

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
5 April 2012 (05.04.2012)

(10) International Publication Number
WO 2012/044492 A1

(51) International Patent Classification:
G06F 11/10 (2006.01) *G06F 3/06* (2006.01)

CA 94043 (US). **MILLER, Ethan** [US/US]; 203 Kalkar Drive, Santa Cruz, CA 95060 (US).

(21) International Application Number:
PCT/US2011/052276

(74) Agent: **MEYERTONS, HOOD, KIVLIN, KOWERT & GOETZEL, P.C.**; RANKIN, Rory D., P.O. Box 398, Austin, TX 78767-0398 (US).

(22) International Filing Date:
20 September 2011 (20.09.2011)

(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/896,669 1 October 2010 (01.10.2010) US

(71) Applicant (for all designated States except US): **PURE STORAGE, INC.** [US/US]; 650 Castro Street, Suite 220, Mountain View, CA 94041 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **COLGROVE, John** [US/US]; 722 Vista Grande Ave., Los Altos, CA 94024 (US). **HAYES, John** [CA/US]; 800 High School Way, #330, Mountain View, CA 94041 (US). **HONG, Bo** [CN/US]; 1555 W. Middlefield Rd., Apt. 95, Mountain View,

(84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ,

[Continued on next page]

(54) Title: RECONSTRUCT READS IN A RAID ARRAY WITH DYNAMIC GEOMETRIES




FIG. 13

(57) **Abstract:** A system and method for dynamic RAID geometries. A computer system comprises client computers and data storage arrays coupled to one another via a network. A data storage array utilizes solid-state drives and Flash memory cells for data storage. A storage controller within a data storage array is configured to configure a first subset of the storage devices for use in a first RAID layout, the first RAID layout including a first set of redundant data. The controller further configures a second subset of the storage devices for use in a second RAID layout, the second RAID layout including a second set of redundant data. Additionally, the controller configures an additional device not included in either the first subset or the second subset to store redundant data for both the first RAID layout and the second RAID layout. The controller is further configured to initiate a reconstruct read corresponding to a given read request directed to a particular storage device of the plurality of storage devices, in response to determining the particular storage device is exhibiting a non-error related relatively slow read response.



TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

**TITLE: RECONSTRUCT READS IN A RAID ARRAY WITH DYNAMIC
GEOMETRIES**

BACKGROUND OF THE INVENTION

5

Field of the Invention

[0001] This invention relates to computer networks and, more particularly, to efficiently distributing data among a plurality of solid-state storage devices.

10

Description of the Related Art

[0002] As computer memory storage and data bandwidth increase, so does the amount and complexity of data that businesses daily manage. Large-scale distributed storage systems, such as data centers, typically run many business operations. A distributed storage system may be coupled to client computers interconnected by one or more networks. If any portion of the distributed storage system has poor performance or becomes unavailable, company operations may be impaired or stopped completely. A distributed storage system therefore is expected to maintain high standards for data availability and high-performance functionality. As used herein, storage disks may be referred to as storage devices as some types of storage technologies do not include disks.

20 **[0003]** To protect against data loss, storage devices often include error detection and correction mechanisms. Often these mechanisms take the form of error correcting codes which are generated by the devices and stored within the devices themselves. In addition, distributed storage systems may also utilize decentralized algorithms to distribute data among a collection of storage devices. These algorithms generally map data objects to storage devices without relying on a central directory. Examples of such algorithms include Replication Under Scalable Hashing (RUSH), and Controlled Replication Under Scalable Hashing (CRUSH). With no central directory, multiple clients in a distributed storage system may simultaneously access data objects on multiple servers. In addition, the amount of stored metadata may be reduced. However, the difficult task remains of distributing data among multiple storage disks with varying capacities, 25 input/output (I/O) characteristics and reliability issues. Similar to the storage devices themselves, these algorithms may also include error detection and correction algorithms such as RAID type algorithms (e.g., RAID5 and RAID6) or Reed-Solomon codes.

[0004] The technology and mechanisms associated with chosen storage devices determine the methods used to distribute data among multiple storage devices, which may be dynamically added and removed. For example, the algorithms described above were developed for systems utilizing hard disk drives (HDDs). The HDDs comprise one or more rotating disks, each coated with a magnetic medium. These disks rotate at a rate of several thousand rotations per minute for several hours daily. In addition, a magnetic actuator is responsible for positioning magnetic read/write devices over the rotating disks. These actuators are subject to friction, wear, vibrations and mechanical misalignments, which result in reliability issues. The above-described data distribution algorithms are based upon the characteristics and behaviors of HDDs.

5 [0005] One example of another type of storage disk is a Solid-State Disk (SSD). A Solid-State Disk may also be referred to as a Solid-State Drive. An SSD may emulate a HDD interface, but an SSD utilizes solid-state memory to store persistent data rather than electromechanical devices as found in a HDD. For example, an SSD may comprise banks of Flash memory. Without moving parts or mechanical delays, an SSD may have a lower access time and latency than a HDD. However, SSD typically have significant write latencies. In addition to different input/output (I/O) characteristics, an SSD experiences different failure modes than a HDD. Accordingly, high performance and high reliability may not be achieved in systems comprising SSDs for storage while utilizing distributed data placement algorithms developed for HDDs.

10 [0006] In view of the above, systems and methods for efficiently distributing data and detecting and correcting errors among a plurality of solid-state storage devices are desired.

15 [0007] Various embodiments of a computer system and methods for efficiently distributing and managing data among a plurality of solid-state storage devices are disclosed.

20 [0008] In one embodiment, a computer system comprises a plurality of client computers configured to convey read and write requests over a network to one or more data storage arrays coupled to receive the read and write requests via the network. Contemplated is a data storage array(s) comprising a plurality of storage locations on a plurality of storage devices. In various

25 embodiments, the storage devices are configured in a redundant array of independent drives (RAID) arrangement for data storage and protection. The data storage devices may include solid-state memory technology for data storage, such as Flash memory cells. The data storage subsystem further comprises a storage controller configured to configure a first subset of the

storage devices for use in a first RAID layout, the first RAID layout including a first set of redundant data. The controller further configures a second subset of the storage devices for use in a second RAID layout, the second RAID layout including a second set of redundant data. Additionally, the controller configures an additional device not included in either the first subset or the second subset to store redundant data for both the first RAID layout and the second RAID layout.

[0009] Also contemplated are embodiments wherein in response to receiving a given read request directed to a particular storage device of either the first RAID layout or the second RAID layout, the controller is configured to initiate a reconstruct read corresponding to the given read request in response to determining the particular storage device is exhibiting a non-error related relatively slow read response. Additionally, each device in the first subset and second subset may be configured to store intra-device redundancy data, where both the first set of redundant data and the second set of redundant data comprise inter-device redundancy data. Further, in various embodiments, the first RAID layout is an $L+x$ layout, and the second RAID layout is an $M+y$ layout, wherein L , x , M , and, y are integers, wherein either or both (1) L is not equal to M , and (2) x is not equal to y .

[0010] These and other embodiments will become apparent upon consideration of the following description and accompanying drawings.

20 BRIEF DESCRIPTION OF THE DRAWINGS

[0011] FIG. 1 is a generalized block diagram illustrating one embodiment of network architecture.

[0012] FIG. 2 is a generalized block diagram of one embodiment of a dynamic intra-device redundancy scheme.

25 [0013] FIG. 3 is a generalized flow diagram illustrating one embodiment of a method for adjusting intra-device protection in a data storage subsystem.

[0014] FIG. 4 is a generalized block diagram of one embodiment of a storage subsystem.

[0015] FIG. 5 is a generalized block diagram of one embodiment of a device unit.

[0016] FIG. 6 is a generalized block diagram illustrating one embodiment of a state table.

30 [0017] FIG. 7 is a generalized block diagram illustrating one embodiment of a flexible RAID data layout architecture.

[0018] FIG. 8 is a generalized block diagram illustrating another embodiment of a flexible RAID data layout architecture.

[0019] FIG. 9 is a generalized flow diagram illustrating one embodiment of a method for dynamically determining a layout in a data storage subsystem.

[0020] FIG. 10 is a generalized block diagram illustrating yet another embodiment of a flexible RAID data layout architecture.

5 [0021] FIG. 11A illustrates one embodiment of a device layout.

[0022] FIG. 11B illustrates one embodiment of a segment.

[0023] FIG. 11C is a generalized block diagram illustrating one embodiment of data storage arrangements within different page types.

10 [0024] FIG. 12 is a generalized block diagram illustrating one embodiment of a hybrid RAID data layout.

[0025] FIG. 13 is a generalized flow diagram illustrating one embodiment of a method for selecting alternate RAID geometries in a data storage subsystem.

15 [0026] While the invention is susceptible to various modifications and alternative forms, specific embodiments are shown by way of example in the drawings and are herein described in detail. It should be understood, however, that drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the invention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

20 DETAILED DESCRIPTION

[0027] In the following description, numerous specific details are set forth to provide a thorough understanding of the present invention. However, one having ordinary skill in the art should recognize that the invention might be practiced without these specific details. In some instances, well-known circuits, structures, signals, computer program instruction, and techniques 25 have not been shown in detail to avoid obscuring the present invention.

[0028] Referring to FIG. 1, a generalized block diagram of one embodiment of network architecture 100 is shown. As described further below, one embodiment of network architecture 100 includes client computer systems 110a-110b interconnected to one another through a network 180 and to data storage arrays 120a-120b. Network 180 may be coupled to a second 30 network 190 through a switch 140. Client computer system 110c is coupled to client computer systems 110a-110b and data storage arrays 120a-120b via network 190. In addition, network 190 may be coupled to the Internet 160 or other outside network through switch 150.

