



- (51) International Patent Classification:
G06F 12/00 (2006.01)
- (21) International Application Number:
PCT/US2011/060417
- (22) International Filing Date:
11 November 2011 (11.11.2011)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
- | | | |
|------------|-------------------------------|----|
| 12/947,393 | 16 November 2010 (16.11.2010) | US |
| 12/947,385 | 16 November 2010 (16.11.2010) | US |
| 12/947,436 | 16 November 2010 (16.11.2010) | US |
| 12/947,418 | 16 November 2010 (16.11.2010) | US |
| 12/947,383 | 16 November 2010 (16.11.2010) | US |
| 12/947,513 | 16 November 2010 (16.11.2010) | US |
| 12/947,438 | 16 November 2010 (16.11.2010) | US |
| 12/947,375 | 16 November 2010 (16.11.2010) | US |
- (71) Applicant (for all designated States except US): **ACTIFIO, INC.** [US/US]; 225 Wyman St., Waltham, MA 02451 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **ASHUTOSH, Ashutosh** [US/US]; 19 Lanes End, Weston, MA 02493 (US). **PROVENZANO, Christopher, A.** [US/US]; 95 Liberty Ave., Somerville, MA 02144 (US). **CHANG, Dav-**

id, F. [US/US]; 336 Highland St., Weston, MA 02493 (US). **ABERCROMBIE, Philip, J.** [US/US]; 79 Winn St., Belmont, MA 02478 (US). **MUTALIK, Madhav** [US/US]; 8 Graystone Way, Southborough, MA 01772 (US). **ROMAN, Mark, A.** [US/US]; 30 Ashland St., Arlington, MA 02476 (US).

- (74) Agent: **DICHIARA, Peter, M.**; Wilmer Cutler Pickering Hale And Dorr LLP, 60 State Street, Boston, MA 02109 (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,

[Continued on next page]

(54) Title: SYSTEMS AND METHODS FOR DATA MANAGEMENT VIRTUALIZATION




FIG. 4

- (57) Abstract: Systems and methods for data management virtualization are disclosed. The systems have a data management engine for performing data management functions, including at least a snapshot function and a back-up function, and a service level policy engine that controls the scheduling of the data management functions using service level agreements in electronic form. Data objects are organized in a deduplicated content store using an organized arrangement of temporal structures to represent states of data over time. Hash signatures are used to track content segments for performing deduplicated copies from a first deduplicated store to a second deduplicated store. Garbage collection is performed on the deduplicated content store for only content segments that have changed relative to an immediately-prior state of a given data object.



SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, —
GW, ML, MR, NE, SN, TD, TG).

*before the expiration of the time limit for amending the
claims and to be republished in the event of receipt of
amendments (Rule 48.2(h))*

Published:

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

SYSTEMS AND METHODS FOR DATA MANAGEMENT VIRTUALIZATION

Cross-References to Related Applications

[0001] This application claims priority to U.S. Patent Application No. 12/947,393, filed on November 16, 2010, entitled System and Method for Performing Backup or Restore Operations Utilizing Difference Information and Timeline State Information; U.S. Patent Application No. 12/947,385, filed on November 16, 2010, entitled System and Method for Managing Data with Service Level Agreements That May Specify Non-Uniform Copying of Data; U.S. Patent Application No. 12/947,436, filed on November 16, 2010, entitled System and Method for Performing a Plurality of Prescribed Data Management Functions in a Manner That Reduces Redundant Access Operations to Primary Storage; U.S. Patent Application No. 12/947,418, filed on November 16, 2010, entitled System and Method for Creating Deduplicated Copies of Data by Tracking Temporal Relationships Among Copies and by Ingesting Difference Data; U.S. Patent Application No. 12/947,375, filed on November 16, 2010, entitled System and Method for Managing Deduplicated Copies of Data Using Temporal Relationships Among Copies; U.S. Patent Application No. 12/947,383, filed on November 16, 2010, entitled System and Method for Improved Garbage Collection Operations in a Deduplicated Store by Tracking Temporal Relationships Among Copies; U.S. Patent Application No. 12/947,947,513, filed on November 16, 2010, entitled System and Method for Creating Deduplicated Copies of Data by Sending Difference Data Between Two Near-Neighbor Temporal States; and U.S. Patent Application No. 12/947,438, filed on November 16, 2010, entitled System and Method for Creating Deduplicated Copies of Data Storing Non-Lossy Encodings of Data Directly in a Content Addressable Store.

Technical Field

[0002] This invention relates generally to data management, data protection, disaster recovery and business continuity. More specifically, this invention relates to systems and methods for performing backup or restore operations utilizing difference information and timeline state information.

Background

[0003] The business requirements for managing the lifecycle of application data have been traditionally met by deploying multiple point solutions, each of which addresses a part of the lifecycle. This has resulted in a complex and expensive infrastructure where multiple copies of data are created and moved multiple times to individual storage repositories. The adoption of server virtualization has become a catalyst for simple, agile and low-cost compute infrastructure. This has led to larger deployments of virtual hosts and storage, further exacerbating the gap between the emerging compute models and the current data management implementations.

[0004] Applications that provide business services depend on storage of their data at various stages of its lifecycle. Figure 1 shows a typical set of data management operations that would be applied to the data of an application such as a database underlying a business service such as payroll management. In order to provide a business service, application 102 requires primary data storage 122 with some contracted level of reliability and availability.

[0005] Backups 104 are made to guard against corruption or the primary data storage through hardware or software failure or human error. Typically backups may be made daily or weekly to local disk or tape 124, and moved less frequently (weekly or monthly) to a remote physically secure location 125.

[0006] Concurrent development and test 106 of new applications based on the same database requires a development team to have access to another copy of the data 126. Such a snapshot might be made weekly, depending on development schedules.

[0007] Compliance with legal or voluntary policies 108 may require that some data be retained for safely future access for some number of years; usually data is copied regularly (say, monthly) to a long-term archiving system 128.

[0008] Disaster Recovery services 110 guard against catastrophic loss of data if systems providing primary business services fail due to some physical disaster. Primary data is copied 130 to a physically distinct location as frequently as is feasible given other constraints (such as cost). In the event of a disaster the primary site can be reconstructed and data moved back from the safe copy.

[0009] Business Continuity services 112 provide a facility for ensuring continued business services should the primary site become compromised. Usually this requires a hot copy 132 of the primary data that is in near-lockstep with the primary data, as well as duplicate systems and applications and mechanisms for switching incoming requests to the Business Continuity servers.

[0010] Thus, data management is currently a collection of point applications managing the different parts of the lifecycle. This has been an artifact of evolution of data management solutions over the last two decades.

Brief Description of the Drawings

[0011] FIG. 1 is a simplified diagram of current methods deployed to manage the data lifecycle for a business service.

[0012] FIG. 2 is an overview of the management of data throughout its lifecycle by a single Data Management Virtualization System.

[0013] FIG. 3 is a simplified block diagram of the Data Management Virtualization system.

[0014] FIG. 4 is a view of the Data Management Virtualization Engine.

[0015] FIG. 5 illustrates the Object Management and Data Movement Engine.

[0016] FIG. 6 shows the Storage Pool Manager.

[0017] FIG. 7 shows the decomposition of the Service Level Agreement.

[0018] FIG. 8 illustrates the Application Specific Module.

[0019] FIG. 9 shows the Service Policy Manager.

[0020] FIG. 10 is a flowchart of the Service Policy Scheduler.

[0021] FIG. 11 is a block diagram of the Content Addressable Storage (CAS) provider.

[0022] FIG. 12 shows the definition of an object handle within the CAS system.

[0023] FIG. 13 shows the data model and operations for the temporal relationship graph stored for objects within the CAS.

[0024] FIG. 14 is a diagram representing the operation of a garbage collection algorithm in the CAS.

[0025] FIG. 15 is a flowchart for the operation of copying an object into the CAS.

[0026] FIG. 16 is a system diagram of a typical deployment of the Data Management Virtualization system.

[0027] FIG. 17 is a schematic diagram of a characteristic physical server device for use with the Data Management Virtualization system.

Detailed Description

[0028] Current Data Management architecture and implementations such as described above involve multiple applications addressing different parts of data lifecycle management, all of them performing certain common functions: (a) make a copy of application data (the frequency of this action is commonly termed the Recovery Point Objective (RPO)), (b) store the copy of data in an exclusive storage repository, typically in a proprietary format, and (c) retain the copy for certain duration, measured as Retention Time. A primary difference in each of the point solutions is in the frequency of the RPO, the Retention Time, and the characteristics of the individual storage repositories used, including capacity, cost and geographic location.

[0029] This disclosure pertains to Data Management Virtualization. Data Management activities, such as Backup, Replication and Archiving are virtualized in that they do not have to be configured and run individually and separately. Instead, the user defines their business requirement with regard to the lifecycle of the data, and the Data Management Virtualization System performs these operations automatically. A snapshot is taken from primary storage to secondary storage; this snapshot is then used for a backup operation to other secondary storage. Essentially an arbitrary number of these backups may be made, providing a level of data protection specified by a Service Level Agreement.

[0030] This disclosure also pertains to a system for backing-up data from a first storage pool to a second storage pool using difference information between time states, incorporating a data management engine for performing data management functions, including at least a back-up function to create a back-up copy of data. The data management engine is operable to execute a sequence of snapshot operations to create point-in-time images of application data on a first storage pool, each successive point-in-time image corresponding to a specific, successive time-state of the application data, and each snapshot operation creating difference information indicating which application data has changed and the content of the changed application data for the corresponding time state. The data management engine is also operable to execute at least one back-up function for the application data that is scheduled for execution at non-consecutive time-states, and is also full of maintain history information having time-state information indicating the time-state of the last back-up function performed on the application data for a corresponding back-up copy of data. The data management engine creates composite difference information from the difference information for each time-state between the time-state of the last back-up function performed on the application data and the time-state of the currently-scheduled back-up function to be performed on the application data, and sends the composite difference information to a second storage pool to be compiled with the back-up copy of data at the last time-state to create a back-up copy of data for the current time-state.

[0031] Data Management Virtualization technology according to this disclosure is based on an architecture and implementation based on the following guiding principles.

[0032] First, define the business requirements of an application with a Service Level Agreement (SLA) for its entire data lifecycle. The SLA is much more than a single RPO, Retention and Recovery Time Objective (RTO). It describes the data protection characteristics for each stage of the data lifecycle. Each application may have a different SLA.

[0033] Second, provide a unified Data Management Virtualization Engine that manages the data protection lifecycle, moving data across the various storage repositories, with improved storage capacity and network bandwidth. The Data Management Virtualization system achieves these improvements by leveraging extended capabilities of modern storage systems by tracking the portions of the data that have changed over time and by data deduplication and compression algorithms that reduce the amount of data that needs to be copied and moved.

[0034] Third, leverage a single master copy of the application data to be the basis for multiple elements within the lifecycle. Many of the Data Management operations such as backup, archival and replication depend on a stable, consistent copy of the data to be protected. The Data Management Virtualization System leverages a single copy of the data for multiple purposes. A single instance of the data maintained by the system may serve as the source, from which each data management function may make additional copies as needed. This contrasts with requiring application data to be copied multiple times by multiple independent data management applications in the traditional approach.

[0035] Fourth, abstracting physical storage resources into a series of data protection storage pools, which are virtualized out of different classes of storage including local and remote disk, solid state memory, tape and optical media, private, public and/or hybrid storage clouds. The storage pools provide access independent of the type, physical location or underlying storage technology. Business requirements for the lifecycle of data may call for copying the data to different types of storage media at different times. The Data Management Virtualization system allows the user to classify and aggregate different storage media into storage pools, for example, a Quick Recovery Pool, which consists of high speed disks, and a Cost Efficient Long-term Storage Pool, which may be a deduplicated store on high capacity disks, or a tape library. The Data Management Virtualization System can move data amongst these pools to take advantage of the unique characteristics of each storage medium. The abstraction of Storage Pools provides access independent of the type, physical location or underlying storage technology.

[0036] Fifth, improve the movement of the data between storage pools and disaster locations utilizing underlying device capabilities and post-deduplicated application data. The Data Management Virtualization System discovers the capabilities of the storage systems that comprise the Storage Pools, and takes advantage of these capabilities to move data efficiently. If the Storage System is a disk array that supports the capability of creating a snapshot or clone of a data volume, the Data Management Virtualization System will take advantage of this capability and use a snapshot to make a copy of the data rather than reading the data from one place and writing it to another. Similarly, if a storage system supports change tracking, the Data Management Virtualization System will update an older copy with just the changes to efficiently create a new copy. When moving data across a network, the Data Management

Virtualization system uses a deduplication and compression algorithm that avoids sending data that is already available on the other side of the network.

[0037] One key aspect of improving data movement is recognizing that application data changes slowly over time. A copy of an application that is made today will, in general, have a lot of similarities to the copy of the same application that was made yesterday. In fact today's copy of the data could be represented as yesterday's copy with a series of delta transformations, where the size of the delta transformations themselves are usually much smaller than all of the data in the copy itself. The Data Management Virtualization system captures and records these transformations in the form of bitmaps or extent lists. In one embodiment of the system, the underlying storage resources - a disk array or server virtualization system - are capable of tracking the changes made to a volume or file; in these environments, the Data Management Virtualization system queries the storage resources to obtain these change lists, and saves them with the data being protected.

