



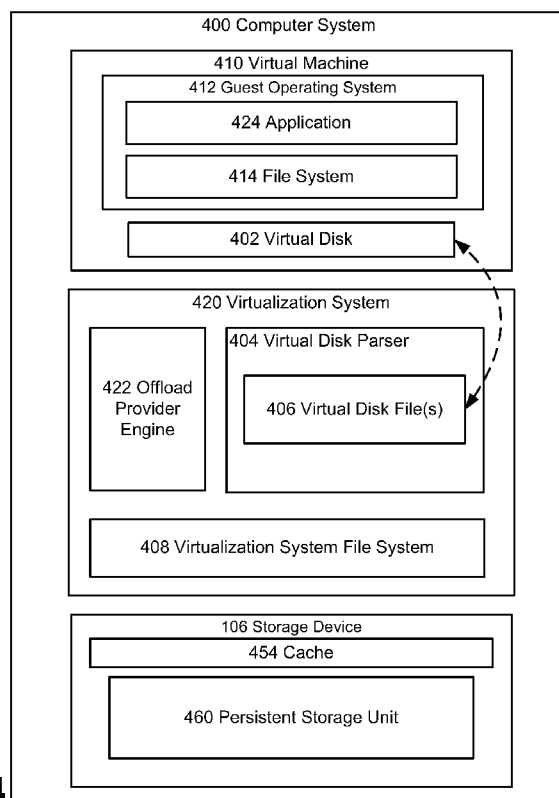
- (51) **International Patent Classification:**  
*G06F 12/08* (2006.01) *G06F 9/06* (2006.01)
- (21) **International Application Number:**  
PCT/US2012/027645
- (22) **International Filing Date:**  
4 March 2012 (04.03.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
13/046,617 11 March 2011 (11.03.2011) US
- (71) **Applicant (for all designated States except US):** **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) **Inventors:** **STARKS, John A.**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **GREEN, Dustin L.**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(US). **HARRIS, Todd William**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **JOHN, Mathew**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **RAJARAM, Senthil**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **MEHRA, Karan**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **CHRISTIANSEN, Neal R.**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **DAI, Chung Lang**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US).

- (81) **Designated States (unless otherwise indicated, for every kind of national protection available):** AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR,

[Continued on next page]

(54) **Title:** VIRTUAL DISK STORAGE TECHNIQUES



**FIG. 4**

(57) **Abstract:** This document describes techniques for storing virtual disk payload data. In an exemplary configuration, each virtual disk extent can be associated with state information that indicates whether the virtual disk extent is described by a virtual disk file. Under certain conditions the space used to describe a virtual disk extent can be reclaimed and state information can be used to determine how read and/or write operations directed to the virtual disk extent are handled. In addition to the foregoing, other techniques are described in the claims, figures, and detailed description of this document.



KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

**(84) Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS,

SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Declarations under Rule 4.17:**

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

**Published:**

- without international search report and to be republished upon receipt of that report (Rule 48.2(g))

## VIRTUAL DISK STORAGE TECHNIQUES

## BACKGROUND

[0001] Storage virtualization technology allows for the separation of logical storage from physical storage. One exemplary use case for storage virtualization is within a virtual machine. A layer of virtualizing software (typically called a hypervisor or virtual machine monitor) is installed on a computer system and controls how virtual machines interact with the physical hardware. Since guest operating systems are typically coded to exercise exclusive control over the physical hardware, the virtualizing software can be configured to subdivide resources of the physical hardware and emulate the presence of physical hardware within the virtual machines. Another use case for storage virtualization is within a computer system configured to implement a storage array. In this case, physical computer systems or virtual machines can be connected to the storage array using the iSCSI protocol, or the like.

[0002] A storage handling module can be used to emulate storage for either a virtual or physical machine. For example, a storage handling module can handle storage IO jobs issued by a virtual or physical machine by reading and writing to one or more virtual disk files, which can be used to describe, i.e., store, the extents of the virtual disk, i.e., a contiguous area of storage such as a block. Likewise, the storage handling program can respond to write requests by writing bit patterns data for the virtual disk to one or more virtual disk files and respond to read requests by reading the bit patterns stored in the one or more virtual disk files.

## SUMMARY

[0003] This document describes techniques for storing data for a virtual disk in one or more virtual disk files. In an exemplary configuration, a virtual disk extent can be associated with state information that indicates whether the virtual disk extent is described by a virtual disk file. Under certain conditions, the space used to describe the virtual disk extent can be reclaimed and state information can be used to determine how to handle subsequent read and/or write operations directed to the virtual disk extent. Reclaimed space, e.g., an extent built from one or more ranges, can be used to describe the same or another virtual disk extent. In addition to the foregoing, other techniques are described in the claims, the detailed description, and the figures.

[0004] It can be appreciated by one of skill in the art that one or more various aspects of the disclosure may include but are not limited to circuitry and/or programming for effecting the herein-referenced aspects; the circuitry and/or programming can be virtually

any combination of hardware, software, and/or firmware configured to effect the herein-referenced aspects depending upon the design choices of the system designer.

[0005] The foregoing is a summary and thus contains, by necessity, simplifications, generalizations and omissions of detail. Those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0006] FIG. 1 depicts a high-level block diagram of a computer system.

[0007] FIG. 2 depicts a high-level block diagram of an exemplary architecture for a virtualizing software program.

[0008] FIG. 3 depicts a high-level block diagram of an alternative architecture for a virtualizing software program.

[0009] FIG. 4 depicts a lower-level block diagram of a computer system configured to effectuate a virtual disk.

[0010] FIG. 5A depicts a lower-level block diagram of a computer system configured to effectuate a virtual disk.

[0011] FIG. 5B illustrates a lower-level block diagram of a computer system configured to effectuate a virtual disk.

[0012] FIG. 6 depicts a high-level block diagram of a differencing disk.

[0013] FIG. 7 depicts a high-level illustration of the relationship between a virtual disk and a virtual disk file.

[0014] FIG. 8 depicts a high-level illustration of the relationship between a virtual disk and a virtual disk file.

[0015] FIG. 9 depicts a high-level illustration of the relationship between a virtual disk and a virtual disk file.

[0016] FIG. 10 depicts a high-level illustration of the relationship between a virtual disk and a virtual disk file.

[0017] FIG. 11 depicts an operational procedure that can be implemented in a computer-readable storage medium and/or executed by a computer system.

[0018] FIG. 12 depicts additional operations that can be executed in conjunction with those illustrated by FIG. 11.

[0019] FIG. 13 depicts additional operations that can be executed in conjunction with those illustrated by FIG. 12.

[0020] FIG. 14 depicts an operational procedure that can be implemented in a computer-readable storage medium and/or executed by a computer system.

[0021] FIG. 15 depicts additional operations that can be executed in conjunction with those illustrated by FIG. 14.

[0022] FIG. 16 depicts an operational procedure that can be implemented in a computer-readable storage medium and/or executed by a computer system.

5 [0023] FIG. 17 depicts additional operations that can be executed in conjunction with those illustrated by FIG. 16.

#### DETAILED DESCRIPTION

[0024] The disclosed subject matter may use one or more computer systems. FIG. 1 and the following discussion are intended to provide a brief general description of a suitable  
10 computing environment in which the disclosed subject matter may be implemented.

[0025] The term circuitry used throughout can include hardware components such as hardware interrupt controllers, hard drives, network adaptors, graphics processors, hardware based video/audio codecs, and the firmware used to operate such hardware. The term circuitry can also include microprocessors, application specific integrated circuits,  
15 and processors, e.g., cores of a multi-core general processing unit that perform the reading and executing of instructions, configured by firmware and/or software. Processor(s) can be configured by instructions loaded from memory, e.g., RAM, ROM, firmware, and/or mass storage, embodying logic operable to configure the processor to perform a function(s). In an example embodiment, where circuitry includes a combination of  
20 hardware and software, an implementer may write source code embodying logic that is subsequently compiled into machine readable code that can be executed by hardware. Since one skilled in the art can appreciate that the state of the art has evolved to a point where there is little difference between hardware implemented functions or software implemented functions, the selection of hardware versus software to effectuate herein  
25 described functions is merely a design choice. Put another way, since one of skill in the art can appreciate that a software process can be transformed into an equivalent hardware structure, and a hardware structure can itself be transformed into an equivalent software process, the selection of a hardware implementation versus a software implementation is left to an implementer.

30 [0026] Referring now to FIG. 1, an exemplary computing system 100 is depicted. Computer system 100 can include processor 102, e.g., an execution core. While one processor 102 is illustrated, in other embodiments computer system 100 may have multiple processors, e.g., multiple execution cores per processor substrate and/or multiple processor substrates that could each have multiple execution cores. As shown by the

figure, various computer-readable storage media 110 can be interconnected by one or more system busses which couples various system components to the processor 102. The system buses may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus

5 architectures. In example embodiments the computer-readable storage media 110 can include for example, random access memory (RAM) 104, storage device 106, e.g., electromechanical hard drive, solid state hard drive, etc., firmware 108, e.g., FLASH RAM or ROM, and removable storage devices 118 such as, for example, CD-ROMs, floppy disks, DVDs, FLASH drives, external storage devices, etc. It should be

10 appreciated by those skilled in the art that other types of computer readable storage media can be used such as magnetic cassettes, flash memory cards, and/or digital video disks.

**[0027]** The computer-readable storage media 110 can provide nonvolatile and volatile storage of processor executable instructions 122, data structures, program modules and other data for the computer system 100 such as executable instructions. A basic

15 input/output system (BIOS) 120, containing the basic routines that help to transfer information between elements within the computer system 100, such as during start up, can be stored in firmware 108. A number of programs may be stored on firmware 108, storage device 106, RAM 104, and/or removable storage devices 118, and executed by processor 102 including an operating system and/or application programs. In exemplary  
20 embodiments, computer-readable storage media 110 can store virtual disk parser 404, which is described in more detail in the following paragraphs, can be executed by processor 102 thereby transforming computer system 100 into a computer system configured for a specific purpose, i.e., a computer system configured according to techniques described in this document.

25 **[0028]** Commands and information may be received by computer system 100 through input devices 116 which can include, but are not limited to, a keyboard and pointing device. Other input devices may include a microphone, joystick, game pad, scanner or the like. These and other input devices are often connected to processor 102 through a serial port interface that is coupled to the system bus, but may be connected by other interfaces,  
30 such as a parallel port, game port, or universal serial bus (USB). A display or other type of display device can also be connected to the system bus via an interface, such as a video adapter which can be part of, or connected to, a graphics processor unit 112. In addition to the display, computers typically include other peripheral output devices, such as speakers and printers (not shown). The exemplary system of FIG. 1 can also include a host adapter,

Small Computer System Interface (SCSI) bus, and an external storage device connected to the SCSI bus.

**[0029]** Computer system 100 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer. The remote  
5 computer may be another computer, a server, a router, a network PC, a peer device or other common network node, and typically can include many or all of the elements described above relative to computer system 100.

**[0030]** When used in a LAN or WAN networking environment, computer system 100 can be connected to the LAN or WAN through network interface card 114. The NIC 114,  
10 which may be internal or external, can be connected to the system bus. In a networked environment, program modules depicted relative to the computer system 100, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections described here are exemplary and other means of establishing a communications link between the computers may be used. Moreover, while it is  
15 envisioned that numerous embodiments of the present disclosure are particularly well-suited for computerized systems, nothing in this document is intended to limit the disclosure to such embodiments.

**[0031]** Turning to FIG. 2, illustrated is an exemplary virtualization platform that can be used to generate virtual machines. In this embodiment, microkernel hypervisor 202 can be  
20 configured to control and arbitrate access to the hardware of computer system 200. Microkernel hypervisor 202 can generate execution environments called partitions such as child partition 1 through child partition N (where N is an integer greater than 1). Here, a child partition is the basic unit of isolation supported by microkernel hypervisor 202. Microkernel hypervisor 202 can isolate processes in one partition from accessing another  
25 partition's resources. In particular, microkernel hypervisor 202 can isolate kernel mode code of a guest operating system from accessing another partition's resources as well as user mode processes. Each child partition can be mapped to a set of hardware resources, e.g., memory, devices, processor cycles, etc., that is under control of the microkernel hypervisor 202. In embodiments, microkernel hypervisor 202 can be a stand-alone  
30 software product, a part of an operating system, embedded within firmware of the motherboard, specialized integrated circuits, or a combination thereof.

**[0032]** Microkernel hypervisor 202 can enforce partitioning by restricting a guest operating system's view of the memory in a physical computer system. When microkernel hypervisor 202 instantiates a virtual machine, it can allocate pages, e.g., fixed

length blocks of memory with starting and ending addresses, of system physical memory (SPM) to the virtual machine as guest physical memory (GPM). Here, the guest's restricted view of system memory is controlled by microkernel hypervisor 202. The term guest physical memory is a shorthand way of describing a page of memory from the viewpoint of a virtual machine and the term system physical memory is shorthand way of describing a page of memory from the viewpoint of the physical system. Thus, a page of memory allocated to a virtual machine will have a guest physical address (the address used by the virtual machine) and a system physical address (the actual address of the page).

**[0033]** A guest operating system may virtualize guest physical memory. Virtual memory is a management technique that allows an operating system to over commit memory and to give an application sole access to a logically contiguous working memory. In a virtualized environment, a guest operating system can use one or more page tables, called guest page tables in this context, to translate virtual addresses, known as virtual guest addresses into guest physical addresses. In this example, a memory address may have a guest virtual address, a guest physical address, and a system physical address.

**[0034]** In the depicted example, parent partition component, which can also be also thought of as similar to domain 0 of Xen's open source hypervisor can include a host environment 204. Host environment 204 can be an operating system (or a set of configuration utilities) and host environment 204 can be configured to provide resources to guest operating systems executing in the child partitions 1-N by using virtualization service providers 228 (VSPs). VSPs 228, which are typically referred to as back-end drivers in the open source community, can be used to multiplex the interfaces to the hardware resources by way of virtualization service clients (VSCs) (typically referred to as front-end drivers in the open source community or para virtualized devices). As shown by the figures, virtualization service clients execute within the context of guest operating systems. However, these drivers are different than the rest of the drivers in the guest in they communicate with host environment 204 via VSPs instead of communicating with hardware or emulated hardware. In an exemplary embodiment the path used by virtualization service providers 228 to communicate with virtualization service clients 216 and 218 can be thought of as the enlightened IO path.

**[0035]** As shown by the figure, emulators 234, e.g., virtualized IDE devices, virtualized video adaptors, virtualized NICs, etc., can be configured to run within host environment 204 and are attached to emulated hardware resources, e.g., IO ports, guest physical address ranges, virtual VRAM, emulated ROM ranges, etc. available to guest operating systems



220 and 222. For example, when a guest OS touches a guest virtual address mapped to a guest physical address where a register of a device would be for a memory mapped device, microkernel hypervisor 202 can intercept the request and pass the values the guest attempted to write to an associated emulator. Here, the emulated hardware resources in this example can be thought of as where a virtual device is located in guest physical address space. The use of emulators in this way can be considered the emulation path. The emulation path is inefficient compared to the enlightened IO path because it requires more CPU time to emulate devices than it does to pass messages between VSPs and VSCs. For example, several actions on memory mapped to registers are required in order to write a buffer to disk via the emulation path, while this may be reduced to a single message passed from a VSC to a VSP in the enlightened IO path, in that the drivers in the VM are designed to access IO services provided by the virtualization system rather than designed to access hardware.

**[0036]** Each child partition can include one or more virtual processors (230 and 232) that guest operating systems (220 and 222) can manage and schedule threads to execute thereon. Generally, the virtual processors are executable instructions and associated state information that provide a representation of a physical processor with a specific architecture. For example, one virtual machine may have a virtual processor having characteristics of an Intel x86 processor, whereas another virtual processor may have the characteristics of a PowerPC processor. The virtual processors in this example can be mapped to processors of the computer system such that the instructions that effectuate the virtual processors will be directly executed by physical processors. Thus, in an embodiment including multiple processors, virtual processors can be simultaneously executed by processors while, for example, other processor execute hypervisor instructions. The combination of virtual processors and memory in a partition can be considered a virtual machine.

**[0037]** Guest operating systems (220 and 222) can be any operating system such as, for example, operating systems from Microsoft®, Apple®, the open source community, etc. The guest operating systems can include user/kernel modes of operation and can have kernels that can include schedulers, memory managers, etc. Generally speaking, kernel mode can include an execution mode in a processor that grants access to at least privileged processor instructions. Each guest operating system can have associated file systems that can have applications stored thereon such as terminal servers, e-commerce servers, email servers, etc., and the guest operating systems themselves. The guest operating systems can

schedule threads to execute on the virtual processors and instances of such applications can be effectuated.

**[0038]** Referring now to FIG. 3, it illustrates an alternative virtualization platform to that described above in FIG. 2. FIG. 3 depicts similar components to those of FIG. 2;

5 however, in this example embodiment hypervisor 302 can include a microkernel component and components similar to those in host environment 204 of FIG. 2 such as the virtualization service providers 228 and device drivers 224, while management operating system 304 may contain, for example, configuration utilities used to configure hypervisor 302. In this architecture, hypervisor 302 can perform the same or similar functions as  
10 microkernel hypervisor 202 of FIG. 2; however, in this architecture hypervisor 304 effectuates the enlightened IO path and includes the drivers for the physical hardware of the computer system. Hypervisor 302 of FIG. 3 can be a stand-alone software product, a part of an operating system, embedded within firmware of the motherboard or a portion of hypervisor 302 can be effectuated by specialized integrated circuits.

15 **[0039]** Turning now to FIG. 4, it describes computer system 400, which illustrates a high-level block diagram of components that can be used to effect the techniques described in this document. Briefly, computer system 400 can include components similar to those described above with respect to FIG. 1 through 3. FIG. 4 shows virtualization system 420, which can be thought of as a high-level representation of the virtualization platform  
20 illustrated by FIG. 2 or FIG. 3. For example, virtualization system 420 can be thought of as a high-level representation of the combination of features provided by microkernel hypervisor 202 and host environment 204. Alternatively, virtualization system 420 can be thought of as a high-level representation of hypervisor 302 and management OS 304.

Thus, use of the term “virtualization system 420” throughout this document signals that  
25 the virtual disk techniques described in the following paragraphs can be implemented within any type of virtualization software layer or in any type of virtualization platform.

**[0040]** Virtualization system 420 can include offload provider engine 422. Briefly, offload provider engine 422 can be configured to service offload read and offload write requests (sometimes called PROXY READ and PROXY WRITE) issued by, for example,  
30 application 424. An offload read request is a request to create a token that represents data that would have been read if the offload read had been a normal read. An offload write is a request to write the data represented by a token to a destination location. In one usage example, an offload read followed by an offload write can be used to copy data from one location to another, e.g., from computer system 400 to a destination computer system

within a domain by using tokens that represent the data to avoid moving the data through local RAM. For example, suppose that computer system 400 and a destination computer system (not shown) can access a common data repository and a request to copy data from computer system to the destination is received. Instead of copying the data to the

5 destination, application 424 can issue a request to offload provider engine 422 to issue a token that represents the data as it exists at the time the token is associated with the data. The token can be sent to the destination and used by a program running on the destination to obtain the data from the common data storage repository and write the data to the destination. Copy-offload techniques are described in more detail in co-pending US  
10 Patent Application No. 12/888,433, entitled "Offload Reads and Writes" and U.S. Patent Application No. 12/938,383, entitled "Virtualization and Offload Reads and Writes," the contents of which are herein incorporated by reference in their entirety to the extent they are consistent with techniques described in this document.

**[0041]** Virtual disk parser 404, which can be a module of executable instructions in a  
15 specific example embodiment, can be used to instantiate virtual disks from virtual disk files and handle storage IO on behalf of a virtual machine. As shown by the figure, virtual disk parser 404 can open one or more virtual disk files such as virtual disk file(s) 406 and generate virtual disk 402.

**[0042]** Virtual disk parser 404 can obtain virtual disk file(s) 406 from storage device 106  
20 via virtualization system file system 408. Briefly, virtualization system file system 408 represents a software module that organizes computer files and data of virtualization system 420, such as virtual disk file(s) 406. Virtualization system file system 408 can store this data in an array of fixed-size physical extents, i.e., contiguous areas of storage on a physical storage device. In a specific example, an extent can be a cluster, which is a  
25 sequence of bytes of bits having a set length. Exemplary cluster sizes are typically a power of 2 between 512 bytes and 64 kilobytes. In a specific configuration, a cluster size can be 4 kilobytes.