[0029] It is noted that in alternative embodiments, the number and type of client computers and servers, switches, networks, data storage arrays, and data storage devices is not limited to those shown in FIG. 1. At various times one or more clients may operate offline. In addition, during operation, individual client computer connection types may change as users connect, disconnect, 5 and reconnect to network architecture 100. A further description of each of the components shown in FIG. 1 is provided shortly. First, an overview of some of the features provided by the data storage arrays 120a-120b is described.

[0030] In the network architecture 100, each of the data storage arrays 120a-120b may be used for the sharing of data among different servers and computers, such as client computer systems 10 110a-110c. In addition, the data storage arrays 120a-120b may be used for disk mirroring, backup and restore, archival and retrieval of archived data, and data migration from one storage device to another. In an alternate embodiment, one or more client computer systems 110a-110c may be linked to one another through fast local area networks (LANs) in order to form a cluster. One or more nodes linked to one another form a cluster, which may share a storage resource, 15 such as a cluster shared volume residing within one of data storage arrays 120a-120b.

[0031] Each of the data storage arrays 120a-120b includes a storage subsystem 170 for data storage. Storage subsystem 170 may comprise a plurality of storage devices 176a-176m. These storage devices 176a-176m may provide data storage services to client computer systems 110a- 20 110c. Each of the storage devices 176a-176m may be configured to receive read and write requests and comprise a plurality of data storage locations, each data storage location being addressable as rows and columns in an array. In one embodiment, the data storage locations within the storage devices 176a-176m may be arranged into logical, redundant storage containers or RAID arrays (redundant arrays of inexpensive/independent disks). However, the storage devices 176a-176m may not comprise a disk. In one embodiment, each of the storage devices 25 176a-176m may utilize technology for data storage that is different from a conventional hard disk drive (HDD). For example, one or more of the storage devices 176a-176m may include or be further coupled to storage consisting of solid-state memory to store persistent data. In other embodiments, one or more of the storage devices 176a-176m may include or be further coupled to storage utilizing spin torque transfer technique, magnetoresistive random access memory 30 (MRAM) technique, or other storage techniques. These different storage techniques may lead to differing reliability characteristics between storage devices.

[0032] The type of technology and mechanism used within each of the storage devices 176a- 176m may determine the algorithms used for data object mapping and error detection and

correction. The logic used in these algorithms may be included within one or more of a base operating system (OS) 116, a file system 140, one or more global RAID engines 178 within a storage subsystem controller 174, and control logic within each of the storage devices 176a-176m.

5 [0033] In one embodiment, the included solid-state memory comprises solid-state drive (SSD) technology. Typically, SSD technology utilizes Flash memory cells. As is well known in the art, a Flash memory cell holds a binary value based on a range of electrons trapped and stored in a floating gate. A fully erased Flash memory cell stores no or a minimal number of electrons in the floating gate. A particular binary value, such as binary 1 for single-level cell (SLC) Flash, is
10 associated with an erased Flash memory cell. A multi-level cell (MLC) Flash has a binary value
11 associated with an erased Flash memory cell. After applying a voltage higher than a given threshold voltage to a controlling gate within a Flash memory cell, the Flash memory cell traps a given range of electrons in the floating gate. Accordingly, another particular binary value, such as binary 0 for SLC Flash, is associated with the programmed (written) Flash memory cell. A
15 MLC Flash cell may have one of multiple binary values associated with the programmed memory cell depending on the applied voltage to the control gate.

[0034] Generally speaking, SSD technologies provide lower read access latency times than HDD technologies. However, the write performance of SSDs is significantly impacted by the availability of free, programmable blocks within the SSD. As the write performance of SSDs is
20 significantly slower compared to the read performance of SSDs, problems may occur with certain functions or operations expecting similar latencies. In addition, the differences in technology and mechanisms between HDD technology and SDD technology lead to differences in reliability characteristics of the data storage devices 176a-176m.

[0035] In various embodiments, a Flash cell within an SSD must generally be erased before it is
25 written with new data. Additionally, an erase operation in various flash technologies must also be performed on a block-wise basis. Consequently, all of the Flash memory cells within a block (an erase segment or erase block) are erased together. A Flash erase block may comprise multiple pages. For example, a page may be 4 kilobytes (KB) in size and a block may include 64 pages, or 256KB. Compared to read operations in a Flash device, an erase operation may have a
30 relatively high latency - which may in turn increase the latency of a corresponding write operation. Programming or reading of Flash technologies may be performed at a lower level of granularity than the erase block size. For example, Flash cells may be programmed or read at a byte, word, or other size.

[0036] A Flash cell experiences wear after repetitive erase-and-program operations. The wear in this case is due to electric charges that are injected and trapped in the dielectric oxide layer between the substrate and the floating gate of the MLC Flash cell. In one example, a MLC Flash cell may have a limit of a number of times it experiences an erase-and-program operation, such as a range from 10,000 to 100,000 cycles. In addition, SSDs may experience program disturb errors that cause a neighboring or nearby Flash cell to experience an accidental state change while another Flash cell is being erased or programmed. Further, SSDs include read disturb errors, wherein the accidental state change of a nearby Flash cell occurs when another Flash cell is being read.

10 [0037] Knowing the characteristics of each of the one or more storage devices 176a-176m may lead to more efficient data object mapping and error detection and correction. In one embodiment, the global RAID engine 178 within the storage controller 174 may detect for the storage devices 176a-176m at least one or more of the following: inconsistent response times for I/O requests, incorrect data for corresponding accesses, error rates and access rates. In response 15 to at least these characteristics, the global RAID engine 178 may determine which RAID data layout architecture to utilize for a corresponding group of storage devices within storage devices 176a-176m. In addition, the global RAID engine 178 may dynamically change both an intra-device redundancy scheme and an inter-device RAID data layout based on the characteristics of the storage devices 176a-176m.

20 [0038] FIG. 1 illustrates an example of a system capable of the described features according to one embodiment. Further details are provided below. Referring to FIG. 1, a further description of the components of network architecture 100 is provided below.

[0039] Components of a Network Architecture

25 [0040] Again, as shown, network architecture 100 includes client computer systems 110a-110c interconnected through networks 180 and 190 to one another and to data storage arrays 120a-120b. Networks 180 and 190 may include a variety of techniques including wireless connection, direct local area network (LAN) connections, storage area networks (SANs), wide area network (WAN) connections such as the Internet, a router, and others. Networks 180 and 190 may 30 comprise one or more LANs that may also be wireless. Networks 180 and 190 may further include remote direct memory access (RDMA) hardware and/or software, transmission control protocol/internet protocol (TCP/IP) hardware and/or software, router, repeaters, switches, grids, and/or others. Protocols such as Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE),

iSCSI, and so forth may be used in networks 180 and 190. Switch 140 may utilize a protocol associated with both networks 180 and 190. The network 190 may interface with a set of communications protocols used for the Internet 160 such as the Transmission Control Protocol (TCP) and the Internet Protocol (IP), or TCP/IP. Switch 150 may be a TCP/IP switch.

5 [0041] Client computer systems 110a-110c are representative of any number of stationary or mobile computers such as desktop personal computers (PCs), workstations, laptops, handheld computers, servers, server farms, personal digital assistants (PDAs), smart phones, and so forth. Generally speaking, client computer systems 110a-110c include one or more processors comprising one or more processor cores. Each processor core includes circuitry for executing 10 instructions according to a predefined general-purpose instruction set. For example, the x86 instruction set architecture may be selected. Alternatively, the Alpha®, PowerPC®, SPARC®, or any other general-purpose instruction set architecture may be selected. The processor cores may access cache memory subsystems for data and computer program instructions. The cache subsystems may be coupled to a memory hierarchy comprising random access memory (RAM) 15 and a storage device.

[0042] Each processor core and memory hierarchy within a client computer system may be in turn connected to a network interface. In addition to hardware components, each of the client computer systems 110a-110c may include a base operating system (OS) stored within the memory hierarchy. The base OS may be representative of any of a variety of specific operating 20 systems, such as, for example, MS-DOS®, MS-WINDOWS®, OS/2®, UNIX®, Linux®, Solaris® or another known operating system. As such, the base OS may be operable to provide various services to the end-user and provide a software framework operable to support the execution of various programs. Additionally, each of the client computer systems 110a-110c may include a hypervisor used to support higher-level virtual machines (VMs). As is well 25 known to those skilled in the art, virtualization may be used in desktops and servers to fully or partially decouple software, such as an OS, from a system's hardware. Virtualization may provide an end-user with an illusion of multiple OSes running on a same machine each having its own resources, such logical storage entities (e.g., logical unit numbers, LUNs) corresponding to the storage devices 176a-176m within each of the data storage arrays 120a-120b.

30 [0043] Each of the data storage arrays 120a-120b may be used for the sharing of data among different servers, such as the client computer systems 110a-110c. Each of the data storage arrays 120a-120b includes a storage subsystem 170 for data storage. Storage subsystem 170 may comprise a plurality of storage devices 176a-176m. Each of these storage devices 176a-176m

may be a SSD. A controller 174 may comprise logic for handling received read/write requests. For example, the algorithms briefly described above may be executed in at least controller 174. A random-access memory (RAM) 172 may be used to batch operations, such as received write requests.

5 [0044] The base OS 132, the file system 134, any OS drivers (not shown) and other software stored in memory medium 130 may provide functionality enabling access to files and LUNs, and the management of these functionalities. The base OS 134 and the OS drivers may comprise program instructions stored on the memory medium 130 and executable by processor 122 to perform one or more memory access operations in storage subsystem 170 that correspond to
10 received requests.

