




- (51) **International Patent Classification:**
G06F 15/16 (2006.01) *H04L 12/28* (2006.01)
- (21) **International Application Number:**
PCT/US2012/031109
- (22) **International Filing Date:**
29 March 2012 (29.03.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (71) **Applicant (for all designated States except US):** HITACHI DATA SYSTEMS CORPORATION [US/US]; 750 Central Expressway, Santa Clara, CA 95050 (US).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** PANNESE, Donald, P. [US/US]; 7 Dell Rd, Westford, MA 10886 (US). ISHERWOOD, Benjamin [US/US]; 298 River Rd, Tewksbury, MA 01878 (US).
- (74) **Agents:** LEUNG, Chun-Pok et al.; c/o Mattingly & Malur, PC, 1800 Diagonal Road, Suite 370, Alexandria, VA 22314 (US).

- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

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

(54) **Title:** DNS ALIAS SYNCHRONIZATION IN REPLICATION TOPOLOGY**FIG. 6**

(57) **Abstract:** A redundant array of independent nodes are networked together. Each node executes an instance of an application that provides object-based storage. The nodes are grouped into a plurality of systems each having multiple nodes. The systems have one or more replication links each being formed to indicate replication of data from one system to another system in a replicated environment where each system is configured as a sub-domain in a Domain Name System (DNS) infrastructure. A DNS alias synchronization method comprises maintaining updated information, within each system, of all replication links involving that system and of DNS aliases of other systems associated with all replication links involving that system (S602-S607). This enables that system to process network-based requests, on behalf of the other systems, without redirecting the requests from the other systems to that system.

DNS ALIAS SYNCHRONIZATION IN REPLICATION TOPOLOGY

BACKGROUND OF THE INVENTION

[0001] The present invention relates generally to storage systems and, more particularly, to system and method for DNS (Domain Name System) alias synchronization in replication topology.

[0002] In a replicated environment of replicated object storage system or content platform, in the event where an object is not available on the primary system, the client application has traditionally had to do the work of redirecting to the replica to access the object. As a result, the client application has the burden of managing additional work. The following are some of the reasons an object is not available on the primary system: (1) the object in the content platform Storage Manager is corrupted; (2) the content platform Storage Manager on which the object exists is down; and (3) the node on which the object exists fails (which in turn caused the Storage Manager on which object existed to go down). For example, if a failover scenario were to occur in a replication topology, the customer has to manually intervene to redirect the client applications from the system that has failed to the system which is the target of the failover.

BRIEF SUMMARY OF THE INVENTION

[0003] Exemplary embodiments of the invention provide a replicated object storage system or content platform (ecosystem) with DNS alias synchronization. In an ecosystem that consists of two or more systems in a replicated environment where each system is configured as a sub-domain in

the Domain Names System (DNS) infrastructure, it is necessary for every participating system to know what the DNS aliases of the other participating systems are in a manner that does not require every participating system to communicate directly with all other systems in the ecosystem. The knowledge of all the DNS aliases in the ecosystem is required to enable any system in the ecosystem to be able to process network-based requests on behalf of another participating system and limiting requests to only those systems in the ecosystem. In this way, individual systems are made aware of all other participating systems in the replication topology. As a result, requests will be serviced without the need for a customer to manually redirect the client applications.

[0004] One solution is to implement an "eventually consistent" method of DNS alias synchronization in replication environments having simple to complex topologies derived from the core chain and many-to-one topologies. When a system joins/leaves the environment by establishing/removing a replication link or a participating system has a DNS alias update, it will notify all directly linked systems (i.e., peers) of its own DNS alias. The peer handling of the notification will include: (1) merging the notifying system's alias with its own known aliases and identifying any possible differences, (2) responding back to the notifying system with all its known valid aliases, and (3) if the peer actually made any change(s) to its known aliases, notifying all its other peers of the change(s) by that peer. Eventually, any change(s) will be propagated throughout the environment to all systems and the systems will no longer have any change(s) to report.

[0005] An aspect of the present invention is related to a redundant array of independent nodes, the nodes being grouped into a plurality of systems each having multiple nodes. The plurality of systems have one or more replication links each being formed to indicate replication of data from one system to another system in a replicated environment where each system is configured as a sub-domain in a Domain Name System (DNS) infrastructure. Each system has a DNS name that is also a DNS alias of said each system to other systems. A DNS alias synchronization method comprises maintaining updated information, within each system of the plurality of systems, of all replication links involving said each system and of DNS aliases of other systems associated with said all replication links involving said each system. The updated information of all replication links involving a system and of DNS aliases of other systems associated with all replication links involving the system enables the system to process network-based requests, on behalf of the other systems associated with all replication links involving the system, without redirecting the requests from the other systems to the system.

[0006] In some embodiments, the method further comprises providing the updated information, from each system, to a DNS server which enables said each system to process network-based requests, on behalf of the other systems associated with all replication links involving said each system, without redirecting the requests from the other systems to said each system. Maintaining updated information comprises updating a known alias list within said each system, the known alias list including mapping of the replication links between the DNS names of the systems associated with all replication

links involving said each system, and any replication links that have recently been deleted from said all replication links involving said each system. For a system to initiate updating the known alias list as an initiating system, the updating comprises: sending a synchronization request to all systems that are each connected to the initiating system by a direct replication link, the synchronization request including a local known alias list of the initiating system and the DNS alias of the initiating system; receiving a response from each system that received the synchronization request from the initiating system; generating a merged known alias list for the initiating system from all the responses received by the initiating system; and storing the merged known alias list generated for the initiating system as the local known alias list of the initiating system.

[0007] In specific embodiments, for a system that receives, as a receiving system, a synchronization request from another system with a known alias list of said another system, the updating comprises: merging the known alias list from the synchronization request with a local known alias list of the receiving system to produce a merged known alias list for the receiving system; storing the merged known alias list for the receiving system as the local known alias list of the receiving system; and returning a response to said another system that sent the synchronization request, the response including the merged known alias list for the receiving system. If the local known alias list of the receiving system has changed based on the synchronization request from said another system, the updating further comprises: propagating the synchronization request to all systems that are each connected to the receiving system by a direct replication link excluding said

another system, the synchronization request including the local known alias list of the receiving system and the DNS alias of the receiving system; receiving a response from each system that received the propagated synchronization request from the receiving system; generating a new merged known alias list for the receiving system from all the responses received by the receiving system; and storing the new merged known alias list as the local known alias list of the receiving system. If the new merged known alias list is different from the local known alias list of the receiving system prior to generating the merged known alias list for the receiving system, the updating further comprises: sending a second synchronization request, with the receiving system as a new initiating system, to all systems that are each connected to the new initiating system by a direct replication link, the second synchronization request including the local known alias list of the new initiating system and the DNS alias of the new initiating system; receiving a response from each system that received the second synchronization request from the new initiating system; generating a merged known alias list for the new initiating system from all the responses received by the new initiating system; and storing the merged known alias list for the new initiating system as the local known alias list of the new initiating system.

[0008] Another aspect of the invention is directed to an apparatus for Domain Name System (DNS) alias synchronization in a redundant array of independent nodes, the nodes being grouped into a plurality of systems each having multiple nodes. The plurality of systems have one or more replication links each being formed to indicate replication of data from one system to another system in a replicated environment where each system is configured

as a sub-domain in a DNS infrastructure. Each system has a DNS name that is also a DNS alias of said each system to other systems. The apparatus comprises a processor, a memory, and a DNS alias synchronization module provided for each system. The DNS alias synchronization module is configured to maintain updated information, within the system in which the DNS alias synchronization module is provided, of all replication links involving the system and of DNS aliases of other systems associated with said all replication links involving the system. The updated information of all replication links involving a system and of DNS aliases of other systems associated with all replication links involving the system enables the system to process network-based requests, on behalf of the other systems associated with all replication links involving the system, without redirecting the requests from the other systems to the system.