**[0043]** When a request to open virtual disk file 406 is received, virtualization system file  
30 system 408 determines where the file is located on disk and issues an IO job to the disk device driver to read the data from one or more physical extents of the disk. The IO job issued by file system 408 determines a disk offset and length that describes the location of the persistent copy of virtual disk file 406 on storage device 106 and issues the IO job to storage device 106. Due to the semantics of how storage devices operate, a write IO job can be buffered in one or more levels of caches of volatile memory, represented by cache

454, until the circuitry of storage device 106 determines to access the location on the persistent storage unit 460, e.g., a platter, a flash memory cell, etc., and write the buffered bit pattern indicative of the new contents of the persistent copy of the virtual disk file(s) 406 to persistent storage unit 460.

- 5   **[0044]** Virtual disk parser 404 can obtain the bit pattern indicative of virtual disk file(s) 406 and expose the payload, e.g., user data, in the virtual disk file(s) 406 as a disk including a plurality of virtual disk extents. In an embodiment, these virtual disk extents can be a fixed-size block 512 kilobytes up to 64 megabytes in size and partitioned into a plurality of sectors; however, in another embodiment the virtual disk extents could be
- 10   variable-sized extents. In an exemplary configuration, prior to booting guest operating system 412, resources related to an emulated or enlightened storage controller and emulated or enlightened aspects of a virtual disk are setup such that an emulated storage controller with memory mapped registers is effected within guest physical address space of the virtual machine 410. Boot code can run and boot guest operating system 412.
- 15   Virtualization system 420 can detect an attempt to access this region of guest physical address space and return a result that causes guest operating system 412 to determine that a storage device is attached to the emulated storage controller. In response, guest operating system 412 can load a driver (either a para virtualization driver or a regular driver) and use the driver to issue storage IO requests to the detected storage device.
- 20   Virtualization system 420 can route the storage IO requests to virtual disk parser 404. **[0045]** After guest operating system 412 is running it can issue IO jobs to virtual disk 402 via file system 414, which is similar to virtualization system file system 414 in that it organizes computer files and data of guest operating system 412 and applications installed on guest operating system 412. Guest operating system 412 can interact with virtual disk
- 25   402 in a way that is similar to how an operating system interacts with a physical storage device and eventually the IO jobs are routed to virtual disk parser 404. Virtual disk parser 404 can include logic for determining how to respond to the IO jobs in a way that emulates a physical storage device. For example, virtual disk parser 404 can read data from virtual disk file(s) 406 and write data to virtual disk file(s) 406. The data written to virtual disk
- 30   file(s) 406 in turn is routed through virtualization system file system 408 and committed to a persistent copy of virtual disk file(s) 406 stored on or in persistent storage unit 460. **[0046]** Referring briefly to FIG. 5A, it illustrates an alternative architecture for implementing techniques described in this document. As shown by FIG. 5, virtual disk parser 404 can also be implemented in an operating system 502 such as an operating

system offered by Microsoft®. In this example, virtual disk parser 404 can be configured to run on storage server 500, which could include components similar to computer system 100 of FIG. 1. In this example, storage server 500 could include an array of physical storage devices 510 and can be configured to make storage available to servers such that the storage appears as locally attached to operating system 508. Virtual disk parser 404 can operate the same as it was described with respect to FIG. 4; the difference being in this configuration read/write IO jobs issued by file system 414 can be routed over a network connection to virtual disk parser 404.

[0047] Referring briefly to FIG. 5B, it illustrates yet another architecture for implementing techniques described in this document. FIG. 5B is similar to FIG. 5A in that virtual disk parser 404 is implemented in operating system 502 and computer system 512 could include components similar to computer system 100 of FIG. 1. The difference in this example; however, is that the figure illustrates a loopback-attached virtual disk 402. File system 414, including applications such as application 424 can be stored in virtual disk 402 and virtual disk file(s) 406 can be stored in computer system file system 514.

[0048] Turning attention now to virtual disk 402, while it can be effected by a single virtual disk file, in other configurations a group of differencing virtual disk files can be used to bring about virtual disk 402. FIG. 6 illustrates exemplary chains of virtual disk files that can be used by virtual disk parser 404 to effect virtual disk 402 as a differencing disk. Generally, a differencing virtual disk file represents the current state of a virtual disk as a set of modified extents in comparison to a parent image. The parent image can be another differencing virtual disk file or a base virtual disk file.

[0049] In an exemplary configuration, the linking between a parent virtual disk file and a child virtual disk file can be stored within the child. In particular, the child can include an identifier of the parent and a value that describes the location of the parent. When starting a virtual machine, virtual disk parser 404 may receive information that describes the last virtual disk file in the chain, i.e., virtual disk file 612 is the last in a chain that includes virtual disk files 612, 610, 606, and 600, and open this file. This file can include an identifier of its parent, i.e., virtual disk file 610, and a path to it. Virtual disk parser 404 can locate and open the parent and so on and so forth until a base virtual disk file is located and opened.

[0050] Virtual disk parser 404 can use information that indicates whether data is present or stored in a parent virtual disk file. Typically, the last virtual disk file in the chain is opened as read/modify and other virtual disk files are opened as read only. Thus, writes

are typically made to the last virtual disk file in the chain. Read operations are similarly directed first to the last virtual disk file in the chain and virtual disk parser 404 will logically search the virtual disk files in logical order from last to base until the data is found in the instance that information about where the data is located is not cached. In a specific example, a block allocation table (not shown) for a virtual disk file, e.g., virtual disk file 612, can include state information that indicates whether the virtual disk extent is defined by a section of the virtual disk file or if this virtual disk extent is transparent, e.g., defined by a different virtual disk file further along the chain. In one implementation, virtual disk parser 404 can determine whether this virtual disk extent is transparent and access the block allocation table for the next virtual disk file in the chain, e.g., virtual disk file 610, and so on and so forth until a virtual disk file in the chain is located that defines the data.

**[0051]** Referring now to FIG. 7, it illustrates virtual disk 402 described at least in part by virtual disk file 702, which could be similar to any virtual disk file described in FIG. 6 that is write/modifiable, e.g., virtual disk file 602, 608, or 612, or a single virtual disk file. As shown by the figure, virtual disk 402 can include N extents of storage (where N is an integer greater than 0) and in this specific example virtual disk 402 includes 10 extents. Virtual disk 402 is illustrated as including the bit patterns for different files and data for guest operating system 412, which are differentiated by the different patterns within the virtual disk extents.

**[0052]** Since virtual disk 402 is not a physical storage device, the underlying payload data for the virtual disk extents can be “described by,” i.e., stored in, different sections within virtual disk file 702. For example, virtual disk extent 1 is described by a section that is defined by a virtual disk file offset value 0 or the first offset that can be used to store payload data. Allocation table 416, which can be stored in random access memory while computer system 400 is in operation, can be persisted in virtual disk file 702 in any section and can span multiple sections. Briefly, allocation table 416 can include information that links virtual disk extents to sections of virtual disk file 702. For example, allocation table 416 can store information that defines the virtual disk file byte offsets that define the section of virtual disk file 702 that stores the data. The arrows signify the relationships stored in allocation table 416.

**[0053]** Described in more detail in the following paragraphs, allocation table 416 can also include state information; however, this configuration is exemplary. In alternate configurations this information can be stored in a different section of virtual disk file 702

and loaded into RAM 104. Allocation table 416 can include an entry for each virtual disk extent; state information indicating what state each extent is in; and a file offset indicating where in virtual disk file 702 each virtual disk extent is described (not illustrated). In an alternative embodiment an extent could also be defined by multiple already-mapped and contiguous (in file offset) table entries. In this configuration, reads and writes that cross block boundaries can be serviced as a single read/write to virtual disk file 702 if the block payloads are contiguous in the file. In a specific example, virtual disk parser 404 can also store information that indicates what type of bit pattern is stored in each unused section of the virtual disk file, i.e., a free space map. In addition to the foregoing, the free-space map can allow be used by virtual disk parser 404 to determine which sectors of virtual disk file 406 are used and which are free. The free space map in this example can be configured to track free space in the file that is non-zero. In an exemplary embodiment, because using a non-zero portion of free space to describe a portion of virtual disk 402, which must be zero or must not disclose information from other virtual disk offsets, the free space is overwritten with zeros or a non-information disclosing pattern (typically zeros), respectively. Virtual disk parser 404 can use this information in order to determine what section of virtual disk file to allocate to a virtual disk extent. For example, if a virtual disk extent in the zero state is written to, virtual disk parser 404 can allocate a section that already has zeros in it to back the virtual disk extent.

**[0054]** As guest operating system 412 or operating system 508 runs it will generate data and files and issue disk writes to virtual disk 402 to store data. When virtual disk file 702 does not have any additional non-used space, virtual disk parser 404 can extend the end of file and use the new space to describe the virtual disk extents. Guest operating system 412 or operating system 508 may use, delete, and reuse sections of virtual disk 402; however, since virtual disk parser 404 is merely storing data on behalf of file system 414, virtual disk parser 404 may be unable to determine whether a section of virtual disk file is still being used by guest operating system 412. Consequently, virtual disk parser 404 may hold allocated space in virtual disk file 702 to describe virtual disk extents that are no longer in use by file system 414. The result of this is that the size of virtual disk file 702 may grow until it reaches the size of virtual disk 402.

**[0055]** In exemplary embodiments, virtual disk parser 404 can be configured to reclaim unused sections of a virtual disk file and optionally reuse them. As such, the frequency at which the virtual disk file needs to be extended is reduced, and the overall size of the virtual disk file is reduced. In an example embodiment, when a file system signals that it

is no longer using a virtual disk extent, virtual disk parser 404 can de-allocate, i.e., unlink, the virtual disk extent from the virtual disk file and associate the virtual disk extent with information that describes how read operations to the virtual disk extent should be treated.

The section of the virtual disk file can then be reused to describe the same or another

5 virtual disk extent.

**[0056]** In an exemplary configuration, virtual disk parser 404 can use TRIM, UNMAP, and/or WRITE SAME of zero commands issued by a file system to determine when a virtual disk extent can be de-allocated from virtual disk file(s) 406. TRIM commands can be issued by guest operating system 412 or operating system 508. For example, as guest  
10 operating system 412 or operating system 508 runs, file system 414 may determine that some sectors are no longer needed and issue a TRIM command. Alternatively or additionally, virtual disk parser 404 can be configured to request that file system 414 issue TRIM commands at predetermined intervals, or when predetermined criteria are satisfied, e.g., when virtual machine 410 is instantiated, when virtual machine 410 is shut down,  
15 under light use, etc.

**[0057]** Briefly, a TRIM command is used to inform the data storage device as to which sectors are no longer considered in use so that the data stored therein can be optionally discarded by the data storage device. One type of TRIM command, called a free space TRIM command, can be used by file system 414 to signal that sectors are no longer in use  
20 by file system 414 and the other, called a standard TRIM command, does not. The difference between the two types of TRIM commands is that when a sector is the subject of a free space TRIM, file system 414 provides security for the sector by preventing user space applications and the like from reading from the sector. The fact that file system 414 secures access to sectors that have been trimmed in this way can be used to increase the  
25 ability to efficiency allocate virtual disk file space. This particular aspect is described in more detail in the following paragraphs.

**[0058]** In an exemplary configuration, virtual disk parser 404 can be configured to execute reclamation operations when a virtual disk extent is fully covered by a TRIM command. Or put another way, virtual disk parser 404 can unlink virtual disk extents from the virtual  
30 disk file in response to receipt of a TRIM command that defines a range of virtual disk sectors that identifies all of the sectors in the virtual disk extent. In the same or an alternative embodiment, when a TRIM command is received that covers a portion of a virtual disk extent, virtual disk parser 404 can determine what portion of the virtual disk file corresponds to the trimmed sectors and send a TRIM command for the portion of the



virtual disk file to storage device 106. In this example, the underlying file system, e.g., virtualization system file system 408, storage server file system 504, or computer system file system 514 can translate the offsets of the TRIM command and send the translated offsets to storage device 106, reclaim space directly via internal data structure updates, or clear data from caches.

**[0059]** In the same or another embodiment, when a TRIM command is received that covers a portion of a virtual disk extent, virtual disk parser 404 can be configured to store information that indicates what sectors have been the subject of the TRIM command and whether the TRIM command was a free space trim or not. In the instance that the remainder of the virtual disk extent is trimmed, virtual disk parser 404 can de-allocate the virtual disk extent from the virtual disk file.

**[0060]** When de-allocating virtual disk extents, virtual disk parser 404 can associate the virtual disk extent with state information that describes how read operations directed to the virtual disk extent can be handled. Table 1 illustrates exemplary state information that virtual disk parser 404 can associate with virtual disk extents and use to optimize the reclamation of the virtual disk file. The ability to reclaim a virtual disk extent can be accomplished in one example by using two states (described and not described); however, since the bit pattern stored in virtual disk file 702 is not typically erased when the data is deleted, additional states can be used to determine when space selected to describe a virtual disk extent needs to be cleared before it can be reused or if it can be reused without overwriting the data previously stored therein. One reason for why the data is not erased upon deletion is that it costs processor cycles to erase data and since some storage devices are configured to perform write operations on a per-block basis, it is more efficient to erase data when over-writing with new data. The following states are exemplary and the disclosure is not limited to using states that are defined by the following table.

State	Description
Mapped	This state indicates that the virtual disk extent is linked to the virtual disk file.
Transparent	This state indicates that the virtual disk extent is defined by a different virtual disk file.
Zero	This state indicates that the virtual disk extent is not described by the virtual disk file. In addition, this state indicates that the virtual disk extent is defined as zero and that the zeros are meaningful.

Unmapped	This state indicates that the virtual disk extent is not described by the virtual disk file. In an embodiment, this state can include sub-states anchored and unanchored.
Uninitialized	This state indicates that the virtual disk extent is not described by the virtual disk file and that the virtual disk extent is defined as free space. In an embodiment, this state can also include sub-states anchored and unanchored.

**Table 1**

[0061] Referring to Table 1 in conjunction with FIG. 7, the first state listed is the “mapped” state, which indicates that the virtual disk extent is described by a section of virtual disk file 702. For example, virtual disk extent 0 is an example virtual disk extent that is illustrated as being in the “mapped” state.

[0062] Continuing with the description of Table 1, a virtual disk extent can be associated with state information that indicates that the virtual disk extent is “transparent,” that is, the virtual disk extent is described by a different virtual disk file. In the instance that a read operation is received by virtual disk parser 404 to a virtual disk extent in the transparent state, virtual disk parser 404 can refer to a different virtual disk file and check its allocation table to determine how to respond to the read. In an instance that virtual disk parser 404 receives a write to the virtual disk extent, virtual disk parser 404 can transition the virtual disk extent from the “transparent” state to the “mapped” state.

[0063] Continuing with the description of Table 1 in conjunction with FIG. 7, a virtual disk extent can also be associated with the “unmapped” state. In this example, the virtual disk extent is not described by virtual disk file 702 nor is it described by any other virtual disk file in a chain. In this example, the unmapped state can be used to describe a virtual disk extent that was subject to a TRIM command that did not indicate that file system 414 would secure access to the virtual disk extent. Or put another way, the TRIM command used to transition this virtual disk extent to this state was a standard TRIM command. In the instance that the virtual disk extent is in the unmapped state and an IO job indicative of a read to the extent is received, virtual disk parser 404 can respond with zeros, the zero token, ones, a token representing all ones, or a non-information disclosing bit pattern, e.g., all zeros, all ones, or a randomly generated pattern of ones and zeros. In this example, if a section of virtual disk file 702 is allocated to back a virtual disk extent in this state, virtual disk parser 404 can write a non-information disclosing bit pattern to the section of virtual

disk file 702 before allocating it or select a section that already includes a non-information disclosing bit pattern to describe the virtual disk extent. Virtual disk extent 6 of FIG. 7 is indicated as in the unmapped state.

**[0064]** In an embodiment, the data defining an unmapped or uninitialized extent can be kept and the unmapped or uninitialized state can include two sub-states: anchored, which means that the data is still present within virtual disk file 702, and unanchored, which means that the data may or may not be kept. In instances where these sub-states are used, virtual disk parser 404 can transition an unmapped but anchored extent to mapped by allocating the section or sections that store the data without zeroing the section or sections. Similarly, while virtual disk parser 404 is configured to treat uninitialized extents as if they were unmapped for at least a portion of virtual disk 402, virtual disk parser 404 can avoid zeroing an uninitialized but anchored extent during transition of that extent to mapped, by allocating the section or sections that store the data without zeroing the section or sections.

**[0065]** Table 1 additionally describes a “zero” state. In this example, the virtual disk extent is not described by virtual disk file 702 nor is it described by any other virtual disk file in a chain; however, the virtual disk extent is required to read as all zeros. In this example, the zero state can be used to describe a virtual disk extent that was subject to either type of TRIM command or to describe a virtual disk extent that a program has written all zeros to. For example, suppose a deletion utility program wrote all zeros to virtual disk extent 4 to ensure that the data it previously stored was completely overwritten. In the instance that the virtual disk extent is in the zeroed state, and an IO job indicative of a read to the extent is received, virtual disk parser 404 can respond to with zeros or the zero token (in an offload read operation). In the instance that a write is directed to a virtual disk extent in this state, virtual disk parser 404 can zero a section of virtual disk file 702 and use it to describe the virtual disk extent or select a section of virtual disk file 702 that is already zero and allocate it to back the virtual disk extent. In this embodiment, zeroed space could be tracked using a data structure or virtual disk file 702. The data structure could be updated periodically, when virtual disk file 702 is opened, when virtual disk file 702 is closed, etc. A read from an extent in the unmapped or uninitialized states may optionally cause virtual disk parser 404 to transition the extent to the zero state in a configuration where virtual disk parser 404 is configured to provide sector stability for extents in the unmapped or uninitialized states.

[0066] Table 1 also describes a state called the “uninitialized” state. The uninitialized state indicates that the virtual disk extent is not described by virtual disk file 702 and file system 414 is securing access to the virtual disk extent. That is, file system 414 is configured to prevent user applications from reading sectors within this virtual disk extent.

5 In this example, the uninitialized state can be used to describe a virtual disk extent that was subject to a free space TRIM command. In the instance that the virtual disk extent is in the uninitialized state and an IO job indicative of a read to the extent is received, virtual disk parser 404 can respond with any data, i.e., a bit pattern from almost anywhere else in virtual disk file 702, zeros, ones, a non-information disclosing bit pattern, etc., because  
10 virtual disk parser 404 is not providing security for the virtual disk extent, beyond the requirement that only virtual disk payload data and non-security-impacting metadata may be exposed to the virtual disk client. In the instance that a write is directed to a virtual disk extent in this state, virtual disk parser 404 can simply allocate a section of the virtual disk file 702 without having to alter any data that may be stored within the section.  
15 Consequently, this state is the most advantageous because space can be allocated within the virtual disk file without clearing it beforehand. Virtual disk extent 5 of FIG. 7 is indicated as in the uninitialized state and virtual disk file 702 is not backing the virtual disk extent.

[0067] Once state information is associated with each virtual disk extent, virtual disk  
20 parser 404 can be configured to provide additional information to an administrator or the like about how virtual disk 402 is arranged. In an example embodiment, virtual disk parser 404 can be configured to respond to offset queries that include certain parameters based on the state information. For example, a user can issue a query to iterate, starting at a given byte offset, through virtual disk 402 and locate ranges that satisfy a specific  
25 criteria such as “mapped,” “unmapped,” “transparent,” etc. In addition, a user can select how “deep” the query should go to take into account differencing virtual disk files. For example, and referring to FIG. 7, a user could set a depth of 2 and execute a query. In response, virtual disk parser 404 will execute the query on the last two virtual disk files in a chain, e.g., virtual disk files 610 and 612. Specific queries can include a query to obtain  
30 the next non-transparent range(s), the next non-zero range(s), the next defined range(s), the next initialized range(s), etc. Briefly, a query for the next defined range can be configured to return the next range(s) which contain defined data (e.g., sectors in the mapped or zeroed state, with transparent sectors resolving to the parent virtual disk file’s state for that sector). A query for the next initialized range(s) can return the next range(s)

which contain data in a state other than the uninitialized state, with transparent sectors resolving to the parent virtual disk file's state for that sector.

**[0068]** Turning now to FIG. 8, it illustrates a specific example of how virtual disk parser

404 can transition virtual disk extents from one state to another in response to a file or  
5 other data being saved to virtual disk 402. For example, suppose that a user uses a  
database management program within virtual machine 410 and creates a database. The  
user can save the database in a file and file system 414 can determine where on virtual  
disk 402 to save file 802. File system 414 can issue one or more disk writes to write file  
802 to, for example, sectors that fall within virtual disk extents 3-5. In this example,  
10 virtual disk extent 3 is "mapped" and virtual disk parser 404 can write the first portion of  
file 802 to the section identified by allocation table 416.

