



- (51) International Patent Classification:
G06F 21/62 (2013.01) *G06F 12/14* (2006.01)
- (21) International Application Number:
PCT/US2016/041920
- (22) International Filing Date:
12 July 2016 (12.07.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/798,371 13 July 2015 (13.07.2015) US
- (71) Applicant: **PURE STORAGE, INC.** [US/US]; Suite 400, 650 Castro St., Mountain View, California 94041 (US).
- (72) Inventors: **HAYES, John**; Suite 400, 650 Castro St., Mountain View, California 94041 (US). **LEE, Robert**; Suite 400, 650 Castro St., Mountain View, California 94041 (US). **OSTROVSKY, Igor**; Suite 400, 650 Castro St., Mountain View, California 94041 (US). **VAJGEL, Peter**; Suite 400, 650 Castro St., Mountain View, California 94041 (US).
- (74) Agent: **GENCARELLA, Michael L.**; Womble Carlyle Sandridge & Rice LLP, P.o. Box 7037, Atlanta, Georgia 30357-0037 (US).
- (81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM,

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

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

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

Published:

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

(54) Title: METHOD AND SYSTEM FOR ACCESSING A FILE





FIG. 4B

(57) Abstract: A method for accessing a file in a storage system is provided. The method includes determining, for each file chunk of the file, an authority among differing storage nodes of the storage system and receiving from the authority having ownership of the file chunk, location information for the file chunk. The method includes accessing file chunks of the file as directed by each of the determined authorities.



METHOD AND SYSTEM FOR ACCESSING A FILE

BACKGROUND

[0001] Solid-state memory, such as flash, is currently in use in solid-state drives (SSD) to augment or replace conventional hard disk drives (HDD), writable CD (compact disk) or writable DVD (digital versatile disk) drives, collectively known as spinning media, and tape drives, for storage of large amounts of data. Flash and other solid-state memories have characteristics that differ from spinning media. Yet, many solid-state drives are designed to conform to hard disk drive standards for compatibility reasons, which makes it difficult to provide enhanced features or take advantage of unique aspects of flash and other solid-state memory. Large files, such as in the gigabyte range and upwards, place a burden on storage system load, and may tie up resources of a storage node thus producing a bottleneck in storage cluster or storage system performance. It is within this context that the embodiments described herein arise.

SUMMARY

[0002] In some embodiments, a method for accessing a file in a storage system is provided. The method includes determining, for each file chunk of the file, an authority among differing storage nodes of the storage system and receiving from the authority having ownership of the file chunk, location information for the file chunk. The method includes accessing file chunks of the file as directed by each of the determined authorities.

[0003] In some embodiments, a method for accessing file chunks of a file in a storage system is provided. The method includes determining one of a plurality of owners in a plurality of storage nodes as having ownership of a file chunk of the file, the determining performed at differing storage nodes for each of a plurality of file chunks of the file. The method includes accessing the plurality of file chunks of the file as directed by determined owners.

[0004] In some embodiments, a storage system having efficient distribution of file chunks is provided. The system includes a plurality of storage nodes coupled as a storage cluster. At least a portion of the plurality of storage nodes is configured to have at least one authority, each authority having ownership of a subset of file chunks, each file chunk including a portion of a file. Each of the plurality of storage nodes is configured to determine which authority in which of the plurality of storage nodes has information regarding location in the storage cluster of each file chunk.

[0005] Other aspects and advantages of the embodiments will become apparent from the following detailed description taken in conjunction with the accompanying drawings which illustrate, by way of example, the principles of the described embodiments.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] The described embodiments and the advantages thereof may best be understood by reference to the following description taken in conjunction with the accompanying drawings. These drawings in no way limit any changes in form and detail that may be made to the described embodiments by one skilled in the art without departing from the spirit and scope of the described embodiments.

[0007] Fig. 1 is a perspective view of a storage cluster with multiple storage nodes and internal storage coupled to each storage node to provide network attached storage, in accordance with some embodiments.

[0008] Fig. 2 is a block diagram showing an interconnect switch coupling multiple storage nodes in accordance with some embodiments.

[0009] Fig. 3 is a multiple level block diagram, showing contents of a storage node and contents of one of the non-volatile solid state storage units in accordance with some embodiments.

[0010] Fig. 4A depicts a file, suitable for storage and access in embodiments of the storage cluster, with file chunks each having a file offset relative to the file.

[0011] Fig. 4B shows an example file, and ownership of various file chunks of the file, in accordance with an embodiment of the present description.

[0012] Fig. 5 is a block diagram of a storage cluster with components for determining ownership of file chunks.

[0013] Fig. 6 is a system action diagram showing use of a hash calculator and a lookup table to point to an authority as an owner of a file chunk.

[0014] Fig. 7A is a flow diagram of a method for distributing ownership of file chunks in a storage system, and writing file chunks of files, which can be practiced by embodiments of the storage cluster, storage nodes and storage units of Figs. 1-6.

[0015] Fig. 7B is a flow diagram of a method for distributing ownership of file chunks in a storage system, and reading file chunks of files, which can be practiced by embodiments of the storage cluster, storage nodes and storage units of Figs. 1-6.

[0016] Fig. 8 is an illustration showing an exemplary computing device which may implement the embodiments described herein.

DETAILED DESCRIPTION

[0017] A storage cluster, with storage nodes and storage units that have storage memory, is herein described as having a mechanism that distributes, to multiple storage nodes, information about physical location of file chunks of files, rather than concentrating all such information in a single storage node, or all such information about file chunks of each file in a storage node that owns that file. Distributing the information as metadata to multiple storage nodes leads to efficient distribution of large files. In various embodiments, storage nodes use a hash calculation, of an inode number of a file and an offset of the file chunk relative to the file, to point via a lookup table to information about the physical location of the file chunk, for each file chunk of the file. Since neither the performance of the hash calculation nor the lookup table are limited to belonging to only one storage node, the storage system avoids a bottleneck situation. Embodiments of storage clusters, storage nodes and storage units are described with reference to Figs. 1-3. Processes and mechanisms for obtaining information about the physical location of a file chunk, in embodiments of storage nodes and embodiments of a storage system, are described with reference to Figs. 4A-7B.

[0018] The embodiments below describe a storage cluster that stores user data, such as user data originating from one or more user or client systems or other sources external to the storage cluster. The storage cluster distributes user data across storage nodes housed within a chassis, using erasure coding and redundant copies of metadata. Erasure coding refers to a method of data protection or reconstruction in which data is stored across a set of different locations, such as disks, storage nodes or geographic locations. Flash memory is one type of solid-state memory that may be integrated with the embodiments, although the embodiments may be extended to other types of solid-state memory or other storage medium, including non-solid state memory. Control of storage locations and workloads are distributed across the storage locations in a clustered peer-to-peer system. Tasks such as mediating communications between the various storage nodes, detecting when a storage node has become unavailable, and balancing I/Os (inputs and outputs) across the various storage nodes, are all handled on a distributed basis. Data is laid out or distributed across multiple storage nodes in data fragments or stripes that support data recovery in some embodiments. Ownership of data can be reassigned within a cluster, independent of input and output patterns. This architecture described in more detail below allows a storage node in the cluster to fail, with the system remaining operational, since the data can be reconstructed from other storage nodes and thus remain available for input and output operations. In various embodiments, a storage node may be referred to as a cluster node, a blade, or a server.

[0019] The storage cluster is contained within a chassis, i.e., an enclosure housing one or more storage nodes. A mechanism to provide power to each storage node, such as a power distribution bus, and a communication mechanism, such as a communication bus that enables communication between the storage nodes are included within the chassis. The storage cluster can run as an independent system in one location according to some embodiments. In one embodiment, a chassis contains at least two instances of both the power distribution and the communication bus which may be enabled or disabled independently. The internal communication bus may be an Ethernet bus, however, other technologies such as Peripheral Component Interconnect (PCI) Express, InfiniBand, and others, are equally suitable. The chassis provides a port for an external communication bus for enabling communication between multiple chassis, directly or through a switch, and with client systems. The external communication may use a technology such as Ethernet, InfiniBand, Fibre Channel, etc. In some embodiments, the external communication bus uses different communication bus technologies for inter-chassis and client communication. If a switch is deployed within or between chassis, the switch may act as a translation between multiple protocols or technologies. When multiple chassis are connected to define a storage cluster, the storage cluster may be accessed by a client using either proprietary interfaces or standard interfaces such as network file system (NFS), common internet file system (CIFS), small computer system interface (SCSI) or hypertext transfer protocol (HTTP). Translation from the client protocol may occur at the switch, chassis external communication bus or within each storage node.

