



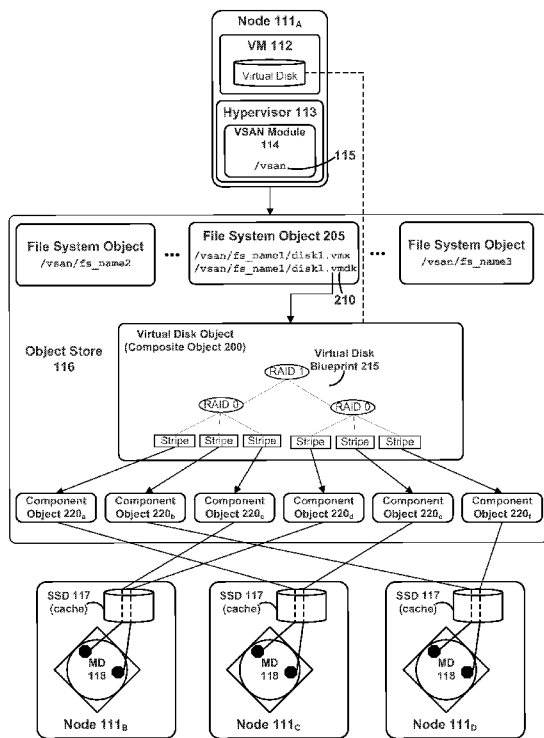
- (51) International Patent Classification:
G06F 17/30 (2006.01)
- (21) International Application Number:
PCT/US2014/041761
- (22) International Filing Date:
10 June 2014 (10.06.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/010,293 26 August 2013 (26.08.2013) US
- (71) Applicant: VMWARE, INC. [US/US]; 3401 Hillview Avenue, Palo Alto, California 94304 (US).
- (72) Inventors: KARAMANOLIS, Christos; 3401 Hillview Avenue, Palo Alto, California 94304 (US). VASANI, Soam; 3401 Hillview Avenue, Palo Alto, California 94304 (US).
- (74) Agents: KIRYAKOZA, James D. et al.; 3401 Hillview Avenue, Palo Alto, California 94304 (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, LT, 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, 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).

[Continued on next page]

(54) Title: SCALABLE DISTRIBUTED STORAGE ARCHITECTURE



(57) Abstract: Techniques are disclosed for providing a file system interface for an object store intended to support simultaneous access to objects stored in the object store by multiple clients. In accordance with one method, an abstraction of a root directory to a hierarchical namespace for the object store is exposed to clients. The object store is backed by a plurality of physical storage devices housed in or directly attached to the plurality of host computers and internally tracks its stored objects using a flat namespace that maps unique identifiers to the stored objects. The creation of top-level objects appearing as subdirectories of the root directory is enabled, wherein each top-level object represents a separate abstraction of a storage device having a separate namespace that can be organized in accordance with any designated file system.

Published:

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

SCALABLE DISTRIBUTED STORAGE ARCHITECTURE

BACKGROUND

[0001] Distributed systems allow multiple clients in a network to access a pool of shared resources. For example, a distributed storage system allows a cluster of host computers to aggregate local disks (e.g., SSD, PCI-based flash storage, SATA, or SAS magnetic disks) located in or attached to each host computer to create a single and shared pool of storage. This pool of storage (sometimes referred to herein as a “datastore” or “store”) is accessible by all host computers in the cluster and may be presented as a single namespace of storage entities (such as a hierarchical file system namespace in the case of files, a flat namespace of unique identifiers in the case of objects, etc.). Storage clients in turn, such as virtual machines spawned on the host computers may use the datastore, for example, to store virtual disks that are accessed by the virtual machines during their operation. Because the shared local disks that make up the datastore may have different performance characteristics (e.g., capacity, input/output per second or IOPS capabilities, etc.), usage of such shared local disks to store virtual disks or portions thereof may be distributed among the virtual machines based on the needs of each given virtual machine.

[0002] This approach provides enterprises with cost-effective performance. For instance, distributed storage using pooled local disks is inexpensive, highly scalable, and relatively simple to manage. Because such distributed storage can use commodity disks in the cluster, enterprises do not need to invest in additional storage infrastructure. However, one challenge that arises relates to developing a mechanism to efficiently track where objects are stored across the commodity disks in the cluster as well as how to efficiently access them when needed. For example, while utilizing a flat namespace may provide a simplistic and efficient means to store and retrieve objects, it does not provide enough flexibility to create hierarchical relationships between objects that may be useful in organizing objects in a manner that is compatible with the existing interfaces of clients or that otherwise satisfies the different storage requirements of different clients. For example, many pre-existing environments that could leverage such a scalable object store (e.g., applications, storage management tools, virtualization hypervisors, etc.) may require that the object

store provide a certain hierarchical file system based storage interface. One example is vSphere Hypervisor from VMware, Inc. which stores virtual machine metadata in the form of files in a hierarchical file system. Furthermore, because the datastore is shared among the cluster of host computers which may simultaneously access the same data stored therein, any file system that is used to manage the single namespace provided by the datastore needs to have mechanisms for concurrency control. Current distributed or clustered file systems typically provide some form of concurrency control. However, due to limitations in their inherent design, such current file systems typically have limitations on the number of simultaneous “clients” (e.g., host computers or virtual machines that access the file system) they can support. If a current file system used to manage the datastore has limits on the number of clients that can simultaneously access it, then the utility of a highly scalable datastore can plateau, even if additional commodity storage can be easily added to the datastore to increase its capacity, since no additional clients can be added to access such additional storage.

SUMMARY

[0003] One embodiment of the present disclosure relates to a method for providing a file system interface for an object store intended to support simultaneous access to objects stored in the object store by multiple clients. In accordance with the method, an abstraction of a root directory to a hierarchical namespace for the object store is exposed to clients. The object store is backed by a plurality of physical storage devices housed in or directly attached to the plurality of host computers and internally tracks its stored objects using a flat namespace that maps unique identifiers to the stored objects. The creation of top-level objects appearing as subdirectories of the root directory is enabled, wherein each top-level object represents a separate abstraction of a storage device having a separate namespace that can be organized in accordance with any designated file system.

[0004] By layering a hierarchical namespace that can be used by clients while using a flat namespace to internally store and access objects, techniques herein are able to offer a framework that supports a large scalable clustered file system using a distribute of commodity storage resources. For example, providing the capability to create top-level objects representing file systems that contain file objects that may ultimately be accessed by clients, the foregoing object store offers higher scalability

than existing distributed or clustered file systems because it is not confined, for example, by any limitations on the number of simultaneous clients inherent in the design of any particular distributed clustered file system (e.g., VMware VMFS, NFS, etc.) configured for any particular file system object in the object store. That is, implementing a clustered file system on top of an object store has scalability advantages, because different parts of the clustered file system can be placed on different objects such that scalability requirements from particular clients need to be achieved only by the subset of such clients that access data in a certain sub-space of the namespace which, in turn, resides on a certain object.

[0005] Other embodiments include, without limitation, a computer-readable medium that includes instructions that enable a processing unit to implement one or more aspects of the disclosed methods as well as a computer system having a processor, memory, and modules configured to implement one or more aspects of the disclosed methods.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] Figure 1 illustrates an example computing environment, according to one embodiment.

[0007] Figure 2 illustrates an example hierarchical structure of objects organized within an object store that represent a virtual disk, according to one embodiment.

[0008] Figure 3 illustrates components of a VSAN module, according to one embodiment.

[0009] Figure 4 illustrates a method flow diagram for creating a virtual disk object based on a defined storage policy, according to one embodiment.