**[0069]** Virtual disk extents 4 and 5, on the other hand, are in the "zero" and "uninitialized"  
state. In this example, virtual disk parser 404 can select an unused section of virtual disk  
file 702 to back virtual disk extent 4 and determine that virtual disk extent 4 is in the  
15 zeroed state. In response to this determination, virtual disk parser 404 can zero the section  
that is going to be used to describe virtual disk extent 4 or locate a section which is already  
all zeros. After locating a zeroed section or the process of zeroing the section is complete,  
virtual disk parser 404 can generate information that identifies the virtual disk file byte  
offset indicative of the first byte of the section that defines where virtual disk extent 4 is  
20 described in virtual disk file 702 and store it in allocation table 416. Virtual disk parser  
404 can then change the state information associated with virtual disk extent 4 to indicate  
that it is "mapped." Then the portion of the write to extent 4 can be written to the located  
section.

**[0070]** Alternatively, for a portion of a write which covers an entire extent of the virtual  
25 disk currently in the zero state, a located section of the virtual disk file may be chosen, the  
portion of the write may be issued to the section, and upon completion of the write the  
virtual disk parser 404 can then change the state information associated with the virtual  
disk extent to indicate that the extent is "mapped". Alternatively, for a portion of a write  
which only covers part of a virtual disk extent currently in the zero state, a located section  
30 of the virtual disk file may be chosen, the portion of the write may be issued to the section,  
a zeroing write may be issued to the remainder of the section, and on completion of the  
write the virtual disk parser 404 can then change the state information associated with the  
virtual disk extent to indicate that the extent is "mapped". Those skilled in the art will

recognize that the given ordering of writes may be enforced using flush or write-through writes, such as force-unit-access writes.

[0071] Similarly, virtual disk parser 404 can select an unused section of virtual disk file 702 to back virtual disk extent 5 and determine that virtual disk extent 5 is in the uninitialized state by consulting allocation table 416. In response to this determination, virtual disk parser 404 can allocate the section to describe virtual disk extent 5 without modifying the contents of the selected section. Virtual disk parser 404 can generate information that identifies the virtual disk file byte offset indicative of the first byte of the section, which indicates where virtual disk extent 4 is described in virtual disk file 702 and store the file byte offset of the section in allocation table 416. Virtual disk parser 404 can then change the state information associated with virtual disk extent 5 to indicate that it is “mapped.”

[0072] FIG. 9 illustrates another specific example of how virtual disk parser 404 can transition virtual disk extents from one state to another in response to a deletion operation on file 802 and an operation that zeros the contents of virtual disk extent 7. For example, a user may have deleted file 802 and file system 414 may have issued a TRIM command. In this example, virtual disk parser 404 may receive a TRIM command that includes a range of virtual disk sectors that fully cover virtual disk extents 4 and 5 and partially cover virtual disk extent 3. In response to a determination that virtual disk extent 4 and 5 are fully trimmed, virtual disk parser 404 can be configured to remove the linking from allocation table 416 and transition virtual disk extent 4 to a state that indicates that virtual disk file 702 is not backing this virtual disk extent. As shown by the allocation table entry for virtual disk extent 4, the state virtual disk parser 404 transitions the virtual disk extent to depends on what states virtual disk parser 404 is configured to use and whether or not file system 414 issues a free space TRIM command or a standard TRIM command. For example, virtual disk parser 404 may be configured to use two states: mapped and zero to describe virtual disk extents. Alternately, virtual disk parser 404 may be configured to use three states: mapped, zero, unmapped to describe virtual disk extents. Alternately, virtual disk parser 404 may be configured to use four states: mapped, zero, unmapped, and uninitialized. The distinction between unmapped and uninitialized corresponds to the distinction between standard TRIM and free space TRIM. If the parser is configured to not use the uninitialized state, then a free space TRIM is treated as a normal TRIM. As shown by the figure, the parts of file 702 are still being stored in virtual disk file 702 since it is inefficient to clear them from virtual disk file 702.

[0073] Since virtual disk extent 5 was partially covered by the TRIM, virtual disk parser 404 can handle this extent in one of a variety of ways. In one configuration, virtual disk parser 404 may leave extent 5 in the mapped state. In this configuration, virtual disk parser 404 may transition extents when TRIM information is received for an entire extent.

5 Alternatively, virtual disk parser 404 may track TRIM information that partially covers extents in the hope that more TRIM information is received that provides an indication that space describing the extent can be de-allocated.

[0074] Similarly, virtual disk extent was also partially covered by the TRIM. In this example, virtual disk parser 404 may leave it in the mapped state and can also be

10 configured to send TRIM information that describes the part of virtual disk file 702 that is no longer in use to the underlying file system, e.g., virtualization file system 408, storage server file system 504, or computer system file system 514.

[0075] In addition to the deletion of file 802, FIG. 9 shows an example where virtual disk extent 7 was zeroed. Virtual disk parser 404 can scan an IO job issued by file system 414

15 that indicates that the entire range of virtual disk extent 7 is zeroed. In response to this determination, virtual disk parser 404 can be configured to remove the linking from extent allocation table 416 and transition virtual disk extent 7 to the zero state. As shown by the figure, the previous contents of virtual disk extent 7 are still being stored in virtual disk file 702.

20 [0076] Turning to FIG. 10, it illustrates virtual disk 402 described at least in part by a group of virtual disk files 1002, 1004, 1006, which could be similar to the chain of virtual disk files defined by virtual disk file 608, 604, and 600. In this exemplary embodiment, the data that represents virtual disk 402 is broken up across multiple virtual disk files. In this exemplary embodiment, when virtual disk parser 404 attempts to read virtual disk

25 extent 1 and 2, virtual disk parser 404 can access the allocation table for virtual disk file 1002 and determine that these extents are transparent. Next, virtual disk parser 404 can access the allocation table for virtual disk file 1004 and determine that these extents are transparent. Finally, virtual disk parser 404 can access allocation table for grandparent virtual disk file 1006 and determine that these virtual disk extents are defined.

30 [0077] The following are a series of flowcharts depicting operational procedures. For ease of understanding, the flowcharts are organized such that the initial flowcharts present implementations via an overall “big picture” viewpoint and subsequent flowcharts provide further additions and/or details that are illustrated in dashed lines. Furthermore, one of

skill in the art can appreciate that the operational procedure depicted by dashed lines are considered optional.

**[0078]** Referring now to FIG. 11, it illustrates an operational procedure for reclaiming space within a virtual disk file including the operations 1100, 1102, 1104, and 1106.

5 Operation 1100 begins the operational procedure and operation 1102 shows instantiating a virtual disk including a virtual disk extent, the virtual disk extent being dissociated from a virtual disk file. Turning briefly to FIG. 4, FIG. 5A or FIG. 5B, virtual disk 402 can be instantiated by virtual disk parser 404, e.g., executable instructions and associated instance data, that exposes the data stored within one or more virtual disk files as a logical hard  
10 drive, which can be configured to handle read/write operations from file system 414 by emulating the behavior of a hard drive. Virtual disk file 406 (which could be one or more files as illustrated in FIG. 6) can store what is typically found on a physical hard drive, i.e., disk partitions, file systems, etc. Turning to FIG. 7, virtual disk 402 is shown including a plurality of extents, some of which are dissociated from any sections of virtual disk file  
15 702.

**[0079]** In a specific example, suppose the extents are blocks. In this example, allocation table 416, which can be loaded from one or more sections in the virtual disk file 702 into random access memory, can be used to store information that links virtual disk blocks in virtual disk 402 to extent sized (e.g., block sized) sections of virtual disk file 702.

20 Allocation table 416 can also store state information for each virtual disk block in virtual disk 402. Virtual blocks that potentially include non-zero data can be associated with state information that indicates that the block is in the mapped state. That is, a section of virtual disk file 702 has been allocated to describe, i.e., store data for, a block of virtual disk 402. Virtual disk blocks 0-3 and 7 are examples of blocks in this state. As shown by the figure,  
25 virtual disk blocks 4 and 5, 6, 8 and 9 may be valid virtual disk blocks; however, these virtual disk blocks may not have any space allocated within virtual disk file 702. Since file system 414 may write to these blocks, in an exemplary embodiment, these virtual disk blocks can be associated with information that can be used by virtual disk parser 404 to determine how to respond to read and/or write operations to them.

30 **[0080]** Referring briefly back to FIG. 11, operation 1104 shows that a computer system can additionally include circuitry for allocating, based on state information associated with the virtual disk, a section of the virtual disk file to describe the virtual disk extent without overwriting a preexisting bit pattern within the section of the virtual disk file. For example, and returning to FIG. 8, virtual disk parser 404 can receive an IO job to write to



a portion of the virtual disk extent. In response to receipt of the write IO job, virtual disk parser 404 can check allocation table 416 and determine that space within virtual disk file 702 has not been allocated to describe the virtual disk extent and allocate a section of virtual disk file 406 to back the virtual disk extent. Thus, the data written by file system 414 to the virtual disk extent will be stored by virtual disk parser 404 in a section of virtual disk file 702.

**[0081]** In this example, virtual disk parser 404 may not overwrite any data already stored in the section of virtual disk file 702 (by writing all zeros, ones, or any other non-information disclosing bit pattern) prior to using it to describe the virtual disk extent based on the state information in allocation table 416. In an exemplary configuration, the state information could indicate that file system 414 is securing access to this virtual disk extent because the virtual disk extent is covered by file system free space. In a specific example, the state information could indicate that the virtual disk extent is in the “uninitialized” state. Allocating the virtual disk extent without clearing it provides an added benefit of saving processor cycles and IO jobs that would be otherwise used to overwrite the section of virtual disk file 702.

**[0082]** In a specific example of operation 1104, and turning to FIG. 7, suppose that an extent is a block and file system 414 sends an IO job to virtual disk 402 to write a bit pattern indicative of file 802 to virtual disk blocks 3-5. In response to receipt of such an IO job, virtual disk parser 404 can determine that virtual disk block 5 is not backed by any sections of virtual disk file 406 and that it is uninitialized. In response to this determination, virtual disk parser 404 can be configured to allocate a section of virtual disk file 702 to describe virtual disk block 5 and write a portion of the bit pattern indicative of file 802 therein without overwriting data that was previously stored in the portion of the section not covered by the IO job.

**[0083]** Turning again to FIG. 11, operation 1106 shows that the computer system can additionally include circuitry configured to modify the state information associated with the virtual disk extent to indicate that the virtual disk extent is described by the virtual disk file. For example, and turning back to FIG. 8, virtual disk parser 404 can modify, e.g., overwrite in memory, the state information associated with virtual disk extent 5 to reflect that virtual disk file 702 is describing the virtual disk extent. In one configuration, the write and the modification of the state information can occur concurrently. For example, virtual disk parser 404 can store information in allocation table 416 that indicates that virtual disk extent 5 is “mapped.” Consequently, subsequent read operations directed to

sectors of virtual disk extent 5 will be handled by virtual disk parser 404 by returning the bit pattern stored at the byte offset identified in allocation table 416. Virtual disk parser 404 can concurrently write data, e.g., a bit pattern associated with a write operation that triggered this procedure, to the section of virtual disk file 702 allocated to describe the virtual disk extent and issue an IO job to write the bit pattern to the section of virtual disk 702 to virtualization system file system 408, storage server file system 504, or computer system file system 514. At some point in time, such as prior to completion of a subsequently issued flush command, the bit pattern will be persisted in persistent storage unit 460.

**[0084]** Turning now to FIG. 12, it shows additional operations that can be executed in conjunction with those illustrated by FIG. 11. Turning to operation 1208, it indicates that the computer system can include circuitry for responding to an offset query command with information that identifies sectors of the virtual disk that are non-zero, sectors of the virtual disk that are in a non-transparent state, sectors of the virtual disk that are in a mapped state, and/or sectors of the virtual disk that are in an initialized state. For example, virtual disk parser 404 can be configured to receive a command to generate information about virtual disk 402 such as the next byte offset on the virtual disk, given a starting byte offset, that is in a non-transparent state, i.e., a state other than transparent, a mapped state, i.e., sectors of the virtual disk 402 that include data in virtual disk file 406, a defined state, i.e., sectors of the virtual disk 402 that are mapped or zero, and/or an initialized state, i.e., a state other than uninitialized. The command can be depth-limited, in that only a specified number of virtual disk files are examined, with any ranges remaining transparent after the specified number of virtual disk files are examined reported back to the requestor in addition to ranges indicated by the state query, regardless of which state query was requested. In response to receipt of such a command, virtual disk parser 404 can start at the initial byte offset on virtual disk 402 and build a response range or set of ranges until the range associated with the command is detected and return the desired information.

**[0085]** Continuing with the description of FIG. 12, operation 1210 shows sending a request to a file system controlling the virtual disk file to issue at least one command selected from a group of commands including a trim command, an unmap command, a write same of zero command, and an offload write of a zero token command. Referring back to FIG. 4, FIG. 5A, or FIG. 5B, virtual disk parser 404 can be configured to issue a request to file system 414. The request in this example can be for file system 414 to issue

a TRIM command. For example, virtual disk parser 404 can issue one or more requests to file system 414 periodically, soon after instantiation of virtual disk 402, and/or prior to shutting down, hibernating, etc., virtual machine 410. In response to such a request, file system 414 can determine what sectors of virtual disk 402 it is no longer using and send one or more TRIM commands identifying these unused sectors to virtual disk parser 404. Virtual disk parser 404 may thereby receive trim information such as a list of ranges of sectors that are no longer in use by file system 414 and whether file system 414 is preventing reads from the ranges of sectors in order to secure access to those sectors.

Virtual disk parser 404 can receive the information and transition virtual disk extents covered by the ranges into states where space within virtual disk file 702 can be reclaimed.

[0086] Continuing with the description of FIG. 12, operation 1212 shows that the computer system can include circuitry for determining a portion of the virtual disk file that corresponds to a portion of a second virtual disk extent in response to receipt of a request to trim a portion of the second virtual disk extent; and circuitry for sending a trim

command for the determined portion of the virtual disk file to a file system configured to store the virtual disk file on a storage device. For example, and referring to FIG. 8, file system 414 may issue a TRIM command that identifies a portion of a virtual disk extent, e.g., the TRIM command may only identify a range of sectors that corresponds to a part of the sectors that form one or more virtual disk blocks. In a specific example, suppose file system 414 trims the space used to store file 802. As such, the trim command may only identify a portion of the sectors that constitute virtual disk extent 3. In this example, virtual disk parser 404 can determine that the range of sectors covers a subsection of the virtual disk extent and use mapping information in allocation table 416 to determine the portion of virtual disk file 702 that corresponds to the trimmed sectors of the virtual disk extent. Virtual disk parser 404 can issue a request to trim the portion of virtual disk file 702 that corresponds to the trimmed sectors of the virtual disk extent to virtualization system file system 408 or storage server file system 504. Virtualization system file system 408 or storage server file system 504 may be configured to use the trim command and benefit from it by trimming a portion of the sectors backing virtual disk file 406, flushing data from a cache, clearing internal buffers, etc.

[0087] Alternatively, virtual disk parser 404 can store information indicating that a portion of the virtual disk extent was trimmed as well as information that indicates whether it was a free space trim or not. As guest operating system 412 or operating system 508 runs, it may eventually zero or trim the remainder of the virtual disk extent. In response to this

occurring, virtual disk parser 404 can determine to transition the virtual disk extent into a state where it is not described by virtual disk file 702 and select a state based on how the different portions of the virtual disk extent were trimmed or zeroed. Virtual disk parser 404 can be configured to select the most restrictive state to transition a virtual disk extent when different portions of a virtual disk extent can be placed in different non-described states, where the zero state is the most restrictive, uninitialized is the least restrictive state, and unmapped is somewhere in between. For example, if a first portion is zeroed and the remainder is uninitialized, virtual disk parser 404 can transition the entire virtual disk extent to the zeroed state.

**[0088]** Continuing with the description of FIG. 12, operation 1214 illustrates that computer system 400 can additionally include circuitry configured to de-allocate the virtual disk extent from the section of the virtual disk file and modify the state information associated with the virtual disk extent to indicate that the virtual disk extent has no associated space in the virtual disk file in response to receipt of a request to trim a range of sectors that covers the virtual disk extent. For example, and turning to FIG. 9, virtual disk parser 404 can remove the linking in allocation table 416 that ties a virtual disk extent to a section of virtual disk file 702. This operation has the effect of dissociating the virtual disk extent from virtual disk file 702. In addition to removing the link, virtual disk parser 404 can modify the state information associated with the virtual disk extent to indicate that the extent has no associated space within virtual disk file 702, i.e., virtual disk parser 404 can place the virtual disk extent into the unmapped, uninitialized, or zeroed state.

**[0089]** Virtual disk parser 404 can remove the linking and update the state information in response to receipt of a request to trim or zero sectors of the virtual disk extent. For example, a request to trim or zero sectors can be received that identifies a range of byte offsets that could cover one or more virtual disk extents. In response to receipt of such an IO job, virtual disk parser 404 can determine that the request covers the sectors of the virtual disk extent and execute the aforementioned operations for removing the linking and updating the state information.

**[0090]** In a specific example, suppose that the IO job indicates that the trim is a free space trim. For example, a user may have deleted file 802, which is stored as a bit pattern across virtual disk extents 3-5 and file system 414 may indicate that the space is no longer being used by file system 414. In response to receipt of a free space TRIM command, virtual disk parser 404 can access allocation table 416 and determine that file system 414 has trimmed a portion of extent 3, 5 and all of extent 4. In this example, virtual disk parser

404 can remove the link mapping virtual disk extent 4 to virtual disk file 702 and modify the state information associated with virtual disk extent 4 to indicate that the virtual disk extent is uninitialized. This section of virtual disk file 702 can now be reused to back other virtual disk extents. In addition, virtual disk parser 404 can determine that virtual disk extent 3 and 5 are the subject of a partial TRIM command. In this example, virtual disk parser 404 can use allocation table 416 to discover the virtual disk file byte offsets that describe the portion of virtual disk file 702 that describes the trimmed portions of virtual disk extent 3 and 5 and issue a TRIM command describing the virtual disk file byte offsets to virtualization system file system 408, storage system file system 504, or computer system file system 514.

**[0091]** In another specific example, suppose that the IO job issued by file system 414 indicates that file 802 was zeroed. For example, file 802 could be a database file storing sensitive information such as credit card numbers and an administrator determined to zero out the contents of the file by writing all zeros to it by issuing a write command with an all-zero buffer, which will write zeros over the data existing in file 802. In response to receipt of such an IO job, virtual disk parser 404 can be configured to determine that virtual disk extent 4 has been zeroed and that this extent can be reclaimed. In this example, virtual disk parser 404 can remove the link mapping virtual disk extent 4 to virtual disk file 702 and modify the state information associated with virtual disk extent 4 to indicate that the virtual disk extent is zeroed. This section of virtual disk file 702 can now be reused to back other virtual disk extents and virtual disk parser 404 can respond to subsequent read operations to virtual disk extent 4 by replying with all zeros.

**[0092]** In another specific example, a user may write bulk zeros to initialize the state of file 802, rather than to overwrite data stored therein. In this example, a command such as a TRIM, in the instance that virtual disk parser 404 reports that trimmed sections read as zero, UNMAP, when virtual disk parser 404 reports that unmapped regions are zero, WRITE SAME of zero, and/or an offload write of a zero token can be used to transition an extent to the zeroed state.

**[0093]** In a specific example, suppose that the IO job indicates that the trim is a standard trim. For example, a user may have deleted file 802, which is stored as a bit pattern across virtual disk extents 3-5; however, the TRIM command may not indicate whether or not the space is being used by file system 414. In response to receipt of a standard TRIM command, virtual disk parser 404 can access allocation table 416 and determine that file system 414 has trimmed a portion of extent 3, 5 and all of extent 4. In this example,

virtual disk parser 404 can remove the link mapping virtual disk extent 4 to virtual disk file 702 and modify the state information associated with virtual disk extent 4 to indicate that the virtual disk extent is unmapped or zero. This section of virtual disk file 702 can now be reused to describe other virtual disk extents. In addition, virtual disk parser 404  
5 can determine that virtual disk extent 3 and 5 are the subject of a partial TRIM command. In this example, virtual disk parser 404 can use allocation table 416 to discover the virtual disk file byte offsets that make up the portion of virtual disk file 702 that describes the trimmed portions of virtual disk extent 3 and 5 and issue a TRIM command specifying the virtual disk file byte offsets, typically in the form of ranges, to virtualization system file  
10 system 408.