[0009] In some embodiments, the DNS alias synchronization module is configured to provide the updated information, from the system in which the DNS alias synchronization module is provided, to a DNS server which enables the system to process network-based requests, on behalf of the other systems associated with all replication links involving the system, without redirecting the requests from the other systems to the system.

[0010] Another aspect of this invention is directed to a computer-readable storage medium storing a plurality of instructions for controlling a data processor to perform Domain Name System (DNS) alias synchronization, in a redundant array of independent nodes, the nodes being grouped into a plurality of systems each having multiple nodes. The plurality of systems have one or more replication links each being formed to indicate

replication of data from one system to another system in a replicated environment where each system is configured as a sub-domain in a DNS infrastructure. Each system has a DNS name that is also a DNS alias of said each system to other systems. The computer-readable storage medium is provided in each system. The plurality of instructions comprise instructions that cause the data processor to maintain updated information, within the system in which the computer-readable storage medium is provided, of all replication links involving the system and of DNS aliases of other systems associated with said all replication links involving the system. The updated information of all replication links involving a system and of DNS aliases of other systems associated with all replication links involving the system enables the system to process network-based requests, on behalf of the other systems associated with all replication links involving the system, without redirecting the requests from the other systems to the system.

[0011] These and other features and advantages of the present invention will become apparent to those of ordinary skill in the art in view of the following detailed description of the specific embodiments.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] FIG. 1 is a simplified block diagram of a fixed content storage archive in which the method and apparatus of the invention may be applied.

[0013] FIG. 2 is a simplified representation of a redundant array of independent nodes each of which is symmetric and supports an archive cluster application.

[0014] FIG. 3 is a high level representation of the various components of the archive cluster application executing on a given node.

[0015] FIG. 4 shows an example of many-to-one replication with disaster recovery.

[0016] FIG. 5 shows the steady state of the replicated system of FIG. 4 after achieving “eventually consistent” DNS alias synchronization.

[0017] FIG. 6 shows an example of DNS alias synchronization workflow in a replicated system.

[0018] FIG. 7 is an example of a flow diagram illustrating DNS alias synchronization for an initiating system.

[0019] FIG. 8 is an example of a flow diagram illustrating DNS alias synchronization for a system that receives a synchronization request from a peer system.

[0020] FIG. 9 shows an example of system addition in a replication link chain.

[0021] FIG. 10 is a table illustrating a process of DNS alias synchronization for the system addition of FIG. 9.

[0022] FIG. 11 shows an example of system addition in a many-to-one replicated system with disaster recovery.

[0023] FIG. 12 is a table illustrating a process of DNS alias synchronization for the system addition of FIG. 11.

[0024] FIG. 13 shows an example of system deletion in a many-to-one replicated system with disaster recovery.

[0025] FIG. 14 is a table illustrating a process of DNS alias synchronization for the system deletion of FIG. 13.

DETAILED DESCRIPTION OF THE INVENTION

[0026] In the following detailed description of the invention, reference is made to the accompanying drawings which form a part of the disclosure, and in which are shown by way of illustration, and not of limitation, exemplary embodiments by which the invention may be practiced. In the drawings, like numerals describe substantially similar components throughout the several views. Further, it should be noted that while the detailed description provides various exemplary embodiments, as described below and as illustrated in the drawings, the present invention is not limited to the embodiments described and illustrated herein, but can extend to other embodiments, as would be known or as would become known to those skilled in the art. Reference in the specification to “one embodiment,” “this embodiment,” or “these embodiments” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the invention, and the appearances of these phrases in various places in the specification are not necessarily all referring to the same embodiment. Additionally, in the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. However, it will be apparent to one of ordinary skill in the art that these specific details may not all be needed to practice the present invention. In other circumstances, well-known structures, materials, circuits, processes and interfaces have not been described in detail, and/or may be illustrated in block diagram form, so as to not unnecessarily obscure the present invention.

[0027] Furthermore, some portions of the detailed description that follow are presented in terms of algorithms and symbolic representations of operations within a computer. These algorithmic descriptions and symbolic representations are the means used by those skilled in the data processing arts to most effectively convey the essence of their innovations to others skilled in the art. An algorithm is a series of defined steps leading to a desired end state or result. In the present invention, the steps carried out require physical manipulations of tangible quantities for achieving a tangible result. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals or instructions capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, instructions, or the like. It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise, as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as “processing,” “computing,” “calculating,” “determining,” “displaying,” or the like, can include the actions and processes of a computer system or other information processing device that manipulates and transforms data represented as physical (electronic) quantities within the computer system’s registers and memories into other data similarly represented as physical quantities within the computer system’s memories or registers or other information storage, transmission or display devices.

[0028] The present invention also relates to an apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may include one or more general-purpose computers selectively activated or reconfigured by one or more computer programs. Such computer programs may be stored in a computer-readable storage medium, such as, but not limited to optical disks, magnetic disks, read-only memories, random access memories, solid state devices and drives, or any other types of media suitable for storing electronic information. The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general-purpose systems may be used with programs and modules in accordance with the teachings herein, or it may prove convenient to construct a more specialized apparatus to perform desired method steps. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein. The instructions of the programming language(s) may be executed by one or more processing devices, e.g., central processing units (CPUs), processors, or controllers.

[0029] Exemplary embodiments of the invention, as will be described in greater detail below, provide apparatuses, methods and computer programs for transparent recovery of damaged or unavailable objects in a replicated object storage system.

[0030] I. Fixed Content Distributed Data Storage

[0031] A need has developed for the archival storage of "fixed content" in a highly available, reliable and persistent manner that replaces or supplements traditional tape and optical storage solutions. The term "fixed content" typically refers to any type of digital information that is expected to be retained without change for reference or other purposes. Examples of such fixed content include, among many others, e-mail, documents, diagnostic images, check images, voice recordings, film and video, and the like. The traditional Redundant Array of Independent Nodes (RAIN) storage approach has emerged as the architecture of choice for creating large online archives for the storage of such fixed content information assets. By allowing nodes to join and exit from a cluster as needed, RAIN architectures insulate a storage cluster from the failure of one or more nodes. By replicating data on multiple nodes, RAIN-type archives can automatically compensate for node failure or removal. Typically, RAIN systems are largely delivered as hardware appliances designed from identical components within a closed system.

[0032] FIG. 1 illustrates one such scalable disk-based archival storage management system. The nodes may comprise different hardware and thus may be considered "heterogeneous." A node typically has access to one or more storage disks, which may be actual physical storage disks, or virtual storage disks, as in a storage area network (SAN). The archive cluster application (and, optionally, the underlying operating system on which that application executes) that is supported on each node may be the same or substantially the same. The software stack (which may include the operating system) on each node is symmetric, whereas the hardware may be heterogeneous. Using the system, as illustrated in FIG. 1, enterprises can

create permanent storage for many different types of fixed content information such as documents, e-mail, satellite images, diagnostic images, check images, voice recordings, video, and the like, among others. These types are merely illustrative, of course. High levels of reliability are achieved by replicating data on independent servers, or so-called storage nodes. Preferably, each node is symmetric with its peers. Thus, because preferably any given node can perform all functions, the failure of any one node has little impact on the archive's availability.