[0038] In the preferred embodiment of the Data Management Virtualization system, there is a mechanism for eavesdropping on the primary data access path of the application, which enables the Data Management Virtualization system to observe which parts of the application data are modified, and to generate its own bitmap of modified data. If, for example, the application modifies blocks 100, 200 and 300 during a particular period, the Data Management Virtualization system will eavesdrop on these events, and create a bitmap that indicates that these particular blocks were modified. When processing the next copy of application data, the Data Management Virtualization system will only process blocks 100, 200 and 300 since it knows that these were the only blocks that were modified.

[0039] In one embodiment of the system, where the primary storage for the application is a modern disk array or storage virtualization appliance, the Data Management Virtualization system takes advantage of a point-in-time snapshot capability of an underlying storage device to make the initial copy of the data. This virtual copy mechanism is a fast, efficient and low-impact technique of creating the initial copy that does not guarantee that all the bits will be copied, or stored together. Instead, virtual copies are constructed by maintaining metadata and data structures, such as copy-on-write volume bitmaps or extents, that allow the copies to be reconstructed at access time. The copy has a lightweight impact on the application and on the

primary storage device. In another embodiment, where the application is based on a Server Virtualization System such as VMware or Xen, the Data Management Virtualization system uses the similar virtual-machine-snapshot capability that is built into the Server Virtualization systems. When a virtual copy capability is not available, the Data Management Virtualization System may include its own built-in snapshot mechanism.

[0040] It is possible to use the snapshot as a data primitive underlying all of the data management functions supported by the system. Because it is lightweight, the snapshot can be used as an internal operation even when the requested operation is not a snapshot per se; it is created to enable and facilitate other operations.

[0041] At the time of creation of a snapshot, there may be certain preparatory operations involved in order to create a coherent snapshot or coherent image, such that the image may be restored to a state that is usable by the application. These preparatory operations need only be performed once, even if the snapshot will be leveraged across multiple data management functions in the system, such as backup copies which are scheduled according to a policy. The preparatory operations may include application quiescence, which includes flushing data caches and freezing the state of the application; it may also include other operations known in the art and other operations useful for retaining a complete image, such as collecting metadata information from the application to be stored with the image.

[0042] Figure 2 illustrates one way that a Virtualized Data Management system can address the data lifecycle requirements described earlier in accordance with these principles.

[0043] To serve local backup requirements, a sequence of efficient snapshots are made within local high-availability storage 202. Some of these snapshots are used to serve development/test requirements without making another copy. For longer term retention of local backup, a copy is made efficiently into long-term local storage 204, which in this implementation uses deduplication to reduce repeated copying. The copies within long-term storage may be accessed as backups or treated as an archive, depending on the retention policy applied by the SLA. A copy of the data is made to remote storage 206 in order to satisfy requirements for remote backup and business continuity – again a single set of copies suffices both purposes. As an alternative for remote backup and disaster recovery, a further copy of the

data may be made efficiently to a repository 208 hosted by a commercial or private cloud storage provider.

The Data Management Virtualization System

[0044] FIG. 3 illustrates the high level components of the Data Management Virtualization System that implements the above principles. Preferably, the system comprises these basic functional components further described below.

[0045] Application 300 creates and owns the data. This is the software system that has been deployed by the user, as for example, an email system, a database system, or financial reporting system, in order to satisfy some computational need. The Application typically runs on a server and utilizes storage. For illustrative purposes, only one application has been indicated. In reality there may be hundreds or even thousands of applications that are managed by a single Data Management Virtualization System.

[0046] Storage Resources 302 is where application data is stored through its lifecycle. The Storage Resources are the physical storage assets, including internal disk drives, disk arrays, optical and tape storage libraries and cloud-based storage systems that the user has acquired to address data storage requirements. The storage resources consist of Primary Storage 310, where the online, active copy of the application data is stored, and Secondary Storage 312 where additional copies of the application data are stored for the purposes such as backup, disaster recovery, archiving, indexing, reporting and other uses. Secondary storage resources may include additional storage within the same enclosure as the primary storage, as well as storage based on similar or different storage technologies within the same data center, another location or across the internet.

[0047] One or more Management Workstations 308 allow the user to specify a Service Level Agreement (SLA) 304 that defines the lifecycle for the application data. A Management workstation is a desktop or laptop computer or a mobile computing device that is used to configure, monitor and control the Data Management Virtualization System. A Service Level Agreement is a detailed specification that captures the detailed business requirements related to the creation, retention and deletion of secondary copies of the application data. The SLA is much more than the simple RTO and RPO that are used in traditional data management

applications to represent the frequency of copies and the anticipated restore time for a single class of secondary storage. The SLA captures the multiple stages in the data lifecycle specification, and allows for non uniform frequency and retention specifications within each class of secondary storage. The SLA is described in greater detail in FIG. 7.

[0048] Data Management Virtualization Engine 306 manages all of the lifecycle of the application data as specified in SLA. It manages potentially a large number of SLAs for a large number of applications. The Data Management Virtualization Engine takes inputs from the user through the Management Workstation and interacts with the applications to discover the applications primary storage resources. The Data Management Virtualization Engine makes decisions regarding what data needs to be protected and what secondary storage resources best fulfill the protection needs. For example, if an enterprise designates its accounting data as requiring copies to be made at very short intervals for business continuity purposes as well as for backup purposes, the Engine may decide to create copies of the accounting data at a short interval to a first storage pool, and to also create backup copies of the accounting data to a second storage pool at a longer interval, according to an appropriate set of SLAs. This is determined by the business requirements of the storage application.

[0049] The Engine then makes copies of application data using advanced capabilities of the storage resources as available. In the above example, the Engine may schedule the short-interval business continuity copy using a storage appliance's built-in virtual copy or snapshot capabilities. Data Management Virtualization Engine moves the application data amongst the storage resources in order to satisfy the business requirements that are captured in the SLA. The Data Management Virtualization Engine is described in greater detail in FIG. 4.

[0050] The Data Management Virtualization System as a whole may be deployed within a single host computer system or appliance, or it may be one logical entity but physically distributed across a network of general-purpose and purpose-built systems. Certain components of the system may also be deployed within a computing or storage cloud.

[0051] In one embodiment of the Data Management Virtualization System the Data Management Virtualization Engine largely runs as multiple processes on a fault tolerant, redundant pair of computers. Certain components of the Data Management Virtualization Engine may run close to the application within the application servers. Some other components

may run close to the primary and secondary storage, within the storage fabric or in the storage systems themselves. The Management stations are typically desktop and laptop computers and mobile devices that connect over a secure network to the Engine.

The Data Management Virtualization Engine

[0052] FIG. 4 illustrates an architectural overview of the Data Management Virtualization Engine 306 according to certain embodiments of the invention. The 306 Engine includes the following modules:

[0053] Application Specific Module 402: This module is responsible for controlling and collecting metadata from the application 300. Application metadata includes information about the application such as the type of application, details about its configuration, location of its datastores, its current operating state. Controlling the operation of the application includes actions such as flushing cached data to disk, freezing and thawing application I/O, rotating or truncating log files, and shutting down and restarting applications. The Application Specific module performs these operations and sends and receives metadata in responses to commands from the Service Level Policy Engine 406, described below. The Application Specific Module is described in more detail in connection with FIG. 8.

[0054] Service Level Policy Engine 406 acts on the SLA 304 provided by the user to make decisions regarding the creation, movement and deletion of copies of the application data. Each SLA describes the business requirements related to protection of one application. The Service Level Policy Engine analyzes each SLA and arrives at a series of actions each of which involve the copying of application data from one storage location to another. The Service Level Policy Engine then reviews these actions to determine priorities and dependencies, and schedules and initiates the data movement jobs. The Service Level Policy Engine is described in more detail in connection with FIG. 9.

[0055] Object Manager and Data Movement Engine 410 creates a composite object consisting of the Application data, the Application Metadata and the SLA which it moves through different storage pools per instruction from the Policy Engine. The Object Manager receives instructions from the Service Policy Engine 406 in the form of a command to create a copy of application data in a particular pool based on the live primary data 413 belonging to the

application 300, or from an existing copy, e.g., 415, in another pool. The copy of the composite object that is created by the Object Manager and the Data Movement Engine is self contained and self describing in that it contains not only application data, but also application metadata and the SLA for the application. The Object Manager and Data Movement Engine are described in more detail in connection with FIG. 5.

[0056] Storage Pool Manager 412 is a component that adapts and abstracts the underlying physical storage resources 302 and presents them as virtual storage pools 418. The physical storage resources are the actual storage assets, such as disk arrays and tape libraries that the user has deployed for the purpose of supporting the lifecycle of the data of the user's applications. These storage resources might be based on different storage technologies such as disk, tape, flash memory or optical storage. The storage resources may also have different geographic locations, cost and speed attributes, and may support different protocols. The role of the Storage Pool Manager is to combine and aggregate the storage resources, and mask the differences between their programming interfaces. The Storage Pool Manager presents the physical storage resources to the Object Manager 410 as a set of storage pools that have characteristics that make these pools suitable for particular stages in the lifecycle of application data. The Storage Pool Manager is described in more detail in connection with FIG. 6.

Object Manager and Data Movement Engine

[0057] FIG. 5 illustrates the Object Manager and Data Movement Engine 410. The Object Manager and Data Movement Engine discovers and uses Virtual Storage Resources 510 presented to it by the Pool Managers 504. It accepts requests from the Service Level Policy Engine 406 to create and maintain Data Storage Object instances from the resources in a Virtual Storage Pool, and it copies application data among instances of storage objects from the Virtual Storage Pools according to the instructions from the Service Level Policy Engine. The target pool selected for the copy implicitly designates the business operation being selected, e.g. backup, replication or restore. The Service Level Policy Engine resides either locally to the Object Manager (on the same system) or remotely, and communicates using a protocol over standard networking communication. TCP/IP may be used in a preferred embodiment, as it is well understood, widely available, and allows the Service Level Policy Engine to be located locally to the Object Manager or remotely with little modification.

[0058] In one embodiment, the system may deploy the Service Level Policy Engine on the same computer system as the Object Manager for ease of implementation. In another embodiment, the system may employ multiple systems, each hosting a subset of the components if beneficial or convenient for an application, without changing the design.

[0059] The Object Manager 501 and the Storage Pool Managers 504 are software components that may reside on the computer system platform that interconnects the storage resources and the computer systems that use those storage resources, where the user's application resides. The placement of these software components on the interconnect platform is designated as a preferred embodiment, and may provide the ability to connect customer systems to storage via communication protocols widely used for such applications (e.g. Fibre Channel, iSCSI, etc.), and may also provide ease of deployment of the various software components.

[0060] The Object Manager 501 and Storage Pool Manager 504 communicate with the underlying storage virtualization platform via the Application Programming Interfaces made available by the platform. These interfaces allow the software components to query and control the behavior of the computer system and how it interconnects the storage resources and the computer system where the user's Application resides. The components apply modularity techniques as is common within the practice to allow replacement of the intercommunication code particular to a given platform.

[0061] The Object Manager and Storage Pool Managers communicate via a protocol. These are transmitted over standard networking protocols, e.g. TCP/IP, or standard Interprocess Communication (IPC) mechanisms typically available on the computer system. This allows comparable communication between the components if they reside on the same computer platform or on multiple computer platforms connected by a network, depending on the particular computer platform. The current configuration has all of the local software components residing on the same computer system for ease of deployment. This is not a strict requirement of the design, as described above, and can be reconfigured in the future as needed.

Object Manager

[0062] Object Manager 501 is a software component for maintaining Data Storage Objects, and provides a set of protocol operations to control it. The operations include creation, destruction, duplication, and copying of data among the objects, maintaining access to objects, and in particular allow the specification of the storage pool used to create copies. There is no common subset of functions supported by all pools; however, in a preferred embodiment, primary pools may be performance-optimized, i.e. lower latency, whereas backup or replication pools may be capacity-optimized, supporting larger quantities of data and content-addressable. The pools may be remote or local. The storage pools are classified according to various criteria, including means by which a user may make a business decision, e.g. cost per gigabyte of storage.

[0063] First, the particular storage device from which the storage is drawn may be a consideration, as equipment is allocated for different business purposes, along with associated cost and other practical considerations. Some devices may not even be actual hardware but capacity provided as a service, and selection of such a resource can be done for practical business purposes.

[0064] Second, the network topological “proximity” is considered, as near storage is typically connected by low-latency, inexpensive network resources, while distant storage may be connected by high-latency, bandwidth limited expensive network resources; conversely, the distance of a storage pool relative to the source may be beneficial when geographic diversity protects against a physical disaster affecting local resources.

[0065] Third, storage optimization characteristics are considered, where some storage is optimized for space-efficient storage, but requires computation time and resources to analyze or transform the data before it can be stored, while other storage by comparison is “performance optimized,” taking more storage resources by comparison but using comparatively little computation time or resource to transform the data, if at all.

[0066] Fourth, “speed of access” characteristics are considered, where some resources intrinsic to a storage computer platform are readily and quickly made available to the user’s Application, e.g. as a virtual SCSI block device, while some can only be indirectly used. These

ease and speed of recovery is often governed by the kind of storage used, and this allows it to be suitably classified.

[0067] Fifth, the amount of storage used and the amount available in a given pool are considered, as there may be benefit to either concentrating or spreading the storage capacity used.

[0068] The Service Level Policy Engine, described below, combines the SLA provided by the user with the classification criteria to determine how and when to maintain the application data, and from which storage pools to draw the needed resources to meet the Service Level Agreement (SLA).