[0020] Each storage node may be one or more storage servers and each storage server is connected to one or more non-volatile solid state memory units, which may be referred to as storage units. One embodiment includes a single storage server in each storage node and between one to eight non-volatile solid state memory units, however this one example is not meant to be limiting. The storage server may include a processor, dynamic random access memory (DRAM) and interfaces for the internal communication bus and power distribution for each of the power buses. Inside the storage node, the interfaces and storage unit share a communication bus, e.g., PCI Express, in some embodiments. The non-volatile solid state memory units may directly access the internal communication bus interface through a storage node communication bus, or request the storage node to access the bus interface. The non-volatile solid state memory unit contains an embedded central processing unit (CPU), solid state storage controller, and a quantity of solid state mass storage, e.g., between 2-32 terabytes (TB) in some embodiments. An embedded volatile storage medium, such as

DRAM, and an energy reserve apparatus are included in the non-volatile solid state memory unit. In some embodiments, the energy reserve apparatus is a capacitor, super-capacitor, or battery that enables transferring a subset of DRAM contents to a stable storage medium in the case of power loss. In some embodiments, the non-volatile solid state memory unit is constructed with a storage class memory, such as phase change or magnetoresistive random access memory (MRAM) that substitutes for DRAM and enables a reduced power hold-up apparatus.

[0021] One of many features of the storage nodes and non-volatile solid state storage is the ability to proactively rebuild data in a storage cluster. The storage nodes and non-volatile solid state storage can determine when a storage node or non-volatile solid state storage in the storage cluster is unreachable, independent of whether there is an attempt to read data involving that storage node or non-volatile solid state storage. The storage nodes and non-volatile solid state storage then cooperate to recover and rebuild the data in at least partially new locations. This constitutes a proactive rebuild, in that the system rebuilds data without waiting until the data is needed for a read access initiated from a client system employing the storage cluster. These and further details of the storage memory and operation thereof are discussed below.

[0022] Fig. 1 is a perspective view of a storage cluster 160, with multiple storage nodes 150 and internal solid-state memory coupled to each storage node to provide network attached storage or storage area network, in accordance with some embodiments. A network attached storage, storage area network, or a storage cluster, or other storage memory, could include one or more storage clusters 160, each having one or more storage nodes 150, in a flexible and reconfigurable arrangement of both the physical components and the amount of storage memory provided thereby. The storage cluster 160 is designed to fit in a rack, and one or more racks can be set up and populated as desired for the storage memory. The storage cluster 160 has a chassis 138 having multiple slots 142. It should be appreciated that chassis 138 may be referred to as a housing, enclosure, or rack unit. In one embodiment, the chassis 138 has fourteen slots 142, although other numbers of slots are readily devised. For example, some embodiments have four slots, eight slots, sixteen slots, thirty-two slots, or other suitable number of slots. Each slot 142 can accommodate one storage node 150 in some embodiments. Chassis 138 includes flaps 148 that can be utilized to mount the chassis 138 on a rack. Fans 144 provide air circulation for cooling of the storage nodes 150 and components thereof, although other cooling components could be used, or an embodiment could be devised without cooling components. A switch fabric 146 couples storage nodes

150 within chassis 138 together and to a network for communication to the memory. In an embodiment depicted in Fig. 1, the slots 142 to the left of the switch fabric 146 and fans 144 are shown occupied by storage nodes 150, while the slots 142 to the right of the switch fabric 146 and fans 144 are empty and available for insertion of storage node 150 for illustrative purposes. This configuration is one example, and one or more storage nodes 150 could occupy the slots 142 in various further arrangements. The storage node arrangements need not be sequential or adjacent in some embodiments. Storage nodes 150 are hot pluggable, meaning that a storage node 150 can be inserted into a slot 142 in the chassis 138, or removed from a slot 142, without stopping or powering down the system. Upon insertion or removal of storage node 150 from slot 142, the system automatically reconfigures in order to recognize and adapt to the change. Reconfiguration, in some embodiments, includes restoring redundancy and/or rebalancing data or load.

[0023] Each storage node 150 can have multiple components. In the embodiment shown here, the storage node 150 includes a printed circuit board 158 populated by a CPU 156, i.e., processor, a memory 154 coupled to the CPU 156, and a non-volatile solid state storage 152 coupled to the CPU 156, although other mountings and/or components could be used in further embodiments. The memory 154 has instructions which are executed by the CPU 156 and/or data operated on by the CPU 156. As further explained below, the non-volatile solid state storage 152 includes flash or, in further embodiments, other types of solid-state memory.

[0024] Referring to Fig. 1, storage cluster 160 is scalable, meaning that storage capacity with non-uniform storage sizes is readily added, as described above. One or more storage nodes 150 can be plugged into or removed from each chassis and the storage cluster self-configures in some embodiments. Plug-in storage nodes 150, whether installed in a chassis as delivered or later added, can have different sizes. For example, in one embodiment a storage node 150 can have any multiple of 4 TB, e.g., 8 TB, 12 TB, 16 TB, 32 TB, etc. In further embodiments, a storage node 150 could have any multiple of other storage amounts or capacities. Storage capacity of each storage node 150 is broadcast, and influences decisions of how to stripe the data. For maximum storage efficiency, an embodiment can self-configure as wide as possible in the stripe, subject to a predetermined requirement of continued operation with loss of up to one, or up to two, non-volatile solid state storage units 152 or storage nodes 150 within the chassis.

[0025] Fig. 2 is a block diagram showing a communications interconnect 170 and power distribution bus 172 coupling multiple storage nodes 150. Referring back to Fig. 1, the

communications interconnect 170 can be included in or implemented with the switch fabric 146 in some embodiments. Where multiple storage clusters 160 occupy a rack, the communications interconnect 170 can be included in or implemented with a top of rack switch, in some embodiments. As illustrated in Fig. 2, storage cluster 160 is enclosed within a single chassis 138. External port 176 is coupled to storage nodes 150 through communications interconnect 170, while external port 174 is coupled directly to a storage node. External power port 178 is coupled to power distribution bus 172. Storage nodes 150 may include varying amounts and differing capacities of non-volatile solid state storage 152 as described with reference to Fig. 1. In addition, one or more storage nodes 150 may be a compute only storage node as illustrated in Fig. 2. Authorities 168 are implemented on the non-volatile solid state storages 152, for example as lists or other data structures stored in memory. In some embodiments the authorities are stored within the non-volatile solid state storage 152 and supported by software executing on a controller or other processor of the non-volatile solid state storage 152. In a further embodiment, authorities 168 are implemented on the storage nodes 150, for example as lists or other data structures stored in the memory 154 and supported by software executing on the CPU 156 of the storage node 150. Authorities 168 control how and where data is stored in the non-volatile solid state storages 152 in some embodiments. This control assists in determining which type of erasure coding scheme is applied to the data, and which storage nodes 150 have which portions of the data. Each authority 168 may be assigned to a non-volatile solid state storage 152. Each authority may control a range of inode numbers, segment numbers, or other data identifiers which are assigned to data by a file system, by the storage nodes 150, or by the non-volatile solid state storage 152, in various embodiments.