[0010] Figure 5 illustrates the handling of an I/O operation originating from a VM, according to one embodiment.

DETAILED DESCRIPTION

[0011] Figure 1 illustrates a computing environment 100, according to one embodiment. As shown, computing environment 100 is a software-based “virtual storage area network” (VSAN) environment that leverages the commodity local storage housed in or directly attached (hereinafter, use of the term “housed” or “housed in” may be used to encompass both housed in or otherwise directly attached)

to host servers or nodes 111 of a cluster 110 to provide an aggregate object store 116 to virtual machines (VMs) 112 running on the nodes. The local commodity storage housed in or otherwise directly attached to the nodes 111 may include combinations of solid state drives (SSDs) 117 and/or magnetic or spinning disks 118. In certain embodiments, SSDs 117 serve as a read cache and/or write buffer in front of magnetic disks 118 to increase I/O performance. As further discussed below, each node 111 may include a storage management module (referred to herein as a “VSAN module”) in order to automate storage management workflows (e.g., create objects in the object store, etc.) and provide access to objects in the object store (e.g., handle I/O operations to objects in the object store, etc.) based on predefined storage policies specified for objects in the object store. For example, because a VM may be initially configured by an administrator to have specific storage requirements for its “virtual disk” depending its intended use (e.g., capacity, availability, IOPS, etc.), the administrator may define a storage profile or policy for each VM specifying such availability, capacity, IOPS and the like. As further described below, the VSAN module may then create an “object” for the specified virtual disk by backing it with physical storage resources of the object store based on the defined policy

[0012] A virtualization management platform 105 is associated with cluster 110 of nodes 111. Virtualization management platform 105 enables an administrator to manage the configuration and spawning of VMs on the various nodes 111. As depicted in the embodiment of Figure 1, each node 111 includes a virtualization layer or hypervisor 113, a VSAN module 114, and hardware 119 (which includes the SSDs 117 and magnetic disks 118 of a node 111). Through hypervisor 113, a node 111 is able to launch and run multiple VMs 112. Hypervisor 113, in part, manages hardware 119 to properly allocate computing resources (e.g., processing power, random access memory, etc.) for each VM 112. Furthermore, as described further below, each hypervisor 113, through its corresponding VSAN module 114, provides access to storage resources located in hardware 119 (e.g., SSDs 117 and magnetic disks 118) for use as storage for virtual disks (or portions thereof) and other related files that may be accessed by any VM 112 residing in any of nodes 111 in cluster 110. In a particular embodiment, vSphere Hypervisor from VMware, Inc. (VMware) may be installed on nodes 111 as hypervisor 113 and vCenter Server from VMware may be used as virtualization management platform 105.

[0013] In one embodiment, VSAN module 114 is implemented as a “VSAN” device driver within hypervisor 113. In such an embodiment, VSAN module 114 provides access to a conceptual “VSAN” 115 through which an administrator can create a number of top-level “device” or namespace objects that are backed by object store 116. In one common scenario, during creation of a device object, the administrator may specify a particular file system for the device object (such device objects hereinafter also thus referred to “file system objects”). For example, in one embodiment, each hypervisor 113 in each node 111 may, during a boot process, discover a `/vsan/` root node for a conceptual global namespace that is exposed by VSAN module 114. By, for example, accessing APIs exposed by VSAN module 114, hypervisor 113 can then determine all the top-level file system objects (or other types of top-level device objects) currently residing in VSAN 115. When a VM (or other client) attempts to access one of the file system objects, hypervisor 113 may dynamically “auto-mount” the file system object at that time. In certain embodiments, file system objects may further be periodically “auto-unmounted” when access to objects in the file system objects cease or are idle for a period of time. A file system object (e.g., `/vsan/fs_name1`, etc.) that is accessible through VSAN 115 may, for example, be implemented to emulate the semantics of a particular file system such as VMware’s distributed or clustered file system, VMFS, which is designed to provide concurrency control among simultaneously accessing VMs. Because VSAN 115 supports multiple file system objects, it is able provide storage resources through object store 116 without being confined by limitations of any particular clustered file system. For example, many clustered file systems (e.g., VMFS, etc.) can only scale to support a certain amount of nodes 111. By providing multiple top-level file system object support, VSAN 115 overcomes the scalability limitations of such clustered file systems.

[0014] As described in further detail in the context of Figure 2 below, a file system object, may, itself, provide access to a number of virtual disk descriptor files (e.g., `.vmdk` files in a vSphere environment, etc.) accessible by VMs 112 running in cluster 110. These virtual disk descriptor files contain references to virtual disk “objects” that contain the actual data for the virtual disk and are separately backed by object store 116. A virtual disk object may itself be a hierarchical or “composite” object that, as described further below, is further composed of “component” objects

(again separately backed by object store 116) that reflect the storage requirements (e.g., capacity, availability, IOPs, etc.) of a corresponding storage profile or policy generated by the administrator when initially creating the virtual disk. As further discussed below, each VSAN module 114 (through a cluster level object management or “CLOM” sub-module, in embodiments as further described below) communicates with other VSAN modules 114 of other nodes 111 to create and maintain an in-memory metadata database (e.g., maintained separately but in synchronized fashion in the memory of each node 111) that contains metadata describing the locations, configurations, policies and relationships among the various objects stored in object store 116. This in-memory metadata database is utilized by a VSAN module 114 on a node 111, for example, when an administrator first creates a virtual disk for a VM as well as when the VM is running and performing I/O operations (e.g., read or write) on the virtual disk. As further discussed below in the context of Figure 3, VSAN module 114 (through a document object manager or “DOM” sub-module, in one embodiment as further described below) traverses a hierarchy of objects using the metadata in the in-memory database in order to properly route an I/O operation request to the node (or nodes) that houses (house) the actual physical local storage that backs the portion of the virtual disk that is subject to the I/O operation.

[0015] Figure 2 illustrates an example hierarchical structure of objects organized within object store 116 that represent a virtual disk, according to one embodiment. As previously discussed above, a VM 112 running on one of nodes 111 may perform I/O operations on a virtual disk that is stored as a hierarchical or composite object 200 in object store 116. Hypervisor 113 provides VM 112 access to the virtual disk by interfacing with the abstraction of VSAN 115 through VSAN module 114 (e.g., by auto-mounting the top-level file system object corresponding to the virtual disk object, as previously discussed, in one embodiment). For example, VSAN module 114, by querying its local copy of the in-memory metadata database, is able to identify a particular file system object 205 (e.g., a VMFS file system object in one embodiment, etc.) stored in VSAN 115 that stores a descriptor file 210 for the virtual disk (e.g., a .vmdk file, etc.). It should be recognized that the file system object 205 may store a variety of other files consistent with its purpose, such as virtual machine configuration files (e.g., .vmx files in a vSphere environment, etc.) and the like when supporting a virtualization environment. In certain embodiments, each file system

object may be configured to support only those virtual disks corresponding to a particular VM (e.g., a “per-VM” file system object).

[0016] Descriptor file 210 includes a reference to composite object 200 that is separately stored in object store 116 and conceptually represents the virtual disk (and thus may also be sometimes referenced herein as a virtual disk object). Composite object 200 stores metadata describing a storage organization or configuration for the virtual disk (sometimes referred to herein as a virtual disk “blueprint”) that suits the storage requirements or service level agreements (SLAs) in a corresponding storage profile or policy (e.g., capacity, availability, IOPs, etc.) generated by an administrator when creating the virtual disk. For example, in the embodiment of Figure 2, composite object 200 includes a virtual disk blueprint 215 that describes a RAID 1 configuration where two mirrored copies of the virtual disk (e.g., mirrors) are each further striped in a RAID 0 configuration. Composite object 225 may thus contain references to a number of “leaf” or “component” objects 220_x corresponding to each stripe (e.g., data partition of the virtual disk) in each of the virtual disk mirrors. The metadata accessible by VSAN module 114 in the in-memory metadata database for each component object 220 (e.g., for each stripe) provides a mapping to or otherwise identifies a particular node 111_x in cluster 110 that houses the physical storage resources (e.g., magnetic disks 118, etc.) that actually store the stripe (as well as the location of the stripe within such physical resource).

