

(12) **DEMANDE DE BREVET CANADIEN**
CANADIAN PATENT APPLICATION

(13) **A1**

(86) Date de dépôt PCT/PCT Filing Date: 2007/05/03
(87) Date publication PCT/PCT Publication Date: 2007/11/08
(85) Entrée phase nationale/National Entry: 2008/10/31
(86) N° demande PCT/PCT Application No.: US 2007/068139
(87) N° publication PCT/PCT Publication No.: 2007/128005
(30) Priorité/Priority: 2006/05/03 (US60/797,127)

(51) **Cl.Int./Int.Cl.** *G06F 3/06* (2006.01),
G06F 17/30 (2006.01)

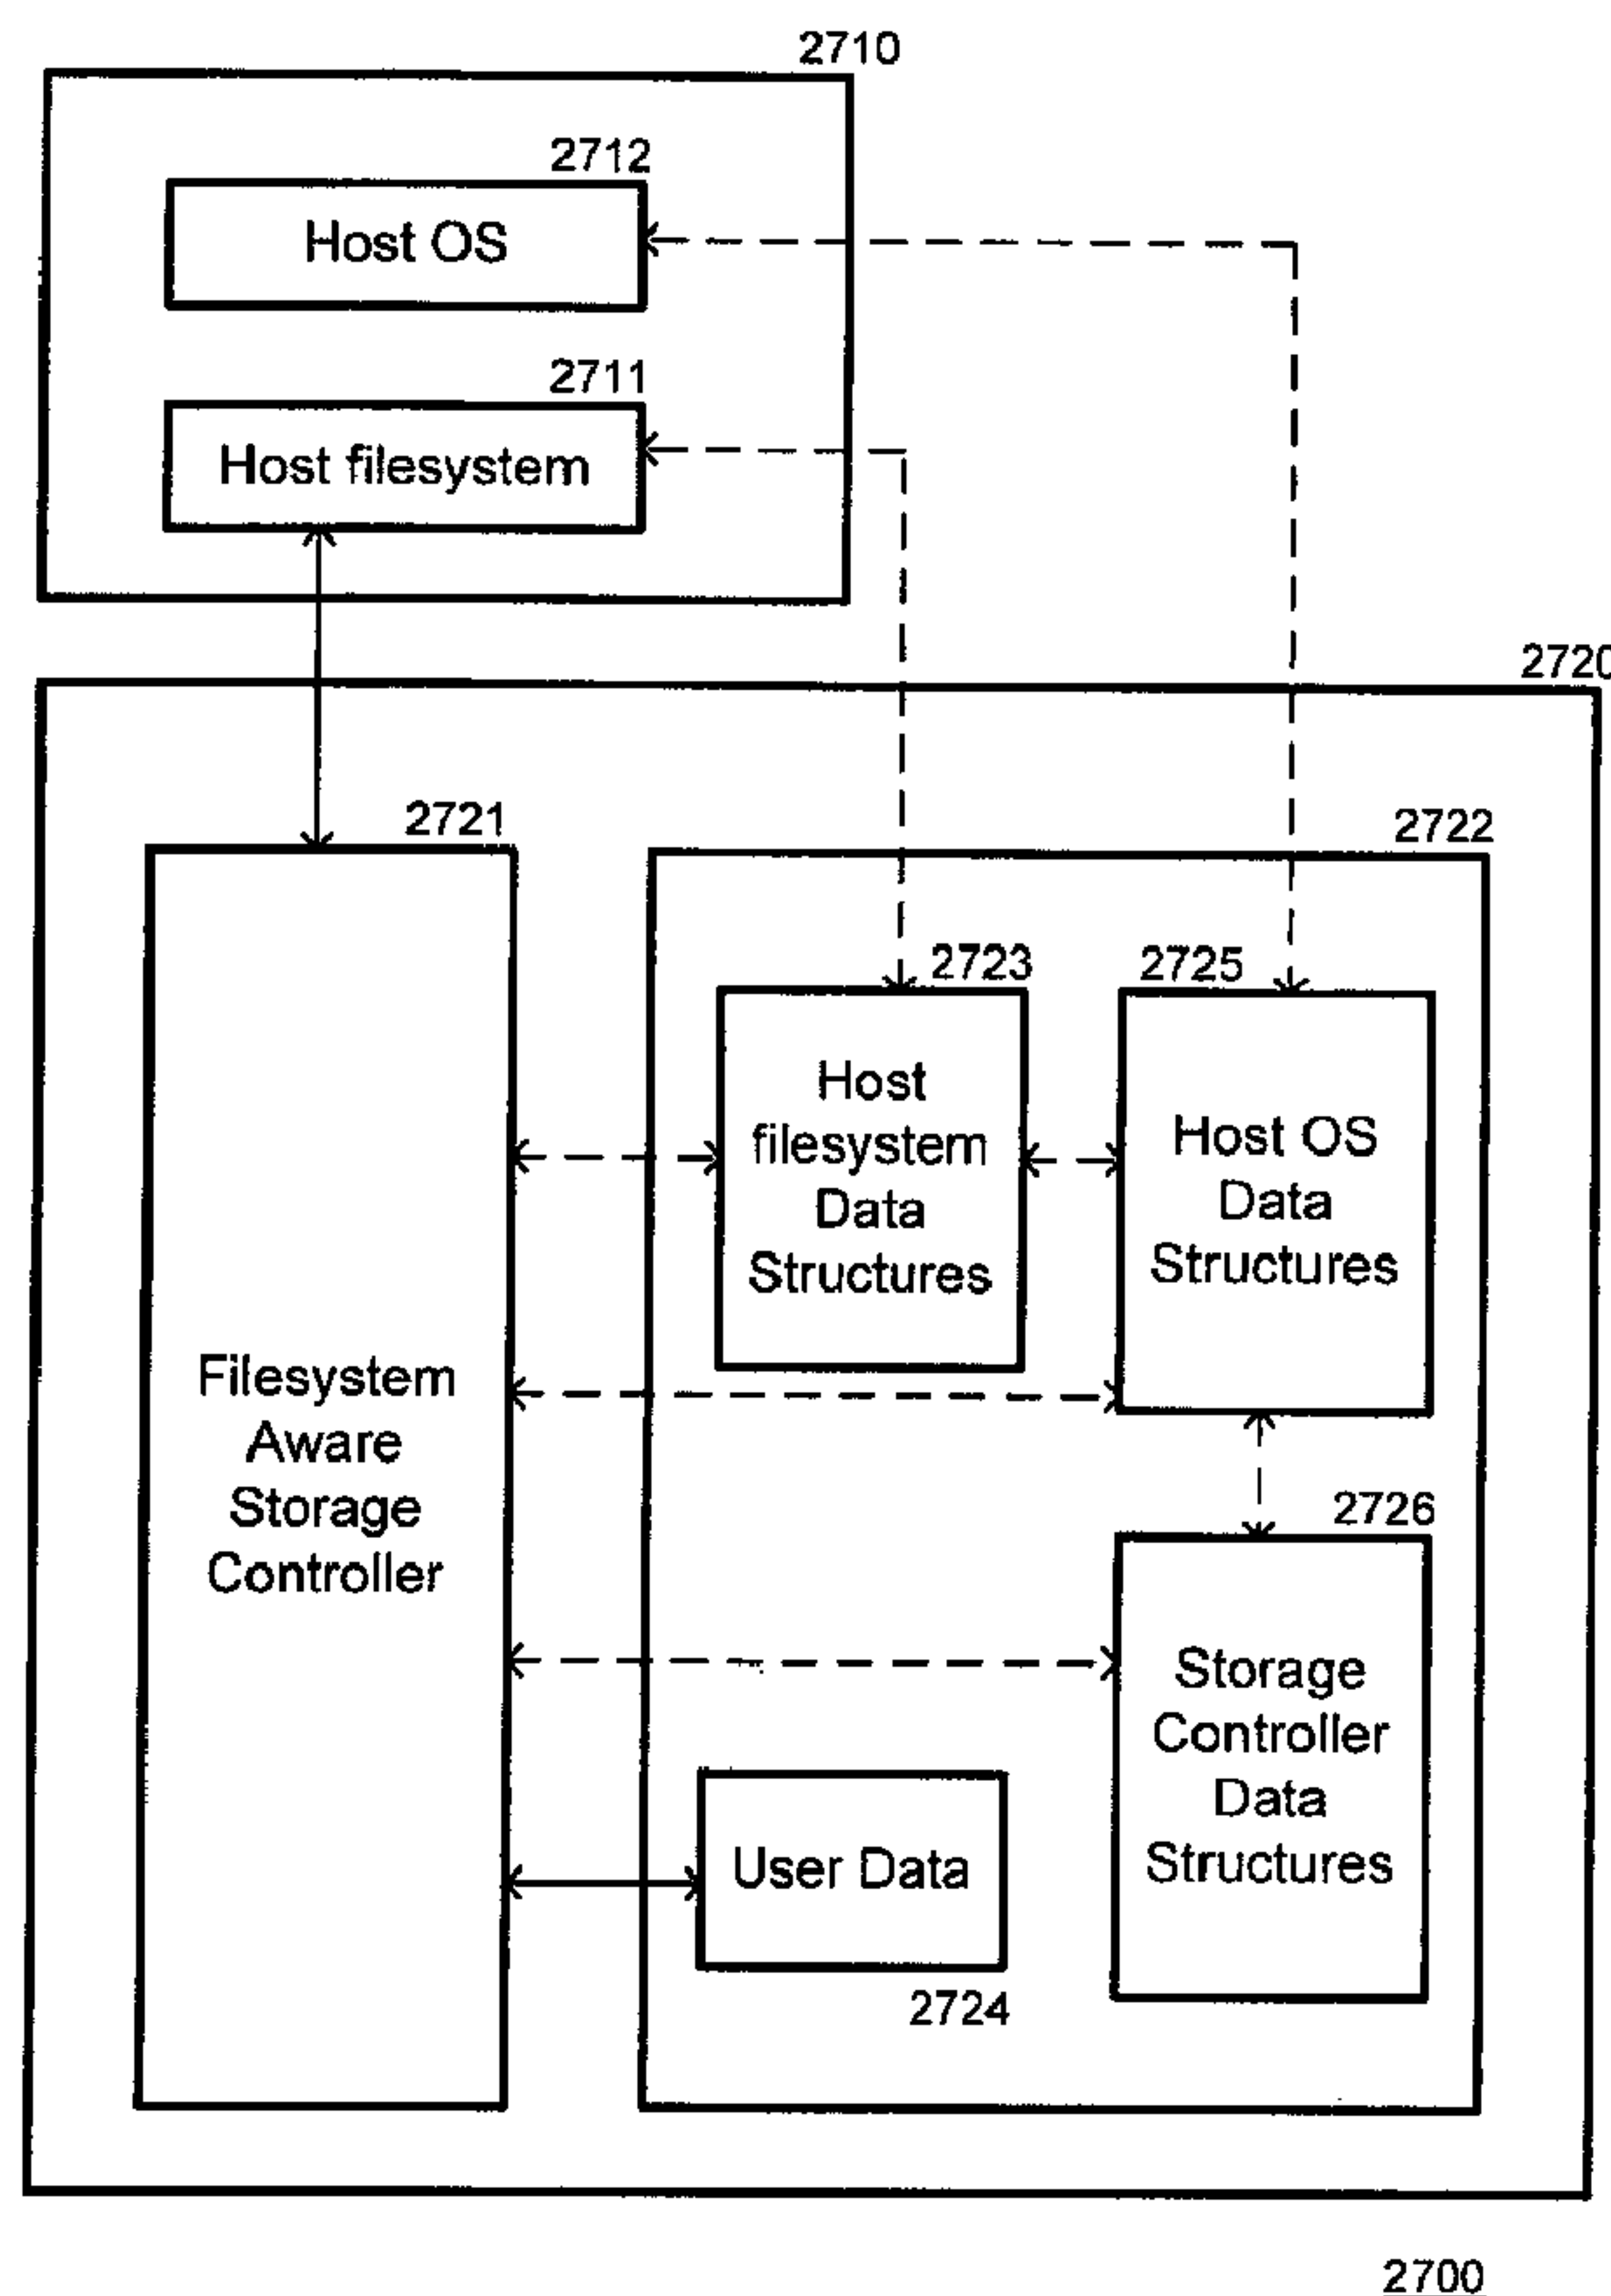
(71) **Demandeur/Applicant:**
DATA ROBOTICS INCORPORATED, US

(72) **Inventeurs/Inventors:**
TERRY, JULIAN M., US;
CLARKSON, NEIL A., GB;
BARRALL, GEOFFREY S., US

(74) **Agent:** GOWLING LAFLEUR HENDERSON LLP

(54) Titre : SYSTEME DE STOCKAGE DE BLOC A RECONNAISSANCE DE SYSTEME DE FICHIERS, APPAREIL ET PROCEDE ASSOCIES

(54) Title: FILESYSTEM-AWARE BLOCK STORAGE SYSTEM, APPARATUS, AND METHOD



(57) **Abrégé/Abstract:**

A filesystem-aware storage system locates and analyzes host filesystem data structures in order to determine storage usage of the host filesystem. To this end, the storage system might locate an operating system partition, parse the operating system partition to

(57) **Abrégé(suite)/Abstract(continued):**

locate its data structures, and parse the operating system data structures to locate the host filesystem data structures. The storage system manages data storage based on the storage usage of the host file system. The storage system can use the storage usage information to identify storage areas that are no longer being used by the host filesystem and reclaim those areas for additional data storage capacity. Also, the storage system can identify the types of data stored by the host filesystem and manage data storage based on the data types, such as selecting a storage layout and/or an encoding scheme for the data based on the data type.



(10) International Publication Number
WO 2007/128005 A3

[Continued on next page]

(57) Abstract: A filesystem-aware storage system locates and analyzes host filesystem data structures in order to determine storage usage of the host filesystem. To this end, the storage system might locate an operating system partition, parse the operating system partition to locate its data structures, and parse the operating system data structures to locate the host filesystem data structures. The storage system manages data storage based on the storage usage of the host file system. The storage system can use the storage usage information to identify storage areas that are no longer being used by the host filesystem and reclaim those areas for additional data storage capacity. Also, the storage system can identify the types of data stored by the host filesystem and manage data storage based on the data types, such as selecting a storage layout and/or an encoding scheme for the data based on the data type.

WO 2007/128005 A3



Published:

- *with international search report*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(88) Date of publication of the international search report:

24 January 2008

**FILESYSTEM-AWARE BLOCK STORAGE
SYSTEM, APPARATUS, AND METHOD**

Priority

This PCT application claims priority from United States Provisional Patent
5 Application No. 60/797,127 entitled **Filesystem-Aware Block Storage System,
Apparatus, and Method** filed on May 3, 2006 in the names of Julian M. Terry, Neil A.
Clarkson, and Geoffrey S. Barrall.

This application is also related to United States Patent Application No.
11/267,938 entitled **Dynamically Expandable and Contractible Fault-Tolerant**
10 **Storage System Permitting Variously Sized Storage Devices and Method** filed on
November 4, 2005 in the name of Geoffrey S. Barrall, which claims priority from United
States Provisional Patent Application No. 60/625,495 filed on November 5, 2004 and
from United States Provisional Patent Application No. 60/718,768 filed on September 20,
2005.

15 All of the above patent applications are hereby incorporated herein by reference in
their entireties.

Technical Field and Background Art

The present invention relates to digital data storage systems and methods, and
more particularly to those providing fault-tolerant storage.

20 It is known in the prior art to provide redundant disk storage in a pattern
according to any one of various RAID (Redundant Array of Independent Disks)
protocols. Typically disk arrays using a RAID pattern are complex structures that require
management by experienced information technologists. Moreover in many array designs
using a RAID pattern, if the disk drives in the array are of non-uniform capacities, the
25 design may be unable to use any capacity on the drive that exceeds the capacity of the
smallest drive in the array.

One problem with a standard RAID system is that it is possible for disc-surface
corruption to occur on an infrequently used area of the disk array. In the event that
another drive fails, it is not always possible to determine that corruption has occurred. In

such a case, the corrupted data may be propagated and preserved when the RAID array rebuilds the failed drive.

In many storage systems, a spare storage device will be maintained in a ready state so that it can be used in the event another storage device fails. Such a spare storage device is often referred to as a “hot spare.” The hot spare is not used to store data during normal operation of the storage system. When an active storage device fails, the failed storage device is logically replaced by the hot spare, and data is moved or otherwise recreated onto the hot spare. When the failed storage device is repaired or replaced, the data is typically moved or otherwise recreated onto the (re-)activated storage device, and the hot spare is brought offline so that it is ready to be used in the event of another failure. Maintenance of a hot spare disk is generally complex, and so is generally handled by a skilled administrator. A hot spare disk also represents an added expense.

Generally speaking, when the host filesystem writes a block of data to the storage system, the storage system allocates a storage block for the data and updates its data structures to indicate that the storage block is in use. From that point on, the storage system considers the storage block to be in use, even if the host filesystem subsequently ceases to use its block.

The host filesystem generally uses a bitmap to track its used disk blocks. Shortly after volume creation, the bitmap will generally indicate that most blocks are free, typically by having all bits clear. As the filesystem is used, the host filesystem will allocate blocks solely through use of its free block bitmap.

When the host filesystem releases some blocks back to its free pool, it simply clears the corresponding bits in its free block bitmap. On the storage system, this is manifested as a write to a cluster that happens to contain part of the host’s free block bitmap, and possibly a write to a journal file; almost certainly no input/output (I/O) to the actual cluster being freed itself. If the host filesystem were running in an enhanced security mode, there might be I/O to the freed block due to overwriting of the current on-disk data by the host so as to reduce the chance of the stale cluster contents being readable by an attacker, but there is no way to identify such writes as being part of a deletion process. Thus, the storage device has no way to distinguish a block that the host filesystem has in use from one that it previously used and has subsequently marked free.

This inability of the storage system to identify freed blocks can lead to a number of negative consequences. For example, the storage system could significantly over-report the amount of storage being used and could prematurely run out of storage space.

Summary of the Invention

5 In accordance with one aspect of the invention there is provided a method of storing data in by a block-level storage system that stores data under control of a host filesystem. The method involves locating host filesystem data structures stored for the host filesystem in the block-level storage system; analyzing the host filesystem data structures to identify a data type associated with the data to be stored; and storing the data
10 using a storage scheme selected based on the data type, whereby data having different data types can be stored using different storage schemes selected based on the data types.

 In accordance with another aspect of the invention there is provided a block-level storage system that stores data under control of a host filesystem. The system comprises a block-level storage in which host filesystem data structures are stored for the host
15 filesystem and a storage controller operably coupled to the block-level storage for locating the host filesystem data structures stored in the block-level storage, analyzing the host filesystem data structures to identify a data type associated with the data to be stored, and storing the data using a storage scheme selected based on the data type, whereby data having different data types can be stored using different storage schemes selected based
20 on the data types.

 In various alternative embodiments, the data may be stored using a storage layout and/or an encoding scheme selected based on the data type. For example, frequently accessed data may be stored so as to provide enhanced accessibility (e.g., in an uncompressed form and in sequential storage), while infrequently access data may be
25 stored so as to provide enhanced storage efficiency (e.g., using data compression and/or non-sequential storage). Additionally or alternatively, the data may compressed and/or encrypted depending on the data type.

 In various alternative embodiments, the host filesystem data structures may be located by maintaining a partition table; parsing the partition table to locate an operating
30 system partition; parsing the operating system partition to identify the operating system

and locate operating system data structures; and parsing the operating system data structures to identify the host filesystem and locate the host filesystem data structures. The operating system data structures may include a superblock, in which case parsing the operating system data structures may include parsing the superblock. The host filesystem data structures may be parsed by making a working copy of a host filesystem data structure and parsing the working copy.

Brief Description of the Drawings

The foregoing features of the invention will be more readily understood by reference to the following detailed description, taken with reference to the accompanying drawings, in which:

Fig. 1 is an illustration of an embodiment of the invention in which an object is parsed into a series of chunks for storage.

Fig. 2 illustrates in the same embodiment how a pattern for fault-tolerant storage for a chunk may be dynamically changed as a result of the addition of more storage.

Fig. 3 illustrates in a further embodiment of the invention the storage of chunks in differing fault-tolerant patterns on a storage system constructed using different sized storage devices.

Fig. 4 illustrates another embodiment of the invention in which indicator states are used to warn of inefficient storage use and low levels of fault tolerance.

Fig. 5 is a block diagram of functional modules used in the storage, retrieval and re-layout of data in accordance with an embodiment of the invention.

Fig. 6 shows an example in which mirroring is used in an array containing more than two drives.

Fig. 7 shows some exemplary zones using different layout schemes to store their data.

Fig. 8 shows a lookup table for implementing sparse volumes.

Fig. 9 shows status indicators for an exemplary array having available storage space and operating in a fault-tolerant manner, in accordance with an exemplary embodiment of the present invention.

Fig. 10 shows status indicators for an exemplary array that does not have enough space to maintain redundant data storage and more space must be added, in accordance with an exemplary embodiment of the present invention.

5 Fig. 11 shows status indicators for an exemplary array that would be unable to maintain redundant data in the event of a failure, in accordance with an exemplary embodiment of the present invention.

Fig. 12 shows status indicators for an exemplary array in which a storage device has failed, in accordance with an exemplary embodiment of the present invention. Slots B, C, and D are populated with storage devices.

10 Fig. 13 shows a module hierarchy representing the different software layers of an exemplary embodiment and how they relate to one another.

Fig. 14 shows how a cluster access table is used to access a data clusters in a Zone, in accordance with an exemplary embodiment of the present invention.

15 FIG. 15 shows a journal table update in accordance with an exemplary embodiment of the present invention.

Fig. 16 shows drive layout in accordance with an exemplary embodiment of the invention.

Fig. 17 demonstrates the layout of Zone 0 and how other zones are referenced, in accordance with an exemplary embodiment of the invention.

20 Fig. 18 demonstrates read error handling in accordance with an exemplary embodiment of the invention.

Fig. 19 demonstrates write error handling in accordance with an exemplary embodiment of the invention.

25 Fig. 20 is a logic flow diagram demonstrating backup of a bad Region by the Error Manager in accordance with an exemplary embodiment of the invention.

Fig. 21 is a schematic block diagram showing the relevant components of a storage array in accordance with an exemplary embodiment of the present invention.

Fig. 22 is a logic flow diagram showing exemplary logic for managing a virtual hot spare in accordance with an exemplary embodiment of the present invention.

Fig. 23 is a logic flow diagram showing exemplary logic for determining a re-layout scenario for each possible disk failure, as in block 2102 of Fig. 22, in accordance with an exemplary embodiment of the present invention.

Fig. 24 is a logic flow diagram showing exemplary logic for invoking the virtual hot spare functionality in accordance with an exemplary embodiment of the present invention.

Fig. 25 is a logic flow diagram showing exemplary logic for automatically reconfiguring the one or more remaining drives to restore fault tolerance for the data, as in block 2306 of Fig. 24, in accordance with an exemplary embodiment of the present invention.

Fig. 26 is a logic flow diagram showing exemplary logic for upgrading a storage device, in accordance with an exemplary embodiment of the present invention.

FIG. 27 is a conceptual block diagram of a computer system in accordance with an exemplary embodiment of the present invention.

FIG. 28 is high-level logic flow diagram for the filesystem-aware storage controller, in accordance with an exemplary embodiment of the present invention.

FIG. 29 is a logic flow diagram for locating the host filesystem data structures, in accordance with an exemplary embodiment of the present invention.

FIG. 30 is a logic flow diagram for reclaiming unused storage space, in accordance with an exemplary embodiment of the present invention.

FIG. 31 is a logic flow diagram for managing storage of the user data based on the data types, in accordance with an exemplary embodiment of the present invention.

FIG. 32 is a schematic block diagram showing the relevant components of a scavenger, in accordance with an exemplary embodiment of the present invention.

FIG. 33 is pseudo code for locating the host filesystem bitmaps, in accordance with an exemplary embodiment of the present invention.

FIG. 34 is high-level pseudo code for the BBUM, in accordance with an exemplary embodiment of the present invention.

FIG. 35 is high-level pseudo code for synchronous processing of an LBA 0 update creating a new partition, in accordance with an exemplary embodiment of the present invention.

FIG. 36 is high-level pseudo code for synchronous processing of an LBA 0 update (re)formatting a partition, in accordance with an exemplary embodiment of the present invention.

FIG. 37 is high-level pseudo code for synchronous processing of an LBA 0 update deleting a partition, in accordance with an exemplary embodiment of the present invention.

FIG. 38 is high-level pseudo code for the asynchronous task, in accordance with an exemplary embodiment of the present invention.

Detailed Description of Specific Embodiments

Definitions. As used in this description and the accompanying claims, the following terms shall have the meanings indicated, unless the context otherwise requires:

A “chunk” of an object is an abstract slice of an object, made independently of any physical storage being used, and is typically a fixed number of contiguous bytes of the object.

A fault-tolerant “pattern” for data storage is the particular which by data is distributed redundantly over one or more storage devices, and may be, among other things, mirroring (e.g., in a manner analogous to RAID1), striping (e.g., in a manner analogous to RAID5), RAID6, dual parity, diagonal Parity, Low Density Parity Check codes, turbo codes, or other redundancy scheme or combination of redundancy schemes.

A hash number for a given chunk is “unique” when the given chunk produces a hash number that generally will differ from the hash number for any other chunk, except when the other chunk has data content identical to the given chunk. That is, two chunks will generally have different hash numbers whenever their content is non-identical. As described in further detail below, the term “unique” is used in this context to cover a hash number that is generated from a hash function occasionally producing a common hash number for chunks that are non-identical because hash functions are not generally perfect at producing different numbers for different chunks.

A “Region” is a set of contiguous physical blocks on a storage medium (e.g., hard drive).

A “Zone” is composed of two or more Regions. The Regions that make up a Zone are generally not required to be contiguous. In an exemplary embodiment as described below, a Zone stores the equivalent of 1GB of data or control information.

5 A “Cluster” is the unit size within Zones and represents a unit of compression (discussed below). In an exemplary embodiment as described below, a Cluster is 4KB (i.e., eight 512-byte sectors) and essentially equates to a Chunk.

A “Redundant set” is a set of sectors/clusters that provides redundancy for a set of data.

10 “Backing up a Region” involves copying the contents of one Region to another Region.

A “first pair” and a “second pair” of storage devices may include a common storage device.

A “first plurality” and a “second plurality” of storage devices may include one or more common storage devices.

15 A “first arrangement” and a “second arrangement” or “different arrangement” of storage devices may include one or more common storage devices.

In embodiments of the present invention, a filesystem-aware storage system analyzes host filesystem data structures in order to determine storage usage of the host filesystem. For example, the block storage device may parse the host filesystem data
20 structures to determine such things as used blocks, unused blocks, and data types. The block storage device manages the physical storage based on the storage usage of the host filesystem.

Such a filesystem-aware block storage device can make intelligent decisions regarding the physical storage of data. For example, the filesystem-aware block storage
25 device can identify blocks that have been released by the host filesystem and reuse the released blocks in order to effectively extend the data storage capacity of the system. Such reuse of released blocks, which may be referred to hereinafter as “scavenging” or “garbage collection,” may be particularly useful in implementing virtual storage, where the host filesystem is configured with more storage than the actual physical storage
30 capacity. The filesystem-aware block storage device can also identify the data types of objects stored by the filesystem and store the objects using different storage schemes

based on the data types (e.g., frequently accessed data can be stored uncompressed and in sequential blocks, while infrequently accessed data can be stored compressed and/or in non-sequential blocks; different encoding schemes such as data compression and encryption can be applied to different objects based on the data types).

5 The filesystem-aware block storage device will generally support a predetermined set of filesystems for which it “understands” the inner workings sufficiently to locate and utilize the underlying data structures (e.g., free block bitmaps). In order to determine the filesystem type (e.g., NTFS, FAT, ReiserFS, ext3), the filesystem-aware block storage device typically parses a partition table to locate the operating system (OS) partition and
10 then parses the OS partition to locate the host filesystem’s superblock and thereby identify the filesystem type. Once the filesystem type is known, the filesystem-aware block storage device can parse the superblock to find the free block bitmaps for the host filesystem, and can then parse the free block bitmaps to identify used and unused blocks.

 In order to detect changes to the data structures (e.g., free block bitmaps) over
15 time, the filesystem-aware block storage device may periodically make a copy of the data structure (e.g., in a private, non-redundant zone) and later compare the currently active version of the data structure with the earlier-made copy to detect changes. For example, any bitmap entries transitioning from allocated to free can be identified, allowing the garbage collection operation to be accurately directed to clusters that are good candidates
20 for reclamation. As each bitmap cluster is processed, the historical copy can be replaced with the current copy to maintain a rolling history of bitmap operations. Over time the copy of the free block bitmap may become a patchwork of temporally disjoint clusters, but since the current copy is used to locate free entries, this should not cause any problems.

25 Exemplary embodiments are described hereinafter with reference to a storage array system.

 Fig. 1 is an illustration of an embodiment of the invention in which an object, in this example, a file, is parsed into a series of chunks for storage. Initially the file **11** is passed into the storage software where it is designated as an object **12** and allocated a
30 unique object identification number, in this case, #007. A new entry **131** is made into the object table **13** to represent the allocation for this new object. The object is now parsed

into “chunks” of data **121**, **122**, and **123**, which are fixed-length segments of the object. Each chunk is passed through a hashing algorithm, which returns a unique hash number for the chunk. This algorithm can later be applied to a retrieved chunk and the result compared with the original hash to ensure the retrieved chunk is the same as that stored.

5 The hash numbers for each chunk are stored in the object table **13** in the entry row for the object **132** so that later the complete object can be retrieved by collection of the chunks.

Also in Fig. 1, the chunk hashes are now compared with existing entries in the chunk table **14**. Any hash that matches an existing entry **141** is already stored and so no action is taken (i.e., the data is not stored again, leading to automatic compression of the
10 objects). A new hash (one which has no corresponding entry in the chunk table **14**) is entered into the chunk table **141**. The data in the chunk is then stored on the available storage devices **151**, **152**, and **153** in the most efficient manner for fault-tolerant storage. This approach may lead to the chunk data’s being stored, for example, in a mirrored fashion on a storage system comprised of one or two devices or parity striped on a system
15 with more than two storage devices. This data will be stored on the storage devices at physical locations **1511**, **1521**, and **1531**, and these locations and the number of locations will be stored in the chunk table columns **143** and **142** so that all physical parts of a chunk may later be located and retrieved.

Fig. 2 illustrates in the same embodiment how a pattern for fault-tolerant storage
20 for a chunk may be dynamically changed as a result of the addition of more storage. In particular, Fig. 2 shows how a chunk physically stored on the storage devices may be laid out in a new pattern once additional storage is added to the overall system. In Fig. 2(a), the storage system comprises two storage devices **221** and **222** and the chunk data is physically mirrored onto the two storage devices at locations **2211** and **2221** to provide
25 fault tolerance. In Figure 2(b) a third storage device **223** is added, and it becomes possible to store the chunk in a parity striped manner, a pattern which is more storage efficient than the mirrored pattern. The chunk is laid out in this new pattern in three physical locations **2311**, **2321**, and **2331**, taking a much lower proportion of the available storage. The chunk table **21** is updated to show the new layout is in three locations **212** and also
30 the new chunk physical locations **2311**, **2321**, and **2331** are recorded **213**.

Fig. 3 shows a mature storage system, in accordance with an embodiment of the present invention, which has been functioning for some time. This illustrates how chunks may be physically stored over time on storage devices with varying storage capacities. The figure shows a storage system comprised of a 40GB storage device **31**, an 80GB storage device **32** and a 120GB storage device **33**. Initially chunks are stored in a fault tolerant stripe pattern **34** until the 40GB storage device **31** became full. Then, due to lack of storage space, new data is stored in a mirrored pattern **36** on the available space on the 80GB **32** and the 120GB **33** storage devices. Once the 80GB storage device **32** is full, then new data is laid out in a single disk fault tolerant pattern **37**. Even though the storage devices comprise a single pool for storing data, the data itself, as stored in the chunks, has been stored in a variety of distinct patterns.

Fig. 4 illustrates another embodiment of the invention in which indicator states are used to warn of inefficient storage use and low levels of fault tolerance. In Fig. 4(a), all three storage devices **41**, **42**, and **45** have free space and the indicator light **44** is green to show data is being stored in an efficient and fault-tolerant manner. In Fig. 4 (b) the 40GB storage device **41** has become full, and thus new data can be stored only on the two storage devices **42** and **43** with remaining free space in a mirrored pattern **46**. In order to show the data is still fully redundant but not efficiently stored, the indicator light **44** has turned amber. In Fig. 4 (c), only the 120GB storage device **43** has free space remaining and so all new data can be stored only in a mirrored pattern on this one device **43**. Because the fault-tolerance is less robust and the system is critically short of space, the indicator light **44** turns red to indicate the addition of more storage is necessary.

In one alternative embodiment, an indicator is provided for each drive/slot in the array, for example, in the form of a three-color light (e.g., green, yellow, red). In one particular embodiment, the lights are used to light the whole front of a disk carrier with a glowing effect. The lights are controlled to indicate not only the overall status of the system, but also which drive/slot requires attention (if any). Each three-color light can be placed in at least four states, specifically off, green, yellow, red. The light for a particular slot may be placed in the off state if the slot is empty and the system is operating with sufficient storage and redundancy so that no drive need be installed in the slot. The light for a particular slot may be placed in the green state if the corresponding drive is

sufficient and need not be replaced. The light for a particular slot may be placed in the yellow state if system operation is degraded such that replacement of the corresponding drive with a larger drive is recommended. The light for a particular slot may be placed in the red state if the corresponding drive must be installed or replaced. Additional states
 5 could be indicated as needed or desired, for example, by flashing the light between on and off states or flashing the light between two different colors (e.g., flash between red and green after a drive has been replaced and re-layout of data is in progress). Additional details of an exemplary embodiment are described below.

Of course, other indication techniques can be used to indicate both system status
 10 and drive/slot status. For example, a single LCD display could be used to indicate system status and, if needed, a slot number that requires attention. Also, other types of indicators (e.g., a single status indicator for the system (e.g., green/yellow/red) along with either a slot indicator or a light for each slot) could be used.

Fig. 5 is a block diagram of functional modules used in the storage, retrieval and
 15 re-layout of data in accordance with an embodiment of the invention, such as discussed above in connections with Figs. 1-3. The entry and exit point for communication are the object interface **511** for passing objects to the system for storage or retrieving objects, the block interface **512**, which makes the storage system appear to be one large storage device, and the CIFS interface **513**, which makes the storage system appear to be a
 20 Windows file system. When these interfaces require the storage of data, the data is passed to the Chunk Parser **52**, which performs the break up of the data into chunks and creates an initial entry into the object table **512** (as discussed above in connection with Fig. 1). These chunks are then passed to the hash code generator **53**, which creates the associated hash codes for each chunk and enters these into the object table so the chunks associated
 25 with each object are listed **512** (as discussed above in connection with in Fig. 1). The chunk hash numbers are compared with the entries in the chunk table **531**. Where a match is found, the new chunk is discarded, as it will be identical to a chunk already stored in the storage system. If the chunk is new, a new entry for it is made in the chunk table **531**, and the hashed chunk is passed to the physical storage manager **54**. The physical storage
 30 manager stores the chunk in the most efficient pattern possible on the available storage devices **571**, **572**, and **573** and makes a corresponding entry in the chunk table **531** to

show where the physical storage of the chunk has occurred so that the contents of the chunk can be retrieved later 512 (as discussed above in connection with Fig. 1).

The retrieval of data in Fig. 5 by the object 511, block 512 or CIFS 513 interface is performed by a request to the retrieval manager 56, which consults the object table 521
 5 to determine which chunks comprise the object and then requests these chunks from the physical storage manager 54. The physical storage manager 54 consults the chunk table 531 to determine where the requested chunks are stored and then retrieves them and passes the completed data (object) back to the retrieval manager 56, which returns the data to the requesting interface. Also included in Fig. 5 is the fault tolerant manager
 10 (FTL) 55, which constantly scans the chunk table to determine if chunks are stored in the most efficient manner possible. (This may change as storage devices 571, 572, and 573 are added and removed.) If a chunk is not stored in the most efficient manner possible, then the FTL will request the physical storage manager 54 to create a new layout pattern for the chunk and update the chunk table 531. This way all data continues to remain
 15 stored in the most efficient manner possible for the number of storage devices comprising the array (as discussed above in connection with Figs. 2 and 3).

The following provides additional details of an exemplary embodiment of the present invention.

20 **Data Layout Scheme - Zones**

Among other things, a Zone has the effect of hiding redundancy and disk re-layout from the actual data being stored on the disk. Zones allow additional layout methods to be added and changed without affecting the user of the zone.

The storage array lays out data on the disk in virtual sections called Zones. A
 25 Zone stores a given and fixed amount of data (for example 1 G Bytes). A zone may reside on a single disk or span across one or more drives. The physical layout of a Zone provides redundancy in the form specified for that zone.

Fig. 6 shows an example in which mirroring is used in an array containing more than two drives. Fig. 7 shows some example zones using different layout schemes to
 30 store their data. The diagram assumes a zone stores 1GB of data. Note the following points:

- i) A zone that spans multiple drives does not necessarily use the same offset into the drive across the set.
- ii) A single drive mirror requires 2G of storage to store 1G of data
- iii) A dual drive mirror requires 2G of storage to store 1G of data.
- 5 iv) A 3 drive stripe requires 1.5G of storage to store 1G of data.
- v) A 4 drive stripe requires 1.33G of storage to store 1G of data.
- vi) Zone A, zone B etc. are arbitrary zone names. In a real implementation each zone would be identified by a unique number.
- vii) Although implied by the diagram, zones are not necessarily contiguous on a
10 disk (see regions later).
- viii) There is no technical reason why mirroring is restricted to 2 drives. For example, in a 3 drive system 1 copy of the data could be stored on 1 drive and half of the mirrored data could be stored on each of the other two drives. Likewise, data could be mirrored across three drives, with half the data on
15 each of two drives and half of the mirror on the other two drives.

Data Layout Scheme - Regions

Each disk is split into a set of equal-sized Regions. The size of a Region is much
20 smaller than a Zone and a Zone is constructed from one or more regions from one or more disks. For efficient use of disk space, the size of a Region is typically a common factor of the different Zone sizes and the different number of disks supported by the array. In an exemplary embodiment, Regions are 1/12 the data size of a Zone. The following table lists the number of Regions/Zone and the number of Regions/disk for
25 various layouts, in accordance with an exemplary embodiment of the invention.

Layout Method	Number of regions/zone	Number of regions/disk
1 drive mirror	24	24
2 drive mirror	24	12
3 drive stripe	18	6
4 drive stripe	16	4

Individual Regions can be marked as used, free or bad. When a Zone is created, a set of free Regions from the appropriate disks are selected and logged in a table. These Regions can be in any arbitrary order and need not be contiguous on the disk. When data is written to or read from a Zone, the access is redirected to the appropriate Region. Among other things, this allows data re-layout to occur in a flexible and efficient manner. Over time, the different sized Zones will likely cause fragmentation to occur, leaving many disk areas too small to hold a complete Zone. By using the appropriate Region size, all gaps left by fragmentation will be at least one Region in size, allowing easy reuse of these small gaps without having to de-fragment the whole disk.

Data Layout Scheme - Re-Layout

In order to facilitate implementation, a fixed sequence of expansion and contraction may be enforced. For example, if two drives are suddenly added, the expansion of a zone may go through an intermediate expansion as though one drive was added before a second expansion is performed to incorporate the second drive. Alternatively, expansion and contraction involving multiple drives may be handled atomically, without an intermediate step.

Before any re-layout occurs, the required space must be available. This should be calculated before starting the re-layout to ensure that unnecessary re-layout does not occur.

Data Layout Scheme - Drive Expansion

The following describes the general process of expanding from single drive mirroring to dual drive mirroring in accordance with an exemplary embodiment of the invention:

- i) Assuming single drive mirror has data 'A' and mirror 'B'
- ii) Allocate 12 regions on drive to expand zone on to 'C'
- iii) Copy mirror 'B' to region set 'C'

- iv) Any writes made to data already copied must be mirrored to the appropriate place in 'C'
- v) When copy is complete, update zone table with new layout type and replace pointers to 'B' with pointers to 'C'
- 5 vi) Mark the regions that make-up 'B' as free.

The following describes the general process of expanding from dual drive mirroring to triple drive striping with parity in accordance with an exemplary embodiment of the invention:

- 10 i) Assume one drive has data 'A' and a second drive has mirror 'B'
- ii) Allocate 6 regions on third drive for parity information 'C'
- iii) Calculate parity information using first 6 regions of 'A' and the second 6 regions of 'B'
- iv) Place parity information in 'C'
- 15 v) Any writes made to data already processed must be parity'd to the appropriate place in 'C'
- vi) When copy is complete, update zone table with new layout type point table to first half of 'A', second half of 'B' and 'C'
- vii) Mark second half of 'A' and first half of 'B' as free.

20

The following describes the general process of expanding from triple drive striping to quad drive striping with parity in accordance with an exemplary embodiment of the invention:

- i) Assume one drive has data 'A', a second drive has data 'B' and a third has parity 'P'
- 25 ii) Allocate four regions on a fourth drive for strip data 'C'
- iii) Copy last two regions of 'A' to the first two regions of 'C'
- iv) Copy first two regions of 'B' to last two regions of 'C'
- v) Allocate four regions on parity drive 'D'
- 30 vi) Calculate parity information using first four regions of A, C and the last four regions of B

- vii) Place parity information in 'D'
- viii) Any writes made to data already processed must be parity'd to the appropriate place in 'D'
- ix) Update zone table with new layout type and point table to first four regions of 'A', 'C', second four regions of 'B' and 'D'
- x) Mark last two regions of 'A' and first two regions of 'B' as free.

Data Layout Scheme - Drive Contraction

Drive contraction takes place when a disk is either removed or fails. In such a case, the array contracts the data to get all zones back into a redundant state if possible. Drive contraction is slightly more complex than expansion as there are more choices to make. However, transitions between layout methods happen in a similar way to expansion, but in reverse. Keeping the amount of data to be reproduced to a minimum allows redundancy to be achieved as quickly as possible. Drive contraction generally precedes one zone at a time while space is available until all zones are re-layed out. Generally speaking, only data which resides on the removed or failed disk will be rebuilt.

