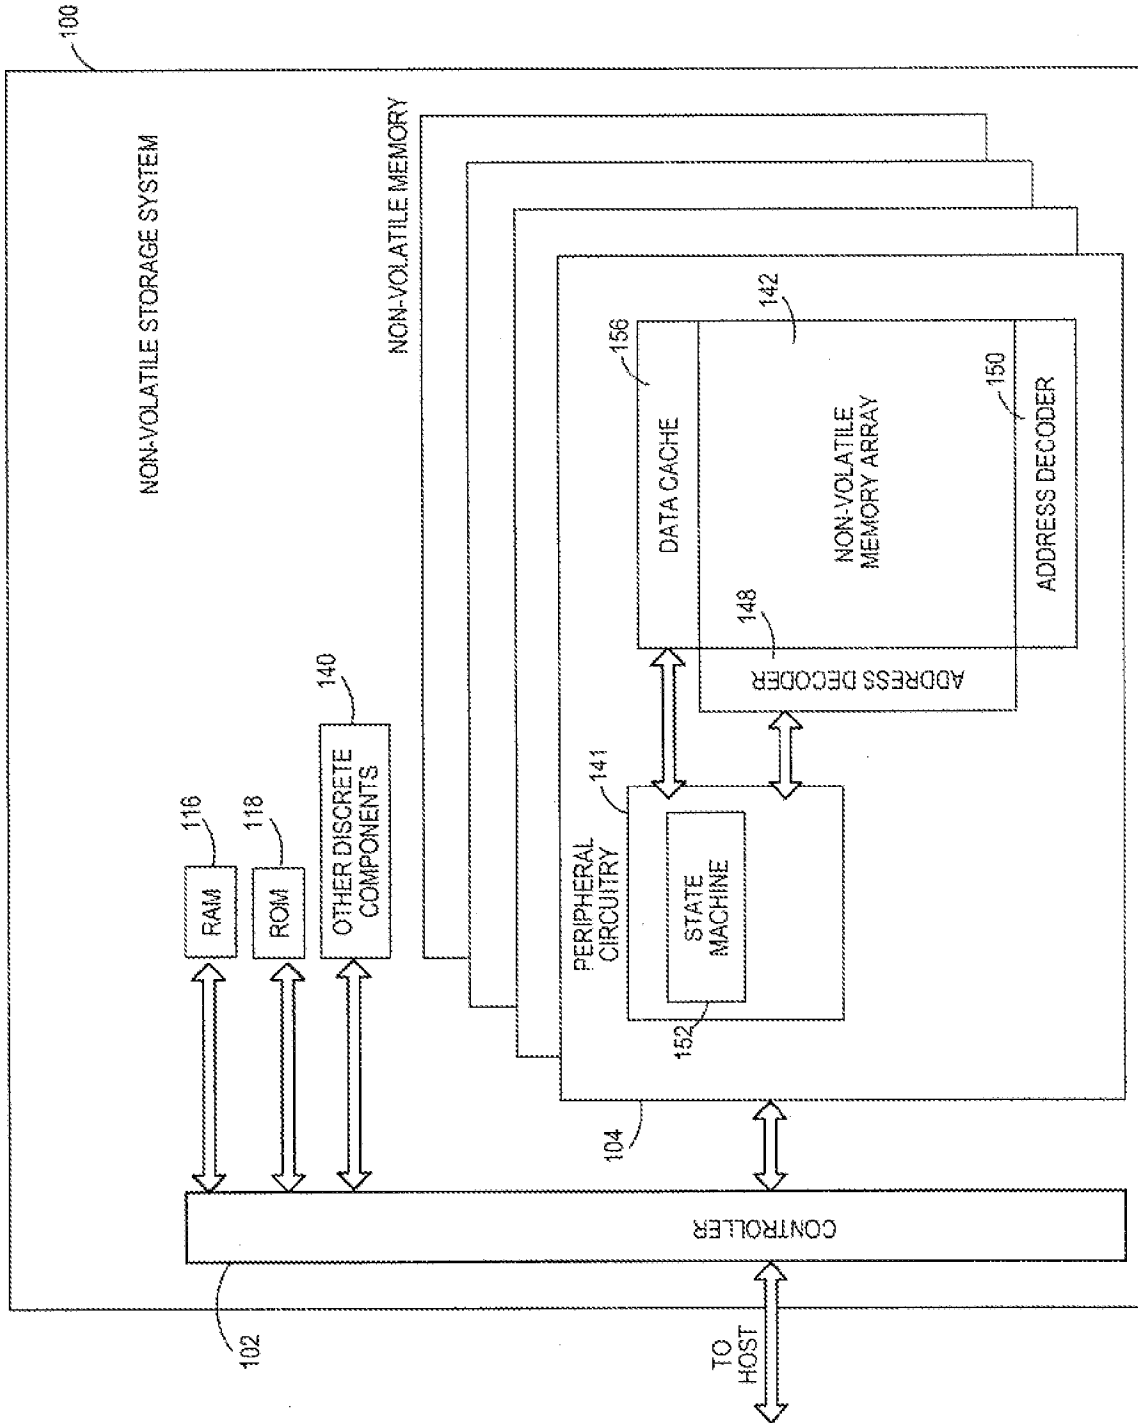


FIG. 2B




FIG. 3

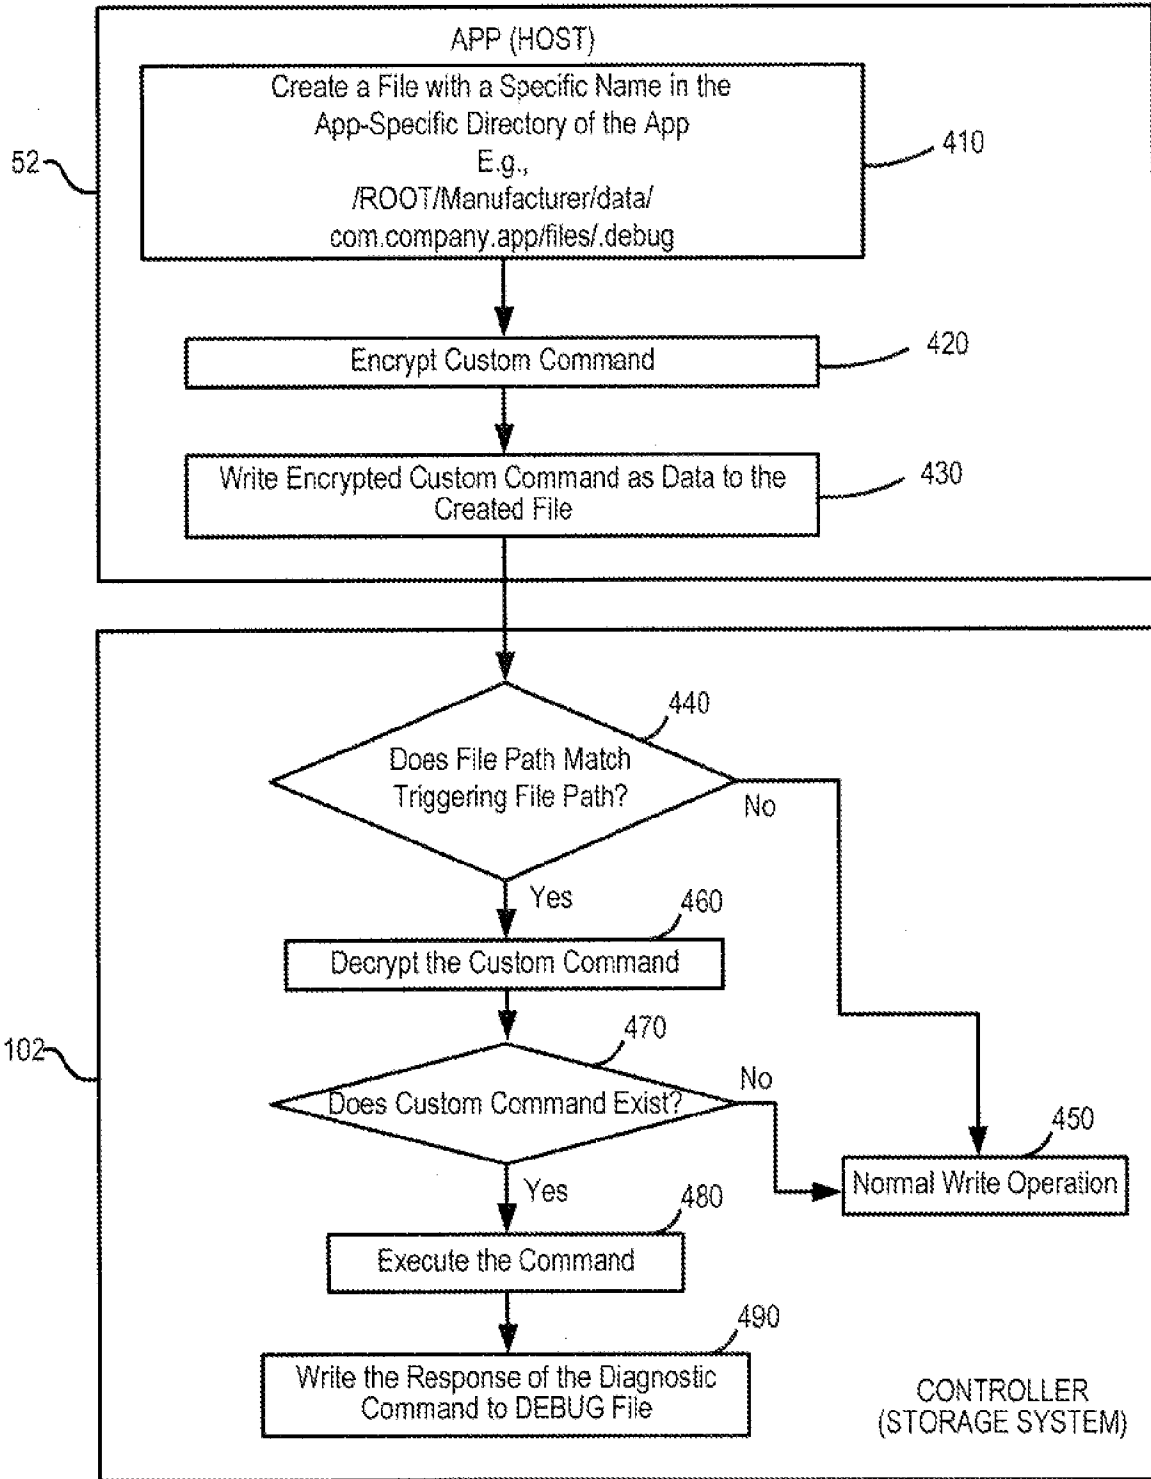


FIG. 4




FIG. 5

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/033448

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F3/06 G06F17/30
ADD.
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
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2008/126584 A1 (MULLIS SAMUEL L [US] ET AL) 29 May 2008 (2008-05-29) abstract paragraphs [0006], [0012], [0019] - paragraph [0027]; figures 1,2 -----	1-20
A	US 2007/050418 A1 (LIN FAN-SHENG [TW] ET AL) 1 March 2007 (2007-03-01) abstract paragraphs [0005] - [0007], [0026], [0028] - paragraph [0034]; figures 1-3 -----	1-20