**[0094]** Referring now to FIG. 13, which illustrates additional operations that can be executed in addition to operation 1214 of FIG. 12. Operation 1316 illustrates that a computer system can include circuitry for receiving a request to write data to the virtual disk extent; circuitry for zeroing an unused section of the virtual disk file based on the  
15 state information associated with the virtual disk extent, the state information indicating that the virtual disk extent was zeroed; and circuitry for allocating the unused section of the virtual disk file to describe the virtual disk extent. Referring to FIG. 9 for context, virtual disk parser 404 can receive a request to write data to the virtual disk extent, e.g., virtual disk extent 4 of FIG. 9, which in this example is associated with state information  
20 that indicates that the virtual disk extent is zeroed. For example, when virtual disk extent 4 was de-allocated virtual disk parser 404 could have determined that the virtual disk extent was zeroed, i.e., an application wrote all zeros to file 602 by using an offload write of a well-known zero token.

**[0095]** In response to determining that the virtual disk extent is in the zeroed state, virtual  
25 disk parser 404 can identify an unused section of virtual disk file 702, i.e., a section that is not actively being used to describe a virtual disk extent and not actively being used to store any allocated metadata, and use the section to back the virtual disk extent. The virtual disk parser further insures that any reads from not-yet-written sectors of the newly allocated extent read as all zeros. The virtual disk parser 404 can write payloads of IO  
30 write jobs to the section; update state information to indicate that the virtual disk extent is mapped; and update information in allocation table 416 to describe the virtual disk file byte offset that identifies the beginning of the section used to store virtual disk extent 4. The virtual disk parser 404 also can create a log entry, which insures that in the event of system failure and re-start prior to writes being flushed, not-yet-written sectors of the

newly allocated extent still read as all zeros, and written sectors of the newly allocated extent read as either all zeros or the written data. Upon the first subsequent flush command, virtual disk parser 404 insures that a system failure subsequent to completion of the flush will result in reads from previously written sectors of the newly allocated extent reading the data that was written, and reads from not-yet-written sectors of the newly allocated extent reading zeros.

[0096] Continuing with the description of FIG. 13, operation 1318 shows that a computer system can include circuitry for receiving a request to write to the virtual disk extent; and circuitry for allocating an unused section of the virtual disk file to describe the virtual disk extent without modifying contents of the unused section of the virtual disk file based on the state information associated with the virtual disk extent, the state information indicating that the file system is securing access to the virtual disk extent. Referring again to FIG. 9 for context, virtual disk parser 404 can receive an IO job to write data to the virtual disk extent, e.g., virtual disk extent 4 of FIG. 9, which in this example is associated with state information that indicates that security for the virtual disk extent is being provided by file system 414. In response to detecting this state information, virtual disk parser 404 can identify an unused section of virtual disk file 702; write the payload of the IO job to the section; update state information to indicate that the virtual disk extent is mapped; and update information in allocation table 416 to describe the virtual disk file byte offset that identifies the beginning of the section used to store virtual disk extent 4.

[0097] Suppose that in this example the extent is a block and the payload for the IO job only covers a portion of the sectors in the virtual disk block. Specifically, the virtual disk block may be 512 kilobytes and the write may cover the first 500 sectors of the virtual disk block. In this example, virtual disk parser 404 can write data in the first 500 sectors of the allocated section of virtual disk file 702 without erasing the data stored in the remaining 524 sectors. Thus, if this section was examined one would find that the first 500 sectors include the payload and the remaining 524 sectors include whatever bit pattern was previously written to the section of virtual disk file 702. In this example, virtual disk parser 404 can use this section without clearing it because file system 414 is configured to deny read operations to sectors that are in file system free space. Since an application will be prevented from reading the remaining 524 sectors of virtual disk block, it can contain any data, which had previously been stored in the virtual disk.

[0098] Turning now to operation 1320 of FIG. 13, it shows that a computer system can be configured to include circuitry for receiving a request to write to the virtual disk extent;

circuitry for logically overwriting an unused section of the virtual disk file with a non-information-disclosing bit pattern based on the state information associated with the virtual disk extent, the state information indicating that the file system is not securing access to the virtual disk extent; and circuitry for allocating the overwritten section of the virtual disk file to describe the virtual disk extent. Referring again to FIG. 9 for context, virtual disk parser 404 can receive a request to write data to the virtual disk extent, which in this example is associated with state information that indicates that file system 414 is not securing access to the virtual disk extent. For example, virtual disk parser 404 may have de-allocated the virtual disk extent in response to receipt of a standard TRIM command and could have stored state information indicating that virtual disk extent is unmapped, i.e., not backed by space in virtual disk file 702, in allocation table 416.

**[0099]** In response to determining that the virtual disk extent is unmapped, virtual disk parser 404 can identify an unused section of virtual disk file 702 to use to describe the virtual extent and logically write a non-information disclosing bit pattern to the section to ensure that reads to the virtual disk extent do not inadvertently reveal any information. In a preferred implementation, the non-information disclosing bit pattern could be all zeros or previously-stored data. After the section is zeroed or some other non-information disclosing bit pattern is logically written to the section such as previously-stored data, virtual disk parser 404 can logically write the payload of an IO job to the section; update state information to indicate that the virtual disk extent is mapped; and update information in allocation table 416 to describe the virtual disk file byte offset that identifies the beginning of the section used to store the virtual disk extent.

**[0100]** Continuing with the description of FIG. 13, operation 1322 shows that the computer system can include circuitry configured to send, based on state information indicating that the virtual disk extent was zeroed, a token representing zeros to a requestor in response to receipt of an offload read request associated with the virtual disk extent. For example, and referring to FIG. 4, offload provider engine 422, e.g., circuitry configured to service offload read and offload write commands, can send a token representing zeros to a requestor, e.g., application 424, in a response to an offload read request issued by the requestor. An offload read request can be used to efficiently copy data from one location to another by generating and sending tokens to requestors, the tokens representing the requested data instead of copying the data into the requestors' memory and then sending the data to the destination. Offload read and offload write commands can be used to achieve copy offload when the destination location recognizes



the token generated by the source location and can logically write the data represented by the token to the destination. In the case of a well-known zero token generated by the source, the destination need not access the underlying storage, e.g., storage device 106, which could be a SAN target in this specific implementation. In this example, the offload  
5 read request can be to perform an offload read operation on one or more files that have data stored in one or more virtual disk extents, one of which is associated with state information indicating that the virtual disk extent is zeroed. In this example, the offload read request may be serviced by generating a well-known zero token value and returning that well-known zero token to the requestor.

10 **[0101]** The offload read request can be routed to offload provider engine 422. Offload provider engine 422 can receive the request and send a message to virtual disk parser 404 for the data stored in the virtual disk extents. Virtual disk parser 404 can receive the request, read the state information for the virtual disk extent, and determine, in this specific example, that the state information indicates that this virtual disk extent is zeroed.

15 Virtual disk parser 404 can send a message back to offload provider engine 422 that indicates that the virtual disk extent is all zeros and offload provider engine 422 can generate a well-known token value that indicates that the requested data is all zeros, e.g., the range of sectors that describes a virtual disk block is all zeros, and send the well-known zero token to the requestor.

20 **[0102]** In a specific example, the offload request can be forward to a SAN instead of being processed by computer system 400, storage service 500, or computers system 512. In this example, the SAN may generate the token and return it back to virtual disk parser 404, which can then send the zero token to the requestor. In yet another example, when offload provider engine 422 receives the message that indicates that the virtual disk extent is all

25 zeros, offload provider engine 422 can generate the well-known-zero token, which in effect achieves logically copying the requested zero data into a separate area that is associated with the token by identifying the data as equivalent to any other zero data and sharing the area associated with the well-known zero token. In the instance that offload provider engine 422 subsequently receives an offload write specifying the token

30 previously sent to the requestor, offload provider engine 422 can logically copy the data from the area associated with the token to offsets specified by the requestor.

**[0103]** Turning now to FIG. 14, it illustrates an operational procedure for reclaiming virtual disk file space including the operations 1400, 1402, 1404, and 1406. As shown by the figure, operation 1400 begins the operational procedure and operation 1402 shows that

a computer system can include circuitry for receiving a signal indicating that a portion of a virtual disk extent is no longer in use, the virtual disk extent being part of a virtual disk, the virtual disk being stored in a virtual disk file. For example, and turning to FIG. 4, virtual disk parser 404 can be configured to instantiate virtual disk 402. File system 414  
5 can send a signal indicating that it is no longer using a portion of virtual disk 402, e.g., a range of sectors of a virtual disk extent, to virtual disk parser 404. In a specific example, the signal could be a TRIM command. In a specific example, the signal received by virtual disk parser 404 could identify byte offset values that define a range of sectors that it is no longer using, which could be the first part of a virtual disk extent.

10 **[0104]** Continuing with the description of FIG. 14, operation 1404 shows that the computer system can also include circuitry configured to identify a portion of the virtual disk file that describes the portion of the virtual disk extent. Referring back to FIG. 7, virtual disk parser 404 can receive the signal and the virtual disk byte offset values that identify, for example, the first portion of virtual disk extent 0. In response to receipt of the  
15 signal, virtual disk parser 404 can check allocation table 416 to determine the portion of virtual disk file 702 that corresponds to the virtual disk byte offset values associated with the signal.

**[0105]** Turning now to operation 1406 of FIG. 14, it shows that the computer system can include circuitry for sending a request to trim the identified portion of the virtual disk file  
20 to a file system configured to store the virtual disk file on a storage device. For example, and again referring to FIG. 7, virtual disk parser 404 can determine that the signal identified less than the entire virtual disk extent. For example, the signal may indicate a range of sectors that does not include all of the sectors of a virtual disk extent. In response to this determination, virtual disk parser 404 can issue a request to trim the portion of  
25 virtual disk file 702 that corresponds to the trimmed portion of the virtual disk extent to a file system hosting virtual disk file 702, e.g., virtualization system file system 408.

Virtualization system file system 408 may be configured to use the trim command and benefit from it by trimming virtual disk file 406, flushing data from a cache, clearing internal buffers, sending the trim to the disk on which the file system data is stored, etc.

30 **[0106]** In a specific example, virtual disk parser 404 can be configured to issue the TRIM command to the underlying file system in response to determining that the request to trim a portion of the virtual disk file does not cover the entire extent. For example, suppose that the signal identifies that the first 600 sectors of a virtual disk extent are no longer in use and virtual disk parser 404 may determine that the 600 sectors of virtual disk extent

are less than the 1024 sectors that constitute the virtual disk extent. In response to this determination, virtual disk parser 404 can access allocation table 416 and determine the virtual disk file byte offsets that describe the first 600 sectors of the section of virtual disk file 702 that describes the virtual disk extent and send a request to trim this part of virtual disk file 702 to a file system that hosts virtual disk file 702.

[0107] Turning now to FIG. 15, it illustrates additional operations that can be executed in conjunction with those depicted by FIG. 14. Turning now to operation 1508, it shows that the computer system can additionally include circuitry for sending, based on state information indicating that the virtual disk extent was zeroed, a token representing zeros to a requestor in response to receipt of an offload read request associated with the virtual disk extent. For example, and referring to FIG. 4, offload provider engine 422, e.g., circuitry configured to service offload read and offload write commands, can send a token representing zeros to a requestor, e.g., application 424, in a response to an offload read request issued by the requestor. An offload read request can be used to efficiently copy data from one location to another by generating and sending tokens to requestors, the tokens representing the requested data instead of copying the data into the requestors' memory and then sending the data to the destination. Offload read and offload write commands can be used to achieve copy offload when the destination location recognizes the token generated by the source location and can logically write the data represented by the token to the destination. In the case of a well-known zero token generated by the source, the destination need not access the underlying storage, e.g., storage device 106, which could be a SAN target in this specific implementation. In this example, the offload read request can be to perform an offload read operation on one or more files that have data stored in one or more virtual disk extents, one of which is associated with state information indicating that the virtual disk extent is zeroed. In this example, the offload read request may be serviced by generating a well-known zero token value and returning that well-known zero token to the requestor.

[0108] Continuing with the description of FIG. 15, operation 1510 shows that the computer system can include circuitry for selecting a sub-group from the group of virtual disk files; and circuitry for generating information that identifies sectors of the sub-group that include data and sectors of the sub-group that are transparent. In an exemplary embodiment, virtual disk 402 can be instantiated from a plurality of virtual disk files. Or put another way, virtual disk 402 can be formed from M virtual disk files (where M is an integer greater than 1). In this exemplary embodiment, virtual disk parser 404 can be

configured to receive a request from, for example, an administrator, to determine the next byte offset on virtual disk 402, starting at a given byte offset, that is associated with a sector defined within a subgroup of the virtual disk files. For example, and referring to FIG. 10, virtual disk parser 404 may receive a request for the next defined byte offset starting at the virtual disk offset corresponding to the first sector of virtual disk extent 2 and information indicating that the subgroup includes virtual disk file 1002 and virtual disk file 1004. In this example, virtual disk parser 404 can start scanning through subgroup and determine that the next defined byte offset is the sector that corresponds to the beginning of virtual disk extent 3. Since in this example, the data in virtual disk extent 2 is backed by a section of virtual disk file 1006 it is outside of the search and is not returned as being defined.

**[0109]** Continuing with the description of FIG. 15, operation 1512 shows that the computer system can include circuitry configured to dissociate the virtual disk extent from the virtual disk file and modify state information associated with the virtual disk extent to indicate that the virtual disk extent has been zeroed in response to determining that the virtual disk extent was zeroed. For example, and turning to FIG. 7, in an embodiment virtual disk parser 404 can determine that the virtual disk extent has been zeroed. For example, virtual disk parser 404 can receive a request to write data represented by a well-known zero token to the virtual disk extent, e.g., virtual disk extent 7. Virtual disk parser 404 can determine from a data structure associated with the request that the request is for the entire virtual disk extent, i.e., the byte offset values can start at the first sector of extent 7 and end at the last sector of extent 7. In response to such a determination, and instead of writing the zeros to the corresponding section of virtual disk file 702, virtual disk parser 404 can be configured to remove the link that maps virtual disk extent 7 to a section of virtual disk file 702 used to describe virtual disk extent 7 and associate the virtual disk extent with information that indicates that the virtual disk extent is all zeros. For example, virtual disk parser 404 can write eight bytes of information in allocation table 416 that indicates that the virtual disk extent includes all zeros. The end result of this operation is that the section of virtual disk file 702 can be reused to store data for other virtual disk extents and the virtual disk extent will read as if it includes all zeros, even though no portion of the virtual disk file is describing the extent on a bit-for-bit basis.

**[0110]** Continuing with the description of FIG. 15, operation 1514 shows that the computer system can additionally include circuitry configured to dissociate the virtual disk extent from the virtual disk file and modify state information associated with the virtual

disk extent to indicate that the virtual disk extent is free space in response to a determination that the virtual disk extent is considered free space by a file system. For example, and again turning to FIG. 7, virtual disk parser 404 can determine that file system 414 has associated the virtual disk extent with information that indicates that it is free space, i.e., space that is not used by file system 414. For example, virtual disk parser 404 can receive a signal from file system 414 indicating a range of sectors that covers the virtual disk extent, e.g., virtual disk extent 3, and information that indicates that the sectors are considered to be free space. In response to receipt of such a signal, virtual disk parser 404 can be configured to remove information that links the virtual disk extent to a section of virtual disk file 702. The result of this operation is that the section of virtual disk file 702 can be reused to store data for other virtual disk extents. Virtual disk parser 404 can additionally associate the virtual disk extent with information that indicates that the virtual disk extent includes arbitrary data, i.e., data previously stored in any part of the virtual disk, all zeros, or all ones. Consequently, read operations directed to this virtual disk extent can be handled by returning arbitrary data which was previously stored in the virtual disk. In addition, the arbitrary data can optionally change each time a read operation is received, if the virtual disk parser 404 is configured to allow the arbitrary data to change each time a read operation is received.

**[0111]** Continuing with the description of FIG. 15, operation 1516 shows that the computer system can additionally include circuitry configured to dissociate the extent from the virtual disk file and modify state information associated with the virtual disk extent to indicate that the virtual disk extent includes a non-information-disclosing bit pattern in response to a determination that the virtual disk extent was trimmed. For example, and again turning to FIG. 7, virtual disk parser 404 can determine that file system 414 has trimmed a range of sectors that compose a virtual disk extent. In response to such a determination, virtual disk parser 404 can remove information in allocation table 416 that links the virtual disk extent to a section of virtual disk file 702. The result of this operation is that the section of virtual disk file 702 can be reused to store data for other virtual disk extents. Virtual disk parser 404 can additionally associate the virtual disk extent with information that indicates that the virtual disk extent includes a non-information-disclosing bit pattern, e.g., all zeros, ones, or a randomly generated bit pattern. Consequently, read operations directed to this virtual disk extent can be handled by returning the non-information-disclosing bit pattern. In a specific preferred implementation, the non-information-disclosing bit pattern can be all zeros. However, this

is different than the zero state described above in that the zero state can be used to represent meaningful zeros, i.e., the instance where the virtual disk extent was intentionally zeroed.

[0112] Referring to operation 1518, it shows that the computer system can additionally include circuitry configured to send a request to issue a trim command to a file system controlling the virtual disk. Referring back to FIG. 7, virtual disk parser 404 can be configured to issue a request that file system 414 issue one or more TRIM commands. In an exemplary configuration, virtual disk parser 404 can be configured to periodically send such a request or to send such a request based on predetermined criteria, e.g., when VM 410 starts or shortly before the VM is to be shut down. In response to such a request, file system 414 can issue one or more TRIM commands that identify the unused sectors of virtual disk 402 to virtual disk parser 404. Virtual disk parser 404 may then receive trim information from the TRIM commands such as the range of sectors that are no longer in use by file system 414 and optionally information that indicates whether the trimmed sectors are considered free space. Virtual disk parser 404 can receive the information and use it to update state information stored in allocation table 416 and to possibly reclaim unused sections of virtual disk file 702.

[0113] Turning now to FIG. 16, it illustrates an operational procedure for storing data for a virtual machine. The operational procedure begins with operation 1600 and transitions to operation 1602, which describes an instance where a computer system can include circuitry for executing a guest operating system including a file system within a virtual machine. For example, and referring to FIG. 4, virtualization system 420, which could be hypervisor 302 of FIG. 3 or the combination of functions executed by host environment 204 and microkernel hypervisor 202 of FIG. 2, can instantiate virtual machine 410 and run a guest operating system (such as guest operating system 412) within it. In this example, guest operating system 412 can include file system 414, which can be executable instructions that organize and control data for guest operating system 412.

[0114] Continuing with the description of FIG. 16, operation 1604 shows that the computer system can include circuitry for exposing a virtual storage device to the guest operating system, the virtual storage device including a virtual disk extent, the virtual disk extent being dissociated from a virtual disk file. Turning back to FIG. 4, virtualization system 420 can expose virtual disk 402 to guest operating system 412. For example, virtual disk parser 404 can be in communication with a storage virtualization service provider that is operable to communicate with a storage virtualization service client

running within guest operating system 410. In a specific example, the storage virtualization service client could be a driver installed within guest operating system 412 that signals to the guest that it can communicate with a storage device. In this example, IO jobs sent by file system 414 are sent first to the storage virtualization service client and then to the storage virtualization service provider via a communication channel, e.g., a region of memory and cross-partition notification facility. Virtual disk 402 can be composed from one or more virtual disk files 406 that can be opened by virtual disk parser 404 and used to store data for virtual disk 402. In a specific example virtual disk 402 can be described at least in part by virtual disk file 702 of FIG. 7. In another specific example, and turning to FIG. 10, virtual disk 402 can be described by a group of virtual disk files (1002 - 1006). In either case, and returning to FIG. 4, virtual disk 402 can include a plurality of virtual disk extents and one of the virtual disk extents can be dissociated, i.e., not described on a bit-for-bit basis by any space within its associated virtual disk file.

**[0115]** Continuing with the description of FIG. 16, operation 1606 shows that the computer system can include circuitry for receiving a request to write data to the virtual disk extent. Turning back to FIG. 7, virtual disk parser 404 can receive a request to write data to the virtual disk extent that has no associated space within virtual disk file 702. For example, an IO job can be received that specifies an offset value indicative of the address of a virtual disk sector, which is within the virtual disk extent.

**[0116]** Turning back to FIG. 16, operation 1608 shows that the computer system can optionally include circuitry for determining that state information associated with the virtual disk extent indicates that the virtual disk extent is free space. In response to receipt of the IO job, virtual disk parser 404 can access allocation table 416 and read state information associated with the virtual disk extent. In this example, the virtual disk extent may be associated with information that indicates that the virtual disk extent is free space, i.e., that file system 414 is not using the virtual disk extent and that read operations to the virtual disk extent can be answered with arbitrary data.