[0069] The object manager 501 creates, maintains and employs a history mechanism to track the series of operations performed on a data object within the performance pools, and to correlate those operations with others that move the object to other storage pools, in particular capacity-optimized ones. This series of records for each data object is maintained at the object manager for all data objects in the primary pool, initially correlated by primary data object, then correlated by operation order: a time line for each object and a list of all such time lines. Each operation performed exploits underlying virtualization primitives to capture the state of the data object at a given point in time.

[0070] Additionally, the underlying storage virtualization appliance may be modified to expose and allow retrieval of internal data structures, such as bitmaps, that indicate the modification of portions of the data within the data object. These data structures are exploited to capture the state of a data object at a point in time: e.g., a snapshot of the data object, and to provide differences between snapshots taken at a specific time, and thereby enables optimal backup and restore. While the particular implementations and data structures may vary among different appliances from different vendors, a data structure is employed to track changes to the data object, and storage is employed to retain the original state of those portions of the object that have changed: indications in the data structure correspond to data retained in the storage. When accessing the snapshot, the data structure is consulted and for portions that have been changed, the preserved data is accessed rather than the current data, as the data object has been modified at the areas so indicated. A typical data structure employed is a bitmap, where each bit corresponds to a section of the data object. Setting the bit indicates that section has been

modified after the point in time of the snapshot operation. The underlying snapshot primitive mechanism maintains this for as long as the snapshot object exists.

[0071] The time line described above maintains a list of the snapshot operations against a given primary data object, including the time an operation is started, the time it is stopped (if at all), a reference to the snapshot object, and a reference to the internal data structure (e.g. bitmaps or extent lists), so that it can be obtained from the underlying system. Also maintained is a reference to the result of copying the state of the data object at any given point in time into another pool – as an example, copying the state of a data object into a capacity-optimized pool 407 using content addressing results in an object handle. That object handle corresponds to a given snapshot and is stored with the snapshot operation in the time line. This correlation is used to identify suitable starting points.

[0072] Optimal backup and restore consult the list of operations from a desired starting point to an end point. A time ordered list of operations and their corresponding data structures (bitmaps) are constructed such that a continuous time series from start to finish is realized: there is no gap between start times of the operations in the series. This ensures that all changes to the data object are represented by the corresponding bitmap data structures. It is not necessary to retrieve all operations from start to finish; simultaneously existing data objects and underlying snapshots overlap in time; it is only necessary that there are no gaps in time where a change might have occurred that was not tracked. As bitmaps indicate that a certain block of storage has changed but not what the change is, the bitmaps may be added or composed together to realize a set of all changes that occurred in the time interval. Instead of using this data structure to access the state at a point in time, the system instead exploits the fact that the data structure represents data modified as time marches forward. Rather, the end state of the data object is accessed at the indicated areas, thus returning the set of changes to the given data object from the given start time to the end time.

[0073] The backup operation exploits this time line, the correlated references, and access to the internal data structures to realize our backup operation. Similarly, it uses the system in a complementary fashion to accomplish our restore operation. The specific steps are described below in the section for "Optimal Backup/Restore."

Virtual Storage Pool Types

[0074] FIG. 5 illustrates several representative storage pool types. Although one primary storage pool and two secondary storage pools are depicted in the figure, many more may be configured in some embodiments.

[0075] Primary Storage Pool 507 – contains the storage resources used to create the data objects in which the user Application stores its data. This is in contrast to the other storage pools, which exist to primarily fulfill the operation of the Data Management Virtualization Engine.

[0076] Performance Optimized Pool 508 – a virtual storage pool able to provide high performance backup (i.e. point in time duplication, described below) as well as rapid access to the backup image by the user Application

[0077] Capacity Optimized Pool 509 – a virtual storage pool that chiefly provides storage of a data object in a highly space-efficient manner by use of deduplication techniques described below. The virtual storage pool provides access to the copy of the data object, but does not do so with high performance as its chief aim, in contrast to the Performance Optimized pool above.

[0078] The initial deployments contain storage pools as described above, as a minimal operational set. The design fully expects multiple Pools of a variety of types, representing various combinations of the criteria illustrated above, and multiple Pool Managers as is convenient to represent all of the storage in future deployments. The tradeoffs illustrated above are typical of computer data storage systems.

[0079] From a practical point of view, these three pools represent a preferred embodiment, addressing most users requirements in a very simple way. Most users will find that if they have one pool of storage for urgent restore needs, which affords quick recovery, and one other pool that is low cost, so that a large number of images can be retained for a large period of time, almost all of the business requirements for data protection can be met with little compromise.

[0080] The format of data in each pool is dictated by the objectives and technology used within the pool. For example, the quick recovery pool is maintained in the form very similar to

the original data to minimize the translation required and to improve the speed of recovery. The long-term storage pool, on the other hand, uses deduplication and compression to reduce the size of the data and thus reduce the cost of storage.

Object Management Operations 505

[0081] The Object Manager 501 creates and maintains instances of Data Storage Objects 503 from the Virtual Storage Pools 418 according to the instructions sent to it by the Service Level Policy Engine 406. The Object Manager provides data object operations in five major areas: point-in-time duplication or copying (commonly referred to as “snapshots”), standard copying, object maintenance, mapping and access maintenance, and collections.

[0082] Object Management operations also include a series of Resource Discovery operations for maintaining Virtual Storage Pools themselves and retrieving information about them. The Pool Manager 504 ultimately supplies the functionality for these.

Point-In-Time Copy (“Snapshot”) Operations

[0083] Snapshot operations create a data object instance representing an initial object instance at a specific point in time. More specifically, a snapshot operation creates a complete virtual copy of the members of a collection using the resources of a specified Virtual Storage Pool. This is called a Data Storage Object. Multiple states of a Data Storage Object are maintained over time, such that the state of a Data Storage Object as it existed at a point in time is available. As described above, a virtual copy is a copy implemented using an underlying storage virtualization API that allows a copy to be created in a lightweight fashion, using copy-on-write or other in-band technologies instead of copying and storing all bits of duplicate data to disk. This may be implemented using software modules written to access the capabilities of an off-the-shelf underlying storage virtualization system such as provided by EMC, vmware or IBM in some embodiments. Where such underlying virtualizations are not available, the described system may provide its own virtualization layer for interfacing with unintelligent hardware.

[0084] Snapshot operations require the application to freeze the state of the data to a specific point so that the image data is coherent, and so that the snapshot may later be used to restore the state of the application at the time of the snapshot. Other preparatory steps may

also be required. These are handled by the Application-Specific Module 302, which is described in a subsequent section. For live applications, therefore, the most lightweight operations are desired.

[0085] Snapshot operations are used as the data primitive for all higher-level operations in the system. In effect, they provide access to the state of the data at a particular point in time. As well, since snapshots are typically implemented using copy-on-write techniques that distinguish what has changed from what is resident on disk, these snapshots provide differences that can also be composed or added together to efficiently copy data throughout the system. The format of the snapshot may be the format of data that is copied by Data Mover 502, which is described below.

Standard Copy Operations

[0086] When a copy operation is not a snapshot, it may be considered a standard copy operation. A standard copy operation copies all or a subset of a source data object in one storage pool to a data object in another storage pool. The result is two distinct objects. One type of standard copy operation that may be used is an initial “baseline” copy. This is typically done when data is initially copied from one Virtual Storage Pool into another, such as from a performance-optimized pool to a capacity-optimized storage pool. Another type of standard copy operation may be used wherein only changed data or differences are copied to a target storage pool to update the target object. This would occur after an initial baseline copy has previously been performed.

[0087] A complete exhaustive version of an object need not be preserved in the system each time a copy is made, even though a baseline copy is needed when the Data Virtualization System is first initialized. This is because each virtual copy provides access to a complete copy. Any delta or difference can be expressed in relation to a virtual copy instead of in relation to a baseline. This has the positive side effect of virtually eliminating the common step of walking through a series of change lists.

[0088] Standard copy operations are initiated by a series of instructions or requests supplied by the Pool Manager and received by the Data Mover to cause the movement of data among the Data Storage Objects, and to maintain the Data Storage Objects themselves. The

copy operations allow the creation of copies of the specified Data Storage Objects using the resources of a specified Virtual Storage Pool. The result is a copy of the source Data Object in a target Data Object in the storage pool.

[0089] The Snapshot and Copy operations are each structured with a preparation operation and an activation operation. The two steps of prepare and activate allow the long-running resource allocation operations, typical of the prepare phase, to be decoupled from the actuation. This is required by applications that can only be paused for a short while to fulfill the point-in-time characteristics of a snapshot operation, which in reality takes a finite but non-zero amount of time to accomplish. Similarly for copy and snapshot operations, this two-step preparation and activation structure allows the Policy Engine to proceed with an operation only if resources for all of the collection members can be allocated.

Object Maintenance

[0090] Object Maintenance operations are a series of operations for maintaining data objects, including creation, destruction, and duplication. The Object Manager and Data Mover use functionality provided by a Pool Request Broker (more below) to implement these operations. The data objects may be maintained at a global level, at each Storage Pool, or preferably both.

Collections

[0091] Collection operations are auxiliary functions. Collections are abstract software concepts, lists maintained in memory by the object manager. They allow the Policy Engine 206 to request a series of operations over all of the members in a collection, allowing a consistent application of a request to all members. The use of collections allows for simultaneous activation of the point-in-time snapshot so that multiple Data Storage Objects are all captured at precisely the same point in time, as this is typically required by the application for a logically correct restore. The use of collections allows for convenient request of a copy operation across all members of a collection, where an application would use multiple storage objects as a logical whole.

Resource Discovery Operations

[0092] The Object Manager discovers Virtual Storage Pools by issuing Object Management Operations 505 to the Pool Manager 504, and uses the information obtained about each of the pools to select one that meets the required criteria for a given request, or in the case where none match, a default pool is selected, and the Object Manager can then create a data storage object using resources from the selected Virtual Storage Pool.

Mapping and Access

[0093] The Object Manager also provides sets of Object Management operations to allow and maintain the availability of these objects to external Applications. The first set is operations for registering and unregistering the computers where the user's Applications reside. The computers are registered by the identities typical to the storage network in use (e.g. Fibre Channel WWPN, iSCSI identity, etc.). The second set is "mapping" operations, and when permitted by the storage pool from which an object is created, the Data Storage Object can be "mapped," that is, made available for use to a computer on which a user Application resides.

[0094] This availability takes a form appropriate to the storage, e.g. a block device presented on a SAN as a Fibre Channel disk or iSCSI device on a network, a filesystem on a file sharing network, etc. and is usable by the operating system on the Application computer. Similarly, an "unmapping" operation reverses the availability of the virtual storage device on the network to a user Application. In this way, data stored for one Application, i.e. a backup, can be made available to another Application on another computer at a later time, i.e. a restore.

502 Data Mover

[0095] The Data Mover 502 is a software component within the Object Manager and Data Mover that reads and writes data among the various Data Storage Objects 503 according to instructions received from the Object Manager for Snapshot (Point in Time) Copy requests and standard copy requests. The Data Mover provides operations for reading and writing data among instances of data objects throughout the system. The Data Mover also provides operations that allow querying and maintaining the state of long running operations that the Object Manager has requested for it to perform.

[0096] The Data Mover uses functionality from the Pool Functionality Providers (see FIG. 6) to accomplish its operation. The Snapshot functionality provider 608 allows creation of a data object instance representing an initial object instance at a specific point in time. The Difference Engine functionality provider 614 is used to request a description of the differences between two data objects that are related in a temporal chain. For data objects stored on content-addressable pools, a special functionality is provided that can provide differences between any two arbitrary data objects. This functionality is also provided for performance-optimized pools, in some cases by an underlying storage virtualization system, and in other cases by a module that implements this on top of commodity storage. The Data Mover 502 uses the information about the differences to select the set of data that it copies between instances of data objects 503.

[0097] For a given Pool, the Difference Engine Provider provides a specific representation of the differences between two states of a Data Storage Object over time. For a Snapshot provider the changes between two points in time are recorded as writes to a given part of the Data Storage Object. In one embodiment, the difference is represented as a bitmap where each bit corresponds to an ordered list of the Data Object areas, starting at the first and ascending in order to the last, where a set bit indicates a modified area. This bitmap is derived from the copy-on-write bitmaps used by the underlying storage virtualization system. In another embodiment, the difference may be represented as a list of extents corresponding to changed areas of data. For a Content Addressable storage provider 610, the representation is described below, and is used to determine efficiently the parts of two Content Addressable Data Objects that differ.

[0098] The Data Mover uses this information to copy only those sections that differ, so that a new version of a Data Object can be created from an existing version by first duplicating it, obtaining the list of differences, and then moving only the data corresponding to those differences in the list. The Data Mover 502 traverses the list of differences, moving the indicated areas from the source Data Object to the target Data Object. (See Optimal Way for Data Backup and Restore.)

506 Copy Operation - Request Translation and Instructions

[0099] The Object Manager 501 instructs the Data Mover 502 through a series of operations to copy data among the data objects in the Virtual Storage Pools 418. The procedure comprises the following steps, starting at the reception of instructions:

[0100] First, create Collection request. A name for the collection is returned.

[0101] Second, add Object to Collection. The collection name from above is used as well as the name of the source Data Object that is to be copied and the name of two antecedents: a Data Object against which differences are to be taken in the source Storage Resource Pool, and a corresponding Data Object in the target Storage Resource Pool. This step is repeated for each source Data Object to be operated on in this set.