Choosing how to contract

The following table describes a decision tree for each zone that needs to be re-laid out, in accordance with an exemplary embodiment of the present invention:

Zone type missing data	Condition	Action
Any	No Space available for zone re-layout	Leave zone in degraded state until new disk added or removed disk is replaced.
Single drive mirror	Data inconsistent	Lock down system and wait for reset or for the missing drive to be replaced
Dual Drive	1 disk left in system	Convert to single drive mirror

Mirror	Space only available on drive that contains remaining data	
	2 or 3 disks left in system with space is available	Reconstruct mirror on another drive
3 Drive Striping	2 disks left in system with space available	Convert to 2 drive mirroring
	3 disks left in system with space available	Reconstruct missing stripe segment on the third drive
4 Drive Striping	3 disks left in system with space available	Convert to 3 drive striping

The following describes the general process of contracting from dual drive mirroring to single drive mirroring in accordance with an exemplary embodiment of the invention:

- 5 i) Assuming single drive mirror has data 'A' and missing mirror 'B' or visa versa
- ii) Allocate 12 regions on the drive that contains 'A' as 'C'
- iii) Copy data 'A' to region set 'C'
- iv) Any writes made to data already copied must be mirrored to the appropriate
- 10 v) When copy is complete, update zone table with new layout type and replace
- pointers to 'B' with pointers to 'C'

The following describes the general process of contracting from triple drive stripe to dual drive mirror (missing parity) in accordance with an exemplary embodiment of the invention:

- i) Assuming that stripe consists of data blocks 'A', 'B' and 'C' on different drives. Parity 'C' is missing.
- ii) Define 'A' as containing the first half of the zone and 'B' as the second half of
- 20 the zone.
- iii) Allocate 6 regions 'D' on the 'A' drive and 6 regions 'E' on the 'B' drive

- iv) Copy 'A' to 'E'.
- v) Copy 'B' to 'D'
- vi) Any writes made to data already copied must be mirrored to the appropriate place in 'D' and 'E'
- 5 vii) When copy is complete, update zone table with new layout type and set pointers to 'A'/'D' and 'E'/'B'

The following describes the general process of contracting from triple drive stripe to dual drive mirror (missing data) in accordance with an exemplary embodiment of the invention:

- i) Assuming that stripe consists of data blocks 'A', 'B' and 'C' on different drives. Data 'C' is missing.
- ii) Define 'A' as containing the first half of the zone and 'C' as the second half of the zone.
- 15 iii) Allocate 6 regions 'D' on the 'A' drive and 12 regions 'E' on the 'B' drive
- iv) Copy 'A' to the first half of 'E'
- v) Reconstruct missing data from 'A' and 'B'. Write data to 'D'
- vi) Copy 'D' to second half of 'E'.
- vii) Any writes made to data already copied must be mirrored to the appropriate place in 'D' and 'E'
- 20 viii) When copy is complete, update zone table with new layout type and set pointers to 'A'/'D' and 'E'
- ix) Mark 'B' regions as free.

The following describes the general process of contracting from quad drive stripe to triple drive stripe (missing parity) in accordance with an exemplary embodiment of the invention:

- i) Assuming that stripe consists of data blocks 'A', 'B', 'C' and 'D' on different drives. Parity 'D' is missing.
- 30 ii) Define 'A' as containing the first third, 'B' as the second third and 'C' as the third third of the zone

- iii) Allocate 2 regions 'G' on the 'A' drive, 2 regions 'E' on the 'C' drive and 6 regions 'F' on the 'B' drive.
- iv) Copy first half of 'B' to 'G'
- v) Copy second half of 'B' to 'E'
- 5 vi) Construct parity from 'A'/'G' and 'E'/'C' and write to 'F'
- vii) Any writes made to data already copied must be mirrored to the appropriate place in 'G', 'E' and 'F'
- viii) When copy is complete, update zone table with new layout type and set pointers to 'A'/'G', 'E'/'C' and 'F'
- 10 ix) Mark 'B' regions as free.

The following describes the general process of contracting from quad drive stripe to triple drive stripe (first 1/3 missing) in accordance with an exemplary embodiment of the invention:

- 15 i) Assuming that stripe consists of data blocks 'A', 'B', 'C' and 'D' on different drives. Data 'A' is missing.
- ii) Define 'A' as containing the 1st third, 'B' as the 2nd third and 'C' as the 3rd third of the zone and 'D' as the parity.
- iii) Allocate 4 regions 'E' on the 'B' drive, 2 regions 'F' on the 'C' drive and 6 regions 'G' on the 'D' drive.
- 20 iv) Copy second half of 'B' to 'F'
- v) Construct missing data from 'B', 'C' and 'D' and write to 'E'
- vi) Construct new parity from 'E'/'1st half B' and 'F'/'C' and write to 'G'
- vii) Any writes made to data already copied must be mirrored to the appropriate place in 'B', 'E', 'F' and 'G'
- 25 viii) When copy is complete, update zone table with new layout type and set pointers to 'E'/'1st half B' and 'F'/'C' and 'G'
- ix) Mark 2nd half 'B' and 'D' regions as free.

The following describes the general process of contracting from quad drive stripe to triple drive stripe (second 1/3 missing) in accordance with an exemplary embodiment of the invention:

- 5 i) Assuming that stripe consists of data blocks 'A', 'B', 'C' and 'D' on different drives. Data 'B' is missing.
- ii) Define 'A' as containing the 1st third, 'B' as the 2nd third and 'C' as the 3rd third of the zone and 'D' as the parity.
- iii) Allocate 2 regions 'E' on the 'A' drive, 2 regions 'F' on the 'C' drive and 6 regions 'G' on the 'D' drive.
- 10 iv) Construct missing data from 1st half 'A', 1st half 'C' and 1st half 'D' and write to 'E'
- v) Construct missing data from 2nd half 'A', 2nd half 'C' and 2nd half 'D' and write to 'F'
- vi) Construct new parity from 'A'/'E' and 'F'/'C' and write to 'G'
- 15 vii) Any writes made to data already copied must be mirrored to the appropriate place in 'E', 'F' and 'G'
- viii) When copy is complete, update zone table with new layout type and set pointers to 'E', 'F' and 'G'
- ix) Mark 'D' regions as free.

20

The following describes the general process of contracting from quad drive stripe to triple drive stripe (third 1/3 missing) in accordance with an exemplary embodiment of the invention:

- 25 i) Assuming that stripe consists of data blocks 'A', 'B', 'C' and 'D' on different drives. Data 'C' is missing.
- ii) Define 'A' as containing the 1st third, 'B' as the 2nd third and 'C' as the 3rd third of the zone and 'D' as the parity.
- iii) Allocate 2 regions 'E' on the 'A' drive, 4 regions 'F' on the 'B' drive and 6 regions 'G' on the 'D' drive.
- 30 iv) Copy 1st half 'B' to 'E'
- v) Construct missing data from 'A', 'B' and 'D' and write to 'F'

- vi) Construct new parity from 'A'/'E' and 2nd half 'B'/'F' and write to 'G'
- vii) Any writes made to data already copied must be mirrored to the appropriate place in 'E', 'F' and 'G'
- viii) When copy is complete, update zone table with new layout type and set pointers to 'A'/'E' and 2nd half 'B'/'F' and 'G'
- ix) Mark 1st half 'B' and 'D' regions as free.

For example, with reference again to Fig. 3, dual drive mirror (Zone B) could be reconstructed on Drive 2 if either Drive 0 or Drive 1 is lost, provided there is sufficient space available on Drive 2. Similarly, three drive stripe (Zone C) could be reconstructed utilizing Drive 3 if any of Drives 0-2 are lost, provided there is sufficient space available on Drive 3.

15 **Data Layout Scheme - Zone Reconstruction**

Zone reconstruction occurs when a drive has been removed and there is enough space on the remaining drives for ideal zone re-layout or the drive has been replaced with a new drive of larger size.

The following describes the general process of dual drive mirror reconstruction in accordance with an exemplary embodiment of the invention:

- i) Assuming single drive mirror has data 'A' and missing mirror 'B'
- ii) Allocate 12 regions 'C' on a drive other than that containing 'A'
- iii) Copy data 'A' to 'C'
- iv) Any writes made to data already copied must be mirrored to the appropriate place in 'C'
- v) When copy is complete, update zone table pointers to 'B' with pointers to 'C'

The following describes the general process of three drive stripe reconstruction in accordance with an exemplary embodiment of the invention:

- i) Assume one drive has data 'A', a second drive has data 'B' and a third has parity 'P'. 'B' is missing. Note it doesn't matter which piece is missing, the required action is the same in all cases.
- ii) Allocate 6 regions 'D' on a drive other than that containing 'A' and 'P'
- 5 iii) Construct missing data from 'A' and 'P'. Write data to 'D'
- iv) Any writes made to data already processed must be parity'd to the appropriate place in 'D'
- v) Update zone table by replacing pointers to 'B' with pointers to 'D'

10 In this exemplary embodiment, four -drive-reconstruction can only occur if the removed drive is replaced by another drive. The reconstruction consists of allocating six regions on the new drive and reconstructing the missing data from the other three region sets.

15 **Data Layout Scheme - The temporarily missing drive problem**

When a drive is removed and there is no room for re-layout, the array will continue to operate in degraded mode until either the old drive is plugged back in or the drive is replaced with a new one. If a new one is plugged in, then the drive set should be rebuilt. In this case, data will be re-laid out. If the old disk is placed back into the array, it will no longer be part of the current disk set and will be regarded as a new disk. However, if a new disk is not placed in the array and the old one is put back in, the old one will still be recognized as being a member of the disk set, albeit an out of date member. In this case, any zones that have already been re-laid out will keep their new configuration and the regions on the old disk will be freed. Any zone that has not been re-laid out will still be pointing at the appropriate regions of the old disk. However, as some writes may have been performed to the degraded zones, these zones need to be refreshed. Rather than logging every write that has occurred, degraded regions that have been modified may be marked. In this way, when the disk is replaced, only the regions that have been modified need to be refreshed.

Furthermore, zones that have been written to may be placed further up the priority list for re-layout. This should reduce the number of regions that need to be refreshed should the disk be replaced. A timeout may also be used, after which point the disk, even if replaced, will be wiped. However, this timeout could be quite large, possibly hours rather than minutes.

Data Layout Scheme - Data Integrity

As discussed above, one problem with a standard RAID system is that it is possible for disc-surface corruption to occur on an infrequently used area of the disk array. In the event that another drive fails, it is not always possible to determine that corruption has occurred. In such a case, the corrupted data may be propagated and preserved when the RAID array rebuilds the failed drive.

The hash mechanism discussed above provides an additional mechanism for data corruption detection over that which is available under RAID. As is mentioned elsewhere, when a chunk is stored, a hash value is computed for the chunk and stored. Any time the chunk is read, a hash value for the retrieved chunk can be computed and compared with the stored hash value. If the hash values do not match (indicating a corrupted chunk), then chunk data can be recovered from redundant data.

In order to minimize the time window in which data corruption on the disk can occur, a regular scan of the disks will be performed to find and correct corrupted data as soon as possible. It will also, optionally, allow a check to be performed on reads from the array.

Data Layout Scheme - Volume

In a sparse volume, regardless of the amount of storage space available on discs in the array, the array always claims to be a fixed size – for example, M Terabytes. Assume that the array contains S bytes of actual storage space, where $S \leq M$, and that data can be requested to be stored at locations L1, L2, L3, etc. within the M Terabyte space. If the requested location $L_n > S$, then the data for L_n must be stored at a location $P_n < S$. This

is managed by including a lookup table to index P_n based on L_n, as shown in Fig. 8. The feature is allows the array to work with operating systems that do not support volume expansion, such as Windows, Linux, and Apple operating systems. In addition, the array can provide multiple Sparse Volumes which all share the same physical storage. Each
 5 sparse volume will have a dedicated lookup table, but will share the same physical space for data storage.

Drive slot indicators

10 As discussed above, the storage array consists of one or more drive slots. Each drive slot can either be empty or contain a hard disk drive. Each drive slot has a dedicated indicator capable of indicating four states: Off, OK, Degraded and Fail. The states are interpreted generally as follows:

Indicator State	Meaning for Array User
Off	Drive slot is empty and is available for an additional drive to be inserted.
OK	Drive in slot is functioning correctly.
Degraded	Action by user recommend: if slot is empty, add a drive to this slot; if slot contains a drive, replace drive with another, higher-capacity drive.
Fail	Action by user required ASAP: if slot is empty, add a drive to this slot; if slot contains a drive, replace drive with another, higher-capacity drive.

15 In this exemplary embodiment, red/amber/green light emitting diodes (LEDs) are used as the indicators. The LEDs are interpreted generally as follows:

LED state	Indicator State	Example circumstances under which state may occur	Figures
Off	Off	Slot is empty. Array has available space.	9,10,12
Green	OK	Drive is functioning correctly, array data is redundant and array has available disk space.	9,10,11,12
Amber	Degraded	Array is approaching a Fail condition; Not enough space to maintain redundant data in the event of a disc failure.	11
Red	Fail	Disk in this slot has failed and must be replaced; the array does not have enough space to maintain redundant data storage and more space must be added.	10,12

Fig. 9 shows an exemplary array having available storage space and operating in a fault-tolerant manner, in accordance with an exemplary embodiment of the present invention. Slots B, C, and D are populated with storage devices, and there is sufficient storage space available to store additional data redundantly. The indicators for slots B, C, and D are green (indicating that these storage devices are operating correctly, the array data is redundant, and the array has available disk space), and the indicator for slot A is off (indicating that no storage device needs to be populated in slot A).

Fig. 10 shows an exemplary array that does not have enough space to maintain redundant data storage and more space must be added, in accordance with an exemplary embodiment of the present invention. Slots B, C, and D are populated with storage devices. The storage devices in slots C and D are full. The indicators for slots B, C, and D are green (indicating that these storage devices are operating correctly), and the

indicator for slot A is red (indicating that the array does not have enough space to maintain redundant data storage and a storage device should be populated in slot A).

Fig. 11 shows an exemplary array that would be unable to maintain redundant data in the event of a failure, in accordance with an exemplary embodiment of the present invention. Slots A, B, C, and D are populated with storage devices. The storage devices in slots C and D are full. The indicators for slots A, B, and C are green (indicating that they are operating correctly), and the indicator for slot D is amber (indicating that the storage device in slot D should be replaced with a storage device having greater storage capacity).

Fig. 12 shows an exemplary array in which a storage device has failed, in accordance with an exemplary embodiment of the present invention. Slots B, C, and D are populated with storage devices. The storage device in slot C has failed. The indicators for slots B and D are green (indicating that they are operating correctly), the indicator for slot C is red (indicating that the storage device in slot C should be replaced), and the indicator for slot A is off (indicating that no storage device needs to be populated in slot A).

The following is a description of the software design for an exemplary embodiment of the present invention. The software design is based on six software layers, which span the logical architecture from physically accessing the disks to communicating with the host computing system.

In this exemplary embodiment, a file system resides on a host server, such as a Windows, Linux, or Apple server, and accesses the storage array as a USB or iSCSI device. Physical disk requests arriving over the host interface are processed by the Host Request Manager (HRM). A Host I/O interface coordinates the presentation of a host USB or iSCSI interface to the host, and interfaces with the HRM. The HRM coordinates data read/write requests from the host I/O interface, dispatches read and write requests, and co-ordinates the retiring of these requests back to the host as they are completed.

An overarching aim of the storage array is to ensure that once data is accepted by the system, it is stored in a reliable fashion, making use of the maximum amount of redundancy the system currently stores. As the array changes physical configuration,

data is re-organized so as to maintain (and possibly maximize) redundancy. In addition, simple hash based compression is used to reduce the amount of storage used.

The most basic layer consists of disk drivers to store data on different disks.

5 Disks may be attached via various interfaces, such as ATA tunneled over a USB interface.

Sectors on the disks are organized into regions, zones, and clusters, each of which has a different logical role.

Regions represent a set of contiguous physical blocks on a disk. On a four drive system, each region is 1/12 GB in size, and represents minimal unit of redundancy. If a
10 sector in a region is found to be physically damaged, the whole region will be abandoned.

Zones represent units of redundancy. A zone will consist of a number of regions, possibly on different disks to provide the appropriate amount of redundancy. Zones will provide 1GB of data capacity, but may require more regions in order to provide the redundancy. 1GB with no redundancy require one set of 12 regions (1GB); a 1GB
15 mirrored zone will require 2 sets of 1GB regions (24 regions); a 1GB 3-disk striped zone will require 3 sets of 0.5GB regions (18 regions). Different zones will have different redundancy characteristics.

Clusters represent the basic unit of compression, and are the unit size within zones. They are currently 4KB: 8 x 512 byte sectors in size. Many clusters on a disk will
20 likely contain the same data. A cluster access table (CAT) is used to track the usage of clusters via a hashing function. The CAT translates between logical host address and the location of the appropriate cluster in the zone.

When writing to disk, a hash function is used to see if the data is already present on the disk. If so, the appropriate entry in the CAT table is set to point to the existing
25 cluster.

The CAT table resides in its own zone. If it exceeds the size of the zone, an additional zone will be used, and a table will be used to map logical address to the zone for that part of the CAT. Alternatively, zones are pre-allocated to contain the CAT table.

In order to reduce host write latency and to ensure data reliability, a journal
30 manager will record all write requests (either to disk, or to NVRAM). If the system is rebooted, journal entries will be committed on reboot.

Disks may come and go, or regions may be retired if they are found to have corruption. In either of these situations, a layout manager will be able to re-organize regions within a zone in order to change its redundancy type, or change the regional composition of a zone (should a region be corrupted).

5 As the storage array provides a virtual disk array, backed by changing levels of physical disk space, and because it presents a block level interface, it is not obvious when clusters are no longer in use by the file system. As a result, the cluster space used will continue to expand. A garbage collector (either located on the host or in firmware) will analyze the file system to determine which clusters have been freed, and remove them
10 from the hash table.

The following table shows the six software layers in accordance with this exemplary embodiment of the invention:

Layer 5 : Garbage collector, Host Interface (USB/iSCSI)
Layer 4 : Host request manager
Layer 3: CAT, HASH, Journal manager
Layer 2 : Zones manager. Allocates/frees chunks of sectors called Zones. Knows about SDM, DDM, SD3 etc in order to deal with errors and error recovery. Layout Manager
Layer 1 : Read/write physical clusters/sectors. Allocates Regions per disk
Layer 0: Disk access drivers

15 Fig. 13 shows a module hierarchy representing the different software layers and how they relate to one another. Software layering is preferably rigid in order to present clear APIs and delineation.

The Garbage Collector frees up clusters which are no longer used by the host file system. For example, when a file is deleted, the clusters that were used to contain the file
20 are preferably freed.

The Journal Manager provides a form of journaling of writes so that pending writes are not lost in the case of a power failure or other error condition.

The Layout Manager provides run-time re-layout of the Zones vis-à-vis their Regions. This may occur as a result of disk insertion/removal or failure.

The Cluster Manager allocates clusters within the set of data Zones. The Disk Utilization Daemon checks for free disk space on a periodic basis.

The Lock Table deals with read after write collision issues.

The Host Request Manager deals with the read/write requests from the Host and
5 Garbage Collector. Writes are passed to the Journal Manager, whereas Reads are processed via the Cluster Access Table (CAT) Management layer.

As discussed above, in typical file systems, some amount of the data will generally be repetitive in nature. In order to reduce disk space utilization, multiple copies of this data are not written out to the disks. Instead, one instance is written, and all other
10 instances of the same data are referenced to this one instance.

In this exemplary embodiment, the system operates on a cluster of data at any time (e.g., 8 physical sectors), and this is the unit that is hashed. The SHA1 algorithm is used to generate a 160-bit hash. This has a number of benefits, including good uniqueness, and being supported on-chip in a number of processors. All 160-bits will be
15 stored in the hash record, but only the least significant 16-bits will be used as an index into a hash table. Other instances matching the lowest 16-bits will be chained via a linked-list.

In this exemplary embodiment, only one read/write operation may occur at a time. For performance purposes, hash analysis is not permitted to happen when writing a
20 cluster to disk. Instead, hash analysis will occur as a background activity by the hash manager.

Write requests are read from the journal's write queue, and are processed to completion. In order to ensure data consistency, writes must be delayed if a write operation is already active on the cluster. Operations on other clusters may proceed un-
25 impeded.

Unless a whole cluster is being written, the data being written will need to be merged with the existing data stored in the cluster. Based on the logical sector address (LSA), the CAT entry for the cluster is located. The hash key, zone and cluster offset information is obtained from this record, which can then be used to search the hash table
30 to find a match. This is the cluster.

It might well be necessary to doubly hash the hash table; once via the SHA1 digest, and then by the zone/cluster offset to improve the speed of lookup of the correct hash entry. If the hash record has already been used, the reference count is decremented. If the reference count is now zero, and there is no snapshot referenced by the hash entry,
 5 the hash entry and cluster can be freed back to their respective free lists.

The original cluster data is now merged with the update section of the cluster, and the data is re-hashed. A new cluster is taken off the free-list, the merged data is written to the cluster, new entry is added to the hash table, and the entry in the CAT table is updated to point to the new cluster.

10 As a result of updating the hash table, the entry is also added to an internal queue to be processed by a background task. This task will compare the newly added cluster and hash entry with other hash entries that match the hash table row address, and will combine records if they are duplicates, freeing up hash entries and CAT table entries as appropriate. This ensures that write latency is not burdened by this activity. If a failure
 15 (e.g., a loss of power) occurs during this processing, the various tables can be deleted, with a resulting loss of data. The tables should be managed in such a way that the final commit is atomic or the journal entry can be re-run if it did not complete fully.

The following is pseudocode for the write logic:

```

20 While (stuff to do)
    writeRecord = journalMgr.read();
    lsa = writeRecord.GetLsa();
    catEntry = catMgr.GetCATEntry(lsa);
    if (catMgr.writeInProgress(catEntry)) delay();
25
    originalCluster = catMgr.readCluster(catEntry);
    originalHash = hashMgr.calcHash(originalCluster);
    hashRecord = hashMgr.Lookup(originalHash, zone, offset);

30 if ((hashRecord.RefCount == 1) && (hashRecord.snapshot == 0))
    hashRecord.free();
  
```

```

        originalCluster.free();
        // Note there are some optimizations here where we can reuse
        // this cluster without having to free & re-allocate it.

5    // otherwise, still users of this cluster, so update & leave it alone

        hashRecord.RefCount--;
        hashRecord.Update(hashRecord);

10   // Now add new record

        mergedCluster = mergeCluster(originalCluster, newCluster);

        newHash = hashMgr.calcHash(mergedCluster);
15
        newCluster = clusterMgr.AllocateCluster(zone, offset);
        clusterMgr.write(cluster, mergedCluster);
        zoneMgr.write(cluster, mergedCluster);
        ...
20   hashMgr.addHash(newHash, newCluster, zone, offset)
        (internal: queue new hash for background processing)
        catMgr.Update(lba, zone, offset, newHash);

        // We've completed the journal entry successfully. Move on to the next one.
25   journalMgr.next();

```

Read requests are also processed one cluster (as opposed to "sector") at a time. Read requests do not go through the hash-related processing outlined above. Instead, the

30 host logical sector address is used to reference the CAT and obtain a Zone number and cluster offset into the Zone. Read requests should look up the CAT table entry in the

CAT Cache, and must be delayed in the write-in-progress bit is set. Other reads/writes may proceed un-impeded. In order to improve data integrity checking, when a cluster is read, it will be hashed, and the hash compared with the SHA1 hash value stored in the hash record. This will require using the hash, zone and cluster offset as a search key into the hash table.

Clusters are allocated to use as few Zones as possible. This is because Zones correspond directly to disk drive usage. For every Zone, there are two or more Regions on the hard drive array. By minimizing the number of Zones, the number of physical Regions is minimized and hence the consumption of space on the hard drive array is reduced.

The Cluster Manager allocates cluster from the set of Data Zones. A linked list is used to keep track of free clusters in a Zone. However, the free cluster information is stored as a bit map (32KB per Zone) on disk. The linked list is constructed dynamically from the bitmap. Initially, a linked list of a certain number of free clusters is created in memory. When clusters are allocated, the list shrinks. At a predetermined low-water mark, new linked list nodes representing free clusters are extracted from the bitmap on disk. In this way, the bitmap does not need to be parsed in order to find a free cluster for allocation.

In this exemplary embodiment, the hash table is a 64K table of records (indexed by the lower 16 bits of the hash) and has the following format:

Offset	Size in bits	Name	Value/Valid Range	Description
0	160	sha1Hash		The complete SHA1 hash digest
	16	refCount		Number of instances of this hash; what do we do if we get beyond 16 bits?
	18	Cluster offset		Cluster offset within zone
	14	Zone #		Zone# containing this cluster
	8	snapshot		One bit per snapshot instance to indicate that this cluster entry is used by that snapshot. This

				model supports 8 snapshots (possible only 7)
--	--	--	--	---

A cluster of all zeros may be fairly common, so the all-zeros case may be treated as a special case, for example, such that it can never be deleted (so wrapping the count
5 would not be a problem).

A linked list of free hash record is used when the multiple hashes have the same least significant hash, or when two hash entries point to different data clusters. In either case, a free hash record will be taken from the list, and linked via the pNextHash pointer.

The hash manager will tidy up entries added to the hash table, and will combine
10 identical clusters on the disk. As new hash records are added to the hash table, a message will be posted to the hash manager. This will be done automatically by the hash manager. As a background activity, the hash manager will process entries on its queue. It will compare the full hash value to see if it matches any existing hash records. If it does, it will also compare the complete cluster data. If the clusters match, the new hash record
15 can be discarded back to the free queue, the hash record count will be incremented, and the duplicate cluster will be returned to the cluster free queue. The hash manager must take care to propagate the snapshot bit forward when combining records.

A Cluster Access Table (CAT) contains indirect pointers. The pointers point to data clusters (with 0 being the first data cluster) within Zones. One CAT entry references
20 a single data cluster (tentatively 4KB in size). CATs are used (in conjunction with hashing) in order to reduce the disk usage requirements when there is a lot of repetitive data. A single CAT always represents a contiguous block of storage. CATs are contained within non-data Zones. Each CAT entry is 48-bits. The following table shows how each entry is laid out (assuming each data Zone contains 1GB of data):

25

Bits 0-17	Bits 18-31	Bits 32-47	Bits 48-63[..]
Offset of data cluster within Zone	Zone# containing data	Hash key	Reserved. Candidates include garbage

			collector write-bit; snapshot bits; snapshot table hash key
--	--	--	---

It is desirable for the CAT to fit into 64 bits, but this is not a necessity. The CAT table for a 2TB array is currently ~ 4GB in size. Each CAT entry points to a Zone which contains the data and the number of the Zone.

5 Fig. 14 shows how the CAT is used to access a data clusters in a Zone. Redundant data is referenced by more than one entry in the CAT. Two logical clusters contain the same data, so their CAT entries are pointed to the same physical cluster.

The Hash Key entry contains the 16-bit extract of the 160-bit SHA1 hash value of the entire cluster. This entry is used to update the hash table during a write operation.

10 There are enough bits in each entry in the CAT to reference 16TB of data. However, if every data cluster is different from another (in terms of contents), then just over 3 Zones' worth of CAT entries are required to reference 2 TB of data (each zone is 1GB in size, and hence can store 1GB/size of CAT entry entries. Assuming 6 byte CAT entries, that is 178956970 CAT entries/zone, i.e. the table references around 682 GB/zone

15 if each cluster is 4K).

A Host Logical Sector Translation Table is used to translate a Host Logical Sector Address into a Zone number. The portion of the CAT that corresponds to the Host Logical Sector Address will reside in this zone. Note that each CAT entry represents a cluster size of 4096 bytes. This is eight 512 byte sectors. The following shows a

20 representation of the host logical sector translation table:

Start Host Logical Sector Address	End Host Logical Sector Address	Zone # of CAT
0 (cluster #0)	1431655759 (cluster #178956969)	
1431655760 (cluster #178956970)	...	

Zones can be pre-allocated to hold the entire CAT. Alternatively, Zones can be allocated for the CAT as more entries to the CAT are required. Since the CAT maps the

2TB virtual disk to the host sector address space, it is likely that a large part of the CAT will be referenced during hard disk partitioning or formatting by the host. Because of this, the Zones may be pre-allocated.

5 The CAT is a large 1GB/zone table. The working set of clusters being used will be a sparse set from this large table. For performance reasons, active entries (probably temporally) may be cached in processor memory rather than always reading them from the disk. There are at least two options for populating the cache – individual entries from the CAT, or whole clusters from the CAT.

10 Because the write-in-progress is combined with the CAT cache table, it is necessary to ensure that all outstanding writes remain in the cache. Therefore, the cache needs to be at least as large as the maximum number of outstanding write requests.

Entries in the cache will be a cluster size (ie. 4K). There is a need to know whether there is a write-in-progress in operation on a cluster. This indication can be stored as a flag in the cache entry for the cluster. The following table shows the format
15 of a CAT cache entry:

Bits 0-17	Bits 18-31	Bits 32-47	Bit 48-63
Offset of data cluster within Zone	Zone# containing data	Hash key	Bit 48: Write-in-progress Bit 49: Dirty

20 The write-in-progress flag in the cache entry has two implications. First, it indicates that a write is in progress, and any reads (or additional writes) on this cluster must be held off until the write has completed. Secondly, this entry in the cache must not be flushed while the bit is set. This is partly to protect the state of the bit, and also to reflect the fact that this cluster is currently in use. In addition, this means that the size of the cache must be at least as large as the number of outstanding write operations.

25 One advantage of storing the write-in-progress indicator in the cache entry for the cluster is that it reflects the fact that the operation is current, saves having another table, and it saves an additional hashed-based lookup, or table walk to check this bit too. The

cache can be a write-delayed cache. It is only necessary to write a cache entry back to disk when the write operation has completed, although it might be beneficial to have it written back earlier. A hash function or other mechanism could be used to increase the number of outstanding write entries that can be hashed.

- 5 An alternate approach is to cache whole clusters of the CAT (i.e., 4K entry of entries). This would generally help performance if there is good spatial locality of access. Care needs to be taken because CAT entries are 48 bits wide, so there will not be a whole number of entries in the cache. The following table shows an example of a clustered CAT cache entry:

10

2 words	2 words	2 words	2 words
CAT entry 1 (partial entry of last 2 words)	CAT Entry 2		
CAT Entry 3			CAT Entry 4
CAT entry 4		CAT entry 5	
CAT Entry 5		CAT Entry 6	
....			
CAT Entry 682			CAT Entry 683 (partial entry of first 2 words)
Write-in-progress bit array [682 bits] : bits 0-255			
Write-in-progress bit array bits 256-511			
Write-in-progress bit array 512-682 + spare bits			Dirty count
			Reserved

The table size would be $4096 + 96$ (4192 bytes). Assuming it is necessary to have a cache size of 250 entries, the cache would occupy approximately 1MB.

It is possible to calculate whether the first and last entry is incomplete or not by appropriate masking of the logical CAT entry address. The caching lookup routine should do this prior to loading an entry and should load the required CAT cluster.

When the host sends a sector (or cluster) read request, it sends over the logical
5 sector address. The logical sector address is used as an offset into the CAT in order to obtain the offset of the cluster in the Zone that contains the actual data that is requested by the host. The result is a Zone number and an offset into that Zone. That information is passed to the Layer 2 software, which then extracts the raw cluster(s) from the drive(s).

In order to deal with clusters that have never been written to by the host, all CAT
10 entries are initialized to point to a "Default" cluster which contain all zeros.

The journal manager is a bi-level write journaling system. An aim of the system is to ensure that write requests can be accepted from the host and quickly indicate back to the host that the data has been accepted while ensuring its integrity. In addition, the system needs to ensure that there will be no corruption or loss of any block level data or
15 system metadata (e.g., CAT and Hash table entries) in the event of a system reset during any disk write.

The J1 journal manager caches all write requests from the hosts to disk as quickly as possible. Once the write has successfully completed (i.e., the data has been accepted by the array), the host can be signaled to indicate that the operation has completed. The
20 journal entry allows recovery of write requests when recovering from a failure. Journal records consist of the data to be written to disk, and the meta-data associated with the write transaction.

In order to reduce disk read/writes, the data associated with the write will be written to free clusters. This will automatically mirror the data. Free clusters will be taken
25 from the free cluster list. Once the data is written, the free cluster list must be written back to disk.

A journal record will be written to a journal queue on a non-mirrored zone. Each record will be a sector in size, and aligned to a sector boundary in order to reduce the risk that a failure during a journal write would corrupt a previous journal entry. Journal
30 entries will contain a unique, incrementing sequence count at the end of the record so that the end of a queue can easily be identified.

Journal write operations will happen synchronously within a host queue processing thread. Journal writes must be ordered as they are written to disk, so only one thread may write to the journal at any time. The address of the journal entry in the J1 table can be used as a unique identifier so that the J1 journal entry can be correlated with entries in the J2 journal. Once the journal entry is written, a transaction completion notification will be posted to the host completion queue. Now the write operation can be executed. It is important to ensure that any subsequent reads to a cluster before the journal write has completed are delayed.

The following table shows the format of the J2 journal record:

10

Size in bits	Name	Details
32	LBA	Logical Block Address
14	Zone	Zone # of associated cluster
18	Offset	Cluster offset of associated cluster
16	Size	Size of data
16	SequenceNumber	An incrementing sequence number so we can easily find the end of the queue

Each journal record will be aligned to a sector boundary. A journal record might contain an array of zone/offset/size tuples.

FIG. 15 shows a journal table update in accordance with an exemplary embodiment of the present invention. Specifically, when a host write request is received, the journal table is updated, one or more clusters are allocated, and data is written to the cluster(s).

Host journal requests are processed. These cause clusters to be written, and also cause updates to meta-data structure which must be shadowed back to disk (for example, the CAT table). It is important to ensure that these meta-data structures are correctly written back to disk even if a system reset occurs. A low level disk I/O write (J2) journal will be used for this.

In order to process a host interface journal entry, the appropriate manipulation of meta-data structures should be determined. The changes should occur in memory and a record of changes to various disk blocks be generated. This record contains the actual changes on disk that should be made. Each data structure that is updated is registered
5 with the J2 journal manager. This record should be recorded to a disk based journal, and stamped with an identifier. Where the record is connected with a J1 journal entry, the identifiers should be linked. Once the record is stored, the changes to disk can be made (or can be done via a background task).

The J2 journal exists logically at layer 3. It is used to journal meta-data updates that
10 would involve writes through the zone manager. When playback of a journal entry occurs, it will use zone manager methods. The journal itself can be stored in a specialized region. Given the short lifespan of journal entries, they will not be mirrored.

Not all meta-data updates need to go through the J2 journal, particularly if updates to
15 structures are atomic. The region manager structure may not use the J2 journal. It would be possible to detect inconsistencies in the region manager bitmap, for example, with an integrity checking background thread.

A simple approach for the J2 journal is to contain a single record. As soon as the record is committed to disk, it is replayed, updating the structures on disk. It is possible
20 to have multiple J2 records and to have a background task committing updating records on disks. In this case, close attention will need to be paid to the interaction between the journal and any caching algorithms associated with the various data structures.

The initial approach will run the journal entry as soon as it has been committed to disk. In principle there could be multiple concurrent users of the J2, but the J2 journal
25 may be locked to one user at a time. Even in this case, journal entries should be committed as soon as they have been submitted.

It is important to ensure that the meta-data structures are repaired before any higher level journal activity occurs. On system reboot, the J2 journal is analyzed, and any records will be replayed. If a journal entry is correlated with a J1 journal entry, the
30 J1 entry will be marked as completed, and can be removed. Once all J2 journal entries

have been completed, the meta-data is in a reliable state and any remaining J1 journal entries can be processed.

The J2 journal record includes the following information:

- Number of operations
- 5 • Each operation contains:
 - J1 record indicator
 - Zone/Data offset to write to
 - Data to write
 - Size of data
 - 10 ○ Offset into data cluster
- Journal record identifier
- End Marker

15 This scheme could operate similarly to the J1 journal scheme, for example, with a sequence number to identify the end of a J2 journal entry and placing J2 journal entries on sector boundaries.

20 If the J1 data pointer indicator is set, then this specific operation would point to a J1 journal record. The host supplied write data would not have to be copied into our journal entry. The operation array should be able to be defined as a fixed size as the maximum number of operations in a journal record is expected to be well understood.

25 In order to permit recovery from corruption of a sector during a low level write operation (e.g., due to a loss of power), the J2 journal could store the whole sector that was written so that the sector could be re-written from this information if necessary. Alternatively or additionally, a CRC calculated for each modified sector could be stored in the J2 record and compared with a CRC computed from the sector on disk (e.g., by the zone manager) in order to determine whether a replay of the write operation is required.

30 The different journals can be stored in different locations, so there will be an interface layer provided to write journal records to backing store. The location should be non-volatile. Two candidates are hard disk and NVRAM. If the J1 journal is stored to hard disk, it will be stored in a J1 journal non-mirrored zone. The J1 journal is a candidate for storing in NVRAM. The J2 journal should be stored on disk, although it

can be stored in a specialized region (i.e., not redundant, as it has a short lifespan). An advantage of storing the J2 journal on disk is that, if there is a system reset during an internal data structure update, the data structures can be returned to a consistent state (even if the unit is left un-powered for a long period of time).

- 5 The Zones Manager (ZM) allocates Zones that are needed by higher level software. Requests to the ZM include:
- a. Allocate Zone
 - b. De-allocate/Free Zone
 - c. Control data read/write pass through to L1 (?)
 - 10 d. Read/Write cluster in a Zone (given the offset of the cluster and the Zone number)

 The ZM manages the redundancy mechanisms (as a function of the number of drives and their relative sizes) and handles mirroring, striping, and other redundancy schemes for data reads/writes.

15 When the ZM needs to allocate a Zone, it will request an allocation of 2 or more sets of Regions. For example, a Zone may be allocated for 1GB of data. The Regions that make up this Zone will be able to contain 1GB of data including redundancy data. For a mirroring mechanism, the Zone will be made up of 2 sets of Regions of 1GB each. Another example, a 3-disk striping mechanism utilize 3 sets of Regions of 1/2 GB each.

20 The ZM uses the ZR translation table (6) to find out the location (drive number and start Region number) of each set of Regions that makes up the Zone. Assuming a 1/12GB Region size, a maximum of 24 Regions will be needed. 24 Regions make up 2x 1GB Zones. So the ZR translation table contains 24 columns that provide drive/region data.

25 The ZM works generally as follows:

- a. In the case of SDM (single drive mirroring), 24 columns are used. The drive numbers are the same in all columns. Each entry corresponds to a physical Region on a physical drive that makes up the Zone. The first 12 entries point to Regions that contain one copy of the data. The last 12 entries point to the Regions containing the second copy of the data.

- b. The case of DDM (dual drive mirroring) is the same as SDM except that the drive number on the first 12 entries is different from that in the last 12 entries.
- c. In the case of striping, three or more columns may be used. For example, if striping is used across three drives, six Regions may be needed from three different drives (i.e., 18 entries are used), with the first six entries containing the same drive number, the next six entries containing another drive number, and the following six entries containing a third drive number; the unused entries are zeroed.

The following table shows a representation of the zone region translation table:

Zone#	Size of Zone	Size of each Region	Usage	Drive/Region (1)	Drive/Region (2)	...	Drive/Region (23)	Drive/Region (24)
0	1GB	1/12	SDM	0,2000	0,1000	...	0,10	0,2000
1	1GB	1/12	DDM	0,8000	0, 3000	...	1,2000	1,10
2	1GB	1/12	SD3	3,4000	3, 3000		4,2000	4,1000
...								
N			Free					

When a read/write request comes in, the ZM is provided with the Zone number and an offset into that Zone. The ZM looks in the ZR translation table to figure out the redundancy mechanism for that Zone and uses the offset to calculate which Drive/Region contains the sector that must be read/written. The Drive/Region information is then provided to the L1 layer to do the actual read/write. An additional possible entry in the Usage column is "Free". "Free" indicates that the Zone is defined but currently not used.

The cluster manager allocates and de-allocates clusters within the set of data Zones.

The Layout Manager provides run-time re-layout of the Zones vis-à-vis their Regions. This may occur as a result of disk insertion/removal or failure.

The Layer 1 (L1) software knows about physical drives and physical sectors. Among other things, the L1 software allocates Regions from physical drives for use by the Zones Manager. In this exemplary embodiment, each Region has a size of 1/12GB (i.e., 174763 sectors) for a four-drive array system. A system with a larger maximum
 5 number of drives (8, 12 or 16) will have a different Region size.

In order to create a Zone containing 1GB of data with SD3 (striping over three drives; two data plus parity), we would end up using six Regions each in three drives ($6 \times 1/12 = 1/2$ GB per drive).

The use of this Region scheme allows us to provide better utilization of disk space
 10 when Zones get moved around or reconfigured e.g., from mirroring to striping. The L1 software keeps track of available space on the physical drives with a bitmap of Regions. Each drive has one bitmap. Each Region is represented by two bits in the bitmap in order to track if the Region is free, used, or bad. When the L2 software (ZM) needs to create a Zone, it gets a set of Regions from the L1 layer. The Regions that make up a Zone are
 15 not contiguous within a disk.

Requests to L1 include:

- a. Data read/write (to a cluster within a group of Regions)
- b. Control data read/write (tables, data structures, DIC etc)
- c. Allocate physical space for a Region (actual physical sectors within 1 drive)
- 20 d. De-allocate Region
- e. Raw read/write to physical clusters within a physical drive
- f. Copy data from one Region to another
- g. Mark region as bad.

25 The free region bitmap may be large, and therefore searches to find the free entry (worst case is that no entries are free) may be slow. In order to improve performance, part of the bitmap can be preloaded into memory, and a linked list of free regions can be stored in memory. There is a list for each active zone. If a low water mark on the list is reached, more free entries can be read from the disk as a background activity.