**[0117]** Referring to FIG. 16, operation 1610 shows that the computer system can optionally include circuitry for allocating a section of the virtual disk file to describe the virtual disk extent without overwriting a preexisting bit pattern within the section of the virtual disk file. For example, and returning to FIG. 7, in response to receipt of a write IO job, virtual disk parser 404 can locate a section in virtual disk file 702 that is not being used and allocate it to store data for the virtual extent. For example, virtual disk parser

404 can write information in allocation table 416 that links the virtual disk extent to byte offset values of the allocated section of virtual disk file 702.

[0118] In this example, virtual disk parser 404 may not overwrite any bit pattern existing within the section, e.g., data from some deleted file and/or arbitrary data, stored in the section of virtual disk file 702 (by writing all zeros, ones, or any other non-information-disclosing bit pattern) prior to using the section to describe the virtual disk extent because state information indicates that file system 414 has identified virtual disk extent 5 as free space. This provides an added benefit of saving processor cycles and IO jobs that would be otherwise used to overwrite the section of the virtual disk extent.

[0119] Referring to operation 1612 of FIG. 16, it shows that the computer system can optionally include circuitry for modifying the state information associated with the virtual disk extent to indicate that the virtual disk extent is mapped to the allocated section of the virtual disk file. For example, and turning back to FIG. 7, virtual disk parser 404 can modify, e.g., overwrite in memory, the state information associated with the virtual disk extent to indicate that it is mapped. Consequently, subsequent read operations directed to sectors of the virtual disk extent will be handled by virtual disk parser 404 by returning the bit pattern stored in corresponding portions of the allocated section.

[0120] Turning now to operation 1614 of FIG. 16, it shows storing the data to the allocated section of the virtual disk file. Turning back to FIG. 6, virtual disk parser 404 can write the data, e.g., a bit pattern, into virtual disk file 702. An IO job indicative of the write to virtual disk file 702 can be issued to virtualization system file system 408 and eventually the change can be persisted by persistent storage unit 460.

[0121] Turning now to FIG. 17, it shows additional operations that can be executed in conjunction with those illustrated by FIG. 16. Turning attention to operation 1716, it shows that the computer system can optionally include circuitry for dissociating the virtual disk extent from the virtual disk file and modifying the state information associated with the virtual disk extent to indicate that the virtual disk extent has been zeroed in response to determining that the virtual disk extent was zeroed. For example, and turning to FIG. 6, in an embodiment virtual disk parser 404 can determine that the virtual disk extent has been zeroed. For example, virtual disk parser 404 can receive an offload write request to write data represented by a well-known zero token to the virtual disk extent, e.g., virtual disk extent 7. Virtual disk parser 404 can determine from a data structure associated with the request that the request is for the entire virtual disk extent, i.e., the byte offset values can start at the first sector of virtual disk extent 7 and end at the last sector of virtual disk



extent 7. In response to such a determination, and instead of writing the zeros to the corresponding section of virtual disk file 702, virtual disk parser 404 can be configured to remove the link from the virtual disk extent to a section of virtual disk file 702 stored in allocation table 416 and associate the virtual disk extent with information that indicates that the virtual disk extent is all zeros.

[0122] Continuing with the description of FIG. 17, operation 1718 shows that the computer system can optionally include circuitry for dissociating the virtual disk extent from the virtual disk file and modifying the state information associated with the virtual disk extent to indicate that the virtual disk extent includes arbitrary data in response to receipt of a signal from a file system identifying the virtual disk extent as free space. For example, and again turning to FIG. 7, virtual disk parser 404 can determine that file system 414 has associated the virtual disk extent with information that indicates that it is free space, i.e., space that is not used by file system 414. For example, virtual disk parser 404 can receive a signal from file system 414 indicating a range of sectors that covers the virtual disk extent, e.g., virtual disk extent 3, and information that indicates that the sectors are free space. In response to such a determination, virtual disk parser 404 can be configured to remove information in allocation table 416 that links the virtual disk extent to a section of virtual disk file 702 and associate the virtual disk extent with information that indicates that arbitrary data, i.e., data previously stored in any part of the virtual disk, all zeros, or all ones, can be returned in response to receipt of a read IO job.

[0123] Operation 1720 of FIG. 17 shows that computer system 400 can optionally include circuitry for dissociating the virtual disk extent from the virtual disk file and modifying the state information associated with the virtual disk extent to indicate that the virtual disk extent includes a non-information disclosing bit pattern in response to receipt of a request to trim all the sectors of the virtual disk extent. For example, and again turning to FIG. 7, virtual disk parser 404 can determine that the sectors that compose a virtual disk extent have been trimmed. For example, virtual disk parser 404 can receive a trim command from file system 414 indicating a range of sectors that covers the virtual disk extent. In response to receipt of such a signal, virtual disk parser 404 can be configured to remove information in allocation table 416 that links the virtual disk extent to a section of virtual disk file 702 and associate the virtual disk extent with information that indicates that the virtual disk extent includes a non-information-disclosing bit pattern.

[0124] The foregoing detailed description has set forth various embodiments of the systems and/or processes via examples and/or operational diagrams. Insofar as such block

diagrams, and/or examples contain one or more functions and/or operations, it will be understood by those skilled in the art that each function and/or operation within such block diagrams, or examples can be implemented, individually and/or collectively, by a wide range of hardware, software, firmware, or virtually any combination thereof.

- 5   **[0125]** While particular aspects of the present subject matter described herein have been shown and described, it will be apparent to those skilled in the art that, based upon the teachings herein, changes and modifications may be made without departing from the subject matter described herein and its broader aspects and, therefore, the appended claims are to encompass within their scope all such changes and modifications as are within the
- 10   true spirit and scope of the subject matter described herein.

What is Claimed:

1. A computer-readable storage medium including instructions that upon execution by a processor cause the processor to:

5        instantiate a virtual disk including a virtual disk extent, the virtual disk extent  
being dissociated from a virtual disk file;  
allocate, based on state information associated with the virtual disk, a section of the  
virtual disk file to describe the virtual disk extent without overwriting a preexisting bit  
pattern within the section of the virtual disk file; and  
10        modify the state information associated with the virtual disk extent to indicate that  
the virtual disk extent is described by the virtual disk file.

2. The computer-readable storage medium of claim 1, further comprising instructions that upon execution cause the processor to:

15        Send a request to a file system controlling the virtual disk file to issue at least one  
command selected from a group of commands including a trim command, an unmap  
command, a write same of zero command, and an offload write of a zero token command.

3. A computer system, comprising:  
a processor;

20        a memory coupled to the processor when the processor and the memory are  
powered, the memory including instructions that upon execution by the processor cause  
the computer system to:

receive a signal indicating that a portion of a virtual disk extent is no longer in use,  
the virtual disk extent being part of a virtual disk, the virtual disk being stored in a virtual  
disk file;

25        identify a portion of the virtual disk file that describes the portion of the virtual  
disk extent; and

send a request to trim the identified portion of the virtual disk file to a file system  
configured to store the virtual disk file on a storage device.

4. The computer system of claim 3, the memory further comprising instructions that upon execution cause the computer system to:

30        send, based on state information indicating that the virtual disk extent was zeroed,  
a token representing zeros to a requestor in response to receipt of an offload read request  
associated with the virtual disk extent.

5. The computer system of claim 3, the virtual disk file being a member of a group of  
virtual disk files that together form a virtual disk that includes the virtual disk extent and

the memory further comprising instructions that upon execution cause the computer system to:

select a sub-group from the group of virtual disk files; and

generate information that identifies sectors of the sub-group that include data and

5 sectors of the sub-group that are transparent.

6. The computer system of claim 3, the memory further comprising instructions that upon execution cause the computer system to:

dissociate the virtual disk extent from the virtual disk file and modify state

information associated with the virtual disk extent to indicate that the virtual disk extent

10 has been zeroed in response to determining that the virtual disk extent was zeroed.

7. The computer system of claim 3, the memory further comprising instructions that upon execution cause the computer system to:

send a request to issue a trim command to a file system controlling the virtual disk.

8. A computer implemented method for storing data for a virtual machine, comprising:

15 executing a guest operating system including a file system within a virtual machine;

exposing a virtual storage device to the guest operating system, the virtual storage device including a virtual disk extent, the virtual disk extent being dissociated from a virtual disk file;

20 receiving a request to write data to the virtual disk extent;

determining that state information associated with the virtual disk extent indicates that the virtual disk extent is free space;

allocating a section of the virtual disk file to describe the virtual disk extent without overwriting a preexisting bit pattern within the section of the virtual disk file;

25 modifying the state information associated with the virtual disk extent to indicate that the virtual disk extent is mapped to the allocated section of the virtual disk file; and

storing the data to the allocated section of the virtual disk file.

9. The method of claim 8, further comprising

dissociating the virtual disk extent from the virtual disk file and modifying the state

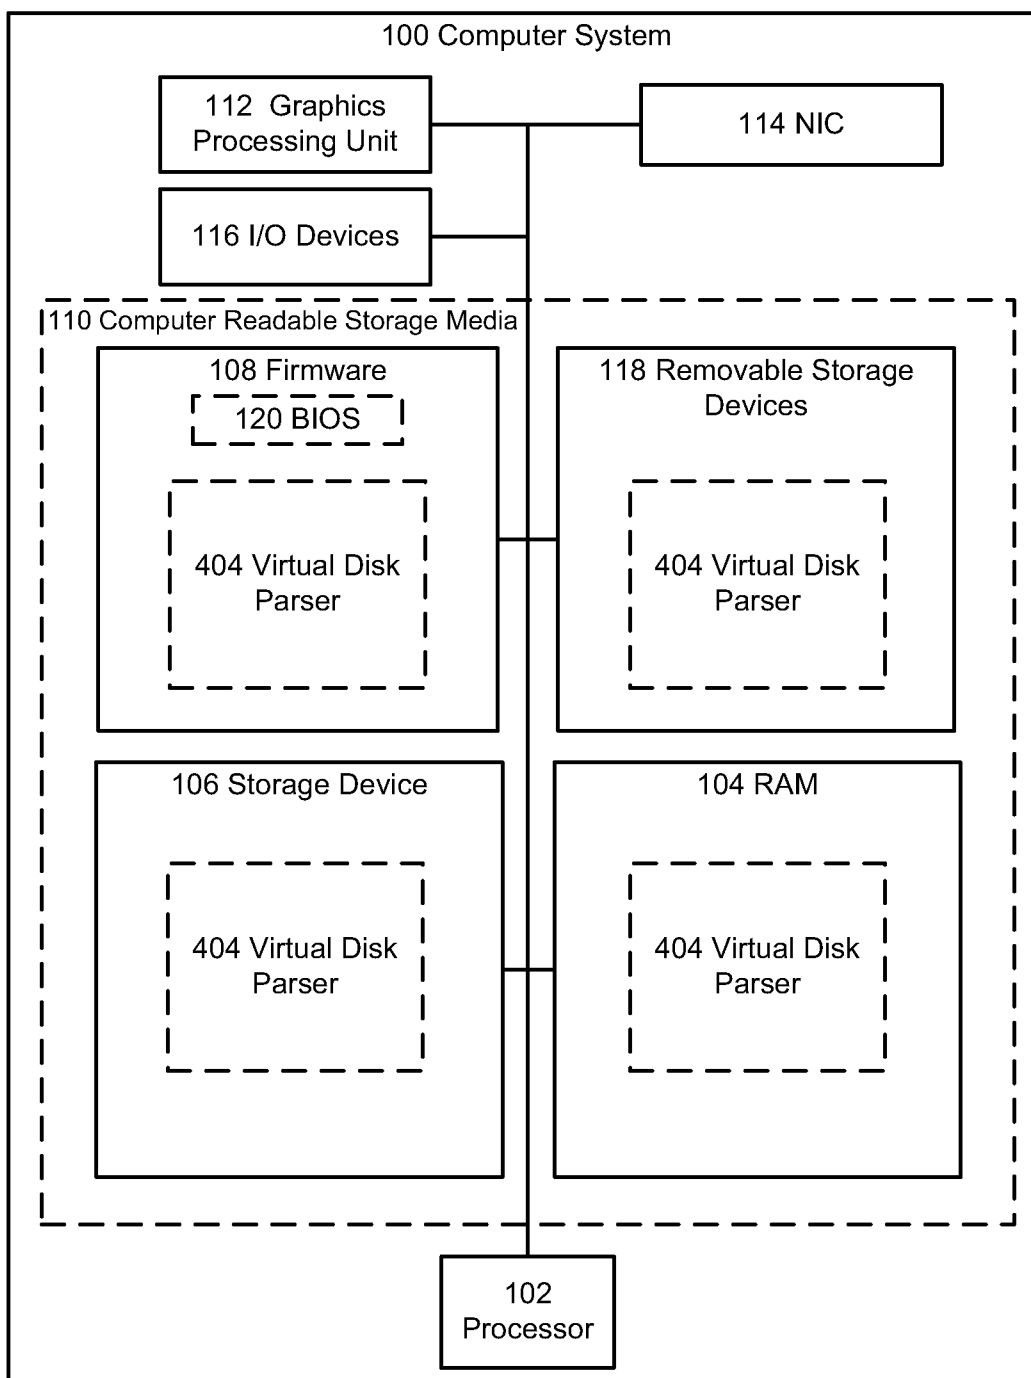
30 information associated with the virtual disk extent to indicate that the virtual disk extent includes arbitrary data in response to receipt of a signal from a file system identifying the virtual disk extent as free space.

10. The method of claim 8, further comprising

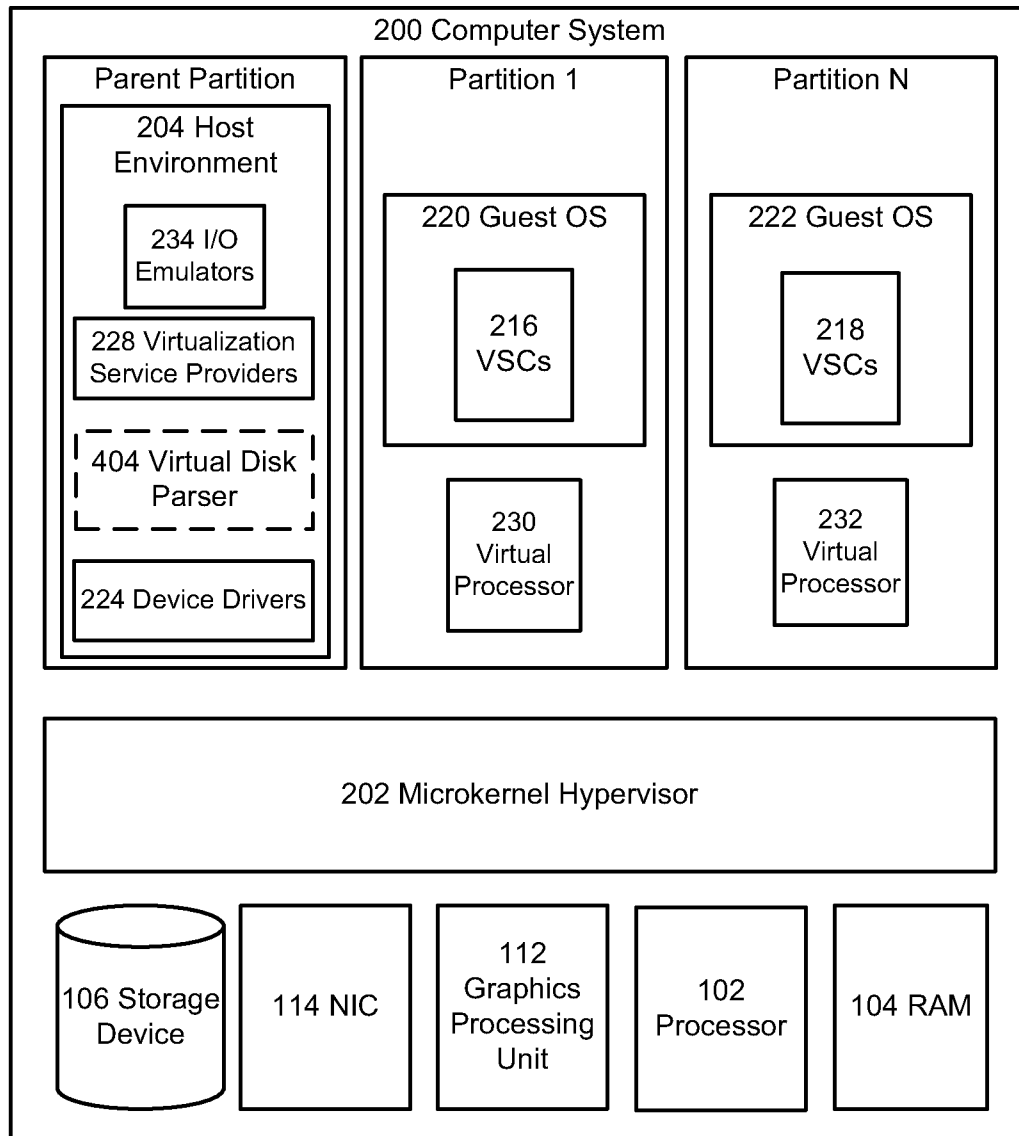
dissociating the virtual disk extent from the virtual disk file and modifying the state information associated with the virtual disk extent to indicate that the virtual disk extent includes a non-information disclosing bit pattern in response to receipt of a request to trim

5 all the sectors of the virtual disk extent.