[0033] As described in commonly owned U.S. Patent No. 7,155,466, it is known in a RAIN-based archival system to incorporate a distributed software application executed on each node that captures, preserves, manages, and retrieves digital assets. FIG. 2 illustrates one such system. A physical boundary of an individual archive is referred to as a cluster (or a system). Typically, a cluster is not a single device, but rather a collection of devices. Devices may be homogeneous or heterogeneous. A typical device is a computer or machine running an operating system such as Linux. Clusters of Linux-based systems hosted on commodity hardware provide an archive that can be scaled from a few storage node servers to many nodes that store thousands of terabytes of data. This architecture ensures that storage capacity can always keep pace with an organization's increasing archive requirements.

[0034] In storage systems such as described above, data typically is distributed across the cluster randomly so that the archive is always protected from device failure. If a disk or node fails, the cluster automatically fails over to other nodes in the cluster that maintain replicas of the same data. While

this approach works well from a data protection standpoint, a calculated mean time to data loss (MTDL) for the cluster may not be as high as desired. In particular, MTDL typically represents a calculated amount of time before the archive will lose data. In a digital archive, any data loss is undesirable, but due to the nature of hardware and software components, there is always a possibility (however remote) of such an occurrence. Because of the random distribution of objects and their copies within an archive cluster, MTDL may end up being lower than required since, for example, a needed copy of an object may be unavailable if a given disk (on which a mirror copy is stored) within a given node fails unexpectedly.

[0035] As shown in FIG. 2, an illustrative cluster in which the present invention is implemented preferably comprises the following general categories of components: nodes 202, a pair of network switches 204, power distribution units (PDUs) 206, and uninterruptible power supplies (UPSs) 208. A node 202 typically comprises one or more commodity servers and contains a CPU (e.g., Intel x86, suitable random access memory (RAM), one or more hard drives (e.g., standard IDE/SATA, SCSI, or the like), and two or more network interface (NIC) cards. A typical node is a 2U rack mounted unit with a 2.4 GHz chip, 512 MB RAM, and six (6) 200 GB hard drives. This is not a limitation, however. The network switches 204 typically comprise an internal switch 205 that enables peer-to-peer communication between nodes, and an external switch 207 that allows extra-cluster access to each node. Each switch requires enough ports to handle all potential nodes in a cluster. Ethernet or GigE switches may be used for this purpose. PDUs 206 are used to power all nodes and switches, and the UPSs 208 are used that protect all

nodes and switches. Although not meant to be limiting, typically a cluster is connectable to a network, such as the public Internet, an enterprise intranet, or other wide area or local area network. In an illustrative embodiment, the cluster is implemented within an enterprise environment. It may be reached, for example, by navigating through a site's corporate DNS name server. Thus, for example, the cluster's domain may be a new sub-domain of an existing domain. In a representative implementation, the sub-domain is delegated in the corporate DNS server to the name servers in the cluster itself. End users access the cluster using any conventional interface or access tool. Thus, for example, access to the cluster may be carried out over any IP-based protocol (HTTP, FTP, NFS, AFS, SMB, a Web service, or the like), via an API, or through any other known or later-developed access method, service, program, or tool.

[0036] Client applications access the cluster through one or more types of external gateways such as standard UNIX file protocols, or HTTP APIs. The archive preferably is exposed through a virtual file system that can optionally sit under any standard UNIX file protocol-oriented facility. These include NFS, FTP, SMB/CIFS, or the like.

[0037] In one embodiment, the archive cluster application runs on a redundant array of independent nodes (H-RAIN) that are networked together (e.g., via Ethernet) as a cluster. The hardware of given nodes may be heterogeneous. For maximum reliability, however, preferably each node runs an instance 300 of the distributed application (which may be the same instance, or substantially the same instance), which comprises several runtime components as now illustrated in FIG. 3. Thus, while hardware may

be heterogeneous, the software stack on the nodes (at least as it relates to the present invention) is the same. These software components comprise a gateway protocol layer 302, an access layer 304, a file transaction and administration layer 306, and a core components layer 308. The "layer" designation is provided for explanatory purposes, as one of ordinary skill will appreciate that the functions may be characterized in other meaningful ways. One or more of the layers (or the components therein) may be integrated or otherwise. Some components may be shared across layers.

[0038] The gateway protocols in the gateway protocol layer 302 provide transparency to existing applications. In particular, the gateways provide native file services such as NFS 310 and SMB/CIFS 312, as well as a Web services API to build custom applications. HTTP support 314 is also provided. The access layer 304 provides access to the archive. In particular, according to the invention, a Fixed Content File System (FCFS) 316 emulates a native file system to provide full access to archive objects. FCFS gives applications direct access to the archive contents as if they were ordinary files. Preferably, archived content is rendered in its original format, while metadata is exposed as files. FCFS 316 provides conventional views of directories and permissions and routine file-level calls, so that administrators can provision fixed-content data in a way that is familiar to them. File access calls preferably are intercepted by a user-space daemon and routed to the appropriate core component (in layer 308), which dynamically creates the appropriate view to the calling application. FCFS calls preferably are constrained by archive policies to facilitate autonomous archive management.

Thus, in one example, an administrator or application cannot delete an archive object whose retention period (a given policy) is still in force.

[0039] The access layer 304 preferably also includes a Web user interface (UI) 318 and an SNMP gateway 320. The Web user interface 318 preferably is implemented as an administrator console that provides interactive access to an administration engine 322 in the file transaction and administration layer 306. The administrative console 318 preferably is a password-protected, Web-based GUI that provides a dynamic view of the archive, including archive objects and individual nodes. The SNMP gateway 320 offers storage management applications easy access to the administration engine 322, enabling them to securely monitor and control cluster activity. The administration engine monitors cluster activity, including system and policy events. The file transaction and administration layer 306 also includes a request manager process 324. The request manager 324 orchestrates all requests from the external world (through the access layer 304), as well as internal requests from a policy manager 326 in the core components layer 308.

[0040] In addition to the policy manager 326, the core components also include a metadata manager 328, and one or more instances of a storage manager 330. A metadata manager 328 preferably is installed on each node. Collectively, the metadata managers in a cluster act as a distributed database, managing all archive objects. On a given node, the metadata manager 328 manages a subset of archive objects, where preferably each object maps between an external file ("EF," the data that entered the archive for storage) and a set of internal files (each an "IF") where the archive data is

physically located. The same metadata manager 328 also manages a set of archive objects replicated from other nodes. Thus, the current state of every external file is always available to multiple metadata managers on several nodes. In the event of node failure, the metadata managers on other nodes continue to provide access to the data previously managed by the failed node. The storage manager 330 provides a file system layer available to all other components in the distributed application. Preferably, it stores the data objects in a node's local file system. Each drive in a given node preferably has its own storage manager. This allows the node to remove individual drives and to optimize throughput. The storage manager 330 also provides system information, integrity checks on the data, and the ability to traverse directly local structures.

[0041] As also illustrated in FIG. 3, the cluster manages internal and external communication through a communications middleware layer 332 and a DNS manager 334. The infrastructure 332 is an efficient and reliable message-based middleware layer that enables communication among archive components. In an illustrated embodiment, the layer supports multicast and point-to-point communications. The DNS manager 334 runs distributed name services that connect all nodes to the enterprise server. Preferably, the DNS manager (either alone or in conjunction with a DNS service) load balances requests across all nodes to ensure maximum cluster throughput and availability.