The Disk Manager operates at layer 0. As shown in the following table, there are two sub-layers, specifically an abstraction layer and the device drivers that communicate with the physical storage array.

Layer 0a : Abstraction
Layer 0b : OS interface to device drivers and device drivers
Physical Storage Array Hardware

5

The Device Drivers layer may also contain several layers. For example, for a storage array using USB drives, there is an ATA or SCSI stack on top of the USB transport layer. The abstraction layer provides basic read/write functions that are independent of the kinds of drives used in the storage array.

10

One or more disk access queues may be used to queue disk access requests. Disk access rates will be one of the key performance bottlenecks in our system. We will want to ensure that the disk interface is kept as busy as possible at all times so as to reduce general system latency and improve performance. Requests to the disk interface should have an asynchronous interface, with a callback handler to complete the operation when the disk operation has finished. Completion of one disk request will automatically initiate the next request on the queue. There may be one queue per drive or one queue for all drives.

15

Layer 1 will reference drives as logical drive numbers. Layer 0 will translate logical drive numbers to physical drive references (e.g., /dev/sda or file device number as a result of an open() call). For flexibility (expansion via USB), there should be a queue for each logical drive.

20

The following are some exemplary object definitions and data flows:

MSG object : incoming from host

25

- Lba
- Length
- LUN
- Data

REPLY object : outgoing to host

Status

Host

5 Length

Data

Data Read

10

Data read flow:

rc=lockm.islocked(MSG)

rc = catm.read(MSG, REPLY)

15 status = zonem.read(zone, offset, length, buffer)

regionm.read(logical_disk, region_number, region_offset, length, buffer)

diskm.read(logical_disk, offset, length, buffer)

20 ***Data Write***

Data write flow:

diskutildaemon.spaceavailable()

25 journalm.write(MSG)

lockm.lock(msg)

zonem.write(journal_zone, offset, length, buffer)

regionm.write - journal entry

diskm.write

30 regionm.write - end marker

diskm.write

```

catm.write(MSG)
    catm.readcluster(lba,offset,length,buffer)    - if need to merge sector into cluster
    - merge
5  "if(lba already allocated)"
    catm.readhashkey(lba)
    hashm.lookup(hashkey,zone,offset)
    "if(refcount==1)"
        hashentry.getrefcount()
10        hashm.remove(hashentry)
        hashm.add(sha1,zone,offset)
        zonem.write(zone,offset,length,buffer)    - write data
    "else"
        hashentry.removeref()
15        clusterm.allocate(zone,offset)          - allocate new cluster
        zonem.createzone(zone)
            regionm.unusedregions(logical_disk)
        regionm.allocate(logical_disk, number_regions, region_list)
        zonem.write(...)                          - write data
20        hashm.add(...)                          - add new entry to hash table
    "endif"
        hashdaemon.add(lba,sha1)                  - add to hash daemon Q

        catm.writehashkey(lba,hashkey)            - copy new hash key to CAT
25    "else"
        catm.update(lba,zone,offset,hashkey)      - update CAT with new entry
    "endif"
    journalm.complete(MSG)
30    lockm.unlock(MSG)
        - update r/w cursors

```

The following is a description of physical disk layout. As discussed above, each disk is divided into Regions of fixed size. In this exemplary embodiment, each Region has a size of 1/12GB (i.e., 174763 sectors) for a four-drive array system. A system with a larger maximum number of drives (8, 12 or 16) will have a different Region size. Initially, Region numbers 0 and 1 are reserved for use by the Regions Manager and are not used for allocation. Region number 1 is a mirror of Region number 0. All internal data used by the Regions Manager for a given hard disk is stored in Region numbers 0 and 1 of this hard disk. This information is not duplicated (or mirrored) to other drives. If there are errors in either Region 0 or 1, other Regions can be allocated to hold the data. The Disk Information Structure points to these Regions.

Each disk will contain a DIS that identifies the disk, the disk set to which it belongs, and the layout information for the disk. The first sector on the hard disk is reserved. The DIS is stored in the first non-bad cluster after the first sector. The DIS is contained in 1KB worth of data. There are two copies of the DIS. The copies of the DIS will be stored on the disk to which it belongs. In addition, every disk in the system will contain a copy of all the DISs of the disks in the system. The following table shows the DIS format:

Offset	Size	Name	Value/Valid Range	Description
0	32 bytes	disStartSignature	“_DISC INFORMATION CLUSTER START_”	Identifies the cluster as being a possible disc information cluster. Cluster must be CRC'd to check that it is valid.
	WORD16	disVersion	Binary non-zero number	Identifies the structure version. This value is only changed when a material change is made to the structure layout or

				content meaning that makes it incompatible with previous versions of the Firmware.
	WORD16	disClusterSize	Binary non-zero number	The number of 512 byte sectors that make a cluster on this disc.
	WORD32	disCRC	CRC-32	CRC of the DIS structure.
	WORD32	disSize ^{!!!}		Size of DIS cluster (in bytes)
	WORD32	disDiskSet		The disk set this disk belongs to
	WORD32	disDriveNumber	0 to 15	The drive number within the disk set
	WORD32	disSystemUUID		UUID of the box this disk belongs to
	WORD64	disDiskSize		Size of the disk in number of sectors
	WORD32	disRegionSize		Size of Regions in number of sectors
	WORD64	disRegionsStart		Sector offset to the start of the first Region on the disk
	WORD64	disCopyOffset		Sector offset to where the copy of this DIS is stored. The disCopyOffset of each DIS reference each other
	WORD64	disDISBackup		Sector offset to the table

				containing the copies of the DISs of all the disks
	WORD32	disDISBackupSize		Number of DISs in the DIS Backup section
	WORD32	disRIS0Region		Region number of where first copy of the RIS is stored
	WORD32	disRIS0Offset		Number of sectors offset within the Region to the sector where the Regions Information Structure is located
	WORD32	disRIS1Region		For the copy of the RIS
	WORD32	disRIS1Offset		For the copy of the RIS
	WORD32	disZIS0Region		Region number of Region where the Zones Information Structure is located. This is ONLY used if there is a ZTR on this disk. Otherwise, it is zero.
	WORD32	disZIS0Offset		Offset to the ZIS within the region
	WORD32	disZIS1Region		Region number of Region where a copy of the ZIS is located. This is ONLY used in a single drive system. In other cases, this entry is 0.
	WORD32	disZIS1Offset		Offset to the ZIS within the region

Regions Manager stores its internal data in a regions information structure. The following table shows the regions information structure format:.

Offset	Size	Name	Value/Valid Range	Description
0	WORD64	risSignature		Indicates that this is a RIS
	WORD32	risSize		Size of this structure (bytes)
	WORD32	risChecksum		Checksum
	WORD32	risVersion		Version of this table (and bitmap)
	WORD32	risDrive		Logical Drive number
	WORD64	risStartSector		Absolute start sector (in disk) of Regions utilization bitmap
	WORD32	risSectorOffset		Sector offset of Regions utilization bitmap within the current Region
	WORD32	risSizeBitmap		Size of bitmap (in bits?)
	WORD64	risNumberRegions		Number of regions on this disk (also implies size of bitmap)

5

The zones information structure provides information on where the Zones Manager can find the Zones Table. The following shows the zones information structure format:

10

Offset	Size	Name	Value/Valid Range	Description
0	WORD64	zisSignature		Indisates that this is a ZIS

8	WORD32	zisSize		Size of this structure (bytes)
12	WORD32	zisChecksum		Checksum
16	WORD32	zisVersion		Version of this table (and bitmap)
20	WORD16	zisFlags		Bit 0 = 1 if this disk is used to contain the Zones info Bits 14-15 : redundancy type (either SDM or DDM only)
22	WORD16	zisOtherDrive		Logical drive number of the drive that contains the other copy of the Zones Table.
24	WORD32	zisNumberRegions		Number of Regions used to contain each copy of the Zones Table. Equal to the number of Zones Table Nodes.
28	WORD32	zisStartOffset		Byte offset pointing to start of linked list of Regions that are used to contain the Zones Table. Each entry in the linked list is called ‘Zones Table Node’
	WORD32	zisNumberofZones		Number of Zones (entries in Zones Table) in the system
	WORD32	zisZoneSize		Size of zones in bytes

High level information zones contain the Zone tables and other tables used by the high level managers. These will be protected using mirroring.

The following table shows the zones table node format:

5

Size	Name	Description
------	------	-------------

WORD32	ztNextEntry	Pointer to next entry in linked list
WORD32	ztCount	Count of this entry
WORD64	ztRegion	Region number

The following is a description of layout of zones information. The linked list of Zones Table Nodes is placed after the ZIS in the following manner:

Zones Information Structure
First Zones Table Node (16bytes)
...
Last Zones Table Node (16 bytes)

5

This information is stored in the Zones Table Region.

Fig. 16 shows the drive layout in accordance with an exemplary embodiment of the invention. The first two regions are copies of one another. A third (optional) Zones Table Region contains the Zone Tables. In a system with more than one drive, only two of the drives contain a ZTR. In a system with only one drive, two Regions are used to hold the two (mirrored) copies of the ZTR. The DIS contains information on the location of the RIS and the ZIS. Note that the first copy of the RIS does not have to be in Region 0 (e.g., could be located in a different Region if Region 0 contains bad sectors).

The Zones Manager needs to load the Zones Tables on system start up. To do that, it extracts the Region number and offset from the DISs. This will point to the start of the ZIS.

Certain modules (e.g., the CAT Manager) store their control structures and data tables in Zones. All control structures for modules in Layer 3 and higher are referenced from structures that are stored in Zone 0. This means, for example, that the actual CAT (Cluster Allocation Tables) locations are referenced from the data structures stored in Zone 0.

The following table shows the zone 0 information table format:

Offset	Size	Name	Value/Valid Range	Description
0	WORD64	zitSignature		Indisates that this is a ZIT
	WORD32	zitSize		Size of this structure (bytes)
	WORD32	zitChecksum		Checksum of this structure
	WORD32	zitVersion		Version of this structure
	WORD32	zitCATLStartOffset		Byte offset (within this Zone) of start of CAT linked list
	WORD32	zitCATSize		Number of nodes in CAT linked list. Equal to number of Zones containing the CAT
	WORD64	zitCATAddressable		The max LBA supported by the CAT. Effectively the size of the CAT
	WORD32	zitHTStartOffset		Byte (within this Zone) of the start of the Hash Table linked list
	WORD32	zitHTNumberNodes		Number of nodes in Hash Table linked list
	WORD64	zitHTSize		Size of Hash Table data in bytes

The CAT linked list is a linked list of nodes describing the Zones that contain the CAT. The following table shows the CAT Linked List node format:

5

Size	Name	Description
WORD32	catllNextEntry	Pointer to next entry in linked list

WORD16	catllCount	Count of this entry
WORD16	catllZone	Zone number containing this portion of the CAT

5 The hash table linked list is a linked list of nodes that describe the Zones which hold the Hash Table. The following table shows the Hash Table Linked List node format:

Size	Name	Description
WORD32	htllNextEntry	Pointer to next entry in linked list
WORD16	htllCount	Count of this entry
WORD16	htllZone	Zone number containing this portion of the hash table

10 Fig. 17 demonstrates the layout of Zone 0 and how other zones are referenced, in accordance with an exemplary embodiment of the invention.

As discussed above, a Redundant set is a set of sectors/clusters that provides redundancy for a set of data. Backing up a Region involves copying the contents of a Region to another Region.

15 In the case of a data read error, the lower level software (Disk Manager or Device Driver) retries the read request two additional times after an initial failed attempt. The failure status is passed back up to the Zones Manager. The Zones Manager then attempts to reconstruct the data that is requested (by the read) from the redundant clusters in the disk array. The redundant data can be either a mirrored cluster (for SDM, DDM) or a set of clusters including parity (for a striped implementation). The reconstructed data is then

passed up back to the host. If the ZM is unable to reconstruct the data, then a read error is passed up back to the host. The Zones Manager sends an Error Notification Packet to the Error Manager. Fig. 18 demonstrates read error handling in accordance with an exemplary embodiment of the invention.

5 In the case of a data write error, the lower level software (Disk Manager or Device Driver) retries the write request two additional times after an initial failed attempt. The failure status is passed back up to the Zones Manager. The Zones Manager sends an Error Notification Packet to the Error Manager.

10 When a data write is performed at this level, the redundancy information is also written to disk. As a result, as long as only one cluster has a write error, a subsequent read will be able to reconstruct the data. If there are multiple disk errors and redundancy information cannot be read or written, then there are at least two possible approaches:

- a. Return a write error status to the host. Back up all the Regions associated with the redundant set to newly allocated Regions that do not contain bad sectors.
- 15 b. Hold off the write. Back up all the Regions associated with the redundant set to newly allocated Regions that do not contain bad sectors. Subsequently, do the write on the appropriate cluster in the newly allocated Regions (along with all redundancy parts e.g., parity etc.). A separate write queue would be used to contain the writes that have been held off.

20

Approach (a) is problematic because a write status would likely have already been sent to the host as a result of a successful write of the Journal, so the host may not know that there has been an error. An alternative is to report a failure with a read, but allow a write. A bit in the CAT could be used to track that the particular LBA should return a
25 bad read.

Fig. 19 demonstrates write error handling in accordance with an exemplary embodiment of the invention.

The Error Manager (EM) checks the cluster to see if it is really bad. If so, the entire region is considered bad. The contents of the Region are copied over to a newly
30 allocated Region on the same disk. The current Region is then marked BAD. While copying over the Region, the Error Manager will reconstruct data where necessary when

it encounters bad sectors. Fig. 20 is a logic flow diagram demonstrating backup of a bad Region by the Error Manager in accordance with an exemplary embodiment of the invention.

5 If there is a data read error and the Error Manager is unable to reconstruct the data for a given cluster (e.g., as a result of read errors across the redundant set) then zeros will be used in place of the data that cannot be reconstructed. In this case, other Regions (from the same Redundant Set) that contain bad sectors will also have to be backed up. Again, zeros will be used in place of the data that cannot be reconstructed.

10 Once a copy of the redundant set is made, the EM disables access to the clusters corresponding to this part of the Zone. It then updates the Zones Table to point to the newly allocated Regions. Subsequently, accesses to the clusters are re-enabled.

This exemplary embodiment is designed to support eight snapshots (which allows use of one byte to indicate whether hash/cluster entries are used by a particular snapshot instance). There are two tables involved with snapshots:

- 15
1. A per-snapshot CAT table will need to exist to capture the relationship between logical sector addresses and the cluster on the disk that contains the data for that LSA. Ultimately the per-snapshot CAT must be a copy of the CAT at the moment the snapshot was taken.
 - 20 2. The system hash table, which maps between hash values and a data cluster. The hash function returns the same results regardless of which snapshot instance is being used, and as a result is common across all snapshots. As a result, this table must understand whether a unique cluster is being used by any snapshots. A hash cluster entry can not be freed, or replaced with new data unless there are no
 - 25 snapshots using the hash entry.

There will always be a snapshot that is current and being added to. When a hash entry is created or updated, we will need to apply the current snapshot number to that hash entry. When a snapshot is made, the current snapshot number will be incremented.

30 Clusters/hash entries that are not longer required by any snapshots are freed by walking through the hash table and find any hash entries with the retiring snapshot bit set

and clearing that bit. If the snapshot byte is now zero, the hash entry can be removed from the table and the cluster can be freed.

To prevent collisions with any new entries being added to the hash tree (because the new snapshot number is the same as the retiring snapshot number), only allow 7
5 snapshots may be permitted to be taken, with the final (eighth) snapshot the one that is being retired. The hash table can be walked as a background activity.

In order to create a snapshot, a second CAT zone could be written whenever the main CAT is being updated. These updates could be queued and the shadow CAT could be updated as another task. In order to snapshot, the shadow CAT becomes the snapshot
10 CAT.

Once the snapshot is done, a background process can be kicked off to copy this snapshot table to a new zone become the new snapshot CAT. A queue could be used so that the shadow CAT queue is not processed until the copy of the CAT had completed. If a failure were to occur before updating the shadow CAT (in which case entries in the
15 queue may be lost), re-shadow from the primary CAT table could be performed before the array is brought online.

Alternatively, when a snapshot is required, a collection of “deltas” plus the initial CAT copy could make up the snapshot. A background task could then reconstitute a full snapshot CAT from this info. This would require little or no downtime to do the
20 snapshot. In the meantime, another set of deltas could be collected for the following snapshot.

FILESYSTEM-AWARE STORAGE SYSTEM

25 As discussed above, embodiments of the present invention analyze host filesystem data structures (i.e., metadata) in order to determine storage usage of the host filesystem and manage physical storage based on the storage usage of the host filesystem. For the sake of convenience, this functionality may be referred to hereinafter as a “scavenger.” A similar function, which may be referred to hereinafter as a “monitor,”
30 could be included to monitor storage usage but not necessarily manage the physical storage. Both scavengers and monitors are discussed below.

FIG. 27 is a conceptual block diagram of a computer system 2700 in accordance with an exemplary embodiment of the present invention. The computer system 2700 includes, among other things, a host computer 2710 and a storage system 2720. The host computer 2710 includes, among other things, a host operating system (OS) 2712 and a host filesystem 2711. The storage system 2720 includes, among other things, a filesystem-aware storage controller 2721 and storage 2722 (e.g., an array including one or more populated disk drives). Storage 2722 is used to store, among other things, storage controller data structures 2726, host OS data structures 2725, host filesystem data structures 2723, and user data 2724. The filesystem-aware storage controller 2721 stores various types of information in the storage controller data structures 2726 (represented by a dashed line between the filesystem-aware storage controller 2721 and the storage controller data structures 2726), such as a partition table including a reference to an OS partition (represented by a dashed line between the storage controller data structures 2726 and the host OS data structures 2725). The host OS 2712 stores various types of information in the host OS data structures 2725 (represented by a dashed line between the host OS 2712 and the host OS data structures 2725), typically including pointers/references to the host filesystem data structures 2723 (represented by a dashed line between the host OS data structures 2725 and the host filesystem data structures 2723). The host filesystem 2711 stores information in the host filesystem data structures 2723 (referred to as metadata, and represented by a dashed line between the host filesystem 2711 and the host filesystem data structures 2723) relating to the user data 2724. The filesystem-aware storage controller 2721 handles storage requests from the host filesystem 2711 (represented by a solid line between the host filesystem 2711 and the filesystem-aware storage controller 2721), utilizes the host OS data structures 2725 and host filesystem data structures 2723 to determine storage usage of the host filesystem 2711 (represented by a dashed line between the filesystem-aware storage controller 2721 and the host OS data structures 2725 and a dashed line between the filesystem-aware storage controller 2721 and the host filesystem data structures 2723), and manages storage of the user data 2724 based on the storage usage of the host filesystem 2711 (represented by a dashed line between the filesystem-aware storage controller 2721 and

the user data 2724). Specifically, the filesystem-aware storage controller 2721 may implement a scavenger and/or a monitor, as discussed below.

The filesystem-aware storage controller 2721 generally needs to have a sufficient understanding of the inner workings of the host filesystem(s) in order to locate and
5 analyze the host filesystem data structures. Of course, different filesystems have different data structures and operate in different ways, and these differences can affect design/implementation choices. Generally speaking, the filesystem-aware storage controller 2721 locates host filesystem data structures 2723 in storage 2722 and analyzes the host filesystem data structures 2723 to determine storage usage of the host filesystem
10 2711. The filesystem-aware storage controller 2721 can then manage the user data storage 2724 based on such storage usage.

FIG. 28 is high-level logic flow diagram for the filesystem-aware storage controller 2721, in accordance with an exemplary embodiment of the present invention. In block 2802, the filesystem-aware storage controller 2721 locates host filesystem data
15 structures 2723 in storage 2722. In block 2804, the filesystem-aware storage controller 2721 analyzes the host filesystem data structures to determine host filesystem storage usage. In block 2806, the filesystem-aware storage controller 2721 manages user data storage based on the host filesystem storage usage.

FIG. 29 is a logic flow diagram for locating the host filesystem data structures
20 2723, in accordance with an exemplary embodiment of the present invention. In block 2902, the filesystem-aware storage controller 2721 locates its partition table in the storage controller data structures 2726. In block 2904, the filesystem-aware storage controller 2721 parses the partition table to locate the OS partition containing the host OS data structures 2725. In block 2906, the filesystem-aware storage controller 2721 parses
25 the OS partition to identify the host OS 2712 and locate the host OS data structures 2725. In block 2908, the filesystem-aware storage controller 2721 parses the host OS data structures 2725 to identify the host filesystem 2711 and locate the host filesystem data structures 2723.

Once the filesystem-aware storage controller 2721 locates the host filesystem data
30 structures 2723, it analyzes the data structures to determine storage usage of the host filesystem 2711. For example, the filesystem-aware storage controller 2721 may use the

host filesystem data structures 2723 for such things as identifying storage blocks no longer being used by the host filesystem 2711 and identifying the types of data stored by the host filesystem 2711. The filesystem-aware storage controller 2721 could then dynamically reclaim storage space no longer being used by the host filesystem 2711 and/or manage storage of the user data 2724 based on data types (e.g., store frequently accessed data uncompressed and in sequential blocks to facilitate access, store infrequently accessed data compressed and/or in non-sequential blocks, and apply different encodingschemes based on data types, to name but a few).

FIG. 30 is a logic flow diagram for reclaiming unused storage space, in accordance with an exemplary embodiment of the present invention. In block 3002, the filesystem-aware storage controller 2721 identifies blocks that are marked as being unused by the host filesystem 2711. In block 3004, the filesystem-aware storage controller 2721 identifies any blocks that are marked as unused by the host filesystem 2711 but are marked as used by the filesystem-aware storage controller 2721. In block 3006, the filesystem-aware storage controller 2721 reclaims any blocks that are marked as used by the filesystem-aware storage controller 2721 but are no longer being used by the host filesystem 2711 and makes the reclaimed storage space available for additional storage.

FIG. 31 is a logic flow diagram for managing storage of the user data 2724 based on the data types, in accordance with an exemplary embodiment of the present invention. In block 3102, the filesystem-aware storage controller 2721 identifies the data type associated with particular user data 2724. In block 3104, the filesystem-aware storage controller 2721 optionally stores the particular user data 2724 using a storage layout selected based on the data type. In block 3106, the filesystem-aware storage controller 2721 optionally encodes the particular user data 2724 using an encoding scheme (e.g., data compression and/or encryption) selected based on the data type. In this way, the filesystem-aware storage controller 2721 can store different types of data using different layouts and/or encoding schemes that are tailored to the data type.

One example of a scavenger is the so-called “garbage collector.” As discussed above, the garbage collector may be used to free up clusters which are no longer used by the host file system (e.g., when a file is deleted). Generally speaking, garbage collection

works by finding free blocks, computing their host LSAs, and locating their CAT entries based on the LSAs. If there is no CAT entry for a particular LSA, then the cluster is already free. If, however, the CAT entry is located, the reference count is decremented, and the cluster is freed if the count hits zero.

5 One concern is that it may be difficult for the garbage collector to distinguish a block that the host filesystem has in use from one that it has previously used and at some point marked free. When the host filesystem writes a block, the storage system allocates a cluster for the data as well as a CAT entry to describe it. From that point on, the cluster will generally appear to be in use, even if the host filesystem subsequently ceases to use
10 its block (i.e., the cluster will still be in use with a valid CAT entry).

For example, certain host filesystems use a bitmap to track its used disk blocks. Initially, the bitmap will indicate all blocks are free, for example, by having all bits clear. As the filesystem is used, the host filesystem will allocate blocks through use of its free block bitmap. The storage system will associate physical storage with these filesystem
15 allocations by allocating clusters and CAT entries as outlined earlier. When the host filesystem releases some blocks back to its free pool, it simply needs to clear the corresponding bits in its free block bitmap. On the storage system, this will generally be manifested as a write to a cluster that happens to contain part of the host's free block
20 I/O to the freed cluster, for example, if the host filesystem were running in some enhanced security mode, in which case it would likely write zeros or a crypto strong hash of random data to the cluster in order to reduce the chance that stale cluster contents can be read by an attacker). Furthermore, there is no guarantee that the host filesystem will reuse blocks that it has previously freed when satisfying new allocation requests. Thus, if
25 the host filesystem continues to allocate what from the storage system's point of view are new, i.e. previously unused, blocks then the storage system will quickly run out of free clusters, subject to whatever space can be reclaimed via compression. For example, assuming a filesystem block is 4k, if the host allocates filesystem blocks 100 through 500, subsequently frees blocks 300 through 500, and then allocates blocks 1000 through
30 1100, the total filesystem usage will be 300 blocks, and yet the array will have 500 clusters in use.

In an exemplary embodiment of the present invention, the storage system may detect the release of host filesystem disk resources by accessing the host filesystem layout, parsing its free block bitmaps, and using that information to identify clusters that are no longer being used by the filesystem. In order for the storage system to be able to
5 identify unused clusters in this way, the storage system must be able to locate and understand the free block bitmaps of the filesystem. Thus, the storage system will generally support a predetermined set of filesystems for which it “understands” the inner working sufficiently to locate and utilize the free block bitmaps. For unsupported filesystems, the storage system would likely be unable to perform garbage collection and
10 should therefore only advertise the real physical size of the array in order to avoid being overcommitted.

In order to determine the filesystem type (e.g., NTFS, FAT, ReiserFS, ext3), the filesystem’s superblock (or an equivalent structure) needs to be located. To find the superblock, the partition table will be parsed in an attempt to locate the OS partition.
15 Assuming the OS partition is located, the OS partition will be parsed in an attempt to locate the superblock and thereby identify the filesystem type. Once the filesystem type is known, the layout can be parsed to find the free block bitmaps.

In order to facilitate searching for free blocks, historical data of the host filesystem bitmap can be kept, for example, by making a copy of the free block bitmap
20 that can be stored in a private, non-redundant zone and performing searches using the copy. Given the size of the bitmap, information may be kept for a relatively small number of clusters at a time rather than for the whole bitmap. When a garbage collection is performed, the current free block bitmap can be compared, cluster-by-cluster, with the historical copy. Any bitmap entries transitioning from allocated to free can be identified,
25 allowing the scavenging operation to be accurately directed to clusters that are good candidates for reclamation. As each bitmap cluster is processed, the historical copy can be replaced with the current copy to maintain a rolling history of bitmap operations. Over time the copy of the free block bitmap will become a patchwork of temporally disjoint clusters, but since the current copy will always be used to locate free entries, this
30 does not cause any problems.

Under certain conditions, there could be a race condition regarding the free block bitmap, for example, if the host filesystem allocates disk blocks using its free block bitmap, then writes its data blocks, then flushes the modified bitmap back to disk. In such a case, the garbage collector might free a cluster even though the filesystem is using the cluster. This could lead to filesystem corruption. The storage system should be implemented to avoid or handle such a condition.

Because garbage collection can be a fairly expensive operation, and since even lightweight scavenging will consume back-end I/O bandwidth, garbage collection should not be overused. The garbage collector should be able to run in several modes ranging from a light background lazy scavenge to an aggressive heavyweight or even high priority scavenge. For example, the garbage collector could be run lightly when 30% of space is used or once per week at a minimum, run slightly more heavily when 50% of space is used, and run at a full high-priority scavenge when 90% or more of disk space is used. The aggressiveness of the garbage collector could be controlled by limiting it to a target number of clusters to reclaim and perhaps a maximum permissible I/O count for each collection run. For example, the garbage collector could be configured to reclaim 1 GB using no more than 10,000 I/Os. Failure to achieve the reclaim request could be used as feedback to the collector to operate more aggressively next time it is run. There may also be a "reclaim everything" mode that gives the garbage collector permission to parse the entire host filesystem free block bitmap and reclaim all blocks that it possibly can. This might be done as a last ditch attempt to reclaim clusters when the array is (almost) completely full. The garbage collector may be run periodically to apply its rules and may or may not decide to perform a scavenge operation. The scavenge operation should also be able to be explicitly requested from another module, for example the region manager when it is struggling to find clusters to build a region.

The garbage collection function can be tied into the status indicator mechanism. For example, at some point, the storage system might be in a "red" condition, although an ongoing garbage collection operation might free up enough space to erase the "red" condition. Additional indicator states could be employed to show related status information (e.g., the red indicator light might be made to blink to indicate that a garbage collection operation is ongoing).

FIG. 21 is a schematic block diagram showing the relevant components of a storage array in accordance with an exemplary embodiment of the present invention. Among other things, the storage array includes a chassis 2502 over which a storage manager 2504 communicates with a plurality of storage devices 2508₁-2508_N, which are
5 coupled to the chassis respectively through a plurality of slots 2506₁-2506_N. Each slot 2506₁-2506_N may be associated with one or more indicators 2507₁-2507_N. Among other things, the storage manager 2504 typically includes various hardware and software components for implementing the functionality described above. Hardware components typically include a memory for storing such things as program code, data structures, and
10 data as well as a microprocessor system for executing the program code.

One concern for implementing the filesystem-aware storage controller 2721 is that many host filesystems do not update the data structures (i.e., metadata) in real time. For example, journaling filesystems do not normally guarantee to preserve all of the user data for transactions that have already happened or guarantee to recover all of the
15 metadata for such transactions, but generally only guarantee the ability to recover to a consistent state. For performance and efficiency, journaling filesystems often deploy some degree of asynchronicity between user data writes and metadata writes. In particular, it is common for the metadata writes to disk to be performed lazily such that there is a delay between a user data update and a corresponding metadata update. Journal
20 writes may also be performed lazily in some filesystems (such as NTFS according to Ed 4 of Microsoft Windows Internals). Furthermore, lazy metadata writes may be performed by play-out of the journal in a transaction-by-transaction manner, and that has considerable potential to push the metadata temporarily into states that are inconsistent with the user data already on-disk. An example of this would be a bitmap update
25 showing a de-allocation after the host had re-allocated the cluster and sent user data corresponding to that reallocated cluster. Thus, the storage system generally needs to cope with metadata updates that do not reliably indicate the present state of the user data. In the previous example, that would generally mean that the storage system could not interpret the de-allocation to mean that the cluster was reclaimable and the user data
30 discardable.

If metadata and journal updates are entirely asynchronous to the user data updates to which they correspond, such that they can happen at any time, the storage system may need to have a relatively detailed understanding of the inner workings of the filesystem, and therefore may need to store extensive state information in order to make appropriate decisions. Specific embodiments of the present invention described below, however, are designed to operate under the assumption that metadata updates will occur within a relatively deterministic window of time after the user data writes to which they pertain (e.g., within one minute). It is understood that such embodiments are essentially trade-offs between complexity and functionality in that they generally are not required to have a detailed understanding of inner workings of the filesystem and storage of extensive state information, although special considerations may need to be made for handling host filesystems that do not adhere to such a window during operation (VxFS may be one example) or boundary conditions that result in a long delay between user data updates and corresponding metadata updates (e.g., a loss of function from the host, which is generally beyond the control of the storage system and might lead to data loss anyway, or a loss of connectivity, which the storage system could detect and thereby suppress any behavior that assumes timely host activity).

In one exemplary embodiment, the scavenger could operate in a purely asynchronous manner. In this embodiment, the scavenger may be a purely asynchronous task that periodically scans the bitmap, either in whole or part, and compares the bitmap with information contained in the CAT to determine whether any of the storage array clusters can be reclaimed. Before checking a bitmap, the system may also check those blocks that contain the location of the bitmap in order to determine whether the bitmap has moved.

One advantage of a purely asynchronous scavenger is that there is essentially no direct impact on processor overhead within the main data path, although it may involve substantial asynchronous disk I/O (e.g., for a 2TB volume logically divided into 4k clusters and having a 64MB bitmap, reading the whole bitmap would involve reading 64+MB of disk data every time the scavenger runs) and therefore may impact overall system performance depending on how often the scavenger runs. Therefore, the scavenger frequency may be varied depending on the amount of available storage space

and/or system load. For example, when available storage space is plentiful or system load is high, the scavenger function could be run less frequently. Decreasing scavenger frequency will generally decrease the rate at which storage space is reclaimed, which is generally acceptable when storage space is plentiful. On the other hand, when available
5 storage space is scarce and system load is low, the scavenger function could be run more frequently in order to increase the rate at which storage space is reclaimed (at the expense of added processing overhead).

In another exemplary embodiment, the scavenger could operate in a partly synchronous, partly asynchronous manner. In this embodiment, the scavenger could
10 monitor changes to the bitmap as they occur, for example, by adding some additional checks to the main write handling path. The scavenger could construct a table at boot time that includes the LBA range(s) of interest (hereinafter referred to as the Bitmap Locator Table or BLT). For an uninitialized disk or an initialized but unpartitioned disk, the BLT would generally include only LBA 0. For a fully initialized and formatted disk,
15 the BLT would generally include LBA 0, the LBA(s) of every partition boot sector, the LBA(s) containing the bitmap metadata, and the LBA range(s) containing the bitmap data itself.

The main write handling path (e.g., HRT) typically calls the scavenger with details of the write being handled, in which case the call would generally internally cross-
20 reference the LBA(s) of the write request with the BLT with a view to identifying those writes which overlap with the LBA range(s) of interest. The scavenger would then need to parse those writes, which could be mostly done with an asynchronous task (in which case key details would generally need to be stored for the asynchronous task, as discussed below), but with critical writes parsed inline (e.g., if an update is potentially indicative of
25 a relocated bitmap, that write could be parsed inline so that the BLT may be updated before any further writes are cross-referenced). As discussed above with reference to a purely asynchronous scavenger, the frequency of the asynchronous task could be varied depending on the amount of available storage space and/or system load.

Storage for the asynchronous task could be in the form of a queue. A simple
30 queue, however, would allow queuing of multiple requests for the same block, which could occur because the semantics of a write cache makes it likely that a number of

requests would point to the same data block in the cache (i.e., the most recent data) and is inefficient because there is generally no reason to hold multiple requests representing the same LBA. This could be alleviated by checking through the queue and removing earlier requests for the same block. Furthermore, assuming the frequency of the asynchronous task is varied depending on the amount of available storage space and/or system load, then the queue should be provisioned with the expectation that it will reach its maximum size during periods of intense activity (which might be sustained for days) in which the asynchronous task is suppressed. Assuming the system is disallowing multiple entries for the same LBA, the maximum theoretical size of the queue is a product of the size of the LBA and the number of LBAs within the bitmap, which could result in very large queue size (e.g., a 2TB volume has a 64MB bitmap (i.e., 128K blocks) and therefore might require a queue size on the order of $128K * 4 = 512K$; a 16TB volume might require a queue size on the order of 4MB.

Alternate storage for the asynchronous task could be in the form of a bitmap of the bitmap (referred to hereinafter as the "Bitmap Block Updates Bitmap" or "BBUB"), with each bit representing one block of the real bitmap. The BBUB inherently avoids multiple requests for the same block, since the same bit is set for each of such requests, so the multiple requests only show up once in the BBUB. Furthermore, the size of the BBUB is essentially fixed, without regard to the frequency of the asynchronous task, and generally occupies less space than a queue (e.g., the BBUB would occupy 16KB of memory for a 2TB volume, or 128KB for a 16TB volume). In the event that the real bitmap moves, the storage system can easily adjust the mapping of the bits in the BBUB, but will generally need to take care not to map pending requests to the new location before the host has copied the data across (in fact, it may be possible to zero out the bitmap on the assumption that the host filesystem will rewrite every LBA anyway). The BBUB may be placed in non-volatile memory (NVRAM) to prevent loss of the current BBUB or may be placed in volatile memory, with the understanding that the current BBUB could be lost and a complete scan of the bitmap would need to be run sometime after reboot to recover the lost information. Since a bitmap does not inherently provide the asynchronous task with a ready measure of the number of requests, the storage system could maintain statistics about the number of bits set in the bitmap so that the

asynchronous task does not have to scan through the entire bitmap just to determine that nothing has been updated. For example, the storage system might maintain a count of how many bits are set in the bitmap and may adjust the frequency of the asynchronous task based on the count (e.g., do not run the asynchronous task unless and until the count reaches a predetermined threshold, which may be user-configurable). Such a strategy may be refined further, for example, by maintaining a separate count of set bits for each of a number of map sections (e.g., 1K chunks) and also keeping track of which map section has the highest count, so that the asynchronous task can parse just the map section(s) that are likely to return the greatest reward.

10 Parsing the updates generally involves different logic for different LBAs. For example, a change to LBA 0 generally means that the partition table has been added, or that a partition has been added to the table, or that a partition has been deleted. An update to the partition boot sector might mean that the bitmap metadata has been relocated. An update to the bitmap metadata might mean that the bitmap has been
15 moved, or that it has been extended. An update to the bitmap itself might indicate allocation or de-allocation of clusters. If the updates are parsed asynchronously, then the system generally cannot readily compare the old data with the new because, by the time the asynchronous task runs, the new data may have overwritten the old. To avoid this problem, the system might keep a separate copy of the old data for comparison or might
20 iterate through the map and compare the unset bits with the CAT (which might require slightly more processor overhead but less disk I/O). Simply comparing the bitmap with the CAT would generally require additional logic and state information, since the bitmap state may not be synchronized with the user data, as discussed above. Furthermore, keeping a copy of the bitmap data would allow the storage system to compare new data
25 with old and thereby determine exactly what changed, but the storage system generally cannot rely on state transitions as an accurate view of current user data state any more than it can rely on the state itself, as discussed above.

In yet another exemplary embodiment, the scavenger could operate in a purely synchronous manner. In this embodiment, the scavenger would process writes as they
30 occur. One advantage of a purely synchronous embodiment is that it avoids the complexities associated with operation of an asynchronous task and its associated

storage, although it interjects overhead on the processor during the critical time of handling a write from the host, and additional logic and state information might be required to compensate for asynchronous metadata updates.

One concern with reclaiming clusters in the context of asynchronous bitmap updates is that the scavenger might free clusters inappropriately, based on bitmap values that do not accurately reflect the state of the user data. To safeguard against such problems, the storage system may keep some history of cluster accesses it performs (e.g., whether or not it has recently accessed the user data in a cluster) and only reclaim a cluster if the cluster has been quiescent over some previous time interval to ensure that no metadata updates are pending for that cluster. For example, the storage system might require a cluster to be quiescent for at least one minute before performing any reclamation of the cluster (generally speaking, increasing the quiescent time reduces the risk of inappropriate reclamation but increases latency in reacting to data deletion, so there is a trade-off here). The storage system could track only cluster writes, although the storage system could additionally track cluster reads for thoroughness in assessing cluster activity, albeit at the expense of additional disk I/O). The quiescent time could be a fixed value or could be different for different filesystems.

Cluster accesses could be tracked, for example, by writing a scavenger cycle number to the CAT as an indicator of access time relative to the scavenger runs.

Cluster accesses could alternatively be tracked by writing bits to the filesystem's bitmap prior to writing the data. Any such modification of the filesystem's metadata would have to be coordinated carefully, though, in order to avoid any adverse interactions with filesystem operation.

Cluster access could alternatively be tracked using a bit for each cluster, block, or chunk (of whatever size). The bit would generally be set when that entity is accessed and might be reset when the scavenger completes its next run or when the scavenger next tries to reclaim the cluster. The scavenger generally would only reclaim the cluster if this bit was already reset when trying to perform the reclamation, which would itself be driven by the corresponding bit in the real host filesystem bitmap being clear. These bits could be kept together as a simple bitmap or could be added to the CAT as a distributed bitmap (requiring an additional one bit per CAT record). The simple bitmap approach

may require an additional read-modify-write on most data write operations, potentially causing a decrease in performance of the main data path unless the bitmap is cached in memory (the bitmap could be cached in volatile memory, which could be problematic if the bitmap is lost due to an unexpected outage, or in non-volatile memory, which might
5 necessitate a smaller bitmap due to memory constraints and therefore less granularity). The CAT approach would generally benefit from the J2 and its free NVRAM caching.

Cluster accesses could alternatively be tracked by maintaining timestamps of when bitmap updates are received and when cluster modifications are performed. Then, if a cluster modification has a later timestamp than the bitmap update, the system
10 generally would not free the cluster. One advantage of this approach over the bit approach is that the scavenger can determine how long ago the last access occurred and, if long enough, reclaim the cluster immediately. A timestamp may also be added to the CAT record. Alternatively, as the field only really needs to indicate age relative to scavenger operation, a global identifier may be assigned to each scavenger run. The
15 system may then use a similar field within the CAT to show the value of the global identifier. The global identifier may identify which scavenger run had most recently completed, or was next due, or when the cluster was last accessed. This information could then be used by the scavenger as a measure of age. To save on space consumption in the CAT record, the identifier could just be a one byte counter. Any incorrect age
20 determinations due to the counter wrapping will be old clusters looking much younger than they are. These clusters will be reclaimed on the next run. The field may be stored in NVRAM to prevent the field from being reset to zero on every reboot, which could cause some cluster accesses to age prematurely.

Thus, for example, every scavenger run may be associated with a one-byte
25 identifier value, which could be implemented as a global counter in NVRAM that increments each time the scavenger wakes such that the identifier for a scavenger run will be the post-increment value of the counter. The CAT manager could use the the current value of the global counter whenever it services an update to a cluster and could store a copy of that value in the corresponding CAT record. Such an implementation would
30 require modification of the CAT manager logic.

Cluster access could alternatively be tracked by keeping a short history of cluster updates in a wrapping list. The scavenger could then search the list to verify that any cluster it was about to free had not recently been accessed by the host. The size of the list would generally be implementation-specific. However long it was, the storage system
5 would generally have to ensure that it could run the asynchronous task before the list could get full, and that would compromise the ability to postpone the task until a quiet period.

In a storage system supporting scavenging, it might be desirable to identify and track premature reclamations, particularly reads that fail because of premature
10 reclamation (i.e., an attempt to read from a cluster that has been freed by the scavenger), but also writes to unallocated clusters (which will generally just result in allocation and should therefore be harmless). In some situations, it may be possible to identify errors based on the filesystem bitmap (e.g., cross reference to the bitmap from the user data write and check that the appropriate bit is set), but only if bitmap updates are guaranteed
15 to be completed ahead of the user data, which is not always the case. Alternatively, when parsing the bitmap, the scavenger might check whether the allocated bits actually correspond to allocated clusters; allocate clusters, or at least CAT records, where they do not; set a bit in each such CAT record indicating that the allocation was forced from the scavenger; a bit that would be reset by a data write to the cluster; and check the bit again
20 on the next scavenger run and scream if its still set. Additional self-diagnostics could be included, such as a count of the number of times that cluster reclamation was aborted due to this preventative measure, to give us a measure of which filesystems do this and how much.