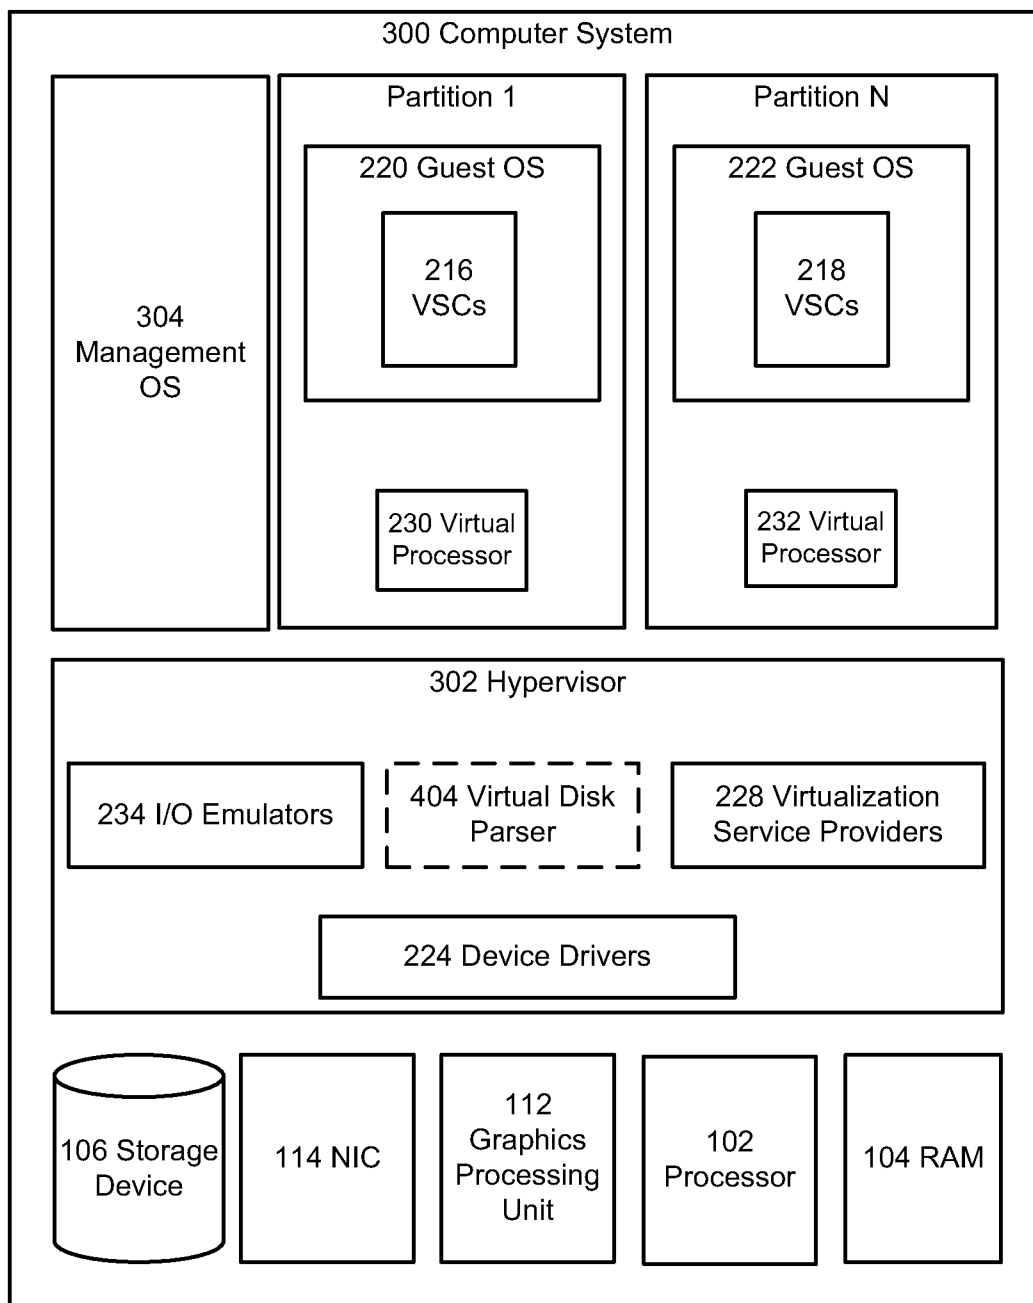
1/18

**FIG. 1**

2/18

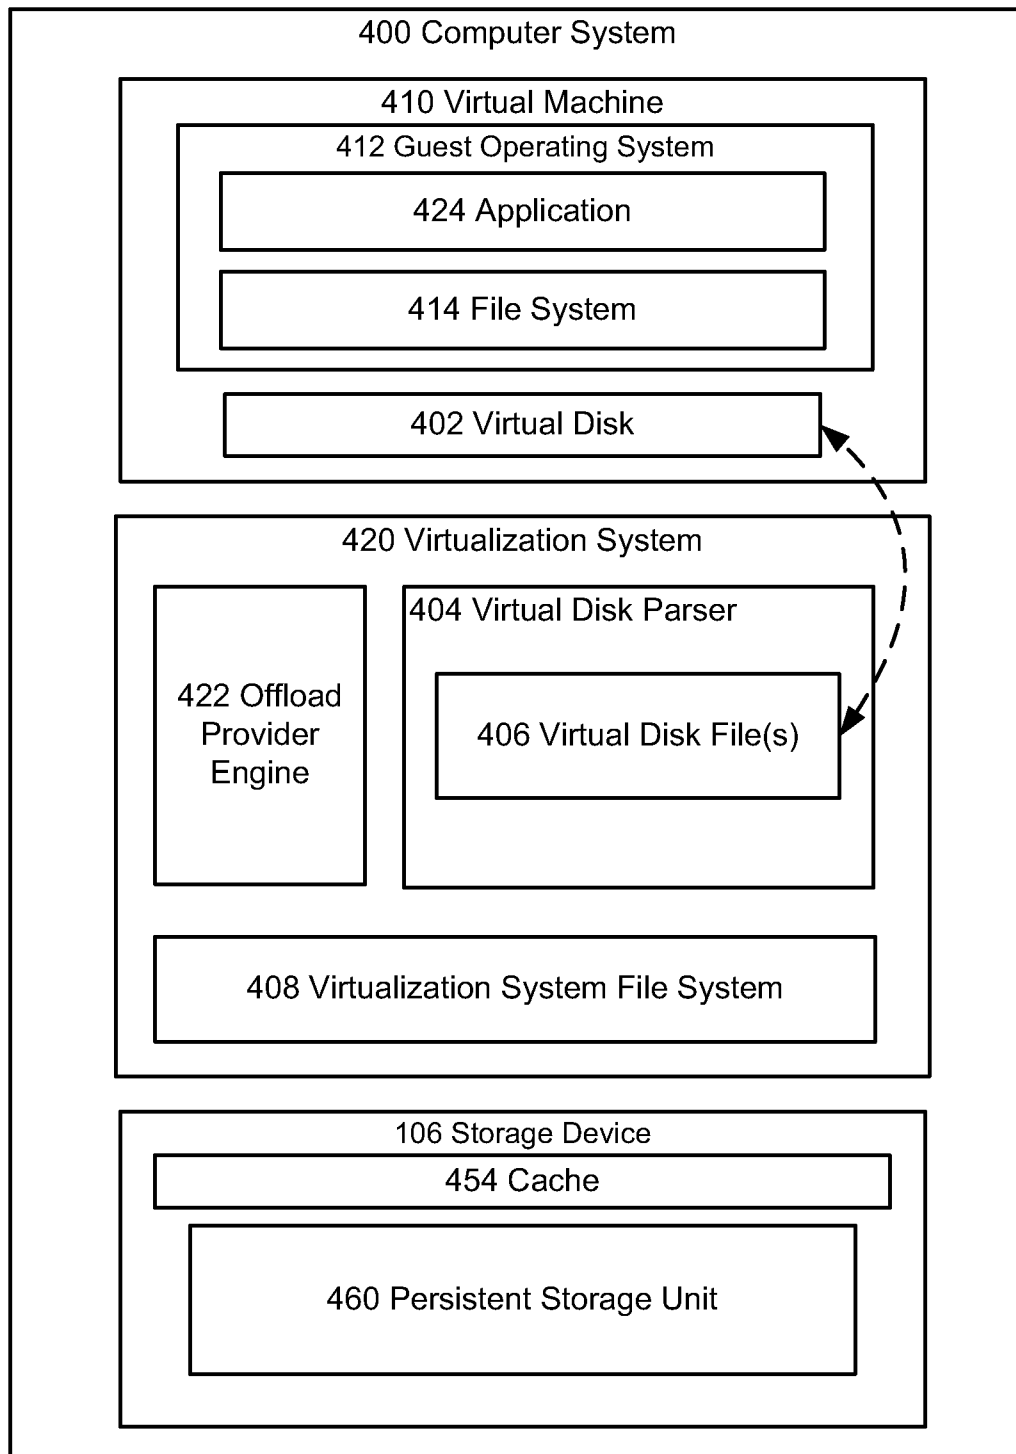
**FIG. 2**

3/18

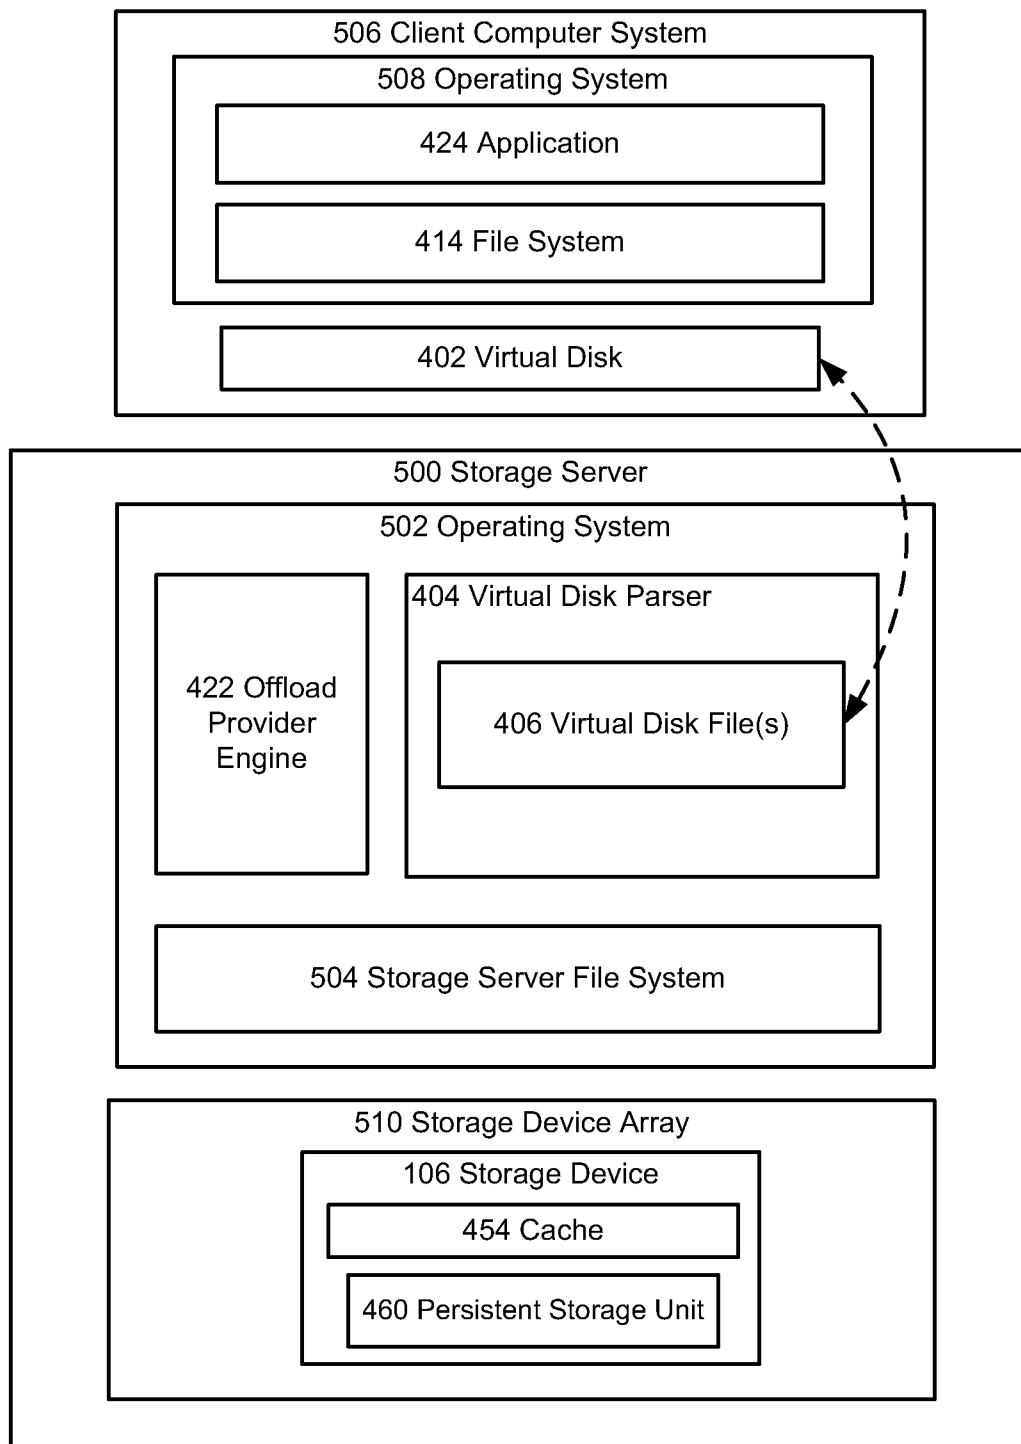
**FIG. 3**



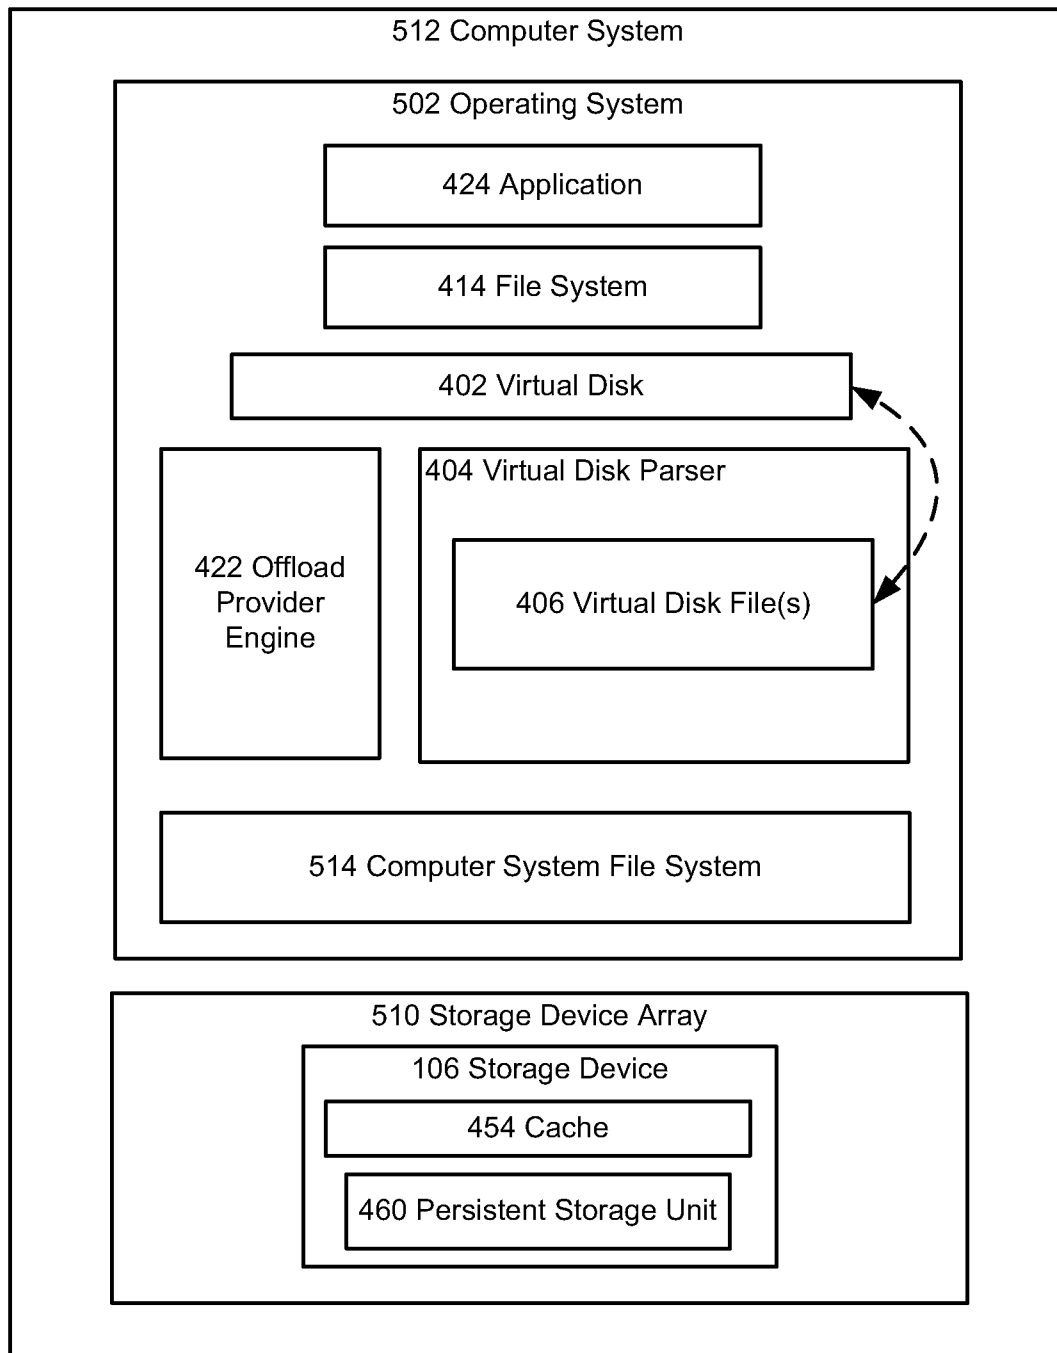
4/18

**FIG. 4**

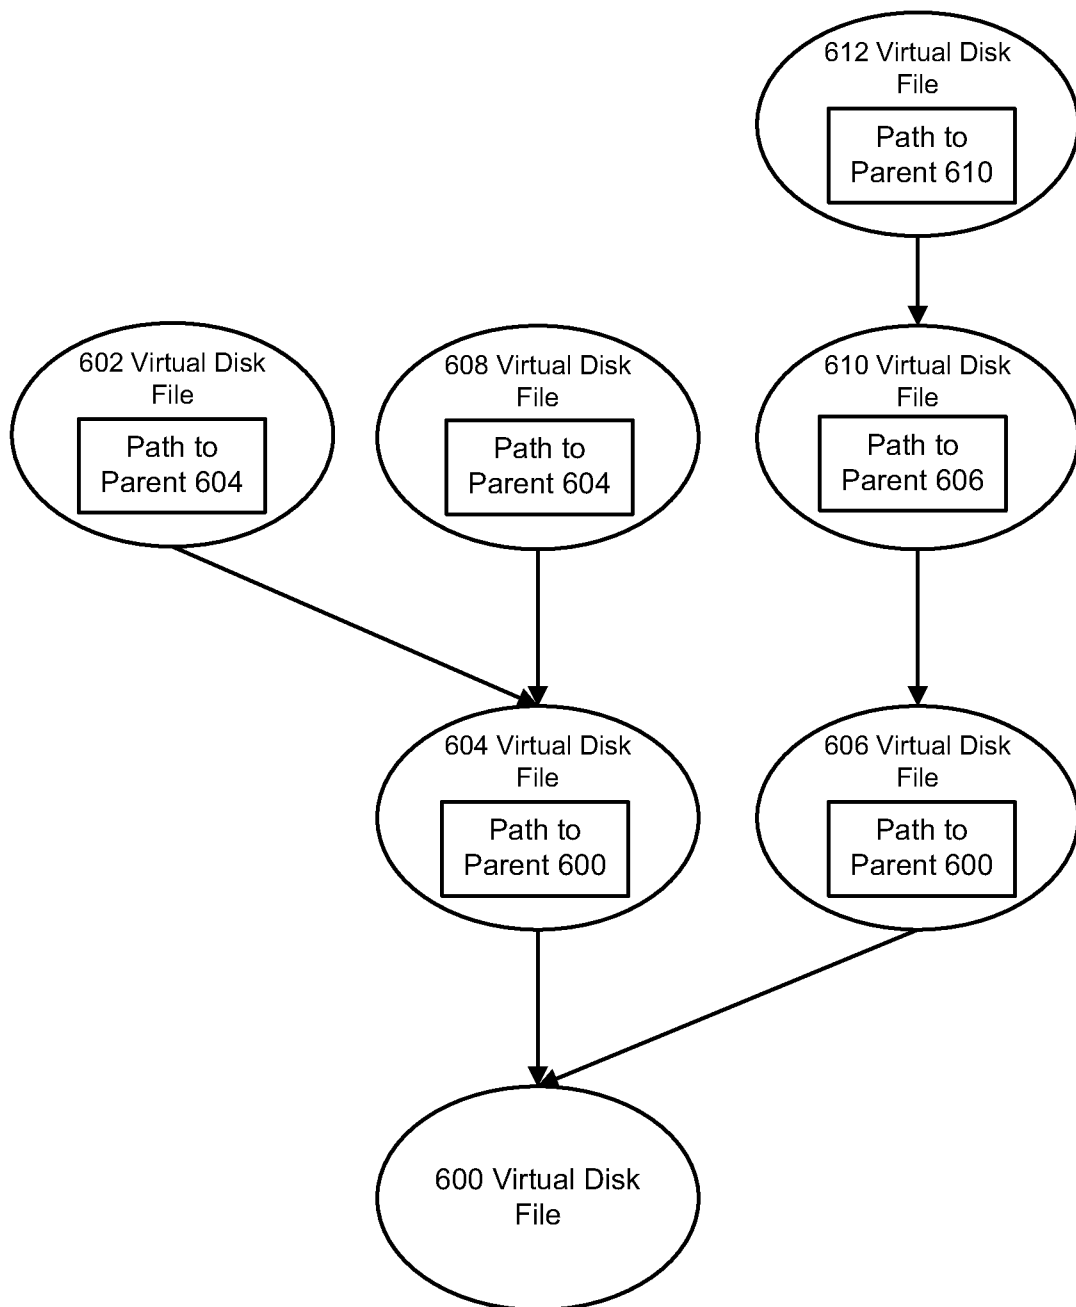
5/18