[0045] Each of the data storage arrays 120a-120b may use a network interface 124 to connect to network 180. Similar to client computer systems 110a-110c, in one embodiment, the functionality of network interface 124 may be included on a network adapter card. The functionality of network interface 124 may be implemented using both hardware and software.

15 Both a random-access memory (RAM) and a read-only memory (ROM) may be included on a network card implementation of network interface 124. One or more application specific integrated circuits (ASICs) may be used to provide the functionality of network interface 124.

[0046] In one embodiment, a data storage model may be developed which seeks to optimize data layouts for both user data and corresponding error correction code (ECC) information. In one
20 embodiment, the model is based at least in part on characteristics of the storage devices within a storage system. For example, in a storage system, which utilizes solid-state storage technologies, characteristics of the particular devices may be used to develop a model for the storage system and may also serve to inform corresponding data storage arrangement algorithms. For example, if particular storage devices being used exhibit a change in reliability over time, such a
25 characteristic may be accounted for in dynamically changing a data storage arrangement.

[0047] Generally speaking, any model which is developed for a computing system is incomplete. Often, there are simply too many variables to account for in a real world system to completely model a given system. In some cases, it may be possible to develop models which are not complete but which are nevertheless valuable. As discussed more fully below, embodiments are
30 described wherein a storage system is modeled based upon characteristics of the underlying devices. In various embodiments, selecting a data storage arrangement is performed based on certain predictions as to how the system may behave. Based upon an understanding of the characteristics of the devices, certain device behaviors are more predictable than others.

However, device behaviors may change over time, and in response, a selected data layout may also be changed. As used herein, characteristics of a device may refer to characteristics of the device as a whole, characteristics of a sub-portion of a device such as a chip or other component, characteristics of an erase block, or any other characteristics related to the device.

5

[0048] Intra-Device Redundancy

[0049] Turning now to FIG. 2, a generalized block diagram illustrating one embodiment of a dynamic intra-device redundancy scheme is shown. As is well known to those skilled in the art, one of several intra-device redundancy schemes may be chosen to reduce the effects of latent 10 sector errors in a storage device. The term “sector” typically refers to a basic unit of storage on a HDD, such as a segment within a given track on the disk. Here, the term “sector” may also refer to a basic unit of allocation on a SSD.

[0050] An allocation unit within an SSD may include one or more erase blocks within an SSD. Referring to FIG. 2, the user data 210 may refer to both stored data to be modified and accessed 15 by end-users and inter-device error-correction code (ECC) data. The inter-device ECC data may be parity information generated from one or more pages on other storage devices holding user data. For example, the inter-device ECC data may be parity information used in a RAID data layout architecture. The user data 210 may be stored within one or more pages included within one or more of the storage devices 176a-176k. In one embodiment, each of the storage devices 20 176a-176k is an SSD.

[0051] An erase block within an SSD may comprise several pages. As described earlier, in one embodiment, a page may include 4KB of data storage space. An erase block may include 64 pages, or 256KB. In other embodiments, an erase block may be as large as 1 megabyte (MB), and include 256 pages. An allocation unit size may be chosen in a manner to provide both 25 sufficiently large sized units and a relatively low number of units to reduce overhead tracking of the allocation units. In one embodiment, one or more state tables may maintain a state of an allocation unit (allocated, free, erased, error), a wear level, and a count of a number of errors (correctable and/or uncorrectable) that have occurred within the allocation unit. In various embodiments, the size of an allocation unit may be selected to balance the number of allocation 30 units available for a give device against the overhead of maintaining the allocation units. For example, in one embodiment the size of an allocation unit may be selected to be approximately 1/100th of one percent of the total storage capacity of an SSD. Other amounts of data storage space for pages, erase blocks and other unit arrangements are possible and contemplated.

[0052] Latent sector errors (LSEs) occur when a given sector or other storage unit within a storage device is inaccessible. A read or write operation may not be able to complete for the given sector. In addition, there may be an uncorrectable error-correction code (ECC) error. An LSE is an error that is undetected until the given sector is accessed. Therefore, any data 5 previously stored in the given sector may be lost. A single LSE may lead to data loss when encountered during RAID reconstruction after a storage device failure. For an SSD, an increase in the probability of an occurrence of another LSE may result from at least one of the following statistics: device age, device size, access rates, storage compactness and the occurrence of previous correctable and uncorrectable errors. To protect against LSEs and data loss within a 10 given storage device, one of a multiple of intra-device redundancy schemes may be used within the given storage device.

[0053] An intra-device redundancy scheme utilizes ECC information, such as parity information, within the given storage device. This intra-device redundancy scheme and its ECC information corresponds to a given device and may be maintained within a given device, but is distinct from 15 ECC that may be internally generated and maintained by the device itself. Generally speaking, the internally generated and maintained ECC of the device is invisible to the system within which the device is included. The intra-device ECC information included within the given storage device may be used to increase data storage reliability within the given storage device. This intra-device ECC information is in addition to other ECC information that may be included 20 within another storage device such as parity information utilized in a RAID data layout architecture.

[0054] A highly effective intra-device redundancy scheme may sufficiently enhance a reliability of a given RAID data layout to cause a reduction in a number of devices used to hold parity information. For example, a double parity RAID layout may be replaced with a single parity 25 RAID layout if there is additional intra-device redundancy to protect the data on each device. For a fixed degree of storage efficiency, increasing the redundancy in an intra-device redundancy scheme increases the reliability of the given storage device. However, increasing the redundancy in such a manner may also increase a penalty on the input/output (I/O) performance of the given storage device.

30 [0055] In one embodiment, an intra-device redundancy scheme divides a device into groups of locations for storage of user data. For example, a division may be a group of locations within a device that correspond to a stripe within a RAID layout as shown by stripes 250a-250c. User data or inter-device RAID redundancy information may be stored in one or more pages within

each of the storage devices 176a-176k as shown by data 210. Within each storage device, intra-device error recovery data 220 may be stored in one or more pages. As used herein, the intra-device error recovery data 220 may be referred to as intra-device redundancy data 220. As is well known by those skilled in the art, the intra-device redundancy data 220 may be obtained by

5 performing a function on chosen bits of information within the data 210. An XOR-based operation may be used to derive parity information to store in the intra-device redundancy data 220. Other examples of intra-device redundancy schemes include single parity check (SPC), maximum distance separable (MDS) erasure codes, interleaved parity check codes (IPC), hybrid SPC and MDS code (MDS+SPC), and column diagonal parity (CDP). The schemes vary in
10 terms of delivered reliability and overhead depending on the manner the data 220 is computed. In addition to the above described redundancy information, the system may be configured to calculate a checksum value for a region on the device. For example, a checksum may be calculated when information is written to the device. This checksum is stored by the system. When the information is read back from the device, the system may calculate the checksum again
15 and compare it to the value that was stored originally. If the two checksums differ, the information was not read properly, and the system may use other schemes to recover the data. Examples of checksum functions include cyclical redundancy check (CRC), MD5, and SHA-1.

[0056] As shown in stripes 250a-250c, the width, or number of pages, used to store the data 210 within a given stripe may be the same in each of the storage devices 176a-176k. However, as
20 shown in stripes 250b-250c, the width, or number of pages, used to store the intra-device redundancy data 220 within a given stripe may not be the same in each of the storage devices 176a-176k. In one embodiment, changing characteristics or behaviors of a given storage device may determine, at least in part, the width used to store corresponding intra-device redundancy data 220. For example, as described above, Flash cells experience program disturb errors and
25 read disturb errors, wherein programming or reading a page may disturb nearby pages and cause errors within these nearby pages. When a storage device is aging and producing more errors, the amount of corresponding intra-device redundancy data 220 may increase. For example, prior to a write operation for stripe 250b, characteristics of each of the storage devices 176a-176k may be monitored and used to predict an increasing error rate. A predicted increase in errors for storage
30 devices 176c and 176j may be detected. In response, the amount of intra-device redundancy data 220 may be increased for storage devices 176c and 176j. In the example of stripes 250a and 250b of FIG. 2, an increase in the amount of protection data stored can be seen for storage devices 176c and 176j for stripes 250a and 250b. For example, now, rather than protecting

storage devices 176c and 176j with single parity, these devices may be protected with double parity or triple parity. It is noted that increasing the amount of intra-device protection for devices 176c and 176j does not necessitate a corresponding increase in other devices of the same stripe. Rather, data for the stripe may have differing levels of protection in each device as 5 desired.

[0057] In various embodiments, increases or decreases in a given level of data protection may occur on a selective basis. For example, in one embodiment, an increase in protection may occur only for storage devices that are detected to generate more errors, such as storage devices 176c and 176j in the above example. In another embodiment, an increase in protection may occur for 10 each of the storage devices 176a-176k when storage devices 176c and 176j are detected to generate more errors. In one embodiment, increasing the amount of intra-device protection on a parity device such as device 176k may require a reduction in the amount of data protected within the stripe. For example, increasing the amount of intra-device data stored on a parity device for a given stripe will necessarily reduce an amount of parity data stored by that device for data within 15 the stripe. If this amount of parity data is reduced to an amount that is less than that needed to protect all of the data in the stripe, then data within the stripe must be reduced if continued parity protection is desired. As an alternative to reducing an amount of data stored within the stripe, a different device could be selected for storing the parity data. Various options are possible and are contemplated. It is also noted that while FIG. 2 and other figures described herein may depict a 20 distinct parity device (e.g., 176k), in various embodiments the parity may be distributed across multiple devices rather than stored in a single device. Accordingly, the depiction of a separate parity device in the figures may generally be considered a logical depiction for ease of discussion.