[0017] Figure 3 illustrates components of a VSAN module 114, according to one embodiment. As previously described, in certain embodiments, VSAN module 114 may execute as a device driver exposing an abstraction of a VSAN 115 to hypervisor 113. Various sub-modules of VSAN module 114 handle different responsibilities and may operate within either user space 315 or kernel space 320 depending on such responsibilities. As depicted in the embodiment of Figure 3, VSAN module 114 includes a cluster level object management (CLOM) sub-module 325 that operates in user space 315. CLOM sub-module 325 generates virtual disk blueprints during creation of a virtual disk by an administrator and ensures that objects created for such virtual disk blueprints are configured to meet storage profile or policy requirements set by the administrator. In addition to being accessed during object creation (e.g., for virtual disks), CLOM sub-module 325 may also be accessed (e.g., to dynamically revise or otherwise update a virtual disk blueprint or the mappings of the virtual disk

blueprint to actual physical storage in object store 116) on a change made by an administrator to the storage profile or policy relating to an object or when changes to the cluster or workload result in an object being out of compliance with a current storage profile or policy.

[0018] In one embodiment, if an administrator creates a storage profile or policy for a composite object such as virtual disk object 200, CLOM sub-module 325 applies a variety of heuristics and/or distributed algorithms to generate virtual disk blueprint 215 that describes a configuration in cluster 110 that meets or otherwise suits the storage policy (e.g., RAID configuration to achieve desired redundancy through mirroring and access performance through striping, which nodes' local storage should store certain portions/partitions/strips of the virtual disk to achieve load balancing, etc.). For example, CLOM sub-module 325, in one embodiment, is responsible for generating blueprint 215 describing the RAID 1/RAID 0 configuration for virtual disk object 200 in Figure 2 when the virtual disk was first created by the administrator. As previously discussed, a storage policy may specify requirements for capacity, IOPS, availability, and reliability. Storage policies may also specify a workload characterization (e.g., random or sequential access, I/O request size, cache size, expected cache hit ration, etc.). Additionally, the administrator may also specify an affinity to VSAN module 114 to preferentially use certain nodes 111 (or the local disks housed therein). For example, when provisioning a new virtual disk for a VM, an administrator may generate a storage policy or profile for the virtual disk specifying that the virtual disk have a reserve capacity of 400GB, a reservation of 150 read IOPS, a reservation of 300 write IOPS, and a desired availability of 99.99%. Upon receipt of the generated storage policy, CLOM sub-module 325 consults the in-memory metadata database maintained by its VSAN module 114 to determine the current state of cluster 110 in order generate a virtual disk blueprint for a composite object (e.g., the virtual disk object) that suits the generated storage policy. As further discussed below, CLOM sub-module 325 may then communicate the blueprint to its corresponding distributed object manager (DOM) sub-module 340 which interacts with object space 116 to implement the blueprint by, for example, allocating or otherwise mapping component objects (e.g., stripes) of the composite object to physical storage locations within various nodes 111 of cluster 110 .

[0019] In addition to CLOM sub-module 325 and DOM sub-module 340, as further depicted in Figure 3, VSAN module 114 may also include a cluster monitoring, membership, and directory services (CMMDS) sub-module 335 that maintains the previously discussed in-memory metadata database to provide information on the state of cluster 110 to other sub-modules of VSAN module 114 and also tracks the general “health” of cluster 110 by monitoring the status, accessibility, and visibility of each node 111 in cluster 110. The in-memory metadata database serves as a directory service that maintains a physical inventory of the VSAN environment, such as the various nodes 111, the storage resources in the nodes 111 (SSD, magnetic disks, etc.) housed therein and the characteristics/capabilities thereof, the current state of the nodes 111 and there corresponding storage resources, network paths among the nodes 111, and the like. As previously discussed, in addition to maintaining a physical inventory, the in-memory metadata database further provides a catalog of metadata for objects stored in object store 116 (e.g., what composite and component objects exist, what component objects belong to what composite objects, which nodes serve as “coordinators” or “owners” that control access to which objects, quality of service requirements for each object, object configurations, the mapping of objects to physical storage locations, etc.). As previously discussed, other sub-modules within VSAN module 114 may access CMMDS sub-module 335 (represented by the connecting lines in Figure 3) for updates to learn of changes in cluster topology and object configurations. For example, as previously discussed, during virtual disk creation, CLOM sub-module 325 accesses the in-memory metadata database to generate a virtual disk blueprint, and in order to handle an I/O operation from a running VM 112, DOM sub-module 340 accesses the in-memory metadata database to determine the nodes 111 that store the component objects (e.g., stripes) of a corresponding composite object (e.g., virtual disk object) and the paths by which those nodes are reachable in order to satisfy the I/O operation.

[0020] As previously discussed, DOM sub-module 340, during the handling of I/O operations as well as during object creation, controls access to and handles operations on those component objects in object store 116 that are stored in the local storage of the particular node 111 in which DOM sub-module 340 runs as well as certain other composite objects for which its node 111 has been currently designated

as the “coordinator” or “owner.” For example, when handling an I/O operation from a VM, due to the hierarchical nature of composite objects in certain embodiments, a DOM sub-module 340 that serves as the coordinator for the target composite object (e.g., the virtual disk object that is subject to the I/O operation) may need to further communicate across the network with a different DOM sub-module 340 in a second node 111 (or nodes) that serves as the coordinator for the particular component object (e.g., stripe, etc.) of the virtual disk object that is stored in the local storage of the second node 111 and which is the portion of the virtual disk that is subject to the I/O operation. If the VM issuing the I/O operation resides on a node 111 that is also different from the coordinator of the virtual disk object, the DOM sub-module 340 of the node running the VM would also have to communicate across the network with the DOM sub-module 340 of the coordinator. In certain embodiments, if the VM issuing the I/O operation resides on node that is different from the coordinator of the virtual disk object subject to the I/O operation, the two DOM sub-modules 340 of the two nodes may to communicate to change the role of the coordinator of the virtual disk object to the node running the VM (e.g., thereby reducing the amount of network communication needed to coordinate I/O operations between the node running the VM and the node serving as the coordinator for the virtual disk object).

[0021] DOM sub-modules 340 also similarly communicate amongst one another during object creation. For example, a virtual disk blueprint generated by CLOM module 325 during creation of a virtual disk may include information that designates which nodes 111 should serve as the coordinators for the virtual disk object as well as its corresponding component objects (stripes, etc.). Each of the DOM sub-modules 340 for such designated nodes is issued requests (e.g., by the DOM sub-module 340 designated as the coordinator for the virtual disk object or by the DOM sub-module 340 of the node generating the virtual disk blueprint, etc. depending on embodiments) to create their respective objects, allocate local storage to such objects (if needed), and advertise their objects to their corresponding CMMDS sub-module 335 in order to update the in-memory metadata database with metadata regarding the object. In order to perform such requests, DOM sub-module 340 interacts with a log structured object manager (LSOM) sub-module 350 that serves as the component in VSAN module 114 that actually drives communication with the local SSDs and magnetic disks of its node 111. In addition to allocating local storage for component objects (as well as to

store other metadata such as policies and configurations for composite objects for which its node serves as coordinator, etc.), LSOM sub-module 350 additionally monitors the flow of I/O operations to the local storage of its node 111, for example, to report whether a storage resource is congested.