[0042] In an illustrated embodiment, an application such as the HCP (Hitachi Content Platform) application instance executes on a base operating system 336, such as Red Hat Linux 9.0, Fedora Core 6, or the like. The

communications middleware is any convenient distributed communication mechanism. Other components may include FUSE (Filesystem in USErspace), which may be used for the Fixed Content File System (FCFS) 316. The NFS gateway 310 may be implemented by a standard nfsd Linux Kernel NFS driver. The database in each node may be implemented, for example, PostgreSQL (also referred to herein as Postgres), which is an object-relational database management system (ORDBMS). The node may include a Web server, such as Jetty, which is a Java HTTP server and servlet container. Of course, the above mechanisms are merely illustrative.

[0043] The storage manager 330 on a given node is responsible for managing the physical storage devices. Preferably, each storage manager instance is responsible for a single root directory into which all files are placed according to its placement algorithm. Multiple storage manager instances can be running on a node at the same time, and each usually represents a different physical disk in the system. The storage manager abstracts the drive and interface technology being used from the rest of the system. When the storage manager instance is asked to write a file, it generates a full path and file name for the representation for which it will be responsible. In a representative embodiment, each object to be stored on a storage manager is received as raw data to be stored, with the storage manager then adding its own metadata to the file as it stores the data to keep track of different types of information. By way of example, this metadata includes: EF length (length of external file in bytes), IF Segment size (size of this piece of the Internal File), EF Protection representation (EF protection mode), IF protection role (representation of this internal file), EF Creation timestamp (external file

timestamp), Signature (signature of the internal file at the time of the write (PUT), including a signature type), and EF Filename (external file filename).

Storing this additional metadata with the internal file data provides for additional levels of protection. In particular, scavenging can create external file records in the database from the metadata stored in the internal files.

Other policies can validate internal file hash against the internal file to validate that the internal file remains intact.

[0044] Internal files may be "chunks" of data representing a portion of the original "file" in the archive object, and they may be placed on different nodes to achieve striping and protection blocks. This breaking apart of an external file into smaller chunked units is not a requirement, however; in the alternative, internal files may be complete copies of the external file.

Typically, one external file entry is present in a metadata manager for each archive object, while there may be many internal file entries for each external file entry. Typically, internal file layout depends on the system. In a given implementation, the actual physical format of this data on disk is stored in a series of variable length records.

[0045] The request manager 324 is responsible for executing the set of operations needed to perform archive actions by interacting with other components within the system. The request manager supports many simultaneous actions of different types, is able to roll-back any failed transactions, and supports transactions that can take a long time to execute. The request manager further ensures that read/write operations in the archive are handled properly and guarantees all requests are in a known state at all times. It also provides transaction control for coordinating multiple read/write

operations across nodes to satisfy a given client request. In addition, the request manager caches metadata manager entries for recently used files and provides buffering for sessions as well as data blocks.

[0046] A cluster's primary responsibility is to store an unlimited number of files on disk reliably. A given node may be thought of as being "unreliable," in the sense that it may be unreachable or otherwise unavailable for any reason. A collection of such potentially unreliable nodes collaborate to create reliable and highly available storage. Generally, there are two types of information that need to be stored: the files themselves and the metadata about the files. Additional details of the fixed content distributed data storage can be found in U.S. Patent Publication No. 2007/0189153 and U.S. Patent No. 7,657,581, which are incorporated herein by reference.

[0047] As used herein, a namespace is a logical partition of the cluster, and essentially serves as a collection of objects particular to at least one defined application. Each namespace has a private filesystem with respect to other namespaces. Moreover, access to one namespace does not grant a user access to another namespace. A cluster/system of nodes is a physical archive instance. A tenant is a grouping of namespace(s) and possibly other subtenants. A cluster/system is a physical archive instance. See commonly assigned U.S. Patent Application Publication No. 2011/0106802, which is incorporated herein by reference in its entirety.

[0048] II. Domain Name Service, DNS Alias, and Replicating Link

[0049] In computer networking, systems are contacted by host domain name specification that uniquely identifies a system that resides in a domain. This organization is known as Domain Name System (DNS). This is a well

known standard for identifying and locating systems on a network. The details of DNS are described in many Request For Comments (RFCs) readily documented on the internet including, for example, RFC1034 and RFC 1035. The host domain name is also known as the DNS alias.

[0050] Within these specifications are the concepts of sub-zones and DNS Zone Files. A sub-zone is a grouping of names that share the same domain name. An example domain is starting network name of the company; for instance, hds.com is the top level domain name for Hitachi Data Systems. These domains may have many child sub-zones or individual systems that need to be identified. To resolve the resources within a domain, a server is deployed that will act as a DNS Server and will utilize a DNS Zone file that identifies all the resources with the domain. The domain and the DNS Server IP address is registered with a higher level DNS Server to advertise itself and the domain it services.

[0051] A content platform contains a DNS Server 111 (FIG. 1) and is assigned a domain name. This domain name is registered as a sub-zone of a higher level domain within the company deploying the system. For example, for a content platform system named "hcp1" that is deployed, at the top level is the "hds.com" domain with the sub-zone name of "hcp1.hds.com." The content platform DNS Server 111 is responsible for resolving addresses for the resources that the content platform provides. These resources are virtual content platform systems organized as tenants and namespaces. An example DNS name for one such resource would be "ns1.ten1.hcp1.hds.com." These resources are configured in a DNS Zone File on each node of a content platform system for which the DNS Server 111

uses to resolve the names. A DNS name for a resource is also referred to as a DNS alias. A content platform system will have many DNS aliases to access the various resources on the system.

[0052] Replication is the function of making multiple copies of data on multiple systems typically in geographically separated locations for the purpose of providing disaster recovery of data in the event of a catastrophic loss of one system. The simplest topology for a replication is one-to-one replication where a first system is replicating content to a second system. More complex topologies can consist of a many-to-one as well as a combination of multiple basic topologies. Within these topologies, data may be copied in only one direction or both directions.

[0053] A content platform can be configured into a complex configuration consisting of a Many-to-One inward replication along with one-to-one outward replication. FIG. 4 shows an example of many-to-one replication with disaster recovery. This topology is useful for consolidating content from edge systems (A, D, E) into a main data center system (B), and then further protecting the main data center system with a geographically deployed disaster recovery system (C).

[0054] Each content platform system in this configuration is a self-sufficient system hosting a DNS Server for the sub-zone for which it is configured. In FIG. 4, there are replication links represented by the arrows between the content platform systems. These arrows indicate the data flow between the systems; for example, the arrow between system A and system B indicates that system A is configured with an outbound replication link and will send data to system B which has an inbound link configuration. While

system A and system B have knowledge about each other, system A has no knowledge about system C, system D, or system E.

[0055] III. DNS Alias Synchronization

[0056] According to exemplary embodiments, the replicated object storage system or content platform performs DNS alias synchronization in order to maintain the knowledge, within each individual system of the replicated system or content platform, of all replication links and DNS aliases of systems in the replicated system or content platform.

[0057] In the replicated system or content platform configuration, each individual system is self-contained and has the responsibility of servicing data requests addressed specifically to it, as identified by the DNS alias provided in the request. A system may have data initially created and alternatively addressed from another system in the ecosystem. If a request addressed to a first system in the ecosystem is received by a second system that is also hosting the data, the second system should service the request on behalf of the first system. However, any request received by any system must be validated to ensure the system specified in the request is actually part of the overall ecosystem.