[0102] Third, prepare Copy Request. The collection name is supplied as well as a Storage Resource Pool to act as a target. The prepare command instructs the Object Manager to contact the Storage Pool Manager to create the necessary target Data Objects, corresponding to each of the sources in the collection. The prepare command also supplies the corresponding Data Object in the target Storage Resource Pool to be duplicated, so the Provider can duplicate the provided object and use that as a target object. A reference name for the copy request is returned.

[0103] Fourth, activate Copy Request. The reference name for the copy request returned above is supplied. The Data Mover is instructed to copy a given source object to its corresponding target object. Each request includes a reference name as well as a sequence number to describe the overall job (the entire set of source target pairs) as well as a sequence number to describe each individual source-target pair. In addition to the source-target pair, the names of the corresponding antecedents are supplied as part of the Copy instruction.

[0104] Fifth, the Copy Engine uses the name of the Data Object in the source pool to obtain the differences between the antecedent and the source from the Difference Engine at the source. The indicated differences are then transmitted from the source to the target. In one embodiment, these differences are transmitted as bitmaps and data. In another embodiment, these differences are transmitted as extent lists and data.

503 Data Storage Objects

[0105] Data Storage Objects are software constructs that permit the storage and retrieval of Application data using idioms and methods familiar to computer data processing equipment and software. In practice these currently take the form of a SCSI block device on a storage network, e.g. a SCSI LUN, or a content-addressable container, where a designator for the content is constructed from and uniquely identifies the data therein. Data Storage Objects are created and maintained by issuing instructions to the Pool Manager. The actual storage for persisting the Application data is drawn from the Virtual Storage Pool from which the Data Storage Object is created.

[0106] The structure of the data storage object varies depending on the storage pool from which it is created. For the objects that take the form of a block device on a storage network, the data structure for a given block device Data Object implements a mapping between the Logical Block Address (LBA) of each of the blocks within the Data Object to the device identifier and LBA of the actual storage location. The identifier of the Data Object is used to identify the set of mappings to be used. The current embodiment relies on the services provided by the underlying physical computer platform to implement this mapping, and relies on its internal data structures, such as bitmaps or extent lists.

[0107] For objects that take the form of a Content Addressable Container, the content signature is used as the identifier, and the Data Object is stored as is described below in the section about deduplication.

504 Pool Manager

[0108] A Pool Manager 504 is a software component for managing virtual storage resources and the associated functionality and characteristics as described below. The Object manager 501 and Data Movement Engine 502 communicate with one or more Pool Managers 504 to maintain Data Storage Objects 503.

510 Virtual Storage Resources

[0109] Virtual Storage Resources 510 are various kinds of storage made available to the Pool Manager for implementing storage pool functions, as described below. In this

embodiment, a storage virtualizer is used to present various external Fibre Channel or iSCSI storage LUNs as virtualized storage to the Pool Manager 504.

The Storage Pool Manager

[0110] FIG. 6 further illustrates the Storage Pool Manager 504. The purpose of the storage pool manager is to present underlying virtual storage resources to the Object Manager/Data Mover as Storage Resource Pools, which are abstractions of storage and data management functionality with common interfaces that are utilized by other components of the system. These common interfaces typically include a mechanism for identifying and addressing data objects associated with a specific temporal state, and a mechanism for producing differences between data objects in the form of bitmaps or extents. In this embodiment, the pool manager presents a Primary Storage Pool, a Performance Optimized Pool, and a Capacity Optimized Pool. The common interfaces allow the object manager to create and delete Data Storage objects in these pools, either as copies of other data storage objects or as new objects, and the data mover can move data between data storage objects, and can use the results of data object differencing operations.

[0111] The storage pool manager has a typical architecture for implementing a common interface to diverse implementations of similar functionality, where some functionality is provided by “smart” underlying resources, and other functionality must be implemented on top of less functional underlying resources.

[0112] Pool request broker 602 and pool functionality providers 604 are software modules executing in either the same process as the Object Manager/Data Mover, or in another process communicating via a local or network protocol such as TCP. In this embodiment the providers comprise a Primary Storage provider 606, Snapshot provider 608, Content Addressable provider 610, and Difference Engine provider 614, and these are further described below. In another embodiment the set of providers may be a superset of those shown here.

[0113] Virtual Storage Resources 510 are the different kinds of storage made available to the Pool Manager for implementing storage pool functions. In this embodiment, the virtual storage resources comprise sets of SCSI logical units from a storage virtualization system that runs on the same hardware as the pool manager, and accessible (for both data and management

operations) through a programmatic interface: in addition to standard block storage functionality additional capabilities are available including creating and deleting snapshots, and tracking changed portions of volumes. In another embodiment the virtual resources can be from an external storage system that exposes similar capabilities, or may differ in interface (for example accessed through a file-system, or through a network interface such as CIFS, iSCSI or CDMI), in capability (for example, whether the resource supports an operation to make a copy-on-write snapshot), or in non-functional aspects (for example, high-speed/limited-capacity such as Solid State Disk versus low-speed/high-capacity such as SATA disk). The capabilities and interface available determine which providers can consume the virtual storage resources, and which pool functionality needs to be implemented within the pool manager by one or more providers: for example, this implementation of a content addressable storage provider only requires “dumb” storage, and the implementation is entirely within content addressable provider 610; an underlying content addressable virtual storage resource could be used instead with a simpler “pass-through” provider. Conversely, this implementation of a snapshot provider is mostly “pass-through” and requires storage that exposes a quick point-in-time copy operation.

[0114] Pool Request Broker 602 is a simple software component that services requests for storage pool specific functions by executing an appropriate set of pool functionality providers against the configured virtual storage resource 510. The requests that can be serviced include, but are not limited to, creating an object in a pool; deleting an object from a pool; writing data to an object; reading data from an object; copying an object within a pool; copying an object between pools; requesting a summary of the differences between two objects in a pool.

[0115] Primary storage provider 606 enables management interfaces (for example, creating and deleting snapshots, and tracking changed portions of files) to a virtual storage resource that is also exposed directly to applications via an interface such as fibre channel, iSCSI, NFS or CIFS.

[0116] Snapshot provider 608 implements the function of making a point-in-time copy of data from a Primary resource pool. This creates the abstraction of another resource pool populated with snapshots. As implemented, the point-in-time copy is a copy-on-write snapshot of the object from the primary resource pool, consuming a second virtual storage resource to

accommodate the copy-on-write copies, since this management functionality is exposed by the virtual storage resources used for primary storage and for the snapshot provider.

[0117] Difference engine provider 614 can satisfy a request for two objects in a pool to be compared that are connected in a temporal chain. The difference sections between the two objects are identified and summarized in a provider-specific way, e.g. using bitmaps or extents. For example, the difference sections might be represented as a bitmap where each set bit denotes a fixed size region where the two objects differ; or the differences might be represented procedurally as a series of function calls or callbacks.

[0118] Depending on the virtual storage resource on which the pool is based, or on other providers implementing the pool, a difference engine may produce a result efficiently in various ways. As implemented, a difference engine acting on a pool implemented via a snapshot provider uses the copy-on-write nature of the snapshot provider to track changes to objects that have had snapshots made. Consecutive snapshots of a single changing primary object thus have a record of the differences that is stored alongside them by the snapshot provider, and the difference engine for snapshot pools simply retrieves this record of change. Also as implemented, a difference engine acting on a pool implemented via a Content Addressable provider uses the efficient tree structure (see below, Fig. 12) of the content addressable implementation to do rapid comparisons between objects on demand.

[0119] Content addressable provider 610 implements a write-once content addressable interface to the virtual storage resource it consumes. It satisfies read, write, duplicate and delete operations. Each written or copied object is identified by a unique handle that is derived from its content. The content addressable provider is described further below (FIG. 11).

Pool Manager Operations

[0120] In operation, the pool request broker 502 accepts requests for data manipulation operations such as copy, snapshot, or delete on a pool or object. The request broker determines which provider code from pool 504 to execute by looking at the name or reference to the pool or object. The broker then translates the incoming service request into a form that can be handled by the specific pool functionality provider, and invokes the appropriate sequence of provider operations.

[0121] For example, an incoming request could ask to make a snapshot from a volume in a primary storage pool, into a snapshot pool. The incoming request identifies the object (volume) in the primary storage pool by name, and the combination of name and operation (snapshot) determines that the snapshot provider should be invoked which can make point-in-time snapshots from the primary pool using the underlying snapshot capability. This snapshot provider will translate the request into the exact form required by the native copy-on-write function performed by the underlying storage virtualization appliance, such as bitmaps or extents, and it will translate the result of the native copy-on-write function to a storage volume handle that can be returned to the object manager and used in future requests to the pool manager.

Optimal Way for Data Backup Using the Object Manager and Data Mover

[0122] Optimal Way for Data Backup is a series of operations to make successive versions of Application Data objects over time, while minimizing the amount of data that must be copied by using bitmaps, extents and other temporal difference information stored at the Object Mover. It stores the application data in a data storage object and associates with it the metadata that relates the various changes to the application data over time, such that changes over time can be readily identified.

[0123] In a preferred embodiment, the procedure comprises the following steps:

1. The mechanism provides an initial reference state, e.g. T0, of the Application Data within a Data Storage Object;
2. Subsequent instances (versions) are created on demand over time of the Data Storage Object in a Virtual Storage Pool that has a Difference Engine Provider;
3. Each successive version, e.g. T4, T5, uses the Difference Engine Provider for the Virtual Storage Pool to obtain the difference between it and the instance created prior to it, so that T5 is stored as a reference to T4 and a set of differences between T5 and T4;
4. The Copy Engine receives a request to copy data from one data object (the source) to another data object (the destination);

5. If the Virtual Storage Pool in which the destination object will be created contains no other objects created from prior versions of the source data object, then a new object is created in the destination Virtual Storage Pool and the entire contents of the source data object are copied to the destination object; the procedure is complete. Otherwise the next steps are followed;
6. If the Virtual Storage Pool in which the destination object is created contains objects created from prior versions of the source data object, a recently created prior version in the destination Virtual Storage Pool is selected for which there exists a corresponding prior version in the Virtual Storage Pool of the source data object. For example, if a copy of T5 is initiated from a snapshot pool, and an object created at time T3 is the most recent version available at the target, T3 is selected as the prior version;
7. Construct a time-ordered list of the versions of the source data object, beginning with an initial version identified in the previous step, and ending with the source data object that is about to be copied. In the above example, at the snapshot pool, all states of the object are available, but only the states including and following T3 are of interest: T3, T4, T5;
8. Construct a corresponding list of the differences between each successive version in the list such that all of the differences, from the beginning version of the list to the end are represented. Difference both, identify which portion of data has changed and includes the new data for the corresponding time. This creates a set of differences from the target version to the source version, e.g. the difference between T3 and T5;
9. Create the destination object by duplicating the prior version of the object identified in Step 6 in the destination Virtual Storage Pool, e.g. object T3 in the target store;
10. Copy the set of differences identified in the list created in Step 8 from the source data object to the destination object; the procedure is complete.

[0124] Each data object within the destination Virtual Storage Pool is complete; that is, it represents the entire data object and allows access to the all of the Application Data at the point in time without requiring external reference to state or representations at other points in time. The object is accessible without replaying all deltas from a baseline state to the present state.

Furthermore, the duplication of initial and subsequent versions of the data object in the destination Virtual Storage Pool does not require exhaustive duplication of the Application Data contents therein. Finally, to arrive at second and subsequent states requires only the transmission of the changes tracked and maintained, as described above, without exhaustive traversal, transmission or replication of the contents of the data storage object.

Optimal Way for Data Restore Using the Object Manager and Data Mover

[0125] Intuitively, the operation of the Optimal Way for Data Restore is the converse of the Optimal Way for Data Backup. The procedure to recreate the desired state of a data object in a destination Virtual Storage Pool at a given point in time comprises the following steps:

1. Identify a version of the data object in another Virtual Storage Pool that has a Difference Engine Provider, corresponding to the desired state to be recreated. This is the source data object in the source Virtual Storage Pool;
2. Identify a preceding version of the data object to be recreated in the destination Virtual Storage Pool;
3. If no version of the data object is identified in Step 2, then create a new destination object in the destination Virtual Storage Pool and copy the data from the source data object to the destination data object. The procedure is complete. Otherwise, proceed with the following steps;
4. If a version of the data object is identified in Step 2, then identify a data object in the source Virtual Storage Pool corresponding to the data object identified in Step 2;
5. If no data object is identified in Step 4, then create a new destination object in the destination Virtual Storage Pool and copy the data from the source data object to the destination data object. The procedure is complete. Otherwise, proceed with the following steps;
6. Create a new destination data object in the Destination Virtual Storage Pool by duplicating the data object identified in Step 2;

7. Employ the Difference Engine Provider for the source Virtual Storage Pool to obtain the set of differences between the data object identified in Step 1 and the data object identified in Step 4;
8. Copy the data identified by the list created in Step 7 from the source data object to the destination data object. The procedure is complete.

[0126] Access to the desired state is complete: it does not require external reference to other containers or other states. Establishing the desired state given a reference state requires neither exhaustive traversal nor exhaustive transmission, only the retrieved changes indicated by the provided representations within the source Virtual Storage Pool.