It should be noted that the three types of scavengers described above are
25 exemplary only and do not limit the present invention to any particular design or implementation. Each scavenger type has certain relative advantages and disadvantages that may make it particularly suitable or unsuitable for a particular implementation. Furthermore, it should be noted that particular implementations could support more than one of the scavenger types and dynamically switch between them as needed, for example,
30 based on such things as the host filesystem, the amount of available storage space, and the system load. A partly synchronous, partly asynchronous scavenger, using a BBUB to

store information for the asynchronous task, and using byte-size scavenger run counter (as a timestamp of sorts) within the CAT to track cluster accesses, is contemplated for a particular implementation.

A separate monitor in addition to, or in lieu of, a scavenger could be used to keep track of how many clusters are being used by the host filesystem (for example, a scavenger might be omitted if the host filesystem is known to reliably reuse de-allocated blocks in preference to using new blocks so that reclamation is not needed and monitoring would be sufficient; a monitor might be omitted as duplicative in systems that implement a scavenger). Generally speaking, the monitor only needs to determine how many bits are set in the bitmap and does not need to know precisely which bits are set and which bits are clear. Furthermore, the monitor may not need a precise bit count, but may only need to determine whether the number of set bits is more or less than certain threshold values or whether the number is more or less than a previous value for the same region. Therefore, the monitor may not need to parse the whole bitmap. As for the scavenger embodiments described above, the monitor function could be implemented in whole or in part using an asynchronous task, which could periodically compare the new data with the CAT or maintain a copy of the bitmap and compare the current bitmap (new data) with the copy (old data) before overwriting the copy with the new data.

For the sake of convenience, various design and operation considerations are discussed below with reference to scavengers, predominantly in the context of NTFS. It should be appreciated, however, that many of the design and operational considerations apply equally to monitors.

FIG. 32 is a schematic block diagram showing the relevant components of a scavenger 3210, in accordance with an exemplary embodiment of the present invention. Among other things, the scavenger 3210 includes a Bitmap Block Updates Monitor (BBUM) 3211, a collection of Bitmap Locator Tables (BLTs) 3212 including a BLT for each LUN, a collection of BBUBs 3213 including one BBUB for each partition, an asynchronous task 3214, and De-allocated Space Tables (DSTs) 3215. Each of these components are discussed in greater detail below. Also as discussed in greater detail below, the BBUM 3211 is informed of write operations through calls received from the HRM 3220.

The scavenger 3210 includes a BLT 3212 for each LUN. Each BLT 3212 contains a series of records that include a partition identifier, an LBA range, an indication of the role of the LBA range, and a flag indicating whether or not that LBA range should be parsed synchronously or asynchronously. Each BLT has an entry for LBA 0, which is partition independent. The BLTs are generally required to provide rapid LBA-based lookup for incoming writes on that LUN (without checking which partition they belong to first) and to provide relatively rapid partition based lookup for LBA 0 writes (which could be achieved, for example, using a sorted vector for storage and, for the lookup, a lower bound binary search of the start LBA plus a check whether the previous element has a lastLBA higher than the LBA being looked up). The BLTs will generally need to be programmed prior to any host writes passing through, for example, during the LoadDiskPack call. A BLT gets programmed with LBA 0, as this is the location of the partition table, and therefore the creation of a partition involves a write to this LBA. LBA 0 will be flagged, within this table, as a location whose updates require immediate parsing.

The scavenger 3210 includes a BBUB 3213 for each partition supported by the storage system. Each BBUB 3213 is sized appropriately for the size of the filesystem bitmap to which it pertains. Each BBUB 3213 is associated with a counter reflecting how many bits are set in the bitmap. The BBUBs 3213 also some mapping information showing how each bitmap pertains to its corresponding filesystem bitmap.

The scavenger 3210 includes a DST 3215 for each LUN. Each DST 3215 includes one LBA range per record. Each LBA range present in the table is part of a deleted or truncated partition that needs to be reclaimed from the CAT. The BBUM 3211 may update the DSTs 3215, for example, when it identifies an unused storage area for reclamation during synchronous processing (in which case the BBUM 3211 adds an LBA range to the DSTs 3215). Similarly, the asynchronous task 3214 may update the DSTs 3215, for example, when it identifies an unused storage area for reclamation during asynchronous processing (in which case it adds an LBA range to the DSTs 3215). The asynchronous task 3214 uses the DSTs 3215 to reclaim unused storage space asynchronously. The DSTs 3215 may be stored persistently in a way that is resilient to unclean shutdown, or else additional logic may be provided to recover from any loss of

the DSTs 3215, e.g., by performing a full scan after boot to find allocated clusters that do not belong to any volume.

Storage of the BBUBs 3213 and DSTs 3215 is an implementation-specific decision. In an exemplary embodiment, the BBUBs 3213 are too large to be stored in
5 NVRAM and therefore may be stored in volatile memory or on disk, while the DSTs 3215 may be stored in non-volatile memory, volatile memory, or on disk. If the DSTs 3215 and BBUBs 3213 are completely volatile, then the scavenger 3210 generally must be capable of recovering from a loss of the DSTs 3215 and BBUBs 3213 (e.g., due to an unexpected shutdown). Recovery might be accomplished, for example, by scanning
10 through the entire CAT and comparing it with current partition and cluster bitmap information to see whether each cluster is mapped to a known partition and whether it is allocated in the cluster bitmap of the corresponding filesystem. Another possibility is to store the DSTs 3215 in NVRAM and leave the BBUBs 3213 in volatile memory so that state information for disk space outside of volumes would be preserved across reboots
15 (potentially preventing the need to query the CATM about clusters outside of partitions), although the current state of the cluster bitmaps would be lost, necessitating a full scan of every cluster bitmap. Such bitmap scans could be reduced or eliminated, for example, by storing all the required state information on disk and merely reloading it on boot. Because the scavenger 3210 cannot expect notification of shutdown, the records would
20 need to be kept closely synchronized with real time state, either by updating them synchronously or by writing them back within just a few milliseconds or seconds. It seems reasonable to assume that intentional shutdowns will be preceded by a few seconds of inactivity from the host, even if the host isn't actually shutdown, so updating within a few seconds is probably adequate for most cases; shutdowns that occurred mid-I/O would
25 likely still require a full scan, though. If the records are updated synchronously (i.e., before writing the bitmap update to disk), then the system might be able to completely eliminate the loss of state and the corresponding need for a full scan on boot (although at the expense of requiring more disk I/O during steady state operation to buy better efficiency on boot). Another option is to write the BBUBs 3213 and DSTs 3215 to disk
30 during the system shutdown procedure so that the information would be available on

reboot (except in the event of an unexpected failure/shutdown, in which case a full scan of the clusters may be needed on reboot).

Generally speaking, the scavenger 3210 does not have much to do until a disk pack is loaded, although, in an exemplary embodiment, it is contemplated that the scavenger 3210 will be initialized by the system manager after initializing the modules that the scavenger depends upon, such as CAT Manager or the Cache Manager (for reading from the DiskPack) and the NVRAM Manager (for incrementing a counter). Alternatively, the scavenger could be initialized lazily, e.g., after a DiskPack is loaded. Since the scavenger could begin reading from the DiskPack almost immediately, the scavenger should not be instructed to load the DiskPack (i.e., LoadDiskPack) until the other components are ready and have loaded the same DiskPack themselves.

During initialization of the scavenger 3210 (or at some other appropriate time), the BBUM 3211 looks for the NTFS Partition Table at LBA 0. The NTFS Partition Table is a 64-byte data structure located in the same LBA as the Master Boot Record, namely LBA 0, and contains information about NTFS primary partitions. Each Partition Table entry is 16 bytes long, making a maximum of four entries available. Each entry starts at a predetermined offset from the beginning of the sector and a predetermined structure. The partition record includes a system identifier that enables the storage system to determine whether the partition type is NTFS or not. It has been found that the Partition Table position and layout is generally somewhat independent of the operating system that writes it, with the same partition table structure serving a range of filesystem formats, not just NTFS, and not just Microsoft formats (HFS+ and other filesystems may use a different structure to locate its partition).

Assuming the NTFS Partition Table is found at LBA 0, the BBUM 3211 reads the Partition Table from LBA 0, and then, for each NTFS partition identified in the Partition Table, reads the boot sector of the partition (the first sector of the partition), and in particular the extended BIOS partition block, which is a structure proprietary to NTFS partitions that will provide the location of the Master File Table (MFT). The BBUM 3211 then reads the resident \$bitmap record of the MFT to get the file attributes, in particular the location(s) and length(s) of the actual bitmap data. The BBUM 3211 also programs the BLTs 3212 with the boot sector LBA of each partition, the LBA(s) of the

bitmap record(s), and the LBAs of the actual bitmaps. Boot sector LBAs and bitmap record LBAs will also be flagged as locations whose updates always require immediate parsing. The actual bitmap generally does not need immediate parsing and will be flagged accordingly. If no partition table is found at LBA 0, then no additional locations are
 5 added to the BLTs 3212.

FIG. 33 is pseudo code for locating the host filesystem bitmaps, in accordance with an exemplary embodiment of the present invention. The filesystem-aware storage controller 2721 first looks for the partition table at LBA 0. Assuming the partition table is found, then the filesystem-aware storage controller 2721 reads the partition table to
 10 identify partitions. Then, for each partition, the filesystem-aware storage controller 2721 reads the boot sector of the partition to find the MFT and reads the resident \$bitmap record of the MFT to get file attributes, such as the location(s) and length(s) of the actual bitmaps. The filesystem-aware storage controller 2721 programs the BLTs with the boot sector LBA of each partition, the LBA(s) of the bitmap record(s), and the LBA(s) of the
 15 actual bitmap(s), and flags the boot sector LBA(s) and the bitmap record LBA(s) to require immediate parsing and flags the actual bitmap(s) to not require immediate parsing. If the filesystem-aware storage controller 2721 is unable to find the partition table at LBA 0, then the filesystem-aware storage controller 2721 ends without adding additional locations to the BLTs.

20 During steady-state operation, all writes will be cross-referenced against the BLTs 3212 through a call to the BBUM 3211 from the HRM 3220. Any write found to be addressed to LBA 0 will be parsed immediately (synchronously), as per the flag instructing that action. Subsequent action depends on the nature of the update.

If a partition is being added, and the partition is of a recognized type, the first
 25 LBA of the new partition will be added to the BLT 3212 and flagged as a location whose updates always require immediate parsing. The DST 3215 will be purged of any LBA ranges that fall within the new partition, in anticipation of there soon being a bitmap with a series of updates that will drive cluster reclamation. One concern is that, if the partition were ever to be written out ahead of the partition table update, then this information is
 30 potentially being written to blocks in the DST 3215, and could be reclaimed incorrectly by the scavenger thread. This could be alleviated, for example, by checking every write

received for coincidence with the ranges in the DST 3215 and removing any written-to block from the DST 3215.

If the partition is being updated to change the partition identifier from the Windows default to NTFS, then the BBUM 3211 will immediately re-examine the LUN
5 at the location of the partition boot sector, as the identifier change tends to occur after the partition boot sector has been written. This is really just part of partition addition.

If an existing partition is being deleted, the BLT will be flushed of records pertaining to the deleted partition, the BBUB 3213 for that partition will be deleted, and the LBA range will be added to the DST 3215 for asynchronous reclamation of clusters.

10 If an existing partition is being relocated, the existing boot sector record in the BLT 3212 will be updated with the new boot sector LBA to monitor. There is potential for the LUN to be immediately re-examined at the new location in case it has already been written, but this is not generally done.

If an existing partition is being truncated, the excised LBA range will be added to
15 the DST 3215. There is potential for the LUN to be immediately re-examined at the location of the partition boot sector in case the new boot sector has already been written, but this is not generally done.

If an existing partition is being enlarged, the DST 3215 will be purged of any LBA ranges that fall within the new partition. There is potential for the LUN to be
20 immediately re-examined at the location of the partition boot sector in case the new boot sector has already been written, but this is not generally done.

Any write found to be addressed to the first LBA of a partition will be parsed immediately (synchronously), as per the flag instructing that action. The starting LBA of the bitmap record will be determined and added to the BLT 3212, and flagged as a
25 location whose updates always require immediate parsing.

FIG. 34 is high-level pseudo code for the BBUM 3211, in accordance with an exemplary embodiment of the present invention. When the BBUM 3211 receives a client request, it gets the LUN from the ClientRequest and finds the right BLT based on the LUN. The BBUM 3211 gets the LBA from the ClientRequest, looks for this LBA in the
30 BLT, and checks the "immediate action" field to see if immediate action is required for this LBA. If immediate action is required, then the BBUM 3211 processes the client

request synchronously. If, however, immediate action is not required, then the BBUM 3211 sets the BBUB bit corresponding to the LBA for asynchronous processing.

FIG. 35 is high-level pseudo code for synchronous processing of an LBA 0 update creating a new partition, in accordance with an exemplary embodiment of the present invention. Specifically, if immediate action is required and the block is the partition table, then the BBUM 3211 compares partitions in new data with partitions in BLT. If a new partition is being added, then the BBUM 3211 gets the start and end of partition from the new data, checks the DSTs 3215 for any overlapping LBA ranges and remove them, adds the start of partition to the BLT, and flags the entry for immediate action.

FIG. 36 is high-level pseudo code for synchronous processing of an LBA 0 update (re)formatting a partition, in accordance with an exemplary embodiment of the present invention. Specifically, if immediate action is required and the block is a partition boot sector, then the BBUM 3211 gets the start of the MFT from the new data and calculates the location of the bitmap record. If there is already an identical bitmap record entry in the BLT for this partition then nothing is required. If, however, the bitmap record is at a different location from the BLT version, then the BBUM 3211 updates the BLT and reads the new location from the disk. If that location does not look like a bitmap record (i.e., it does not have a \$bitmap string), then nothing is required. If, however, the location does look like a bitmap record, then the BBUM 3211 gets the new bitmap location(s) and compares them with the BLT. If the new bitmap location(s) are identical, then nothing is required. If the new bitmaps are at a different location, then the BBUM 3211 sets all BBUB bits, updates the BBUB mappings, and moves the LBA ranges in the BLT. If the new bitmap is smaller than the existing bitmap, then the BBUM 3211 contracts the BBUB, adds the unmapped LBA range into the DST, and contracts the LBA range in the BLT. If the new bitmap is bigger than the existing bitmap, then the BBUM 3211 sets all the additional BBUB bits, enlarges the BBUB, and enlarges the LBA range in the BLT.

FIG. 37 is high-level pseudo code for synchronous processing of an LBA 0 update deleting a partition, in accordance with an exemplary embodiment of the present invention. Specifically, if immediate action is required and the block is a partition table, then the BBUM 3211 compares partitions in new data with partitions in BLT. If a

partition is being deleted, then the BBUM 3211 deletes the BBUB, deletes the boot sector from the BLT, deletes the Bitmap record from the BLT, deletes the Bitmap ranges from the BLT, and adds the partition range to the DST.

FIG. 38 is high-level pseudo code for the asynchronous task 3214, in accordance with an exemplary embodiment of the present invention. The asynchronous task 3214 parses the BBUB and then, for each bit set in the BBUB, the asynchronous task 3214 checks whether the corresponding cluster is marked unused by the host filesystem. If the cluster is marked unused by the host filesystem, then the asynchronous task 3214 checks whether the cluster is marked used by the storage controller. If the cluster is marked used by the storage controller, then the asynchronous task 3214 adds the LBA range to the DST. The asynchronous task 3214 also reclaims the storage space for each LBA range in the DST.

After receiving a boot sector update, it is generally not sufficient to wait for the write of the bitmap record (it is generally not know what order an NTFS format occurs in, and it could change in a minor patch anyway), since the bitmap record may already have been written to disk. If the bitmap record is written before the extended BPB, the BBUM 3211 will not catch it because the location is not present in the BLT 3212; an exception to this is when the location of the bitmap record has not changed. The exception notwithstanding, the BBUM 3211 generally has to immediately read the bitmap record location from the disk at this point to see if the bitmap record is present, and it generally needs to be able to distinguish random noise from an initialized bitmap record (checking for the \$bitmap Unicode string is a possibility). If it has not been written, it can wait for the write. If it is already on disk, it generally must be parsed immediately. Parsing generally requires that the record be decoded for the location(s) of the bitmap, and those locations are added to the BLT 3212 and flagged as not requiring immediate parsing. Parsing generally also requires that, if the size of the bitmap has changed, a new BBUB 3213 be instantiated, based on the new size and location(s) of bitmap; otherwise it is generally sufficient to update the existing BBUB 3213 with the new location. It also seems appropriate to set all the bits, as it is generally not know whether the host writes out the new bitmap before or after writing the bitmap record (the likelihood is after, but if that happens, the bits will just be harmlessly set a second time within a few seconds as

the bitmap writes come in). A danger is if the writes happen before, in which case the writes will have been missed due to being in a different location; setting all the bits ensures that the bitmap gets parsed.

In the case of a boot sector update (e.g., due to reformat of the partition), the
5 bitmap would likely be the same size and occupy the same place, so there generally would be no change to the BLT 3212 or BBUB 3213. The new bitmap would presumably be rewritten with most blocks being all zero, so the asynchronous task 3214 should be able to get on with processing them in order to reclaim the unallocated clusters from the CAT. The volume serial number of the boot sector could be checked to determine
10 whether the update was the result of a reformat.

The bitmap record could also be updated at any time, for reasons independent of the boot sector. The scavenger 3210 may have to be able to cope with the bitmap moving or changing size on the fly; it is not clear whether the bitmap could ever change in size without creating a different sized partition, but future versions of NTFS may support this
15 for whatever reason. In this situation, the new location(s) of the bitmap generally must be programmed into the BLT 3212, with the old entries removed and the new ones added. The BBUB 3213 has to be enlarged or contracted accordingly. Any LBA ranges freed up by a contraction can be added to the DST 3215, although strictly they still map to the partition.

20 Another concern is that, if the time of last update field of the bitmap record is frequently modified to reflect ongoing modification of the bitmap, the result could be a substantial amount of inline parsing.

All subsequent writes to the bitmap itself are pushed to the asynchronous task 3214 via the BBUBs 3213.

25 The basic strategy here is that all allocated clusters will be represented either in a BBUB 3213 or the DST 3215, and either way, unallocated ones will be reclaimed. An alternative solution would be to have a volume identifier for each volume, known to both the BBUM 3211 and the CAT such that each write would have to be mapped to a volume by the BBUM 3211 and tagged with that identifier before going to the CAT Manager,
30 where the volume identifier would be stored in the CAT record. A new volume would typically get a different identifier from the old volume it was overwriting, so the

asynchronous task 3214 could reclaim records with the old volume identifier without danger of reclaiming clusters that have been overwritten with data from the new volume. Obviously this would consume space in the CAT record. It is also dependent on the order of writes in a reformat. Since the system generally cannot know about a new volume
5 until the new volume serial number is seen in the boot partition, any other writes to the volume that precede this will be tagged with the old volume identifier.

The majority of the work will be done by the dedicated scavenger task 3214 that nominally wakes up once a minute, collects some work from the BBUB 3213, and executes it by paging-in bitmap blocks through the cache and comparing the bits with the
10 CAT. In an exemplary embodiment, the BBUB 3213 will be logically segmented (1k segment sizes), with a counter for each segment showing the number of updates for that segment, and a global counter that reflects the highest value held by any counter; these counters will be incremented by the work producer (the BBUM 3211) and decremented by the work consumer (the scavenger task 3214). The scavenger task 3214, on waking,
15 will check the global counter and decide whether the value therein is high enough to justify paging-in the bitmap. If it is, then the task 3214 will determine which segment that value corresponds to (e.g., by iterating through the counter array) and then begin iterating through the bits of the appropriate BBUB segment. When it finds a set bit, it will page-in that block of the bitmap and compare it with the CAT.

As discussed above, operation of the scavenger task 3214 could be dynamically adjusted, for example, by changing the frequency with which it runs or the thresholds at which it decides to do some work. In an exemplary embodiment of the present invention, however, such dynamic adjustment is generally not done because the architecture is somewhat coupled to the frequency with which the scavenger runs. Specifically, because
25 the proposed architecture has been designed with the assumption that the host filesystem 2711 will update its metadata within a maximum time window (e.g., one minute), the proposed architecture could not realistically run the scavenger task 3214 more frequently than that maximum time window. The proposed architecture could make the runs less frequently without actually breaking that rule, but it could then make cluster reclamation
30 less efficient by making some cluster updates look more recent than they actually are (e.g., if the age computations are designed with the assumption that the runs occur every

minute but they actually occur, say, every three minutes, then the age calculations could be off by a factor of three). Furthermore, task priority is generally fixed at compile time and therefore is generally not changed during system operation.

It should be noted that, in an exemplary embodiment of the present invention, the storage system implements clusters of size 4K. Thus, if the filesystem is formatted with a cluster size other than 4K, a bit in the filesystem bitmap would not correlate neatly with a cluster in the storage system. For example, if the filesystem cluster size is less than 4K, then multiple bits of the bitmap will generally have to be clear to bother cross-referencing with the CAT. If, however, the filesystem cluster size is greater than 4K, then one clear bit of the bitmap will generally require multiple lookups into the CAT, one for each 4K.

Another concern is how to handle situations in which the scavenger encounters a cluster that is too young to reclaim. In such situations, the scavenger could leave the bit set in the BBUB, thereby requiring one or more subsequent scans to parse through the whole 512 bits again (e.g., the next scan might go through the 512 bits only to find that the cluster is still too young to reclaim). Alternatively, the scavenger could clear the bit and add the cluster to a list of young blocks / bitoffsets that need rechecking. From an implementation standpoint, the latter approach would only be practical if the list could be kept fairly small.

From an implementation standpoint, the scavenger and BBUB will both read from disk through the CAT Manager. Cluster reclamation will be performed through a special API provided by the CAT Manager.

VIRTUAL HOT SPARE

As discussed above, in many storage systems, a hot spare storage device will be maintained in a ready state so that it can be brought online quickly in the event another storage device fails. In certain embodiments of the present invention, rather than maintaining a physically separate hot spare, a virtual hot spare is created from unused storage capacity across a plurality of storage devices. Unlike a physical hot spare, this

unused storage capacity is available if and when a storage device fails for storage of data recovered from the remaining storage device(s).

The virtual hot spare feature requires that enough space be available on the array to ensure that data can be re-laid out redundantly in the event of a disk failure. Thus, on an ongoing basis, the storage system typically determines the amount of unused storage capacity that would be required for implementation of a virtual hot spare (e.g., based on the number of storage devices, the capacities of the various storage devices, the amount of data stored, and the manner in which the data is stored) and generates a signal if additional storage capacity is needed for a virtual hot spare (e.g., using green/yellow/red lights to indicate status and slot, substantially as described above). As zones are allocated, a record is kept of how many regions are required to re-layout that zone on a per disk basis. The following table demonstrates a virtual hot spare with four drives used:

Zone	Type	Stored On Disks	Comments	Regions required if disk fails			
				Disk 0	Disk 1	Disk 2	Disk 3
2	Dual Drive Mirror	0, 1	Reconstruct on disk 2 or 3 if 0 or 1 fails	12	12	0	0
3	Dual Drive Mirror	0, 3	Reconstruct on disk 1 or 2 if 1 or 2 fails	12	0	0	12
5	Triple Drive Stripe	1, 2, 3	Reconstruct on disk 0 if 1, 2, or 3 fails	0	6	6	6
10	Four Drive Stripe	0, 1, 2, 3	Convert to triple drive strip across other three disks	2, 2, 2	2, 2, 2	2, 2, 2	2, 2, 2

The following table demonstrates a virtual hot spare with three drives used:

Zone	Type	Stored On Disks	Comments	Regions required if disk fails
------	------	-----------------	----------	--------------------------------

				Disk 0	Disk 1	Disk 2
2	Dual Drive Mirror	0, 1	Reconstruct on disk 3	12	12	0
3	Dual Drive Mirror	0, 3	Reconstruct on disk 1	12	0	12
5	Triple Drive Stripe	1, 2, 3	Convert to dual drive mirror	6, 6	6, 6	6, 6

In this exemplary embodiment, virtual hot spare is not available on an array with only 1 or 2 drives. Based on the information for each zone and the number of disks in the array, the array determines a re-layout scenario for each possible disk failure and ensure
5 that enough space is available on each drive for each scenario. The information generated can be fed back into the re-layout engine and the zone manager so that the data can be correctly balanced between the data storage and the hot spare feature. Note that the hot spare feature requires enough spare working space regions on top of those calculated from the zone layout data so that re-layout can occur.

Fig. 22 is a logic flow diagram showing exemplary logic for managing a virtual
10 hot spare in accordance with an exemplary embodiment of the present invention. In block 2102, the logic determines a re-layout scenario for each possible disk failure. In block 2104, the logic determines the amount of space needed on each drive for re-layout of data redundantly in a worst case scenario. In block 2106, the logic determines the
15 amount of spare working space regions needed for re-layout of data redundantly in a worst case scenario. In block 2108, the logic determines the total amount of space needed on each drive in order to permit re-layout of data redundantly in a worst case scenario (essentially the sum of the amount of space needed for re-layout and the amount of spare working space regions needed). In block 2110, the logic determines whether the
20 storage system contains an adequate amount of available storage. If there is an adequate amount of available storage (YES in block 2112), then the logic iteration terminates in block 2199. If, however, there is an inadequate amount of available storage (NO in block 2112), then the logic determines which drive/slot requires upgrade, in block 2114. Then,

in block 2116, the logic signals that additional storage space is needed and indicates which drive/slot requires upgrade. The logic iteration terminates in block 2199.

Fig. 23 is a logic flow diagram showing exemplary logic for determining a re-layout scenario for each possible disk failure, as in block 2102 of Fig. 22, in accordance
5 with an exemplary embodiment of the present invention. In block 2202, the logic allocates a zone. Then, in block 2204, the logic determines how many regions are required to re-layout that zone on a per-disk basis. The logic iteration terminates in block 2299.

Fig. 24 is a logic flow diagram showing exemplary logic for invoking the virtual
10 hot spare functionality in accordance with an exemplary embodiment of the present invention. In block 2302, the logic maintains a sufficient amount of available storage to permit re-layout of data redundantly in the event of a worst case scenario. Upon determining loss of a drive (e.g., removal or failure), in block 2304, the logic automatically reconfigures the one or more remaining drives to restore fault tolerance for
15 the data, in block 2306. The logic iteration terminates in block 2399.

Fig. 25 is a logic flow diagram showing exemplary logic for automatically reconfiguring the one or more remaining drives to restore fault tolerance for the data, as
in block 2306 of Fig. 24, in accordance with an exemplary embodiment of the present invention. In block 2402, the logic may convert a first striped pattern across four or more
20 storage devices to a second striped pattern across three or more remaining storage devices. In block 2404, the logic may convert a striped pattern across three storage devices to a mirrored pattern across two remaining storage devices. Of course, the logic may convert patterns in other ways in order to re-layout the data redundantly following loss of a drive. The logic iteration terminates in block 2499.

25 With reference again to FIG. 21, the storage manager 2504 typically includes appropriate components and logic for implementing the virtual hot spare functionality as described above.

DYNAMIC UPGRADE

30

The logic described above for handling dynamic expansion and contraction of storage can be extended to provide a dynamically upgradeable storage system in which storage devices can be replaced with a larger storage devices as needed, and existing data is automatically reconfigured across the storage devices in such a way that redundancy is maintained or enhanced and the additional storage space provided by the larger storage devices will be included in the pool of available storage space across the plurality of storage devices. Thus, when a smaller storage device is replaced by a larger storage device, the additional storage space can be used to improve redundancy for already stored data as well as to store additional data. Whenever more storage space is needed, an appropriate signal is provided to the user (e.g., using green/yellow/red lights substantially as described above), and the user can simply remove a storage device and replace it with a larger storage device.

Fig. 26 is a logic flow diagram showing exemplary logic for upgrading a storage device, in accordance with an exemplary embodiment of the present invention. In block 2602, the logic stores data on a first storage device in a manner that the data stored thereon appears redundantly on other storage devices. In block 2604, the logic detects replacement of the first storage device with a replacement device having greater storage capacity than the first storage device. In block 2606, the logic automatically reproduces the data that was stored on the first device onto the replacement device using the data stored redundantly on other devices. In block 2608, the logic makes the additional storage space on the replacement device available for storing new data redundantly. In block 2610, the logic may store new data redundantly within the additional storage space on the replacement device if no other device has a sufficient amount of available storage capacity to provide redundancy for the new data. In block 2612, the logic may store new data redundantly across multiple storage devices if at least one other device has a sufficient amount of available storage capacity to provide redundancy for the new data.

With reference again to FIG. 21, the storage manager 2504 typically includes appropriate components and logic for implementing the dynamic upgrade functionality as described above.

30

MISCELLANEOUS

Embodiments of the present invention may be employed to provide storage capacity to a host computer, e.g., using a peripheral connect protocol in the manner described in my United States Provisional Patent Application No. 60/625,495, which was
5 filed on November 5, 2004 in the name of Geoffrey S. Barrall, and is hereby incorporated herein by reference in its entirety.

It should be noted that a hash algorithm may not produce hash values that are strictly unique. Thus, is it conceivable for the hash algorithm to generate the same hash value for two chunks of data having non-identical content. The hash function (which
10 generally incorporates the hash algorithm) typically includes a mechanism for confirming uniqueness. For example, in an exemplary embodiment of the invention as described above, if the hash value for one chunk is different than the hash value of another chunk, then the content of those chunks are considered to be non-identical. If, however, the hash value for one chunk is the same as the hash value of another chunk, then the hash
15 function might compare the contents of the two chunks or utilize some other mechanism (e.g., a different hash function) to determine whether the contents are identical or non-identical.

It should be noted that the logic flow diagrams are used herein to demonstrate various aspects of the invention, and should not be construed to limit the present
20 invention to any particular logic flow or logic implementation. The described logic may be partitioned into different logic blocks (e.g., programs, modules, functions, or subroutines) without changing the overall results or otherwise departing from the true scope of the invention. Often times, logic elements may be added, modified, omitted, performed in a different order, or implemented using different logic constructs (e.g., logic
25 gates, looping primitives, conditional logic, and other logic constructs) without changing the overall results or otherwise departing from the true scope of the invention.

The present invention may be embodied in many different forms, including, but in no way limited to, computer program logic for use with a processor (e.g., a microprocessor, microcontroller, digital signal processor, or general purpose computer),
30 programmable logic for use with a programmable logic device (e.g., a Field Programmable Gate Array (FPGA) or other PLD), discrete components, integrated

circuitry (*e.g.*, an Application Specific Integrated Circuit (ASIC)), or any other means including any combination thereof.

Computer program logic implementing all or part of the functionality previously described herein may be embodied in various forms, including, but in no way limited to, a source code form, a computer executable form, and various intermediate forms (*e.g.*, forms generated by an assembler, compiler, linker, or locator). Source code may include a series of computer program instructions implemented in any of various programming languages (*e.g.*, an object code, an assembly language, or a high-level language such as Fortran, C, C++, JAVA, or HTML) for use with various operating systems or operating environments. The source code may define and use various data structures and communication messages. The source code may be in a computer executable form (*e.g.*, via an interpreter), or the source code may be converted (*e.g.*, via a translator, assembler, or compiler) into a computer executable form.

The computer program may be fixed in any form (*e.g.*, source code form, computer executable form, or an intermediate form) either permanently or transitorily in a tangible storage medium, such as a semiconductor memory device (*e.g.*, a RAM, ROM, PROM, EEPROM, or Flash-Programmable RAM), a magnetic memory device (*e.g.*, a diskette or fixed disk), an optical memory device (*e.g.*, a CD-ROM), a PC card (*e.g.*, PCMCIA card), or other memory device. The computer program may be fixed in any form in a signal that is transmittable to a computer using any of various communication technologies, including, but in no way limited to, analog technologies, digital technologies, optical technologies, wireless technologies (*e.g.*, Bluetooth), networking technologies, and internetworking technologies. The computer program may be distributed in any form as a removable storage medium with accompanying printed or electronic documentation (*e.g.*, shrink wrapped software), preloaded with a computer system (*e.g.*, on system ROM or fixed disk), or distributed from a server or electronic bulletin board over the communication system (*e.g.*, the Internet or World Wide Web).

Hardware logic (including programmable logic for use with a programmable logic device) implementing all or part of the functionality previously described herein may be designed using traditional manual methods, or may be designed, captured, simulated, or documented electronically using various tools, such as Computer Aided Design (CAD), a

hardware description language (*e.g.*, VHDL or AHDL), or a PLD programming language (*e.g.*, PALASM, ABEL, or CUPL).

Programmable logic may be fixed either permanently or transitorily in a tangible storage medium, such as a semiconductor memory device (*e.g.*, a RAM, ROM, PROM, 5 EEPROM, or Flash-Programmable RAM), a magnetic memory device (*e.g.*, a diskette or fixed disk), an optical memory device (*e.g.*, a CD-ROM), or other memory device. The programmable logic may be fixed in a signal that is transmittable to a computer using any of various communication technologies, including, but in no way limited to, analog technologies, digital technologies, optical technologies, wireless technologies (*e.g.*, 10 Bluetooth), networking technologies, and internetworking technologies. The programmable logic may be distributed as a removable storage medium with accompanying printed or electronic documentation (*e.g.*, shrink wrapped software), preloaded with a computer system (*e.g.*, on system ROM or fixed disk), or distributed from a server or electronic bulletin board over the communication system (*e.g.*, the 15 Internet or World Wide Web).

This application is related to the following United States Patent Applications, which are hereby incorporated herein by reference in their entireties:

Attorney Docket No. 2950/104 entitled **Dynamically Upgradeable Fault-Tolerant Storage System Permitting Variously Sized Storage Devices and Method;**

20 Attorney Docket No. 2950/105 entitled **Dynamically Expandable and Contractible Fault-Tolerant Storage System With Virtual Hot Spare;** and

Attorney Docket No. 2950/107 entitled **Storage System Condition Indicator and Method.**

The present invention may be embodied in other specific forms without departing 25 from the true scope of the invention. The described embodiments are to be considered in all respects only as illustrative and not restrictive.

What is claimed is:

1. A method of storing data in by a block-level storage system that stores data under control of a host filesystem, the method comprising:
 - locating host filesystem data structures stored for the host filesystem in the block-level storage system;
 - analyzing the host filesystem data structures to identify a data type associated with the data to be stored; and
 - storing the data using a storage scheme selected based on the data type, whereby data having different data types can be stored using different storage schemes selected based on the data types.
2. A method according to claim 1, wherein storing the data using a storage scheme selected based on the data type comprises:
 - storing the data using a storage layout selected based on the data type.
3. A method according to claim 2, wherein storing the data using a storage layout selected based on the data type comprises:
 - storing frequently accessed data so as to provide enhanced accessibility.
4. A method according to claim 3, wherein storing frequently accessed data so as to provide enhanced accessibility comprises:
 - storing the frequently accessed data uncompressed and in sequential storage.
5. A method according to claim 2, wherein storing the data using a storage layout selected based on the data type comprises:
 - storing infrequently accessed data so as to provide enhanced storage efficiency.
6. A method according to claim 5, wherein storing infrequently accessed data so as to provide enhanced storage efficiency comprises:

storing the infrequently accessed data using at least one of data compression and non-sequential storage.

7. A method according to claim 1, wherein storing the data using a storage scheme selected based on the data type comprises:

storing the data using an encoding scheme selected based on the data type.

8. A method according to claim 7, wherein the encoding scheme comprises at least one of:

data compression; and
encryption.

9. A method according to claim 1, wherein locating host filesystem data structures in storage comprises:

maintaining a partition table;
parsing the partition table to locate an operating system partition;
parsing the operating system partition to identify the operating system and locate operating system data structures; and
parsing the operating system data structures to identify the host filesystem and locate the host filesystem data structures.

10. A method according to claim 9, wherein the operating system data structures include a superblock, and wherein parsing the operating system data structures includes parsing the superblock.

11. A method according to claim 9, wherein parsing the host filesystem data structures comprises:

making a working copy of a host filesystem data structure; and
parsing the working copy.

12. A block-level storage system that stores data under control of a host filesystem, the system comprising:

a block-level storage in which host filesystem data structures are stored for the host filesystem;

a storage controller operably coupled to the block-level storage for locating the host filesystem data structures stored in the block-level storage, analyzing the host filesystem data structures to identify a data type associated with the data to be stored, and storing the data using a storage scheme selected based on the data type, whereby data having different data types can be stored using different storage schemes selected based on the data types.

13. A system according to claim 12, wherein the storage controller is operably coupled to store the data using a storage layout selected based on the data type.

14. A system according to claim 13, wherein the storage controller is operably coupled to store frequently accessed data so as to provide enhanced accessibility.

15. A system according to claim 14, wherein the storage controller is operably coupled to store frequently accessed data uncompressed and in sequential storage.

16. A system according to claim 13, wherein the storage controller is operably coupled to store infrequently accessed data so as to provide enhanced storage efficiency.

17. A system according to claim 16, wherein the storage controller is operably coupled to store infrequently accessed data using at least one of data compression and non-sequential storage.

18. A system according to claim 12, wherein the storage controller is operably coupled to store the data using an encoding scheme selected based on the data type.

19. A system according to claim 18, wherein the encoding scheme comprises at least one of:

data compression; and
encryption.

20. A system according to claim 12, wherein the storage controller is operably coupled to maintain a partition table, parse the partition table to locate an operating system partition, parse the operating system partition to identify the operating system and locate operating system data structures, parse the operating system data structures to identify the host filesystem and locate the host filesystem data structures, and parse the host filesystem data structures to identify the data type.

21. A system according to claim 20, wherein the operating system data structures include a superblock, and wherein the storage controller is operably coupled to parse the superblock.

22. A system according to claim 20, wherein the storage controller is operably coupled to make a working copy of a host filesystem data structure and parse the working copy.







