

METHOD AND APPARATUS FOR A NETWORK CONNECTED STORAGE SYSTEM

BACKGROUND

[0001] Data Storage Devices (DSDs) are often used to record data onto or to reproduce data from a storage media such as rotating magnetic disks or solid-state memories such as flash. DSDs are conventionally used by a host that processes or transforms data and stores data in the DSD or retrieves data from the DSD. The DSD often accesses data stored in the DSD with a file system that organizes the data into files used by an application running on the host.

[0002] The growth of distributed computing, mobile applications, social media applications, and big data applications (i.e., the collection of large amounts of data or complex data) has led to an increase in object based storage which generally allows for better scalability than file based storage. In many cases, the data objects may never be accessed again, but they are expected to remain accessible if needed. Examples of such data objects can include photos, movies, e-commerce data, or archival data stored across a network as in cloud storage. The growth of object based storage has created a need for a storage environment that can accommodate both file based storage and object based storage.

[0003] In addition, distributed computing, mobile applications, social media applications, and big data applications have led to the need for larger storage capacities to retain the data used by these applications. The traditional model of a host retrieving data from a DSD, processing the retrieved data, and storing the processed data back in the DSD may not provide the most efficient use of resources in terms of processing efficiency or network traffic when dealing with large amounts of data distributed among different DSDs. In addition, data storage systems will need to provide for more flexibility in handling data, such as allowing for the use of new interfaces and applications being developed in fields such as distributed computing.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] The features and advantages of the embodiments of the present disclosure will become more apparent from the detailed description set forth below when taken in conjunction with the drawings. The drawings and the associated descriptions are

provided to illustrate embodiments of the disclosure and not to limit the scope of what is claimed.

[0005] FIG. 1 is a block diagram of a Data Storage Device (DSD) according to an embodiment.

[0006] FIG. 2 is a block diagram of a System on a Chip (SoC) of the DSD of FIG. 1 according to an embodiment.

[0007] FIG. 3 illustrates a user space and a kernel space according to an embodiment.

[0008] FIG. 4 is a flowchart for a processor interface process according to an embodiment.

[0009] FIG. 5 is a flowchart for an application or interface installation process according to an embodiment.

[0010] FIG. 6 is a diagram depicting interactions between a host and multiple interfaces of a DSD according to an embodiment.

[0011] FIG. 7A is a flowchart for performing a command for storing data in a Non-Volatile Memory (NVM) or retrieving data from an NVM using a file interface or an object interface according to an embodiment.

[0012] FIG. 7B is a diagram depicting the receipt of commands with different transport attributes according to an embodiment.

[0013] FIG. 8A is a flowchart for a file representation process according to an embodiment.

[0014] FIG. 8B is a flowchart for a data object representation process according to an embodiment.

[0015] FIG. 9 is a flowchart for an off-load interface process to transform data according to an embodiment.

[0016] FIG. 10A is a flowchart for transforming data received from a network using an off-load interface according to an embodiment.

[0017] FIG. 10B is a flowchart for transforming data retrieved from an NVM using an off-load interface according to an embodiment.

[0018] FIG. 11 is a diagram depicting a cluster of DSDs used to transform data in series according to an embodiment.

[0019] FIG. 12 is a diagram depicting a cluster of DSDs used to perform a MapReduce function according to an embodiment.

[0020] FIG. 13 is a diagram depicting a cluster of DSDs where each DSD serves as a compute node according to an embodiment.

[0021] FIG. 14 is a block diagram of a data storage system including a plurality of Hard Disk Assemblies (HDAs) according to an embodiment.

[0022] FIG. 15 is a block diagram of a data storage system including a plurality of DSDs according to an embodiment.

DETAILED DESCRIPTION

[0023] In the following detailed description, numerous specific details are set forth to provide a full understanding of the present disclosure. It will be apparent, however, to one of ordinary skill in the art that the various embodiments disclosed may be practiced without some of these specific details. In other instances, well-known structures and techniques have not been shown in detail to avoid unnecessarily obscuring the various embodiments.

Example System Overview

[0024] FIG. 1 shows system 100 according to an embodiment which includes host 101 and Data Storage Devices (DSDs) 106 and 107 in communication with each other through network 102. System 100 can be, for example, a data storage center or other computer system that uses multiple DSDs. In addition, system 100 may be a stand-alone system or part of a network, such as network 102, which can, for example, be a local or wide area network, or the Internet. Those of ordinary skill in the art will appreciate that system 100 can include more or less than those elements shown in FIG. 1 and that the disclosed processes can be implemented in other environments. For example, in other embodiments system 100 may only include host 101 and DSD 106 without DSD 107. In such embodiments, system 100 can be, for example, a computer system (e.g., server, desktop, mobile/laptop, tablet, smartphone, etc.) or another electronic device such as a Digital Video Recorder (DVR).

[0025] In the example of FIG. 1, DSD 106 includes one or both of solid-state memory 128 and rotating magnetic disk 150 as Non-Volatile Memory (NVM) for storing data. If both are included, DSD 106 can be considered a Solid-State Hybrid Drive (SSHD) in that it includes both solid-state and disk NVM. In other embodiments,

DSDs 106 or 107 can include different types of NVM such as, for example, a magnetic tape drive, or the NVM of DSDs 106 or 107 may only disk NVM or solid-state NVM. In yet other embodiments, each of disk 150 or solid-state memory 128 may be replaced by multiple Hard Disk Drives (HDDs) or multiple Solid-State Drives (SSDs), respectively, so that DSD 106 includes pools of HDDs or SSDs. Along these lines, and as an example of a further example environment, the embodiment of FIG. 14 depicts data storage system 206 with multiple Hard Disk Assemblies (HDAs), each including at least one rotating magnetic disk and at least one head for accessing data on a corresponding disk surface.

[0026] In the example environment of FIG. 1, DSD 106 includes System on a Chip (SoC) 120 which includes circuitry for executing instructions and can include a microcontroller, a Digital Signal Processor (DSP), an Application Specific Integrated Circuit (ASIC), a Field Programmable Gate Array (FPGA), hard-wired logic, analog circuitry and/or a combination thereof. An example arrangement of SoC 120 is discussed in more detail below with reference to FIG. 2.

[0027] Network interface 126 is configured to connect DSD 106 with network 102 using, for example, an Ethernet connection or a WiFi wireless connection. Network interface 126 allows DSD 106 to interface with other devices on network 102 (e.g., host 101 or DSD 107) using a protocol such as TCP/IP. As will be appreciated by those of ordinary skill in the art, network interface 126 can be included as part of SoC 120. In other embodiments, network interface 126 may be replaced with an interface for communicating on a data bus according to a standard such as Serial Advanced Technology Attachment (SATA), PCI express (PCIe), Small Computer System Interface (SCSI), or Serial Attached SCSI (SAS).

[0028] As shown in FIG. 1, DSD 106 also includes sensor 122 for obtaining environmental information about an environmental condition of DSD 106. Sensor 122 can include one or more environmental sensors such as, for example, a mechanical shock sensor, a vibration sensor, an accelerometer (e.g., XYZ or YPR accelerometer), a temperature sensor, a humidity sensor, or an air pressure sensor. In addition, one type of sensor can be used to indicate multiple environmental conditions. For example, an accelerometer can be used to indicate both vibration and mechanical shock conditions or an air pressure sensor can be used to indicate changes in altitude

and changes in air pressure. In other embodiments, DSD 106 may obtain data from an external sensor such as a camera, a radio frequency sensor, or Radar.

[0029] In the example of FIG. 1, disk 150 is rotated by a Spindle Motor (SM) 154. DSD 106 also includes head 136 connected to the distal end of actuator 130 which is rotated by Voice Coil Motor (VCM) 132 to position head 136 in relation to disk 150. SoC 120 can control the position of head 136 and the rotation of disk 150 using VCM control signal 34 and SM control signal 38, respectively.

[0030] As appreciated by those of ordinary skill in the art, disk 150 may form part of a disk pack with additional disks radially aligned below disk 150. In addition, head 136 may form part of a head stack assembly including additional heads with each head arranged to read data from and write data to a corresponding surface of a disk in a disk pack.

[0031] Disk 150 includes a number of radial spaced, concentric tracks 152 for storing data on a surface of disk 150. Tracks 152 can be grouped together into zones of tracks with each track divided into a number of sectors that are spaced circumferentially along the tracks.

[0032] In some implementations, some or all of tracks 152 can be written by a write element of head 136 using Shingled Magnetic Recording (SMR) so as to overlap adjacent tracks. SMR provides a way of increasing the amount of data that can be stored in a given area on disk 150 by overlapping tracks like roof shingles. The non-overlapping portion then serves as a narrow track that can be read by a read element of head 136. In other implementations, all of tracks 152 may be written such that they do not overlap by using Conventional Magnetic Recording (CMR).

[0033] Although a higher number of tracks per inch is ordinarily possible with SMR, the overlap in tracks can generally prevent new writes to a previously overlapped track since such new writes would affect data written in the overlapping track. For this reason, tracks are usually sequentially written in SMR implementations to avoid affecting previously written data.

[0034] In addition to, or in lieu of disk 150, the NVM media of DSD 106 may also include solid-state memory 128 for storing data. While the description herein refers to solid-state memory generally, it is understood that solid-state memory may comprise one or more of various types of 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 chips, or any combination thereof.

[0035] Memory 140 can represent a volatile memory of DSD 106, such as Dynamic Random Access Memory (DRAM), for temporarily storing data used by SoC 120. In other embodiments, memory 140 can be an NVM such as MRAM. In addition, memory 140 can be included as part of SoC 120 in other embodiments. Those of ordinary skill in the art will also appreciate that other embodiments may include less than all of the items depicted as being stored in memory 140.

[0036] In operation, a processor of SoC 120 (e.g., processor 138 shown in FIG. 2) accesses information stored in memory 140 during execution of computer-executable instructions in software programs such as application Operating System (OS) 12, file system(s) 14, file interface(s) 16, object interface(s) 18, off-load interface(s) 20, driver(s) 21, or application(s) 22. More specifically, SoC 120 can load computer-executable instructions from an NVM of DSD 106 (e.g., disk 150 or solid-state memory 128) into a region of memory 140 and then execute the stored process instructions from memory 140.

[0037] Application OS 12 is an embedded OS of DSD 106 in the sense that application OS 12 is executed on DSD 106 and not executed on a host such as host 101. Application OS 12 manages hardware and software resources of DSD 106 and can include, for example, a Linux OS, Android OS, Windows OS, Mac OS, or a customized OS. Hardware resources managed by application OS 12 can include, for example, network interface 126, solid-state memory 128, disk 150, memory 140, and one or more processors in SoC 120 (e.g., processor 138 shown in FIG. 2). Software resources managed by application OS 12 can include, for example, file system(s) 14, file interface(s) 16, object interface(s) 18, off-load interface(s) 20, driver(s) 21, or application(s) 22.

[0038] File system(s) 14 can include one or more file systems for accessing or organizing files stored in NVM of DSD 106. By executing a file system on DSD 106, it is ordinarily possible to tailor the file system to a particular storage media used by DSD 106 to store data. In one example, file system(s) 14 can include a file system

that may be well suited to sequentially writing data on SMR media, such as Linear Tape File System (LTFS) or a log-structured file system like New Implementation of a Log-structured File System (NILFS). Other file systems of file system(s) 14 can include, for example, B-tree file system (Btrfs), ext2, ext3, ext4, or XFS. File system(s) 14 can operate in a kernel space of application OS 12, as shown in the example of the kernel space of FIG. 3.

[0039] File interface(s) 16 provide one or more file system interfaces for retrieving or storing data as files in NVM of DSD 106 via network 102. Such file interfaces can include, for example, network file systems such as Common Internet File System (CIFS), Network File System (NFS), or Server Message Block (SMB). File interface(s) 16 can operate in a kernel space of application OS 12.

[0040] Object interface(s) 18 provide an object based storage interface for retrieving or storing data as data objects in NVM of DSD 106 via network 102. Such an object interface can include, for example, Amazon Web Services Simple Storage Service (AWS S3). Object interface(s) 18 can operate in a user space as shown in FIG. 3.

[0041] Object based storage at DSD 106 can provide several advantages over file based storage such as a better use of the full capacity of the storage media and more efficient accessing of data by making the DSD more “data aware” of what is being accessed. In one example, object interface(s) 18 can provide additional information about data objects with metadata pertaining to, for example, the ownership of the data object, the creation date of the data object, or the type of data object so that certain data objects or related data can be cached (e.g., stored in memory 140) for quicker access.

[0042] Off-load interface(s) 20 provide a way of shifting processing from a host on network 102 (e.g., host 101) to DSD 106. In one implementation, host 101 sends computer-executable instructions to DSD 106 via network interface 126. A processor of SoC 120 (e.g., processor 138) executes the computer-executable instructions to create an off-load interface for accessing or transforming data stored in NVM of DSD 106. Off-load interface(s) 20 can operate in a kernel space of application OS 12 as shown in FIG. 3.

[0043] Examples of an off-load interface can include a Ceph or OpenStack software platform or a Hadoop software framework. In some implementations, a

processor of SoC 120 may partition its available storage capacity in NVM (e.g., available data capacity in disk 150 and solid-state memory 128) based on areas reserved for different interfaces. For example, a processor of SoC 120 may reserve 20% of its NVM storage capacity to AWS S3 storage, 30% to Ceph functions, and 50% to OpenStack functions. The NVM of DSD 106 could also be dynamically allocated as needed. In one example, 10% of NVM could be reserved for AWS S3 storage, 10% for Ceph functions, 10% for OpenStack functions, and 70% to be allocated among AWS S3, Ceph, and OpenStack as needed. As discussed in more detail below with reference to FIGS. 11 to 13 below, off-load interface(s) 20 can also provide for distributed compute applications that allow DSD 106 to act as a compute node in a cluster of DSDs on network 102.