[0058] Referring now to FIG. 3, one embodiment of a method 300 for adjusting intra-device 25 protection in a data storage subsystem is shown. The components embodied in network architecture 100 and data storage arrays 120a-120b described above may generally operate in accordance with method 300. The steps in this embodiment are shown in sequential order. However, some steps may occur in a different order than shown, some steps may be performed concurrently, some steps may be combined with other steps, and some steps may be absent in 30 another embodiment.

[0059] In block 302, a first amount of space for storing user data in a storage device is determined. This user data may be data used in end-user applications or inter-device parity information used in a RAID architecture as described earlier regarding data 210. This first

amount of space may comprise one or more pages within a storage device as described earlier.

In one embodiment, a global RAID engine 178 within the storage controller 174 receives behavioral statistics from each one of the storage devices 176a-176m. For a given device group comprising two or more of the storage devices 176a-176m, the global RAID engine 178 may

5 determine both a RAID data layout and an initial amount of intra-device redundancy to maintain within each of the two or more storage devices. In block 304, the RAID engine 178 may determine a second amount of space for storing corresponding intra-device protection data in a storage device. This second amount of space may comprise one or more pages within a storage device. The intra-device protection data may correspond to the to intra-device redundancy data

10 220 described earlier.

[0060] In block 306, data is written in the first amount of space within each storage device included within a given device group. In one embodiment, both user data and inter-device parity information is written as a single RAID stripe across multiple storage devices included within the given device group. Referring again to FIG. 2, the width for the corresponding data being

15 written is the same within each storage device. In block 308, the intra-device protection data is generated by an ECC algorithm, an XOR-based algorithm, or any other suitable algorithm. In addition, the system may generate a checksum to help identify data that has not been retrieved properly. In block 310, the generated intra-device protection data is written in the second amount of space in the storage devices.

20 **[0061]** In block 312, the RAID engine 178 may monitor behavior of the one or more storage devices. In one embodiment, the RAID engine 178 may include a model of a corresponding storage device and receive behavioral statistics from the storage device to input to the model. The model may predict behavior of the storage device by utilizing known characteristics of the storage device. For example, the model may predict an upcoming increasing error rate for a

25 given storage device. If the RAID engine 178 detects characteristics of a given storage device which affect reliability (conditional block 314), then in block 316, the RAID engine may adjust the first amount and the second amount of space for storing data and corresponding intra-device redundancy data. For example, the RAID engine may be monitoring the statistics described earlier such as at least device age, access rate and error rate. Referring again to FIG. 2, the RAID

30 engine 178 may detect storage devices 176c and 176j have an increase in a number of errors. Alternatively, the RAID engine may predict an increase in a number of errors for storage devices 176c and 176j. Accordingly, prior to writing the second stripe 250b, the RAID engine 178 may adjust a number of pages used to store data 210 and data 220 in each of the storage devices 176a-

176k. Similarly, the RAID engine 178 may detect storage device 176b has decreased reliability. Therefore, prior to writing the third stripe 250c, the RAID engine 178 may again adjust a number of pages used to store data 210 and data 220 in each of the storage devices 176a-176k.

5 **[0062] Monitoring Storage Device Characteristics**

[0063] Turning now to FIG. 4, a generalized block diagram of one embodiment of a storage subsystem is shown. Each of the one or more storage devices 176a-176m may be partitioned in one of one or more device groups 173a-173m. Other device groups with other devices may be present as well. One or more corresponding operation queues and status tables for each storage 10 device may be included in one of the device units 400a-400w. These device units may be stored in RAM 172. A corresponding RAID engine 178a-178m may be included for each one of the device groups 173a-173m. Each RAID engine 178 may include a monitor 410 that tracks statistics for each of the storage devices included within a corresponding device group. Data layout logic 420 may determine an amount of space to allocate within a corresponding storage 15 device for user data, inter-device redundancy data and intra-device redundancy data. The storage controller 174 may comprise other control logic 430 to perform at least one of the following tasks: wear leveling, garbage collection, I/O scheduling, deduplication and protocol conversion for incoming and outgoing packets.

[0064] Turning now to FIG. 5, a generalized block diagram of one embodiment of a device unit 20 is shown. A device unit may comprise a device queue 510 and tables 520. Device queue 510 may include a read queue 512, a write queue 514 and one or more other queues such as other operation queue 516. Each queue may comprise a plurality of entries for storing one or more corresponding requests 530a-530d. For example, a device unit for a corresponding SSD may include queues to store at least read requests, write requests, trim requests, erase requests and so 25 forth. Tables 520 may comprise one or more state tables 522a-522b, each comprising a plurality of entries for storing state data, or statistics, 530. It is also noted that while the queues and tables are shown to include a particular number of entries in this and other figures, the entries themselves do not necessarily correspond to one another. Additionally, the number of queues, tables, and entries may vary from that shown in the figure and may differ from one another.

30 [0065] Referring now to FIG. 6, a generalized block diagram illustrating one embodiment of a state table corresponding to a given device is shown. In one embodiment, such a table may include data corresponding to state, error and wear level information for a given storage device, such as an SSD. A corresponding RAID engine may have access to this information, which may

allow the RAID engine to dynamically change space allocated for data storage and schemes used for both inter-device protection and intra-device protection. In one embodiment, the information may include at least one or more of a device age 602, an error rate 604, a total number of errors detected on the device 606, a number of recoverable errors 608, a number of 5 unrecoverable errors 610, an access rate of the device 612, an age of the data stored 614 and one or more allocation states for allocation spaces 616a-616n. The allocation states may include filled, empty, error and so forth.

[0066] Flexible RAID Layout

10 **[0067]** Turning now to FIG. 7, a generalized block diagram illustrating one embodiment of a flexible RAID data layout architecture is shown. A RAID engine may determine a level of protection to use for storage devices 176a-176k. For example, a RAID engine may determine to utilize RAID double parity for the storage devices 176a-176k. The inter-device redundancy data 240 may represent the RAID double parity values generated from corresponding user data. In 15 one embodiment, storage devices 176j and 176k may store the double parity information. It is understood other levels of RAID parity protection are possible and contemplated. In addition, in other embodiments, the storage of the double parity information may rotate between the storage devices rather than be stored within storage devices 176j and 176k for each RAID stripe. The storage of the double parity information is shown to be stored in storage devices 176j and 176k 20 for ease of illustration and description.

25 **[0068]** Referring now to FIG. 8, a generalized block diagram illustrating another embodiment of a flexible RAID data layout architecture is shown. Similar to the example shown in FIG. 7, double parity may be used for the storage devices 176a-176k. Although a RAID double parity is described in this example, any amount of redundancy in a RAID data layout architecture may be chosen.

30 **[0069]** During operation, the RAID engine 178 may monitor characteristics of the storage devices 176a-176k and determine the devices are exhibiting a reliability level higher than an initial or other given reliability level. In response, the RAID engine 178 may change the RAID protection from a RAID double parity to a RAID single parity. In other RAID data layout architectures, another reduction in the amount of supported redundancy may be used. In other embodiments, the monitoring of storage devices 176a-176k and changing a protection level may be performed by other logic within storage controller 174.

[0070] Continuing with the above example, only single parity information may be generated and stored for subsequent write operations executing on a given RAID stripe. For example, storage device 176k may not be used in subsequent RAID stripes for write operations after the change in the amount of supported redundancy. In addition, data stored in storage device 176k may be 5 invalidated, thereby freeing the storage. Pages corresponding to freed data in storage device 176k may then be reallocated for other uses. The process of reducing an amount of parity protection and freeing space formerly used for storing parity protection data may be referred to as “parity shredding”. In addition, in an embodiment wherein storage device 176k is an SSD, one or more erase operations may occur within storage device 176k prior to rewriting the pages within 10 stripe 250a.

[0071] Continuing with the above example of parity shredding, the data stored in the reallocated pages of storage device 176k within stripe 250a after parity shredding may hold user data or corresponding RAID single parity information for other RAID stripes that do not correspond to stripe 250a. For example, the data stored in storage devices 176a-176j within stripe 250a may 15 correspond to one or more write operations executed prior to parity shredding. The data stored in storage device 176k within stripe 250a may correspond to one or more write operations executed after parity shredding. Similarly, the data stored in storage devices 176a-176j within stripe 250b may correspond to one or more write operations executed prior to parity shredding. The pages in storage device 176k within stripe 250b may be freed, later erased, and later rewritten with data 20 corresponding to one or more write operations executed after the change in the amount of supported redundancy. It is noted that this scheme may be even more effective when redundancy information is rotated across storage devices. In such an embodiment, space that is freed by shredding will likewise be distributed across the storage devices.

[0072] Referring again to FIG. 8, the deallocated pages shown in storage device 176k within stripe 250c represent storage locations that may have previously stored RAID double parity information prior to parity shredding. However, now these pages are invalid and have not yet been reallocated. Particular characteristics of an SSD determine the manner and the timing of both freeing and reallocating pages within storage device 176k in the above example. Examples of these characteristics include at least erasing an entire erase block prior to reprogramming 30 (rewriting) one or more pages. As can be seen from FIG. 8, when parity is shredded, it is not necessary to shred an entire device. Rather, parity may be shredded for individual stripes as desired. Similarly, parity protection for a stripe may be increased by adding protection data stored on an additional device to a stripe.