[0058] As seen in FIG. 4, data written to system A will be replicated to system B and further replicated to system C. In the event that system A is unavailable, requests originally addressed to system A should be able to be satisfied by either system B or system C. While system B knows about system A through the replication link configuration ($A \rightarrow B$), system C does not know about system A. Therefore, any requests for content addressed to system A will be rejected if requested by system C.

[0059] This invention overcomes the problem by maintaining the knowledge of all replication links and DNS aliases of systems in the ecosystem within each individual system. In each system, this knowledge is called the Known Alias List (KAL). Each system is responsible for notifying all its replication-link-connected systems of its KAL. These KAL notifications are periodic or are triggered either by changes to its own replication link configuration or by receiving KAL notification from a system replication peer. With this method, the KAL on all systems will eventually become synchronized and notifications of changes will then discontinue. This method provides both periodic synchronization insurance and DNS alias synchronization for explicit addition/removal of system from the ecosystem.

[0060] III.A. Overall Process

[0061] FIG. 5 shows the steady state of the replicated system of FIG. 4 after achieving "eventually consistent" DNS alias synchronization. In such steady state, each system has knowledge of all other systems in the KAL. Below each system in the diagram of FIG. 5 is a simplified representation of the KAL showing what other systems are in the replicated system.

[0062] The KAL contains an alias list of all system(s) currently configured in the replicated system and a list of what replication links have recently been deleted. The KALs throughout the replicated system will be updated when one of the following triggers initiates a KAL synchronization (KAL-sync): (1) create/restore/deletion of a replication link, (2) update of replication link system name configuration, (3) change of the DNS name of a system, and (4) timer expiration for periodic DNS alias synchronization (e.g., every 2 hours). This processing method produces an "eventually consistent"

synchronization. All systems in the replicated system may not immediately have their KAL synchronized, but after multiple updates all systems will eventually be updated for any given change. All KAL-sync operations apply between only 2 clusters/systems (although they may trigger other KAL-sync operations between other system pairs). This approach is more scalable, and does not require that all systems in the replicated system be aware of each other. The result is that each cluster in a replication topology will contain a list of all other DNS aliases for systems in the replication topology (e.g., stored in memory), without ever directly communicating with them. This allows any given cluster in a replication topology to service requests which have been redirected from any other cluster.

[0063] This processing method provides a number of benefits. There is only ever one alias per cluster in the topology. Updates automatically remove any previous (now stale) aliases. Alias discovery/replacement continues to remove any clusters that are no longer relevant. Each cluster only communicates with those directly connected to it. The DNS server 111 preferably updates automatically the DNS alias information based on periodic inputs from the clusters.

[0064] III.B. DNS Alias Synchronization Workflow

[0065] FIG. 6 shows an example of DNS alias synchronization workflow in a replicated system. Upon synchronization trigger, the triggered system initiates a replicated system wide KAL-sync as the initiating system. The initiating system does so by sending its local KAL to each system directly connected to it by replication links (S601). Each remote system (set I) that receives the sent KAL will update its own KAL and immediately send the

result back to the initiating system (S602). The initiating system merges all remote system KALs which it receives in response, and sets it as the local KAL (S603).

[0066] A remote system, upon receiving a KAL-sync request and if changes are made to its local KAL, will continue to propagate the KAL-sync by initiating a KAL-sync with all other remote systems over replication links other than the one that initiated the KAL-sync (S604). The next set (set II) of replication linked systems perform steps S605 and S607 similar to S602 and S604. In S606, the current system updates its local KAL if necessary in response to S605. If this KAL-sync propagation causes changes to its local KAL, it will then trigger a new KAL-sync with the current system as a new initiating system (S608). Note that the KAL-sync for the current system will stop if no changes are made to its local KAL based on either the initial KAL-sync request in S602 or the response(s) from the remote system(s) as a result of propagating the KAL-sync in S606.

[0067] The following describes the KAL update processing. In specific embodiments, a KAL has two lists. The first is a mapping of replication link identifiers and domain names of systems known to be active in the replicated system (e.g., A→B, D→B, E→B, and B→C in FIG. 5). The second list has recently deleted replication link identifier(s). This feature is described further in connection with FIG. 13 below. In the context of any individual system within the replicated system environment, there are two activities that can cause the system to update its local KAL. The first is initiating a KAL-sync request. The second is receiving a KAL-sync request from a peer system.

[0068] FIG. 7 is an example of a flow diagram illustrating DNS alias synchronization for an initiating system. When a system is triggered to initiate a KAL-sync request, it will perform the following steps. For all replication links connected locally/directly, the system will send a KAL-sync request to each replication link remote system, the KAL-sync request containing both the local KAL and the local system DNS alias, and will store the response from the remote system (S701). See S601 and S603. The system will generate a merged KAL based on the response(s) from all remote system(s) (S702). See S603. The system will store the merged KAL as its local KAL (S703). See S603. This will remove any systems that were in the prior local KAL that were not rediscovered by the remote systems responses. The deleted replication links list will be cleared.

[0069] FIG. 8 is an example of a flow diagram illustrating DNS alias synchronization for a system that receives a synchronization request from a peer system. When a system receives a KAL-sync request from a peer, it will perform the following steps. The system will merge the KAL information from the request with its local KAL and persist it within the system (S801). See S602. The merge includes processing any deleted replication link entries. The merged KAL is stored as the local KAL. The system will immediately return a response to the KAL-sync request to the requestor and the response will include its local KAL less the deleted link entries received in the KAL-sync request (S802). See S602. For all other replication links connected (excluding the link from which the initial request was received), if its local KAL has changed based on the initial request, the system will propagate the KAL-sync request to each system directly connected by replication link except the

requestor system, and merge response(s) from the replication link system(s) with the local KAL while removing any deleted links (S803). See S604 and S606. Otherwise, the KAL-sync ends for the current system. Following S803, if there are any changes to the KAL as compared with the local KAL generated in S801, as a result of S803, it will trigger an event to initiate a new, full, local KAL-sync for the current system as a new initiating system according to FIG. 7 (S804). See S608. Otherwise, the KAL-sync ends for the current system.

[0070] Each system/cluster has a DNS alias synchronization module 222 (FIG. 2) that is configured to perform DNS alias synchronization including, for instance, the steps in FIGS. 6-8. The module 333 may be implemented in any suitable manner to maintain updated information of all replication links involving the system on which the module 333 resides and of DNS aliases of other systems associated with all replication links involving the system. In specific embodiments, the module 333 performs KAL-sync and maintains KAL-sync tables such as those shown in FIGS. 10, 12, and 14 below.

[0071] III.C. System Addition in Replication Link Chain

[0072] FIG. 9 shows an example of system addition in a replication link chain. The existing replicated system has systems A and B coupled by replication link A→B. A new system C is added to the replicated system with a replication link B→C. This triggers a KAL-sync on system B since a new link is created there.

[0073] FIG. 10 is a table illustrating a process of DNS alias synchronization for the system addition of FIG. 9. The table shows the process flow to update all the KALs in the replicated system. The changes to

the values in the table are marked in underlined. The initial state of the KAL shows an alias list of AB for systems A and B and C for system C. After adding the replication link $B \rightarrow C$, there is a KAL-sync between systems B and C to produce ABC for both systems B and C. Next, system B notifies system A of the change to produce ABC for system A. System A responds to the notification from system B with known aliases. This produces no changes to the KAL. System B finds no changes and the synchronization process ends.