[0044] Driver(s) 21 can include software for interfacing with a firmware of DSD 106 (e.g., controller firmware 11 or servo firmware 10 shown in FIG. 2, or a firmware of network interface 126) and can execute in a kernel space of application OS 12 as shown in FIG. 3. Driver(s) 21 can also allow for interfacing with other devices on network 102 such as DSD 107.

[0045] Application(s) 22 can include applications developed by a manufacturer of DSD 106 and/or independently developed applications that have been downloaded from network 102. For example, DSD 106 may receive computer-executable instructions from host 101 via network interface 126 and then execute the computer-executable instructions to create an application 22. In some implementations, a Software Development Kit (SDK) (e.g., SDK 19 shown in FIG. 3) could be made available to allow users on network 102 to develop their own applications to run on DSD 106.

[0046] Application(s) 22 or driver(s) 21 can include data storage related applications such as “sniff and store” applications to collect information about data use (e.g., identify file types, attributes, IP addresses and store traffic information), a user interface for operating DSD 106, DSD health monitoring for monitoring a reliability of DSD 106 and/or migrating data to another DSD or NVM within DSD 106 before losing data, data encryption, data compression, erasure coding or error correction, directing data for storage on disk 150 or solid-state memory 128 based on attributes of the data (e.g., tiered storage), deduplication of data stored in DSD 106, or mirroring data (e.g., data backup).

[0047] In addition, application(s) 22 or driver(s) 21 can customize DSD 106 for specific uses such as working with sensor data (e.g., where application 22 might include MATLAB or LabVIEW), streaming certain types of media over network 102, configuring DSD 106 to operate as a DVR or media server, managing the synching or backup of compute devices, providing a Bluetooth connection such as for an A/V system, providing a WiFi hotspot, or configuring DSD 106 to operate as a Network-Attached Storage (NAS). Application(s) 22 or driver(s) 21 can also provide for certain network functions on network 102 such as bit coin mining, identification of another DSD on network 102 and identification of data stored on another DSD, migration or copying of data from a failed DSD NVM to another DSD NVM in the network, or the rerouting of data to other DSDs in the network based on environmental conditions such as vibration or temperature detected by sensor 122.

[0048] In one implementation, an application 22 can allow a processor of DSD 106 (e.g., processor 138 of SoC 120 shown in FIG. 2) to monitor at least one indicator of a reliability of another DSD on network 102 (e.g., DSD 107) for storing data with the use of network interface 126. An indicator of reliability could include, for example, a number of errors in retrieving or storing data on the other DSD or an amount of use of a particular NVM of the other DSD, such as a number of Program/Erase (P/E) cycles for a flash memory of the other DSD. In some cases, the application 22 may also cause a processor of DSD 106 to copy data from another DSD to an NVM of DSD 106 if one or more indicators of reliability for the other DSD reach an unsafe or undesirable threshold.

[0049] In another implementation, an application 22 can cause a processor of DSD 106 to receive an input from sensor 122 indicating an environmental condition of DSD 106 such as a vibration condition, an air pressure condition, a humidity condition, or a temperature condition. The processor can then determine whether the input exceeds an unsafe or undesirable threshold. If the input exceeds the threshold, the processor can redirect at least one command to store data in NVM of DSD 106 to another DSD on network 102 (e.g., DSD 107). The processor may also request environmental condition information from other DSDs on network 102 to identify another DSD to receive one or more redirected commands. In addition, the processor may stop redirecting commands if a subsequent input from sensor 122 indicates that the

subsequent input has fallen below a threshold, thereby indicating that it is safe to store data in an NVM of DSD 106.

[0050] Data included in mapping 24, write pointers 26, command queue 28, buffer 30, or data to be stored in or retrieved from NVM can also be stored in memory 140 so that the data can be accessed by a processor of DSD 106 (e.g., processor 138 shown in FIG. 2) during execution of software programs to the extent that such software programs have a need to access and/or modify the data. In addition to shifting processing from host 101 to DSD 106 with off-load interface(s) 20, running application OS 12 at DSD 106 can allow for shifting tasks to file system(s) 14, driver(s) 20, and/or application(s) 22 that may otherwise have been performed by firmware (e.g., controller firmware 11 shown in FIG. 2) of DSD 106. In addition to the tasks mentioned above of data encryption, data compression, erasure coding or other error correction, the direction of data for storage on disk 150 or solid-state memory 128, deduplication of data, and data mirroring, a processor of DSD 106 (e.g., processor 138 shown in FIG. 2) executing file system(s) 14, driver(s) 21, or application(s) 22 can use mapping 24, write pointers 26, command queue 28, and buffer 30 to perform other tasks related to the management of storage media in DSD 106. By moving one or more such tasks from a firmware of DSD 106 to a file system 14, application 22, or driver 21 executed on a processor of DSD 106, the tasks can be more easily modified in the field and the firmware of DSD 106 can be simplified.

[0051] In one example, mapping 24 can map logical addresses (e.g., Logical Block Addresses (LBAs)) for data to physical addresses (e.g., block addresses) corresponding to locations in NVM of DSD 106 that store the data. This can allow for a processor executing application OS 12 to direct data to a particular NVM (e.g., disk 150 or solid-state memory 128) or particular zones within an NVM based on attributes of the data. Mapping 24 may also be stored in an NVM of DSD 106 such as disk 150 or solid-state memory 128 so that mapping 24 is available after DSD 106 has been powered off. Alternatively, memory 140 may be an NVM such as MRAM.

[0052] In some cases, a processor of DSD 106 executing application OS 12 may perform address indirection using mapping 24 to make better use of storage media such as zones of overlapping tracks (i.e., SMR zones) on disk 150 or cells (not shown) in solid-state memory 128. In the case of solid-state memory 128, the processor may use a mapping indirection system to facilitate the rewriting of data

associated with a particular logical address at different physical locations on the storage media to provide for wear leveling so that solid-state memory 128 is more evenly used to prolong a usable life of solid-state memory 128.

[0053] In the case of SMR tracks on disk 150, a processor executing application OS 12 may use indirection to facilitate the rewriting of data associated with a particular logical address at different physical locations on disk 150 while maintaining a sequential writing. When data is updated for a particular logical address, the update is often written in a different location on disk 150 than where the data for the logical address was previously written to avoid having to rewrite an entire SMR zone of overlapping tracks. Mapping 24 can be used to keep track of where the current versions of the data are stored for a particular logical address.

[0054] Write pointers 26 can be used by a processor executing application OS 12 to keep track of a location for performing a next write in a sequentially written zone such as an SMR zone on disk 150. Write pointers 26 may also be stored in an NVM of DSD 106 such as disk 150 or solid-state memory 128 so that write pointers 26 are available after DSD 106 has been powered off.

[0055] In another example, a processor executing application OS 12 may reorder the performance of pending write commands in command queue 28 so that the data to be written in performing the pending write commands follows an increasing or decreasing order of logical addresses to be written. In one example, the processor may group the order of performance of commands by host so that all of the commands from a particular host are performed in a sequential manner. Such an ordering of commands can allow for an improved overall performance in processing the commands. In addition, the processor can manage command queue 28 and mapping 24 so that all of the write commands from a particular host are written in a localized area of NVM such as in a particular zone of tracks on disk 150. The processor may also reorder the performance of pending read and write commands for disk 150 as part of a Rotational Position Optimization (RPO) algorithm to reduce the amount of movement needed in positioning head 136 over disk 150 in performing the commands.

[0056] Buffer 30 may be used by a processor executing application OS 12 in performing Read-Modify-Write (RMW) operations on disk 150, such as garbage collection to reclaim portions of disk 150 storing invalid or obsolete data. In such an

operation, the processor may read a zone of tracks and store the valid data from the zone in buffer 30 before rewriting the valid data on disk 150. In addition, the processor may identify or prioritize a particular zone for garbage collection by determining a level of invalid or obsolete data stored in the zone and/or a frequency of use of the zone.

Embedded Application OS Example Environments

[0057] FIG. 2 is a block diagram depicting SoC 120 of DSD 106 according to an embodiment. In FIG. 2, SoC 120 includes processors 138, 141, and 142. In other embodiments, one or more of processors 138, 141, and 142 may be located on different chips such that not all of processors 138, 141, and 142 are located in SoC 120. In addition, processors 141 and 142 may be combined into one processor or processors 141 and 142 may serve as co-processors of processor 138 in other embodiments.

[0058] In one implementation, each of processors 138, 141, and 142 is a processor core such as an ARM M3 processor. In another implementation, processor 138 can include an ARM A5 or A6 processor while processors 141 and 142 can be ARM M3 processors. In yet other implementations, different types of processors can be used such as Intel processors.

[0059] As shown in FIG. 2, processor 138 executes application OS 12, file system(s) 14, file interface(s) 16, object interface(s) 18, off-load interface(s) 20, driver(s) 21, and application(s) 22. Processor 138 interfaces with memory 140 in executing such software and interfaces with network interface 126 in communicating via network 102.

[0060] As discussed above, having application OS 12 embedded or running on DSD 106 can provide several advantages over conventional DSDs that do not locally execute an embedded application OS. Such advantages include the ability to support TCP/IP over Ethernet or WiFi interfaces (e.g., via network interface 126), the ability to embed a file system (e.g., file system(s) 14) that is better adapted to a particular storage media of DSD 106, and to allow for new applications (e.g., application(s) 22) to be developed for different uses of DSD 106. As will be appreciated by those of ordinary skill in the art, one or both of processors 141 and 142 may run a Real-Time Operating System (RTOS) that is intended to perform real-time processes for servo

firmware 10 or controller firmware 11. In contrast, processor 138 runs application OS 12 which allows for the execution of file system(s) 14, file interface(s) 16, object interface(s) 18, off-load interface(s) 20, driver(s) 21, and/or application(s) 22 as discussed in more detail below.

[0061] In storing or retrieving data from NVM of DSD 106, processor 138 executes application OS 12 to interface with processor 141, which executes controller firmware 11. Controller firmware 11 controls operation of the NVM of DSD 106, and may be stored in a dedicated memory of SoC 120 (e.g., a flash memory not shown) or may be stored on other NVM of DSD 106 such as disk 150 or solid-state memory 128.

[0062] As noted above, the use of an application OS at processor 138 can allow for a simplified firmware of DSD 106. In more detail, many of the tasks conventionally performed by executing DSD firmware may be shifted to software executed by processor 138. As a result, controller firmware 11 in some implementations may primarily serve only to store or retrieve data in NVM with many of the maintenance operations for the NVM being performed by file system(s) 14, driver(s) 21, and/or application(s) 22. Tasks that may be shifted to processor 138 can include, for example, data encryption, data compression, erasure coding or other error correction, data deduplication, data mirroring, the direction of data for storage on disk 150 or solid-state memory 128 based on attributes of the data, the direction of data for storage in a CMR zone (i.e., a zone of non-overlapping tracks) or an SMR zone (i.e., a zone of overlapping tracks) of disk 150 based on attributes of the data, address mapping with mapping 24, maintenance of write pointers 26, ordering of command queue 28, or garbage collection using buffer 30.

[0063] In operation, processor 138 can execute application OS 12 to interface with processor 141 and send a command to processor 141 to retrieve data from or store data in disk 150 or solid-state memory 128. The interface between processor 138 and processor 141 can be object based, use a standard such as SAS or SATA, or be a custom interface. In the case of an object based interface, processor 138 can use application OS 12 to execute object interface 18 to send a command to retrieve, store, or delete particular data objects stored in disk 150 or solid-state memory 128. In the case of using a standard such as SAS or SATA, processor 138 can use a file system 14, a file interface 16, or a driver 21 to send read, write, or trim commands for particular LBAs associated with the data. In the case of a custom interface, a

manufacturer may provide a customized file system 14 or a driver 21 to send commands to processor 141.

[0064] If access to disk 150 is needed, processor 141 communicates with processor 142, which executes servo firmware 10. In this regard, processor 142 controls SM 154 via SM control signal 38 to rotate disk 150. Processor 142 also controls VCM 132 via VCM control signal 34 to position head 136 over disk 150.

[0065] FIG. 3 depicts a user space and a kernel space of DSD 106 according to an embodiment. Application OS 12 can segregate the memory used for the applications and libraries shown in the user space from the memory used for the drivers, kernel, and file system shown in the kernel space to provide for better fault tolerance and security of the software executed by processor 138. As noted above, application OS 12 can include various types of OS such as a Linux OS, Android OS, Windows OS, Mac OS, or a customized embedded OS. Accordingly, the example OS environment of FIG. 3 can vary in other embodiments and can include different components than those shown in FIG. 3.

[0066] As shown in FIG. 3, object interface(s) 18 and application(s) 22 can run in a user space, and can access libraries 44 and 46 in the user space to make system calls to a kernel of application OS 12 (e.g., kernel 37). Libraries 44 and 46 can include, for example, specific OS libraries such as Android's SQLite, or protocols such as Google Buffer Protocols, or AWS protocol.

[0067] Hardware drivers such as network interface driver 39, HDD driver 40, and Solid-State Memory (SSM) driver 42 can operate in the kernel space. Off-load interface(s) 20, file interface(s) 16, and file system(s) 14 can also operate in the kernel space. In other implementations, a driver 21 or file system 14 may be treated like application(s) 22 and run in user space instead of kernel space (e.g., a customized DSD driver or a File System in User Space (FUSE)).

[0068] FIG. 3 also depicts SDK 19 which can be used outside of DSD 106 to create new applications 22 for DSD 106. In one implementation, SDK 19 can include an Application Programming Interface (API) to allow access to an application framework and certain libraries of the user space so that a developer can create a new application 22 to operate on processor 138.

