



(19) **United States**

(12) **Patent Application Publication**
Prahlad et al.

(10) **Pub. No.: US 2020/0228598 A1**

(43) **Pub. Date: Jul. 16, 2020**

(54) **DATA TRANSFER TECHNIQUES WITHIN DATA STORAGE DEVICES, SUCH AS NETWORK ATTACHED STORAGE PERFORMING DATA MIGRATION**

Publication Classification

(51) **Int. Cl.**
H04L 29/08 (2006.01)
G06F 3/06 (2006.01)
(52) **U.S. Cl.**
CPC *H04L 67/1095* (2013.01); *G06F 3/0605* (2013.01); *G06F 3/067* (2013.01); *G06F 12/0866* (2013.01); *H04L 67/1097* (2013.01); *H04L 67/2852* (2013.01); *G06F 3/0649* (2013.01)

(71) Applicant: **Commvault Systems, Inc.**, Tinton Falls, NJ (US)

(72) Inventors: **Anand Prahlad**, Bangalore (IN); **Rajiv Kottomtharayil**, Marlboro, NJ (US); **Manoj Kumar Vijayan**, Marlboro, NJ (US); **Parag Gokhale**, Marlboro, NJ (US); **Deepak Raghunath Attarde**, Marlboro, NJ (US); **Kamleshkumar K. Lad**, Dublin, CA (US)

(21) Appl. No.: **16/732,262**

(22) Filed: **Dec. 31, 2019**


Related U.S. Application Data

(63) Continuation of application No. 14/963,954, filed on Dec. 9, 2015, now Pat. No. 10,547,678, which is a continuation of application No. 12/558,640, filed on Sep. 14, 2009, now abandoned.

(60) Provisional application No. 61/097,176, filed on Sep. 15, 2008.

(57) **ABSTRACT**

A stand-alone, network accessible data storage device, such as a filer or NAS device, is capable of transferring data objects based on portions of the data objects. The device transfers portions of files, folders, and other data objects from a data store within the device to external secondary storage based on certain criteria, such as time-based criteria, age-based criteria, and so on. A portion may be one or more blocks of a data object, or one or more chunks of a data object, or other segments that combine to form or store a data object. For example, the device identifies one or more blocks of a data object that satisfy a certain criteria, and migrates the identified blocks to external storage, thereby freeing up storage space within the device. The device may determine that a certain number of blocks of a file have not been modified or called by a file system in a certain time period, and migrate these blocks to secondary storage.