[0074] FIG. 11 shows an example of system addition in a many-to-one replicated system with disaster recovery. The existing replicated system has systems A, B, C, P, and R with replication links $A \rightarrow P$, $B \rightarrow P$, $C \rightarrow P$, and $P \rightarrow R$. This example shows an initial replication topology with a Many-to-One configuration ($A, B, C \rightarrow P$) with a disaster recovery system R. A new system Z is added to the replicated system with a replication link $Z \rightarrow R$. This triggers a KAL-sync on system R since a new link is created there.

[0075] FIG. 12 is a table illustrating a process of DNS alias synchronization for the system addition of FIG. 11. The table shows the process flow to update all the KALs in the replicated system. It shows the transitions to the KALs on each system as the systems process the replicated-system-wide KAL-sync. The changes to the values in the table are marked in underlined. The initial state of the KAL shows an alias list of PRABC for systems A, B, C, P, and R and Z for system Z. After adding the replication link $Z \rightarrow R$, there is a KAL-sync between systems R and Z by trading aliases to produce PRABCZ for both systems R and Z. Next, system R notifies system P of the change to produce PRABCZ for system P. System P responds to the notification from system R with known aliases. This

produces no changes to the KAL. System P then notifies systems A, B, and C of the change to produce PRABCZ for systems A, B, and C. Systems A, B, and C each respond to the notification from system P. System P finds no changes and the synchronization process ends.

[0076] III.D. System Deletion in Replication Link Chain

[0077] FIG. 13 shows an example of system deletion in a many-to-one replicated system with disaster recovery. The existing replicated system has systems A, B, C, P, R, and Z with replication links $A \rightarrow P$, $B \rightarrow P$, $C \rightarrow P$, $P \rightarrow R$, and $Z \rightarrow R$. The replication link $C \rightarrow P$ is deleted. This can result from the deletion of system C or unavailability of system C for various reasons, or the like. This triggers a KAL-sync on system P since a link is deleted there.

[0078] This introduces the concept of providing the KALs with replication link delete information to aid in cleaning up systems from the individual KALs (recently deleted replication link identifier(s)). The life cycles of the delete entries are short lived and are used to avoid the situation where a system DNS alias is re-introduced into a KAL in error because a system in the replicated system still thinks the alias is valid and reports it to other systems as being valid, thus introducing it to the KAL of other systems in error. This erroneous reintroduction can occur because the process of performing KAL-sync operations is asynchronous. When a system receives a KAL-sync with a replication delete (i.e., deleted replication link), it will use the information to filter all future incoming and outgoing KALs. Once any given system has received notification that all other systems directly connected via local replication links are free of the deleted replication link, the local record of replication link deletion can be cleared.

[0079] FIG. 14 is a table illustrating a process of DNS alias synchronization for the system deletion of FIG. 13. The table shows the process flow to remove system C from the replication topology and to propagate that change to the other systems by updating all the KALs in the replicated system. It shows the transitions to the KALs on each system as the systems process the replicated-system-wide KAL-sync. The changes to the values in the table are marked in underlined. The initial state of the KAL shows an alias list of PRABCZ for all systems A, B, C, P, R, and Z. After deleting the replication link C→P, there is a KAL-sync between systems C and P to produce an alias list of PRABZ with Del-CP (indicating recently deleted link C→P) for system P and C for system C. Next, system P triggers a KAL-sync on systems A, B, and R. Systems A, B, and R process the KAL-sync request and responds to system P to produce an alias list of PRABZ for systems A, B, and R, and Del-CP for system R. Next, system R triggers a KAL-sync on system Z. System Z responds to system R with an updated alias of PRABZ. This produces no changes to the KAL. System R finds no changes and the synchronization process ends.

[0080] With all systems in the replication topology having knowledge about all the replication links and DNS aliases for all other systems in the topology, this will lay the groundwork for more advanced operations within replication topologies. Some potential uses would be for the core of implementing an active-active replication topology and/or providing the ability to publish/discover the topology to features that allow geographically dispersed systems to act as one system, thus allowing localization of data transfer to and from the replicated system. In addition, this information about

the overall topology will help in making decisions with objects to create more intelligence with the life cycle of objects. For instance, an object could be tagged with metadata to reside in a geographical location as its final resting place within the replicated system without having to be ingested into a system in that geographical location.

[0081] Of course, the system configuration illustrated in FIG. 1 is purely exemplary of content platforms or replicated object storage systems in which the present invention may be implemented, and the invention is not limited to a particular hardware configuration. The computers and storage systems implementing the invention can also have known I/O devices (e.g., CD and DVD drives, floppy disk drives, hard drives, etc.) which can store and read the modules, programs and data structures used to implement the above-described invention. These modules, programs and data structures can be encoded on such computer-readable media. For example, the data structures of the invention can be stored on computer-readable media independently of one or more computer-readable media on which reside the programs used in the invention. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include local area networks, wide area networks, e.g., the Internet, wireless networks, storage area networks, and the like.

[0082] In the description, numerous details are set forth for purposes of explanation in order to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that not all of these specific details are required in order to practice the present invention. It

is also noted that the invention may be described as a process, which is usually depicted as a flowchart, a flow diagram, a structure diagram, or a block diagram. Although a flowchart may describe the operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be re-arranged.

[0083] As is known in the art, the operations described above can be performed by hardware, software, or some combination of software and hardware. Various aspects of embodiments of the invention may be implemented using circuits and logic devices (hardware), while other aspects may be implemented using instructions stored on a machine-readable medium (software), which if executed by a processor, would cause the processor to perform a method to carry out embodiments of the invention. Furthermore, some embodiments of the invention may be performed solely in hardware, whereas other embodiments may be performed solely in software. Moreover, the various functions described can be performed in a single unit, or can be spread across a number of components in any number of ways. When performed by software, the methods may be executed by a processor, such as a general purpose computer, based on instructions stored on a computer-readable medium. If desired, the instructions can be stored on the medium in a compressed and/or encrypted format.

[0084] From the foregoing, it will be apparent that the invention provides methods, apparatuses and programs stored on computer readable media for DNS alias synchronization in a replicated object storage system. Additionally, while specific embodiments have been illustrated and described in this specification, those of ordinary skill in the art appreciate that any

arrangement that is calculated to achieve the same purpose may be substituted for the specific embodiments disclosed. This disclosure is intended to cover any and all adaptations or variations of the present invention, and it is to be understood that the terms used in the following claims should not be construed to limit the invention to the specific embodiments disclosed in the specification. Rather, the scope of the invention is to be determined entirely by the following claims, which are to be construed in accordance with the established doctrines of claim interpretation, along with the full range of equivalents to which such claims are entitled.

[0085]

[Citation List]

[Patent Literature]

[0086]

[PTL 1]

U.S. Patent No. 7,155,466

[PTL 2]

U.S. Patent Publication No. 2007/0189153

[PTL 3]

U.S. Patent No. 7,657,581

[PTL 4]

U.S. Patent Publication No. 2011/0106802

WHAT IS CLAIMED IS:

1. In a redundant array of independent nodes, the nodes being grouped into a plurality of systems each having multiple nodes, the plurality of systems having one or more replication links each being formed to indicate replication of data from one system to another system in a replicated environment where each system is configured as a sub-domain in a Domain Name System (DNS) infrastructure, each system having a DNS name that is also a DNS alias of said each system to other systems, a DNS alias synchronization method comprising:

maintaining updated information, within each system of the plurality of systems, of all replication links involving said each system and of DNS aliases of other systems associated with said all replication links involving said each system;

wherein the updated information of all replication links involving a system and of DNS aliases of other systems associated with all replication links involving the system enables the system to process network-based requests, on behalf of the other systems associated with all replication links involving the system, without redirecting the requests from the other systems to the system.