[0069] FIG. 4 is a flowchart for a processor interface process that can be performed by processor 138 executing application OS 12 according to an

embodiment. In block 402, processor 138 executes application OS 12 to interface with processor 141 executing firmware 11.

[0070] In block 404, processor 138 uses application OS 12 to send a command to processor 141 to retrieve data from NVM (e.g., disk 150 or solid-state memory 128) or to store data in NVM. As discussed above, the command to processor 141 can be object based in specifying device block addresses, or the command to processor 241 can specify LBAs for the data with a standard such as SAS or SATA. In this regard, application OS 12 may execute a file interface 16, an object interface 18, or a driver 21 such as HDD driver 40 or SSM driver 42 to retrieve or store data in the NVM.

[0071] FIG. 5 is a flowchart for an application or interface installation process that can be performed by processor 138 according to an embodiment. In block 502, processor 138 receives computer-executable instructions via network interface 126 from a device on network 102 such as host 101.

[0072] In block 504, processor 138 executes the computer-executable instructions received in block 502 to install an application 22 or install an interface such as an object interface 18 or off-load interface 20. As discussed above, application(s) 22 can include applications developed by a manufacturer of DSD 106 and/or applications developed by others. In some implementations, an SDK (e.g., SDK 19 shown in FIG. 3) could be made available to allow users on network 102 to develop their own applications 22 to run on DSD 106 with access to certain libraries such as library 44 or 46.

Multiple Interface Examples

[0073] FIG. 6 is a diagram depicting interactions between host 101 and different software interfaces executed by processor 138 according to an embodiment. As shown in FIG. 6, host 101 can communicate with file interface(s) 16, object interface(s) 18, or off-load interface(s) 20 of DSD 106 via network 102. In other embodiments, host 101 may only communicate with one type of interface of DSD 106 while another host on network 102 may communicate with a different type of interface of DSD 106. This can ordinarily allow for greater flexibility in the types of devices on network 102 or the software that can be executed by devices on network 102.

[0074] In the example shown in FIG. 6, host applications 48 can interface with DSD 106 via host file system 50 and NFS/SMB mount 52 or through client library 56 with or

without FUSE 54. When host 101 interfaces through NFS/SMB mount 52, a file interface 16 such as CIFS/SMB mounts to host file system 50, which can send commands from processor 138 to processor 141 of DSD 106 to retrieve data or store data in NVM such as disk 150 or solid-state memory 128.

[0075] When host 101 interfaces with DSD 106 through client library 56, an object interface 18 such as AWS S3 or an off-load interface 20 such as Ceph, Hadoop, or OpenStack executing in processor 138 can perform object commands such as Create, Read, Update, and Delete (CRUD) commands on objects stored in NVM of DSD 106. In this regard, the object interface 18 or the off-load interface 20 may send a read or write command to processor 141 to access the data stored in NVM.

[0076] Each of the example interfaces executed at DSD 106 (i.e., CFS/NFS, AWS S3, Ceph, Hadoop, OpenStack) is implemented by application OS 12 within a software container 58 such as a docker in a Linux OS to provide an additional layer of abstraction and resource isolation. This can be accomplished, for example, with different namespaces that allow for different software containers to operate in a single instance while generally avoiding the overhead associated with running a Virtual Machine (VM).

[0077] FIG. 6 also depicts an example interaction between a file interface 16 and an off-load interface 20 with an arrow from file interface(s) 16 to off-load interface(s) 20. In such an example, data received from NFS/SMB mount 52 may include an executable image as part of a certain directory known to a file interface 16 to include data to be handled by an off-load interface 20. If so, the file interface 16 can pass the executable image to the off-load interface 20 for further handling or processing.

[0078] FIG. 7 is a flowchart for performing a command for storing data in NVM (e.g., disk 150 or solid-state memory 128) of DSD 106 or retrieving data from NVM of DSD 106 using a file interface or an object interface according to an embodiment. The process of FIG. 7 can be performed by processor 138 using application OS 12 to determine whether to execute a file interface 16 or an object interface 18 based on a command received via network interface 126.

[0079] In block 702, processor 138 receives a command for storing data in NVM or retrieving data from NVM. The data is received by processor 138 from network 102 via network interface 126 and may come from a host or client (e.g., host 101) on network 102 or from another DSD (e.g., DSD 107) on network 102.

[0080] In block 704, processor 138 determines whether to execute a file interface 16 or an object interface 18 based on a transport attribute of the command received in block 702. FIG. 7B provides an example of the receipt of commands with different transport attributes. In the example of FIG. 7B, DSD 106 receives multiple commands (i.e., commands 1 to N) from network 102. Each of the commands are directed to an appropriate interface such as a file interface 16, an object interface 18, or an off-load interface 20 based on a socket (i.e., sockets 1 to N) that receives the command. In one implementation, an interface (e.g., file interface 16, object interface 18, or off-load interface 20) may create a socket and a service (e.g. services 1, 2, or 3) of application OS 12 routes an incoming command to the appropriate interface based on a socket address included with the command that may specify a protocol, local address, local port, remote address, and remote port. In such an implementation, the socket address of the command can serve as a transport attribute that can be used by application OS 12 to determine whether to use a file interface 16, an object interface 18, or an off-load interface 20.

[0081] Returning to FIG. 7A, processor 138 in block 706 performs the received command using the interface determined in block 704. In the case where a file interface 16 is used, processor 138 can interface with processor 141 using a standard such as SATA or SAS. In the case where an object interface 18 is used, processor 138 can either provide processor 141 with LBAs for the data object or can provide device block addresses for the data object without specifying LBAs. Processor 138 may also use a custom interface with either of file interface 16 or object interface 18.

[0082] FIG. 8A is a flowchart for a file representation process that can be performed by processor 138 according to an embodiment. In block 802, processor 138 receives a first data object and a data second object via network interface 126. An object interface 18 can be used by processor 138 to store the first data object and the second data object in NVM of DSD 106 as described above for block 706 of FIG. 7.

[0083] In block 804, processor 138 executes a file interface 16 to represent data from the first and second data objects as a file. In one example, the file may be represented to a device on network 102 based on a request received from the device to retrieve data stored in NVM of DSD 106. In this regard, different devices on network 102 may interface with DSD 106 using different interfaces of DSD 106. For

example, one host on network 102 executing a particular file system may retrieve and store data on DSD 106 by specifying LBAs associated with the data in the host's file system while another host on network 102 may retrieve and store data on DSD 106 by specifying Object IDs (OIDs) that are uniquely associated with particular data objects stored in network 102. By using both file interface 16 and object interface 18, it is ordinarily possible for DSD 106 to perform commands received from a device using a file system and from a device using object based storage. This arrangement can allow for more flexibility in the devices and software used on network 102 and the distributed functions that can be performed on network 102.

[0084] FIG. 8B is a flowchart for a data object creation process that can be performed by processor 138 according to an embodiment. In block 808, processor 138 receives a first file and a second file via network interface 126. A file interface 16 can be used by processor 138 to store the first file and the second file in NVM of DSD 106 as described above for block 706 of FIG. 7.

[0085] In block 810, processor 138 executes an object interface 18 to represent data from the first file and the second file as a data object. In one example, the data object may be represented based on a request received from a device such as host 101 to retrieve data stored in NVM of DSD 106. As noted above with respect to FIG. 8A, different devices on network 102 may interface with DSD 106 using different interfaces of DSD 106.

Off-Load Interface Examples

[0086] As described in more detail with reference to FIGS. 9 to 13 below, the use of an off-load interface executed on DSD 106 can allow for processing at DSD 106 that might otherwise be performed by host 101 or by other hosts/clients in network 102. By off-loading or shifting processing closer to where data is stored, it is ordinarily possible to process data more efficiently and reduce network traffic since the data to be processed does not need to be retrieved by a host or client. In addition, data that is locally processed by a DSD in some cases may not need to be sent back to the host or client, which further improves overall processing efficiency and network traffic. These improvements are particularly apparent where large amounts of data need to be processed.

[0087] FIG. 9 is a flowchart for an off-load interface process that can be performed by processor 138 according to an embodiment. In block 902, processor 138 receives computer-executable instructions from a device on network 102 (e.g., host 101) via network interface 126. The computer-executable instructions can include, for example, an executable image, protocol rules, or other executable instructions set by host 101.

[0088] In block 904, processor 138 executes the computer-executable instructions using an off-load interface 20 to transform data stored in NVM of DSD 106 or data received from network 102 via network interface 126. Examples of the off-load interface 20 can include, for example, a Ceph software platform, a Hadoop software framework, or an OpenStack software platform. The computer-executable instructions received in block 902 can then correspond to the particular off-load interface used by processor 138.

[0089] In block 906, processor 138 using the off-load interface 20 sends the transformed data to another device on network 102 and/or stores the transformed data in NVM of DSD 106. As discussed in more detail below with reference to FIGS. 11 to 13, in some implementations, the transformed data can be sent to another DSD in a cluster of DSDs for further processing or transformation of the transformed data.

[0090] FIG. 10A is a flowchart for transforming data received from network 102 using an off-load interface 20 executed by processor 138 according to an embodiment. In block 1002, data is received from network 102 via network interface 126. The data can be received from a device on network 102 such as host 101 or DSD 107.

[0091] In block 1004, processor 138 optionally authenticates the data received in block 1002 to verify that the data came from a device on network 102. In one implementation, an off-load interface 20 or an application 22 can be used to apply a public key to verify a digital signature encrypted by a device on network 102 that sent the data using the device's private key. Other methods of authentication may be used in other implementations. In other embodiments, block 1004 can be omitted.

[0092] In block 1006, processor 138 using the off-load interface 20 optionally identifies a portion of the data received in block 1002 for transformation based on a predetermined criterion. The predetermined criterion can result from execution of computer-executable instructions received from host 101 via network interface 126.

For example, processor 138 may receive data corresponding to a set of digital photographs. The computer-executable instructions received from host 101 can configure processor 138 to identify portions of the data corresponding color photographs that are to be transformed into black and white photographs. This may be based on attributes of the data, or in the case of data objects, metadata sent with the data objects indicating that a particular data object corresponds to a color photograph.

[0093] In block 1008, processor 138 executes computer-executable instructions received from host 101 via network interface 126 to transform the data received in block 1002. In one implementation, the computer-executable instructions may cause processor 138 to transform the received data by combining the received data with data stored in NVM of DSD 106. For example, the computer-executable instructions can cause processor 138 to add a zip code from a data structure stored in NVM when the received data includes an address that is missing a zip code.

[0094] In block 1010, processor 138 optionally generates metadata concerning the transformed data. In one example, the metadata may provide statistical information about the transformed data such as a count of a certain type of data received by DSD 106. The metadata can be stored in NVM of DSD 106 and/or sent to a device on network 102 via network interface 126. In other embodiments, block 1010 can be omitted.

[0095] In block 1012, processor 138 executing the off-load interface 20 stores the transformed data in NVM of DSD 106 and/or sends the transformed data to another device on network 102 via network interface 126. As discussed below with reference to FIGS. 11 to 13 the transformed data can be sent to another device on network 102 for further processing or storage of the transformed data.

[0096] FIG. 10B is a flowchart that can be performed by processor 138 for transforming data read from NVM of DSD 106 using an off-load interface according to an embodiment. In block 1014, processor 138 retrieves data from NVM of DSD 106 using an off-load interface 20. The retrieval of data can be, for example, in response to a command received from a device on network 102 or as a function triggered by computer-executable instructions received from host 101.

[0097] In block 1015, processor 138 optionally authenticates the data retrieved in block 1014 to verify that the data is the correct data or that the data has not been

altered. In one implementation, an off-load interface 20 or an application 22 can be used to authenticate the retrieved data with a cryptographic hash function. Other methods of authentication may be used in other implementations. In other embodiments, block 1015 can be omitted.

[0098] In block 1016, processor 138 using the off-load interface 20 optionally identifies a portion of the data retrieved in block 1014 for transformation based on a predetermined criterion. In some implementations, the predetermined criterion can result from execution of computer-executable instructions received from host 101 via network interface 102. In other embodiments, block 1016 can be omitted.

[0099] In block 1018, processor 138 using the off-load interface 20 transforms the data retrieved from NVM in block 1014. In block 1020, processor 138 optionally generates metadata concerning the transformed data and sends the metadata to a device on network 102 and/or stores the metadata in NVM of DSD 106. In one example, the generated metadata provides statistical information concerning the transformed data. In other embodiments, block 1018 can be omitted.

[00100] In block 1022, processor 138 executing the off-load interface 20 stores the transformed data in NVM of DSD 106 and/or sends the transformed data to another device on network 102 via network interface 126. As discussed below with reference to FIGS. 11 to 13 the transformed data can be sent to another device on network 102 for further processing or storage of the transformed data.

[00101] FIG. 11 is a diagram depicting cluster 200 of DSDs used to transform data in series according to an embodiment. As shown in FIG. 11, host 101 provides cluster 200 with executable instructions for processing data. The executable instructions can include, for example, an executable image, protocol rules, or other executable instructions set by host 101 that can be used by an off-load interface 20 at each of the DSDs in cluster 200.

[00102] Raw data is fed to DSD 106 from a device on network 102 (e.g., a host/client or DSD) where DSD 106 executes the executable instructions received from host 101 to transform the raw data into first transformed data. The transformation of the raw data can follow a process such as the process described above for FIG. 10A.

[0100] The first transformed data is sent to DSD 106, which further transforms the first transformed data into second transformed data before sending the second

transformed data to another DSD in cluster 200. This process may continue for a set number of transformations or until a particular result is obtained. In other embodiments, the second transformed data can be stored at DSD 107 without sending the second transformed data on to a next DSD in cluster 200.