[0073] Referring now to FIG. 9, one embodiment of a method for dynamically determining a RAID layout is shown. The components embodied in network architecture 100 and data storage arrays 120a-120b described above may generally operate in accordance with method 900. In FIG. 9, two processes 910 and 920 are shown. Each of the processes may operate concurrently, or in a given order. Further, the steps in this embodiment are shown in sequential order. However, some steps may occur in a different order than shown, some steps may be performed concurrently, some steps may be combined with other steps, and some steps may be absent in another embodiment. Block 910 illustrates a process whereby a storage control system monitors the characteristics and behaviors of storage devices in the system (block 912). For example, characteristics such as those described in FIG. 6 may be observed and/or recorded. If a particular condition is detected, such as a change in reliability (decision block 914), then a change in the amount of protection used for stored data may be made (block 916). For example, when given devices are relatively young in age, the reliability of the devices may not be known (e.g., the devices may suffer “infant mortality” and fail at a relatively young age). Therefore, one or more extra storage devices per RAID stripe may be used to store parity information. At a later time, this extra protection may be removed when the devices prove over time that they are reliable. In various embodiments, characteristics regarding error rates may be maintained for devices. For example, characteristics concerning correctable and/or uncorrectable errors may be maintained and used to make decisions regarding the reliability of a given device. Based upon this information, the storage controller may dynamically alter various levels of protection for a device or stripe.

[0074] Block 920 of FIG. 9 generally illustrates a process whereby at the time a stripe or other portion of storage is to be allocated (decision block 922), a determination regarding the layout and protection level to use for the data may be made (block 924). It is noted that the process of block 910 could be performed at this time. Alternatively, levels of protection may have been determined by process 910 and stored. The determination of block 924 could then be based upon that stored data. In one embodiment, once a given layout has been determined, the particular devices to be used for the layout may be selected from a group of devices (block 925). For example, in one embodiment a group of 20 devices may be available for use. If a layout of 5+2 is determined, then any seven devices may be selected for use from the group of 20. Additionally, it is noted that a subsequent write with a selected 5+2 layout need not use the same 7 devices. Subsequent to determining the layout, protection level, and devices for the stripe, the stripe may be written (block 926).

[0075] In various embodiments, the RUSH algorithm may be utilized to determine which devices on which the data and redundancy information for a given stripe will reside . For example, the RUSH algorithm may be used to select the particular devices to utilize for an 8+2 RAID layout for a given stripe in storage devices 176a-176k. Generally speaking, as used 5 herein, an M+N layout may generally describe a layout which includes M data devices and N parity devices for a given data stripe. Additionally, as discussed above, parity may be distributed across the devices rather than fully located within particular devices. Accordingly, an 8+2 layout may include data and parity striped across 10 devices – with 8 of the devices storing data and two of the devices storing parity. On a subsequent occasion, a layout of 12+2 may be selected. In this 10 manner, the desired layout and protection characteristics may be determined dynamically at the time a write (e.g., a stripe) is to be written. In one embodiment, storage devices 176a-176k may include more than 10 storage devices, such as 30, 50 or more storage devices. However, for a stripe with an 8+2 layout, only 10 of the storage devices are utilized . It is noted that any 10 of the devices may be selected and any suitable algorithm may be used for selecting the 10 devices 15 for use in storing the stripe. For example, the CRUSH algorithm could be used to select which 10 of the storage devices 176a-176k to utilize for a given 8+2 RAID layout.

[0076] In one example of a chosen 8+2 RAID layout for storage devices 176a-176k, 2 of the storage devices may be used to store error correcting code (ECC) information, such as parity information. This information may be used to perform reconstruct read requests. Referring 20 again to FIG. 8, the storage devices 176j and 176k may be selected to store RAID double parity information in this example. Again, the parity information may be stored in a rotated fashion between each of the storage devices 176a-176k included within the RAID array, rather than consistently stored in the same storage devices. For ease of illustration and description, the storage devices 176j and 176k are described as storing RAID double parity.

[0077] In block 926, during execution of a write operation, metadata, user data, intra-device 25 parity information and inter-device parity information may be written as a RAID stripe across multiple storage devices included within the RAID array. In block 912, the RAID engine 178 may monitor behavior of the one or more storage devices within the RAID array. In one embodiment, the RAID engine 178 may include a monitor 410 and data layout logic 420 as 30 shown in FIG. 4. The RAID engine 178 may monitor at least an age of a given storage device, a number and a type of errors, detected configuration changes since a last allocation of data, an age of given data, a current usage of storage space in the RAID array, and so forth.

[0078] The data, which is monitored by the RAID engine 178, may be stored in RAM 172, such as in one of the device units 400a-400w shown in FIG. 4. Tables may be used to store this data, such as the examples shown in FIG. 5 and FIG. 6. The logic included within a corresponding RAID engine may both detect and predict behavior of storage devices by monitoring updated 5 statistics of the storage devices. For example, the model may predict an upcoming increasing error rate for a given storage device.

[0079] If increased reliability of the storage device(s) is detected (conditional block 908), then in block 910, the RAID engine may decrease the level of data protection within the system. For example, in one embodiment the amount of parity information stored in the storage subsystem 10 may be reduced. Regarding the above example, the RAID engine may decrease the RAID double parity to RAID single parity for the corresponding 8+2 RAID array, converting it to an 8+1 RAID array. In other examples a given RAID array may be utilizing an N-level amount of redundancy, or parity, in a RAID architecture prior to block 916. In block 916, the RAID engine may determine to utilize an (N-m)-level amount of redundancy, wherein $N > 1$ and $1 \leq m < N$. 15 Therefore, during subsequent write operations for a given RAID stripe, there will be m fewer storage devices written to within the given RAID stripe.

[0080] In order to reduce the level of data protection within the system, the RAID engine (or another component) may perform parity shredding as described earlier. Subsequently, the storage controller 174 may reallocate those pages which were freed as a result of the shredding 20 operation to be used in subsequent write operations.

[0081] As each of the storage devices 176a-176k both age and fill up with data, extra parity information may be removed from the RAID array as described above. The metadata, the user data, corresponding intra-device redundancy information and some of the inter-device redundancy information remains. Regarding the above example with an 8+2 RAID array, the 25 information stored in storage devices 176a-176j remains. However, extra inter-device redundancy information, or extra parity information, may be removed from the RAID array. For example, extra parity information stored in storage device 176k may be removed from the RAID stripes.

[0082] The information that remains, such as the information stored in storage devices 176a-176j 30 in the above example, may remain in place. The storage space storing the extra parity information, such as the corresponding pages in storage device 176k in the above example, may be reused and reallocated for subsequent write operations. In one embodiment, each new allocation receives a new virtual address. Each new allocation may have any given size, any

given alignment or geometry, and may fit in any given storage space (either virtual or physical).

In one embodiment, each one of the storage devices 176a-176k and each allocated page within a storage device have a header comprising identification information. This identification information may allow the reuse of storage space for freed extra parity information without 5 changing a given configuration.

[0083] In an embodiment wherein one or more of the storage devices 176a-176k is an SSD, an erase block is erased prior to reprogramming one or more pages within the erase block. Therefore, in an embodiment wherein storage device 176k is an SSD, corresponding erase blocks are erased prior to reprogramming freed pages in storage device 176k. Regarding the above 10 example with an original 8+2 RAID array, one or more erase blocks are erased in storage device 176k within stripes 250a-250b prior to reprogramming pages with data 210. The original 8+2 RAID array is now an 8+1 RAID array with storage device 176j providing the single parity information for RAID stripes written prior to the parity shredding.

[0084] As is well known to those skilled in the art, during a read or write failure for a given 15 storage device, data may be reconstructed from the supported inter-device parity information within a corresponding RAID stripe. The reconstructed data may be written to the storage device. However, if the reconstructed data fails to be written to the storage device, then all the data stored on the storage device may be rebuilt from corresponding parity information. The rebuilt data may be relocated to another location. With Flash memory, a Flash Translation Layer 20 (FTL) remaps the storage locations of the data. In addition, with Flash memory, relocation of data includes erasing an entire erase block prior to reprogramming corresponding pages within the erase block. Maintaining mapping tables at a granularity of erase blocks versus pages allows the remapping tables to be more compact. Further, during relocation, extra pages that were freed during parity shredding may be used.

25

[0085] Offset Parity

[0086] Turning now to FIG. 10, a generalized block diagram illustrating yet another embodiment 30 of a flexible RAID data layout architecture is shown. Similar to the generalized block diagram shown in FIG. 8, a flexible RAID data layout architecture may be used. The storage devices 176a-176k comprise multiple RAID stripes laid out across multiple storage devices. Although each of the storage devices 176a-176k comprises multiple pages, only page 1010 and page 1020 are labeled for ease of illustration. In the example shown, a double parity RAID data layout is chosen, wherein storage devices 176j and 176k store double parity information.

[0087] Each of the pages in the storage devices 176a-176k stores a particular type of data. Some pages store user data 210 and corresponding generated inter-device parity information 240. Other pages store corresponding generated intra-device parity information 220. Yet other pages store metadata 242. The metadata 242 may include page header information, RAID stripe 5 identification information, log data for one or more RAID stripes, and so forth. In addition to inter-device parity protection and intra-device parity protection, each of the pages in storage devices 176a-176k may comprise additional protection such as a checksum stored within each given page. In various embodiments, the single metadata page at the beginning of each stripe may be rebuilt from the other stripe headers. Alternatively, this page could be at a different offset 10 in the parity shard so the data can be protected by the inter-device parity. A “shard” represents a portion of a device. Accordingly, a parity shard refers to a portion of a device storing parity data.