2. The method according to claim 1, further comprising:

providing the updated information, from each system, to a DNS server which enables said each system to process network-based requests, on behalf of the other systems associated with all replication links involving said

each system, without redirecting the requests from the other systems to said each system.

3. The method according to claim 1, wherein maintaining updated information comprises:

updating a known alias list within said each system, the known alias list including mapping of the replication links between the DNS names of the systems associated with all replication links involving said each system, and any replication links that have recently been deleted from said all replication links involving said each system.

4. The method according to claim 3, wherein for a system to initiate updating the known alias list as an initiating system, the updating comprises:

sending a synchronization request to all systems that are each connected to the initiating system by a direct replication link, the synchronization request including a local known alias list of the initiating system and the DNS alias of the initiating system;

receiving a response from each system that received the synchronization request from the initiating system;

generating a merged known alias list for the initiating system from all the responses received by the initiating system; and

storing the merged known alias list generated for the initiating system as the local known alias list of the initiating system.

5. The method according to claim 3, wherein for a system that receives, as a receiving system, a synchronization request from another system with a known alias list of said another system, the updating comprises:

merging the known alias list from the synchronization request with a local known alias list of the receiving system to produce a merged known alias list for the receiving system;

storing the merged known alias list for the receiving system as the local known alias list of the receiving system; and

returning a response to said another system that sent the synchronization request, the response including the merged known alias list for the receiving system.

6. The method according to claim 5, wherein if the local known alias list of the receiving system has changed based on the synchronization request from said another system, the updating further comprises:

propagating the synchronization request to all systems that are each connected to the receiving system by a direct replication link excluding said another system, the synchronization request including the local known alias list of the receiving system and the DNS alias of the receiving system;

receiving a response from each system that received the propagated synchronization request from the receiving system;

generating a new merged known alias list for the receiving system from all the responses received by the receiving system; and

storing the new merged known alias list as the local known alias list of the receiving system.

7. The method according to claim 6, wherein if the new merged known alias list is different from the local known alias list of the receiving system prior to generating the merged known alias list for the receiving system, the updating further comprises:

sending a second synchronization request, with the receiving system as a new initiating system, to all systems that are each connected to the new initiating system by a direct replication link, the second synchronization request including the local known alias list of the new initiating system and the DNS alias of the new initiating system;

receiving a response from each system that received the second synchronization request from the new initiating system;

generating a merged known alias list for the new initiating system from all the responses received by the new initiating system; and

storing the merged known alias list for the new initiating system as the local known alias list of the new initiating system.

8. An apparatus for Domain Name System (DNS) alias synchronization in a redundant array of independent nodes, the nodes being grouped into a plurality of systems each having multiple nodes, the plurality of systems having one or more replication links each being formed to indicate replication of data from one system to another system in a replicated environment where each system is configured as a sub-domain in a DNS infrastructure, each system having a DNS name that is also a DNS alias of said each system to other systems, the apparatus comprising a processor, a memory, and a DNS

alias synchronization module provided for each system, the DNS alias synchronization module being configured to:

maintain updated information, within the system in which the DNS alias synchronization module is provided, of all replication links involving the system and of DNS aliases of other systems associated with said all replication links involving the system;

wherein the updated information of all replication links involving a system and of DNS aliases of other systems associated with all replication links involving the system enables the system to process network-based requests, on behalf of the other systems associated with all replication links involving the system, without redirecting the requests from the other systems to the system.

9. The apparatus according to claim 8, wherein the DNS alias synchronization module is configured to:

provide the updated information, from the system in which the DNS alias synchronization module is provided, to a DNS server which enables the system to process network-based requests, on behalf of the other systems associated with all replication links involving the system, without redirecting the requests from the other systems to the system.

10. The apparatus according to claim 8, wherein maintaining updated information comprises:

updating a known alias list within the system, the known alias list including mapping of the replication links between the DNS names of the

systems associated with all replication links involving the system, and any replication links that have recently been deleted from said all replication links involving the system.

11. The apparatus according to claim 10, wherein for a system to initiate updating the known alias list as an initiating system, the updating to be performed by the DNS alias synchronization module on the initiating system comprises:

- sending a synchronization request to all systems that are each connected to the initiating system by a direct replication link, the synchronization request including a local known alias list of the initiating system and the DNS alias of the initiating system;

- receiving a response from each system that received the synchronization request from the initiating system;

- generating a merged known alias list for the initiating system from all the responses received by the initiating system; and

- storing the merged known alias list generated for the initiating system as the local known alias list of the initiating system.

12. The apparatus according to claim 10, wherein for a system that receives, as a receiving system, a synchronization request from another system with a known alias list of said another system, the updating to be performed by the DNS alias synchronization module on the receiving system comprises:

merging the known alias list from the synchronization request with a local known alias list of the receiving system to produce a merged known alias list for the receiving system;

storing the merged known alias list for the receiving system as the local known alias list of the receiving system; and

returning a response to said another system that sent the synchronization request, the response including the merged known alias list for the receiving system.

13. The apparatus according to claim 12, wherein if the local known alias list of the receiving system has changed based on the synchronization request from said another system, the updating to be performed by the DNS alias synchronization module on the receiving system further comprises:

propagating the synchronization request to all systems that are each connected to the receiving system by a direct replication link excluding said another system, the synchronization request including the local known alias list of the receiving system and the DNS alias of the receiving system;

receiving a response from each system that received the propagated synchronization request from the receiving system;

generating a new merged known alias list for the receiving system from all the responses received by the receiving system; and

storing the new merged known alias list as the local known alias list of the receiving system.

14. The apparatus according to claim 13, wherein if the new merged known alias list is different from the local known alias list of the receiving system prior to generating the merged known alias list for the receiving system, the updating to be performed by the DNS alias synchronization module on the receiving system further comprises:

sending a second synchronization request, with the receiving system as a new initiating system, to all systems that are each connected to the new initiating system by a direct replication link, the second synchronization request including the local known alias list of the new initiating system and the DNS alias of the new initiating system;

receiving a response from each system that received the second synchronization request from the new initiating system;

generating a merged known alias list for the new initiating system from all the responses received by the new initiating system; and

storing the merged known alias list for the new initiating system as the local known alias list of the new initiating system.

15. A computer-readable storage medium storing a plurality of instructions for controlling a data processor to perform Domain Name System (DNS) alias synchronization, in a redundant array of independent nodes, the nodes being grouped into a plurality of systems each having multiple nodes, the plurality of systems having one or more replication links each being formed to indicate replication of data from one system to another system in a replicated environment where each system is configured as a sub-domain in a DNS infrastructure, each system having a DNS name that is also a DNS alias of

said each system to other systems, wherein the computer-readable storage medium is provided in each system, the plurality of instructions comprising:

instructions that cause the data processor to maintain updated information, within the system in which the computer-readable storage medium is provided, of all replication links involving the system and of DNS aliases of other systems associated with said all replication links involving the system;

wherein the updated information of all replication links involving a system and of DNS aliases of other systems associated with all replication links involving the system enables the system to process network-based requests, on behalf of the other systems associated with all replication links involving the system, without redirecting the requests from the other systems to the system.

16. The computer-readable storage medium according to claim 15, wherein the instructions that cause the data processor to maintain updated information comprise:

instructions that cause the data processor to update a known alias list within the system, the known alias list including mapping of the replication links between the DNS names of the systems associated with all replication links involving the system, and any replication links that have recently been deleted from said all replication links involving the system.