[0101] FIG. 12 is a diagram depicting cluster 300 of DSDs used to perform a MapReduce function according to an embodiment. As with the example of FIG. 11, host 101 can provide cluster 300 with executable instructions for processing data. The executable instructions can include, for example, an executable image, protocol rules, or other executable instructions set by host 101 that can be used by an off-load interface 20 at each of the DSDs in cluster 300.

[0102] As shown in FIG. 12, the DSDs in cluster 300 are arranged in three tiers to perform a MapReduce function. A MapReduce function can provide for parallel processing of a large data set to search, analyze, log, and/or return a reduced data set. In the example of FIG. 12, raw data is mapped using a first tier of DSDs to output transformed data to a second tier of DSDs which shuffles or redistributes the transformed data to a third tier of DSDs to perform a reduce function. Other embodiments may include more than three tiers of DSDs or may include more DSDs than those shown in FIG. 12 for each tier.

[0103] In one example, DSDs 106, 107, and 108 store the raw data and transforms the raw data using a process such as the process described above in FIG. 10B and sends the transformed data to DSDs 109 and 110 in a second tier of cluster 300. DSDs 109 and 110 store the transformed data and further transform the data and directs the transformed data to DSD 111 in a third tier of cluster 300 which reduces the transformed data to a result that is stored in DSD 112.

[0104] In one example, the MapReduce function can be used to identify certain combinations of DNA from sets of raw DNA sequences received in the first tier of cluster 300. The second tier of cluster 300 can be used to further refine the identified DNA combinations based on a narrower set of characteristics before sending the refined data set to the third tier for further reducing. Since each set of DNA combinations is stored on a DSD in cluster 300, a different reduction can subsequently be performed using a different set of characteristics at different tiers. A similar application of a MapReduce function can be used, for example, with chemical, bacterial, or viral analysis.

[0105] In another example, the MapReduce function can be used for photo analytics where raw photos are stored in the first tier, and search criteria at the first tier is used to filter the raw photos to a reduced set of photos that are then stored in the second tier. The reduced set stored in the second tier can be further reduced for example by location using cell phone logistical information applied by an off-load interface at DSD 111.

[0106] In another example, the MapReduce function can be used to analyze shopping behaviors. Raw data stored in the first tier can include data for all shoppers, including items purchased, and the time of purchase. The transformed data stored in the second tier can include data for shoppers who bought a particular item during a certain time period. The third tier of DSD 111 can store transformed data identifying which of the shopping data stored in the second tier includes items purchased with a coupon. In such an example, the resulting data stored in DSD 112 can be used to target marketing to particular shoppers.

[0107] In another example, cluster 300 can be used for an Internet or Intranet search engine function where the first tier of DSDs store a raw set of searches parsed by user, the second tier stores transformed data reduced to the most frequent searches and the most frequent users to enable faster retrieval, and the third tier can be used to further reduce the transformed data to very high frequency searches such as "viral" hits.

[0108] In yet another example, the MapReduce function of cluster 300 can be used to perform sound analysis. The first tier could store all sounds, from all sources and locations, the second tier could store a subset of the raw data based upon search criteria (i.e., known voice matching or recognition, background noise identification, or machine noise recognition), and the third tier of DSD 111 may match the data stored in the second tier to a particular sound or location for a criminal investigation.

[0109] As demonstrated by the five examples provided above, many more example applications of cluster 300 with a MapReduce are possible within the scope of this disclosure.

[0110] FIG. 13 is a diagram depicting a cluster 400 of DSDs where each DSD serves as a compute node according to an embodiment. As shown in FIG. 13, host 101 provides cluster 400 with executable instructions to configure each of the DSDs in cluster 400 (i.e., DSD 106 and DSDs 107 to DSD N) to serve as a compute node.

The executable instructions can include, for example, an executable image, protocol rules, or other executable instructions set by host 101 that can be used by an off-load interface 20 at each of the DSDs in cluster 400, and can cause a processor of each DSD to transform the DSD into a compute node such as, for example, a Hadoop node or an OpenStack node. The compute nodes can then operate independently or in coordination with one or more other nodes to form a distributed system.

[0111] As discussed above, by using an off-load interface at a DSD it is ordinarily possible to improve processing efficiency and network traffic in a distributed system by taking advantage of local data processing. In addition, the use of a off-load interface as described above can allow for a wider variety of devices, software, and functions to be performed in a network of DSDs.

Other Embodiments

[0112] FIGS. 14 and 15 depict alternative data storage environments for implementing the processes described above. Those of ordinary skill in the art will appreciate that other embodiments may include more or less components than those depicted in FIGS. 14 and 15.

[0113] FIG. 14 is a block diagram depicting data storage system 206 where SoC 220 is shared among a plurality of HDAs including HDAs 244, 246, 248, and 250 according to an embodiment. Each HDA can include at least one rotating magnetic disk and at least one head for accessing data from a corresponding disk surface without certain circuitry that is shared among the HDAs. By sharing processing and/or memory resources among a plurality of HDAs, it is ordinarily possible to reduce the cost of storing data and to make better use of processing and memory resources.

[0114] As shown in FIG. 14, bridge 224 includes network interface 226 for communicating on network 102, memory 240, and processor 238 for executing certain software discussed above. Network interface 226 can, for example, allow processor 238 to communicate on network 102 using TCP/IP over Ethernet or WiFi. Memory 240 can represent a volatile memory of data storage system 206 (e.g., DRAM) for temporarily storing data used by bridge 224 or SoC 220. In other embodiments, memory 240 can be an NVM such as MRAM.

[0115] In the example of FIG. 14, processor 238 is configured to execute application OS 12, file system(s) 14, file interface(s) 16, object interface(s) 18, off-load

interface(s) 20, driver(s) 21, and application(s) 22. Application(s) 22 can include applications developed by a manufacturer of data storage system 206 and/or independently developed applications that have been downloaded from network 102. For example, data storage system 206 may receive computer-executable instructions via network interface 226 and then execute the computer-executable instructions to create an application 22. In some implementations, an SDK (e.g., SDK 19 shown in FIG. 3) could be made available to allow users on network 102 to develop their own applications to run on data storage system 206.

[0116] Driver(s) 21 or application(s) 22 can include data storage related applications such as “sniff and store” applications to collect information about data use, a user interface for operating data storage system 206, an HDA health monitoring for monitoring HDA reliability and/or the migration of data to another DSD on network 102 or NVM within data storage system 206 before losing data, data encryption, data compression, an erasure coding or error correction, directing data for storage in a particular NVM of data storage system 206 (e.g., tiered storage), deduplication of data stored in data storage system 206, or mirroring data (e.g., data backup) among HDAs of data storage system 206.

[0117] In addition, driver(s) 21 or application(s) 22 can customize data storage system 206 for specific uses such as streaming certain types of media over network 102, or configuring data storage system 206 to operate as a DVR or media server. Driver(s) 21 or application(s) 22 can also provide for certain network functions on network 102 such as bit coin mining, identification of another DSD or data storage system on network 102 and identification of data stored in another DSD or data storage system, migration or copying of data from a failed DSD or HDA NVM to another DSD or HDA NVM on network 102, or the rerouting of data to other DSDs in the network based on environmental conditions such as vibration or temperature.

[0118] Data included in mapping 24, write pointers 26, command queue 28, buffer 30, or data to be stored in or retrieved from HDAs can be stored in memory 240 so that the data can be accessed by processor 238 during execution of software programs to the extent that such software programs have a need to access and/or modify the data. In addition to shifting processing from a host on network 102 to data storage system 206 with off-load interface(s) 20, running application OS 12 at processor 238 can allow for shifting tasks to file system(s) 14, driver(s) 21, and/or

application(s) 22 that may otherwise have been performed by firmware of data storage system 206 (e.g., controller firmware 11). In addition to the tasks mentioned above of data encryption, data compression, erasure coding or other error correction, the direction of data for storage in HDAs based on data attributes, the direction of data for storage in different zones of NVM (e.g., SMR zones and CMR zones), deduplication of data, and data mirroring, processor 238 can use mapping 24, write pointers 26, command queue 28, and buffer 30 as described above for the embodiment of FIG. 1 to perform other tasks related to the management of HDAs in data storage system 206. By moving such tasks from controller firmware 11 to a file system 14, application 22, or driver 21 executed on processor 238, the tasks can be more easily modified in the field. For example, a new HDA can be added as part of data storage system 206 that may require or benefit from a new driver 21 or new application 22 that can be more easily added to bridge 224 than by changing a firmware of data storage system 206. In addition, shifting more tasks to software operating on processor 238 from firmware of data storage system 206 can ordinarily allow for a simplified firmware.

[0119] Unlike the example embodiment of FIG. 1, processor 238 interfaces with processor 239 of SoC 220 that executes interface firmware 15 for interfacing with processor 238. The interface between processor 238 and processor 239 can be accomplished using a bus standard, such as for example, SATA, PCIe, SCSI, or SAS. In other implementations, processor 238 may use an object interface 18 or a custom interface such as a customized file system 14 or off-load interface 20 to provide processor 239 with, for example, device block addresses indicating a physical location for retrieving or storing data objects in HDAs 244, 246, 248, or 250.

[0120] SoC 220 also includes processor 241 for executing controller firmware 11 and processor 242 for executing servo firmware 10. In other embodiments, two or more of processors 239, 241, or 242 can be combined into a single processor. In the example of FIG. 14, each of processor 241 and 242 communicate with HDAs 244, 246, 248, and 250 to retrieve data from the HDAs and store data in the HDAs. In this regard, SoC 220 may include one or more multiplexers (not shown) for sending signals such as VCM control signals 34, SM control signals 38, or read/write signals 36 to each of HDAs 244, 246, 248, and 250. Processors 241 and 242 can be configured to interface with one HDA at a time or with multiple HDAs at the same time depending on different design criteria for data storage system 206. In one example,

processors 241 and 242 may only interface with one HDA at a time to conserve power in a data archive implementation. In another example, processors 241 and 242 may interface with multiple HDAs at the same time to offer quicker access to mirrored data stored on rotationally synchronized disks in different HDAs.

[0121] In this regard, each of HDAs 244, 246, 248, and 250 can include one or more rotating magnetic disks such as disk 150 in FIG. 1, together with one or more heads for reading and writing data on a disk surface. As with the embodiment of FIG. 1, the heads used in HDAs 244, 246, 248, and 250 can be positioned relative to a corresponding disk surface by moving an actuator with a VCM.

[0122] FIG. 15 is a block diagram depicting data storage system 306 where controller 324 is shared among a plurality of DSDs including DSDs 307, 308, 309, and 310 according to an embodiment. In the example embodiment of FIG. 15, processor 338 interfaces with DSDs 307, 308, 309, and 310. The interface between processor 338 and the DSDs can be accomplished using a bus standard, such as for example, SATA, PCIe, SCSI, or SAS. In other implementations, processor 338 may use an object interface 18 or a custom interface such as a customized file system 14, off-load interface 20, or driver 21 to interface with one or more of DSDs 307, 308, 309, and 310. In this regard, processor 338 may interface with different DSDs using different interfaces. For example, processor 338 may use a file system 14 to interface with DSD 307, a file interface 16 to interface with DSD 308, an object interface to interface with DSD 309, and an off-load interface 20 or driver 21 to interface with DSD 310.

[0123] As shown in FIG. 15, controller 324 includes network interface 326 for communicating on network 102, memory 340, and processor 338. Network interface 326 can, for example, allow processor 338 to communicate on network 102 using TCP/IP over Ethernet or WiFi. Memory 340 can represent a volatile memory of data storage system 206 (e.g., DRAM) for temporarily storing data used by controller 324. In other embodiments, memory 340 can be an NVM such as MRAM.

[0124] In the example of FIG. 15, processor 338 is configured to execute application OS 12, file system(s) 14, file interface(s) 16, object interface(s) 18, off-load interface(s) 20, driver(s) 21, and application(s) 22. Application(s) 22 can include applications developed by a manufacturer of data storage system 306 and/or independently developed applications that have been downloaded from network 102.

[0125] Similar to the driver(s) 21 and application(s) 22 discussed above, driver(s) 21 or application(s) 22 can include data storage related applications such as “sniff and store” applications, DSD health monitoring and/or the migration of data to another DSD on network 102 or DSD within data storage system 306, data encryption, data compression, erasure coding or error correction, directing data for storage in a particular NVM of data storage system 306 based on attributes of the data (e.g., tiered storage), deduplication of data stored in data storage system 306, or mirroring data (e.g., data backup or Redundant Array of Independent Disks (RAID) applications).

[0126] In addition, driver(s) 21 or application(s) 22 can customize data storage system 306 for specific uses such as streaming certain types of media over network 102, or configuring data storage system 306 to operate as a DVR or media server. Driver(s) 21 or application(s) 22 can also provide for certain network functions on network 102 such as bit coin mining, identification of another DSD or data storage system on network 102 and identification of data stored in another DSD or data storage system, migration or copying of data from a failed DSD NVM to another DSD NVM on network 102, or the rerouting of data to other DSDs in the network or within data storage system 306 based on environmental conditions such as vibration or temperature. In this case, a sensor such as sensor 122 can be located at each of DSDs 307, 308, 309, and 310, or at selected DSDs in data storage system 306, to determine environmental conditions at different DSDs.

[0127] Data included in mapping 24, write pointers 26, command queue 28, buffer 30, or data to be stored in or retrieved from NVM can be stored in memory 340 so that the data can be accessed by processor 338 during execution of software programs to the extent that such software programs have a need to access and/or modify the data. In addition to shifting processing from a host on network 102 to data storage system 306 with off-load interface(s) 20, running application OS 12 at processor 338 can allow for shifting tasks to file system(s) 14, driver(s) 21, and/or application(s) 22 that may otherwise have been performed by firmware (e.g., controller firmware 11) at DSDs in data storage system 306. By moving such tasks from a controller firmware to a file system 14, application 22, or driver 21 executed on processor 338, the DSD firmware can be simplified and tasks for maintaining data storage system 306 can be more easily modified in the field. In one such example, a new DSD with a different type of NVM can be added to data storage system 306 and a new driver 21 or

application 22 executed by processor 338 can be added for managing the different type of NVM.