Physical Layer

[0088] In various embodiments, the systems described herein may include a physical layer 15 through which other elements of the system communicate with the storage devices. For example, scheduling logic, RAID logic, and other logic may communicate with the storage devices via a physical layer comprising any suitable combination of software and/or hardware. In general, the physical layer performs a variety of functions including providing access to persistent storage, and performing functions related to integrity of data storage.

[0089] FIG. 11A illustrates one embodiment of a hypothetical device layout for a 500GB device. In various embodiments, the storage devices described herein may be formatted with a partition 20 table 1101 at the beginning of the device, and a copy of the partition table at the end of the device. Additionally, a device header 1103 may be stored in the first and last blocks. For example, in a flash based storage device, a device header may be stored in the first and last erase 25 blocks. As previously discussed, an erase block is a flash construct that is typically in the range of 256KB-1MB. Additional unused space in the first erase block may be reserved (padding 1105). The second erase block in each device may be reserved for writing logging and diagnostic information 1107. The rest of the erase blocks in between are divided into Allocation Units (AUs) 1109 of a multiple erase blocks. The AU size may be chosen so there are a reasonable 30 number of AUs per device for good allocation granularity. In one embodiment, there may be something in the range of 10,000 AUs on a device to permit allocation in large enough units to avoid overhead, but not too many units for easy tracking. Tracking of the state of an AU (allocated/free/erased/bad) may be maintained an AU State Table. The wear level of an AU may

be maintained in a Wear Level Table, and a count of errors may be maintained in an AU Error Table.

[0090] In various embodiments, the physical layer allocates space in segments which include one segment shard in each device across a set of devices (which could be on different nodes). FIG.

5 11B depicts one embodiment of a segment and various identifiable portions of that segment in one possible segment layout. In the embodiment shown, a single segment is shown stored in multiple devices. Illustrated are data devices Data 0 – Data N, and parity devices Parity P and Parity Q. In one embodiment, each segment shard includes one or more allocation units on a device such that the size of the shard is equal on each device. Segment shard 1123 is called out to 10 illustrate a segment shard. Also illustrated in FIG. 11B, is an I/O read size 1127 which in one embodiment corresponds to a page. Also shown is an I/O parity chunk 1129 which may include one or more pages of page parity for the I/O shard.

[0091] In one embodiment, each segment will have its own geometry which may include one or more of the following parameters:

15 (1) RAID level - The RAID level used for cross device protection in the segment. This may determine mirroring, parity, or ECC RAID and how many segment shards contain parity.

20 (2) Device Layout I/O shard size - This represents the size used to stripe across each device during a write. This will typically be in the range of 256KB to 1MB and probably be a multiple of the erase block size on each device. FIG. 11B calls out I/O shard size 1125 for purposes of illustration.

25 (3) I/O read size - This is a logical read size. Each I/O shard may be formatted as a series of logical pages. Each page may in turn include a header and a checksum for the data in the page. When a read is issued it will be for one or more logical pages and the data in each page may be validated with the checksum.

30 (4) I/O shard RAID level - The I/O shard has intra-shard parity to handle latent errors found during a rebuild. This parameter determines what type of parity is used for intra-shard protection and thus how many copies of the intra-shard parity will be maintained.

(5) I/O parity chunk – In various embodiments, the storage devices may do ECC on a page basis. Consequently, if an error is seen it is likely to indicate failure of an entire physical page. The I/O parity chunk is the least common multiple of the physical page size on each device in the segment and the intra-shard parity is calculated by striping

down the I/O shard in the larger of the I/O parity chunks or the I/O read size.

Included may be one or more pages of page parity. In various embodiments, this parity may be used to rebuild data in the event of a failed checksum validation.

[0092] In various embodiments, as each new segment is written a RAID geometry for the segment will be selected. Selection of the RAID geometry may be based on factors such as the current set of active nodes and devices, and the type of data in the segment. For example if 10 nodes or devices are available then an (8+2) RAID 6 geometry may be chosen and the segment striped across the nodes to withstand two device or node failures. If a node then fails, the next segment may switch to a (7+2) RAID 6 geometry. Within the segment some of the segment shards will contain data and some will contain ECC (e.g., parity).

[0093] In one embodiment, there are five types of segments. Three of these segments correspond to the AU State Table, the AU Error Table, and the Wear Level Table. In some embodiments, these three segments may be mirrored for additional protection. In addition to these three segments, there are metadata segments which may also be additionally protected through mirroring. Finally there are Data segments which hold client blocks and log information. The log information contains update information associated with the client blocks in the segment. The data segments will likely be protected by RAID 6 as illustrated in FIG. 11B with Parity P and Parity Q shards. In addition to the above, a segment table is maintained as an in memory data structure that is populated at startup with information from the headers of all the segment shards.

20 In some embodiments, the table may be cached completely on all nodes so any node can translate a storage access to a physical address. However, in other embodiments an object storage model may be used where each node may have a segment table that can take a logical reference and identify the segment layout node where the data is stored. Then the request would be passed to the node to identify the exact storage location on the node. FIG. 11B also depicts segment tail data which identifies any (volume, snapshot) combinations that take up a significant amount of space in the segment. When snapshots are removed, a data scrubber may help identify segments for garbage collection based on this data.

25 **[0094]** In one embodiment, the basic unit of writing is the segio which is one I/O shard on each of the devices in the segment. Each logical page in the segio is formatted with a page header that contains a checksum (which may be referred to as a “media” checksum) of the page so the actual page size for data is slightly smaller than one page. For pages in the parity shards of a segment the page header is smaller so that the page checksums in the data page are protected by the parity page. The last page of each I/O shard is a parity page that again has a smaller header and protects

all the checksums and page data in the erase block against a page failure. The page size referred to here is the I/O read size which may be one or more physical flash pages. For some segments, a read size smaller than a physical page may be used. This may occur for metadata where reads to lookup information may be index driven and smaller portion of data may be read while still

5 obtaining the desired data. In such a case, reading half a physical page would mean tying up the I/O bus (and network) with less data and validating (e.g., checksumming) less data. To support a read size smaller than a physical page, an embodiment may include multiple parity pages at the end of the erase block such that the total size of all the parity pages is equal to the flash page size.

[0095] As the wear level of an erase block increases, the likelihood of an error increases. In 10 addition to tracking wear levels, data may be maintained regarding observed how often errors are seen on an erase block and blocks with a higher probability of error identified. For some erase blocks, it may be decided to keep double or triple error correcting parity at the end of the erase block instead of the single RAID 5 parity. In this case, the data payload of the segio may be reduced accordingly. It may only be necessary to reduce the poor erase block within the segio, 15 rather than all the erase blocks. The page headers in the erase block may be used to identify which pages are parity and which are data.

[0096] Whenever a page is read from storage, the contents may be validated using the page checksum. If the validation fails, a rebuild of the data using the erase block parity may be attempted. If that fails, then cross device ECC for the segment may be used to reconstruct the 20 data.

[0097] In data segments the payload area may be divided into two areas. There will be pages 25 formatted as log data which may include updates related to stored client blocks. The remainder of the payload area may contain pages formatted as client blocks. The client block data may be stored in a compressed form. Numerous compression algorithms are possible and are contemplated. Additionally, in various embodiments Intel® Advanced Encryption Standard instructions may be used for generating checksums. Additionally, there may be a header for the client block that resides in the same page as the data and contains information needed to read the client block, including an identification of the algorithm used to compress the data. Garbage collection may utilize both the client block header and the log entries in the segio. In addition, 30 the client block may have a data hash which may be a checksum of the uncompressed data used for deduplication and to check the correctness of the decompressed data.

[0098] In some embodiments, segments and segios may have a monotonically increasing ID number used to order them. As part of writing a segio, a logical layer can record dependencies

on prior flushes. At startup, the physical layer may build an ordered list of segments and segios and if a segio is dependent on another uncompleted segio it may be rolled back and not considered to have been written.

5 Wear Level Table

[0099] The Wear Level Table (WLT) for each device may be stored in a segment local to each device. The information may also be stored in the header of each segment shard. In one embodiment, the wear information is an integer that represents the number of times the allocation unit has been erased and reused. As the wear information may not be accurate, a flush of the 10 table to the device may be performed when there has been a certain amount of activity or when the system has been idle for a reasonable period. The WLT may also be responsible for cleaning up old WLT segments as it allocates new ones. To add an extra layer of protection, old copies may be maintained before freeing them. For example, a table manager may ensure that it keeps the previous erase block and the current erase block of WLT entries at all times. when it allocates 15 a new segment it won't free the old segment until it has written into the second erase block of the new segment.

AU State Table

[00100] The AU State Table (AST) tracks the state of each AU. The states include Free, 20 Allocated, Erased and Bad. The AST may be stored in a segment on the device. Changing a state to Allocated or Free may be a synchronous update, while changing a state to Bad or Erased may be an asynchronous update. This table may generally be small enough and have enough updates that updates may be logged in NVRAM. The AST may be responsible for cleaning up old AST segments as it allocates new ones. Since the AST can be completely recovered by scanning the 25 first block of each AU on the drive, there is no need to keep old copies of the AST.

AU Error Table

[00101] The AU Error Table (AET) may be used to track the number of recoverable errors and 30 unrecoverable errors within each AU. The AET is stored in a segment on the device and each field may be a two byte integer. With four bytes per AU the entire table may be relatively small.

[00102] Referring now to FIG. 11C, a generalized block diagram illustrating one embodiment of data storage arrangements within different page types is shown. In the embodiment shown,

three page types are shown although other types are possible and contemplated. The shown page types include page 1110 comprising metadata 1150, page 1120 comprising user data 1160, and page 1130 comprising parity information 1170 (inter-device or intra-device). Each of the pages 1110-1130 comprises metadata 1140, which may include header and identification information. In addition, each of the pages 1110-1130 may comprise intra-page error recovery data 1142, such as a corresponding checksum or other error detecting and/or correcting code. This checksum value may provide added protection for data stored in storage devices 176a-176k in a given device group.