[0026] Every piece of data, and every piece of metadata, has redundancy in the system in some embodiments. In addition, every piece of data and every piece of metadata has an owner, which may be referred to as an authority. If that authority is unreachable, for example through failure of a storage node, there is a plan of succession for how to find that data or that metadata. In various embodiments, there are redundant copies of authorities 168. Authorities 168 have a relationship to storage nodes 150 and non-volatile solid state storage 152 in some embodiments. Each authority 168, covering a range of data segment numbers or other identifiers of the data, may be assigned to a specific non-volatile solid state storage 152. In some embodiments the authorities 168 for all of such ranges are distributed over the non-volatile solid state storages 152 of a storage cluster. Each storage node 150 has a network port that provides access to the non-volatile solid state storage(s) 152 of that storage node

150. Data can be stored in a segment, which is associated with a segment number and that segment number is an indirection for a configuration of a RAID (redundant array of independent disks) stripe in some embodiments. The assignment and use of the authorities 168 thus establishes an indirection to data. Indirection may be referred to as the ability to reference data indirectly, in this case via an authority 168, in accordance with some embodiments. A segment identifies a set of non-volatile solid state storage 152 and a local identifier into the set of non-volatile solid state storage 152 that may contain data. In some embodiments, the local identifier is an offset into the device and may be reused sequentially by multiple segments. In other embodiments the local identifier is unique for a specific segment and never reused. The offsets in the non-volatile solid state storage 152 are applied to locating data for writing to or reading from the non-volatile solid state storage 152 (in the form of a RAID stripe). Data is striped across multiple units of non-volatile solid state storage 152, which may include or be different from the non-volatile solid state storage 152 having the authority 168 for a particular data segment.

[0027] If there is a change in where a particular segment of data is located, e.g., during a data move or a data reconstruction, the authority 168 for that data segment should be consulted, at that non-volatile solid state storage 152 or storage node 150 having that authority 168. In order to locate a particular piece of data, embodiments calculate a hash value for a data segment or apply an inode number or a data segment number. The output of this operation points to a non-volatile solid state storage 152 having the authority 168 for that particular piece of data. In some embodiments there are two stages to this operation. The first stage maps an entity identifier (ID), e.g., a segment number, inode number, or directory number to an authority identifier. This mapping may include a calculation such as a hash or a bit mask. The second stage is mapping the authority identifier to a particular non-volatile solid state storage 152, which may be done through an explicit mapping. The operation is repeatable, so that when the calculation is performed, the result of the calculation repeatably and reliably points to a particular non-volatile solid state storage 152 having that authority 168. The operation may include the set of reachable storage nodes as input. If the set of reachable non-volatile solid state storage units changes the optimal set changes. In some embodiments, the persisted value is the current assignment (which is always true) and the calculated value is the target assignment the cluster will attempt to reconfigure towards. This calculation may be used to determine the optimal non-volatile solid state storage 152 for an authority in the presence of a set of non-volatile solid state storage 152 that are reachable and constitute the same cluster. The calculation also determines an ordered set of peer non-

volatile solid state storage 152 that will also record the authority to non-volatile solid state storage mapping so that the authority may be determined even if the assigned non-volatile solid state storage is unreachable. A duplicate or substitute authority 168 may be consulted if a specific authority 168 is unavailable in some embodiments.

[0028] With reference to Figs. 1 and 2, two of the many tasks of the CPU 156 on a storage node 150 are to break up write data, and reassemble read data. When the system has determined that data is to be written, the authority 168 for that data is located as above. When the segment ID for data is already determined the request to write is forwarded to the non-volatile solid state storage 152 currently determined to be the host of the authority 168 determined from the segment. The host CPU 156 of the storage node 150, on which the non-volatile solid state storage 152 and corresponding authority 168 reside, then breaks up or shards the data and transmits the data out to various non-volatile solid state storage 152. The transmitted data is written as a data stripe in accordance with an erasure coding scheme. In some embodiments, data is requested to be pulled, and in other embodiments, data is pushed. In reverse, when data is read, the authority 168 for the segment ID containing the data is located as described above. The host CPU 156 of the storage node 150 on which the non-volatile solid state storage 152 and corresponding authority 168 reside requests the data from the non-volatile solid state storage and corresponding storage nodes pointed to by the authority. In some embodiments the data is read from flash storage as a data stripe. The host CPU 156 of storage node 150 then reassembles the read data, correcting any errors (if present) according to the appropriate erasure coding scheme, and forwards the reassembled data to the network. In further embodiments, some or all of these tasks can be handled in the non-volatile solid state storage 152. In some embodiments, the segment host requests the data be sent to storage node 150 by requesting pages from storage and then sending the data to the storage node making the original request.

[0029] In some systems, for example in UNIX-style file systems, data is handled with an index node or inode, which specifies a data structure that represents an object in a file system. The object could be a file or a directory, for example. Metadata may accompany the object, as attributes such as permission data and a creation timestamp, among other attributes. A segment number could be assigned to all or a portion of such an object in a file system. In other systems, data segments are handled with a segment number assigned elsewhere. For purposes of discussion, the unit of distribution is an entity, and an entity can be a file, a directory or a segment. That is, entities are units of data or metadata stored by a storage system. Entities are grouped into sets called authorities. Each authority has an authority

owner, which is a storage node that has the exclusive right to update the entities in the authority. In other words, a storage node contains the authority, and that the authority, in turn, contains entities.

[0030] A segment is a logical container of data in accordance with some embodiments. A segment is an address space between medium address space and physical flash locations, i.e., the data segment number, are in this address space. Segments may also contain meta-data, which enable data redundancy to be restored (rewritten to different flash locations or devices) without the involvement of higher level software. In one embodiment, an internal format of a segment contains client data and medium mappings to determine the position of that data. Each data segment is protected, e.g., from memory and other failures, by breaking the segment into a number of data and parity shards, where applicable. The data and parity shards are distributed, i.e., striped, across non-volatile solid state storage 152 coupled to the host CPUs 156 (See Fig. 5) in accordance with an erasure coding scheme. Usage of the term segments refers to the container and its place in the address space of segments in some embodiments. Usage of the term stripe refers to the same set of shards as a segment and includes how the shards are distributed along with redundancy or parity information in accordance with some embodiments.

[0031] A series of address-space transformations takes place across an entire storage system. At the top are the directory entries (file names) which link to an inode. Inodes point into medium address space, where data is logically stored. Medium addresses may be mapped through a series of indirect mediums to spread the load of large files, or implement data services like deduplication or snapshots. Medium addresses may be mapped through a series of indirect mediums to spread the load of large files, or implement data services like deduplication or snapshots. Segment addresses are then translated into physical flash locations. Physical flash locations have an address range bounded by the amount of flash in the system in accordance with some embodiments. Medium addresses and segment addresses are logical containers, and in some embodiments use a 128 bit or larger identifier so as to be practically infinite, with a likelihood of reuse calculated as longer than the expected life of the system. Addresses from logical containers are allocated in a hierarchical fashion in some embodiments. Initially, each non-volatile solid state storage 152 may be assigned a range of address space. Within this assigned range, the non-volatile solid state storage 152 is able to allocate addresses without synchronization with other non-volatile solid state storage 152.

[0032] Data and metadata is stored by a set of underlying storage layouts that are optimized for varying workload patterns and storage devices. These layouts incorporate multiple redundancy schemes, compression formats and index algorithms. Some of these layouts store information about authorities and authority masters, while others store file metadata and file data. The redundancy schemes include error correction codes that tolerate corrupted bits within a single storage device (such as a NAND flash chip), erasure codes that tolerate the failure of multiple storage nodes, and replication schemes that tolerate data center or regional failures. In some embodiments, low density parity check (LDPC) code is used within a single storage unit. Reed-Solomon encoding is used within a storage cluster, and mirroring is used within a storage grid in some embodiments. Metadata may be stored using an ordered log structured index (such as a Log Structured Merge Tree), and large data may not be stored in a log structured layout.

[0033] In order to maintain consistency across multiple copies of an entity, the storage nodes agree implicitly on two things through calculations: (1) the authority that contains the entity, and (2) the storage node that contains the authority. The assignment of entities to authorities can be done by pseudorandomly assigning entities to authorities, by splitting entities into ranges based upon an externally produced key, or by placing a single entity into each authority. Examples of pseudorandom schemes are linear hashing and the Replication Under Scalable Hashing (RUSH) family of hashes, including Controlled Replication Under Scalable Hashing (CRUSH). In some embodiments, pseudo-random assignment is utilized only for assigning authorities to nodes because the set of nodes can change. The set of authorities cannot change so any subjective function may be applied in these embodiments. Some placement schemes automatically place authorities on storage nodes, while other placement schemes rely on an explicit mapping of authorities to storage nodes. In some embodiments, a pseudorandom scheme is utilized to map from each authority to a set of candidate authority owners. A pseudorandom data distribution function related to CRUSH may assign authorities to storage nodes and create a list of where the authorities are assigned. Each storage node has a copy of the pseudorandom data distribution function, and can arrive at the same calculation for distributing, and later finding or locating an authority. Each of the pseudorandom schemes requires the reachable set of storage nodes as input in some embodiments in order to conclude the same target nodes. Once an entity has been placed in an authority, the entity may be stored on physical devices so that no expected failure will lead to unexpected data loss. In some embodiments, rebalancing algorithms attempt to store the copies of all entities within an authority in the same layout and on the same set of machines.