FIG. 1

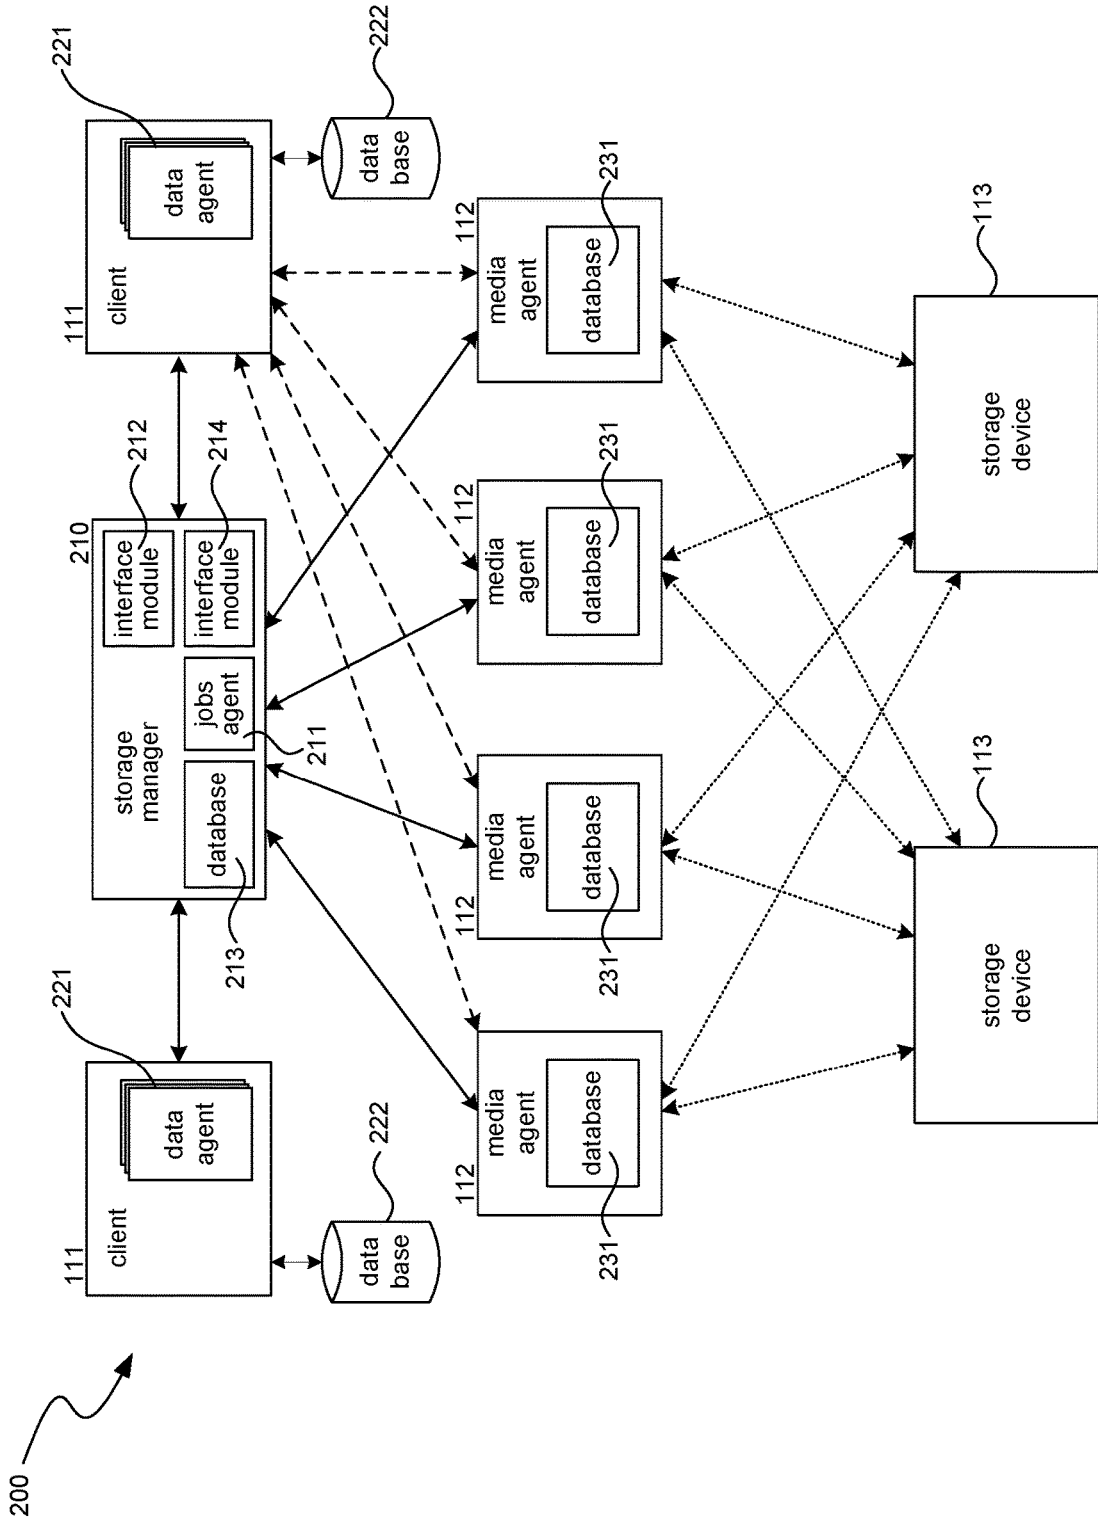


FIG. 2




FIG. 3

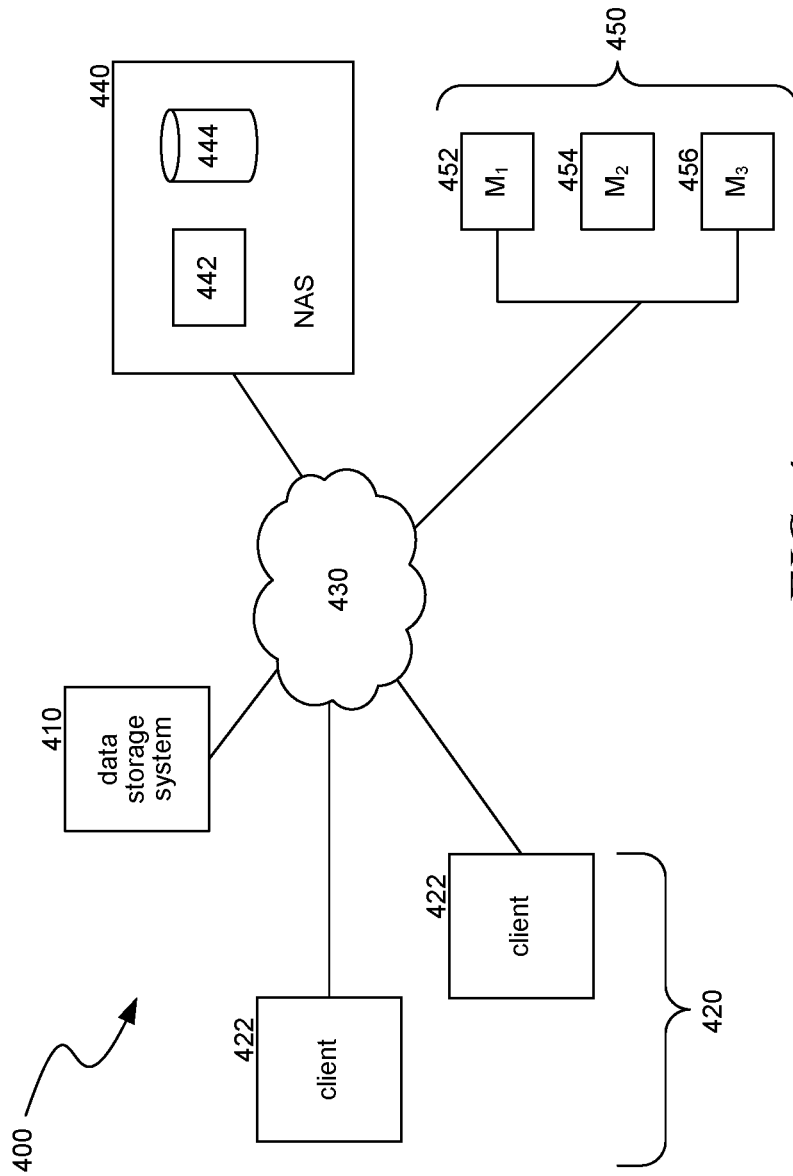


FIG. 4

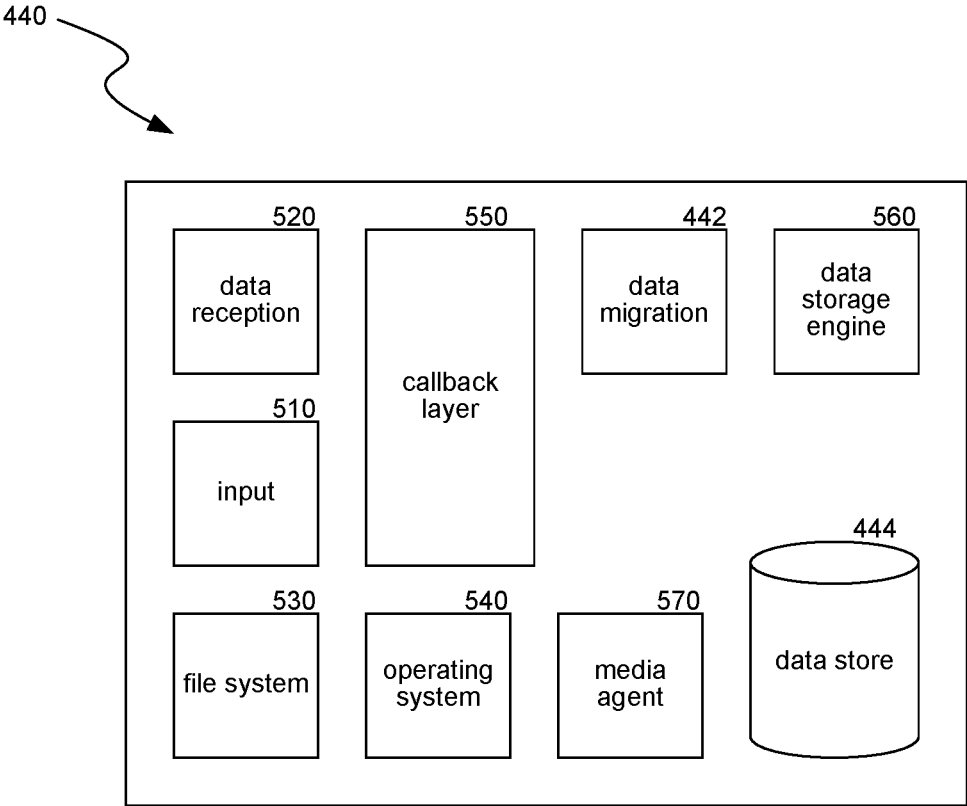


FIG. 5




FIG. 6B

FIG. 6A

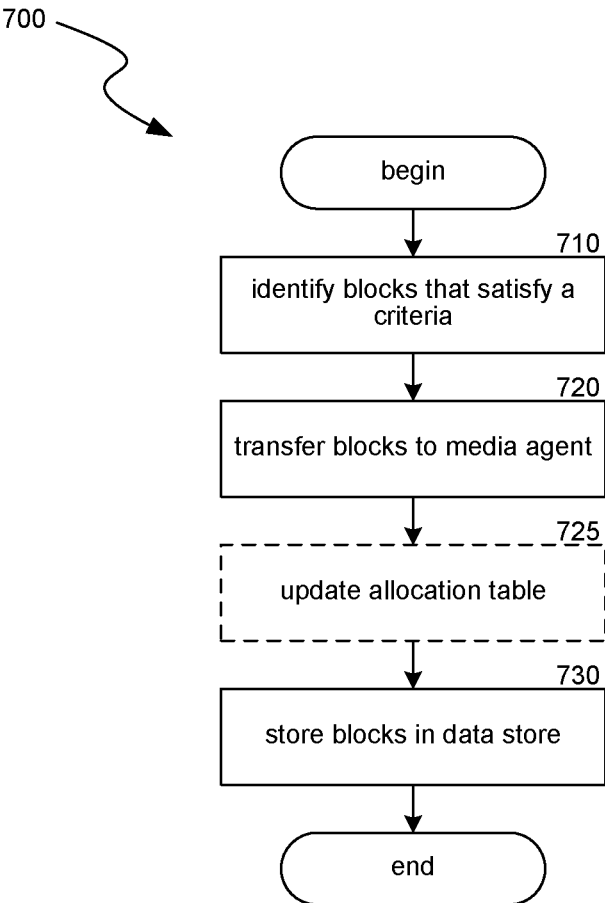


FIG. 7

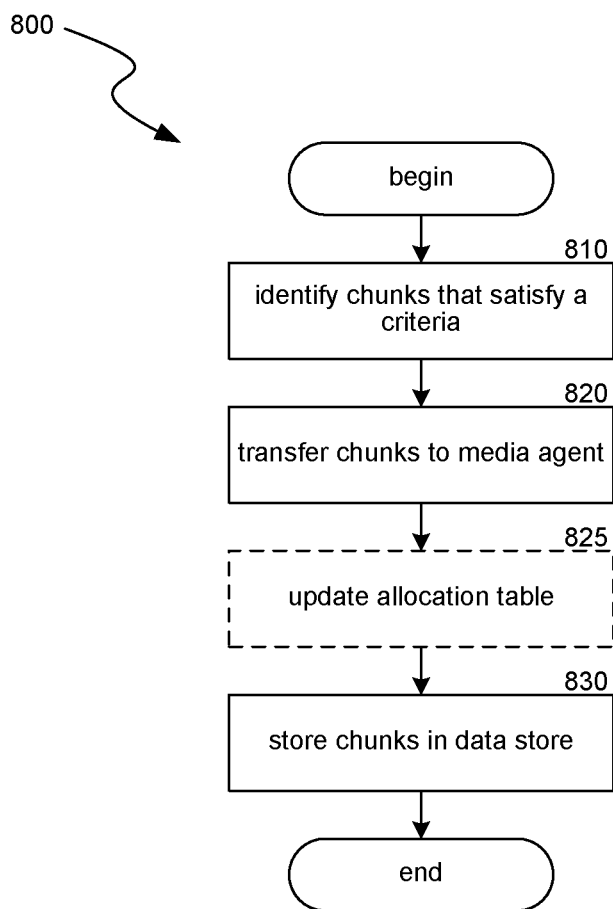


FIG. 8

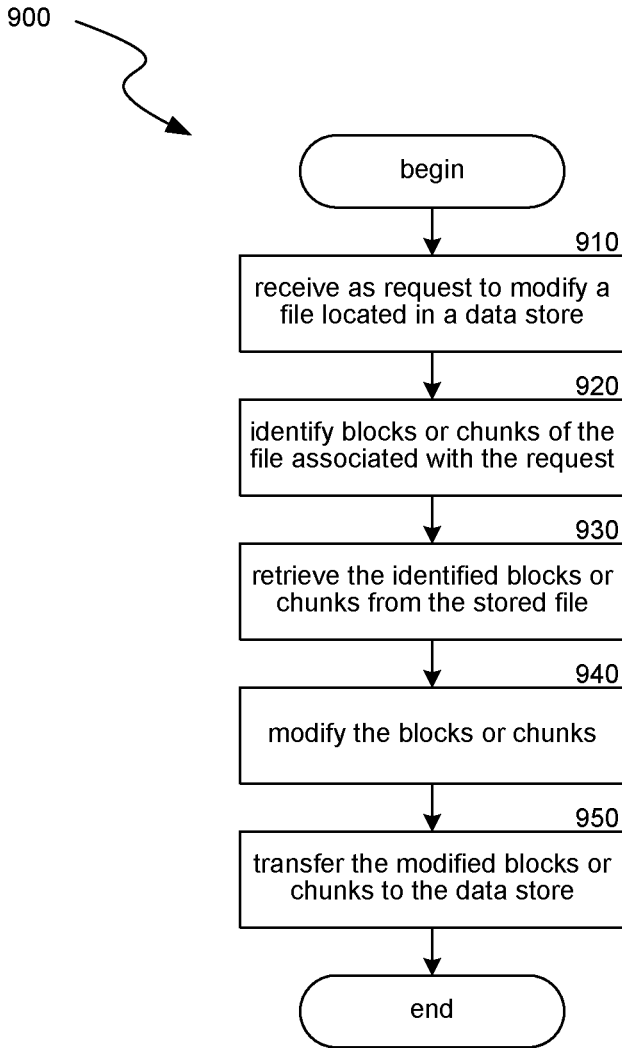


FIG. 9

**DATA TRANSFER TECHNIQUES WITHIN
DATA STORAGE DEVICES, SUCH AS
NETWORK ATTACHED STORAGE
PERFORMING DATA MIGRATION**

**CROSS-REFERENCE TO RELATED
APPLICATION**

[0001] This application is a continuation of U.S. application Ser. No. 14/963,954 filed Dec. 9, 2015, which is a continuation of U.S. patent