The Service Level Agreement

[0127] FIG. 7 illustrates the Service Level Agreement. The Service Level Agreement captures the detailed business requirements with respect to secondary copies of the application data. In the simplest description, the business requirements define when and how often copies are created, how long they are retained and in what type of storage pools these copies reside. This simplistic description does not capture several aspects of the business requirements. The frequency of copy creation for a given type of pool may not be uniform across all hours of the day or across all days of a week. Certain hours of the day, or certain days of a week or month may represent more (or less) critical periods in the application data, and thus may call for more (or less) frequent copies. Similarly, all copies of application data in a particular pool may not be required to be retained for the same length of time. For example, a copy of the application data created at the end of monthly processing may need to be retained for a longer period of time than a copy in the same storage pool created in the middle of a month.

[0128] The Service Level Agreement 304 of certain embodiments has been designed to represent all of these complexities that exist in the business requirements. The Service Level Agreement has four primary parts: the name, the description, the housekeeping attributes and a collection of Service Level Policies. As mentioned above, there is one SLA per application.

[0129] The name attribute 701 allows each Service Level Agreement to have a unique name.

[0130] The description attribute 702 is where the user can assign a helpful description for the Service Level Agreement.

[0131] The Service Level agreement also has a number of housekeeping attributes 703 that enable it to be maintained and revised. These attributes include but are not limited to the owner's identity, the dates and times of creation, modification and access, priority, enable/disable flags.

[0132] The Service Level Agreement also contains a plurality of Service Level Policies 705. Some Service level Agreements may have just a single Service Level Policy. More typically, a single SLA may contain tens of policies.

[0133] Each Service Level Policy consists of at least the following, in certain embodiments: the source storage pool location 706 and type 708; the target storage pool location 710 and type 712; the frequency for the creation of copies 714, expressed as a period of time; the length of retention of the copy 716, expressed as a period of time; the hours of operation 718 during the day for this particular Service Level Policy; and the days of the week, month or year 720 on which this Service Level Policy applies.

[0134] Each Service Level Policy specifies a source and target storage pool, and the frequency of copies of application data that are desired between those storage pools. Furthermore, the Service Level Policy specifies its hours of operation and days on which it is applicable. Each Service Level Policy is the representation of one single statement in the business requirements for the protection of application data. For example, if a particular application has a business requirement for an archive copy to be created each month after the monthly close and retained for three years, this might translate to a Service level Policy that requires a copy from the Local Backup Storage Pool into the Long-term Archive Storage Pool at midnight on the last day of the month, with a retention of three years.

[0135] All of the Service Level Policies with a particular combination of source and destination pool and location, say for example, source Primary Storage pool and destination local Snapshot pool, when taken together, specify the business requirements for creating copies into that particular destination pool. Business requirements may dictate for example that snapshot copies be created every hour during regular working hours, but only once every four

hours outside of these times. Two Service Level Policies with the same source and target storage pools will effectively capture these requirements in a form that can be put into practice by the Service Policy Engine.

[0136] This form of a Service Level Agreement allows the representation of the schedule of daily, weekly and monthly business activities, and thus captures business requirements for protecting and managing application data much more accurately than traditional RPO and RPO based schemes. By allowing hour of operation and days, weeks, and months of the year, scheduling can occur on a “calendar basis.”

[0137] Taken together, all of the Service Level Policies with one particular combination of source and destinations, for example, "source: local primary and destination: local performance optimized", captures the non-uniform data protection requirements for one type of storage. A single RPO number, on the other hand, forces a single uniform frequency of data protection across all times of day and all days. For example, a combination of Service Level Policies may require a large number of snapshots to be preserved for a short time, such as 10 minutes, and a lesser number of snapshots to be preserved for a longer time, such as 8 hours; this allows a small amount of information that has been accidentally deleted can be reverted to a state not more than 10 minutes before, while still providing substantial data protection at longer time horizons without requiring the storage overhead of storing all snapshots taken every ten minutes. As another example, the backup data protection function may be given one Policy that operates with one frequency during the work week, and another frequency during the weekend.

[0138] When Service Level Policies for all of the different classes of source and destination storage are included, the Service Level Agreement fully captures all of the data protection requirements for the entire application, including local snapshots, local long duration stores, off-site storage, archives, etc. A collection of policies within a SLA is capable of expressing when a given function should be performed, and is capable of expressing multiple data management functions that should be performed on a given source of data.

[0139] Service Level Agreements are created and modified by the user through a user interface on a management workstation. These agreements are electronic documents stored by the Service Policy Engine in a structured SQL database or other repository that it manages.

The policies are retrieved, electronically analyzed, and acted upon by the Service Policy Engine through its normal scheduling algorithm as described below.

[0140] FIG. 8 illustrates the Application Specific Module 402. The Application Specific module runs close to the Application 300 (as described above), and interacts with the Application and its operating environment to gather metadata and to query and control the Application as required for data management operations.

[0141] The Application Specific Module interacts with various components of the application and its operating environment including Application Service Processes and Daemons 801, Application Configuration Data 802, Operating System Storage Services 803 (such as VSS and VDS on Windows), Logical Volume Management and Filesystem Services 804, and Operating System Drivers and Modules 805.

[0142] The Application Specific Module performs these operations in response to control commands from the Service Policy Engine 406. There are two purposes for these interactions with the application: Metadata Collection and Application Consistency.

[0143] Metadata Collection is the process by which the Application Specific Module collects metadata about the application. In some embodiments, metadata includes information such as: configuration parameters for the application; state and status of the application; control files and startup/shutdown scripts for the application; location of the datafiles, journal and transaction logs for the application; and symbolic links, filesystem mount points, logical volume names, and other such entities that can affect the access to application data.

[0144] Metadata is collected and saved along with application data and SLA information. This guarantees that each copy of application data within the system is self contained and includes all of the details required to rebuild the application data.

[0145] Application Consistency is the set of actions that ensure that when a copy of the application data is created, the copy is valid, and can be restored into a valid instance of the application. This is critical when the business requirements dictate that the application be protected while it is live, in its online, operational state. The application may have interdependent data relations within its data stores, and if these are not copied in a consistent state will not provide a valid restorable image.

[0146] The exact process of achieving application consistency varies from application to application. Some applications have a simple flush command that forces cached data to disk. Some applications support a hot backup mode where the application ensures that its operations are journalled in a manner that guarantees consistency even as application data is changing. Some applications require interactions with operating system storage services such as VSS and VDS to ensure consistency. The Application Specific Module is purpose-built to work with a particular application and to ensure the consistency of that application. The Application Specific Module interacts with the underlying storage virtualization device and the Object Manager to provide consistent snapshots of application data.

[0147] For efficiency, the preferred embodiment of the Application Specific Module 402 is to run on the same server as Application 300. This assures the minimum latency in the interactions with the application, and provides access to storage services and filesystems on the application host. The application host is typically considered primary storage, which is then snapshotted to a performance-optimized store.

[0148] In order to minimize interruption of a running application, including minimizing preparatory steps, the Application Specific Module is only triggered to make a snapshot when access to application data is required at a specific time, and when a snapshot for that time does not exist elsewhere in the system, as tracked by the Object Manager. By tracking which times snapshots have been made, the Object Manager is able to fulfill subsequent data requests from the performance-optimized data store, including for satisfying multiple requests for backup and replication which may issue from secondary, capacity-optimized pools. The Object Manager may be able to provide object handles to the snapshot in the performance-optimized store, and may direct the performance-optimized store in a native format that is specific to the format of the snapshot, which is dependent on the underlying storage appliance. In some embodiments this format may be application data combined with one or more LUN bitmaps indicating which blocks have changed; in other embodiments it may be specific extents. The format used for data transfer is thus able to transfer only a delta or difference between two snapshots using bitmaps or extents.

[0149] Metadata, such as the version number of the application, may also be stored for each application along with the snapshot. When a SLA policy is executed, application metadata is

read and used for the policy. This metadata is stored along with the data objects. For each SLA, application metadata will only be read once during the lightweight snapshot operation, and preparatory operations which occur at that time such as flushing caches will only be performed once during the lightweight snapshot operation, even though this copy of application data along with its metadata may be used for multiple data management functions.

The Service Policy Engine

[0150] FIG. 9 illustrates the Service Policy Engine 406. The Service Policy Engine contains the Service Policy Scheduler 902, which examines all of the Service Level Agreements configured by the user and makes scheduling decisions to satisfy Service Level Agreements. It relies on several data stores to capture information and persist it over time, including, in some embodiments, a SLA Store 904, where configured Service Level Agreements are persisted and updated; a Resource Profile Store 906, storing Resource Profiles that provide a mapping between logical storage pool names and actual storage pools; Protection Catalog Store 908, where information is cataloged about previous successful copies created in various pools that have not yet expired; and centralized History Store 910.

[0151] History Store 910 is where historical information about past activities is saved for the use of all data management applications, including the timestamp, order and hierarchy of previous copies of each application into various storage pools. For example, a snapshot copy from a primary data store to a capacity-optimized data store that is initiated at 1 P.M. and is scheduled to expire at 9 P.M. will be recorded in History Store 910 in a temporal data store that also includes linked object data for snapshots for the same source and target that have taken place at 11 A.M. and 12 P.M.

[0152] These stores are managed by the Service Policy Engine. For example, when the user, through the Management workstation creates a Service Level Agreement, or modifies one of the policies within it, it is the Service Policy Engine that persists this new SLA in its store, and reacts to this modification by scheduling copies as dictated by the SLA. Similarly, when the Service Policy Engine successfully completes a data movement job that results in a new copy of an application in a Storage Pool, the Storage Policy Engine updates the History Store, so that this copy will be factored into future decisions.

[0153] The preferred embodiment of the various stores used by the Service Policy Engine is in the form of tables in a relational database management system in close proximity to the Service Policy Engine. This ensures consistent transactional semantics when querying and updating the stores, and allows for flexibility in retrieving interdependent data.

[0154] The scheduling algorithm for the Service Policy Scheduler 902 is illustrated in FIG. 10. When the Service Policy Scheduler decides it needs to make a copy of application data from one storage pool to another, it initiates a Data Movement Requestor and Monitor task, 912. These tasks are not recurring tasks and terminate when they are completed. Depending on the way that Service Level Policies are specified, a plurality of these requestors might be operational at the same time.

[0155] The Service Policy Scheduler considers the priorities of Service Level Agreements when determining which additional tasks to undertake. For example, if one Service Level Agreement has a high priority because it specifies the protection for a mission-critical application, whereas another SLA has a lower priority because it specifies the protection for a test database, then the Service Policy Engine may choose to run only the protection for the mission-critical application, and may postpone or even entirely skip the protection for the lower priority application. This is accomplished by the Service Policy Engine scheduling a higher priority SLA ahead of a lower priority SLA. In the preferred embodiment, in such a situation, for auditing purposes, the Service Policy Engine will also trigger a notification event to the management workstation.

The Policy Scheduling Algorithm

[0156] FIG. 10 illustrates the flowchart of the Policy Schedule Engine. The Policy Schedule Engine continuously cycles through all the SLAs defined. When it gets to the end of all of the SLAs, it sleeps for a short while, e.g. 10 seconds, and resumes looking through the SLAs again. Each SLA encapsulates the complete data protection business requirements for one application; thus all of the SLAs represent all of the applications.

[0157] For each SLA, the schedule engine collects together all of the Service Level Policies that have the same source pool and destination pool 1004 the process state at 1000 and iterates to the next SLA in the set of SLAs in 1002. Taken together, this subset of the Service Level

Policies represent all of the requirements for a copy from that source storage pool to that particular destination storage pool.

[0158] Among this subset of Service Level Policies, the Service Policy Scheduler discards the policies that are not applicable to today, or are outside their hours of operation. Among the policies that are left, find the policy that has the shortest frequency 1006, and based on the history data and in history store 910, the one with the longest retention that needs to be run next 1008.

[0159] Next, there are a series of checks 1010-1014 which rule out making a new copy of application data at this time – because the new copy is not yet due, because a copy is already in progress or because there is not new data to copy. If any of these conditions apply, the Service Policy Scheduler moves to the next combination of source and destination pools 1004. If none of these conditions apply, a new copy is initiated. The copy is executed as specified in the corresponding service level policy within this SLA 1016.

[0160] Next, the Scheduler moves to the next Source and Destination pool combination for the same Service Level agreement 1018. If there are no more distinct combinations, the Scheduler moves on to the next Service Level Agreement 1020.

[0161] After the Service Policy Scheduler has been through all source/destination pool combinations of all Service Level Agreements, it pauses for a short period and then resumes the cycle.

[0162] A simple example system with a snapshot store and a backup store, with only 2 policies defined, would interact with the Service Policy Scheduler as follows. Given two policies, one stating “backup every hour, the backup to be kept for 4 hours” and another stating “backup every 2 hours, the backup to be kept for 8 hours,” the result would be a single snapshot taken each hour, the snapshots each being copied to the backup store but retained a different amount of time at both the snapshot store and the backup store. The “backup every 2 hours” policy is scheduled to go into effect at 12:00 P.M by the system administrator.