[0022] Figure 3 also depicts a reliable datagram transport (RDT) sub-module 345 that delivers datagrams of arbitrary size between logical endpoints (e.g., nodes, objects, etc.), where the endpoints may potentially be over multiple paths. In one embodiment, the underlying transport is TCP. Alternatively, other transports such as RDMA may be used. RDT sub-module 345 is used, for example, when DOM sub-modules 340 communicate with one another, as previously discussed above to create objects or to handle I/O operations. In certain embodiments, RDT module 345 interacts with CMMDS module 335 to resolve the address of logical endpoints dynamically in order to maintain up-to-date location information in the in-memory metadata database as well as to create, remove, or reestablish connections based on link health status. For example, if CMMDS module 335 reports a link as unhealthy, RDT sub-module 345 may drop the connection in favor of a link in better condition.

[0023] Figure 4 illustrates a method flow diagram for creating a virtual disk object based on a defined storage policy, according to one embodiment. For example, in step 400, an administrator may interact with a user interface of virtual management platform 105 to create a virtual disk having capacity, availability and IOPS requirements (e.g., the defined storage policy). In one embodiment, virtual management platform 105 may then request a “master” node 111 to create an object for the virtual disk in step 405. In step 410, such a master node 111 may generate a virtual disk blueprint through its CLOM sub-module 325 in VSAN module. As previously discussed, CLOM sub-module 35 generates a virtual disk blueprint for the creation of a virtual disk object (e.g., a composite object) based on the status of cluster 110 as determined by consulting the in-memory metadata database of CMMS sub-module 335. The virtual disk blueprint may identify a particular node that should serve as the coordinator or owner of the virtual disk object. In step 415, the DOM sub-module 340 of the master node 111 may request the DOM sub-module 340 of the identified node to create the virtual disk object. In step 420, the DOM sub-module 340 of the identified node receives the request and creates the virtual disk object, by, for example, communicating with its corresponding LSOM sub-module 350 to

persistently store metadata describing the virtual disk object in its local storage. In step 425, the DOM sub-module 340, based on the virtual disk object blueprint, identifies those other nodes in cluster 110 that have been designated to serve as the coordinator or owner for any component objects in the virtual disk blueprint. The DOM sub-module 340 communicates (e.g., using its RTP sub-module 345) with the DOM sub-modules 340 of the other nodes that will serve as coordinators for the component objects and store the data backing such component objects in their local storage. When such DOM sub-modules 340 receive a request from the DOM sub-module 340 of the coordinator of the virtual disk object to create their respective component objects, they, in turn in step 430, communicate with their respective LSOM modules 350 to allocate local storage for the component object (and its related metadata). Once such component objects have been created, their DOM sub-modules 340 advertise the creation of the components to the in-memory metadata database of its CMMS sub-module 335 in step 435. In step 440, in turn, the DOM sub-module 340 for the coordinator of the virtual disk object also advertises its creation to its CMMS sub-module 335 to update the in-memory metadata database and ultimately transmits an acknowledgement to the administrator (e.g., via the master node communications back to virtual management platform 105).

[0024] Figure 5 illustrates the handling of an I/O operation originating from a VM, according to one embodiment. When a VM running on a particular node performs I/O operations to its virtual disk, the VM's guest operating system, in step 500, transmits an I/O operation request intended for its virtual disk (through a device driver of the guest operating system) which, in step 505, is received by hypervisor 113 and ultimately transmitted and transformed through various layers of an I/O stack in hypervisor 113 to DOM sub-module 340 of VSAN module 114. In step 510, the I/O request received by DOM sub-module 340 includes a unique identifier for an object representing the virtual disk that DOM sub-module 340 uses to identify the coordinator node of the virtual disk object by accessing the in-memory metadata database of CMMS sub-module 335 (in certain embodiments, accessing the in-memory metadata database to look up a mapping of the identity of the coordinator node to the unique identifier occurs only when the virtual disk object is initially accessed, with such mapping persisting for future I/O operations such that subsequent lookups are not needed). Upon identifying the coordinator node for the virtual disk

object, the DOM sub-module 340 of the node running the VM communicates (e.g., using its RTP sub-module 345) with the DOM sub-module 340 of the coordinator node to request that it perform the I/O operation in step 515. As previously discussed, in certain embodiment, if the node running the VM and the node serving as coordinator of the virtual disk object are different, the two DOM sub-modules will communicate to update the role of the coordinator of the virtual disk object to be the node of the running VM. Upon the coordinator's receipt of the I/O request, in step 520, its DOM sub-module identifies (e.g., by again referencing the in-memory metadata database, in certain embodiments) those coordinator nodes for the particular component objects (e.g., stripes) of the virtual disk object that are subject to the I/O operation. For example, if the I/O operation spans multiple stripes (e.g., multiple component objects) of a RAID 0 configuration, DOM sub-module 340 may split the I/O operation and appropriately transmit correspond I/O requests to the respective coordinate nodes for the relevant component objects that correspond to the two stripes. In step 525, the DOM sub-module of the coordinator node for the virtual disk object requests that the DOM sub-modules for the coordinator nodes of the identified component objects perform the I/O operation request and, in step 530, the DOM sub-modules of such coordinator nodes for the identified component objects interact with their corresponding LSOM sub-modules to perform the I/O operation in the local storage resource where the component object is stored.

[0025] Although one or more embodiments have been described in some detail for clarity of understanding, it will be apparent that certain changes and modifications may be made within the scope of the claims. Accordingly, the described embodiments are to be considered as illustrative and not restrictive, and the scope of the claims is not to be limited to details given herein, but may be modified within the scope and equivalents of the claims. For example, although a number of foregoing described embodiments describe virtual machines as the clients that access the virtual disks provided by the VSAN module, it should be recognized that any clients, such as a cluster of non-virtualized host servers and/or non-virtualized applications running therein may similarly utilize the VSAN module in alternative embodiment. Similarly, alternative embodiments of the VSAN module may enable creation of high level storage objects other than virtual disks, such as, without limitation, REST objects, files, file systems, blob (binary large objects) and other objects. Similarly, while

VSAN module 114 has been generally depicted as embedded in hypervisor 113, alternative embodiments may implement VSAN module separate from hypervisor 113, for example as a special virtual machine or virtual appliance, a separate application or any other “pluggable” module or driver that can be inserted into computing platform in order to provide and manage a distributed object store. Similarly, while the foregoing embodiments have referred to RAID configurations as one technique to organize a blueprint, it should be recognized that other configurations may be utilized in other embodiments, including, without limitation, using erasure codes and other similar techniques. While descriptions herein have discussed using “unique identifiers” to reference objects in the objects, it should be recognized that techniques to generate unique identifiers (hashes, etc.) may not necessarily be guaranteed to generate truly unique identifiers and therefore certain embodiments may further implement techniques for handling name collisions in case there are instances where identifiers for objects are not truly unique. In one such embodiment, in addition to the purported unique identifier, an additional administrator or user specified identifier is also assigned to an object during its creation (or otherwise mapped to the object’s unique identifier). When there is a name collision due to a pre-existing object having the same unique identifier as a newly created object, the administrator or user-specified name can be transformed into the newly created object’s user identifier.