[0034] Examples of expected failures include device failures, stolen machines, datacenter fires, and regional disasters, such as nuclear or geological events. Different failures lead to different levels of acceptable data loss. In some embodiments, a stolen storage node impacts neither the security nor the reliability of the system, while depending on system configuration, a regional event could lead to no loss of data, a few seconds or minutes of lost updates, or even complete data loss.

[0035] In the embodiments, the placement of data for storage redundancy is independent of the placement of authorities for data consistency. In some embodiments, storage nodes that contain authorities do not contain any persistent storage. Instead, the storage nodes are connected to non-volatile solid state storage units that do not contain authorities. The communications interconnect between storage nodes and non-volatile solid state storage units consists of multiple communication technologies and has non-uniform performance and fault tolerance characteristics. In some embodiments, as mentioned above, non-volatile solid state storage units are connected to storage nodes via PCI express, storage nodes are connected together within a single chassis using Ethernet backplane, and chassis are connected together to form a storage cluster. Storage clusters are connected to clients using Ethernet or fiber channel in some embodiments. If multiple storage clusters are configured into a storage grid, the multiple storage clusters are connected using the Internet or other long-distance networking links, such as a “metro scale” link or private link that does not traverse the internet.

[0036] Authority owners have the exclusive right to modify entities, to migrate entities from one non-volatile solid state storage unit to another non-volatile solid state storage unit, and to add and remove copies of entities. This allows for maintaining the redundancy of the underlying data. When an authority owner fails, is going to be decommissioned, or is overloaded, the authority is transferred to a new storage node. Transient failures make it non-trivial to ensure that all non-faulty machines agree upon the new authority location. The ambiguity that arises due to transient failures can be achieved automatically by a consensus protocol such as Paxos, hot-warm failover schemes, via manual intervention by a remote system administrator, or by a local hardware administrator (such as by physically removing the failed machine from the cluster, or pressing a button on the failed machine). In some embodiments, a consensus protocol is used, and failover is automatic. If too many failures or replication events occur in too short a time period, the system goes into a self-preservation mode and halts replication and data movement activities until an administrator intervenes in accordance with some embodiments.

[0037] As authorities are transferred between storage nodes and authority owners update entities in their authorities, the system transfers messages between the storage nodes and non-volatile solid state storage units. With regard to persistent messages, messages that have different purposes are of different types. Depending on the type of the message, the system maintains different ordering and durability guarantees. As the persistent messages are being processed, the messages are temporarily stored in multiple durable and non-durable storage hardware technologies. In some embodiments, messages are stored in RAM, NVRAM and on NAND flash devices, and a variety of protocols are used in order to make efficient use of each storage medium. Latency-sensitive client requests may be persisted in replicated NVRAM, and then later NAND, while background rebalancing operations are persisted directly to NAND.

[0038] Persistent messages are persistently stored prior to being transmitted. This allows the system to continue to serve client requests despite failures and component replacement. Although many hardware components contain unique identifiers that are visible to system administrators, manufacturer, hardware supply chain and ongoing monitoring quality control infrastructure, applications running on top of the infrastructure address virtualize addresses. These virtualized addresses do not change over the lifetime of the storage system, regardless of component failures and replacements. This allows each component of the storage system to be replaced over time without reconfiguration or disruptions of client request processing.

[0039] In some embodiments, the virtualized addresses are stored with sufficient redundancy. A continuous monitoring system correlates hardware and software status and the hardware identifiers. This allows detection and prediction of failures due to faulty components and manufacturing details. The monitoring system also enables the proactive transfer of authorities and entities away from impacted devices before failure occurs by removing the component from the critical path in some embodiments.

[0040] Fig. 3 is a multiple level block diagram, showing contents of a storage node 150 and contents of a non-volatile solid state storage 152 of the storage node 150. Data is communicated to and from the storage node 150 by a network interface controller (NIC) 202 in some embodiments. Each storage node 150 has a CPU 156, and one or more non-volatile solid state storage 152, as discussed above. Moving down one level in Fig. 3, each non-volatile solid state storage 152 has a relatively fast non-volatile solid state memory, such as nonvolatile random access memory (NVRAM) 204, and flash memory 206. In some embodiments, NVRAM 204 may be a component that does not require program/erase cycles (DRAM, MRAM, PCM), and can be a memory that can support being written vastly more

often than the memory is read from. Moving down another level in Fig. 3, the NVRAM 204 is implemented in one embodiment as high speed volatile memory, such as dynamic random access memory (DRAM) 216, backed up by energy reserve 218. Energy reserve 218 provides sufficient electrical power to keep the DRAM 216 powered long enough for contents to be transferred to the flash memory 206 in the event of power failure. In some embodiments, energy reserve 218 is a capacitor, super-capacitor, battery, or other device, that supplies a suitable supply of energy sufficient to enable the transfer of the contents of DRAM 216 to a stable storage medium in the case of power loss. The flash memory 206 is implemented as multiple flash dies 222, which may be referred to as packages of flash dies 222 or an array of flash dies 222. It should be appreciated that the flash dies 222 could be packaged in any number of ways, with a single die per package, multiple dies per package (i.e. multichip packages), in hybrid packages, as bare dies on a printed circuit board or other substrate, as encapsulated dies, etc. In the embodiment shown, the non-volatile solid state storage 152 has a controller 212 or other processor, and an input output (I/O) port 210 coupled to the controller 212. I/O port 210 is coupled to the CPU 156 and/or the network interface controller 202 of the flash storage node 150. Flash input output (I/O) port 220 is coupled to the flash dies 222, and a direct memory access unit (DMA) 214 is coupled to the controller 212, the DRAM 216 and the flash dies 222. In the embodiment shown, the I/O port 210, controller 212, DMA unit 214 and flash I/O port 220 are implemented on a programmable logic device (PLD) 208, e.g., a field programmable gate array (FPGA). In this embodiment, each flash die 222 has pages, organized as sixteen kB (kilobyte) pages 224, and a register 226 through which data can be written to or read from the flash die 222. In further embodiments, other types of solid-state memory are used in place of, or in addition to flash memory illustrated within flash die 222.

[0041] Storage cluster 160, in various embodiments as disclosed herein, can be contrasted with storage arrays in general. The storage nodes 150 are part of a collection that creates the storage cluster 160. Each storage node 150 owns a slice of data and computing required to provide the data. Multiple storage nodes 150 cooperate to store and retrieve the data. Storage memory or storage devices, as used in storage arrays in general, are less involved with processing and manipulating the data. Storage memory or storage devices in a storage array receive commands to read, write, or erase data. The storage memory or storage devices in a storage array are not aware of a larger system in which they are embedded, or what the data means. Storage memory or storage devices in storage arrays can include various types of storage memory, such as RAM, solid state drives, hard disk drives, etc. The storage units

152 described herein have multiple interfaces active simultaneously and serving multiple purposes. In some embodiments, some of the functionality of a storage node 150 is shifted into a storage unit 152, transforming the storage unit 152 into a combination of storage unit 152 and storage node 150. Placing computing (relative to storage data) into the storage unit 152 places this computing closer to the data itself. The various system embodiments have a hierarchy of storage node layers with different capabilities. By contrast, in a storage array, a controller owns and knows everything about all of the data that the controller manages in a shelf or storage devices. In a storage cluster 160, as described herein, multiple controllers in multiple storage units 152 and/or storage nodes 150 cooperate in various ways (e.g., for erasure coding, data sharding, metadata communication and redundancy, storage capacity expansion or contraction, data recovery, and so on).