17. The computer-readable storage medium according to claim 16, wherein for a system to initiate updating the known alias list as an initiating

system, the instructions that cause the data processor to update a known alias list within the system comprise:

instructions that cause the data processor to send a synchronization request to all systems that are each connected to the initiating system by a direct replication link, the synchronization request including a local known alias list of the initiating system and the DNS alias of the initiating system;

instructions that cause the data processor to receive a response from each system that received the synchronization request from the initiating system;

instructions that cause the data processor to generate a merged known alias list for the initiating system from all the responses received by the initiating system; and

instructions that cause the data processor to store the merged known alias list generated for the initiating system as the local known alias list of the initiating system.

18. The computer-readable storage medium according to claim 16, wherein for a system that receives, as a receiving system, a synchronization request from another system with a known alias list of said another system, the instructions that cause the data processor to update a known alias list within the system comprise:

instructions that cause the data processor to merge the known alias list from the synchronization request with a local known alias list of the receiving system to produce a merged known alias list for the receiving system;

instructions that cause the data processor to store the merged known alias list for the receiving system as the local known alias list of the receiving system; and

instructions that cause the data processor to return a response to said another system that sent the synchronization request, the response including the merged known alias list for the receiving system.

19. The computer-readable storage medium according to claim 18, wherein if the local known alias list of the receiving system has changed based on the synchronization request from said another system, the instructions that cause the data processor to update a known alias list within the system further comprise:

instructions that cause the data processor to propagate the synchronization request to all systems that are each connected to the receiving system by a direct replication link excluding said another system, the synchronization request including the local known alias list of the receiving system and the DNS alias of the receiving system;

instructions that cause the data processor to receive a response from each system that received the propagated synchronization request from the receiving system;

instructions that cause the data processor to generate a new merged known alias list for the receiving system from all the responses received by the receiving system; and

instructions that cause the data processor to store the new merged known alias list as the local known alias list of the receiving system.


20. The computer-readable storage medium according to claim 19, wherein if the new merged known alias list is different from the local known alias list of the receiving system prior to generating the merged known alias list for the receiving system, the instructions that cause the data processor to update a known alias list within the system further comprise:


instructions that cause the data processor to send a second synchronization request, with the receiving system as a new initiating system, to all systems that are each connected to the new initiating system by a direct replication link, the second synchronization request including the local known alias list of the new initiating system and the DNS alias of the new initiating system;


instructions that cause the data processor to receive a response from each system that received the second synchronization request from the new initiating system;

instructions that cause the data processor to generate a merged known alias list for the new initiating system from all the responses received by the new initiating system; and

instructions that cause the data processor to store the merged known alias list for the new initiating system as the local known alias list of the new initiating system.

**FIG. 1**

**FIG. 2**






FIG. 4








FIG. 5

**FIG. 6**


**FIG. 7**

**FIG. 8**

**FIG. 9**


EVENT	ACTION	KNOWN ALIAS LIST (KAL)		
		A	B	C
INITIAL STATE		AB	AB	C
LINK B->C ADDED	B AND C SYNC KAL WITH EACH OTHER	AB	ABC	ABC
	B NOTIFIES A OF CHANGE	ABC	ABC	ABC
	A RESPONDS TO B WITH KNOWN ALIASES	ABC	ABC	ABC
	B FINDS NO CHANGES AND COMPLETES	ABC	ABC	ABC

FIG. 10

**FIG. 11**

EVENT	ACTION	KNOWN ALIAS LIST (KAL) ON SYSTEMS					
		P	R	A	B	C	Z
INITIAL STATE		PRABC	PRABC	PRABC	PRABC	PRABC	Z
ADD LINK Z -> R	Z & R TRADE ALIASES	PRABC	PRABC <u>Z</u>	PRABC	PRABC	PRABC	<u>PRABCZ</u>
	R NOTIFIES P OF CHANGE	PRABC <u>Z</u>	PRABCZ	PRABC	PRABC	PRABC	PRABCZ
	P RESPONDS TO R WITH ALIASES. NO UPDATES	PRABCZ	PRABCZ	PRABC	PRABC	PRABC	PRABCZ
	P NOTIFIES A,B,C OF CHANGE	PRABCZ	PRABCZ	PRABC <u>Z</u>	PRABC <u>Z</u>	PRABC <u>Z</u>	PRABCZ
	A,B,C RESPONDS TO P. P FINDS NO CHANGES AND STOPS	PRABCZ	PRABCZ	PRABCZ	PRABCZ	PRABCZ	PRABCZ

FIG. 12



EVENT	ACTION	KNOWN ALIAS LIST (KAL) ON SYSTEMS					
		P	R	A	B	C	Z
INITIAL STATE		PRABCZ	PRABCZ	PRABCZ	PRABCZ	PRABCZ	PRABCZ
DELETE LINK C->P	C & P UPDATE KAL	PRABZ <u>DEL-CP</u>	PRABCZ	PRABCZ	PRABCZ	C	PRABCZ
	P TRIGGERS KAL-SYNC ON A,B AND R	PRABZ <u>DEL-CP</u>	PRABCZ	PRABCZ	PRABCZ	C	PRABCZ
	A,B, AND R PROCESS AND RESPONDS TO P	PRABZ	PRABZ <u>DEL-CP</u>	PRABZ	PRABZ	C	PRABCZ
	R TRIGGERS KAL-SYNC ON Z	PRABZ	PRABZ <u>DEL-CP</u>	PRABZ	PRABZ	C	PRABCZ
	Z RESPONDS TO R	PRABZ	PRABZ	PRABZ	PRABZ	C	PRABZ

FIG. 14

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 2012/031109

A. CLASSIFICATION OF SUBJECT MATTER

G06F 15/16 (2006.01)**H04L 12/28 (2006.01)**

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 12/00, 13/00, 13/38-13/42, 15/00, 15/16, 15/163, 15/173, 17/00, 17/40, H04L 12/00, 12/28-12/46, 12/66, 12/70, 12/70
12/741, 12/749, 12/751, 12/757, 29/00-29/12, H04W 16/00, 28/00, 40/00, 56/00, 84/00

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)

PatSearch (RUPTO internal), USPTO, PAJ, Esp@cenet

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No
A	US 2011/0106757 A1 (DAVID B. PICKNEY et al.) 05.05.2011, par. [0029]-[0031], [0045], [0080]-[0082], фиг. 1, 2	1-20
A	WO 2003/107204 A1 (NOKIA CORPORATION) 24.12.2003, abstract, par.[0008], [0018]-[0022]	1-20
A	WO 2011/091646 A1 (COMPUTER NETWORK INFORMATION CENTRE, CHINESE ACADEMY OF SCIENCES) 04.08.2011	1-20
A	WO 2008/157267 A1 (QUALCOMM INCORPORATED) 24.12.2008	1-20
A	US 2009/0113075 A1 (FRANCE TELECOM) 30.04.2009	1-20
A	RU 2310903 C2 (MOTOROLA INC) 20.11.2007	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

26 November 2012 (26.11.2012)

Date of mailing of the international search report

06 December 2012 (06.12.2012)

Name and mailing address of the ISA/ FIPS
Russia, 123995, Moscow, G-59, GSP-5,
Berezhkovskaya nab., 30-1

Authorized officer

S. Ilyasov

Facsimile No. +7 (499) 243-33-37

Telephone No. 499-240-25-91