[0128] Those of ordinary skill in the art will appreciate that the various illustrative logical blocks, modules, and processes described in connection with the examples disclosed herein may be implemented as electronic hardware, computer software, or combinations of both. Furthermore, the foregoing processes can be embodied on a computer readable medium which causes a processor or computer to perform or execute certain functions.

[0129] To clearly illustrate this interchangeability of hardware and software, various illustrative components, blocks, and modules have been described above generally in terms of their functionality. Whether such functionality is implemented as hardware or software depends upon the particular application and design constraints imposed on the overall system. Those of ordinary skill in the art may implement the described functionality in varying ways for each particular application, but such implementation decisions should not be interpreted as causing a departure from the scope of the present disclosure.

[0130] The various illustrative logical blocks, units, modules, and controllers described in connection with the examples disclosed herein may be implemented or performed with a general purpose processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, discrete hardware components, or any combination thereof designed to perform the functions described herein. A general purpose processor may be a microprocessor, but in the alternative, the processor may be any conventional processor, controller, microcontroller, or state machine. A processor may also be implemented as a combination of computing devices, e.g., a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DSP core, or any other such configuration.

[0131] The activities of a method or process described in connection with the examples disclosed herein may be embodied directly in hardware, in a software module executed by a processor, or in a combination of the two. The steps of the method or algorithm may also be performed in an alternate order from those provided in the examples. A software module may reside in RAM memory, flash memory, ROM

memory, EPROM memory, EEPROM memory, registers, hard disk, a removable media, an optical media, or any other form of storage medium known in the art. An exemplary storage medium is coupled to the processor such that the processor can read information from, and write information to, the storage medium. In the alternative, the storage medium may be integral to the processor. The processor and the storage medium may reside in an Application Specific Integrated Circuit (ASIC).

[0132] The foregoing description of the disclosed example embodiments is provided to enable any person of ordinary skill in the art to make or use the embodiments in the present disclosure. Various modifications to these examples will be readily apparent to those of ordinary skill in the art, and the principles disclosed herein may be applied to other examples without departing from the spirit or scope of the present disclosure. The described embodiments are to be considered in all respects only as illustrative and not restrictive and the scope of the disclosure is, therefore, indicated by the following claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

CLAIMS

What is claimed is:

1. A Data Storage Device (DSD), comprising:
a Non-Volatile Memory (NVM) for storing data;
a network interface for communicating on a network; and
a processor configured to:
receive a command via the network interface for storing data in the NVM or retrieving data from the NVM;
based on a transport attribute of the command, determine whether to perform the command using a file interface for accessing files stored in the NVM or an object interface for accessing data objects stored in the NVM; and
perform the command to store data in the NVM or retrieve data from the NVM using the determined interface.
2. The DSD of Claim 1, wherein the processor is further configured to use the file interface and the object interface to retrieve the same data stored in the NVM.
3. The DSD of Claim 1, wherein the processor is further configured to:
receive a first data object and a second data object via the network interface using the object interface; and
represent data from the first data object and the second data object as a file using the file interface.
4. The DSD of Claim 1, wherein the processor is further configured to:
receive a first file and a second file via the network interface using the file interface; and
represent data from the first file and the second file as a data object using object interface.
5. The DSD of Claim 1, wherein the processor is further configured to execute an off-load interface for transforming data stored in the NVM.

6. The DSD of Claim 5, wherein the off-load interface includes at least one of a Ceph software platform, a Hadoop software framework, and an OpenStack software platform.

7. The DSD of Claim 5, wherein the processor is further configured to isolate the off-load interface from other interfaces executed by the processor using a software container.

8. The DSD of Claim 1, wherein the object interface includes Amazon Web Services Simple Storage Service (AWS S3).

9. A Data Storage Device (DSD) in communication with a device, the DSD comprising:

a Non-Volatile Memory (NVM) for storing data;

a network interface for communicating with the device on a network; and

a processor configured to:

receive computer-executable instructions from the device via the network interface;

execute the computer-executable instructions using an off-load interface executed by the processor, wherein the computer-executable instructions cause the processor to:

transform data stored in the NVM or data received from the network via the network interface; and

send the transformed data to another device on the network via the network interface or store the transformed data in the NVM.

10. The DSD of Claim 9, wherein the computer-executable instructions further cause the processor to send the transformed data to at least one other DSD via the network interface for further transformation of the transformed data.

11. The DSD of Claim 9, wherein the computer-executable instructions configure the DSD to form part of a tiered cluster of DSDs used to perform a MapReduce function.

12. The DSD of Claim 9, wherein the computer-executable instructions configure the DSD to form part of a series of DSDs used to transform data.

13. The DSD of Claim 9, wherein the processor is further configured to isolate the off-load interface from other interfaces executed by the processor using a software container.

14. The DSD of Claim 9, wherein the computer-executable instructions configure the DSD to serve as a compute node in a cluster of DSDs.

15. The DSD of Claim 9, wherein the computer-executable instructions further cause the processor to:

- receive data from the network via the network interface;
- retrieve stored data from the NVM in response to the received data; and
- combine the stored data with the received data to form the transformed data.

16. The DSD of Claim 9, wherein the computer-executable instructions further cause the processor to generate metadata concerning the transformed data.

17. The DSD of Claim 9, wherein the processor is further configured to:

- receive data from the network via the network interface; and
- authenticate the received data as coming from a device on the network.

18. The DSD of Claim 9, wherein the computer-executable instructions further cause the processor to:

- receive data from the network via the network interface;
- identify a portion of the received data for transformation based on a predetermined criterion; and
- transform only the identified portion of the received data.

19. The DSD of Claim 9, wherein the computer-executable instructions further cause the processor to:

identify a portion of the data stored in the NVM for transformation based on a predetermined criterion; and

transform only the identified portion of the data stored in the NVM.

20. A method for operating a Data Storage Device (DSD) including a Non-Volatile Memory (NVM) for storing data, a network interface for communicating on a network, and a processor, the method comprising:

receiving a command via the network interface for storing data in the NVM or retrieving data from the NVM;

based on a transport attribute of the command, determining whether to execute a file interface for accessing files stored in the NVM or an object interface for accessing data objects stored in the NVM; and

performing the command to store the data in the NVM or retrieve the data from the NVM using the determined interface.

21. The method of Claim 20, further comprising using the file interface and the object interface to retrieve the same data stored in the NVM.

22. The method of Claim 20, further comprising:

receiving a first data object and a second data object via the network interface using the object interface; and

representing data from the first data object and the second data object as a file using the file interface.

23. The method of Claim 20, further comprising:

receiving a first file and a second file via the network interface using the file interface; and

representing data from the first file and the second file as a data object using object interface.

24. The method of Claim 20, further comprising using the processor to execute an off-load interface for transforming data stored in the NVM.

25. The method of Claim 24, wherein the off-load interface includes at least one of a Ceph software platform, a Hadoop software framework, and an OpenStack software platform.

26. The method of Claim 24, further comprising isolating the off-load interface from other interfaces executed by the processor using a software container.

27. The method of Claim 20, wherein the object interface includes Amazon Web Services Simple Storage Service (AWS S3).

28. A method for operating a Data Storage Device (DSD) including a Non-Volatile Memory (NVM) for storing data, a network interface for communicating on a network, and a processor, the method comprising:

receiving computer-executable instructions from a device via the network interface; and

executing the computer-executable instructions using an off-load interface executed by the processor, wherein the computer-executable instructions cause the processor to:

transform data stored in the NVM or data received from the network via the network interface; and

send the transformed data to another device on the network via the network interface or store the transformed data in the NVM.

29. The method of Claim 28, wherein executing the computer-executable instructions further cause the processor to send the transformed data to at least one other DSD via the network interface for further transformation of the transformed data.

30. The method of Claim 28, wherein executing the computer-executable instructions configures the DSD to form part of a tiered cluster of DSDs used to perform a MapReduce function.

31. The method of Claim 28, wherein executing the computer-executable instructions configures the DSD to form part of a series of DSDs used to transform data.

32. The method of Claim 28, further comprising isolating the off-load interface from other interfaces executed by the processor using a software container.

33. The method of Claim 28, wherein executing the computer-executable instructions configures the DSD to serve as a compute node in a cluster of DSDs.

34. The method of Claim 28, further comprising:
receiving data from the network via the network interface;
retrieving stored data from the NVM in response to the received data; and
combining the stored data with the received data to form the transformed data.

35. The method of Claim 28, further comprising generating metadata concerning the transformed data.

36. The method of Claim 28, further comprising:
receiving data from the network via the network interface; and
authenticating the received data as coming from a device on the network.

37. The method of Claim 28, wherein executing the computer-executable instructions further cause the processor to:
receive data from the network via the network interface;
identify a portion of the received data for transformation based on a predetermined criterion; and
transform only the identified portion of the received data.

38. The method of Claim 28, wherein executing the computer-executable instructions further cause the processor to:

identify a portion of the data stored in the NVM for transformation based on a predetermined criterion; and

transform only the identified portion of the data stored in the NVM.