[0042] Fig. 4A depicts a file 402, suitable for storage and access in embodiments of the storage cluster 160, with file chunks 406 each having a file offset 404 relative to the file 402. The file 402 is identified by an inode number 410, for example as a file system object. Directories may also be identified by inode numbers, in some embodiments. The file offset 404 of a specific file chunk 406 is the address of the beginning of the file chunk 406 relative to the beginning of the file 402 in this example, although other definitions for offsets could be used. A file 402 of a given file size 408 could have few or many file chunks 406, which could be various sizes. For example, a file chunk 406 could be a number of kilobytes, a number of megabytes, a number of gigabytes, or conceivably a number of terabytes, etc. In typical, known storage systems, ownership of all of the file chunks 406 for a specific file, or even all of the file chunks 406 for all files in the system, is found in one location in the storage system, such as in a storage node. This creates a bottleneck, especially for large files, which ties up system resources and inhibits throughput. By contrast, ownership of file chunks 406 is distributed throughout storage nodes 150 in present embodiments of the storage cluster 160. This more evenly distributes or partitions system resource load across a file.

[0043] Fig. 4B shows an example file 402, and ownership of various file chunks 406 of the file 402, in accordance with an embodiment of the present description. In this example, the file 402 has a file size 408 of 1 TB (terabyte), and chunk size of 16 kB (kilobytes) although other file sizes 408 and sizes for file chunks 406 could be used. The file 402 is identified by “inode 65” as the inode number 410. File chunks 406 are owned by various authorities 168. For example, the first file chunk 406 in the example is shown as owned by an authority 168 identified as “authority 42”, the second file chunk 406 is shown as owned by

an authority 168 identified as “authority 50”, and the third file chunk 406 is shown as owned by an authority 168 identified as “authority 40”. It is likely most file chunks 406 of the file 402 are owned by different authorities 168, but it is statistically possible and indeed in very large files, statistically likely, that two or more file chunks 406 could be owned by the same authority 168. It should be appreciated that this statistical likelihood depends upon the total number of authorities 168 in the storage cluster 160 as compared to the total number of file chunks 406 in the file 402.

[0044] Fig. 5 is a block diagram of a storage cluster 160 with components for determining ownership of file chunks 406. In some embodiments, each storage node 150 of the storage cluster 160 has one or more authorities 168, a hash calculator 502, and storage memory 506, although fewer than all of the storage nodes 150 could have these components, with some storage nodes 150 as compute nodes. User data is stored in storage memory 506 of the storage nodes 150, for example using erasure coding and distribution in accordance with authorities 168 as described above with reference to Figs. 1-3. Operation of the hash calculator 502 is described below with reference to Fig. 6. The hash calculator 502 could be implemented in firmware, hardware, or software executing on a processor, or combinations thereof. Operation of the hash calculator 502 could be performed by a processor 156 of a storage node 150 (see Fig. 2). It should be appreciated that storage nodes 150 could include one or more logical constructs of a compute agent holding storage information, and/or one or more logical constructs of a compute agent configured to service a client request, in a storage system.

[0045] Fig. 6 is a system action diagram showing use of a hash calculator 502 that identifies an authority 168 as an owner of a file chunk 406. The inode number 410, as an identifier of a file 402, and the file offset 404 of a file chunk 406 relative to the file 402, are fed into the hash calculator 502, which produces a hash calculation result 602. This could be implemented by having a processor 156 of a storage node 150 perform a hash calculation on the file identifier and the offset 404. The hash calculation result 602 identifies an authority 168. The messaging system of the storage system enables communication to the authority identified by the hash calculation result. Inode numbers 410 or other file identifiers, file offsets 404 or other information about file chunks 406, file sizes 408, and information handled by authorities 168 relating to files 402, file chunks 406, directories, etc., may be referred to as types of metadata. Figs. 4A-6 thus describe a mechanism for delegating the location of portions of metadata for files, which can also be used for various further objects. It should be appreciated that some storage nodes of the system may not include an authority

in some embodiments as a portion or subset of the storage nodes may include an authority, as opposed to each storage node including an authority in other embodiments.

[0046] Fig. 7A is a flow diagram of a method for distributing ownership of file chunks in a storage system, and writing file chunks of files, which can be practiced by embodiments of the storage cluster, storage nodes and storage units of Figs. 1-6. The method can be practiced by one or more processors in a storage system. In an action 702, a request is received to write a file to storage memory. The request could be received by a storage node. In an action 704, the file is portioned or segmented into file chunks. In an action 706, for each file chunk of the file, a hash calculation is performed on the file offset of the file chunk and a file identifier, e.g., an inode number of the file. In an action 708, for each file chunk the authority is identified based on the hash calculation result for that file chunk. In an action 710, the authority directs storage of the file chunk. The file chunks are stored in accordance with direction by the associated authorities, in an action 712.

[0047] Fig. 7B is a flow diagram of a method for distributing ownership of file chunks in a storage system, and accessing file chunks of files, which can be practiced by embodiments of the storage cluster, storage nodes and storage units of Figs. 1-6. The method can be practiced by one or more processors in a storage system. In an action 714, a request is received to read a file from storage memory. The request could be received by a storage node. In an action 716, an authority is determined or identified among differing storage nodes of the storage system for each file chunk of the file. In some embodiments, a hash calculation is performed on the file offset of the file chunk and a file identifier, e.g., an inode number of the file, for each file chunk of the file. In an action 718, location information for the file chunk for each file chunk is received from the authority having ownership of the file chunk. In an action 720, the file chunks of the file are accessed as directed by each of the determined authorities. It should be appreciated that the embodiments determine one of a plurality of owners in a plurality of storage nodes as having ownership of a file chunk of the file, and the determining is performed at differing storage nodes for each of a plurality of file chunks of the file. Thus, the distribution of ownership of chunks of a file enables the efficient access to those file chunks without the need of a centralized look up table or agent to manage the process. That is, differing authorities distributed among differing storage nodes have ownership for various chunks or portions of a file according to the mechanisms described above.

[0048] It should be appreciated that the methods described herein may be performed with a digital processing system, such as a conventional, general-purpose computer system.

Special purpose computers, which are designed or programmed to perform only one function may be used in the alternative. Fig. 8 is an illustration showing an exemplary computing device which may implement the embodiments described herein. The computing device of Fig. 8 may be used to perform embodiments of the functionality for distributing ownership of file chunks in a storage system in accordance with some embodiments. The computing device includes a central processing unit (CPU) 801, which is coupled through a bus 805 to a memory 803, and mass storage device 807. Mass storage device 807 represents a persistent data storage device such as a disc drive, which may be local or remote in some embodiments. The mass storage device 807 could implement a backup storage, in some embodiments. Memory 803 may include read only memory, random access memory, etc. Applications resident on the computing device may be stored on or accessed via a computer readable medium such as memory 803 or mass storage device 807 in some embodiments. Applications may also be in the form of modulated electronic signals modulated accessed via a network modem or other network interface of the computing device. It should be appreciated that CPU 801 may be embodied in a general-purpose processor, a special purpose processor, or a specially programmed logic device in some embodiments.

[0049] Display 811 is in communication with CPU 801, memory 803, and mass storage device 807, through bus 805. Display 811 is configured to display any visualization tools or reports associated with the system described herein. Input/output device 809 is coupled to bus 805 in order to communicate information in command selections to CPU 801. It should be appreciated that data to and from external devices may be communicated through the input/output device 809. CPU 801 can be defined to execute the functionality described herein to enable the functionality described with reference to Figs. 1-7B. The code embodying this functionality may be stored within memory 803 or mass storage device 807 for execution by a processor such as CPU 801 in some embodiments. The operating system on the computing device may be MS-WINDOWS™, UNIX™, LINUX™, iOS™, CentOS™, Android™, Redhat Linux™, z/OS™, or other known operating systems. It should be appreciated that the embodiments described herein may also be integrated with a virtualized computing system implemented with physical computing resources.

[0050] Detailed illustrative embodiments are disclosed herein. However, specific functional details disclosed herein are merely representative for purposes of describing embodiments. Embodiments may, however, be embodied in many alternate forms and should not be construed as limited to only the embodiments set forth herein. It should be

appreciated that while the embodiments are described with respect to file systems, the embodiments may be extended to object based systems also.

[0051] It should be understood that although the terms first, second, etc. may be used herein to describe various steps or calculations, these steps or calculations should not be limited by these terms. These terms are only used to distinguish one step or calculation from another. For example, a first calculation could be termed a second calculation, and, similarly, a second step could be termed a first step, without departing from the scope of this disclosure. As used herein, the term “and/or” and the “/” symbol includes any and all combinations of one or more of the associated listed items.