FIG. 1

**FIG. 2A****FIG. 2B**







CHUNKS #1, #2, #3 STORED USING PATTERN 3 (STRIPING USING 3 DEVICES)


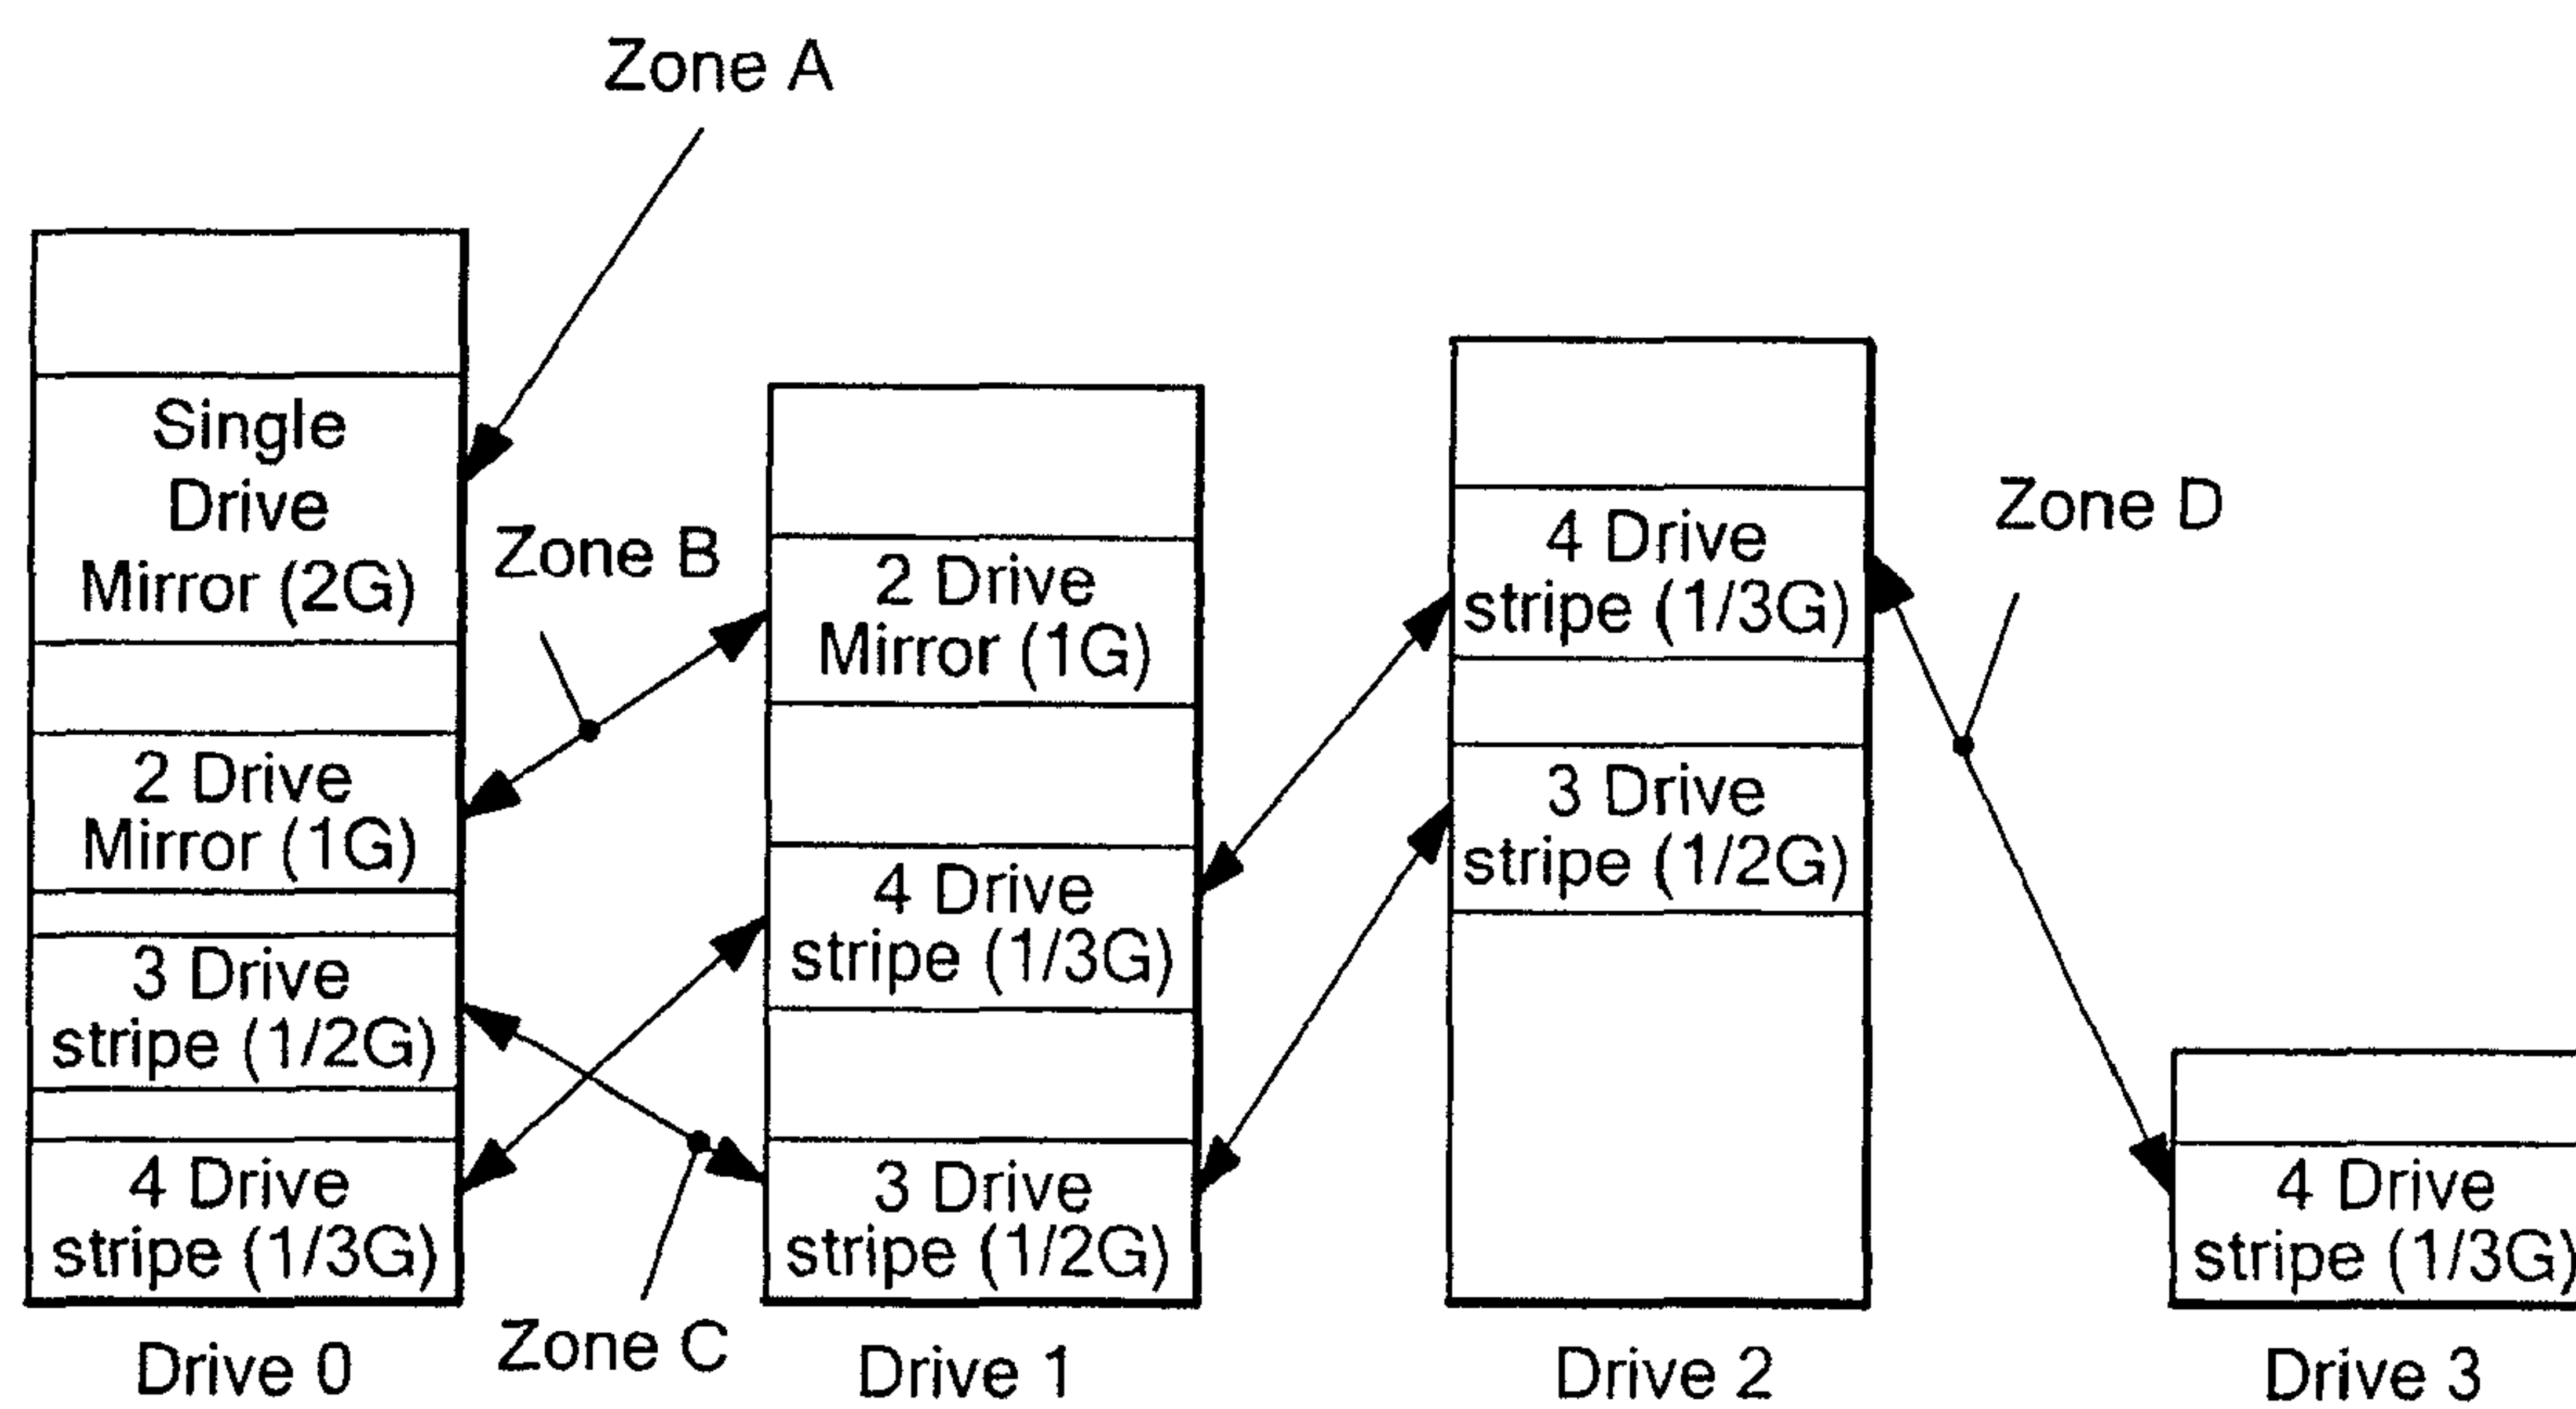
CHUNKS #10, #11, #12 STORED USING PATTERN 2 (MIRRORED USING 2 DEVICES)

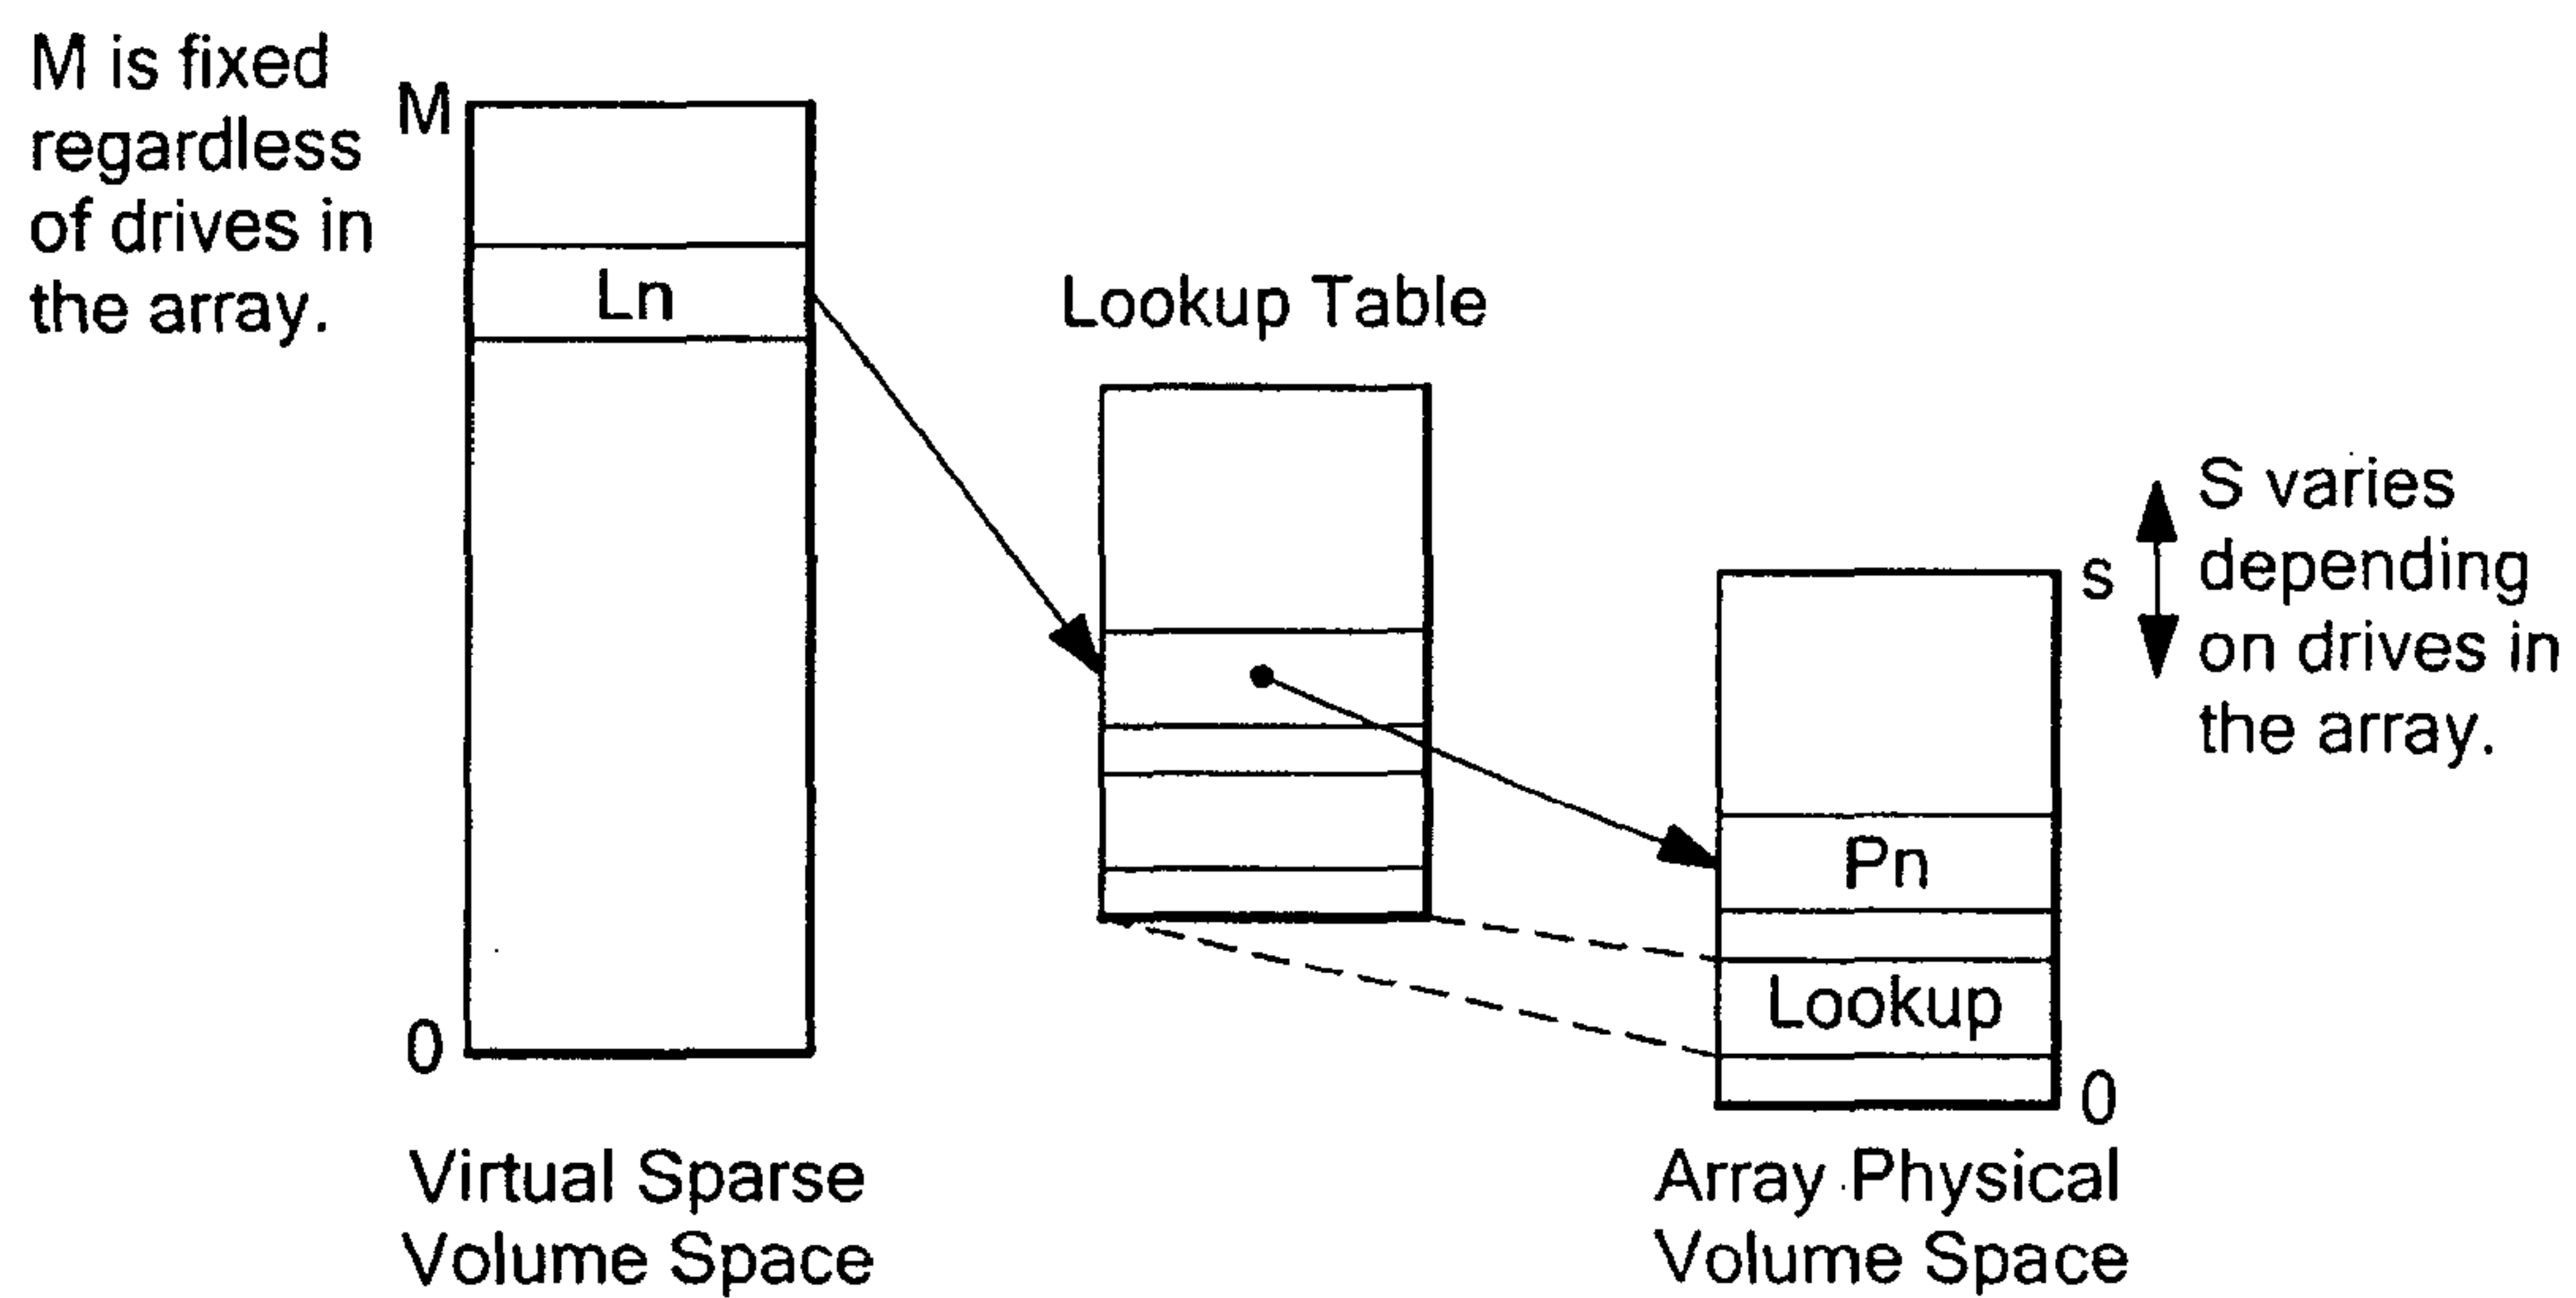

CHUNKS #20, #21 STORED USING PATTERN 1 (MIRRORED USING 1 DEVICE)

FIG. 3




**FIG. 5**

**FIG. 6****FIG. 7**

**FIG. 8**


a) All Drives OK with Array space available.

FIG. 9




b) Slot C & Slot D full - data not redundant: Add drive to empty Slot A

FIG. 10



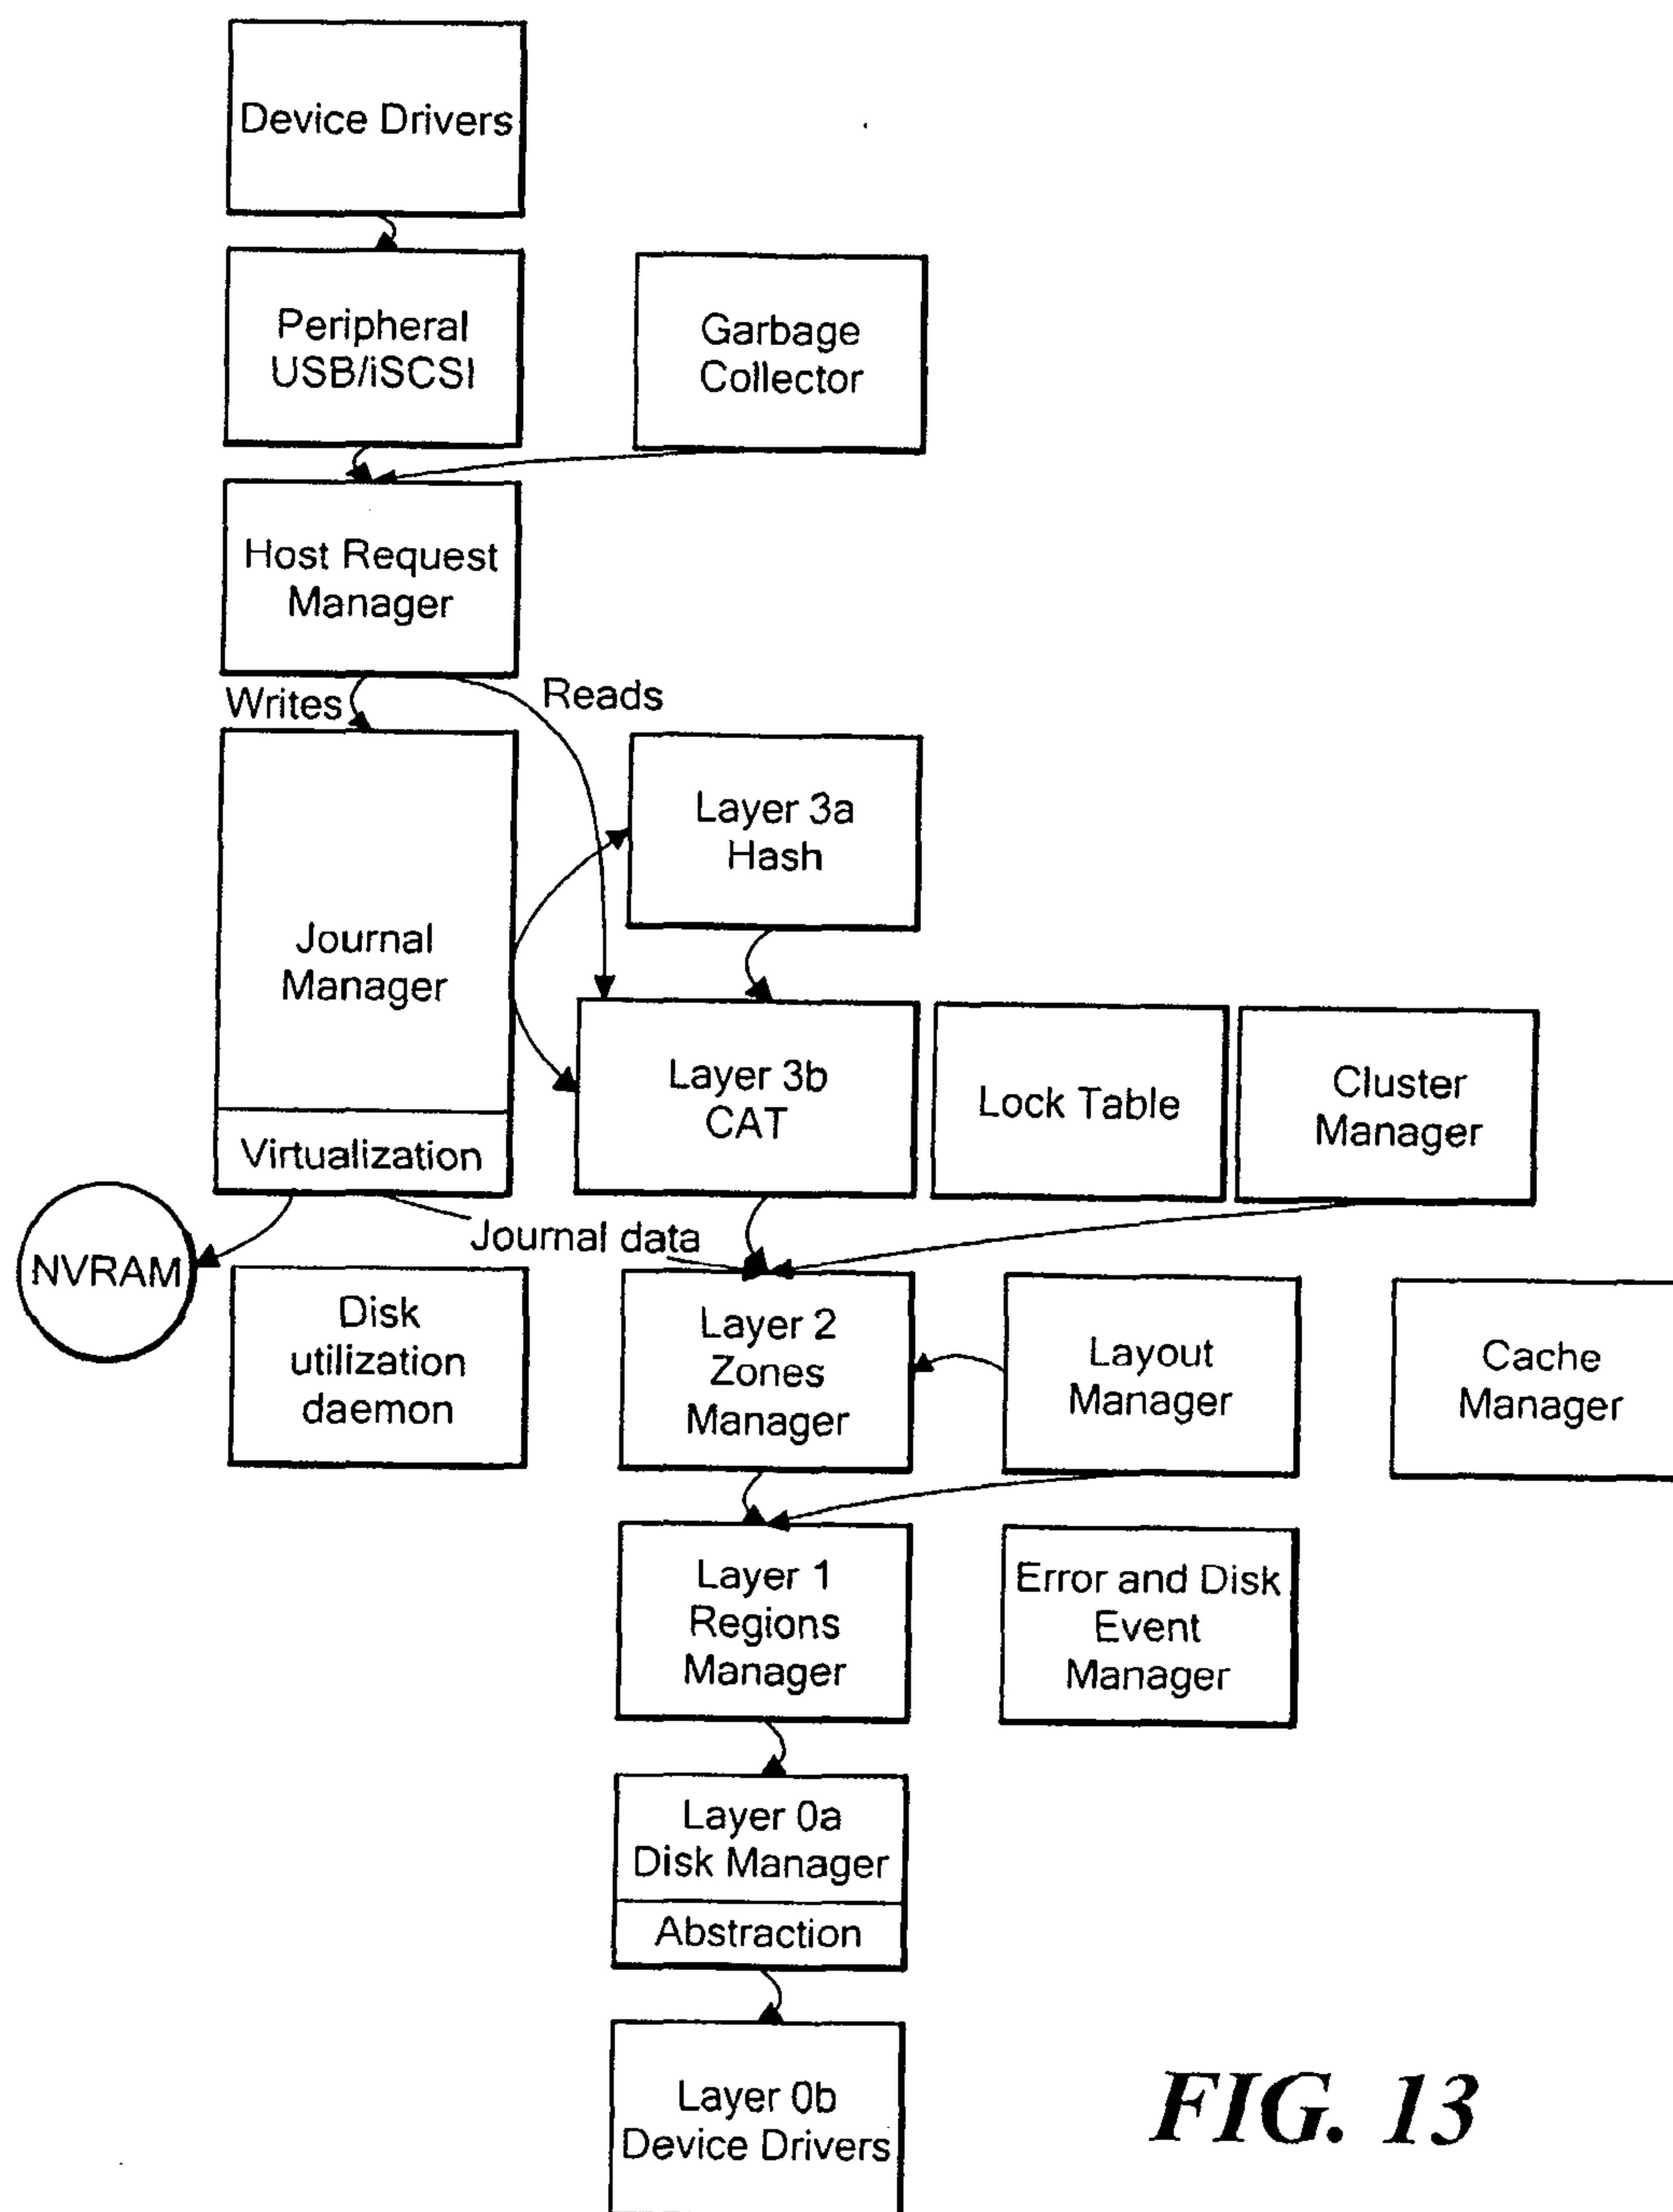
c) Array cannot maintain redundant data in the event of a power failure:
Replace drive in Slot D