**FIG. 5A**

6/18

**FIG. 5B**

7/18

**FIG. 6**

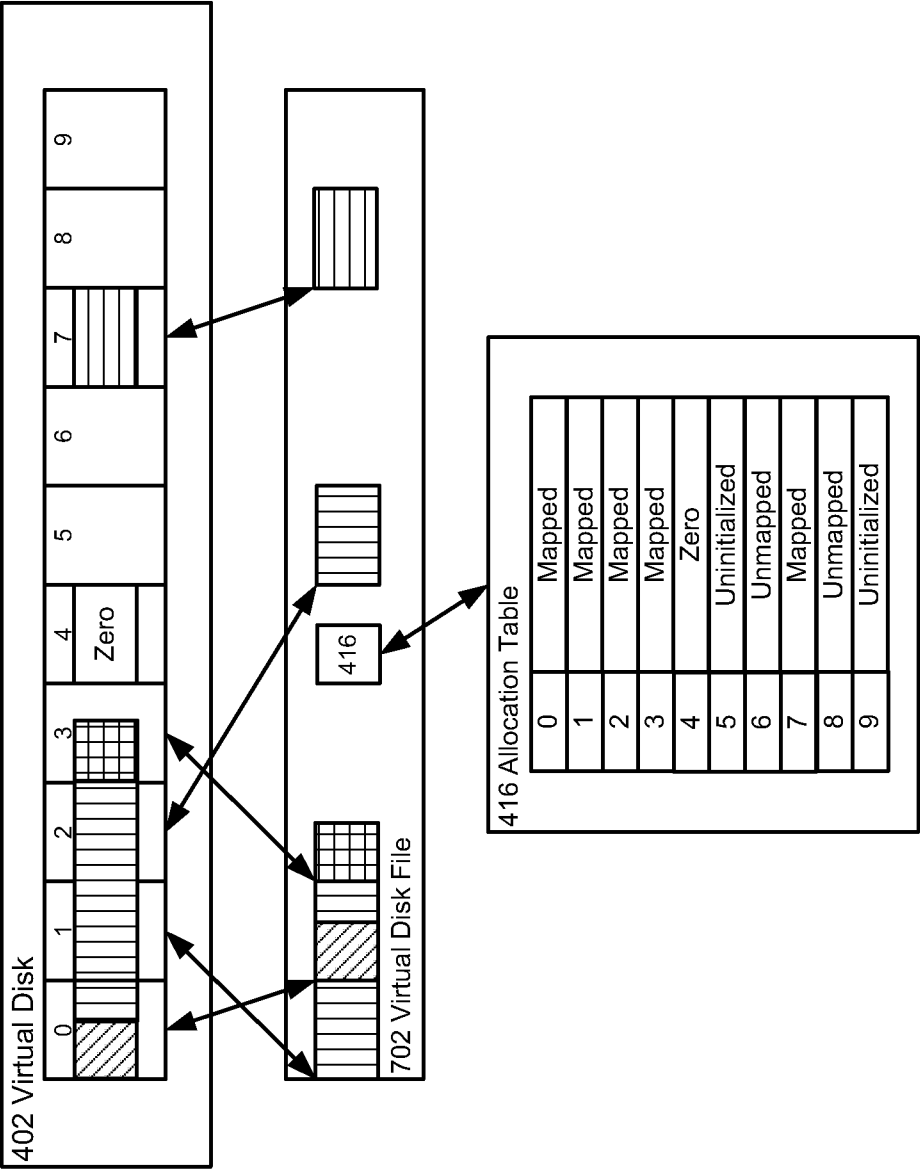


FIG. 7

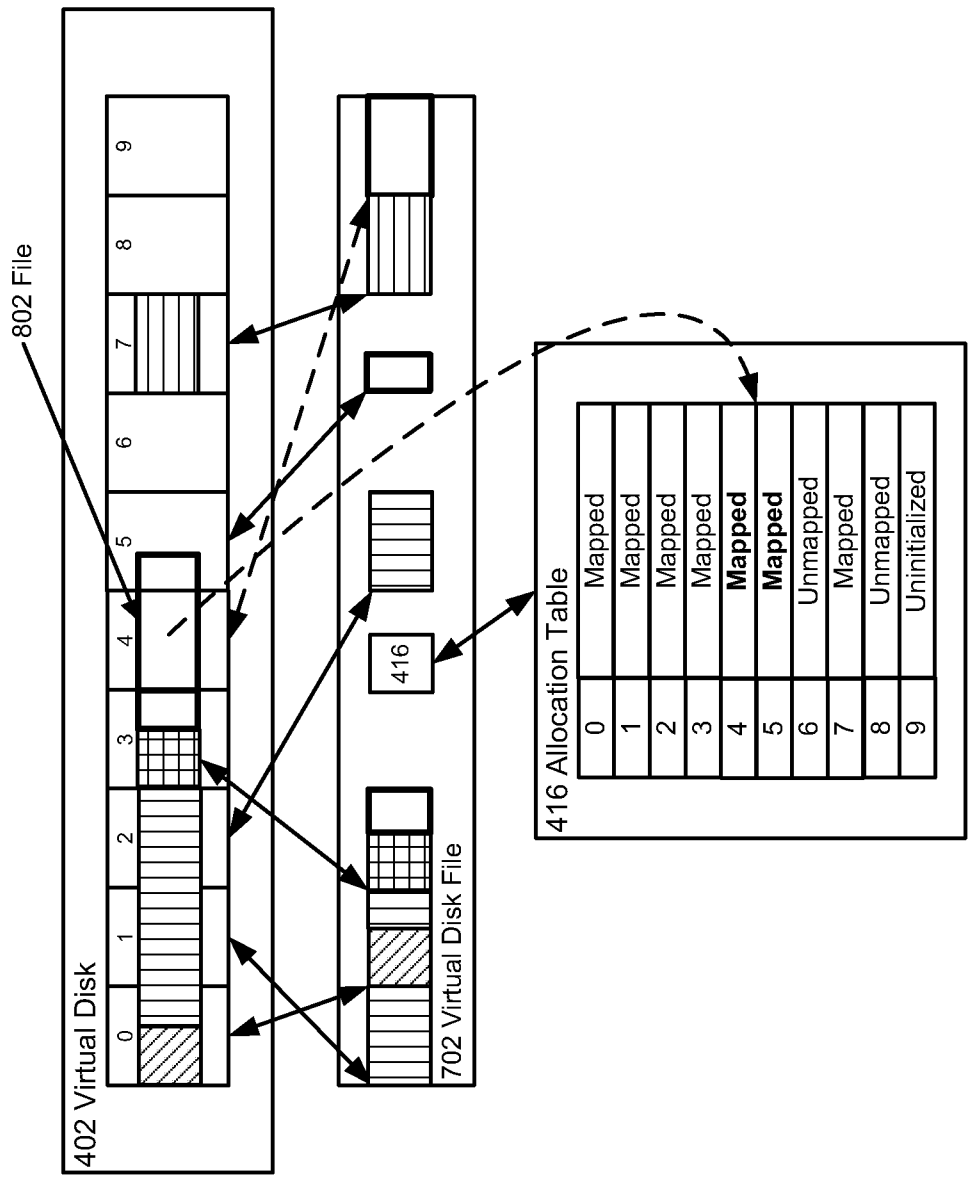


FIG. 8

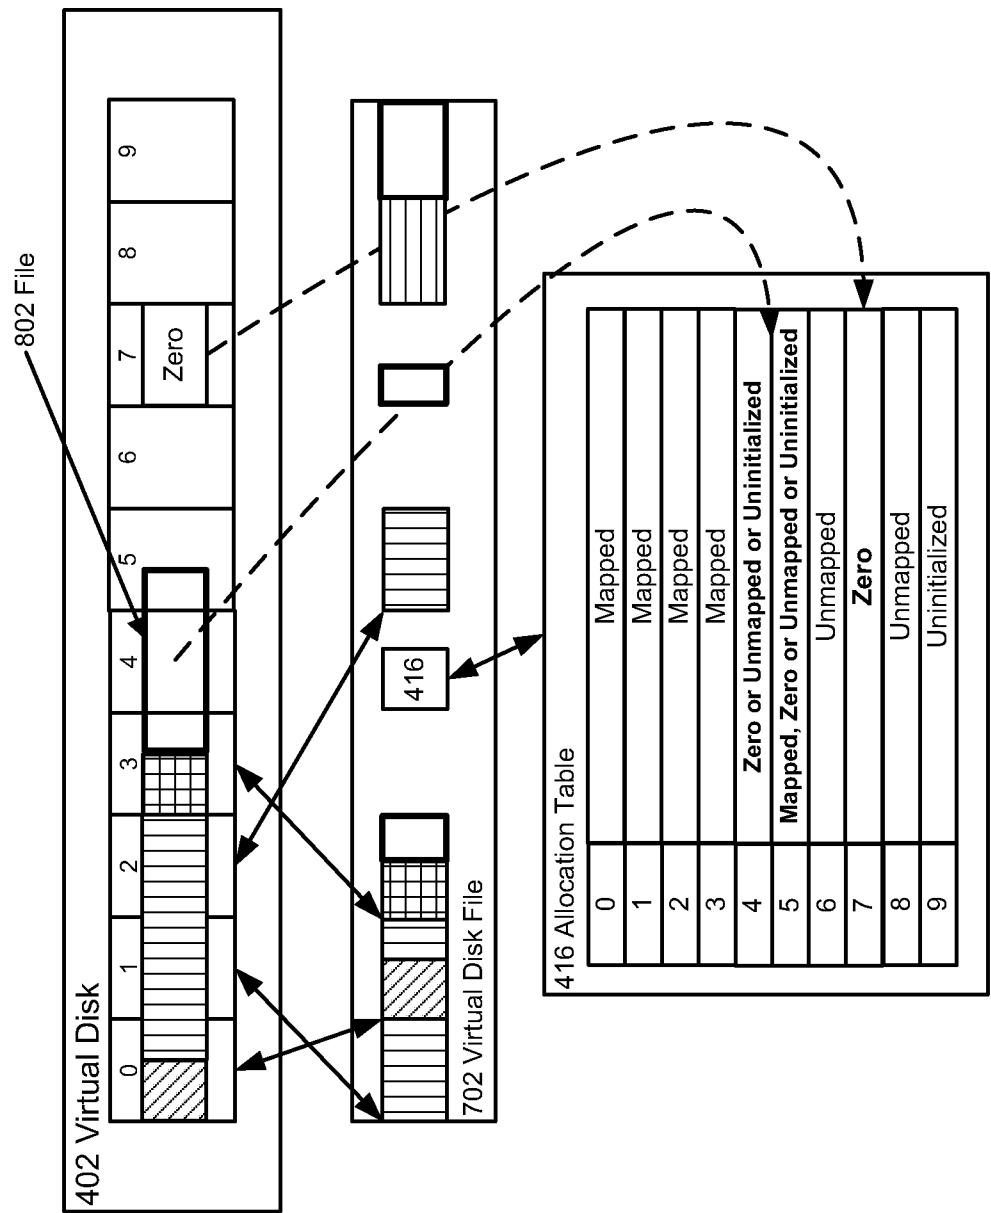


FIG.9

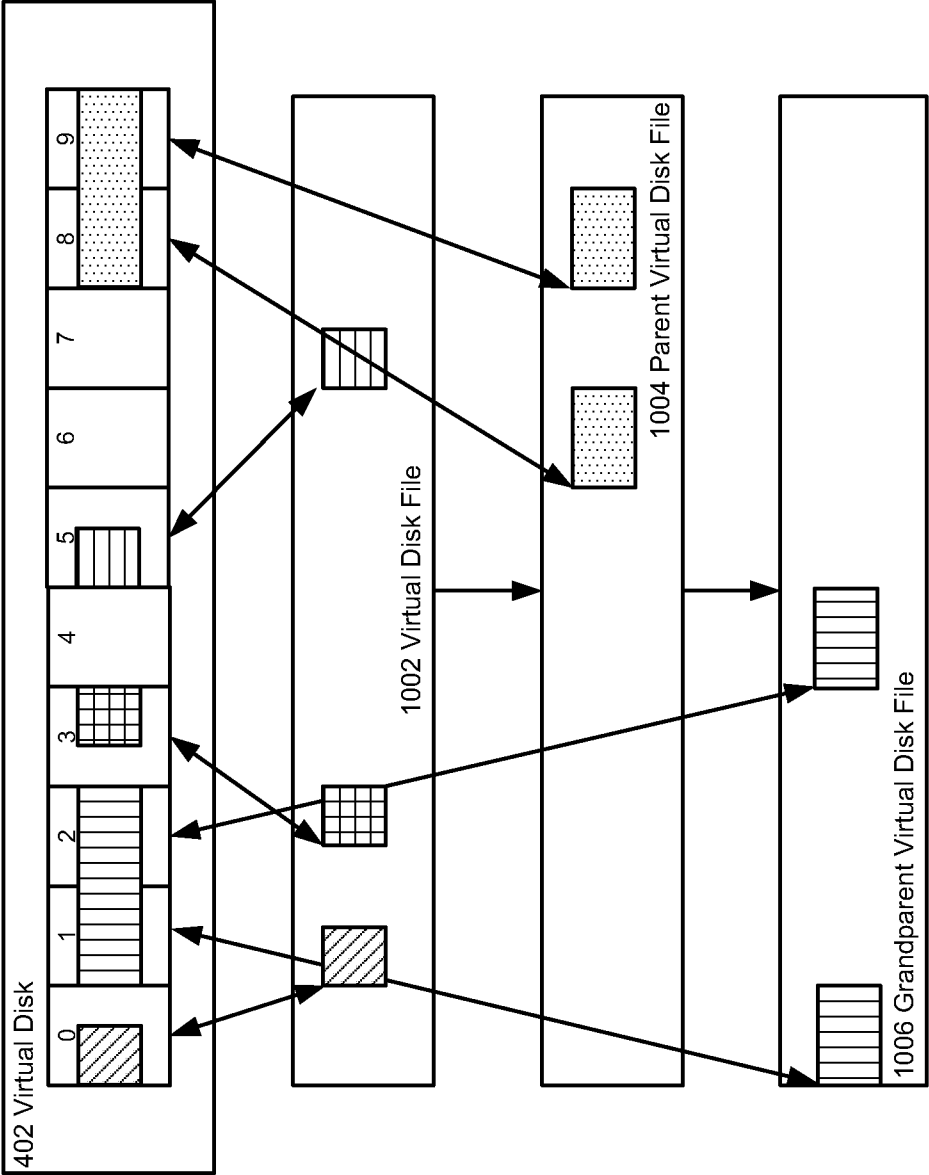
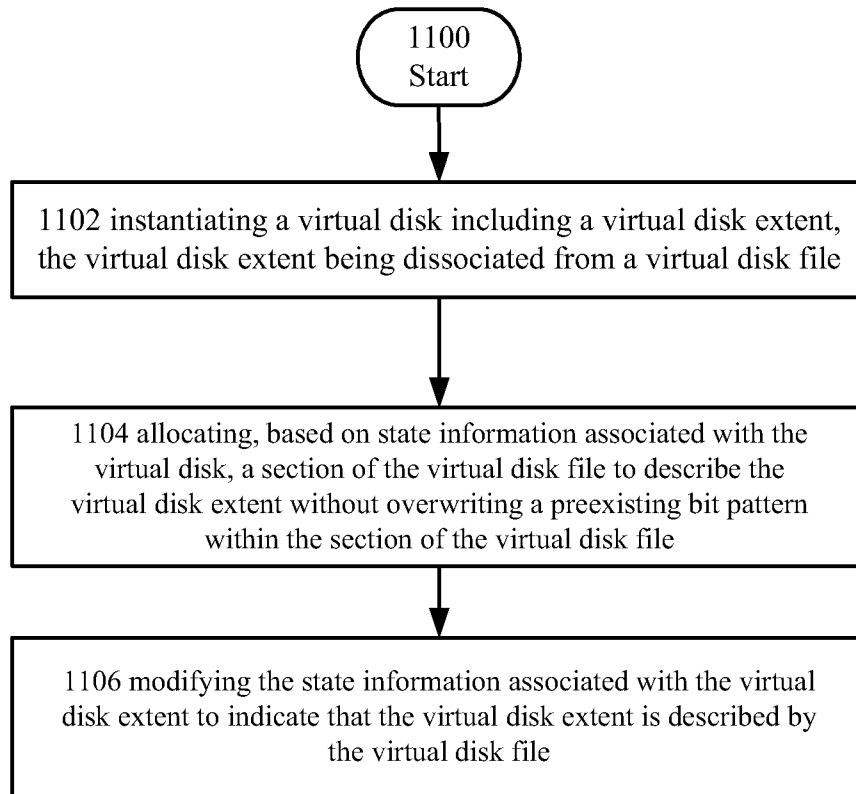