[0026] The various embodiments described herein may employ various computer-implemented operations involving data stored in computer systems. For example, these operations may require physical manipulation of physical quantities usually, though not necessarily, these quantities may take the form of electrical or magnetic signals where they, or representations of them, are capable of being stored, transferred, combined, compared, or otherwise manipulated. Further, such manipulations are often referred to in terms, such as producing, identifying, determining, or comparing. Any operations described herein that form part of one or more embodiments may be useful machine operations. In addition, one or more embodiments also relate to a device or an apparatus for performing these operations. The apparatus may be specially constructed for specific required purposes, or it may be a general purpose computer selectively activated or configured by a computer program stored in the computer. In particular, various general purpose machines may

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.

[0027] The various embodiments described herein may be practiced with other computer system configurations including hand-held devices, microprocessor systems, microprocessor-based or programmable consumer electronics, minicomputers, mainframe computers, and the like.

[0028] One or more embodiments may be implemented as one or more computer programs or as one or more computer program modules embodied in one or more computer readable media. The term computer readable medium refers to any data storage device that can store data which can thereafter be input to a computer system. Computer readable media may be based on any existing or subsequently developed technology for embodying computer programs in a manner that enables them to be read by a computer. Examples of a computer readable medium include a hard drive, network attached storage (NAS), read-only memory, random-access memory (e.g., a flash memory device), a CD (Compact Discs), CD-ROM, a CD-R, or a CD-RW, a DVD (Digital Versatile Disc), a magnetic tape, 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.

[0029] In addition, while described virtualization methods have generally assumed that virtual machines present interfaces consistent with a particular hardware system, the methods described may be used in conjunction with virtualizations that do not correspond directly to any particular hardware system. Virtualization systems in accordance with the various embodiments, implemented as hosted embodiments, non-hosted embodiments, or as embodiments that tend to blur distinctions between the two, are all envisioned. Furthermore, various virtualization operations may be wholly or partially implemented in hardware. For example, a hardware implementation may employ a look-up table for modification of storage access requests to secure non-disk data.

[0030] Many variations, modifications, additions, and improvements are possible, regardless the degree of virtualization. The virtualization software can therefore

include components of a host, console, or guest operating system that performs virtualization functions. Plural instances may be provided for components, operations or structures described herein as a single instance. Finally, boundaries between various components, operations and data stores are somewhat arbitrary, and particular operations are illustrated in the context of specific illustrative configurations. Other allocations of functionality are envisioned and may fall within the scope of one or more embodiments. In general, structures and functionality presented as separate components in exemplary configurations may be implemented as a combined structure or component. Similarly, structures and functionality presented as a single component may be implemented as separate components. These and other variations, modifications, additions, and improvements may fall within the scope of the appended claims(s). In the claims, elements and/or steps do not imply any particular order of operation, unless explicitly stated in the claims.

CLAIMS:

1. A method for providing a file system interface for an object store intended to support simultaneous access to objects stored in the object store by multiple clients, the steps comprising:

exposing to clients an abstraction of a root directory to a hierarchical namespace for the object store, wherein the object store (i) is backed by a plurality of physical storage devices housed in or directly attached to the plurality of host computers, and (ii) internally tracks its stored objects using a flat namespace that maps unique identifiers to the stored objects; and

enabling creation of top-level objects appearing as subdirectories of the root directory, wherein each top-level object represents a separate abstraction of a storage device having a separate namespace that can be organized in accordance with any designated file system.

2. The method of claim 1, wherein at least one top-level object is contains a clustered file system designed to support simultaneous access by clients up to a maximum limit.

3. The method of claim 2, wherein the top-level object is auto-mounted as a subdirectory at a client's local version of the abstraction of the root directory when the client attempts to access a file stored within the clustered file system.

4. The method of claim 2, wherein the top-level object is configured to store metadata for file objects representing files stored in the clustered file system, wherein the file objects are stored in the object store separate from the top-level object.

5. The method of claim 4, wherein a file object is configured to store unique identifiers usable by the flat namespace to find locations of corresponding component objects that represent separate portions of the files and are stored in the object store separate from the file object.

6. The method of claim 4, wherein the metadata stored in the top-level object for a file object includes at least a file descriptor that comprises a unique identifier usable by the flat namespace to find a location of the file object in the object store.
7. The method of claim 6, wherein the file represented by the file object is a virtual disk file used by a virtual machine running on one of the plurality host computers.
8. The method of claim 7, wherein the top-level object is configured to store metadata for file objects that are only utilized by one of a plurality of virtual machines capable of being instantiated on any of the plurality of host computers.
9. The method of claim 3, wherein the top-level objects is automatically unmounted after a period of time in which no files stored within the clustered file system are accessed by the client.
10. The method of claim 1, wherein the flat namespace is implemented as an in-memory database, a synchronized copy of which is stored in the memory of each of the host computers.
11. The method of claim 1, wherein a pathname in the designated file system of a top level object can be mapped to an object configured to represent a separate abstraction of a storage device having a separate namespace that can be organized in accordance with any designated file system and can be auto-mounted at a client when the client attempts to access a file stored within the designated file system.
12. The method of claim 1, wherein an administrator specified name for an object is transformed into a unique identifier for the object when the flat namespace reports a conflict with a auto-generated unique identifier for the object.
13. A non-transitory computer readable storage medium storing instructions, which, when executed on a processor, performs an operation for providing a file

system interface for an object store intended to support simultaneous access to objects stored in the object store by multiple clients, the operation comprising:

exposing an abstraction of a storage device to a plurality of host computers in a cluster, wherein the storage device is backed by an object store comprising, in aggregate, a plurality of physical storage devices housed in the plurality of host computers; and

enabling creation of multiple top-level file system objects in the object store that each represent separate file systems, wherein each top-level file system object is configured to (i) respond according to the semantics of its corresponding file system, and (ii) store metadata for file objects representing files stored in the corresponding file system, wherein the file objects are stored in the object store separate from their corresponding top-level file system objects.

14. The computer readable storage medium of claim 13, wherein at least one top-level object contains a clustered file system designed to support simultaneous access by clients up to a maximum limit.

15. The computer readable storage medium of claim 14, wherein the top-level object is auto-mounted as a subdirectory at a client's local version of the abstraction of the root directory when the client attempts to access a file stored within the clustered file system.

16. The computer readable storage medium of claim 14, wherein the top-level object is configured to store metadata for file objects representing files stored in the clustered file system, wherein the file objects are stored in the object store separate from the top-level object.

17. The computer readable storage medium of claim 16, wherein a file object is configured to store unique identifiers usable by the flat namespace to find locations of corresponding component objects that represent separate portions of the files and are stored in the object store separate from the file object.

18. The computer readable storage medium of claim 16, wherein the metadata stored in the top-level object for a file object includes at least a file descriptor that comprises a unique identifier usable by the flat namespace to find a location of the file object in the object store.

19. The computer readable storage medium of claim 19, wherein the file represented by the file object is a virtual disk file used by a virtual machine running on one of the plurality host computers.

20. The computer readable storage medium of claim 19, wherein the top-level object is configured to store metadata for file objects that are only utilized by one of a plurality of virtual machines capable of being instantiated on any of the plurality of host computers.

21. The computer readable storage medium of claim 13, wherein the flat namespace is implemented as an in-memory database, a synchronized copy of which is stored in the memory of each of the host computers.

22. The computer readable storage medium of claim 13, wherein a pathname in the designated file system of a top level object can be mapped to an object configured to represent a separate abstraction of a storage device having a separate namespace that can be organized in accordance with any designated file system and can be auto-mounted at a client when the client attempts to access a file stored within the designated file system.