FIG. 11

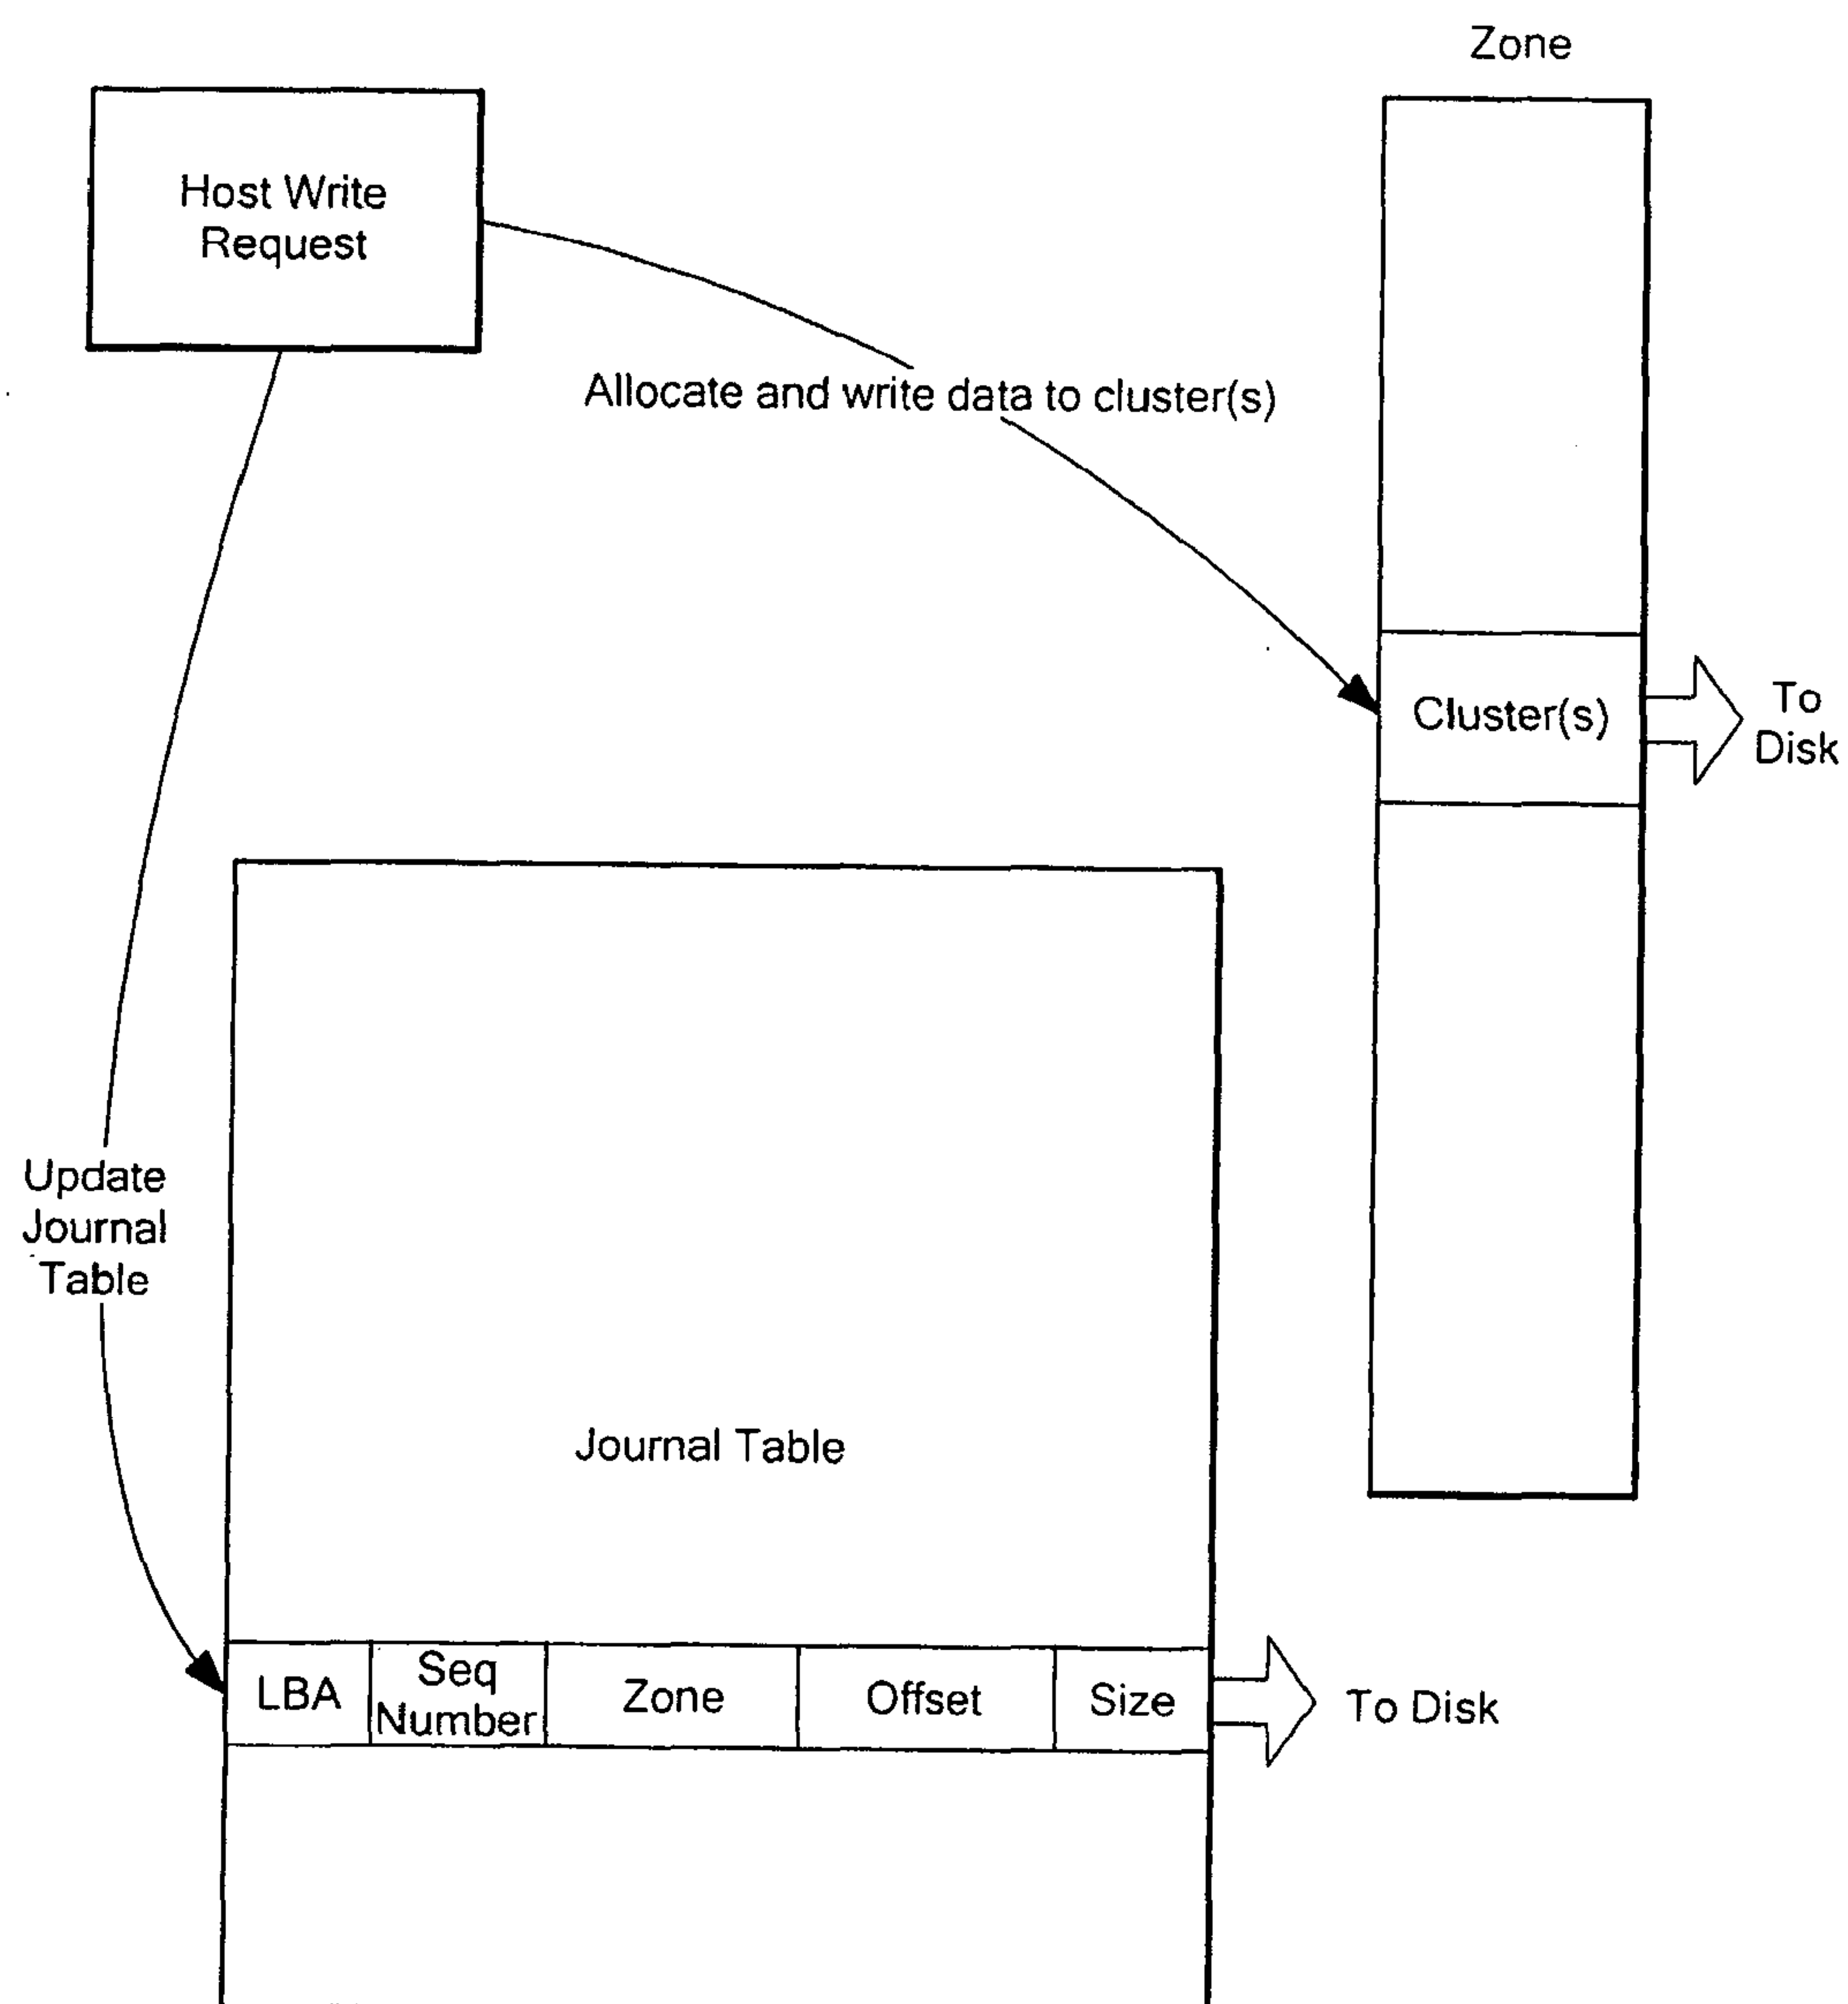



d) Drive in Slot C has failed


FIG. 12

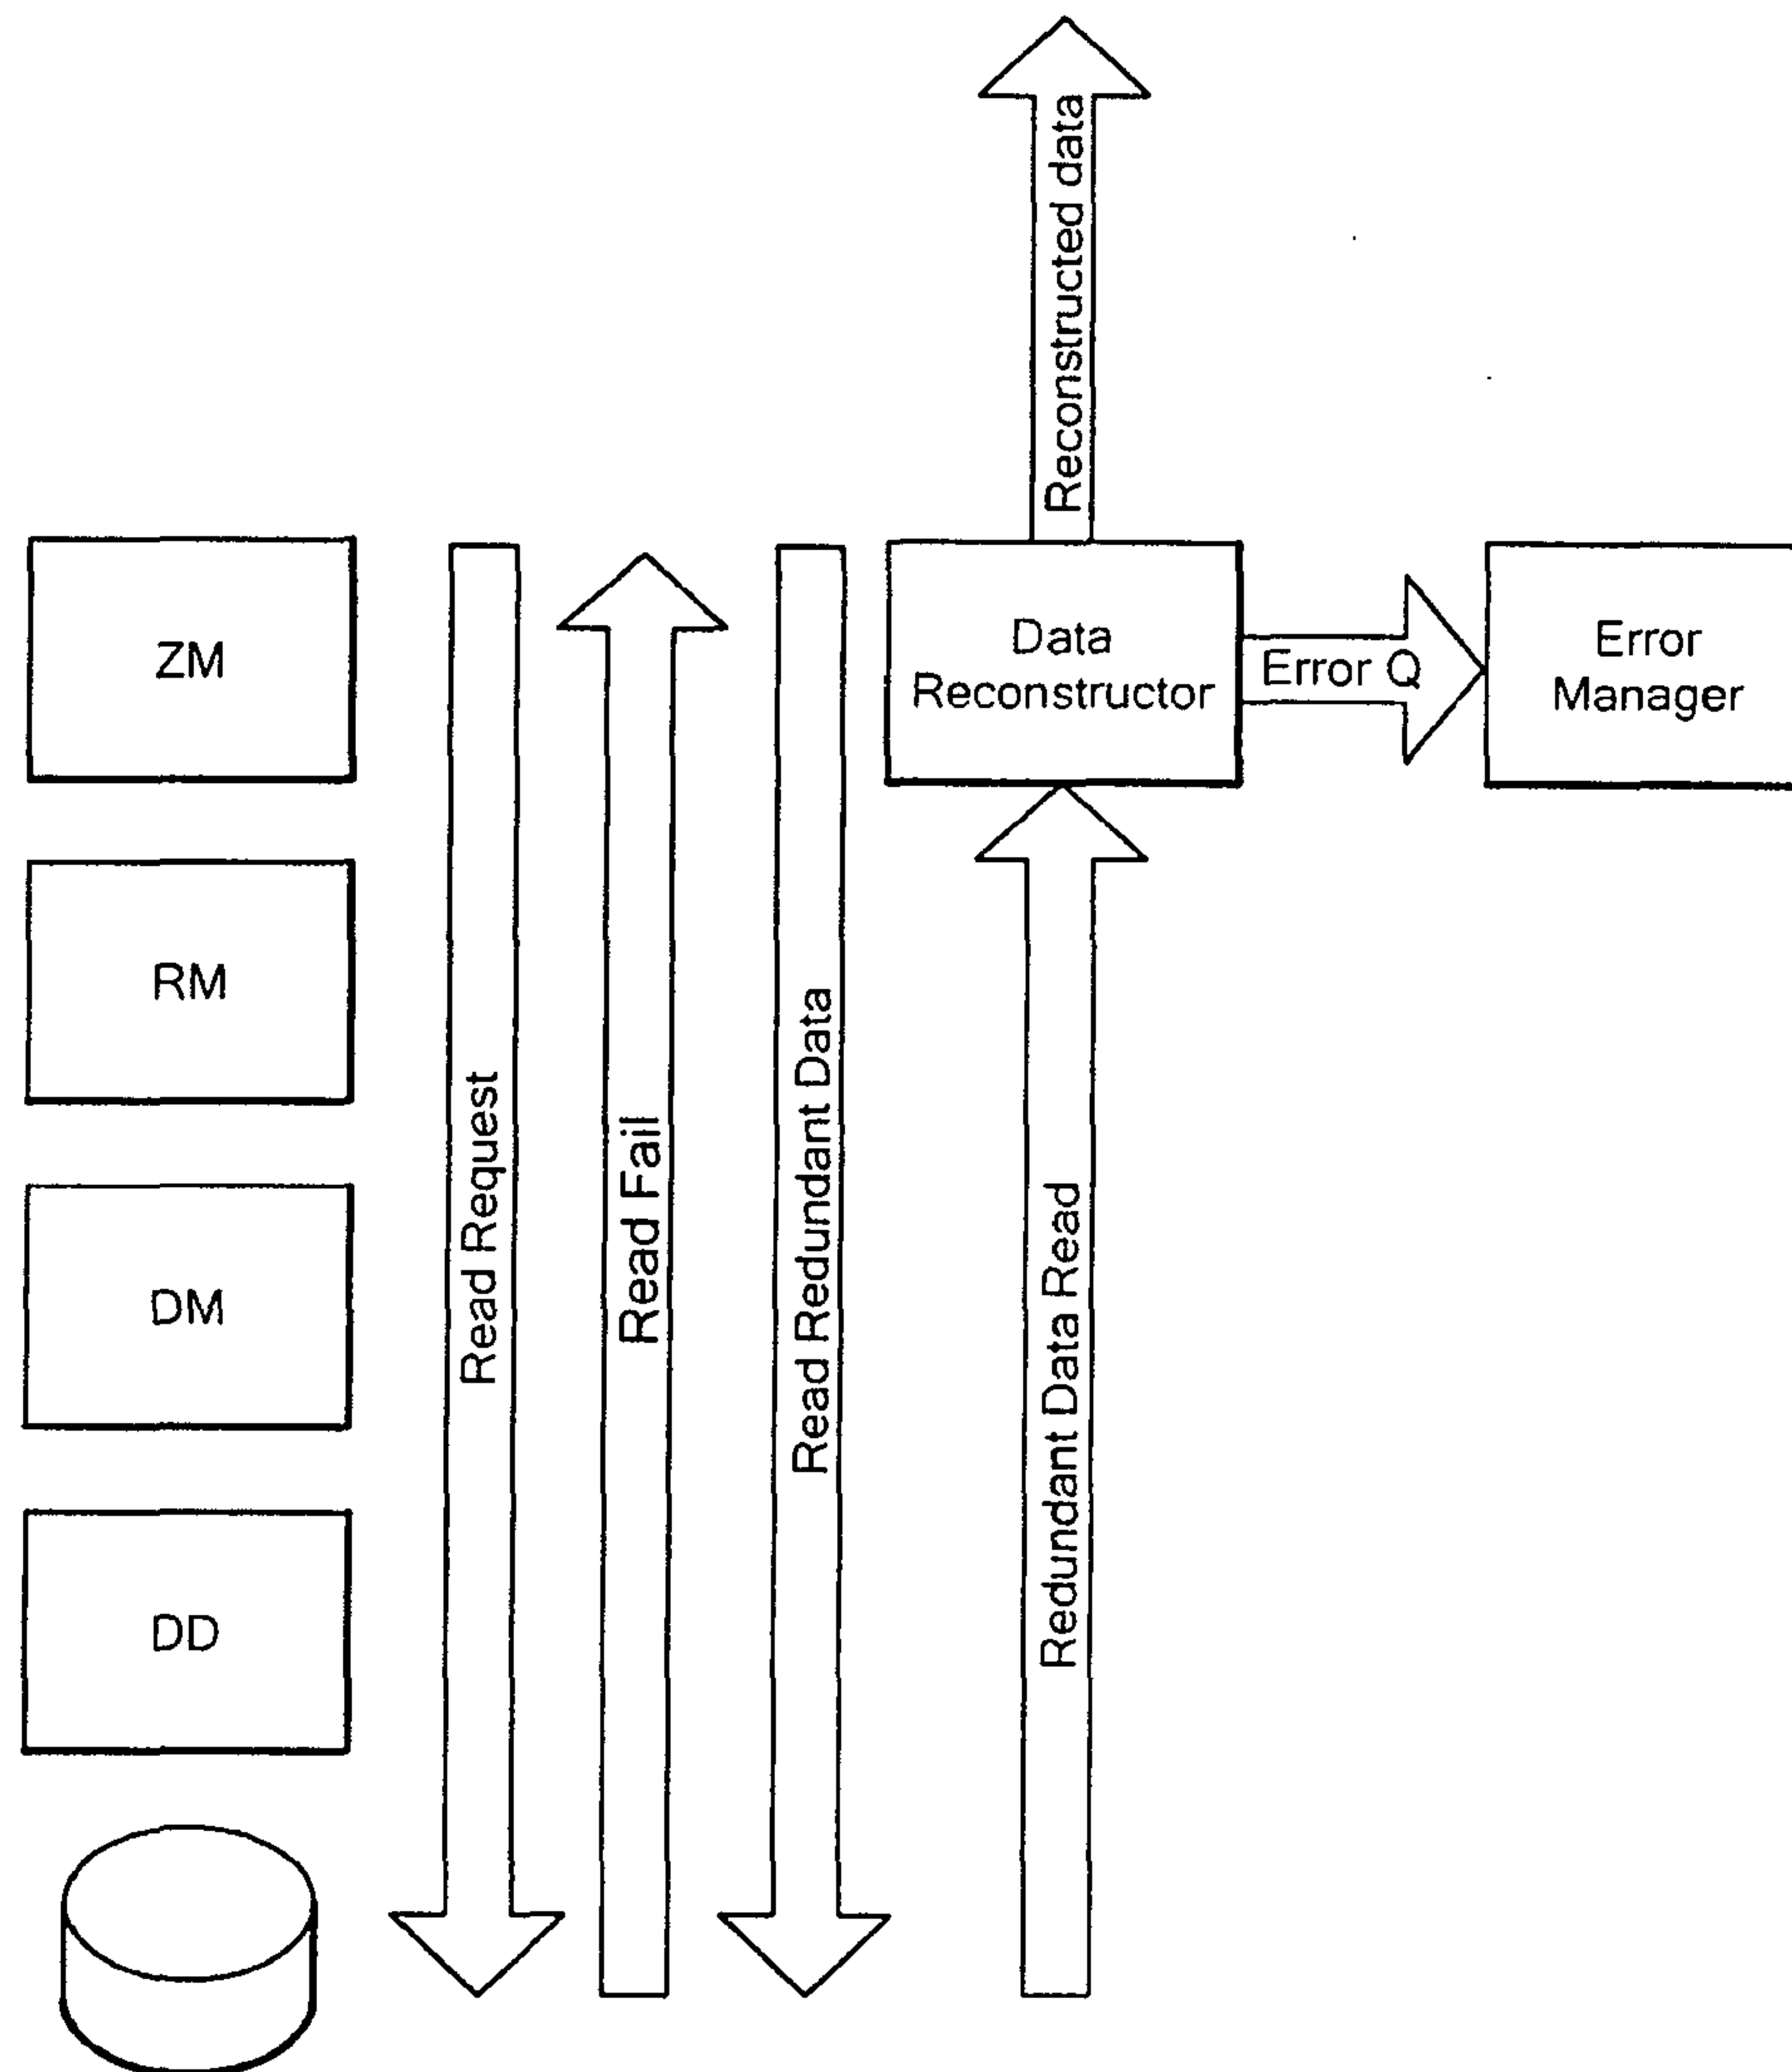
**FIG. 13**


**FIG. 14**

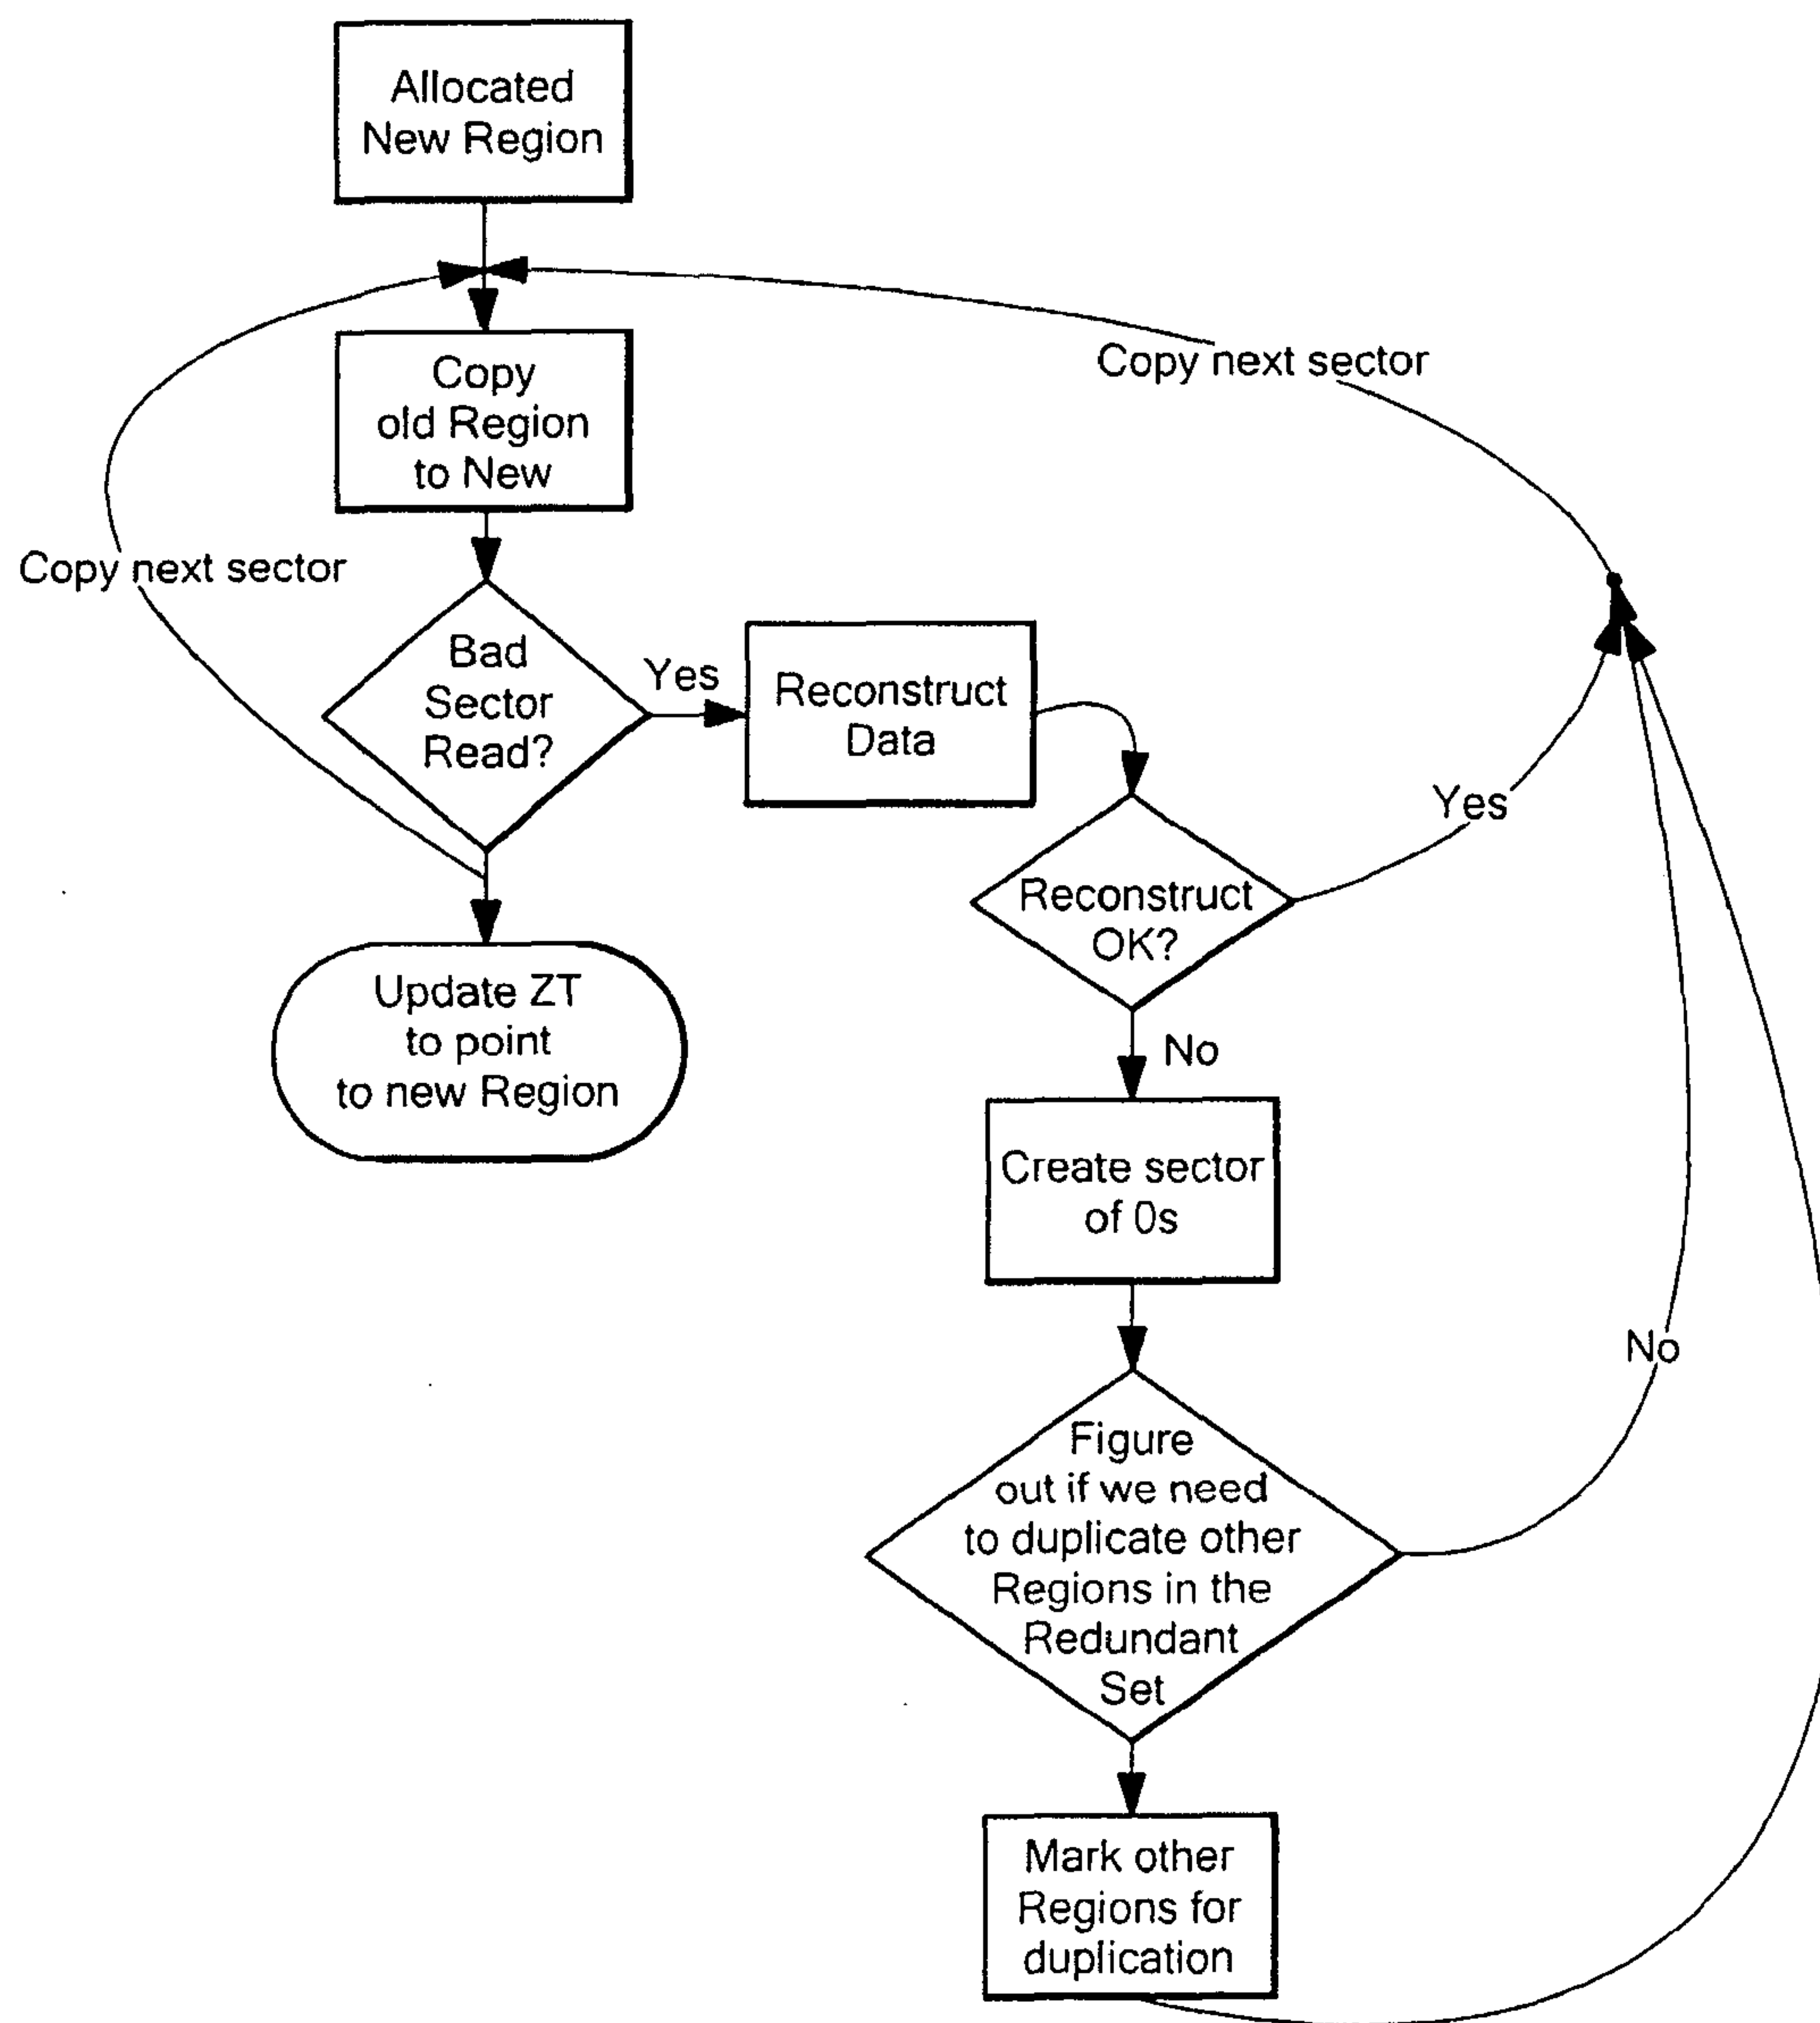
**FIG. 15**

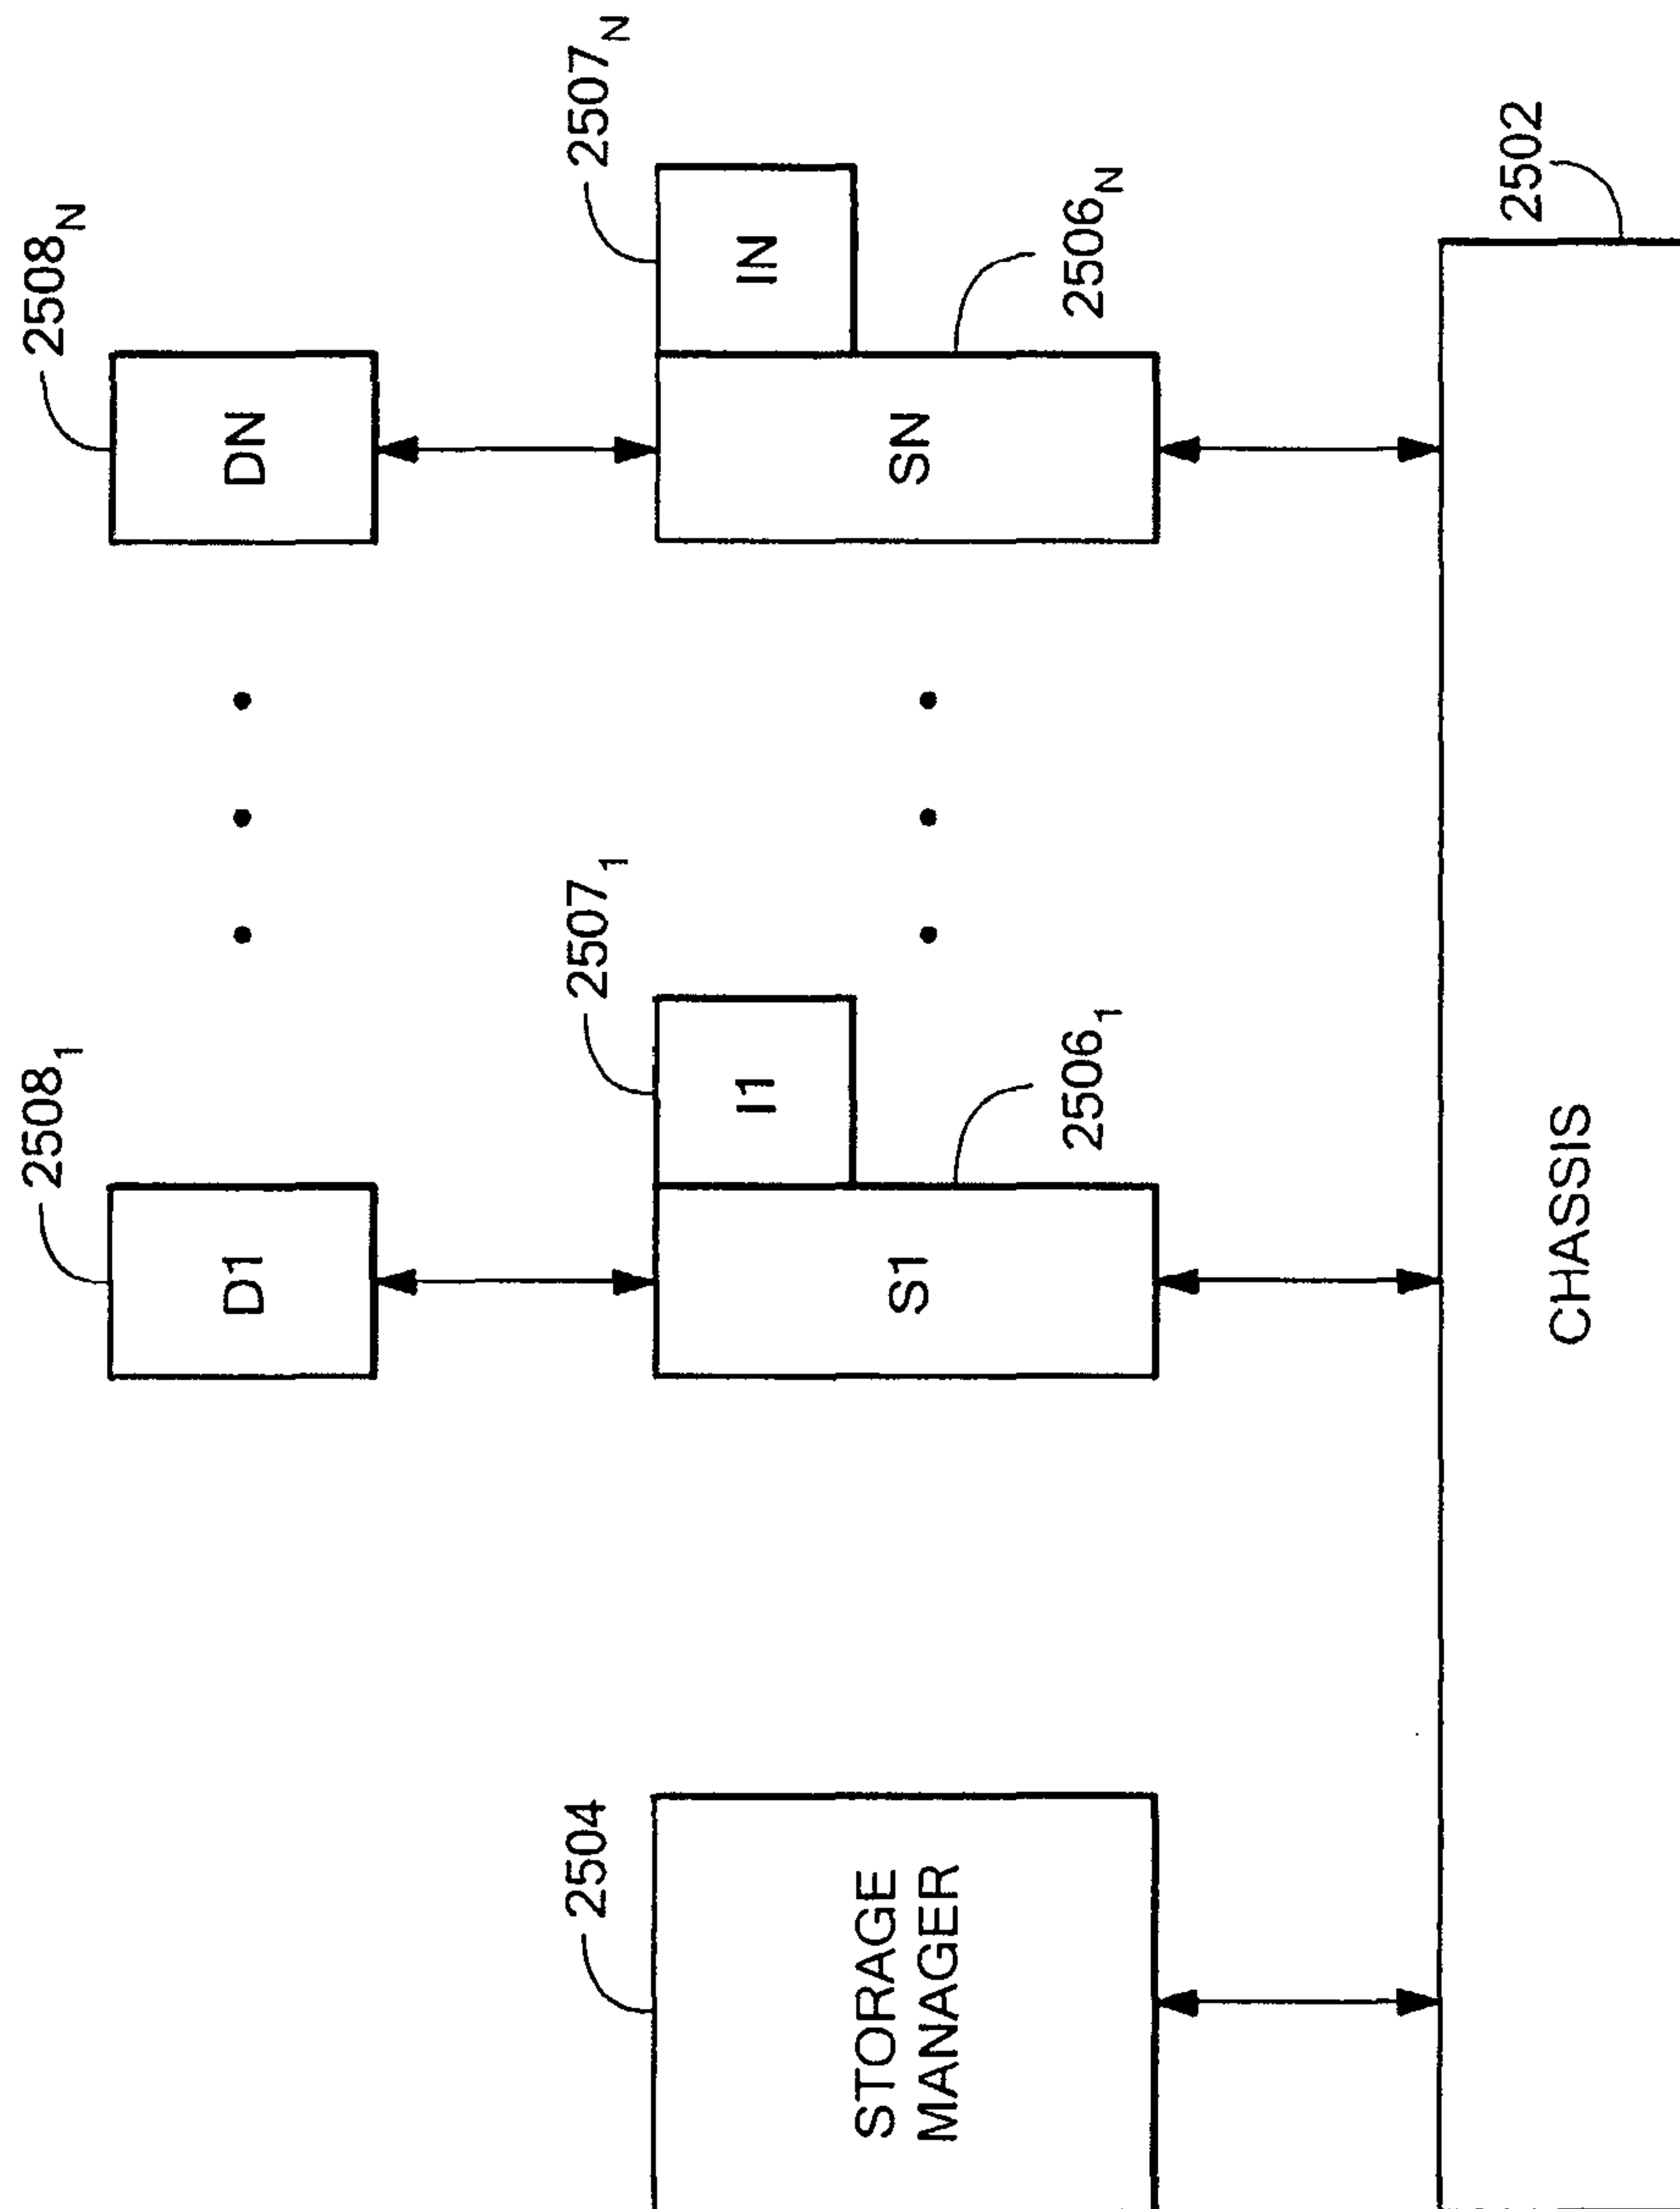


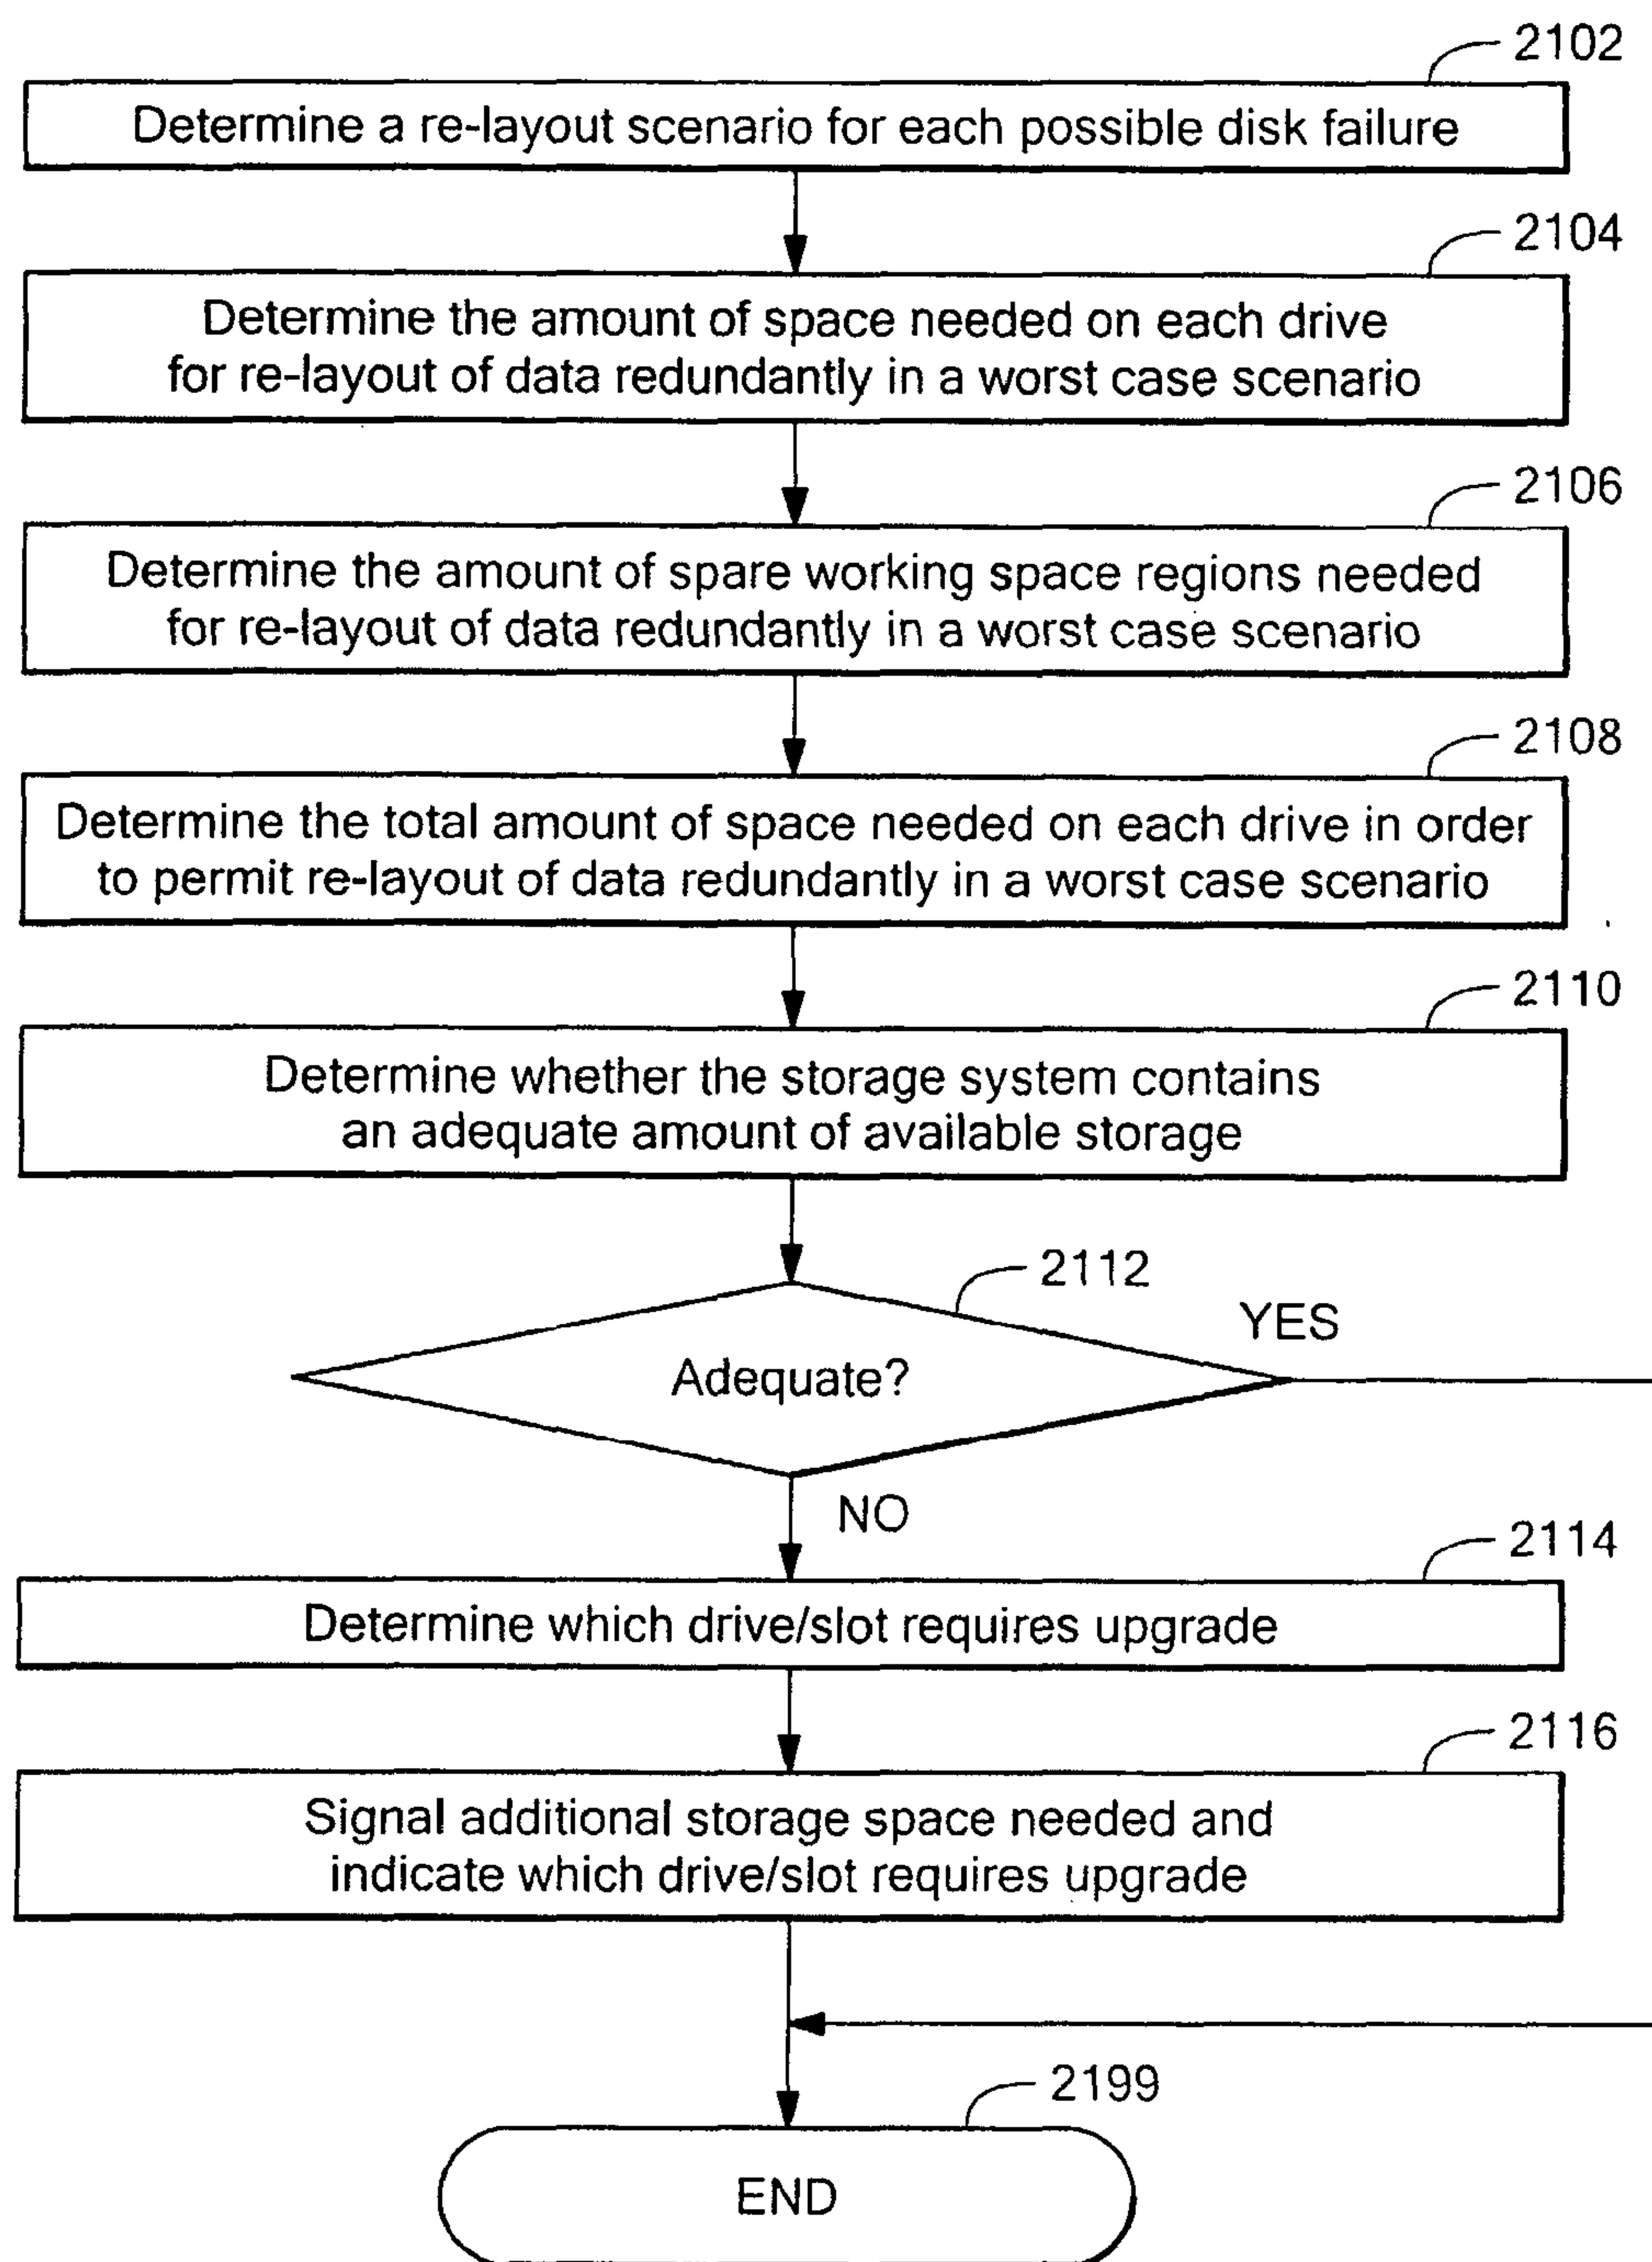
**FIG. 17**

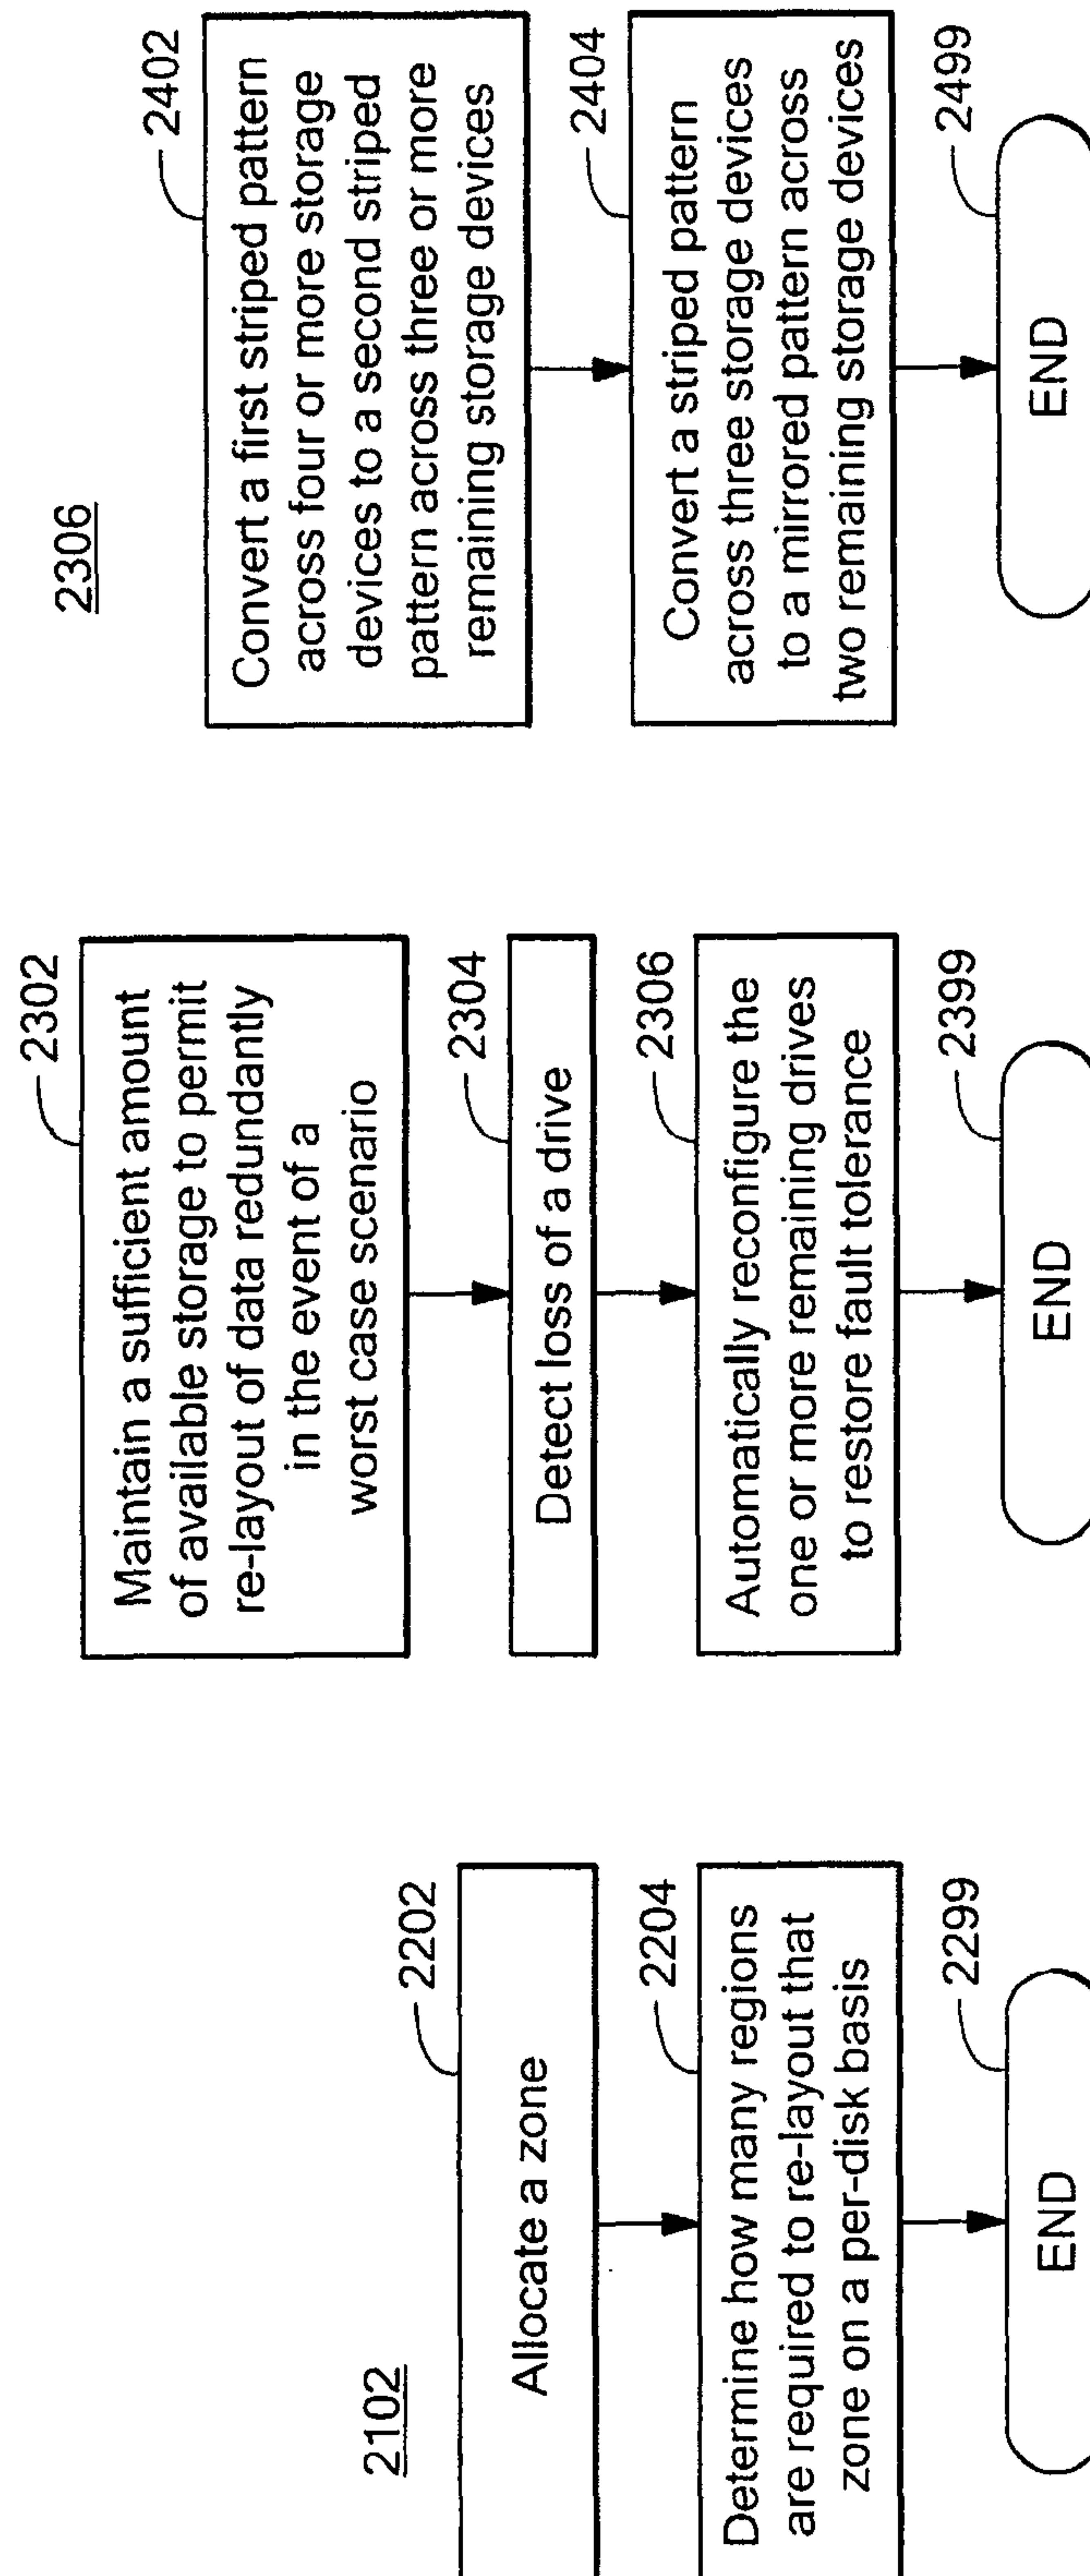
**FIG. 18**

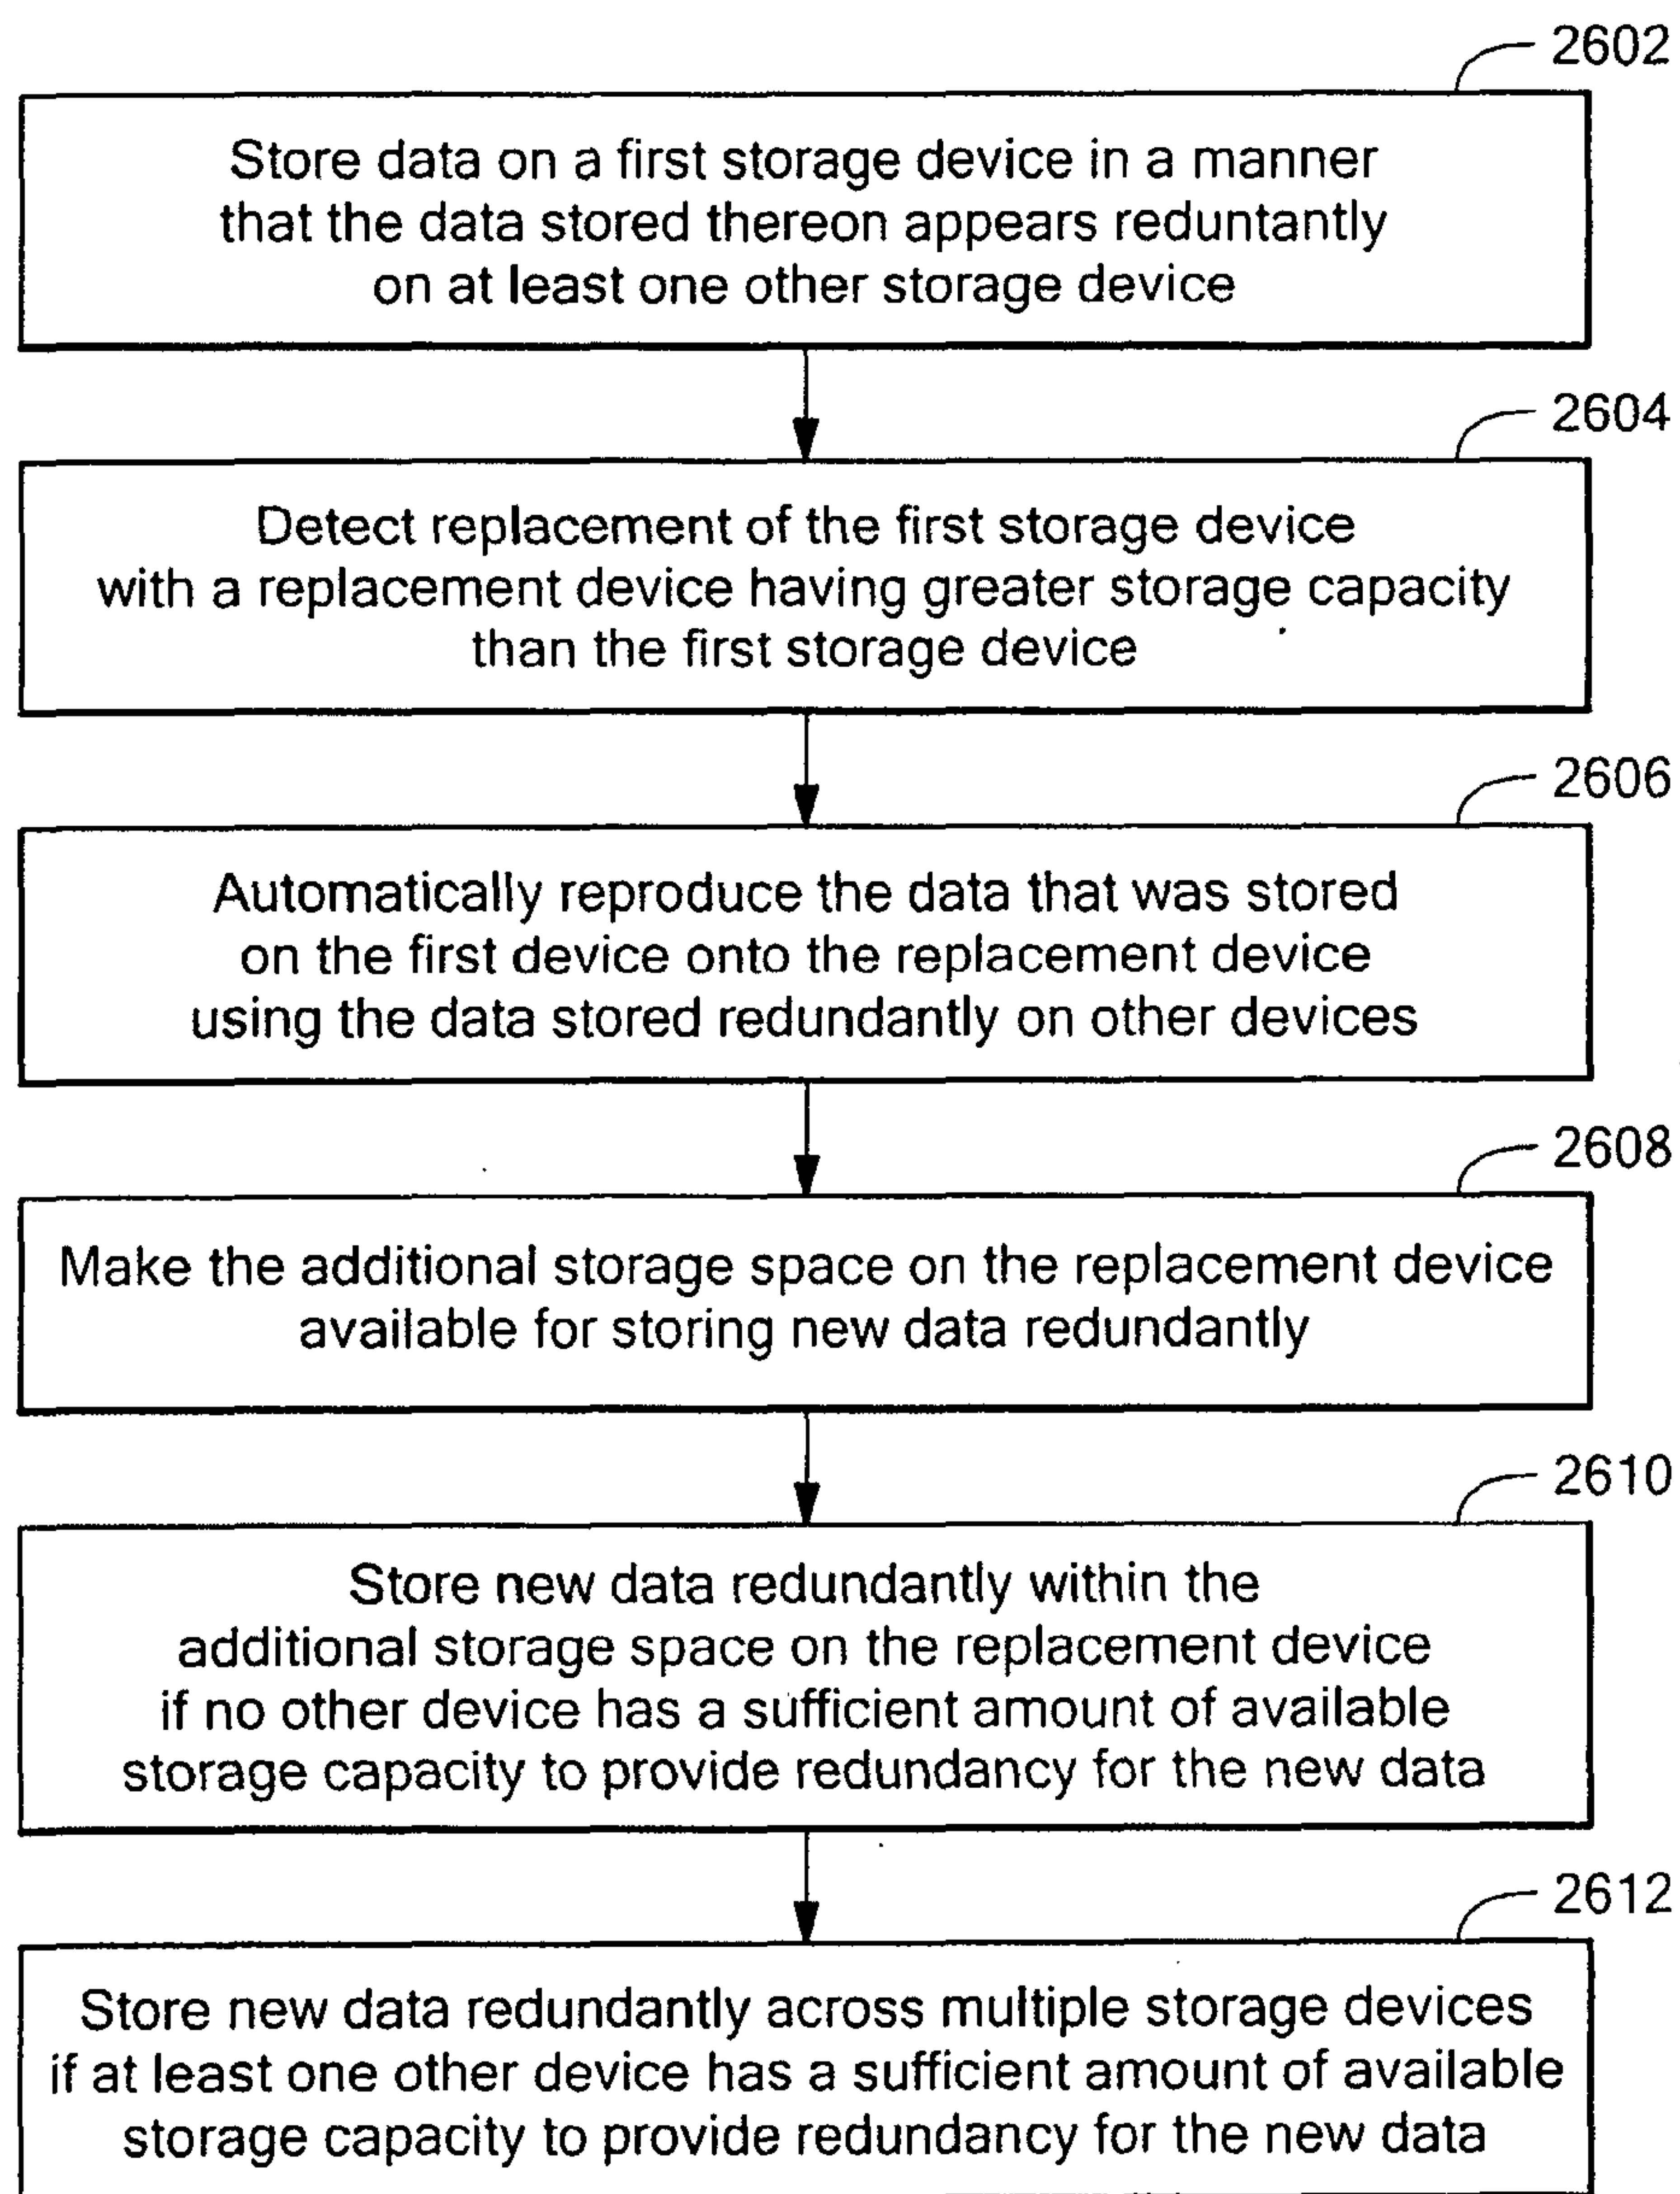
**FIG. 19**


**FIG. 20**

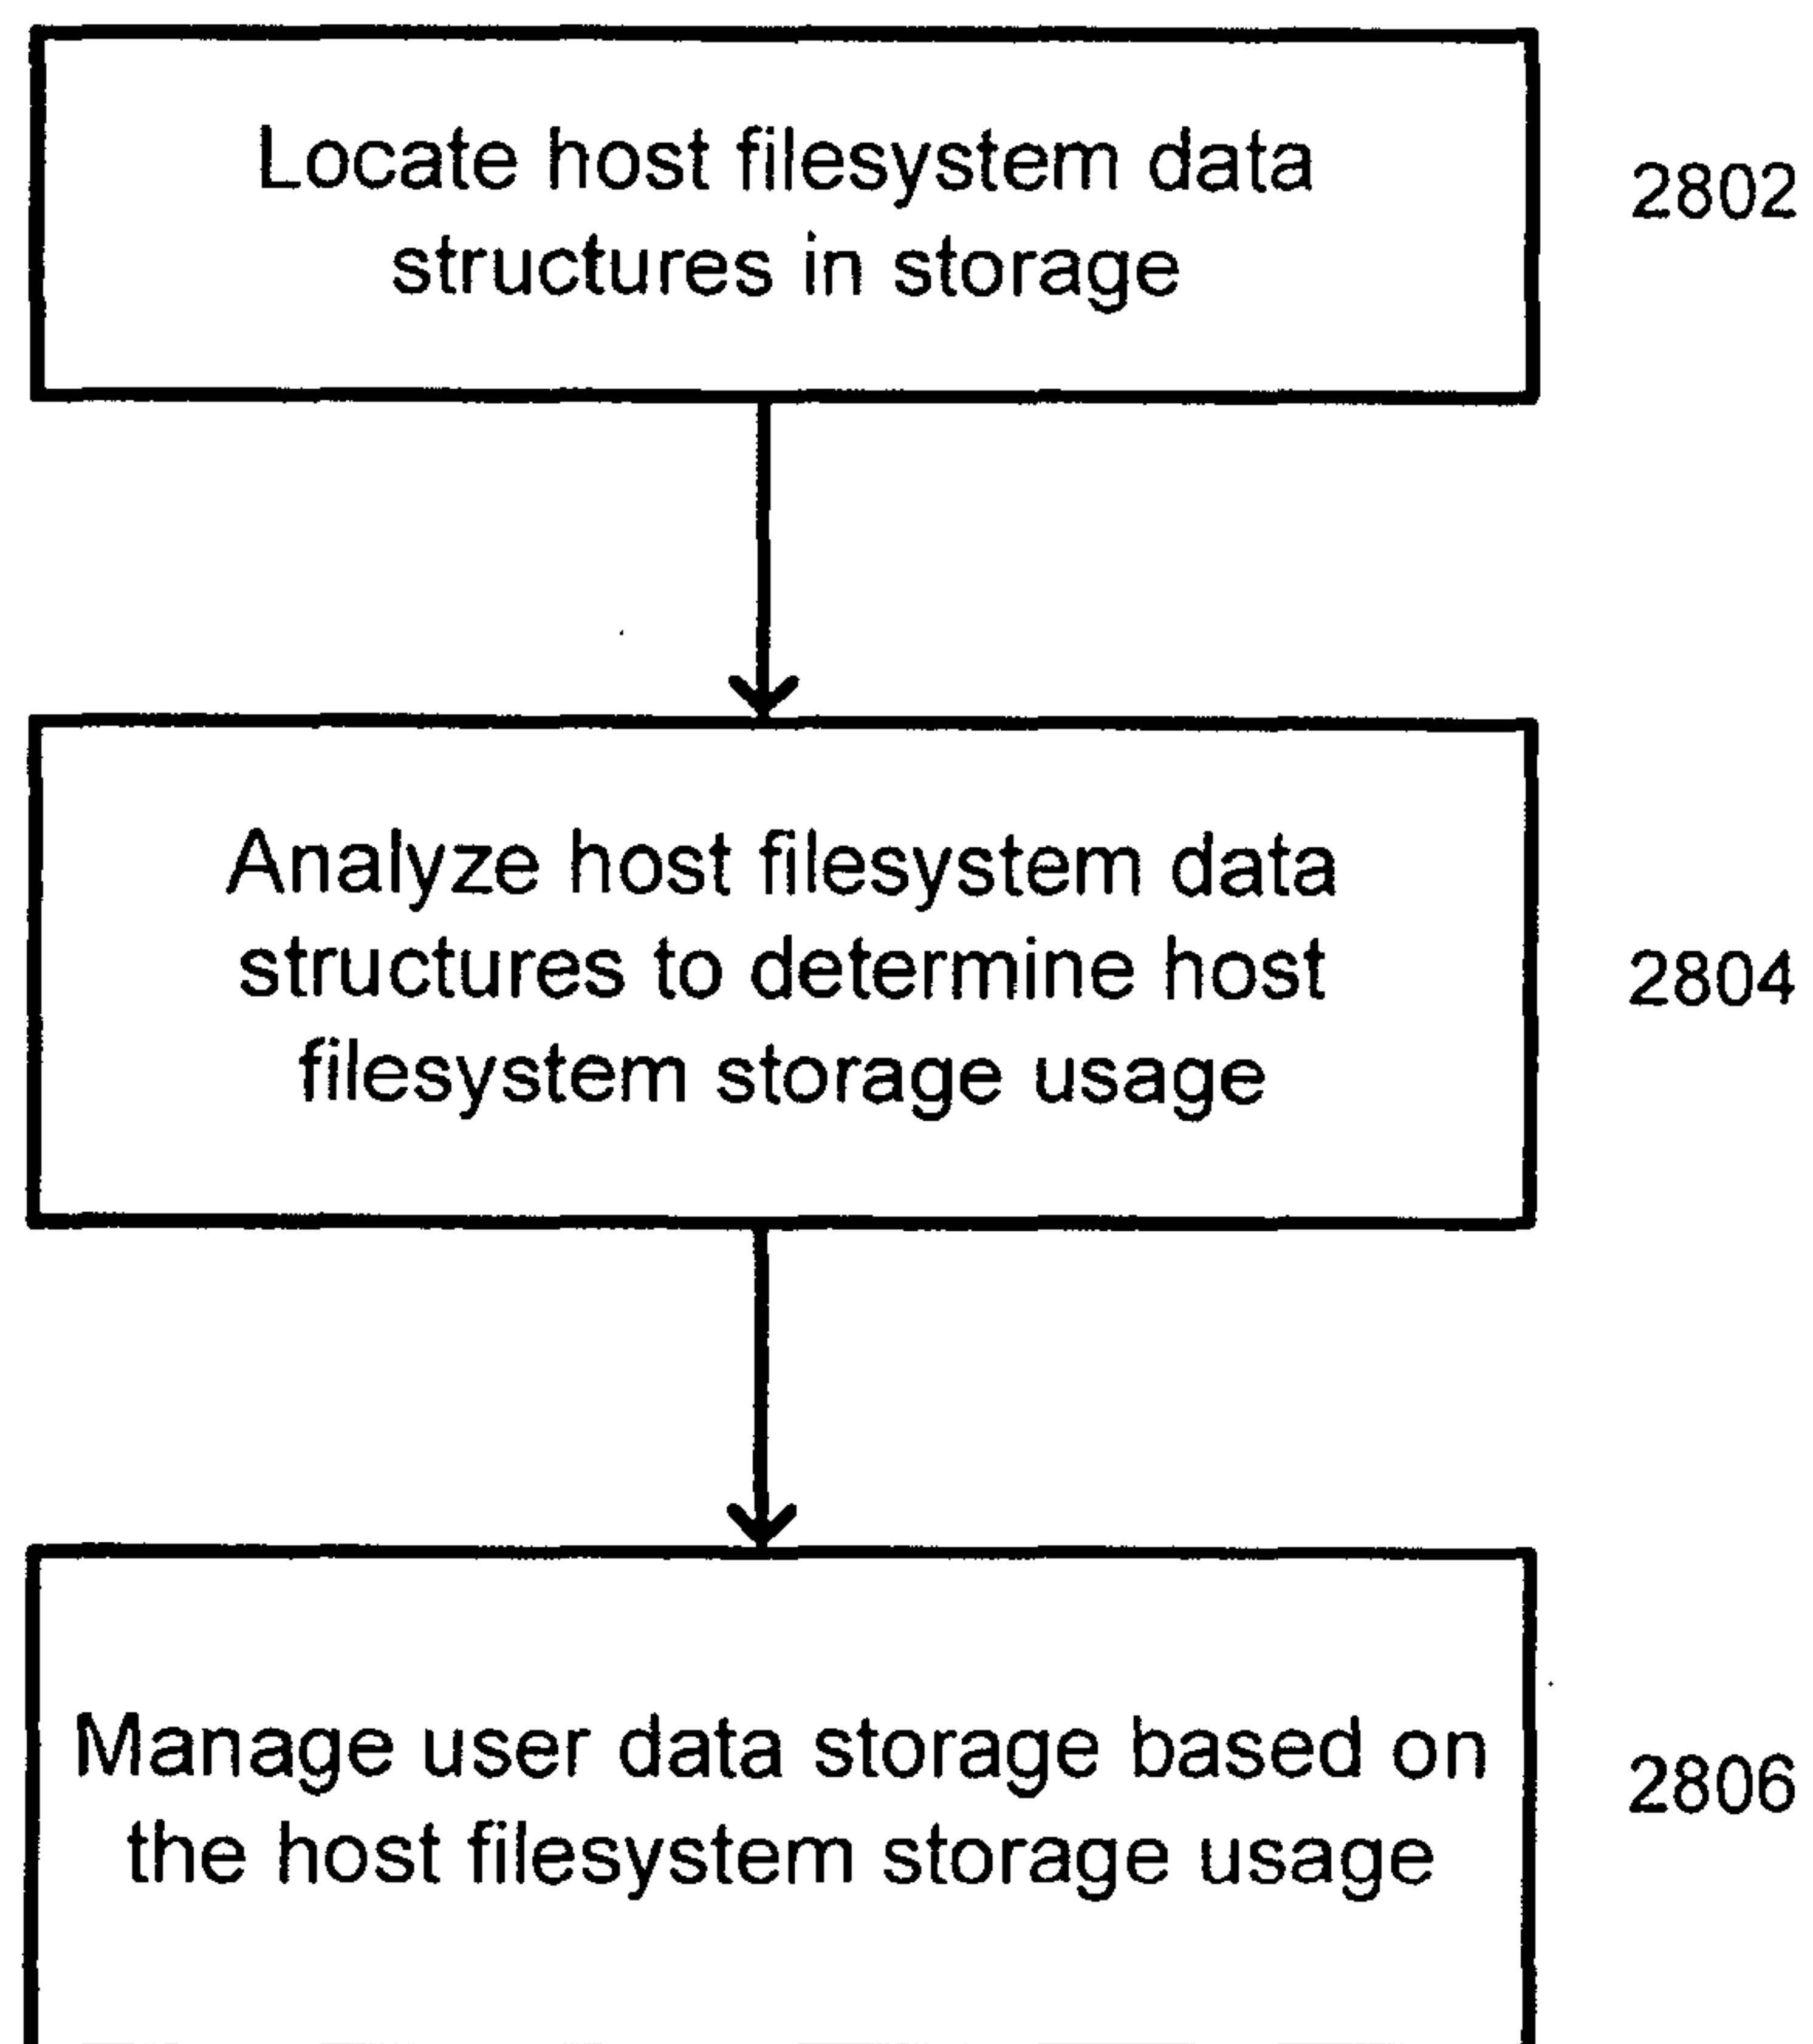
**FIG. 21**

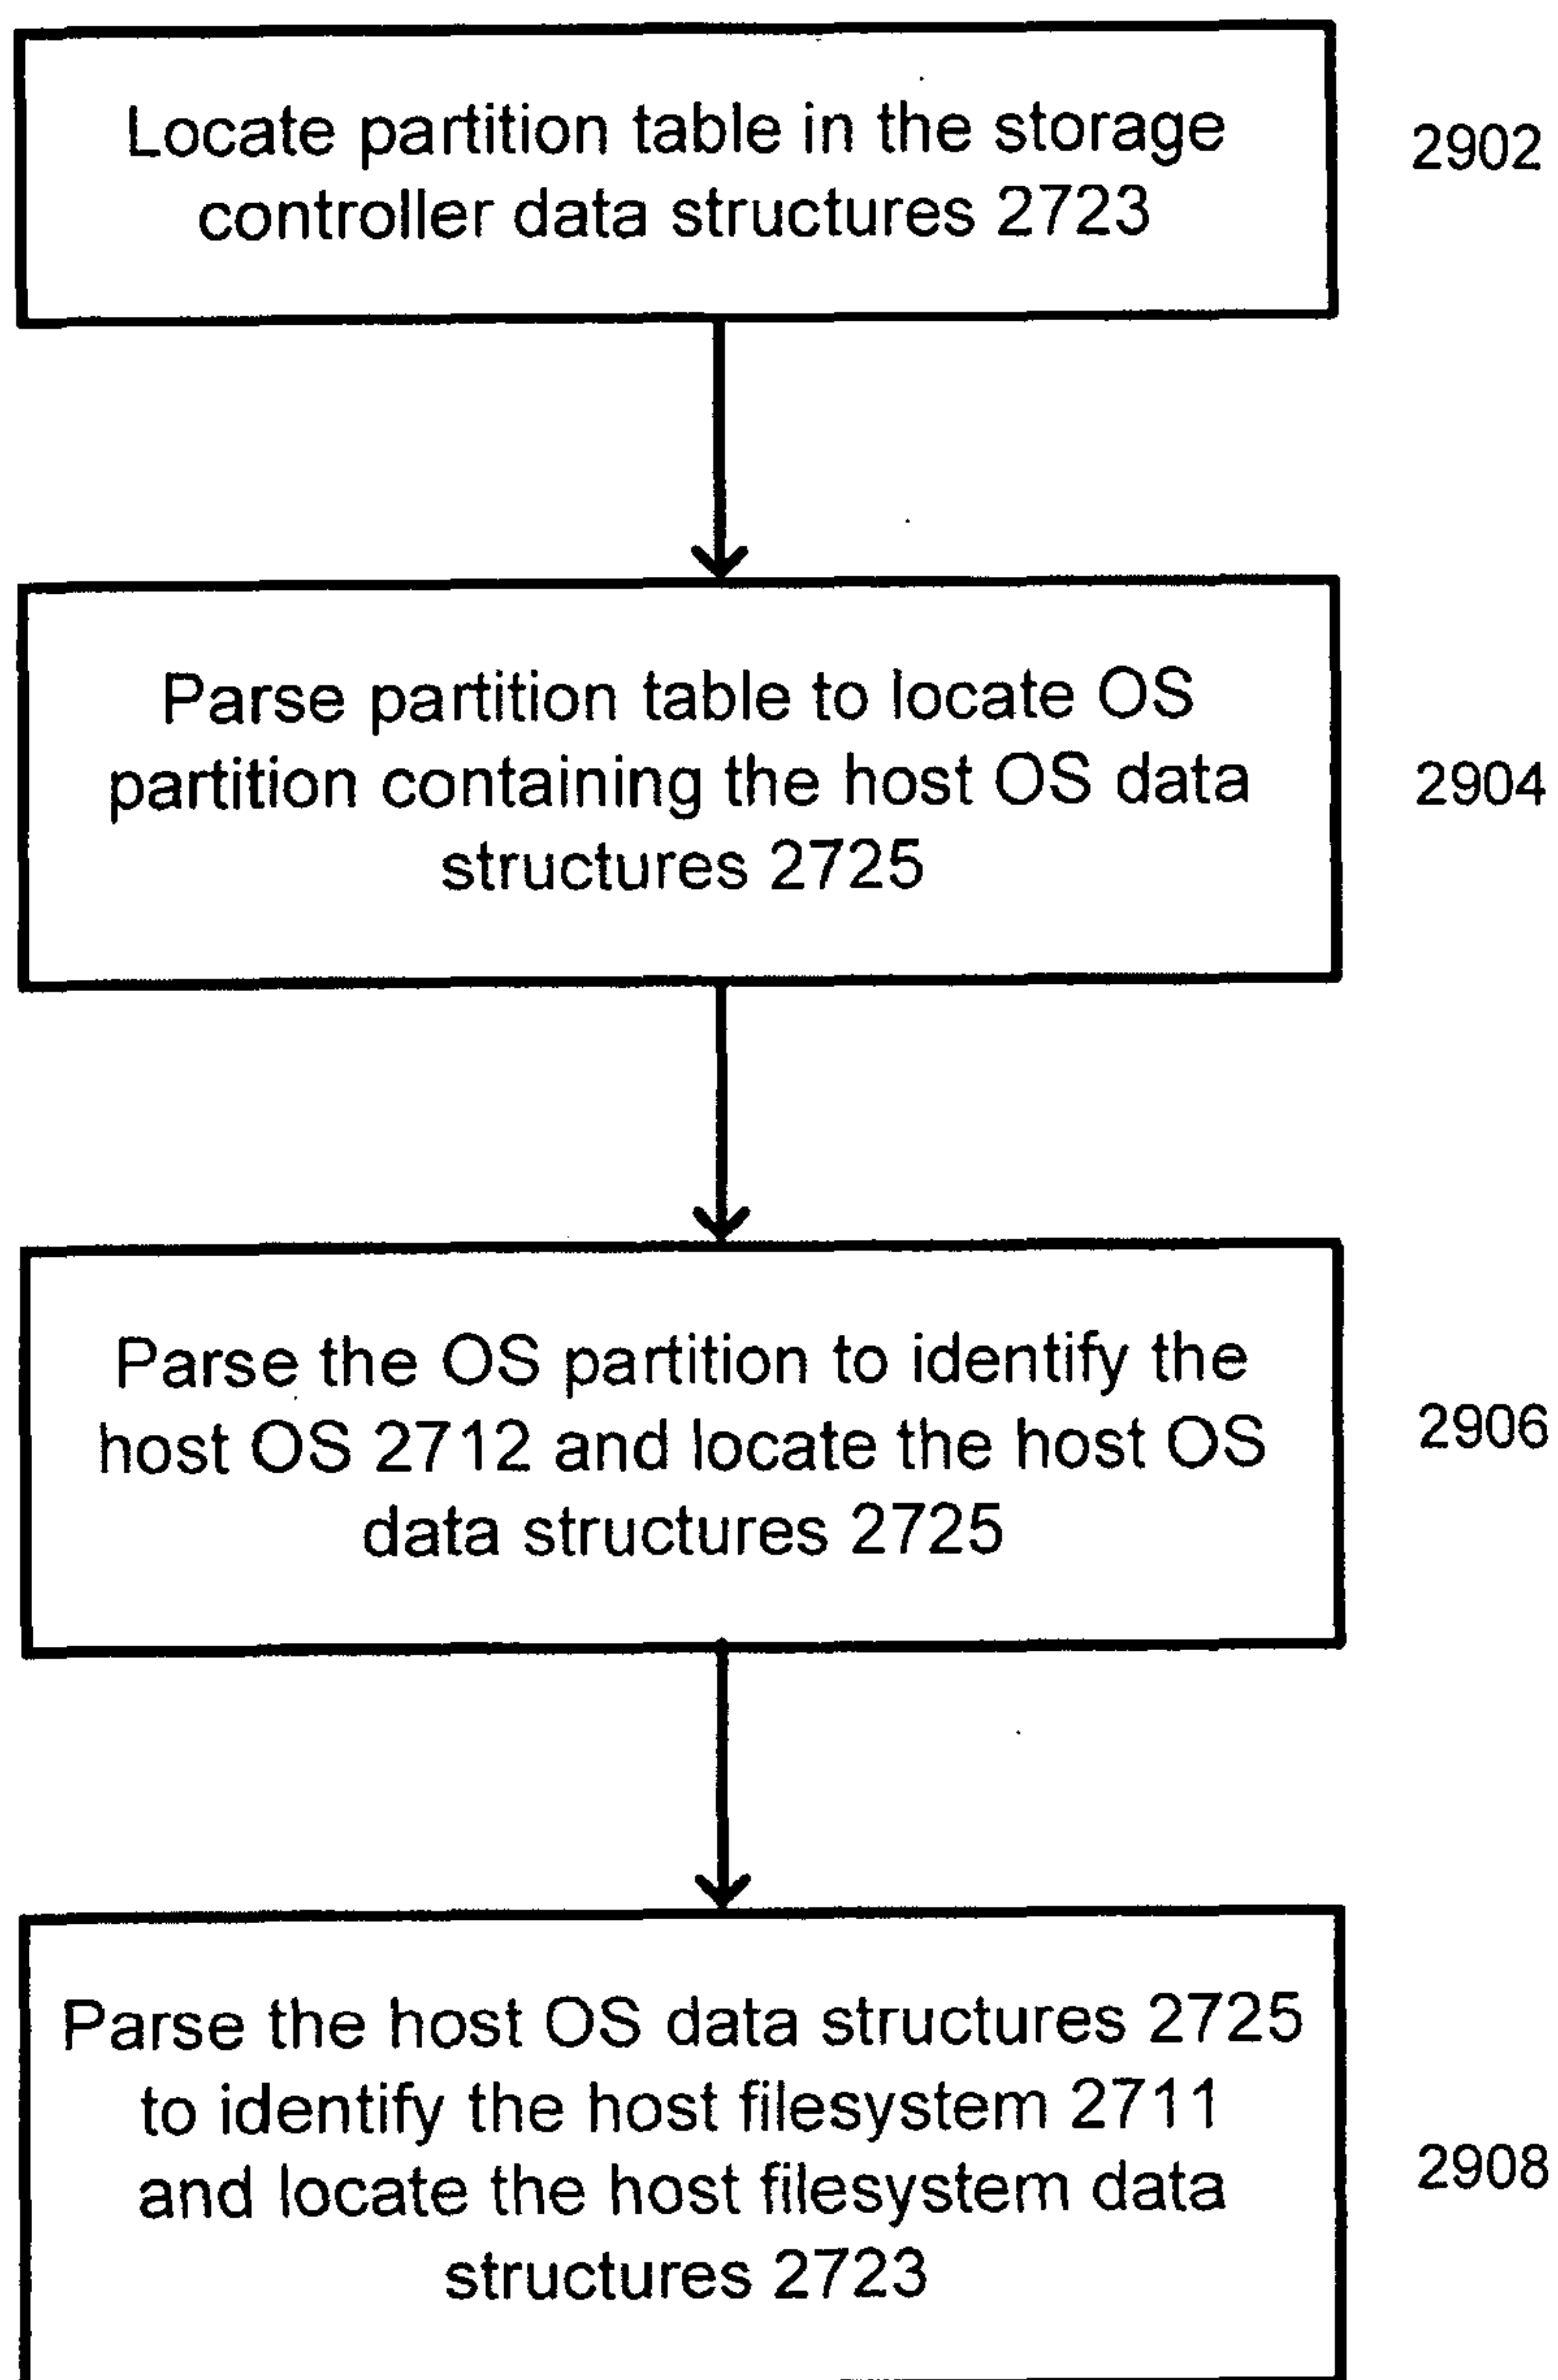
**FIG. 22**

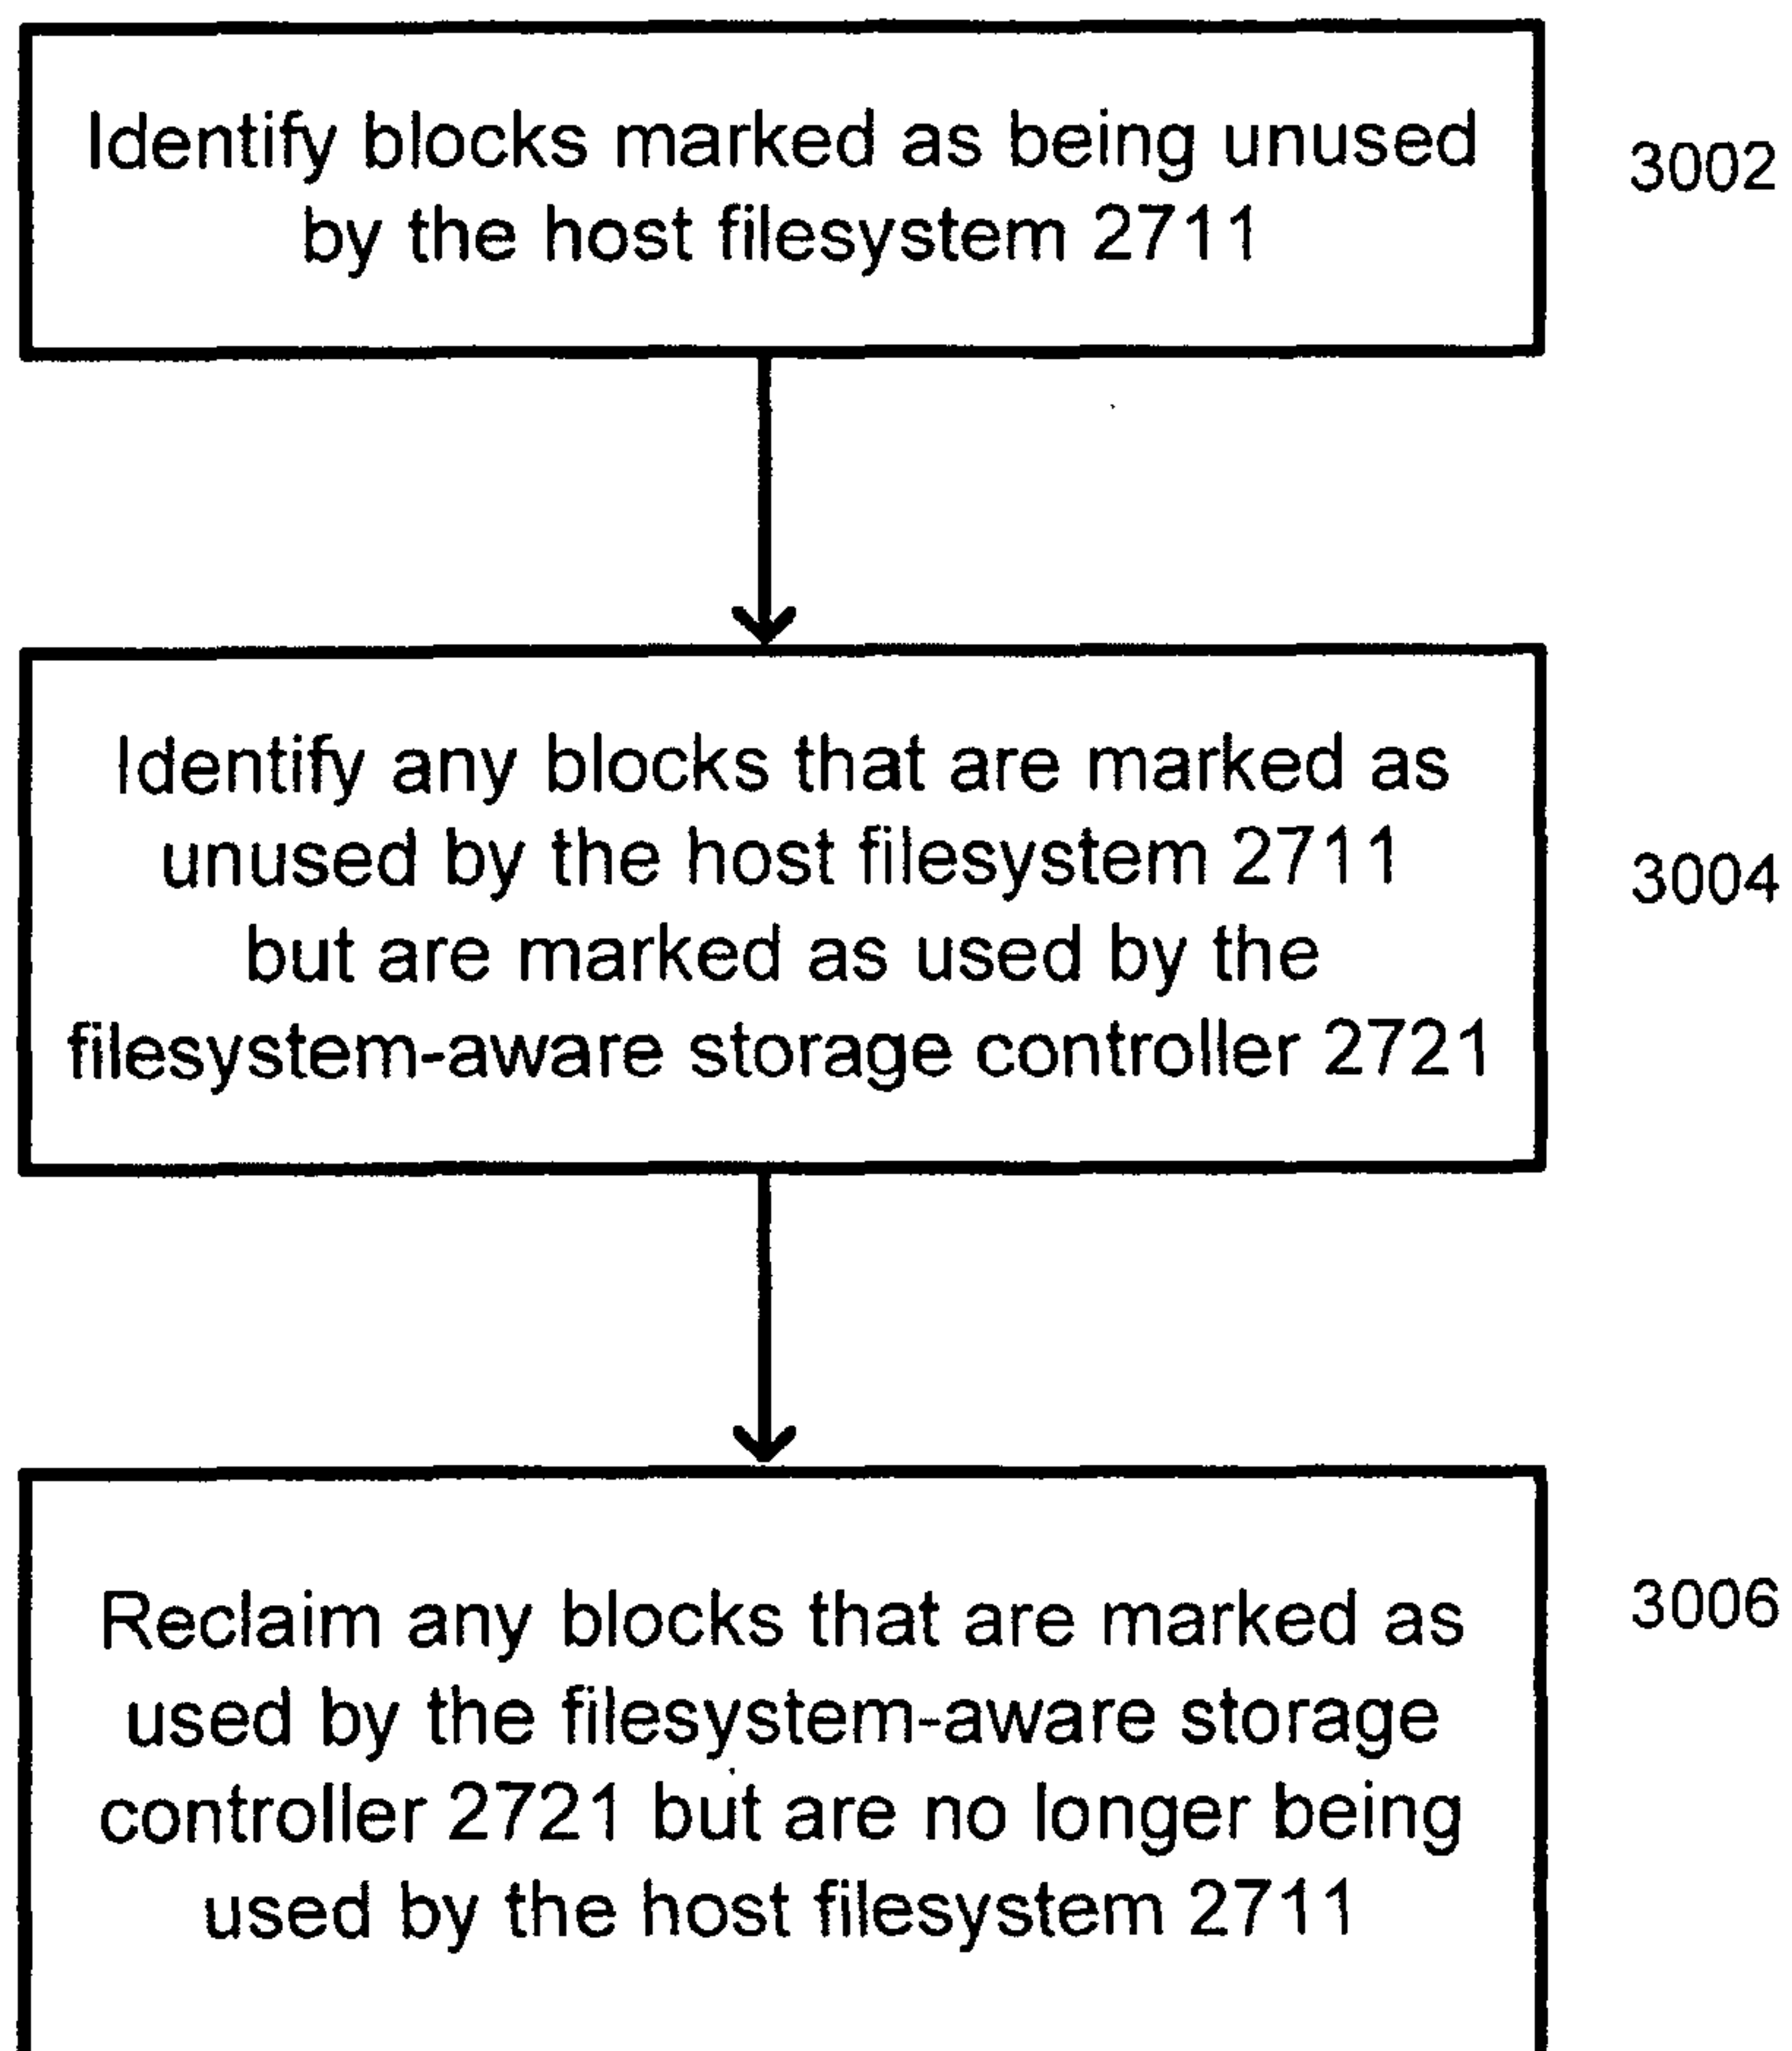


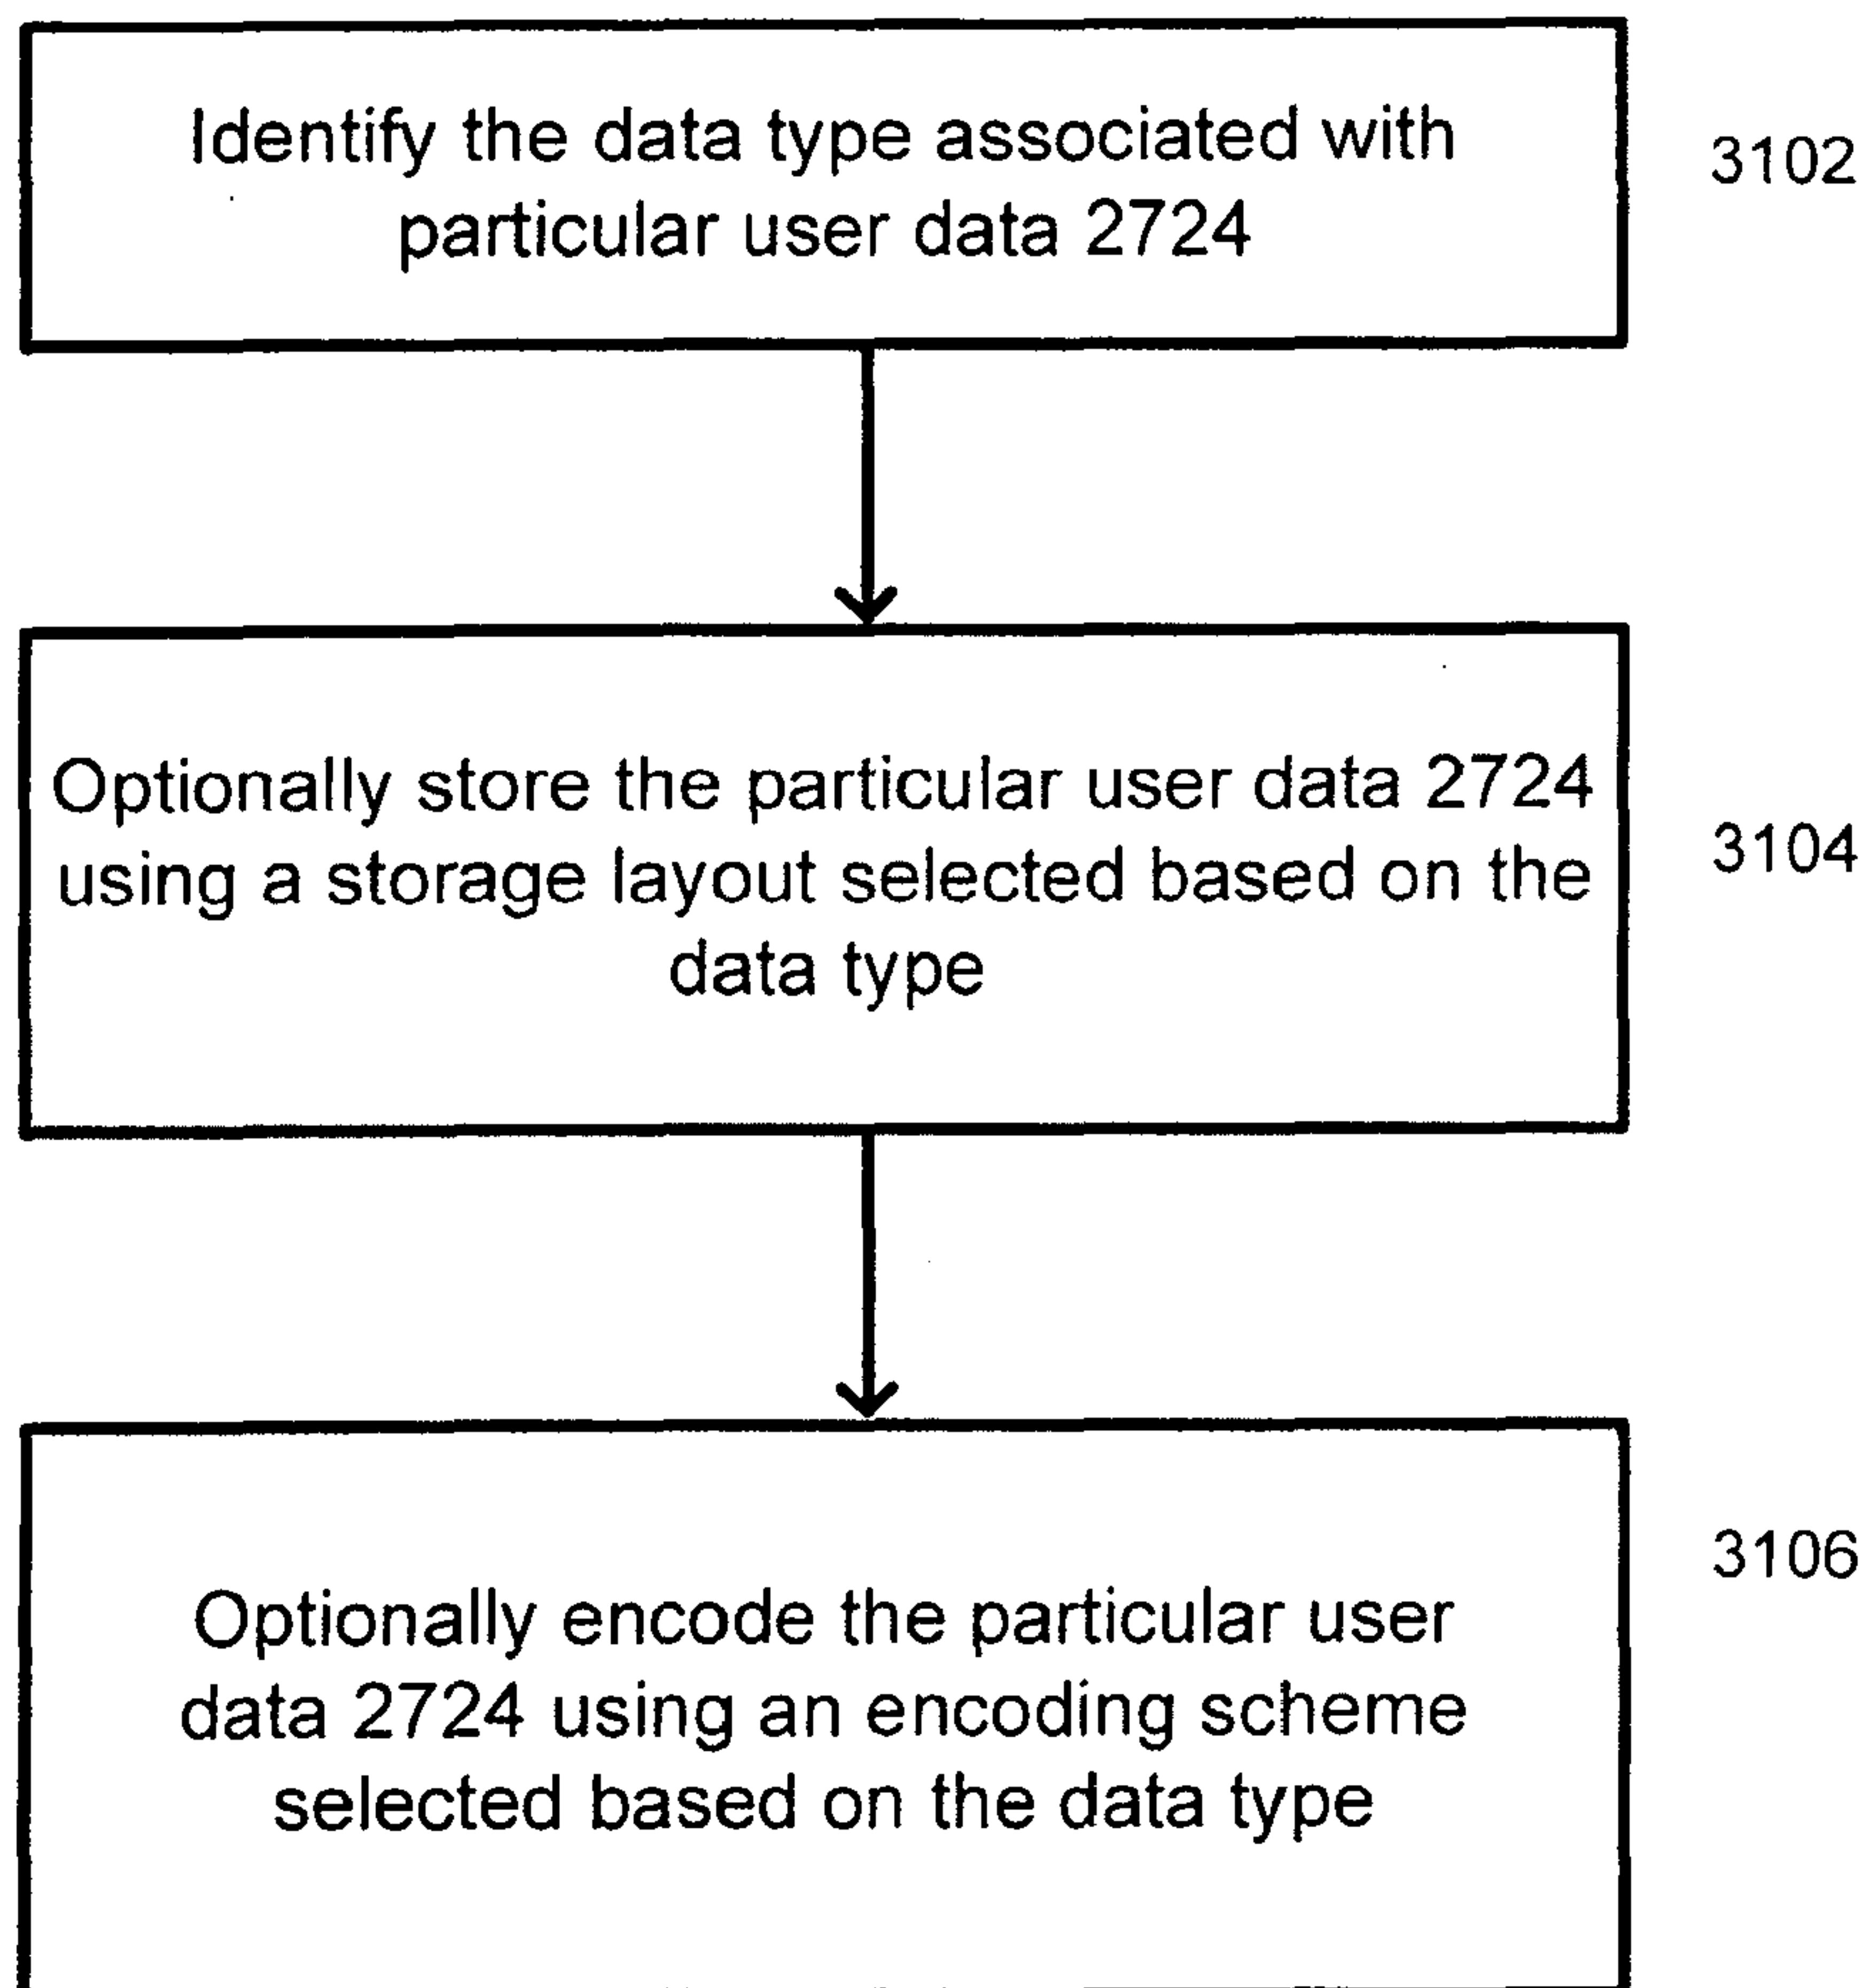
**FIG. 26**

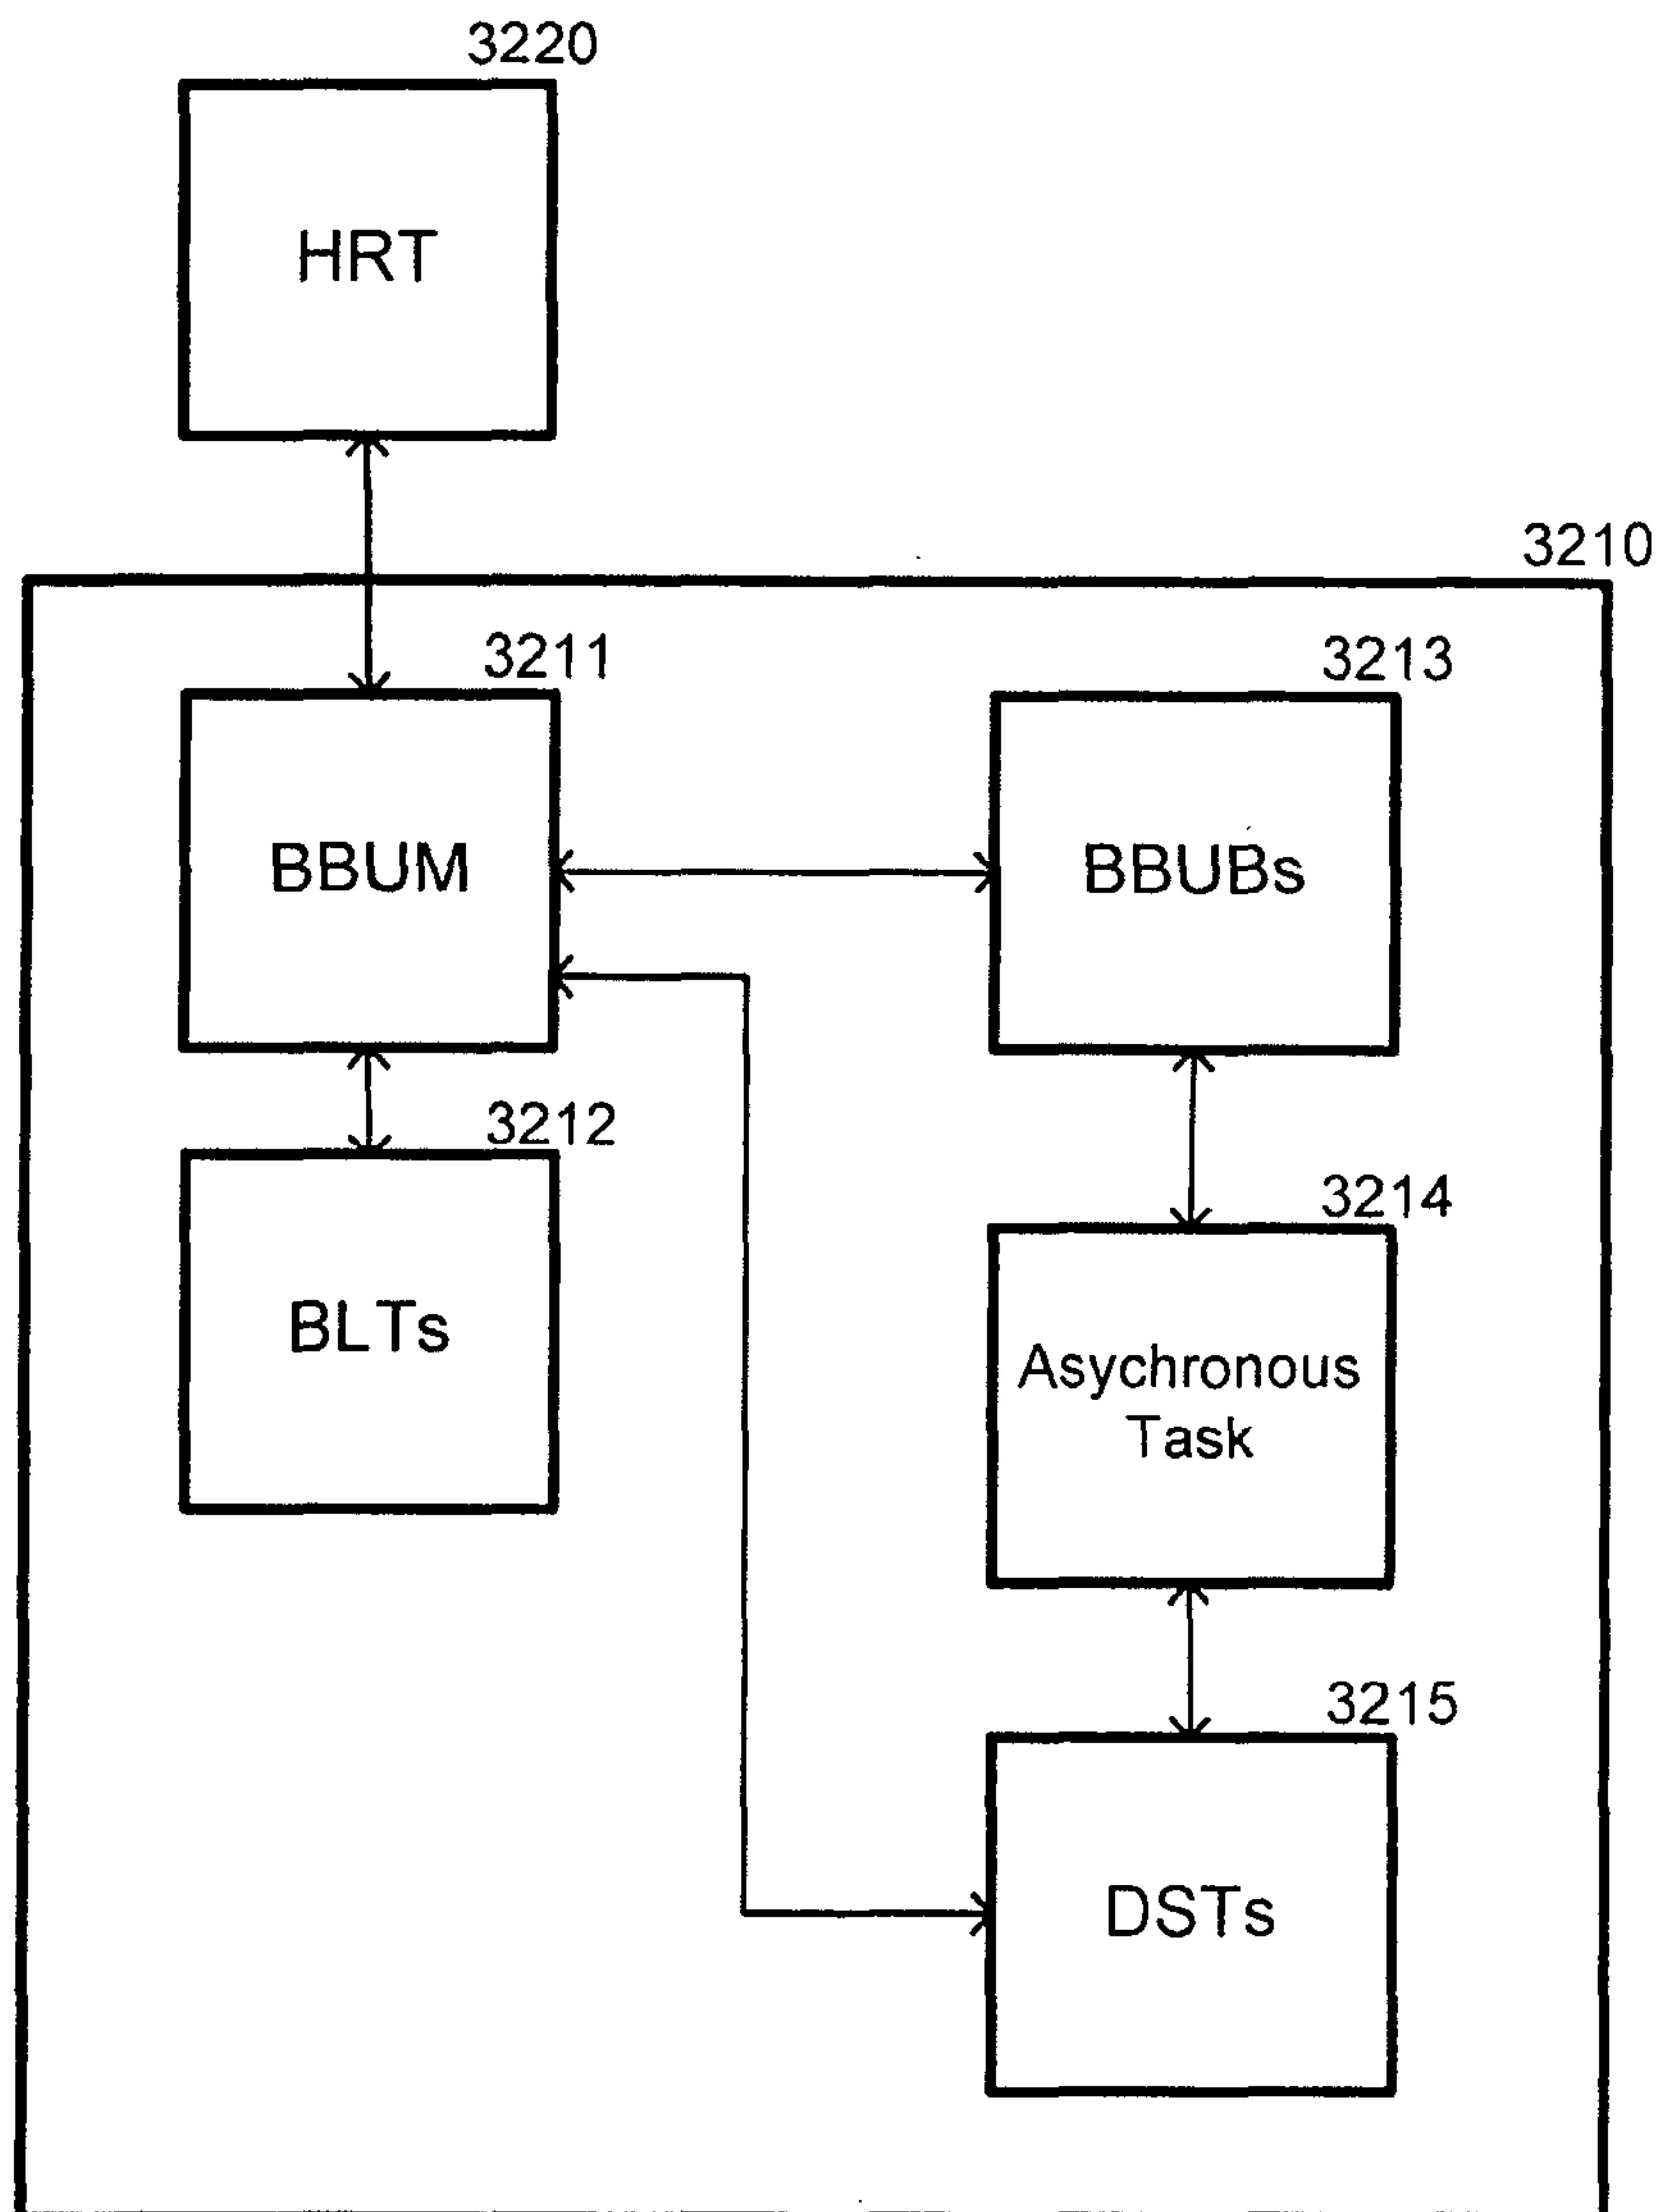
*Fig. 27*

*Fig. 28*

*Fig. 29*

*Fig. 30*

*Fig. 31*

*Fig. 32*

Look for partition table at LBA 0

If (partition table found)

 Read the partition table to identify partitions

 For each partition

 Read the boot sector of the partition to find the MFT

 Read the resident \$bitmap record of the MFT to get file attributes

 Program the BLTs with the boot sector LBA of each partition

 Program the BLTs with the LBA(s) of the bitmap record(s)

 Program the BLTs with the LBA(s) of the actual bitmap(s)

 Flag boot sector LBA(s) to require immediate parsing