[00103] Further, page 1130 may comprise inter-page error recovery data 1144. The data 1144 may be ECC information derived from the intra-page data 1142 stored in other storage devices. For example, referring again to FIG. 10, each page within storage device 176j, which stores inter-device parity information 240, may also store inter-page error recovery data 1144. The data 1144 may be a parity, a checksum, or other value generated from intra-page error recovery data 1142 stored in one or more of the storage devices 176a-176i. In one embodiment, the data 1144 is a checksum value generated from one or more other checksum values 1142 stored in other storage devices. In order to align data 1144 in a given page in storage device 176j with data 1142 in a corresponding page in one or more of the storage devices 176a-176i, padding 1146 may be added to the corresponding pages.

[00104] In one embodiment, end-user applications perform I/O operations on a sector-boundary, wherein a sector is 512 bytes for HDDs. In order to add extra protection, an 8-byte checksum may be added to form a 520-byte sector. In various embodiments, compression and remapping may be used in a flash memory based system to allow user data to be arranged on a byte boundary rather than a sector boundary. In addition, a checksum (8 byte, 4 byte, or otherwise) may be placed inside a page after a header and before the user data, which may be compressed. This placement is shown in each of pages 1110-1130.

[00105] When an end-user application reads a 512-byte sector, a corresponding page, which may be 2KB-8KB in size in one embodiment, has extra protection with an 8-byte checksum at the beginning of the page. In various embodiments, the page may not be formatted for a non-power of 2 sector size. As shown in pages 1110-1120, the checksum may be offset a few bytes into the page. This offset allows a parity page, such as page 1130, to store both a checksum that covers the parity page and ECC to protect checksums of the other pages.

[00106] For yet another level of protection, data location information may be included when calculating a checksum value. The data 1142 in each of pages 1110-1130 may include this

information. This information may include both a logical address and a physical address. Sector numbers, data chunk and offset numbers, track numbers, plane numbers, and so forth may be included in this information as well.

5 [00107] Alternate Geometries

[00108] Turning now to FIG. 12, a generalized block diagram illustrating one embodiment of a hybrid RAID data layout 1200 is shown. Three partitions are shown although any number of partitions may be chosen. Each partition may correspond to a separate device group, such as device groups 713a-173b shown in FIG. 1. Each partition comprises multiple storage devices.

10 In one embodiment, an algorithm such as the CRUSH algorithm may be utilized to select which devices to use in a RAID data layout architecture to use for data storage.

[00109] In the example shown, an L+1 RAID array, M+1 RAID array, and N+1 RAID array are shown. In various embodiments, L, M, and N may all be different, the same, or a combination thereof. For example, RAID array 1210 is shown in partition 1. The other storage devices 1212 are candidates for other RAID arrays within partition 1. Similarly, RAID array 1220 illustrates a given RAID array in partition 2. The other storage devices 1222 are candidates for other RAID arrays within partition 2. RAID array 1230 illustrates a given RAID array in partition 3. The other storage devices 1232 are candidates for other RAID arrays within partition 3.

[00110] Within each of the RAID arrays 1210, 1220 and 1230, a storage device P1 provides RAID single parity protection within a respective RAID array. Storage devices D1-DN store user data within a respective RAID array. Again, the storage of both the user data and the RAID single parity information may rotate between the storage devices D1-DN and P1. However, the storage of user data is described as being stored in devices D1-DN. Similarly, the storage of RAID single parity information is described as being stored in device P1 for ease of illustration and description.

[00111] One or more storage devices among each of the three partitions may be chosen to provide an additional amount of supported redundancy for one or more given RAID arrays. For example, storage device Q1 in partition 3 may be combined with each of the RAID arrays 1210, 1220 and 1230. The storage device Q1 may provide RAID double parity information for each of the RAID arrays 1210, 1220 and 1230. This additional parity information is generated and stored when a stripe is written to one of the arrays 1210, 1220, or 1230. Further this additional parity information may cover stripes in each of the arrays 1210, 1220, and 1230. Therefore, the ratio of a number of storage devices storing RAID parity information to a total number of storage devices

is lower. For example, if each of the partitions used N+2 RAID arrays, then the ratio of a number of storage devices storing RAID parity information to a total number of storage devices is $3(2)/(3(N+2))$, or $2/(N+2)$. In contrast, the ratio for the hybrid RAID layout 1200 is $(3+1)/(3(N+1))$, or $4/(3(N+1))$.

5 [00112] It is possible to reduce the above ratio by increasing a number of storage devices used to store user data. For example, rather than utilize storage device Q1, each of the partitions may utilize a 3N+2 RAID array. In such a case, the ratio of a number of storage devices storing RAID parity information to a total number of storage devices is $2/(3N+2)$. However, during a reconstruct read operation, $(3N + 1)$ storage devices receive a reconstruct read request for a 10 single device failure. In contrast, for the hybrid RAID layout 1200, only N storage devices receive a reconstruct read request for a single device failure.

[00113] It is noted each of the three partitions may utilize a different RAID data layout architecture. A selection of a given RAID data layout architecture may be based on a given ratio number of storage devices storing RAID parity information to a total number of storage devices.

15 In addition, the selection may be based on a given number of storage devices, which may receive a reconstruct read request during reconstruction. For example, the RAID arrays 1210, 1220 and 1230 may include geometries such as L+a, M+b and N+c, respectively.

[00114] In addition, one or more storage devices, such as storage device Q1, may be chosen based on the above or other conditions to provide an additional amount of supported redundancy 20 for one or more of the RAID arrays within the partitions. In an example with three partitions comprising the above RAID arrays and a number Q of storage devices providing extra protection for each of the RAID arrays, a ratio of a number of storage devices storing RAID parity information to a total number of storage devices is $(a+b+c+Q)/(L+a+M+b+N+c+Q)$. For a single device failure, a number of storage devices to receive a reconstruct read request is L, M and N,

25 respectively, for partitions 1 to 3 in the above example. It is noted that the above discussion generally describes 3 distinct partitions in FIG. 12. In such an embodiment, this type of “hard” partitioning where a given layout is limited to a particular group of devices may guarantee that reconstruct reads in one partition will not collide with those in another partition. However, in other embodiments the partitions may not be hard as described above. Rather, given a pool of 30 devices, layouts may be selected from any of the devices. For example, treating the devices as on big pool it is possible to configure layouts such as $(L+1, M+1, N+1) + 1$. Consequently, there is a chance that geometries overlap and reconstruct reads could collide. If L, M, and N are small

relative to the size of the pool then the percentage of reconstruct reads relative to normal reads may be kept low.

[00115] Referring now to FIG. 13, one embodiment of a method 1300 for selecting alternate RAID geometries in a data storage subsystem is shown. The components embodied in network architecture 100 and data storage arrays 120a-120b described above may generally operate in accordance with method 1300. The steps in this embodiment are shown in sequential order. However, some steps may occur in a different order than shown, some steps may be performed concurrently, some steps may be combined with other steps, and some steps may be absent in another embodiment.

10 [00116] In block 1302, a RAID engine 178 or other logic within a storage controller 174 determines to use a given number of devices to store user data in a RAID array within each partition of a storage subsystem. A RUSH or other algorithm may then be used to select which devices are to be used. In one embodiment, each partition utilizes a same number of storage devices. In other embodiments, each partition may utilize a different, unique number of storage 15 devices to store user data. In block 1304, the storage controller 174 may determine to support a number of storage devices to store corresponding Inter-Device Error Recovery (parity) data within each partition of the subsystem. Again, each partition may utilize a same number or a different, unique number of storage devices for storing RAID parity information.

20 [00117] In block 1306, the storage controller may determine to support a number Q of storage devices to store extra Inter-Device Error Recovery (parity) data across the partitions of the subsystem. In block 1308, both user data and corresponding RAID parity data may be written in selected storage devices. Referring again to FIG. 12, when a given RAID array is written, such as RAID array 1210 in partition 1, one or more bits of parity information may be generated and stored in storage device Q1 in partition 3.

25 [00118] If the storage controller 174 detects a condition for performing read reconstruction in a given partition (conditional block 1310), and if the given partition has a sufficient number of storage devices holding RAID parity information to handle a number of unavailable storage devices (conditional block 1312), then in block 1314, the reconstruct read operation(s) is performed with one or more corresponding storage devices within the given partition. The 30 condition may include a storage device within a given RAID array is unavailable due to a device failure or the device operates below a given performance level. The given RAID array is able to handle a maximum number of unavailable storage devices with the number of storage devices storing RAID parity information within the given partition. For example, if RAID array 1210 in

partition 1 in the above example is an L+a RAID array, then RAID array 1210 is able to perform read reconstruction utilizing only storage devices within partition 1 when k storage devices are unavailable, where $1 \leq k \leq a$.

[00119] If the given partition does not have a sufficient number of storage devices holding

5 RAID parity information to handle a number of unavailable storage devices (conditional block 1312), and if there is a sufficient number of Q storage devices to handle the number of unavailable storage devices (conditional block 1316), then in block 1318, the reconstruct read operation(s) is performed with one or more corresponding Q storage devices. One or more storage devices in other partitions, which are storing user data, may be accessed during the read
10 reconstruction. A selection of these storage devices may be based on a manner of a derivation of the parity information stored in the one or more Q storage devices. For example, referring again to FIG. 12, storage device D2 in partition 2 may be accessed during the read reconstruction, since this storage device may have been used to generate corresponding RAID parity information stored in storage device Q1. If there are not a sufficient number of Q storage devices to handle
15 the number of unavailable storage devices (conditional block 1316), then in block 1320, the corresponding user data may be read from another source or be considered lost.