[0052] As used herein, the singular forms “a”, “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms “comprises”, “comprising”, “includes”, and/or “including”, when used herein, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof. Therefore, the terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting.

[0053] It should also be noted that in some alternative implementations, the functions/acts noted may occur out of the order noted in the figures. For example, two figures shown in succession may in fact be executed substantially concurrently or may sometimes be executed in the reverse order, depending upon the functionality/acts involved.

[0054] With the above embodiments in mind, it should be understood that the embodiments might employ various computer-implemented operations involving data stored in computer systems. These operations are those requiring physical manipulation of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. Further, the manipulations performed are often referred to in terms, such as producing, identifying, determining, or comparing. Any of the operations described herein that form part of the embodiments are useful machine operations. The embodiments also relate to a device or an apparatus for performing these operations. The apparatus can be specially constructed for the required purpose, or the apparatus can be a general-purpose computer selectively activated or configured by a computer program stored in the computer. In particular, various general-purpose machines can be used with computer programs written

in accordance with the teachings herein, or it may be more convenient to construct a more specialized apparatus to perform the required operations.

[0055] A module, an application, a layer, an agent or other method-operable entity could be implemented as hardware, firmware, or a processor executing software, or combinations thereof. It should be appreciated that, where a software-based embodiment is disclosed herein, the software can be embodied in a physical machine such as a controller. For example, a controller could include a first module and a second module. A controller could be configured to perform various actions, e.g., of a method, an application, a layer or an agent.

[0056] The embodiments can also be embodied as computer readable code on a non-transitory computer readable medium. The computer readable medium is any data storage device that can store data, which can be thereafter read by a computer system. Examples of the computer readable medium include hard drives, network attached storage (NAS), read-only memory, random-access memory, CD-ROMs, CD-Rs, CD-RWs, magnetic tapes, and other optical and non-optical data storage devices. The computer readable medium can also be distributed over a network coupled computer system so that the computer readable code is stored and executed in a distributed fashion. Embodiments described herein may be practiced with various computer system configurations including hand-held devices, tablets, microprocessor systems, microprocessor-based or programmable consumer electronics, minicomputers, mainframe computers and the like. The embodiments can also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a wire-based or wireless network.

[0057] Although the method operations were described in a specific order, it should be understood that other operations may be performed in between described operations, described operations may be adjusted so that they occur at slightly different times or the described operations may be distributed in a system which allows the occurrence of the processing operations at various intervals associated with the processing.

[0058] 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, 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 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.

[0059] Various units, circuits, or other components may be described or claimed as “configured to” perform a task or tasks. In such contexts, the phrase “configured to” is used to connote structure by indicating that the units/circuits/components include structure (e.g., circuitry) that performs the task or tasks during operation. As such, the unit/circuit/component can be said to be configured to perform the task even when the specified unit/circuit/component is not currently operational (e.g., is not on). The units/circuits/components used with the “configured to” language include hardware--for example, circuits, memory storing program instructions executable to implement the operation, etc. Reciting that a unit/circuit/component is “configured to” perform one or more tasks is expressly intended not to invoke 35 U.S.C. 112, sixth paragraph, for that unit/circuit/component. Additionally, “configured to” can include generic structure (e.g., generic circuitry) that is manipulated by software and/or firmware (e.g., an FPGA or a general-purpose processor executing software) to operate in manner that is capable of performing the task(s) at issue. “Configured to” may also include adapting a manufacturing process (e.g., a semiconductor fabrication facility) to fabricate devices (e.g., integrated circuits) that are adapted to implement or perform one or more tasks.

[0060] The foregoing description, for the purpose of explanation, has been described with reference to specific embodiments. However, the illustrative discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The embodiments were chosen and described in order to best explain the principles of the embodiments and its practical applications, to thereby enable others skilled in the art to best utilize the embodiments and various modifications as may be suited to the particular use contemplated. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

CLAIMS

What is claimed is:

1. A method for accessing a file in a storage system, comprising:
 - determining, for each file chunk of the file, an authority among differing storage nodes of the storage system;
 - receiving from the authority having ownership of the file chunk, location information for the file chunk; and
 - accessing file chunks of the file as directed by each of the determined authorities.
2. The method of claim 1, wherein the determining includes performing a hash calculation on an identifier of the file and an offset of the file chunk.
3. The method of claim 1, wherein each file chunk of the file is associated with an offset value and wherein differing authorities are associated with differing file chunks of the file based on the offset value corresponding to the differing file chunks.
4. The method of claim 1, wherein:
 - the determining includes performing a hash calculation;
 - a result of the hash calculation indicates the authority; and
 - differing ones of the plurality of storage nodes of the storage system contain differing authorities for differing file chunks of the file.
5. The method of claim 1, wherein the determining includes, for each file chunk of the file, performing a hash calculation based on an offset of the file chunk and an inode number of the file, and wherein a result of the hash calculation indicates the authority having ownership of the file chunk.
6. The method of claim 3, wherein the differing authorities are distributed among differing storage nodes.
7. The method of claim 1, wherein the determining for each file chunk is not centralized in the storage system.
8. A method for accessing file chunks of a file in a storage system, comprising:
 - determining one of a plurality of owners in a plurality of storage nodes as having ownership of a file chunk of the file, the determining performed at differing storage nodes for each of a plurality of file chunks of the file; and
 - accessing the plurality of file chunks of the file as directed by determined owners.
9. The method of claim 8, wherein each file chunk of the file is associated with an offset value and wherein differing authorities are associated with differing file chunks of the file based on the offset value corresponding to the differing file chunks.

10. The method of claim 8, wherein the determining comprises:
 - performing a hash calculation on a file offset for the file chunk and on an inode number of the file, wherein a result of the hash calculation identifies one of the plurality of owners as having the ownership of the file chunk.
11. The method of claim 8, wherein the determining comprises:
 - performing a hash calculation having a result that indicates an authority of one of the plurality of storage nodes having the ownership of the file chunk.
12. The method of claim 8, the determining is not centralized in the storage system.
13. The method of claim 11, wherein distributing the ownership of differing subsets of file chunks among differing owners in differing storage nodes partitions a load of the storage system across the file for a plurality of files of the storage system.
14. A storage system having efficient distribution of file chunks, comprising:
 - a plurality of storage nodes coupled as a storage cluster;
 - at least a portion of the plurality of storage nodes configured to have at least one authority, each authority having ownership of a subset of file chunks, each file chunk including a portion of a file; and
 - each of the plurality of storage nodes configured to determine which authority in which of the plurality of storage nodes has information regarding location in the storage cluster of each file chunk.
15. The storage cluster of claim 14, further comprising:
 - each of the plurality of storage nodes configured to perform a hash calculation based on a file offset of a file chunk and an inode number, wherein a result of the hash calculation indicates an authority having the information regarding the location of the file chunk.
16. The storage cluster of claim 14, further comprising:
 - each of the plurality of storage nodes configured to perform a hash calculation, based on a file offset of a file chunk and an identifier of the file, and to apply a result of the hash calculation to a lookup table, to determine the authority having the information regarding the location of the file chunk.
17. The storage cluster of claim 14, further comprising:
 - each of the plurality of storage nodes configured to perform a hash calculation on an inode number of the file and on the file offset, to obtain a result of the hash calculation that indicates an authority, and one of the plurality of storage nodes having the authority.

18. The storage cluster of claim 14, wherein determining which authority in which of the plurality of storage nodes has information regarding location in the storage cluster of each file chunk is not centralized in the storage system.
19. The storage cluster of claim 14, further comprising:
 - each of the plurality of storage nodes including one or more of a logical construct of a compute agent holding storage information, or a logical construct of a compute agent configured to service a client request.
20. The storage cluster of claim 14, wherein the storage cluster is housed within a single chassis and wherein each storage node of the plurality of storage nodes includes flash memory.