FIG. 1

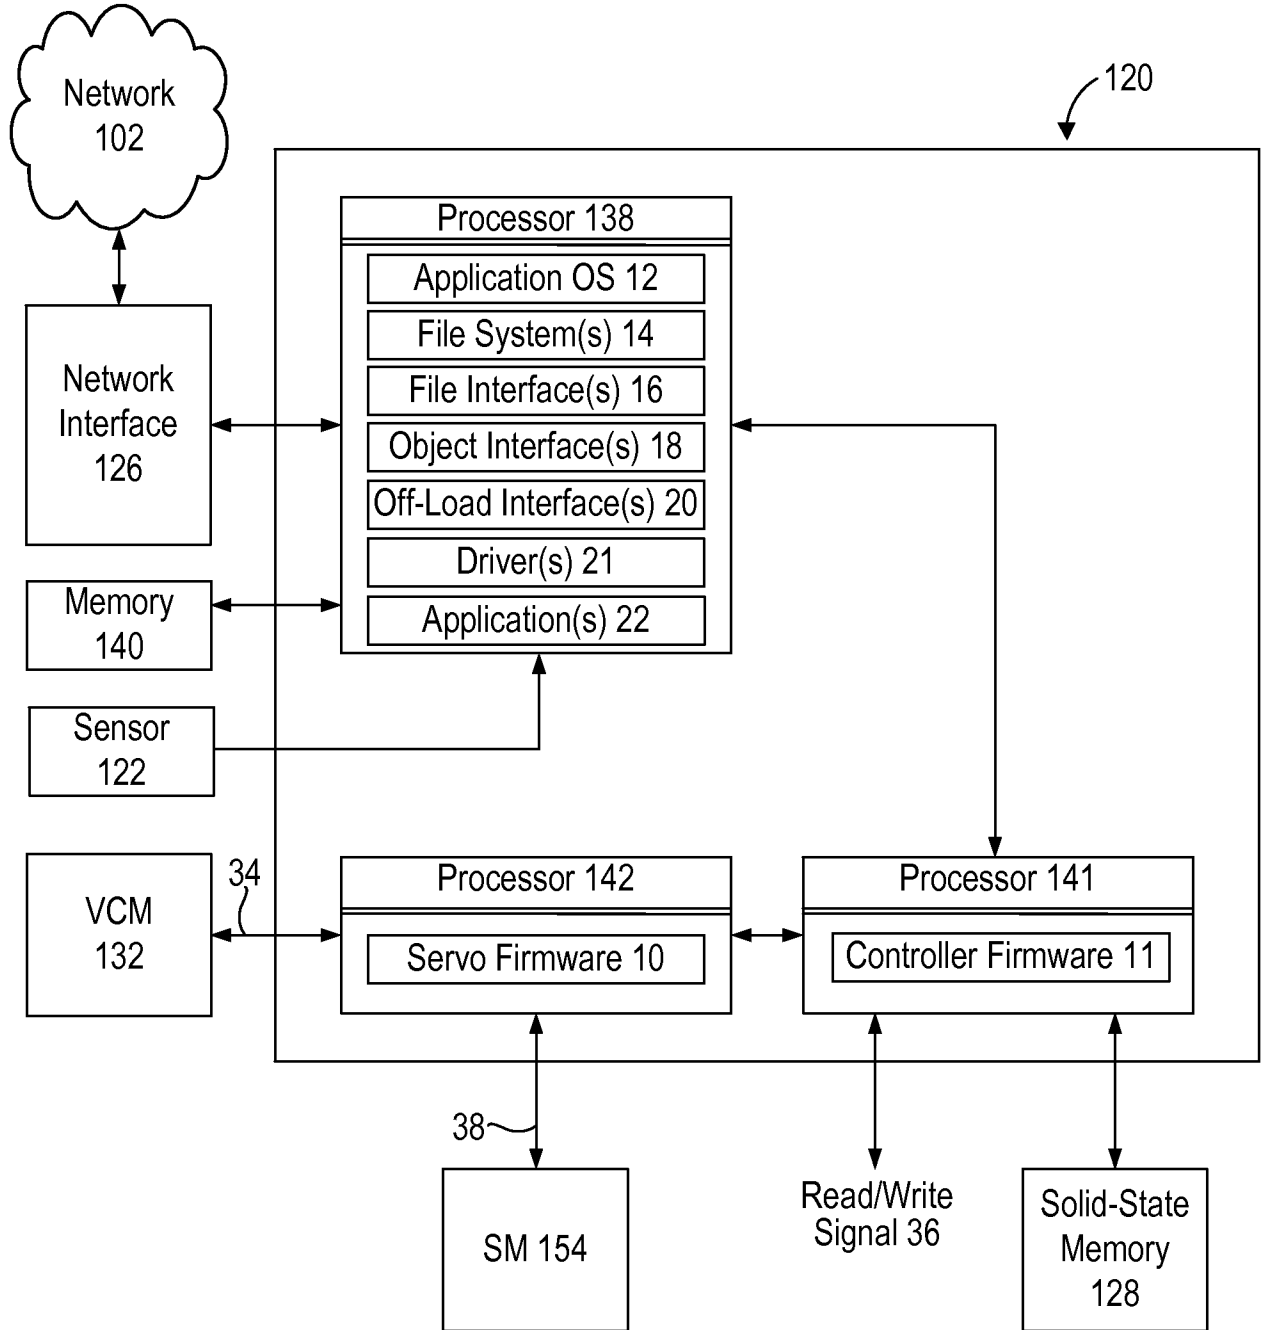


FIG. 2




FIG. 3

4/15

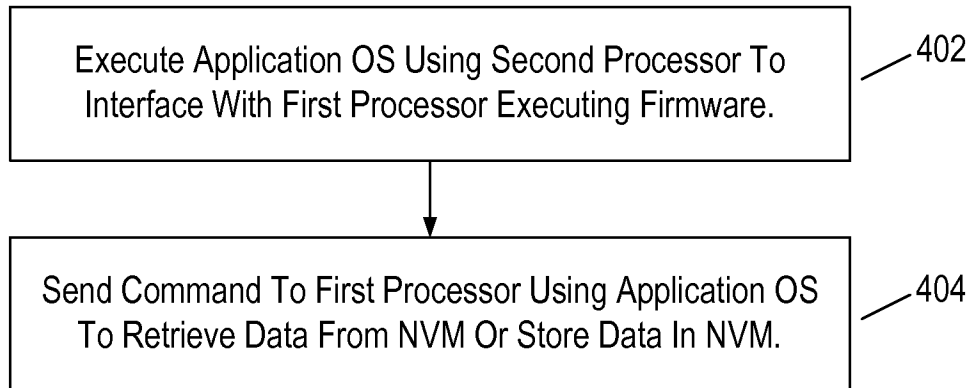


FIG. 4

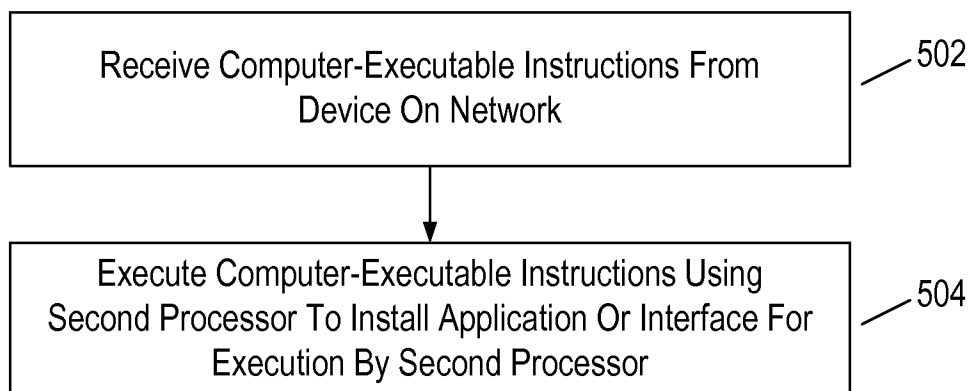


FIG. 5

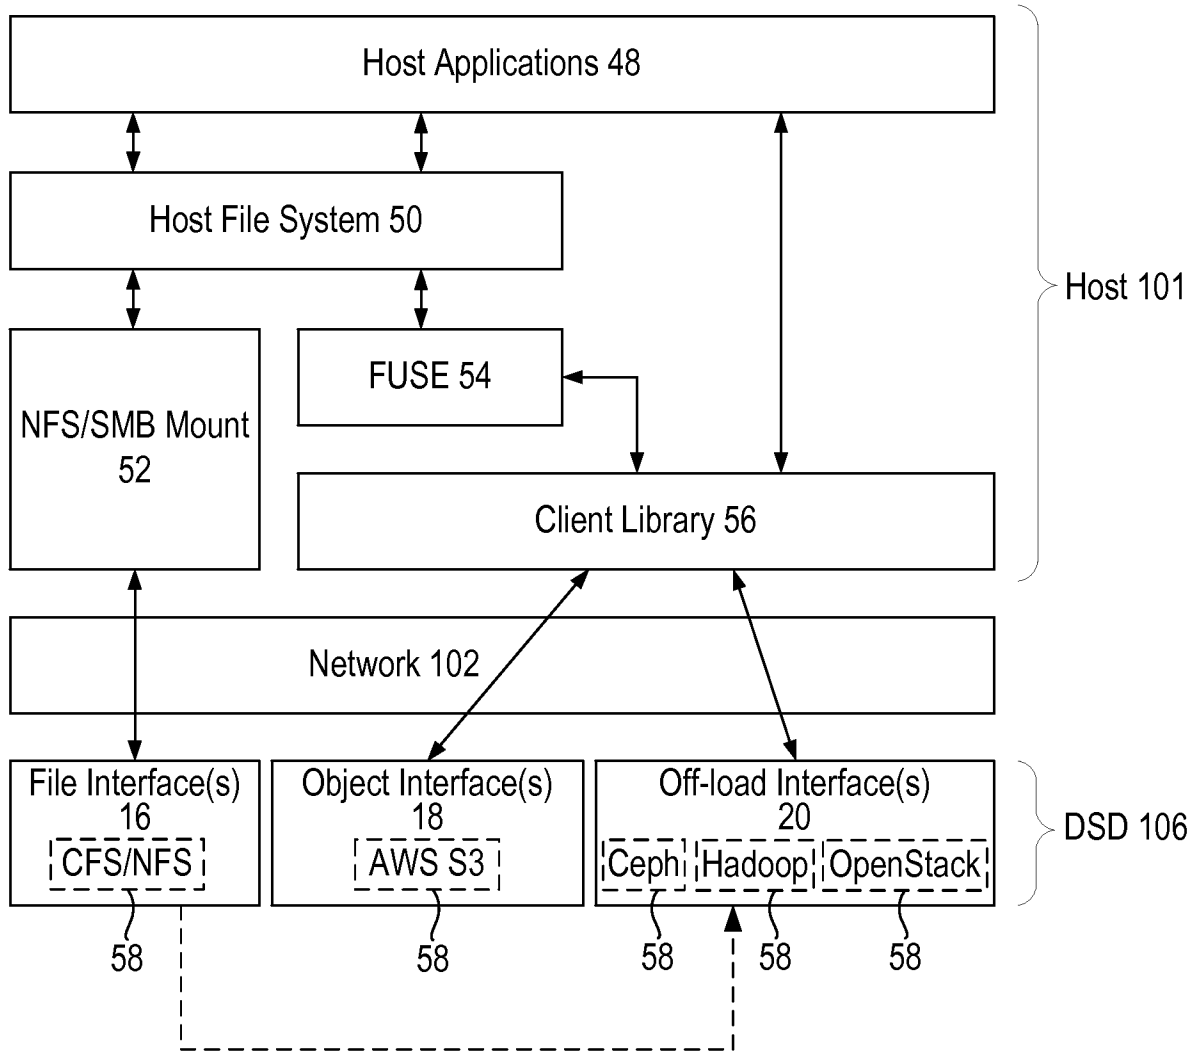


FIG. 6

6/15

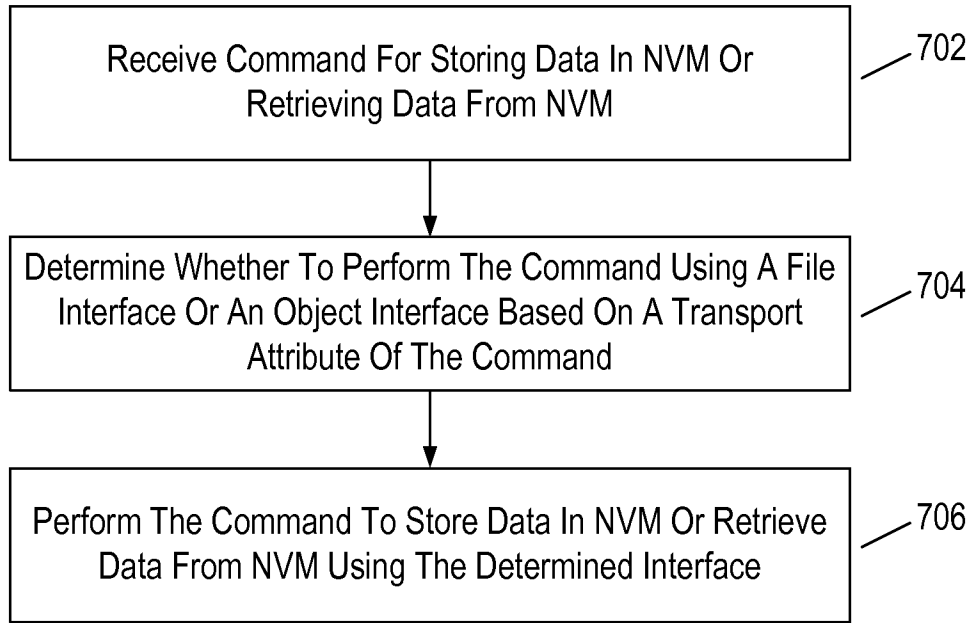


FIG. 7A

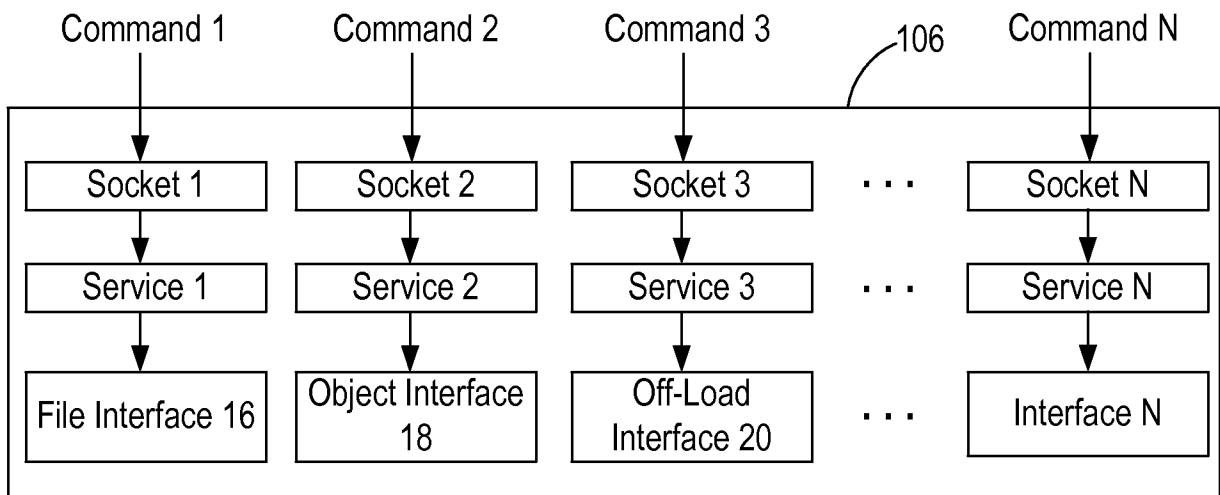


FIG. 7B

7/15

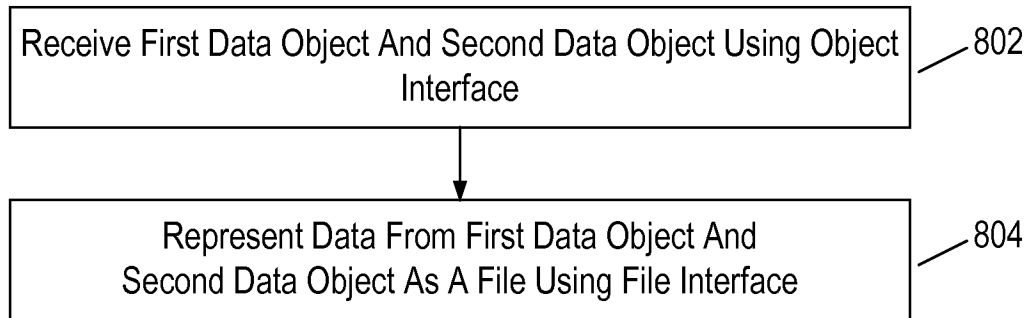


FIG. 8A

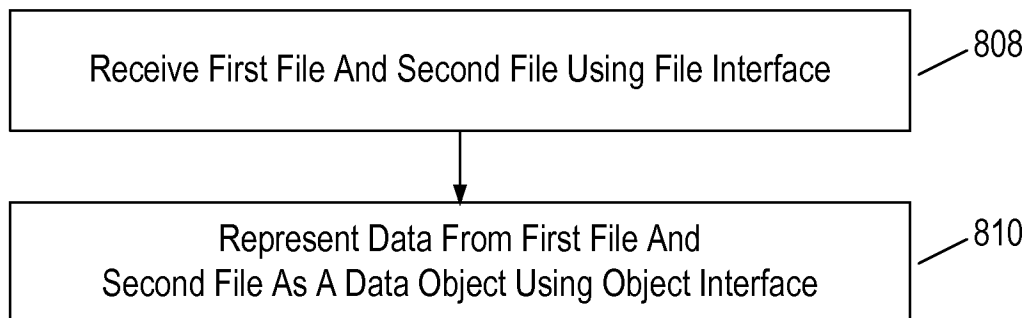


FIG. 8B

8/15

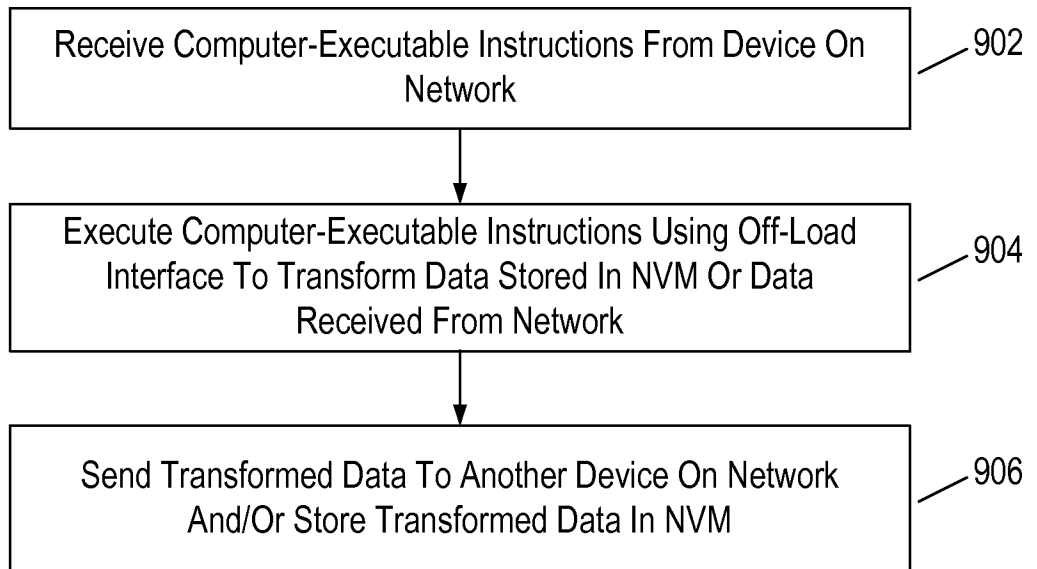