[0163] At 4:00 P.M., when the Service Policy Scheduler begins operating at step 1000, it finds the two policies at step 1002. (Both policies apply because a multiple of two hours has elapsed since 12:00 P.M.) There is only one source and destination pool combination at step

1004. There are two frequencies at step 1006, and the system selects the 1-hour frequency because it is shorter than the 2-hour frequency. There are two operations with different retentions at step 1008, and the system selects the operation with the 8-hour retention, as it has the longer retention value. Instead of one copy being made to satisfy the 4-hour requirement and another copy being made to satisfy the 8-hour requirement, the two requirements are coalesced into the longer 8-hour requirement, and are satisfied by a single snapshot copy operation. The system determines that a copy is due at step 1010, and checks the relevant objects at the History Store 910 to determine if the copy has already been made at the target (at step 912) and at the source (at step 914). If these checks are passed, the system initiates the copy at step 916, and in the process triggers a snapshot to be made and saved at the snapshot store. The snapshot is then copied from the snapshot store to the backup store. The system then goes to sleep 1022 and wakes up again after a short period, such as 10 seconds. The result is a copy at the backup store and a copy at the snapshot store, where every even-hour snapshot lasts for 8 hours, and every odd-hour snapshot lasts 4 hours. The even-hour snapshots at the backup store and the snapshot store are both tagged with the retention period of 8 hours, and will be automatically deleted from the system by another process at that time.

[0164] Note that there is no reason to take two snapshots or make two backup copies at 2 o'clock, even though both policies apply, because both policies are satisfied by a single copy. Combining and coalescing these snapshots results in the reduction of unneeded operations., while retaining the flexibility of multiple separate policies. As well, it may be helpful to have two policies active at the same time for the same target with different retention. In the example given, there are more hourly copies kept than two-hour copies, resulting in more granularity for restore at times that are closer to the present. For example, in the previous system, if at 7:30 P.M. damage is discovered from earlier in the afternoon, a backup will be available for every hour for the past four hours: 4, 5, 6, 7 P.M. As well, two more backups will have been retained from 2 P.M. and 12 P.M.

The Content Addressable Store

[0165] FIG. 11 is a block diagram of the modules implementing the content addressable store for the Content Addressable Provider 510.

[0166] The content addressable store 510 implementation provides a storage resource pool that is optimized for capacity rather than for copy-in or copy-out speed, as would be the case for the performance-optimized pool implemented through snapshots, described earlier, and thus is typically used for offline backup, replication and remote backup. Content addressable storage provides a way of storing common subsets of different objects only once, where those common subsets may be of varying sizes but typically as small as 4 KiBytes. The storage overhead of a content addressable store is low compared to a snapshot store, though the access time is usually higher. Generally objects in a content addressable store have no intrinsic relationship to one another, even though they may share a large percentage of their content, though in this implementation a history relationship is also maintained, which is an enabler of various optimizations to be described. This contrasts with a snapshot store where snapshots intrinsically form a chain, each storing just deltas from a previous snapshot or baseline copy. In particular, the content addressable store will store only one copy of a data subset that is repeated multiple times within a single object, whereas a snapshot-based store will store at least one full-copy of any object.

[0167] The content addressable store 510 is a software module that executes on the same system as the pool manager, either in the same process or in a separate process communicating via a local transport such as TCP. In this embodiment, the content addressable store module runs in a separate process so as to minimize impact of software failures from different components.

[0168] This module's purpose is to allow storage of Data Storage Objects 403 in a highly space-efficient manner by deduplicating content (i.e., ensuring repeated content within single or multiple data objects is stored only once).

[0169] The content addressable store module provides services to the pool manager via a programmatic API. These services comprise the following:

[0170] Object to Handle mapping 1102: an object can be created by writing data into the store via an API; once the data is written completely the API returns an object handle determined by the content of the object. Conversely, data may be read as a stream of bytes from an offset within an object by providing the handle. Details of how the handle is constructed are explained in connection with the description of FIG. 12.

[0171] Temporal Tree Management 1104 tracks parent/child relationships between data objects stored. When a data object is written into the store 510, an API allows it to be linked as a child to a parent object already in the store. This indicates to the content addressable store that the child object is a modification of the parent. A single parent may have multiple children with different modifications, as might be the case for example if an application's data were saved into the store regularly for some while; then an early copy were restored and used as a new starting point for subsequent modifications. Temporal tree management operations and data models are described in more detail below.

[0172] Difference Engine 1106 can generate a summary of difference regions between two arbitrary objects in the store. The differencing operation is invoked via an API specifying the handles of two objects to be compared, and the form of the difference summary is a sequence of callbacks with the offset and size of sequential difference sections. The difference is calculated by comparing two hashed representations of the objects in parallel.

[0173] Garbage Collector 1108 is a service that analyzes the store to find saved data that is not referenced by any object handle, and to reclaim the storage space committed to this data. It is the nature of the content addressable store that much data is referenced by multiple object handles, i.e., the data is shared between data objects; some data will be referenced by a single object handle; but data that is referenced by no object handles (as might be the case if an object handle has been deleted from the content addressable system) can be safely overwritten by new data.

[0174] Object Replicator 1110 is a service to duplicate data objects between two different content addressable stores. Multiple content addressable stores may be used to satisfy additional business requirements, such as offline backup or remote backup.

[0175] These services are implemented using the functional modules shown in FIG. 11. The Data Hash module 1112 generates fixed length keys for data chunks up to a fixed size limit. For example, in this embodiment the maximum size of chunk that the hash generator will make a key for is 64 KiB. The fixed length key is either a hash, tagged to indicate the hashing scheme used, or a non-lossy algorithmic encoding. The hashing scheme used in this embodiment is SHA-1, which generates a secure cryptographic hash with a uniform

distribution and a probability of hash collision near enough zero that no facility need be incorporated into this system to detect and deal with collisions.

[0176] The Data Handle Cache 1114 is a software module managing an in-memory database that provides ephemeral storage for data and for handle-to-data mappings.

[0177] The Persistent Handle Management Index 1104 is a reliable persistent database of CAH-to-data mappings. In this embodiment it is implemented as a B-tree, mapping hashes from the hash generator to pages in the persistent data store 1118 that contain the data for this hash. Since the full B-tree cannot be held in memory at one time, for efficiency, this embodiment also uses an in-memory bloom filter to avoid expensive B-tree searches for hashes known not to be present.

[0178] The Persistent Data Storage module 1118 stores data and handles to long-term persistent storage, returning a token indicating where the data is stored. The handle/token pair is subsequently used to retrieve the data. As data is written to persistent storage, it passes through a layer of lossless data compression 1120, in this embodiment implemented using zlib, and a layer of optional reversible encryption 1122, which is not enabled in this embodiment.

[0179] For example, copying a data object into the content addressable store is an operation provided by the object/handle mapper service, since an incoming object will be stored and a handle will be returned to the requestor. The object/handle mapper reads the incoming object, requests hashes to be generated by the Data Hash Generator, stores the data to Persistent Data Storage and the handle to the Persistent Handle Management Index. The Data Handle Cache is kept updated for future quick lookups of data for the handle. Data stored to Persistent Data Storage is compressed and (optionally) encrypted before being written to disk. Typically a request to copy in a data object will also invoke the temporal tree management service to make a history record for the object, and this is also persisted via Persistent Data Storage.

[0180] As another example, copying a data object out of the content addressable store given its handle is another operation provided by the object/handle mapper service. The handle is looked up in the Data Handle Cache to locate the corresponding data; if the data is missing in the cache the persistent index is used; once the data is located on disk, it is retrieved via

persistent data storage module (which decrypts and decompresses the disk data) and then reconstituted to return to the requestor.

The Content Addressable Store Handle

[0181] FIG. 12 shows how the handle for a content addressed object is generated. The data object manager references all content addressable objects with a content addressable handle. This handle is made up of three parts. The first part 1201 is the size of the underlying data object the handle immediately points to. The second part 1202 is the depth of object it points to. The third 1203 is a hash of the object it points to. Field 1203 optionally includes a tag indicating that the hash is a non-lossy encoding of the underlying data. The tag indicates the encoding scheme used, such as a form of run-length encoding (RLE) of data used as an algorithmic encoding if the data chunk can be fully represented as a short enough RLE. If the underlying data object is too large to be represented as a non-lossy encoding, a mapping from the hash to a pointer or reference to the data is stored separately in the persistent handle management index 1104.

[0182] The data for a content addressable object is broken up into chunks 1204. The size of each chunk must be addressable by one content addressable handle 1205. The data is hashed by the data hash module 1102, and the hash of the chunk is used to make the handle. If the data of the object fits in one chunk, then the handle created is the final handle of the object. If not, then the handles themselves are grouped together into chunks 1206 and a hash is generated for each group of handles. This grouping of handles continues 1207 until there is only one handle 1208 produced which is then the handle for the object.

[0183] When an object is to be reconstituted from a content handle (the copy-out operation for the storage resource pool), the top level content handle is dereferenced to obtain a list of next-level content handles. These are dereferenced in turn to obtain further lists of content handles until depth-0 handles are obtained. These are expanded to data, either by looking up the handle in the handle management index or cache, or (in the case of an algorithmic hash such as run-length encoding) expanding deterministically to the full content.

Temporal Tree Management

[0184] FIG. 13 illustrates the temporal tree relationship created for data objects stored within the content addressable store. This particular data structure is utilized only within the content addressable store. The temporal tree management module maintains data structures 1302 in the persistent store that associate each content-addressed data object to a parent (which may be null, to indicate the first in a sequence of revisions). The individual nodes of the tree contain a single hash value. This hash value references a chunk of data, if the hash is a depth-0 hash, or a list of other hashes, if the hash is a depth-1 or higher hash. The references mapped to a hash value is contained in the Persistent Handle Management Index 1104. In some embodiments the edges of the tree may have weights or lengths, which may be used in an algorithm for finding neighbors.

[0185] This is a standard tree data structure and the module supports standard manipulation operations, in particular: 1310 Add: adding a leaf below a parent, which results in a change to the tree as between initial state 1302 and after-add state 1304; and 1312 Remove: removing a node (and reparenting its children to its parent), which results in a change to the tree as between after-add state 1304 and after-remove state 1306.

[0186] The “Add” operation is used whenever an object is copied-in to the CAS from an external pool. If the copy-in is via the Optimal Way for Data Backup, or if the object is originating in a different CAS pool, then it is required that a predecessor object be specified, and the Add operation is invoked to record this predecessor/successor relationship.

[0187] The “Remove” operation is invoked by the object manager when the policy manager determines that an object’s retention period has expired. This may lead to data stored in the CAS having no object in the temporal tree referring to it, and therefore a subsequent garbage collection pass can free up the storage space for that data as available for re-use.

[0188] Note that it is possible for a single predecessor to have multiple successors or child nodes. For example, this may occur if an object is originally created at time T1 and modified at time T2, the modifications are rolled back via a restore operation, and subsequent modifications are made at time T3. In this example, state T1 has two children, state T2 and state T3.

[0189] Different CAS pools may be used to accomplish different business objectives such as providing disaster recovery in a remote location. When copying from one CAS to another CAS, the copy may be sent as hashes and offsets, to take advantage of the native deduplication capabilities of the target CAS. The underlying data pointed to by any new hashes is also sent on an as-needed basis.

[0190] The temporal tree structure is read or navigated as part of the implementation of various services:

- Garbage Collection navigates the tree in order to reduce the cost of the “mark” phase, as described below
- Replication to a different CAS pool finds a set of near-neighbors in the temporal tree that are also known to have been transferred already to the other CAS pool, so that only a small set of differences need to be transferred additionally
- Optimal-Way for data restore uses the temporal tree to find a predecessor that can be used as a basis for the restore operation. In the CAS temporal tree data structure, children are subsequent versions, e.g., as dictated by archive policy. Multiple children are supported on the same parent node; this case may arise when a parent node is changed, then used as the basis for a restore, and subsequently changed again.

CAS Difference Engine

[0191] The CAS difference engine 1106 compares two objects identified by hash values or handles as in FIGs. 11 and 12, and produces a sequence of offsets and extents within the objects where the object data is known to differ. This sequence is achieved by traversing the two object trees in parallel in the hash data structure of FIG. 12. The tree traversal is a standard depth- or breadth-first traversal. During traversal, the hashes at the current depth are compared. Where the hash of a node is identical between both sides, there is no need to descend the tree further, so the traversal may be pruned. If the hash of a node is not identical, the traversal continues descending into the next lowest level of the tree. If the traversal reaches a depth-0 hash that is not identical to its counterpart, then the absolute offset into the data object being compared where the non-identical data occurs, together with the data length, is emitted into the output sequence. If one object is smaller in size than another, then its traversal

will complete earlier, and all subsequent offsets encountered in the traversal of the other are emitted as differences.

Garbage Collection via Differencing

[0192] As described under FIG. 11, Garbage Collector is a service that analyzes a particular CAS store to find saved data that is not referenced by any object handle in the CAS store temporal data structure, and to reclaim the storage space committed to this data. Garbage collection uses a standard “Mark and Sweep” approach. Since the “mark” phase may be quite expensive, the algorithm used for the mark phase attempts to minimize marking the same data multiple times, even though it may be referenced many times; however the mark phase must be complete, ensuring that no referenced data is left unmarked, as this would result in data loss from the store as, after a sweep phase, unmarked data would later be overwritten by new data.

[0193] The algorithm employed for marking referenced data uses the fact that objects in the CAS are arranged in graphs with temporal relationships using the data structure depicted in FIG. 13. It is likely that objects that share an edge in these graphs differ in only a small subset of their data, and it is also rare that any new data chunk that appears when an object is created from a predecessor should appear again between any two other objects. Thus, the mark phase of garbage collection processes each connected component of the temporal graph.