FIG. 1




FIG. 2

3/7

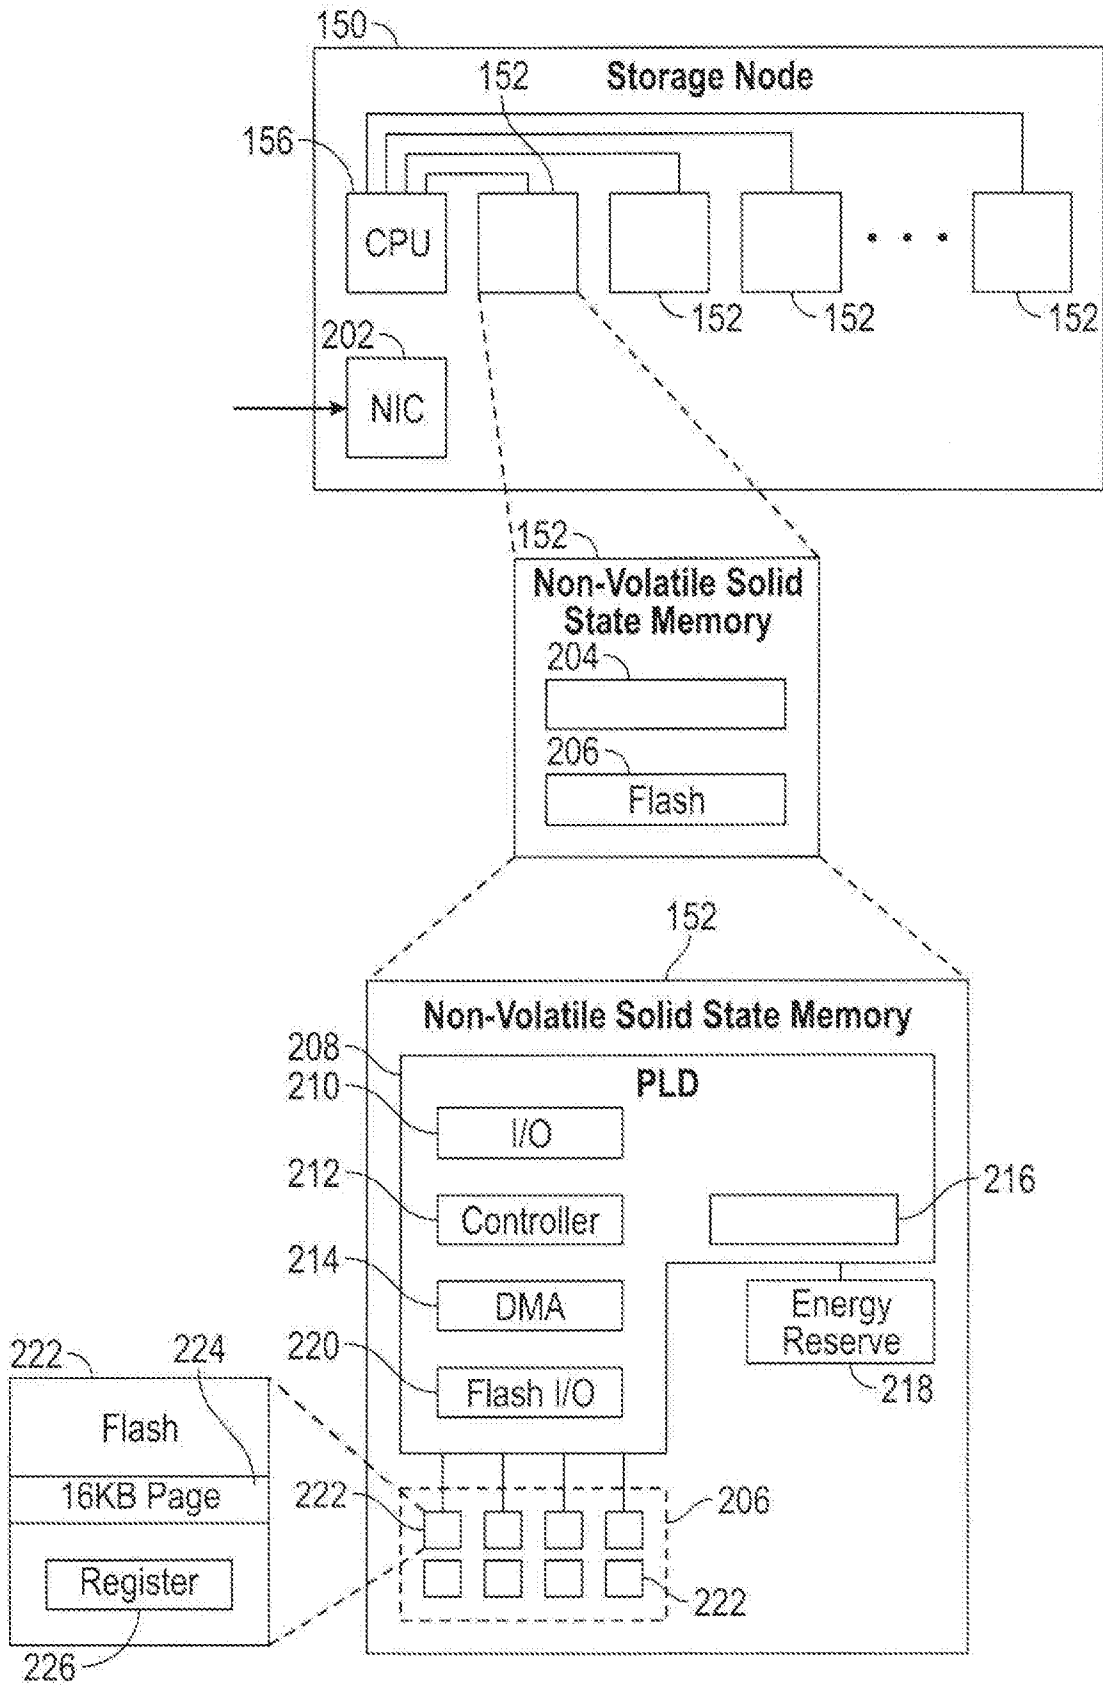


FIG. 3




FIG. 4A




FIG. 4B

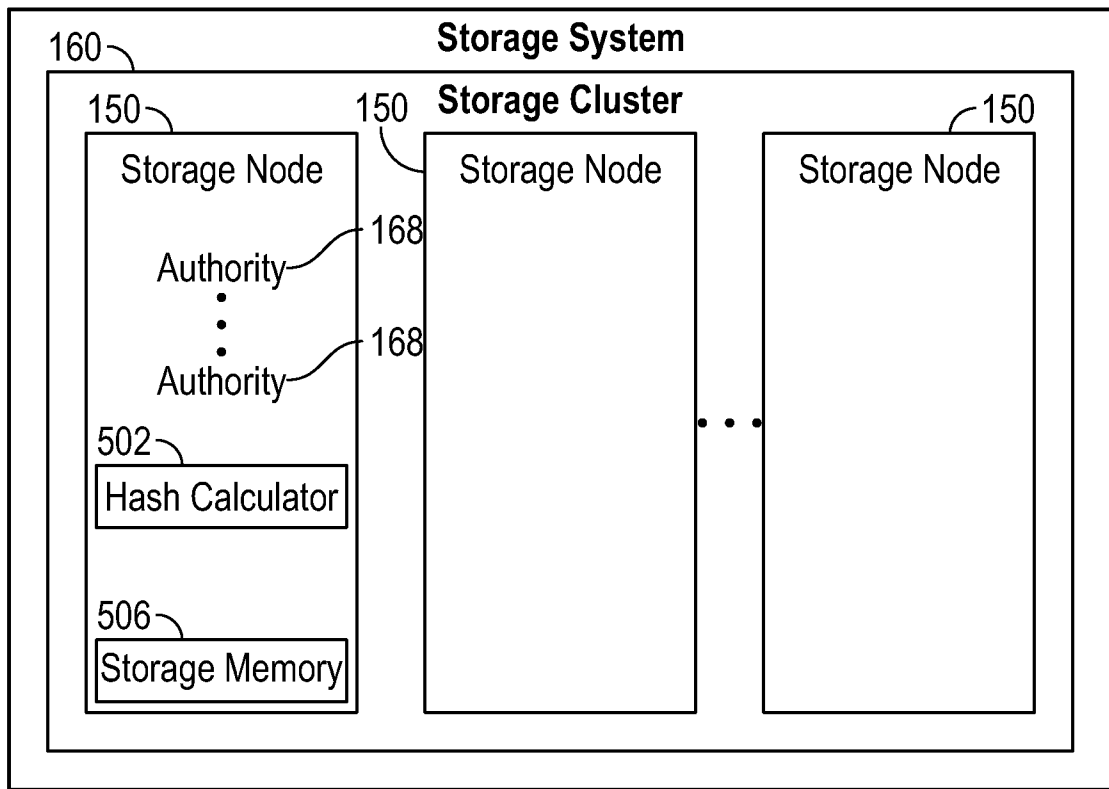


FIG. 5

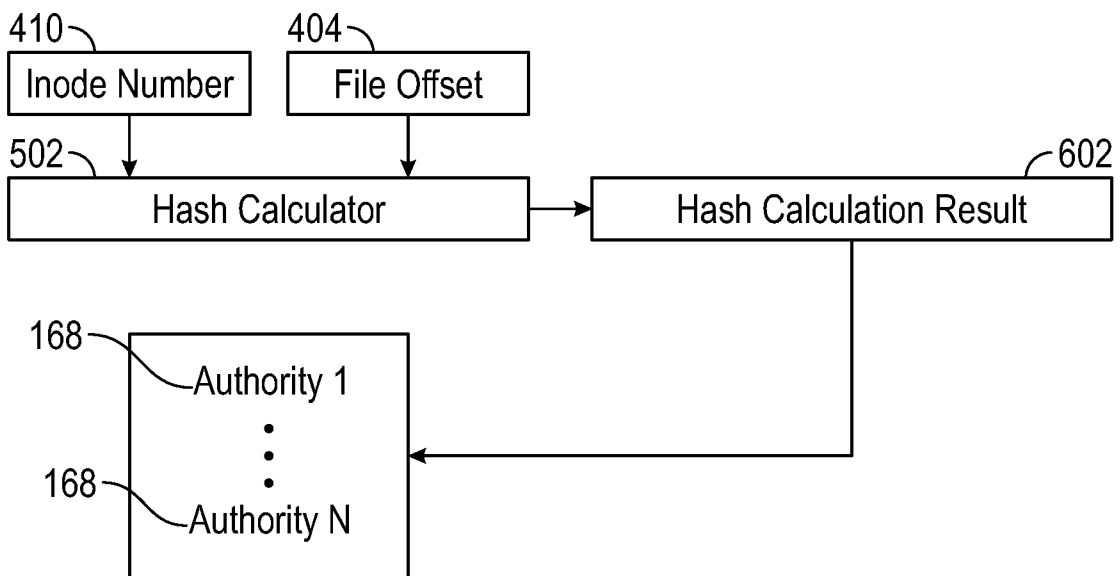


FIG. 6

6/7

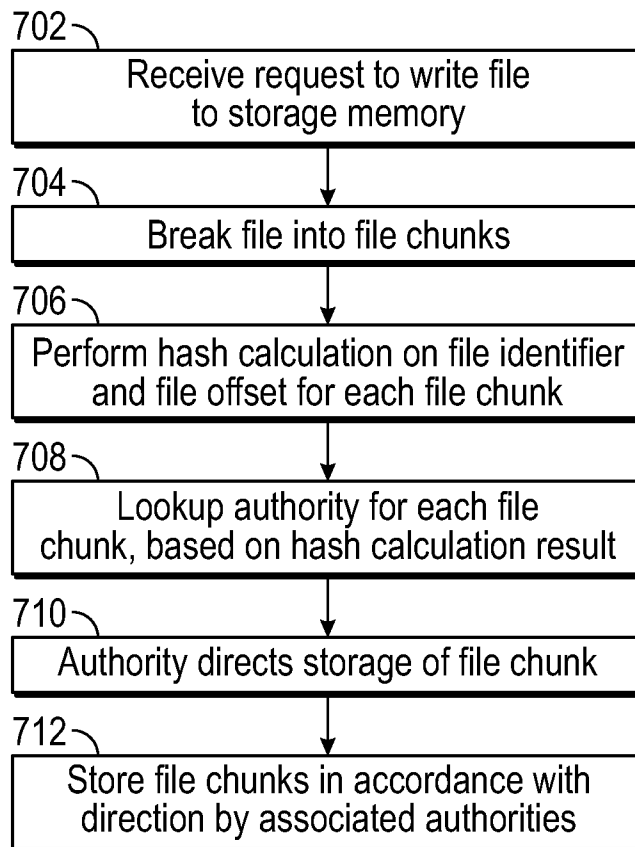


FIG. 7A

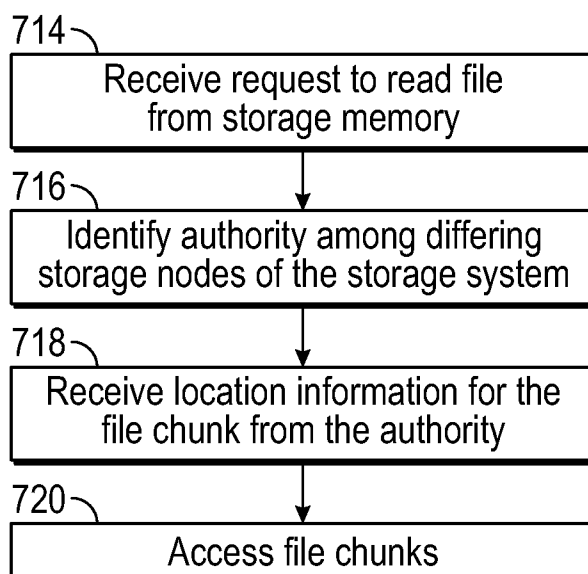


FIG. 7B




FIG. 8

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/041920**A. CLASSIFICATION OF SUBJECT MATTER****G06F 21/62(2013.01)i, G06F 12/14(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 21/62; G06F 7/00; G06F 17/30; G06F 12/00; G06F 12/08; G06F 12/14

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: data chunk, storage system, ownership, authority, distribute

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2012-0078844 A1 (TAE WOONG KIM) 29 March 2012 See paragraphs [0010], [0037]-[0041], [0078]; and figures 2-3.	1-20
Y	US 2012-0166401 A1 (JIN LI et al.) 28 June 2012 See paragraphs [0024]-[0025], [0029], [0043], [0079]-[0103], [0106]; and figures 1-3.	1-20
A	US 2011-0313973 A1 (MANDAYAM C. SRIVAS et al.) 22 December 2011 See paragraphs [0082]-[0091]; claims 1, 5; and figures 3, 6-9.	1-20
A	US 2013-0103708 A1 (CHANG-SOO KIM et al.) 25 April 2013 See paragraphs [0049]-[0063]; and figures 1-2.	1-20