FIG. 9

9/15

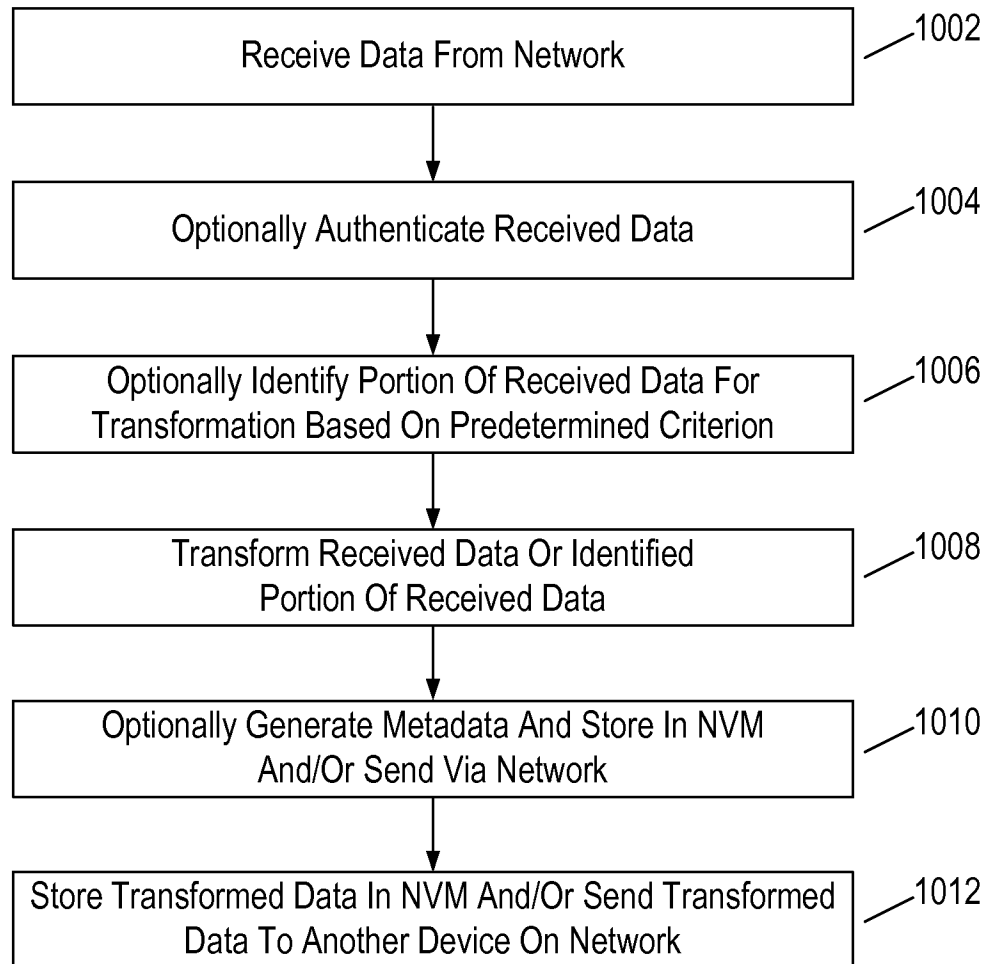
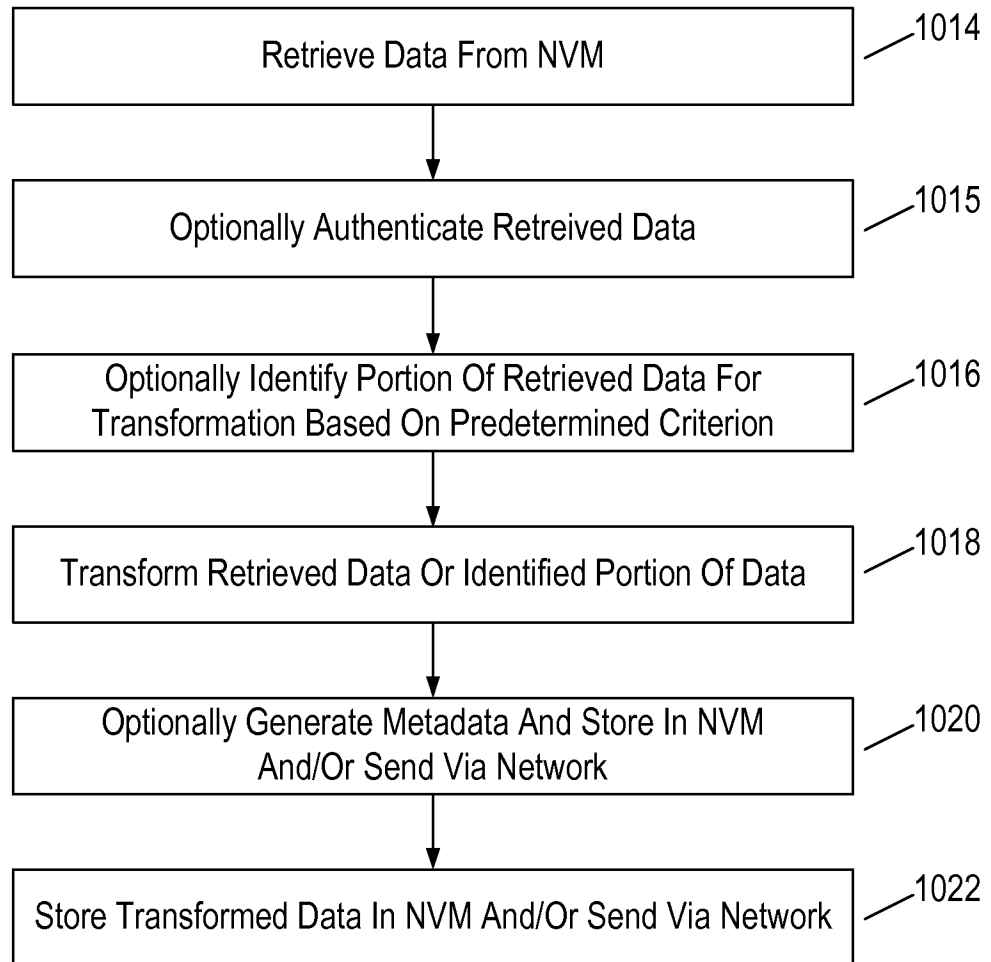


FIG. 10A

10/15

**FIG. 10B**




FIG. 11

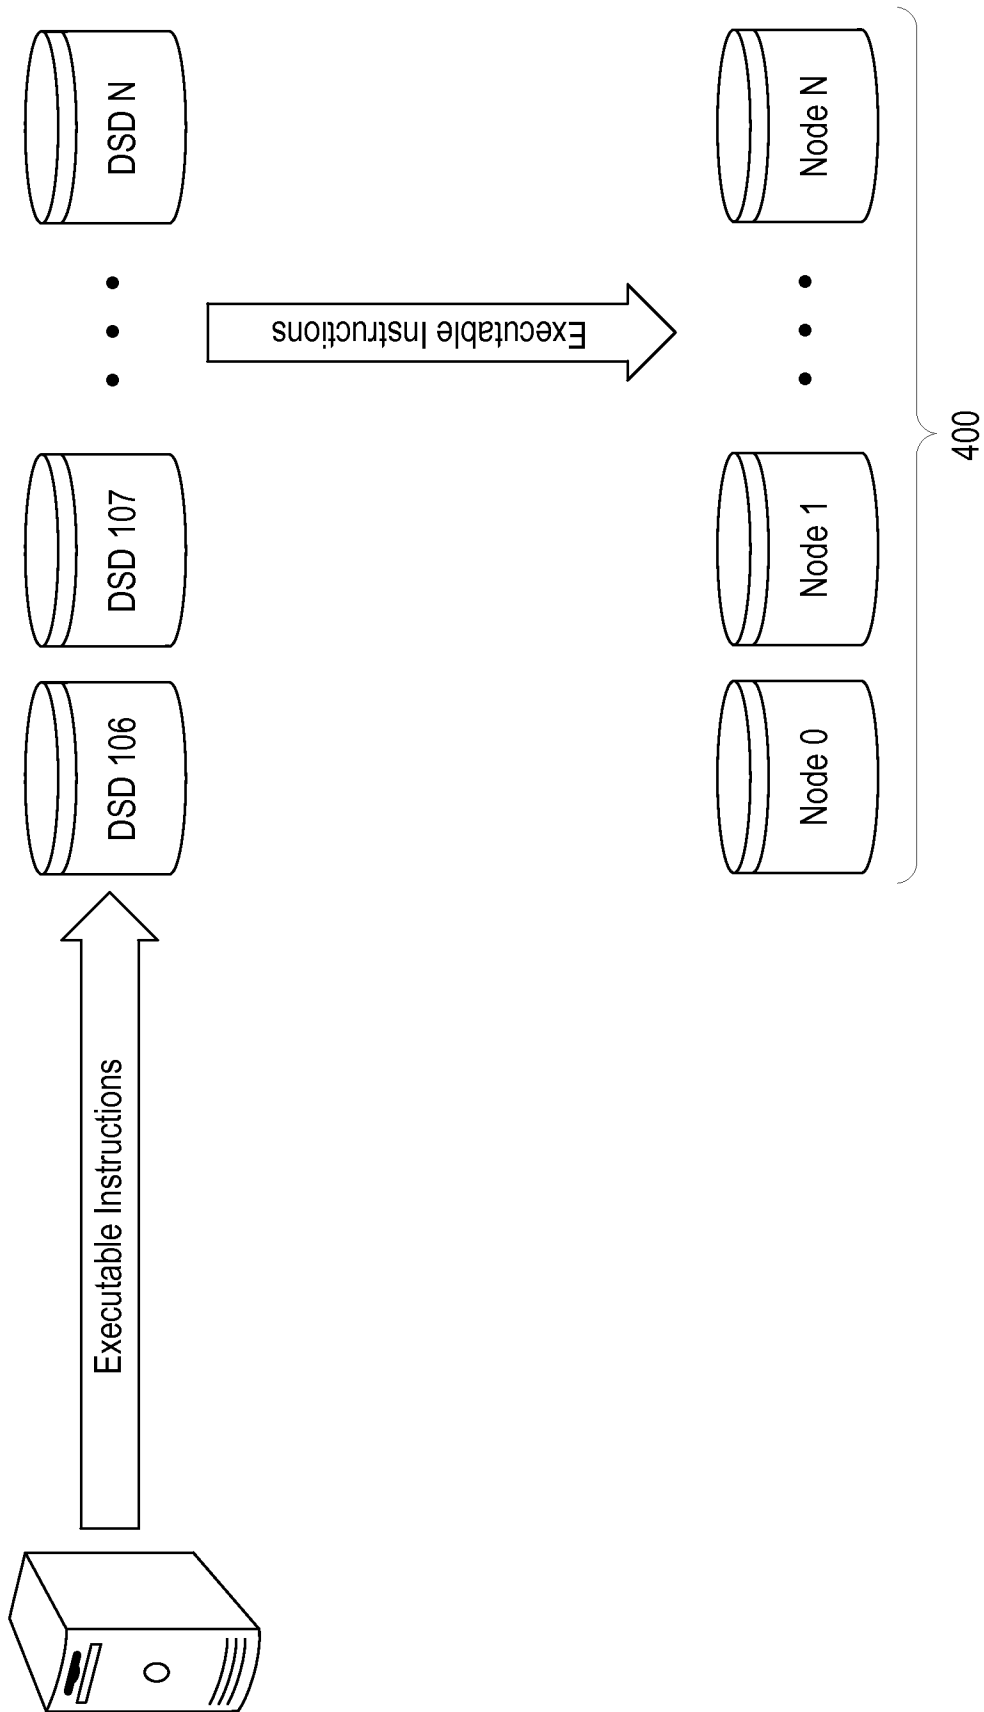


FIG. 13




FIG. 14

15/15




FIG. 15

A. CLASSIFICATION OF SUBJECT MATTER**G06F 13/38(2006.01)i, G06F 9/06(2006.01)i, G06F 12/06(2006.01)i, H04L 29/08(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 13/38; G06F 12/02; G06F 17/30; G06F 12/00; G06F 13/28; G06F 12/08; G06F 9/06; G06F 12/06; H04L 29/08