 Flag bitmap record LBA(s) to require immediate parsing

 Flag actual bitmap(s) to not require immediate parsing

 Endfor

Else

 End without adding additional locations to the BLTs

Endif

FIG. 33

```
Receive ClientRequest
Get the LUN from the ClientRequest
Find the right BLT based on the LUN
Get the LBA from the ClientRequest
Look for this LBA in the BLT
Check "immediate action" field
If (immediate action is required)
    Process ClientRequest synchronously
Else
    Set BBUB bit corresponding to the LBA for asynchronous processing
Endif
```

FIG. 34

If (block is partition table)

Compare partitions in new data with partitions in BLT

If (new partition being added)

Get start and end of partition from new data

Check DST for any overlapping LBA ranges

Remove overlapping LBA ranges

Add start of partition to BLT

Flag for immediate action

FIG. 35

If (block is partition boot sector)

 Get start of MFT from new data

 Calculate location of bitmap record

 If (identical bitmap record entry already in BLT for this partition)

 Nothing required

 Endif

 If (bitmap record is different location from BLT version)

 Update the BLT

 Read the new location from the disk

 If (that looks like a bitmap record ie. Has \$bitmap string)

 Get the new bitmap location(s)

 Compare with BLT

 If (identical)

 Nothing required

 If (different location)

 Set all BBUB bits

 Update BBUB mappings

 Move the LBA ranges in BLT

 If (smaller)

 Contract the BBUB

 Add the unmapped LBA range into the DST

 Contract the LBA range in the BLT

 If (bigger)

 Set all the additional BBUB bits

 Enlarge the BBUB

 Enlarge the LBA range in the BLT

 Else

 Nothing required

 Endif

 Endif

Endif

FIG. 36


```
If (block is partition table)
  Compare partitions in new data with partitions in BLT
  If(partition being deleted)
    Delete BBUB
    Delete boot sector from BLT
    Delete Bitmap record from BLT
    Delete Bitmap ranges from BLT
    Add partition range to DST
  Endif
Endif
```

FIG. 37

Parse the BBUB

For each bit set in BBUB

 Check whether the corresponding cluster is marked unused by the host filesystem

 If (cluster is marked unused by the host filesystem)

 Check whether the cluster is marked used by the storage controller

 If (cluster is marked used by the storage controller)

 Add LBA range to DST

 Endif

 Endif

Endfor

For each LBA range in the DST

 Reclaim the storage space

Endfor

FIG. 38