[00120] It is noted that the above-described embodiments may comprise software. In such an

embodiment, the program instructions that implement the methods and/or mechanisms may be conveyed or stored on a computer readable medium. Numerous types of media which are
20 configured to store program instructions are available and include hard disks, floppy disks, CD-ROM, DVD, flash memory, Programmable ROMs (PROM), random access memory (RAM), and various other forms of volatile or non-volatile storage.

[00121] In various embodiments, one or more portions of the methods and mechanisms described herein may form part of a cloud-computing environment. In such embodiments,

25 resources may be provided over the Internet as services according to one or more various models. Such models may include Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). In IaaS, computer infrastructure is delivered as a service. In such a case, the computing equipment is generally owned and operated by the service provider. In the PaaS model, software tools and underlying equipment used by developers to develop software
30 solutions may be provided as a service and hosted by the service provider. SaaS typically includes a service provider licensing software as a service on demand. The service provider may host the software, or may deploy the software to a customer for a given period of time. Numerous combinations of the above models are possible and are contemplated.

[00122] Although the embodiments above have been described in considerable detail, numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

WHAT IS CLAIMED IS

1. A computer system comprising:
 - 5 a client computer configured to send read and write requests over a network;
 - 10 a data storage subsystem coupled to the network configured to receive the read and write requests, wherein the subsystem comprises a plurality of data storage locations on a plurality of storage devices;
 - 15 wherein the data storage subsystem further comprises a storage controller configured to: configure a first subset of the storage devices for use in a first RAID layout, the first RAID layout including a first set of redundant data; configure a second subset of the storage devices for use in a second RAID layout, the second RAID layout including a second set of redundant data; configure an additional device not included in either the first subset or the second subset to store redundant data for both the first RAID layout and the second RAID layout;
 - 20 wherein in response to receiving a given read request directed to a particular storage device of either the first RAID layout or the second RAID layout, the controller is configured to initiate a reconstruct read corresponding to the given read request in response to determining the particular storage device is exhibiting a non-error related relatively slow read response.
- 25 2. The computer system as recited in claim 1, wherein each device in the first subset and second subset is configured to store intra-device redundancy data, and wherein both the first set of redundant data and the second set of redundant data comprise inter-device redundancy data.
- 30 3. The computer system as recited in claim 2, wherein the plurality of storage devices are solid state storage devices.
4. The computer system as recited in claim 1, wherein the first RAID layout is an L+x layout, and the second RAID layout is an M+y layout, wherein L, x, M, and, y are integers, and wherein either or both (1) L is not equal to M, and (2) x is not equal to y.

5. The computer system as recited in claim 4, wherein the storage controller is configured to change one or more of L, x, M, and y in response to detecting a characteristic of the data storage subsystem, the characteristic being one or more of an error rate, an access rate, a device age, an occupancy level, an age of stored data, and an expected lifetime of stored data.
6. The computer system as recited in claim 1, wherein both the first RAID layout and the second RAID layout are selected from a single device group.
7. The computer system as recited in claim 1, wherein the first RAID layout is selected from a first device group, and the second RAID layout is selected from a second device group, wherein the first device group and the second device group do not share a device.
8. The computer system as recited in claim 1, wherein the first RAID layout and the second RAID layout include at least one device that has a larger storage capacity than other devices included in the first subset and the second subset.
9. The computer system as recited in claim 1, wherein the additional device is selected from the first subset, the second subset, or otherwise.
- 20 10. A method for use in a computing system, the method comprising:
 - receiving the read and write requests in a data storage subsystem, wherein the subsystem comprises a plurality of data storage locations on a plurality of storage devices;
 - configuring a first subset of the storage devices for use in a first RAID layout, the first RAID layout including a first set of redundant data;
 - 25 configuring a second subset of the storage devices for use in a second RAID layout, the second RAID layout including a second set of redundant data; and
 - configuring an additional device not included in either the first subset or the second subset to store redundant data for both the first RAID layout and the second RAID layout;
- 30 initiating a reconstruct read corresponding to a given read request directed to a particular storage device of the plurality of storage devices, in response to determining the particular storage device is exhibiting a non-error related relatively slow read response.

11. The method as recited in claim 10, wherein each device in the first subset and second subset is configured to store intra-device redundancy data, and wherein both the first set of redundant data and the second set of redundant data comprise inter-device redundancy data.

5 12. The method as recited in claim 11, wherein the plurality of storage devices are solid state storage devices.

10 13. The method as recited in claim 10, wherein the first RAID layout is an $L+x$ layout, and the second RAID layout is an $M+y$ layout, wherein L , x , M , and, y are integers, and wherein either or both (1) L is not equal to M , and (2) x is not equal to y .

15 14. The method as recited in claim 13, further comprising changing one or more of L , x , M , and y in response to detecting a characteristic of the data storage subsystem, the characteristic being one or more of an error rate, an access rate, a device age, an occupancy level, an age of stored data, and an expected lifetime of stored data

15. The method as recited in claim 10, wherein both the first RAID layout and the second RAID layout are selected from a single device group.

20 16. The method as recited in claim 10, wherein the first RAID layout is selected from a first device group, and the second RAID layout is selected from a second device group, wherein the first device group and the second device group do not share a device.

25 17. A computer readable storage medium storing program instructions, wherein the program instructions are executable to:

receive the read and write requests in a data storage subsystem, wherein the subsystem comprises a plurality of data storage locations on a plurality of storage devices; configure a first subset of the storage devices for use in a first RAID layout, the first RAID layout including a first set of redundant data;

30 configure a second subset of the storage devices for use in a second RAID layout, the second RAID layout including a second set of redundant data; and

configure an additional device not included in either the first subset or the second subset to store redundant data for both the first RAID layout and the second RAID layout;

5 initiate a reconstruct read corresponding to a given read request directed to a particular storage device of the plurality of storage devices, in response to determining the particular storage device is exhibiting a non-error related relatively slow read response.

18. The computer readable storage medium as recited in claim 17, wherein each device in the

10 first subset and second subset is configured to store intra-device redundancy data, and wherein both the first set of redundant data and the second set of redundant data comprise inter-device redundancy data.

19. The computer readable storage medium as recited in claim 18, wherein the plurality of

15 storage devices are solid state storage devices.

20. The computer readable storage medium as recited in claim 17, wherein the first RAID layout

is an $L+x$ layout, and the second RAID layout is an $M+y$ layout, wherein L , x , M , and, y are integers, and wherein either or both (1) L is not equal to M , and (2) x is not equal to y .

20




FIG. 1




FIG. 2

3/15




FIG. 3





FIG. 4



6/15

State Table 522

Device Age <u>602</u>
Error Rate <u>604</u>
Total Errors <u>606</u>
Number of Recoverable Errors <u>608</u>
Number of Unrecoverable Errors <u>610</u>
Access Rate <u>612</u>
Data Age <u>614</u>
Allocation State of a First Allocation Space <u>616a</u>
⋮
Allocation State of an nth Allocation Space <u>616n</u>

FIG. 6




FIG. 7




FIG. 8




FIG. 9

10/15




FIG. 10

11/15

Sample Device Layout




FIG. 11A




FIG. 11B

13/15





FIG. 11C




FIG. 12

15/15




FIG. 13

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2011/052276

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F11/10 G06F3/06
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

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

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

EPO-Internal, WPI Data, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2009/210742 A1 (ADARSHAPPANAVAR MALLIKARJUN [IN] ET AL) 20 August 2009 (2009-08-20) the whole document ----- WO 00/52802 A1 (SEACHANGE SYSTEMS INC [US]) 8 September 2000 (2000-09-08) abstract; claim 1 page 2, line 3 - line 13 ----- -/-	1-20
Y		1-20

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
2 December 2011	14/12/2011
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Salsa, Francesco

INTERNATIONAL SEARCH REPORT

International application N

PCT/US2011/052276

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	AMER A ET AL: "Progressive Parity-Based Hardening of Data Stores",	5,14
	PERFORMANCE, COMPUTING AND COMMUNICATIONS CONFERENCE, 2008. IPCCC 2008. IEEE INTERNATIONAL, IEEE, PISCATAWAY, NJ, USA, 7 December 2008 (2008-12-07), pages 34-42, XP031404354,	
	ISBN: 978-1-4244-3368-1 the whole document	1-4, 6-13, 15-20
Y	-----	8
	AVANI WILDANI ET AL: "Protecting against rare event failures in archival systems",	
	MODELING, ANALYSIS&SIMULATION OF COMPUTER AND TELECOMMUNICATION SYSTEMS, 2009.	
	MASCOTS '09. IEEE INTERNATIONAL SYMPOSIUM ON, IEEE, PISCATAWAY, NJ, USA, 21 September 2009 (2009-09-21), pages	
	1-11, XP031591252, ISBN: 978-1-4244-4927-9	
	the whole document	1-7,9-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2011/052276

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
US 2009210742	A1	20-08-2009	NONE	

WO 0052802	A1	08-09-2000	AT 304237 T AU 3713000 A CA 2365694 A1 CN 1391716 A DE 60022488 D1 EP 1166418 A1 JP 2003503766 A US 6321345 B1 US 2002032882 A1 WO 0052802 A1	15-09-2005 21-09-2000 08-09-2000 15-01-2003 13-10-2005 02-01-2002 28-01-2003 20-11-2001 14-03-2002 08-09-2000