A	US 2011-0029730 A1 (COLIN D. DUROCHER et al.) 03 February 2011 See paragraphs [0027]-[0044]; and figures 3(a)-3(b)	1-20

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

07 October 2016 (07.10.2016)

Date of mailing of the international search report

11 October 2016 (11.10.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

CHIN, Sang Bum

Telephone No. +82-42-481-8398



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/041920

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012-0078844 A1	29/03/2012	JP 2012-074039 A KR 10-1374655 B1 KR 10-2012-0032920 A	12/04/2012 24/03/2014 06/04/2012
US 2012-0166401 A1	28/06/2012	CN 102591946 A EP 2659376 A2 US 2016-0012098 A1 US 9110936 B2 WO 2012-092212 A2 WO 2012-092212 A3	18/07/2012 06/11/2013 14/01/2016 18/08/2015 05/07/2012 04/10/2012
US 2011-0313973 A1	22/12/2011	EP 2583180 A1 JP 2013-541057 A JP 5671615 B2 US 9323775 B2 WO 2011-160054 A1	24/04/2013 07/11/2013 18/02/2015 26/04/2016 22/12/2011
US 2013-0103708 A1	25/04/2013	KR 10-1601877 B1 US 9378218 B2	09/03/2016 28/06/2016
US 2011-0029730 A1	03/02/2011	EP 2460079 A1 JP 2013-501272 A JP 5373199 B2 US 8627015 B2 WO 2011-014724 A1	06/06/2012 10/01/2013 18/12/2013 07/01/2014 03/02/2011