[0194] FIG. 14 is an example of garbage collection using temporal relationships in certain embodiments. A depth-first search is made, represented by arrows 1402, of a data structure containing temporal relationships. Take a starting node 1404 from which to begin the tree traversal. Node 1404 is the tree root and references no objects. Node 1406 contains references to objects H_1 and H_2 , denoting a hash value for object 1 and a hash value for object 2. All depth-0, depth-1 and higher data objects that are referenced by node 1406, here H_1 and H_2 , are enumerated and marked as referenced.

[0195] Next, node 1408 is processed. As it shares an edge with node 1406, which has been marked, the difference engine is applied to the difference between the object referenced by 1406 and the object referenced by 1408, obtaining a set of depth-0, depth-1 and higher hashes that exist in the unmarked object but not in the marked object. In the figure, the hash that

exists in node 1408 but not in node 1406 is H_3 , so H_3 is marked as referenced. This procedure is continued until all edges are exhausted.

[0196] A comparison of the results produced by a prior art algorithm 1418 and the present embodiment 1420 shows that when node 1408 is processed by the prior art algorithm, previously-seen hashes H_1 and H_2 are emitted into the output stream along with new hash H_3 . Present embodiment 1420 does not emit previously seen hashes into the output stream, resulting in only new hashes H_3 , H_4 , H_5 , H_6 , H_7 being emitted into the output stream, with a corresponding improvement in performance. Note that this method does not guarantee that data will not be marked more than once. For example, if hash value H_4 occurs independently in node 1416, it will be independently marked a second time.

Copy an Object into the CAS

[0197] Copying an object from another pool into the CAS uses the software modules described in FIG. 11 to produce a data structure referenced by an object handle as in FIG. 12. The input to the process is (a) a sequence of chunks of data at specified offsets, sized appropriately for making depth-0 handles, and optionally (b) a previous version of the same object. Implicitly, the new object will be identical to the previous version except where the input data is provided and itself differs from the previous version. The algorithm for the copy-in operation is illustrated in a flowchart at FIG. 15.

[0198] If a previous version (b) is provided, then the sequence (a) may be a sparse set of changes from (b). In the case that the object to be copied is known to differ from a previous object at only a few points, this can greatly reduce the amount of data that needs to be copied in, and therefore reduce the computation and i/o activity required. This is the case, for example, when the object is to be copied in via the optimal way for data backup described previously.

[0199] Even if the sequence (a) includes sections that are largely unchanged from a predecessor, identifying the predecessor (b) allows the copy-in procedure to do quick checks as to whether the data has indeed changed and therefore to avoid data duplication at a finer level of granularity than might be possible for the difference engine in some other storage pool providing input to a CAS.

[0200] Implicitly then, the new object will be identical to the previous version except where the input data is provided and itself differs from the previous version. The algorithm for the copy-in operation is illustrated in a flowchart at FIG. 15.

[0201] The process starts at step 1500 as an arbitrarily-sized data object in the temporal store is provided, and proceeds to 1502, which enumerates any and all hashes (depth-0 through the highest level) referenced by the hash value in the predecessor object, if such is provided. This will be used as a quick check to avoid storing data that is already contained in the predecessor.

[0202] At step 1504, if a predecessor is input, create a reference to a clone of it in the content-addressable data store temporal data structure. This clone will be updated to become the new object. Thus the new object will become a copy of the predecessor modified by the differences copied into the CAS from the copying source pool.

[0203] At steps 1506, 1508, the Data Mover 502 pushes the data into the CAS. The data is accompanied by an object reference and an offset, which is the target location for the data. The data may be sparse, as only the differences from the predecessor need to be moved into the new object. At this point the incoming data is broken into depth-0 chunks sized small enough that each can be represented by a single depth-0 hash.

[0204] At step 1510, the data hash module generates a hash for each depth-0 chunk.

[0205] At step 1512, read the predecessor hash at the same offset. If the hash of the data matches the hash of the predecessor at the same offset, then no data needs to be stored and the depth-1 and higher objects do not need to be updated for this depth-0 chunk. In this case, return to accept the next depth-0 chunk of data. This achieves temporal deduplication without having to do expensive global lookups. Even though the source system is ideally sending only the differences from the data that has previously been stored in the CAS, this check may be necessary if the source system is performing differencing at a different level of granularity, or if the data is marked as changed but has been changed back to its previously-stored value. Differencing may be performed at a different level of granularity if, for example, the source system is a snapshot pool which creates deltas on a 32KiB boundary and the CAS store creates hashes on 4 KiB chunks.

[0206] If a match is not found, the data may be hashed and stored. Data is written starting at the provided offset and ending once the new data has been exhausted. Once the data has been stored, at step 1516, if the offset is still contained within the same depth-1 object, then depth-1, depth-2 and all higher objects 1518 are updated, generating new hashes at each level, and the depth-0, depth-1 and all higher objects are stored at step 1514 to a local cache.

[0207] However, at step 1520, if the amount of data to be stored exceeds the depth-1 chunk size and the offset is to be contained in a new depth-1 object, the current depth-1 must be flushed to the store, unless it is determined to be stored there already. First look it up in the global index 1116. If it is found there, remove the depth-1 and all associated depth-0 objects from the local cache and proceed with the new chunk 1522.

[0208] At step 1524, as a quick check to avoid visiting the global index, for each depth-0, depth-1 and higher object in the local cache, lookup its hash in the local store established in 1502. Discard any that match.

[0209] At step 1526, for each depth-0, depth-1 and higher object in the local cache, lookup its hash in the global index 1116. Discard any that match. This ensures that data is deduplicated globally.

[0210] At step 1528: store all remaining content from the local cache into the persistent store, then continue to process the new chunk.

[0211] Reading an object out of the CAS is a simpler process and is common across many implementations of CAS. The handle for the object is mapped to a persistent data object via the global index, and the offset required is read from within this persistent data. In some cases it may be necessary to recurse through several depths in the object handle tree.

CAS Object Network Replication

[0212] As described under FIG. 11, the Replicator 1110 is a service to duplicate data objects between two different content addressable stores. The process of replication could be achieved through reading out of one store and writing back into another, but this architecture allows more efficient replication over a limited bandwidth connection such as a local- or wide-area network.

[0213] A replicating system operating on each CAS store uses the difference engine service described above together with the temporal relationship structure as described in FIG. 13, and additionally stores on a per-object basis in the temporal data structure used by the CAS store a record of what remote store the object has been replicated to. This provides definitive knowledge of object presence at a certain data store.

[0214] Using the temporal data structure, it is possible for the system to determine which objects exist on which data stores. This information is leveraged by the Data Mover and Difference Engine to determine a minimal subset of data to be sent over the network during a copy operation to bring a target data store up to date. For example, if data object O has been copied at time T3 from a server in Boston to a remote server in Seattle, Protection Catalog Store 908 will store that object O at time T3 exists both in Boston and Seattle. At time T5, during a subsequent copy from Boston to Seattle, the temporal data structure will be consulted to determine the previous state of object O in Seattle that should be used for differencing on the source server in Boston. The Boston server will then take the difference of T5 and T3, and send that difference to the Seattle server.

[0215] The process to replicate an object A is then as follows: Identify an object A0 that is recorded as having already been replicated to the target store and a near neighbor of A in the local store. If no such object A0 exists then send A to the remote store and record it locally as having been sent. To send a local object to the remote store, a typical method as embodied here is: send all the hashes and offsets of data chunks within the object; query the remote store as to which hashes represent data that is not present remotely; send the required data to the remote store (sending the data and hashes is implemented in this embodiment by encapsulating them in a TCP data stream).

[0216] Conversely, if A0 is identified, then run the difference engine to identify data chunks that are in A but not in A0. This should be a superset of the data that needs to be sent to the remote store. Send hashes and offsets for chunks that are in A but not in A0. Query the remote store as to which hashes represent data that is not present remotely; send the required data to the remote store.

Sample Deployment Architecture

[0217] FIG. 16 shows the software and hardware components that comprise one embodiment of the Data Management Virtualization (DMV) system. The software that comprises the system executes as three distributed components:

[0218] The Host Agent software 1602a, 1602b, 1602c implements some of the application-specific module described above. It executes on the same servers 1610a, 1610b, 1610c as the application whose data is under management.

[0219] The DMV server software 1604a, 1604b implements the remainder of the system as described here. It runs on a set of Linux servers 1612, 1614 that also provide highly available virtualized storage services.

[0220] The system is controlled by Management Client software 1606 that runs on a desktop or laptop computer 1620.

[0221] These software components communicate with one another via network connections over an IP network 1628. Data Management Virtualization systems communicate with one another between primary site 1622 and data replication (DR) site 1624 over an IP network such as a public internet backbone.

[0222] The DMV systems at primary and DR sites access one or more SAN storage systems 1616, 1618 via a fibre-channel network 1626. The servers running primary applications access the storage virtualized by the DMV systems access the storage via fibre-channel over the fibre-channel network, or iSCSI over the IP network. The DMV system at the remote DR site runs a parallel instance of DMV server software 1604c on Linux server 1628. Linux server 1628 may also be an Amazon Web Services EC2 instance or other similar cloud computational resource.

[0223] FIG. 17 is a diagram that depicts the various components of a computerized system upon which certain elements may be implemented, according to certain embodiments of the invention. The logical modules described may be implemented on a host computer 1701 that contains volatile memory 1702, a persistent storage device such as a hard drive, 1708, a processor, 1703, and a network interface, 1704. Using the network interface, the system

computer can interact with storage pools 1705, 1706 over a SAN or Fibre Channel device, among other embodiments. Although FIG. 17 illustrates a system in which the system computer is separate from the various storage pools, some or all of the storage pools may be housed within the host computer, eliminating the need for a network interface. The programmatic processes may be executed on a single host, as shown in FIG. 17, or they may be distributed across multiple hosts.

[0224] The host computer shown in FIG. 17 may serve as an administrative workstation, or may implement the application and Application Specific Agent 402, or may implement any and all logical modules described in this specification, including the Data Virtualization System itself, or may serve as a storage controller for exposing storage pools of physical media to the system. Workstations may be connected to a graphical display device, 1707, and to input devices such as a mouse 1709 and a keyboard 1710. Alternately, the active user's workstation may comprise a handheld device.

[0225] Throughout this specification we refer to software components, but all references to software components are intended to apply to software running on hardware. Likewise, objects and data structures referred to in the specification are intended to apply to data structures actually stored in memory, either volatile or non-volatile. Likewise, servers are intended to apply to software, and engines are intended to apply to software, all running on hardware such as the computer systems described in FIG. 17.

[0226] The foregoing has outlined some of the more pertinent features of the subject matter. These features should be construed to be merely illustrative. Many other beneficial results can be attained by applying the disclosed subject matter in a different manner or by modifying the subject matter as will be described. For example, the garbage collection algorithm that is disclosed may incorporate other garbage collection optimization techniques such as tri-color marking.

[0227] What is claimed is:

1. A system for performing a plurality of prescribed data management functions in a manner that reduces redundant access operations to primary storage, said system comprising:

a data management engine for performing data management functions, including at least a snapshot function operable to create a point-in-time image of primary storage data to secondary storage, and at least one back-up function operable to create at least one back-up copy of data, said data management engine being responsive to an electronic service level agreement (SLA) that specifies a schedule for performing data management functions,

wherein point-in-time images of data include a reference to a complete baseline image of data at a specific point in time and difference data indicating changes to the data at a later, specific point in time, and

wherein, in response to the schedule requiring at least some data management functions to be performed concurrently, the data management engine creates a point-in-time image of the primary storage data and communicates the difference information of that point-in-time image to the secondary storage to update the at least one of a back-up copy of the primary data, such that the primary storage is accessed only once for all corresponding sets of updates to the secondary storage.
2. The system of claim 1, wherein the point-in-time image of primary storage data at secondary storage is stored on performance optimized secondary storage.
3. The system of claim 1, wherein said back-up copy of the point-in-time image of primary storage data is stored on remote storage.
4. The system of claim 1, wherein said back-up copy of the point-in-time image of primary storage data is stored at capacity optimized storage.
5. The system of claim 4, wherein said back-up copy of the point-in-time image of primary storage data is stored as a deduplicated image on capacity optimized storage.

6. The system of claim 1, wherein difference data includes bitmap information with each bit of the bitmap corresponding to a portion of primary storage data, and including new data for those portions of the bitmap which are set to indicate that data has changed.
7. The system of claim 1, wherein difference data includes extent information.
8. The system of claim 1, wherein the data management engine includes logic to invoke the primary storage to provide a point-in-time image of data and includes logic to retrieve the point-in-time image from the primary storage.
9. A system for managing data in accordance with service level agreements (SLAs) that specify schedules on a calendar basis for performing prescribed data management functions and for reducing inter-function redundancy, said system comprising:

a data management engine for performing data management functions, including at least a snapshot function, and at least one back-up function, said data management engine including a service level policy engine that receives SLAs in electronic form and which controls the scheduling of the data management functions in accordance therewith,

wherein each electronic SLA is associated with a corresponding application that uses data, and wherein each SLA specifies at least one service level policy, each policy specifying a source pool for data, a destination pool where a copy should be made of the source pool data, copy frequency indicating the frequency of operation for that policy, retention period indicating how long a given copy should be retained before being allowed to expire, and schedule information indicating hours and days of operation when the policy is in place, such that a collection of policies within a SLA is capable of expressing a non-uniform schedule for when a given function should be performed and is capable of expressing multiple data management functions that should be performed on a given source of data, and

wherein said data management engine is operable to perform preparatory operations with the application and with the source pool so that the source pool of data has a coherent image of data to be copied and wherein said preparatory operations are

performed once even if the SLA specifies multiple data management functions to be performed on that source pool at the current time.