A	US 6 493 811 B1 (BLADES JERRY A [US] ET AL) 10 December 2002 (2002-12-10) abstract column 6, line 22 - line 56 column 8, line 53 - column 10, line 44; figures 3,4 -----	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 21 August 2018	Date of mailing of the international search report 28/08/2018
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Coenen, Jean Pierre

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/033448

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2017/068686 A1 (BROIDO JACOB [IL]) 9 March 2017 (2017-03-09) abstract paragraph [0063] - paragraph [0075] -----	1-20
A	US 7 688 867 B1 (KIZHEPAT GOVIND [US]) 30 March 2010 (2010-03-30) abstract column 5, line 3 - line 11 column 6, line 5 - line 25; figure 3 -----	1-20
A	US 5 327 553 A (JEWETT DOUGLAS E [US] ET AL) 5 July 1994 (1994-07-05) abstract column 2, line 61 - column 3, line 4 column 22, line 20 - column 24, line 35 -----	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2018/033448

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2008126584	A1	29-05-2008	EP 2097811 A1 09-09-2009 US 2008126584 A1 29-05-2008 WO 2008067432 A1 05-06-2008

US 2007050418	A1	01-03-2007	TW I275944 B 11-03-2007 US 2007050418 A1 01-03-2007

US 6493811	B1	10-12-2002	US 6493811 B1 10-12-2002 US 2003033497 A1 13-02-2003

US 2017068686	A1	09-03-2017	NONE

US 7688867	B1	30-03-2010	US 7688867 B1 30-03-2010 US 2004028068 A1 12-02-2004

US 5327553	A	05-07-1994	NONE