23. A computer system, comprising:

a processor and

a memory hosting a module, which, when executed on the processor, performs an operation for providing a file system interface for an object store intended to support simultaneous access to objects stored in the object store by multiple clients, the operation comprising:

exposing an abstraction of a storage device to a plurality of host computers in a cluster, wherein the storage device is backed by an object store comprising, in

aggregate, a plurality of physical storage devices housed in the plurality of host computers; and

enabling creation of multiple top-level file system objects in the object store that each represent separate file systems, wherein each top-level file system object is configured to (i) respond according to the semantics of its corresponding file system, and (ii) store metadata for file objects representing files stored in the corresponding file system, wherein the file objects are stored in the object store separate from their corresponding top-level file system objects.

24. The computer system of claim 23, wherein at least one top-level object contains a clustered file system designed to support simultaneous access by clients up to a maximum limit.

25. The computer system of claim 24, wherein the top-level object is auto-mounted as a subdirectory at a client's local version of the abstraction of the root directory when the client attempts to access a file stored within the clustered file system.

26. The computer system of claim 24, wherein the top-level object is configured to store metadata for file objects representing files stored in the clustered file system, wherein the file objects are stored in the object store separate from the top-level object.

27. The computer system of claim 26, wherein a file object is configured to store unique identifiers usable by the flat namespace to find locations of corresponding component objects that represent separate portions of the files and are stored in the object store separate from the file object.

28. The computer system of claim 26, wherein the metadata stored in the top-level object for a file object includes at least a file descriptor that comprises a unique identifier usable by the flat namespace to find a location of the file object in the object store.

29. The computer system of claim 28, wherein the file represented by the file object is a virtual disk file used by a virtual machine running on one of the plurality of host computers.

30. The computer system of claim 29, wherein the top-level object is configured to store metadata for file objects that are only utilized by one of a plurality of virtual machines capable of being instantiated on any of the plurality of host computers.

31. The computer system of claim 23, wherein the flat namespace is implemented as an in-memory database, a synchronized copy of which is stored in the memory of each of the host computers.

32. The computer system of claim 23, wherein a pathname in the designated file system of a top level object can be mapped to an object configured to represent a separate abstraction of a storage device having a separate namespace that can be organized in accordance with any designated file system and can be auto-mounted at a client when the client attempts to access a file stored within the designated file system.