FIG. 10



12/18

**FIG. 11**

13/18

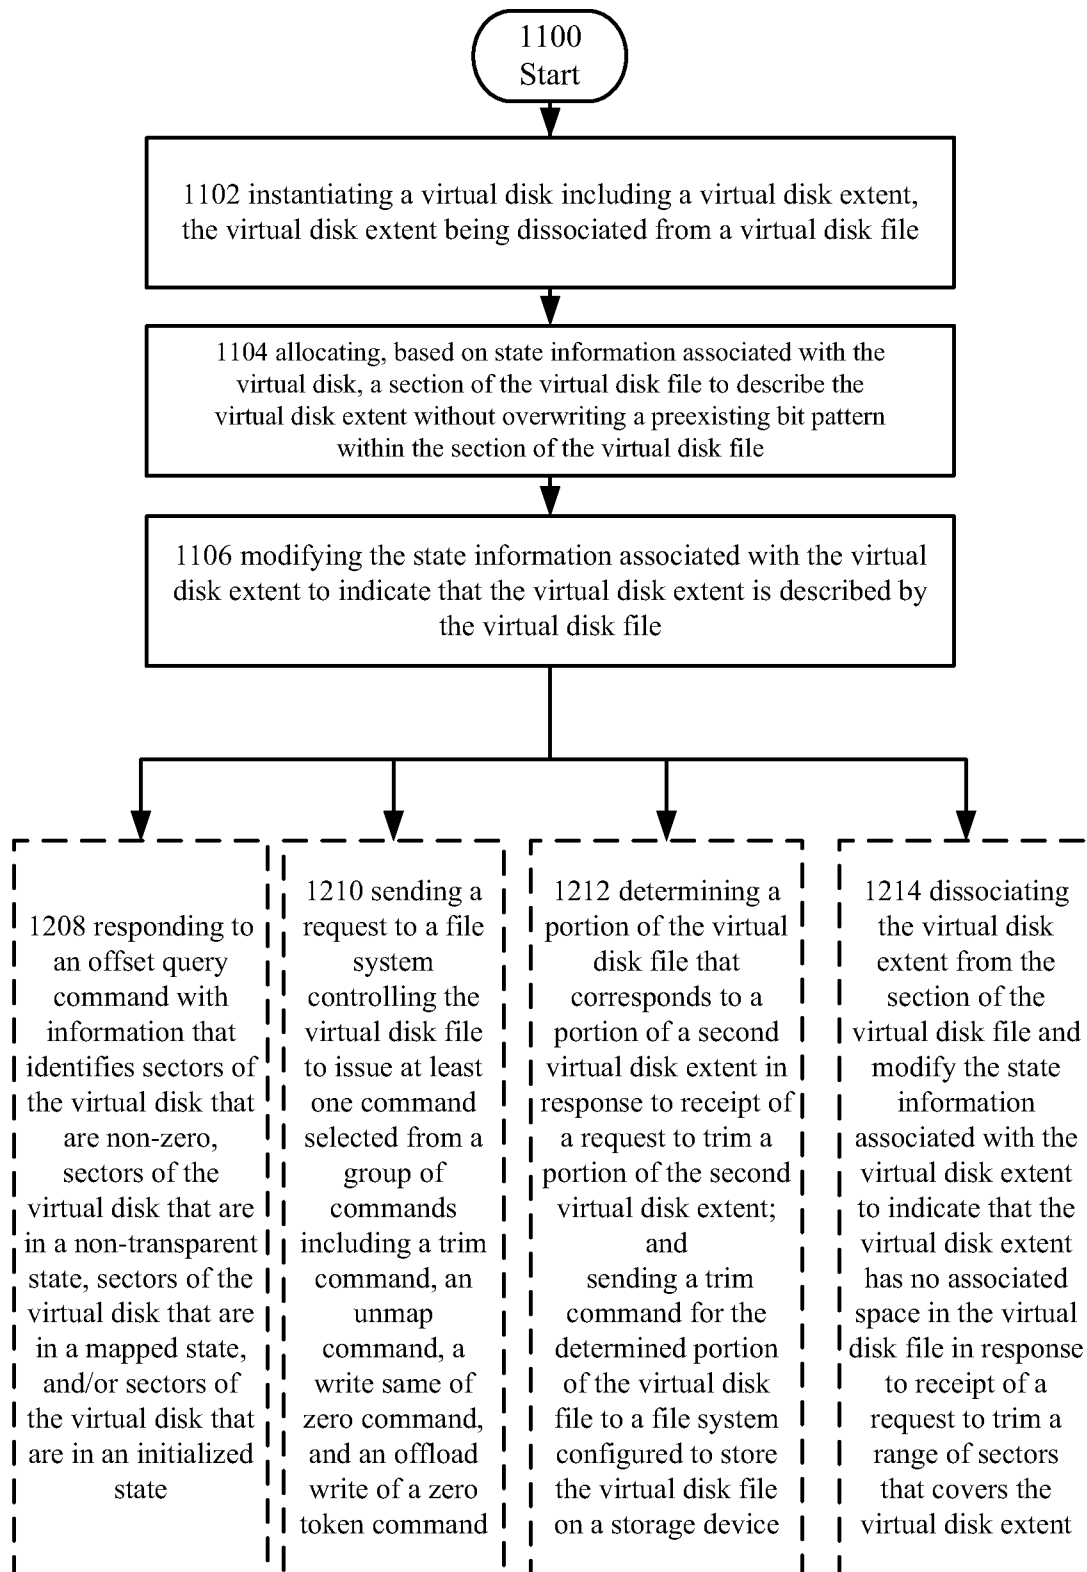


FIG. 12

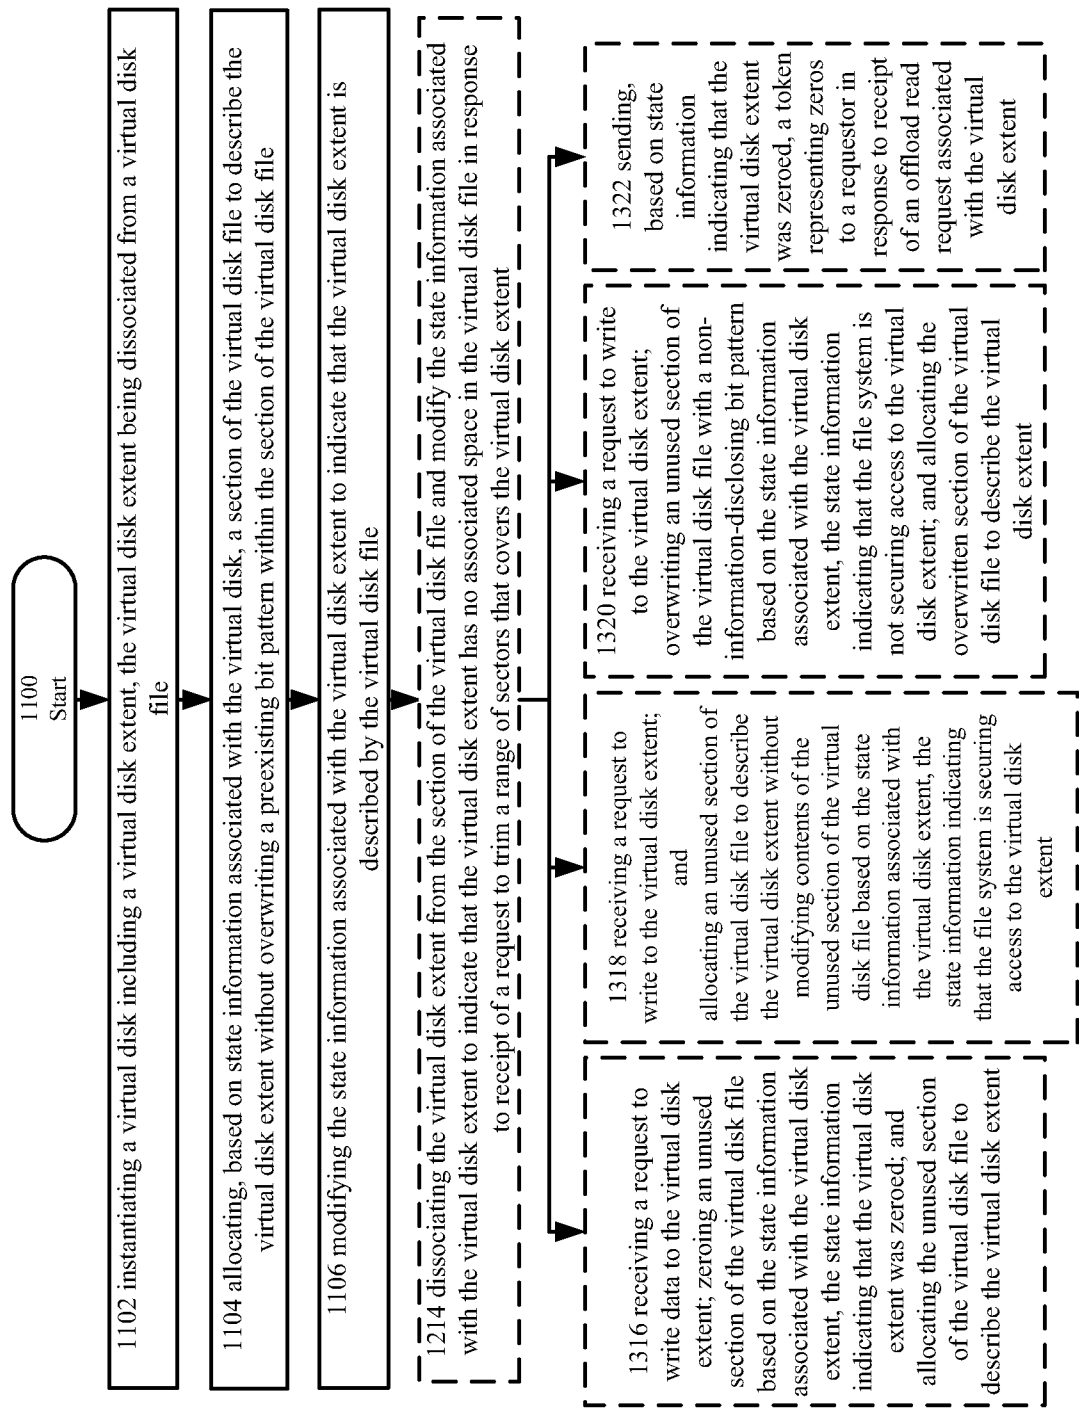
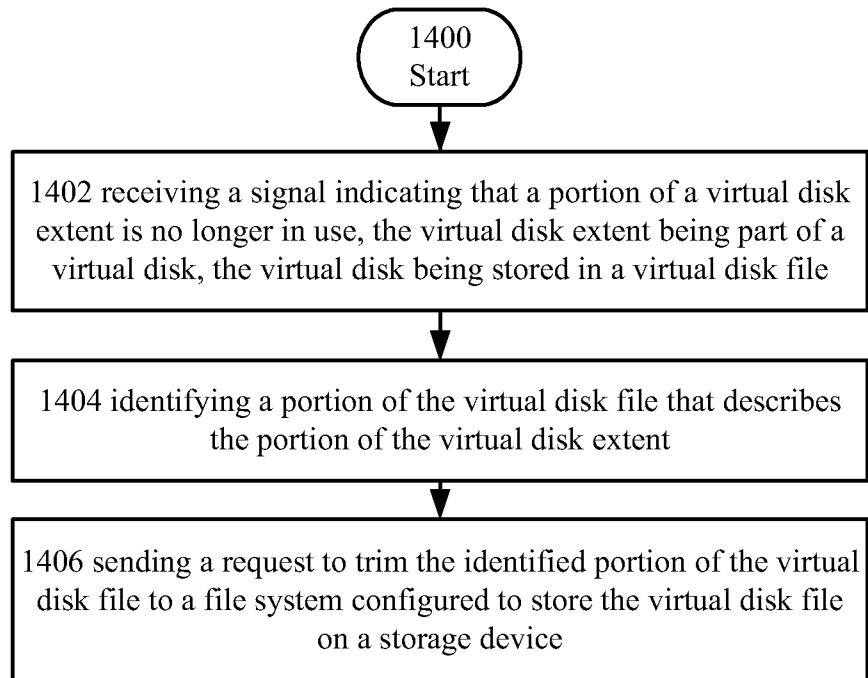


FIG. 13

15/18

**FIG. 14**

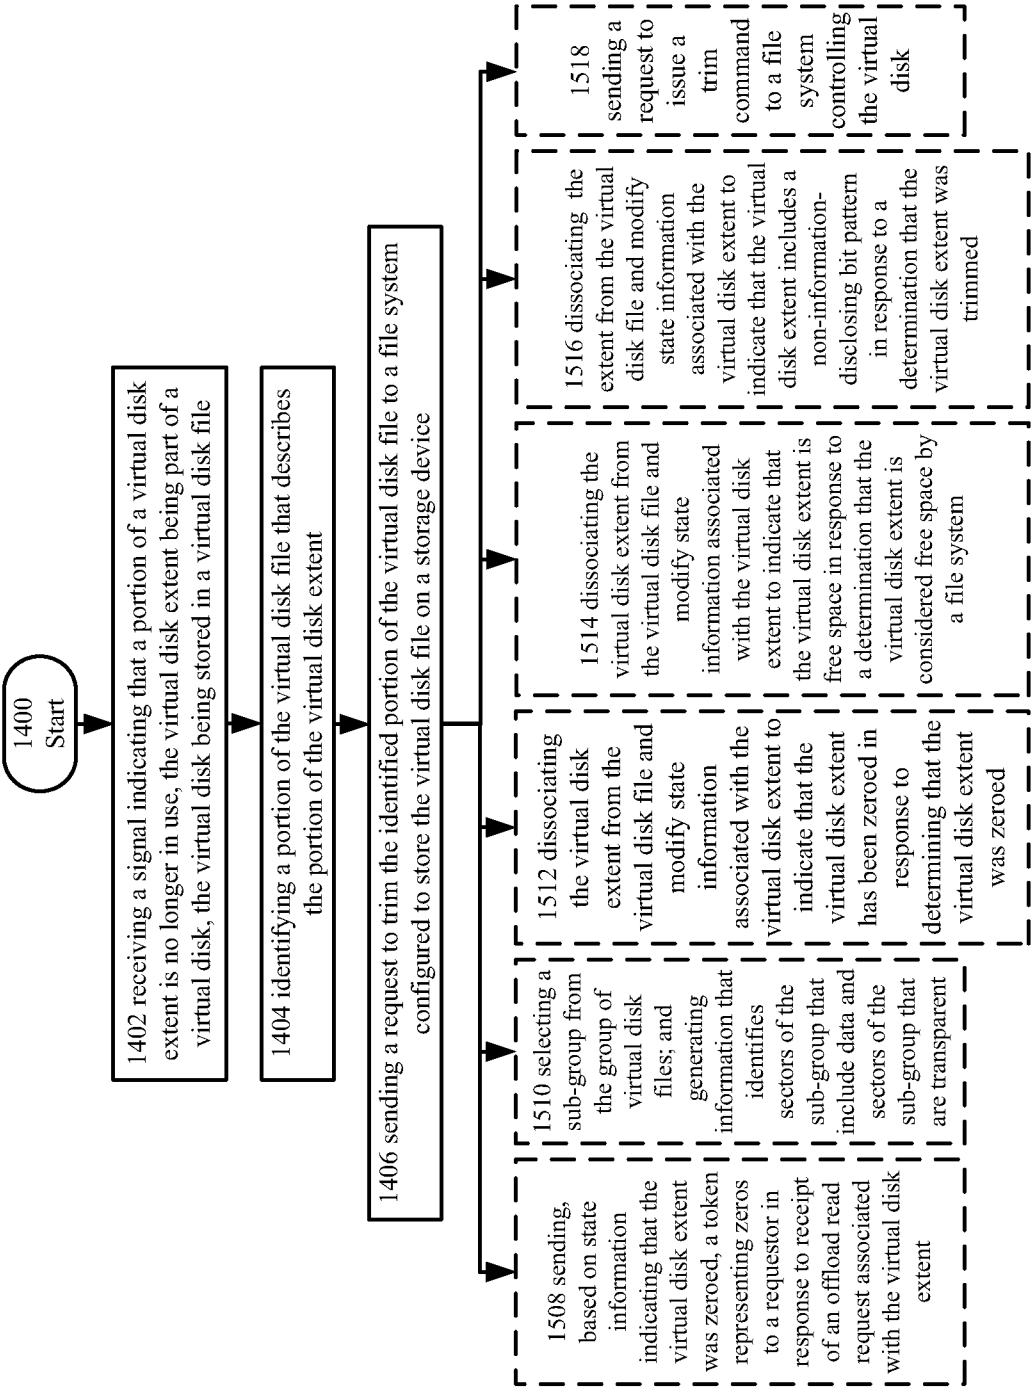
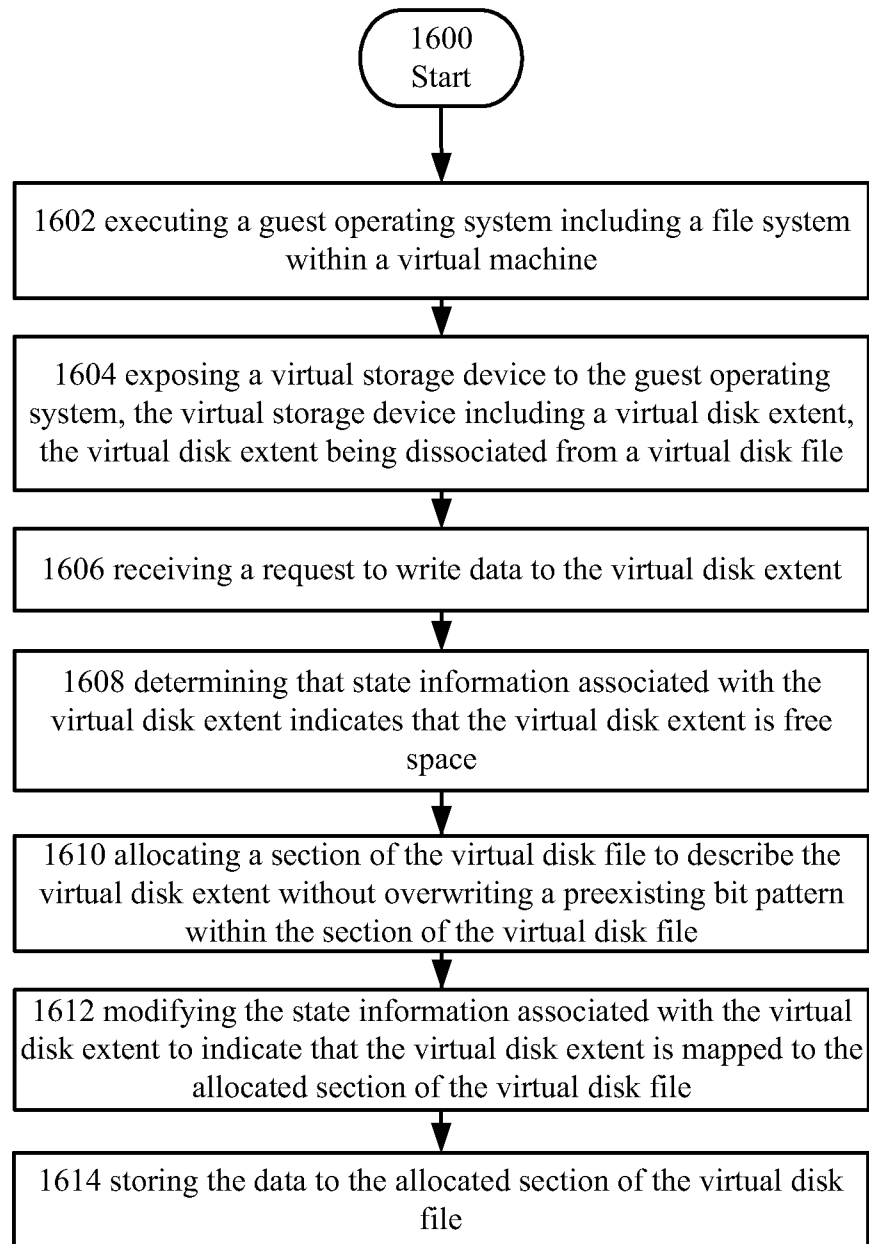


FIG. 15

17/18

**FIG. 16**

18/18

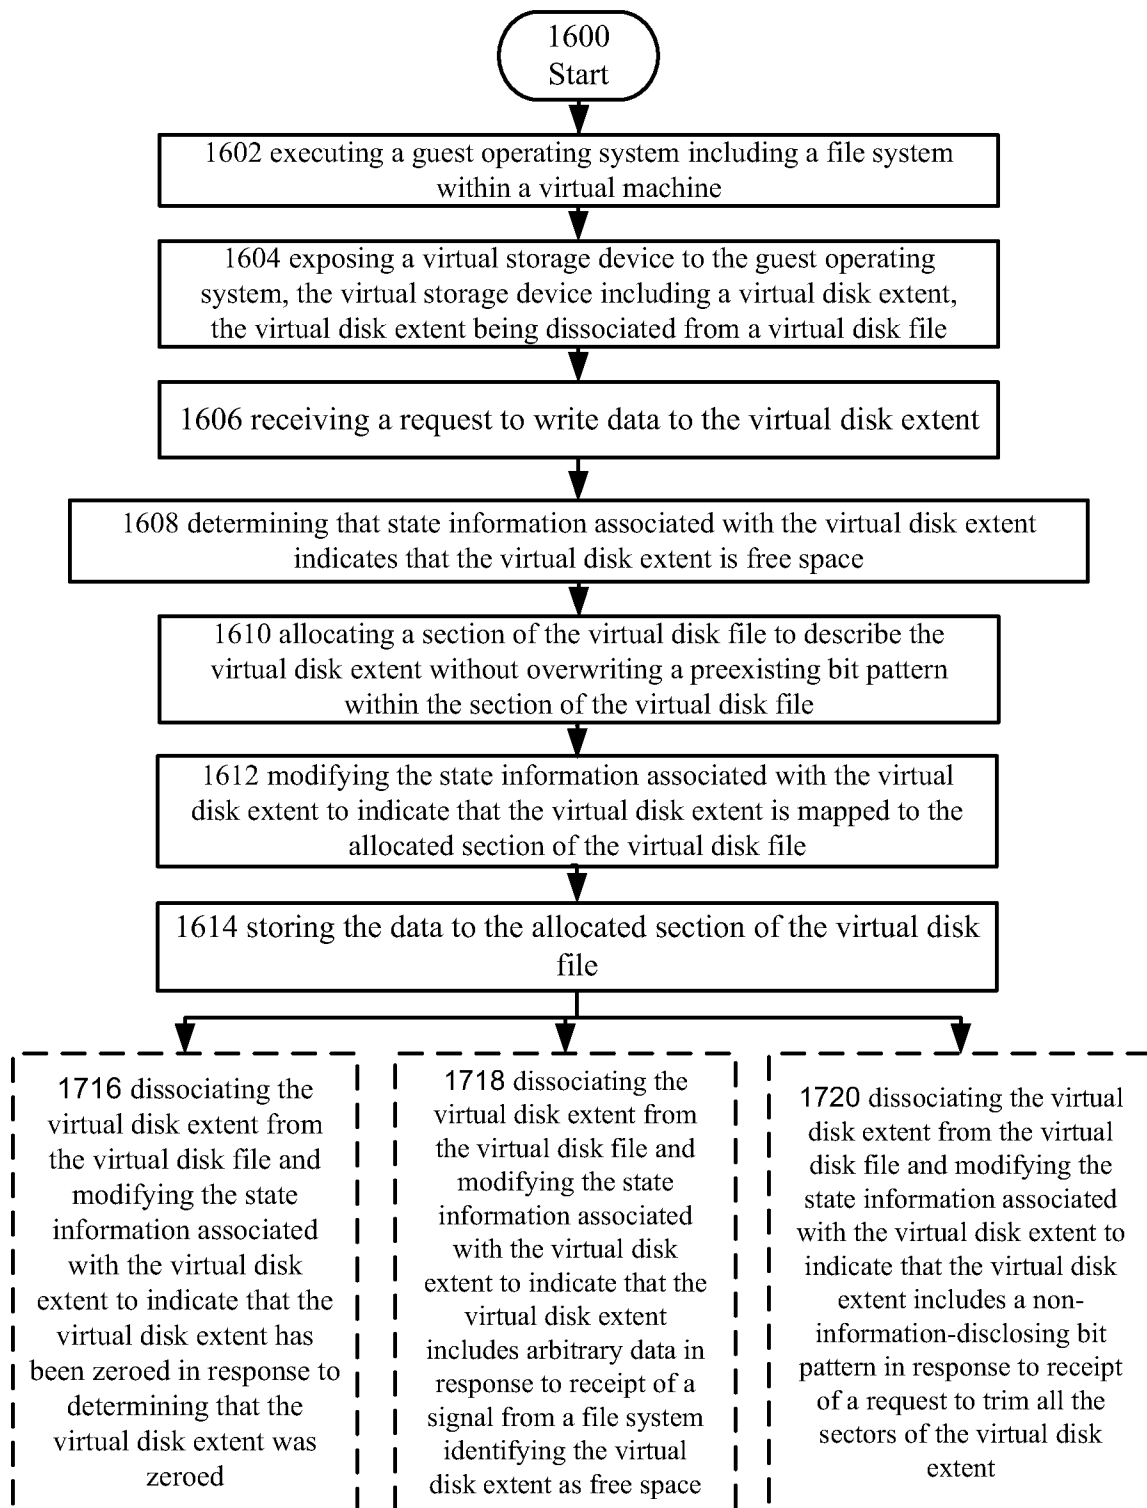


FIG. 17