10. The system of claim 9, wherein if two or more copy operations are scheduled to occur at the same instant between the same source pool and destination pool, only one of the two or more copy operations is performed by the data management engine and that copy is associated with the longest retention time corresponding to the two or more scheduled copy operations.
11. The system of claim 9, wherein preparatory operations include the data management engine collecting metadata about the application to store in conjunction with application data.
12. The system of claim 9, wherein preparatory operations include application quiescing operations.
13. The system of claim 12, wherein the application quiescing operations include freezing the application from further updating application data.
14. The system of claim 12, wherein the application quiescing operations include flushing the I/O cache of the application server of application data.
15. A system for backing-up data from a first storage pool to a second storage pool using difference information between time states, said system comprising:

a data management engine for performing data management functions, including at least a back-up function to create a back-up copy of data,

said data management engine operable to execute a sequence of snapshot operations to create point-in-time images of application data on a first storage pool, each successive point-in-time image corresponding to a specific, successive time-state of the application data, and each snapshot operation creating difference information indicating which application data has changed and the content of the changed application data for the corresponding time state;

said data management engine operable to execute at least one back-up function for the application data wherein the backup operation is scheduled for execution at non-consecutive time-states

wherein said data management engine is operable to maintain history information having time-state information indicating the time-state of the last back-up function performed on the application data for a corresponding back-up copy of data; and

wherein the data management engine is operable to create composite difference information from the difference information for each time-state between the time-state of the last back-up function performed on the application data and the time-state of the currently-scheduled back-up function to be performed on the application data, and wherein the data management engine is operable to send the composite difference information to a second storage pool to be compiled with the back-up copy of data at the last time-state to create a back-up copy of data for the current time-state.

16. The system of claim 15, wherein difference information includes bitmap information with each bit of the bitmap corresponding to a portion of primary storage data, and including new data for those portions of the bitmap which are set to indicate that data has changed.
17. The system of claim 15, wherein difference information includes extent information.
18. The system of claim 15, wherein multiple back-up functions are scheduled to occur simultaneously, each with different gaps of non-consecutive time-states, and each with different composite difference information generated corresponding to the different gaps
19. A system for restoring data in a storage pool from a back-up copy of the data using difference information between time states, said system comprising:

a data management engine wherein said data management engine is operable to maintain history information indicating the time-states for which storage pools have point-in-time images of application data; and

wherein said data management engine includes logic for restoring application data in a storage pool to a point-in-time image of the data for a specified time-state;

said data management engine operable to identify the existence of a point-in-time image of the data at the storage pool for a time-state prior to the specified time-state and sending difference information from the back-up copy of data to the storage pool, said difference information indicating which application data has changed and the content of the changed application data for the time between the specified time state and the time state prior to the specified time-state.

20. The system of claim 19, wherein difference information includes bitmap information with each bit of the bitmap corresponding to a portion of primary storage data, and includes back-up data for those portions of the bitmap which are set to indicate that data has changed.
21. The system of claim 19, wherein difference information includes extent information.
22. The system of claim 19, wherein the prior time-state and the specified time-state are non-consecutive time-states.
23. A method of forming deduplicated images of a data object that changes over time using difference information between temporal states of the data object, said method comprising:

organizing the content of the data object for a first temporal state as a plurality of content segments and storing the content segments in a data store;

creating an organized arrangement of hash structures to represent the data object in its first temporal state wherein for a subset of the hash structures, each structure includes a hash signature for a corresponding content segment and is associated with a reference to the corresponding content segment, and wherein the logical organization of the arrangement represents the organization of the content segments as they are represented within the data object;

receiving difference information for the data object, said difference information indicating the changed content for the data object for a second temporal state relative to

the first temporal state, and said difference information indicating the location of the changed content within the data object;

forming at least one hash signature for the changed content;

storing the changed content that is unique in the data store as content segments

modifying the organized arrangement of hash structures to incorporate new structures for the at least one hash signature for the changed content, incorporating the new structures in the organized arrangement of structures at a position corresponding to the location of the changed content within the data object as indicated within said difference information, and associating the hash signatures for the new structures with references to the corresponding content segments for the changed content; and

associating the new structures with the second temporal state, whereby a deduplicated image of the data object for a second temporal state is stored without requiring reception of a complete image of the data object for the second temporal state.

24. The method of claim 23, wherein after forming the at least one hash signature for changed content, the formed signature is compared to at least one hash signature in the organized arrangement of hash structures to determine if the formed structure already exists in the organized arrangement.
25. The method of claim 23, wherein the comparison first occurs with a hash structure in the organized arrangement at a position corresponding to the location of the changed content as indicated in the difference information.
26. The method of claim 23, wherein the organized arrangement of hash structures is an organized tree structure.
27. The method of claim 23, wherein an organized arrangement of temporal structures are maintained, each temporal structure associated with a time state and each including information indicative of the hash structures corresponding to the associated time state.
28. A method of managing deduplicated images of data objects that change over time, said method comprising:

organizing unique content of each data object as a plurality of content segments and storing the content segments in a data store;

for each data object, creating an organized arrangement of hash structures, wherein each structure, for a subset of the hash structures, includes a hash signature for a corresponding content segment and is associated with a reference to the corresponding content segment, wherein the logical organization of the arrangement represents the logical organization of the content segments as they are represented within the data object, and wherein another subset of the hash structures includes a hierarchy of hash signatures of said hash signatures for corresponding content segments so that the organized arrangement may be traversed to determine if content is represented by said organized arrangement of hash structures; and

for each data object, maintaining an organized arrangement of temporal structures to represent a corresponding data object over time, wherein each structure is associated with a temporal state of the data object and wherein the logical arrangement of structures is indicative of the changing temporal states of the data object, and wherein each temporal state is associated with the hash structures representing the content of the data object during that temporal state.

29. The method of claim 28, wherein the temporal structures for a data object at a given temporal state is associated with hash structures for data object content that has changed relative to a prior temporal state of the data object.
30. The method of claim 28, wherein the hash structures for the changed data content is organized as a graph separate from the organized arrangement of hash structures for a prior temporal state.
31. The method of claim 28, wherein differences in content of a data object from one temporal state to another is determined by referencing the organized arrangement of temporal structures for the other temporal state and for all other temporal states intervening between the other temporal state and the one state so that differences may be determined over multiple temporal states.

32. A method of storing deduplicated images in which a portion of the image is stored in encoded form directly in a hash table, the method comprising:
- organizing unique content of each data object as a plurality of content segments and storing the content segments in a data store;
- for each data object, creating an organized arrangement of hash structures, wherein each structure, for a subset of the hash structures, includes a field to contain a hash signature for a corresponding content segment and is associated with a reference to the corresponding content segment, wherein the logical organization of the arrangement represents the logical organization of the content segments as they are represented within the data object;
- receiving content to be included in the deduplicated image of the data object;
- determining if the received content may be encoded using a predefined non-lossy encoding technique and in which the encoded value would fit within the field for containing a hash signature;
- if so, placing the encoding in the field and marking the hash structure to indicate that the field contains encoded content for the deduplicated image;
- if not, generating a hash signature for the received content and placing the hash signature in the field and placing the received content in a corresponding content segment in said data store if it is unique.
33. The method of claim 32, wherein said non-lossy encoding is a run-length encoding.
34. The method of claim 32, wherein hash signatures for each data object are created using the SHA-1 cryptographic hash function.
35. The method of claim 32, further comprising subsequently reconstituting the content from the encoded content.

36. A method for using a first deduplicating store to update a second deduplicating store with information representing how data objects change over time, said method comprising:

at a first deduplicating store, organizing unique content of each data object as a plurality of content segments and storing the content segments in a data store;

at a first deduplicating store, for each data object, creating an organized arrangement of hash structures, wherein each structure, for a subset of the hash structures, includes a hash signature for a corresponding content segment and is associated with a reference to the corresponding content segment, wherein the logical organization of the arrangement represents the logical organization of the content segments as they are represented within the data object;

at a first deduplicating store, for each data object, maintaining an organized arrangement of temporal structures to represent a corresponding data object over time, wherein each structure is associated with a temporal state of the data object and wherein the logical arrangement of structures is indicative of the changing temporal states of the data object, and wherein each temporal state is associated with the hash structures representing the content of the data object that has changed relative to a prior temporal state;

at a second deduplicating store, organizing unique content of each data object as a plurality of content segments and storing the content segments in a data store;

at a second deduplicating store, for each data object, maintaining an organized arrangement of hash structures, that is at least a subset of the hash structures at said first deduplicating store;

at a second deduplicating store, for each data object, maintaining an organized arrangement of temporal structures to represent a corresponding data object over time, wherein said organized arrangement of temporal structures is at least a subset of the temporal structures at said first deduplicating store, representing a subset of the temporal states;

in response to a request to update the second deduplicating store with information from said first deduplicating store, finding a temporal state that is common to the first and second deduplicating stores and that is in temporal proximity to the current state of the second deduplicating store; and

compiling a set of hash signatures for the content that has changed from the common state to the current temporal state of the first deduplicating store and sending that set of hash signatures to the second deduplicating store so it can update its organized arrangement of hash structures to represent the content of the data object up to the current temporal state of the first deduplicating store.

37. The method of claim 36, further including maintaining a history of the hash signatures that each deduplicating store contains, and for hash signatures in the set of hash signatures that are new to the second deduplicating store, sending the corresponding content segments from the first deduplicating store so that the second deduplicating store may update its data store with the new content.
38. The method of claim 36, wherein the temporal state that is in proximity to the current state of the second deduplicating store is a nearest neighbor state of the current state.
39. The method of claim 36, wherein the temporal state that is in proximity to the current state of the second deduplicating store is an ancestor state of the current state.
40. The method of claim 36, wherein the temporal state that is in proximity to the current state of the second deduplicating store is a child state of the current state.
41. The method of claim 38, wherein the nearest neighbor state is the state which is connected by a set of edges whose sum is lower than the sum of any other set of edges.
42. The method of claim 36, wherein the logical arrangement of structures contains branching.
43. The method of claim 37, further comprising recording, at the current state, to what states the content segments corresponding to the current temporal state has been sent.

44. A method of performing garbage collection to identify content segments no longer referenced in a deduplicating storage system in which redundant mark operations in a mark-and-sweep technique are avoided, the method comprising:

organizing unique content of each data object as a plurality of content segments in the deduplicating storage system;

for each data object, creating an organized arrangement of hash structures, wherein each structure, for a subset of the hash structures, includes a hash signature for a corresponding content segment and is associated with a reference to the corresponding content segment, wherein the logical organization of the arrangement represents the logical organization of the content segments as they are represented within the data object, and wherein another subset of the hash structures includes a hierarchy of hash signatures of said hash signatures for corresponding content segments so that the organized arrangement may be traversed to determine if content is represented by said organized arrangement of hash structures;

for each data object, maintaining an organized arrangement of temporal structures to represent a corresponding data object over time, wherein each structure is associated with a temporal state of the data object and wherein the logical arrangement of structures is indicative of the changing temporal states of the data object, and wherein each temporal state is associated with the hash structures representing the content of the data object that has changed relative to an immediately prior temporal state of the data object;

for each content segment in the deduplicating storage system, clearing its corresponding garbage collection state;

iterating over the temporal structures and, for each temporal structure, marking the garbage collection state for the associated content segments for only the content segments that have changed relative to an immediately prior temporal state of the data object; and

returning any content segments to a free pool of storage that has a cleared garbage collection state after the iteration step.

45. The method of claim 44, further comprising iterating over the temporal structures using a depth first search.
46. The method of claim 44, further comprising repeating the method at periodic intervals.
47. The method of claim 44, further comprising performing the method subsequent to the addition of a new temporal state of a data object.
48. The method of claim 44, further comprising performing the method subsequent to the removal of a temporal state of a data object.
49. The method of claim 44, further comprising maintaining a global reference list of all content segments that have been allocated in the deduplicated storage system.




FIG. 1 (Prior Art)

2/17




FIG. 2




FIG. 3

4/17




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9

10/17




FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14

15/17




FIG. 15

16/17




FIG. 16




FIG. 17

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2011/060417

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 12/00 (2012.01)

USPC - 707/649

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)

IPC(8) - G06F 7/00; 12/00; 12/02; 12/08 (2012.01)

USPC - 707/640, 644-45, 648-54, 674-79, 790-92; 711/100, 161-62, E12.103

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)

MicroPatent, Google Patents, ProQuest, Google Scholar

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2002/0049778 A1 (BELL et al) 25 April 2002 (25.04.2002) entire document	1-49
A	US 2008/0034016 A1 (CISLER et al) 07 February 2008 (07.02.2008) entire document	1-49
A	US 6,883,073 B2 (ARAKAWA et al) 19 April 2005 (19.04.2005) entire document	1-49
A	US 2010/0077013 A1 (CLEMENTS et al) 25 March 2010 (25.03.2010) entire document	1-49
A	US 2010/0138827 A1 (FRANK et al) 03 June 2010 (03.06.2010) entire document	1-49

☐ Further documents are listed in the continuation of Box C.

* 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

09 March 2012

Date of mailing of the international search report

16 MAR 2012

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774