FIGURE 1

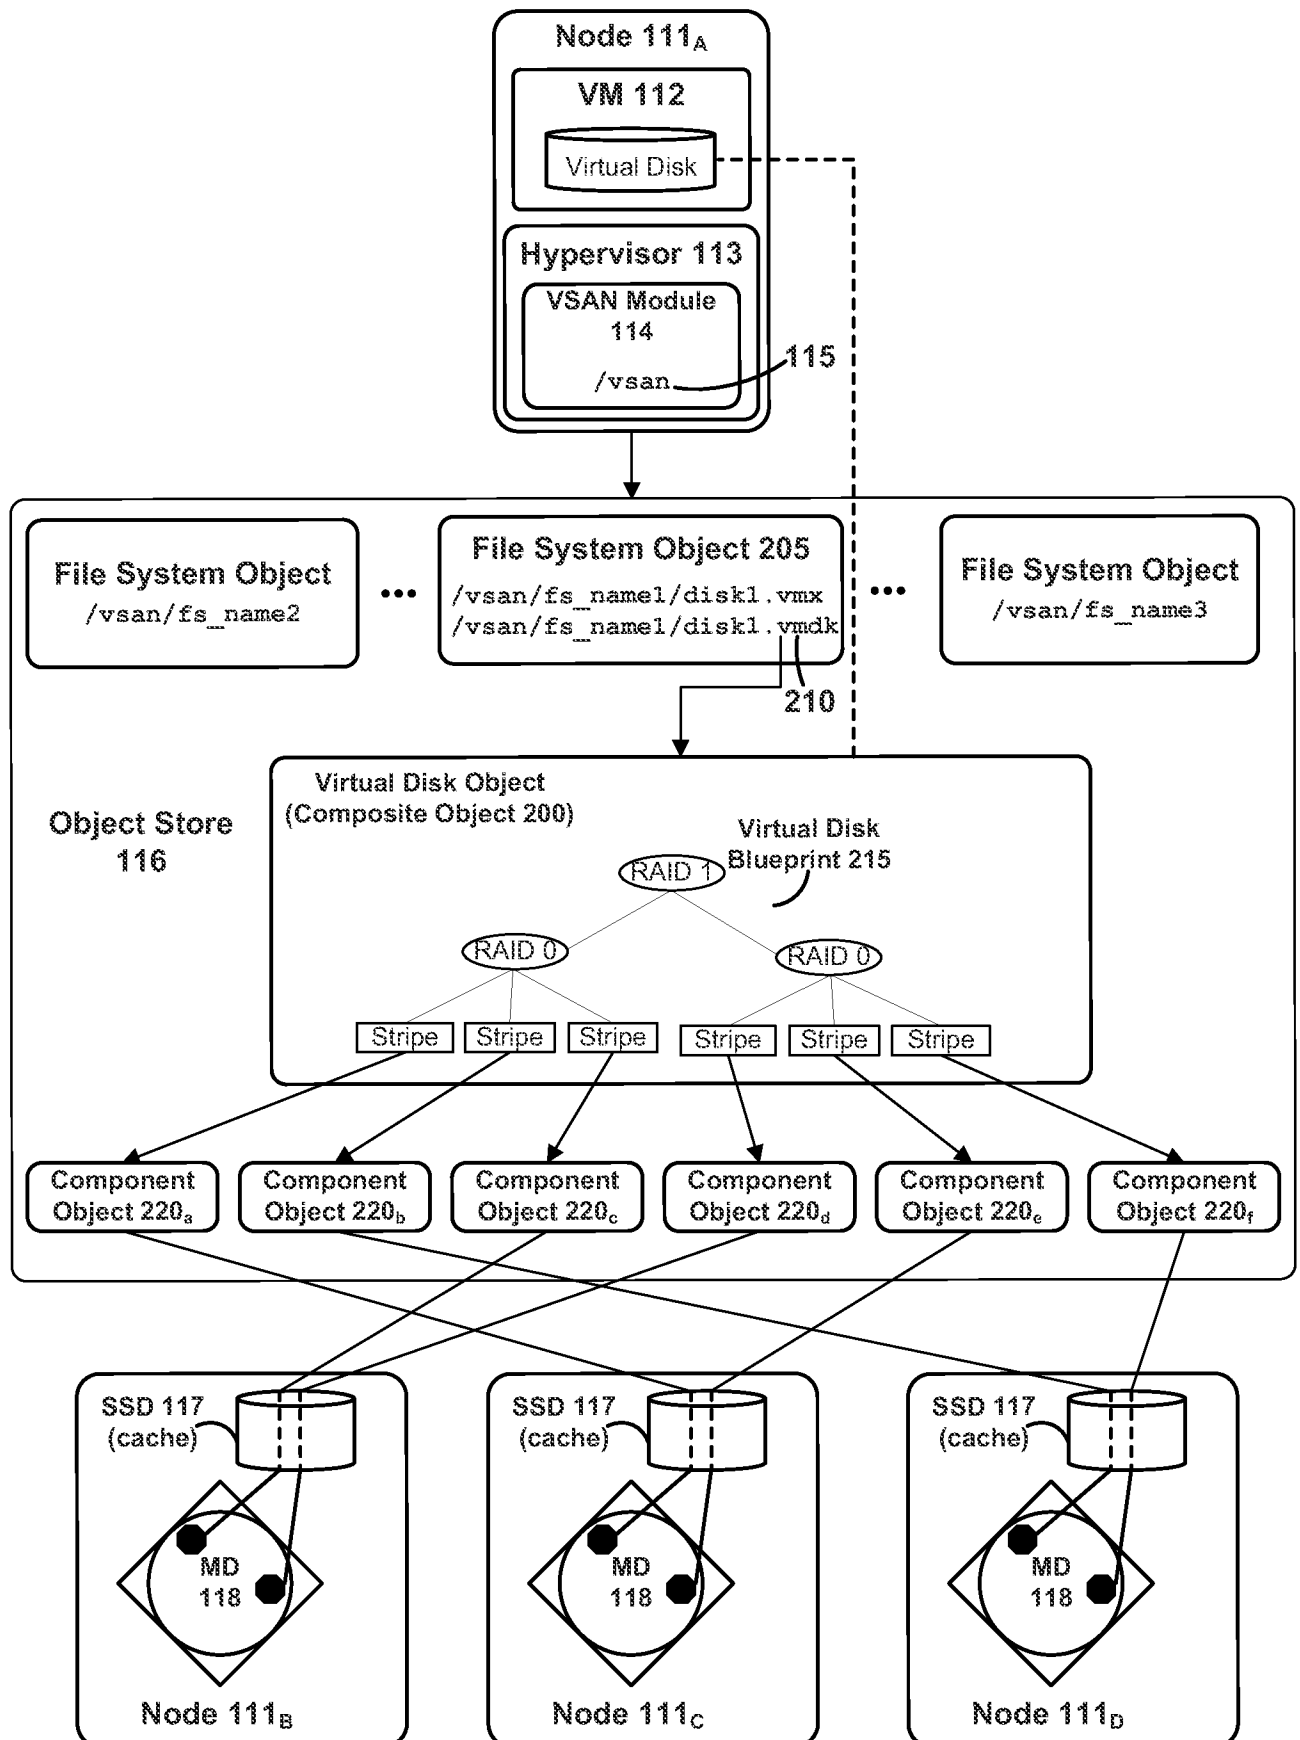


FIGURE 2

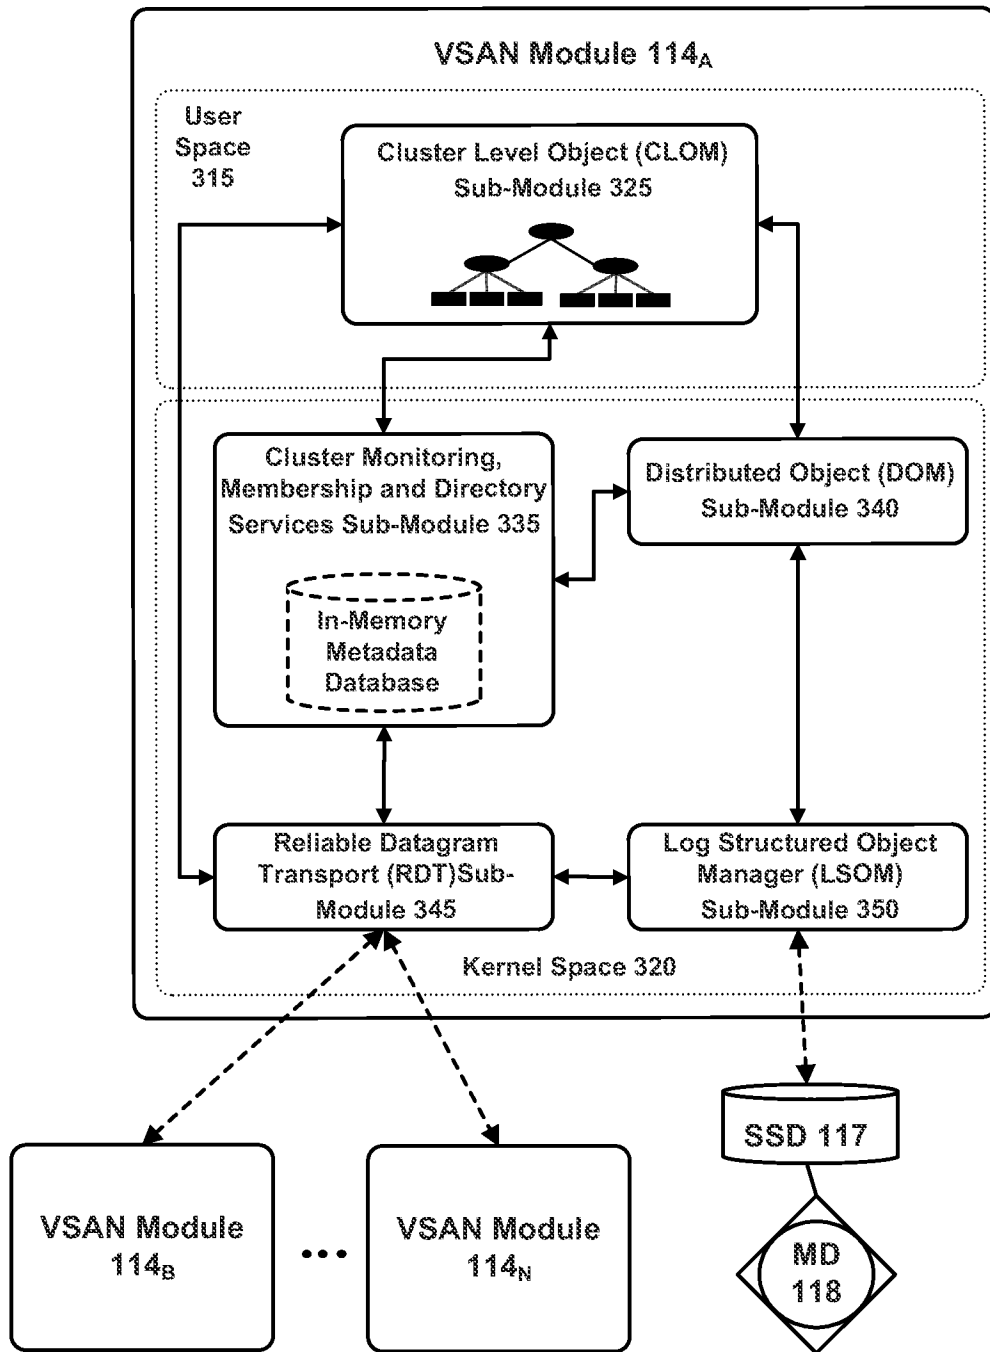


FIGURE 3

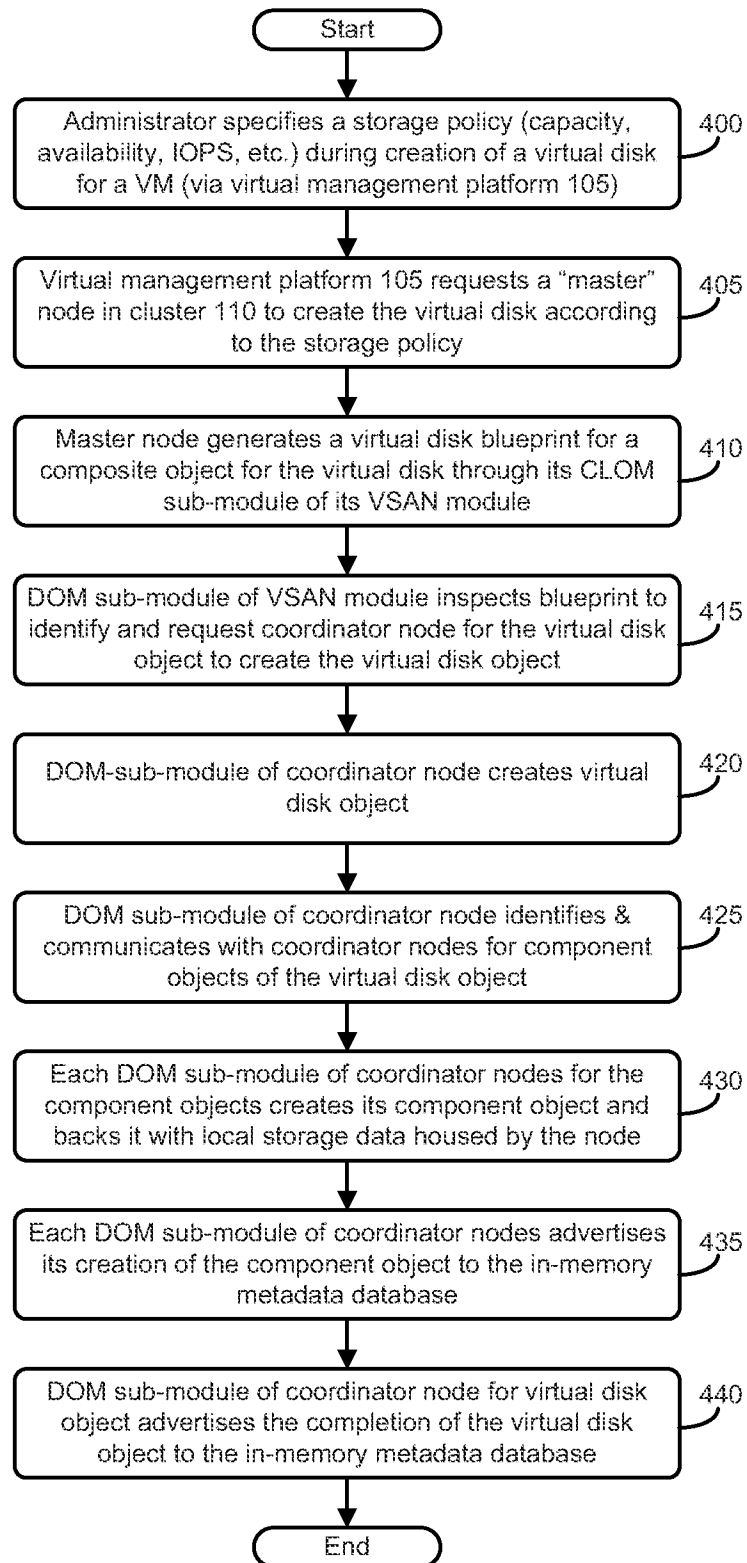


FIGURE 4

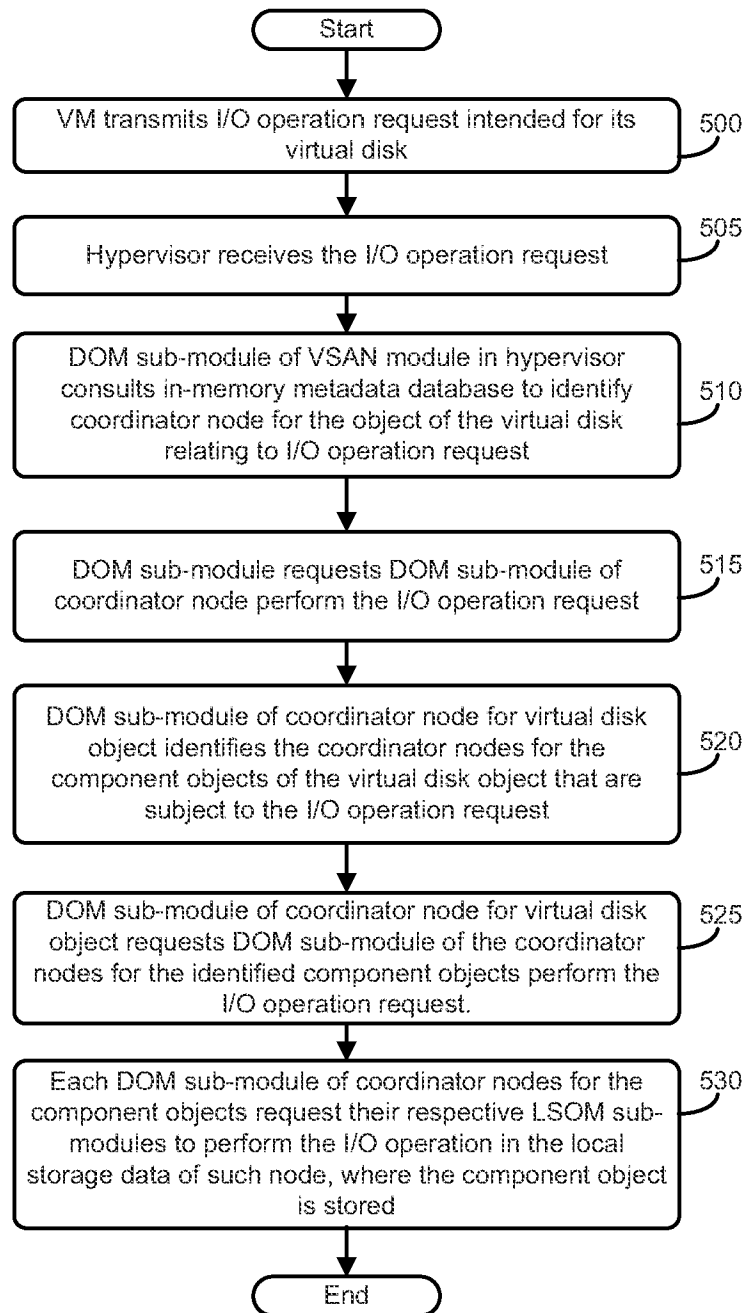


FIGURE 5

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2014/041761

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F17/30

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 practicable, search terms used)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2003/158836 A1 (VENKATESH DINESH [US] ET AL) 21 August 2003 (2003-08-21) paragraph [0009] paragraph [0031] paragraph [0033] paragraph [0035] paragraph [0054] -----	1-32
A	US 2009/271412 A1 (LACAPRA FRANCESCO [US] ET AL) 29 October 2009 (2009-10-29) the whole document -----	1-32
A	US 2005/097073 A1 (MAIR DAVID A [US] ET AL) 5 May 2005 (2005-05-05) the whole document ----- -/-	1-32



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

13 October 2014

Date of mailing of the international search report

17/10/2014

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

de Castro Palomares

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2014/041761

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2005/240714 A1 (MCCAULEY RODNEY [US] ET AL) 27 October 2005 (2005-10-27) the whole document -----	1-32
A	US 2011/238814 A1 (PITTS WILLIAM [US]) 29 September 2011 (2011-09-29) the whole document -----	1-32

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2014/041761

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2003158836	A1	21-08-2003	NONE
US 2009271412	A1	29-10-2009	TW 201007489 A 16-02-2010
		US 2009271412 A1	29-10-2009
		US 2013013619 A1	10-01-2013
		US 2013013639 A1	10-01-2013
		US 2013013654 A1	10-01-2013
		US 2013013655 A1	10-01-2013
		US 2013013675 A1	10-01-2013
		US 2013018928 A1	17-01-2013
		US 2013018930 A1	17-01-2013
		US 2013066830 A1	14-03-2013
		US 2013066931 A1	14-03-2013
		WO 2009134772 A2	05-11-2009
US 2005097073	A1	05-05-2005	NONE
US 2005240714	A1	27-10-2005	NONE
US 2011238814	A1	29-09-2011	NONE