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: network attached storage, object, NVM, Hadoop, AWS S3, storage device, object based storage, and similar terms

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2005-0154841 A1 (GAUTHAM SASTRI et al.) 14 July 2005 See paragraphs [0015], [0066]-[0067], [0069], and [0248]; and figure 3.	1-38
A	US 2013-0166568 A1 (NATHAN BINKERT et al.) 27 June 2013 See paragraphs [0011] and [0020].	1-38
A	US 2012-0036287 A1 (JAESOO LEE) 09 February 2012 See paragraphs [0007], [0025]-[0027], and [0046]; and figures 1 and 5.	1-38
A	US 2009-0204872 A1 (FRANK YU et al.) 13 August 2009 See paragraphs [0036]-[0047] and figures 1A-1B.	1-38
A	US 2013-0067147 A1 (KAZUHITO OKITA) 14 March 2013 See paragraphs [0013]-[0016] and figures 1-2.	1-38

 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

28 April 2015 (28.04.2015)

Date of mailing of the international search report

28 April 2015 (28.04.2015)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea

Facsimile No. ++82 42 472 7140

Authorized officer

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/015877

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2005-0154841 A1	14/07/2005	US 2003-0037061 A1	20/02/2003
US 2013-0166568 A1	27/06/2013	CA 2860322 A1	27/06/2013
		CN 104160394 A	19/11/2014
		EP 2795487 A1	29/10/2014
		US 2014-181141 A1	26/06/2014
		US 8732213 B2	20/05/2014
		WO 2013-096887 A1	27/06/2013
US 2012-0036287 A1	09/02/2012	KR 10-2012-0021490 A	09/03/2012
		US 8631166 B2	14/01/2014
US 2009-0204872 A1	13/08/2009	CN 100487678 C	13/05/2009
		CN 100501702 C	17/06/2009
		CN 101097551 A	02/01/2008
		CN 101118783 A	06/02/2008
		CN 101122865 A	13/02/2008
		CN 101122865 B	14/07/2010
		CN 101122887 A	13/02/2008
		CN 101398764 A	01/04/2009
		CN 101398785 A	01/04/2009
		CN 101399075 A	01/04/2009
		CN 101399076 A	01/04/2009
		CN 101403997 A	08/04/2009
		CN 101409111 A	15/04/2009
		CN 101425324 A	06/05/2009
		CN 101630375 A	20/01/2010
		CN 101644995 A	10/02/2010
		CN 101644995 B	28/12/2011
		CN 101727976 A	09/06/2010
		CN 101727976 B	19/09/2012
		CN 101923512 A	22/12/2010
		CN 101923512 B	20/03/2013
		CN 102469715 A	23/05/2012
		CN 2859750 Y	17/01/2007
		CN 2886681 Y	04/04/2007
		JP 2001-005945 A	12/01/2001
		JP 2001-118046 A	27/04/2001
		JP 3274433 B2	15/04/2002
		JP 3338417 B2	28/10/2002
		TW 200813713 A	16/03/2008
		TW 200842887 A	01/11/2008
		TW 200915080 A	01/04/2009
		TW 200917017 A	16/04/2009
		TW 200917256 A	16/04/2009
		TW 200926394 A	16/06/2009
		TW 201005652 A	01/02/2010
		TW I351599 B	01/11/2011
		TW I351605 B	01/11/2011

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/015877

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 2003-061474 A1	27/03/2003
		US 2004-236980 A1	25/11/2004
		US 2005-055481 A1	10/03/2005
		US 2005-059273 A1	17/03/2005
		US 2005-059301 A1	17/03/2005
		US 2005-070138 A1	31/03/2005
		US 2005-085129 A1	21/04/2005
		US 2005-085133 A1	21/04/2005
		US 2005-114587 A1	26/05/2005
		US 2005-120146 A1	02/06/2005
		US 2005-120157 A1	02/06/2005
		US 2005-120163 A1	02/06/2005
		US 2005-127735 A1	16/06/2005
		US 2005-138288 A1	23/06/2005
		US 2005-156333 A1	21/07/2005
		US 2005-160213 A1	21/07/2005
		US 2005-160218 A1	21/07/2005
		US 2005-164532 A1	28/07/2005
		US 2005-181645 A1	18/08/2005
		US 2005-182881 A1	18/08/2005
		US 2005-193161 A1	01/09/2005
		US 2005-193162 A1	01/09/2005
		US 2005-197017 A1	08/09/2005
		US 2005-201148 A1	15/09/2005
		US 2005-204187 A1	15/09/2005
		US 2005-223158 A1	06/10/2005
		US 2006-002096 A1	05/01/2006
		US 2006-030080 A1	09/02/2006
		US 2006-067054 A1	30/03/2006
		US 2006-075395 A1	06/04/2006
		US 2006-161725 A1	20/07/2006
		US 2006-286865 A1	21/12/2006
		US 2006-294272 A1	28/12/2006
		US 2007-076387 A1	05/04/2007
		US 2007-079043 A1	05/04/2007
		US 2007-118688 A1	24/05/2007
		US 2007-130414 A1	07/06/2007
		US 2007-130436 A1	07/06/2007
		US 2007-143509 A1	21/06/2007
		US 2007-147157 A1	28/06/2007
		US 2007-150963 A1	28/06/2007
		US 2007-156587 A1	05/07/2007
		US 2007-168614 A1	19/07/2007
		US 2007-178769 A1	02/08/2007
		US 2007-180264 A1	02/08/2007
		US 2007-183209 A1	09/08/2007
		US 2007-184685 A1	09/08/2007
		US 2007-184719 A1	09/08/2007
		US 2007-197101 A1	23/08/2007
		US 2007-198856 A1	23/08/2007

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/015877

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 2007-201274 A1	30/08/2007
		US 2007-204128 A1	30/08/2007
		US 2007-204206 A1	30/08/2007
		US 2007-233955 A1	04/10/2007
		US 2007-250564 A1	25/10/2007
		US 2007-255891 A1	01/11/2007
		US 2007-262155 A1	15/11/2007
		US 2007-268754 A1	22/11/2007
		US 2007-274032 A1	29/11/2007
		US 2007-276987 A1	29/11/2007
		US 2007-276988 A1	29/11/2007
		US 2007-283428 A1	06/12/2007
		US 2007-292009 A1	20/12/2007
		US 2007-293088 A1	20/12/2007
		US 6547130 B1	15/04/2003
		US 6854984 B1	15/02/2005
		US 6874044 B1	29/03/2005
		US 7004794 B2	28/02/2006
		US 7021971 B2	04/04/2006
		US 7035110 B1	25/04/2006
		US 7044802 B2	16/05/2006
		US 7069369 B2	27/06/2006
		US 7073010 B2	04/07/2006
		US 7082056 B2	25/07/2006
		US 7094074 B2	22/08/2006
		US 7095617 B1	22/08/2006
		US 7103684 B2	05/09/2006
		US 7103765 B2	05/09/2006
		US 7104848 B1	12/09/2006
		US 7108560 B1	19/09/2006
		US 7125287 B1	24/10/2006
		US 7130958 B2	31/10/2006
		US 7174628 B1	13/02/2007
		US 7182646 B1	27/02/2007
		US 7186147 B1	06/03/2007
		US 7215551 B2	08/05/2007
		US 7243185 B2	10/07/2007
		US 7249978 B1	31/07/2007
		US 7257714 B1	14/08/2007
		US 7259967 B2	21/08/2007
		US 7264992 B2	04/09/2007
		US 7269004 B1	11/09/2007
		US 7296345 B1	20/11/2007
		US 7297024 B2	20/11/2007
		US 7299316 B2	20/11/2007
		US 7301776 B1	27/11/2007
		US 7318117 B2	08/01/2008
		US 7333364 B2	19/02/2008
		US 7347736 B2	25/03/2008
		US 7365985 B1	29/04/2008

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/015877

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 7383362 B2	03/06/2008
		US 7394661 B2	01/07/2008
		US 7407393 B2	05/08/2008
		US 7420803 B2	02/09/2008
		US 7427217 B2	23/09/2008
		US 7428605 B2	23/09/2008
		US 7438562 B2	21/10/2008
		US 7440286 B2	21/10/2008
		US 7440287 B1	21/10/2008
		US 7447037 B2	04/11/2008
		US 7457897 B1	25/11/2008
		US 7466556 B2	16/12/2008
		US 7467290 B2	16/12/2008
		US 7471556 B2	30/12/2008
		US 7475174 B2	06/01/2009
		US 7476105 B2	13/01/2009
		US 7479039 B2	20/01/2009
		US 7483329 B2	27/01/2009
		US 7507119 B2	24/03/2009
		US 7517252 B2	14/04/2009
		US 7524198 B2	28/04/2009
		US 7535088 B2	19/05/2009
		US 7535719 B2	19/05/2009
		US 7544073 B2	09/06/2009
		US 7547218 B2	16/06/2009
		US 7552251 B2	23/06/2009
		US 7576990 B2	18/08/2009
		US 7606111 B2	20/10/2009
		US 7609523 B1	27/10/2009
		US 7610438 B2	27/10/2009
		US 7620769 B2	17/11/2009
		US 7628622 B2	08/12/2009
		US 7631195 B1	08/12/2009
		US 7643334 B1	05/01/2010
		US 7649742 B2	19/01/2010
		US 7649743 B2	19/01/2010
		US 7657692 B2	02/02/2010
		US 7660938 B1	09/02/2010
		US 7660941 B2	09/02/2010
		US 7664902 B1	16/02/2010
		US 7673080 B1	02/03/2010
		US 7676640 B2	09/03/2010
		US 7680977 B2	16/03/2010
		US 7690030 B1	30/03/2010
		US 7690031 B2	30/03/2010
		US 7702831 B2	20/04/2010
		US 7702984 B1	20/04/2010
		US 7707321 B2	27/04/2010
		US 7707354 B2	27/04/2010
		US 7708570 B2	04/05/2010

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/015877

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 7740493 B2	22/06/2010
		US 7744387 B2	29/06/2010
		US 7761653 B2	20/07/2010
		US 7768785 B2	03/08/2010
		US 7768789 B2	03/08/2010
		US 7769944 B2	03/08/2010
		US 7771215 B1	10/08/2010
		US 7788553 B2	31/08/2010
		US 7789680 B2	07/09/2010
		US 7795714 B2	14/09/2010
		US 7802155 B2	21/09/2010
		US 7804163 B2	28/09/2010
		US 7806705 B2	05/10/2010
		US 7809862 B2	05/10/2010
		US 7814337 B2	12/10/2010
		US 7818492 B2	19/10/2010
		US 7827348 B2	02/11/2010
		US 7830666 B2	09/11/2010
		US 7836236 B2	16/11/2010
		US 7844763 B2	30/11/2010
		US 7849242 B2	07/12/2010
		US 7850082 B1	14/12/2010
		US 7850468 B2	14/12/2010
		US 7855099 B2	21/12/2010
		US 7861312 B2	28/12/2010
		US 7865630 B2	04/01/2011
		US 7865809 B1	04/01/2011
		US 7866562 B2	11/01/2011
		US 7869218 B2	11/01/2011
		US 7869219 B2	11/01/2011
		US 7872871 B2	18/01/2011
		US 7872873 B2	18/01/2011
		US 7873837 B1	18/01/2011
		US 7873885 B1	18/01/2011
		US 7874067 B1	25/01/2011
		US 7877542 B2	25/01/2011
		US 7878852 B2	01/02/2011
		US 7886108 B2	08/02/2011
		US 7889544 B2	15/02/2011
		US 7890846 B2	15/02/2011
		US 7930531 B2	19/04/2011
		US 7934037 B2	26/04/2011
		US 7934074 B2	26/04/2011
		US 7941916 B1	17/05/2011
		US 7944702 B2	17/05/2011
		US 7944703 B2	17/05/2011
		US 7953931 B2	31/05/2011
		US 7962836 B1	14/06/2011
		US 7965546 B2	21/06/2011
		US 7966429 B2	21/06/2011

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/015877

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 7966462 B2	21/06/2011
		US 7984303 B1	19/07/2011
		US 7987006 B2	26/07/2011
		US 8014130 B1	06/09/2011
		US 8015348 B2	06/09/2011
		US 8021166 B1	20/09/2011
		US 8037234 B2	11/10/2011
		US 8043099 B1	25/10/2011
		US 8060670 B2	15/11/2011
		US 8073985 B1	06/12/2011
		US 8078794 B2	13/12/2011
		US 8095971 B2	10/01/2012
		US 8102658 B2	24/01/2012
		US 8102662 B2	24/01/2012
		US 8108590 B2	31/01/2012
		US 8112574 B2	07/02/2012
		US 8116083 B2	14/02/2012
		US 8141240 B2	27/03/2012
		US 8166221 B2	24/04/2012
		US 8171204 B2	01/05/2012
		US 8180931 B2	15/05/2012
		US 8200862 B2	12/06/2012
		US 8241047 B2	14/08/2012
		US 8254134 B2	28/08/2012
		US 8262416 B2	11/09/2012
		US 8266367 B2	11/09/2012
		US 8296467 B2	23/10/2012
		US 8341332 B2	25/12/2012
		US 8452912 B2	28/05/2013
		US RE40115 E1	26/02/2008
US 2013-0067147 A1	14/03/2013	JP 2013-061795 A	04/04/2013
		US 2014-0237170 A1	21/08/2014