application Ser. No. 12/558,640 filed Sep. 14, 2009 (entitled DATA TRANSFER TECHNIQUES WITHIN DATA STORAGE DEVICES, SUCH AS NETWORK ATTACHED STORAGE PERFORMING DATA MIGRATION, Attorney Docket No. 060692-8066.US01), which claims priority to U.S. Provisional Patent Application No. 61/097,176 filed Sep. 15, 2008 (entitled DATA TRANSFER TECHNIQUES WITHIN DATA STORAGE DEVICES, SUCH AS NETWORK ATTACHED STORAGE PERFORMING DATA MIGRATION, Attorney Docket No. 060692-8066.US00), which are hereby incorporated by reference in their entireties.

BACKGROUND

[0002] Networked attached storage (NAS), often refers to a computing system, attached to a network, which provides file-based data storage services to other devices on the network. A NAS system, or NAS device, may include a file system (e.g., under Microsoft Windows) that manages the data storage services, but is generally controlled by other resources via an IP address or other communication protocol. A NAS device may also include an operating system, although the operating system is often configured only to facilitate operations performed by the NAS system. Mainly, a NAS device includes one or more redundantly arranged hard disks, such as RAID arrays. A NAS device works with various file-based and/or communication protocols, such as NFS (Network File System) for UNIX or LINUX systems, SMB/CIFS (Server Message Block/Common Internet File System) for Windows systems, or iSCSI (Internet SCSI) for IP communications.

[0003] NAS devices provide a few similar functionalities to Storage Area Networks (SANs), although typical NAS devices only facilitate file level storage. Some hybrid systems exist, which provide both NAS and SAN functionalities. However, in these hybrid systems, such as Openfiler on LINUX, the NAS device serves the SAN device at the file level, and not at a file system level, such as at the individual file level. For example, the assignee's U.S. Pat. No. 7,546,324, entitled Systems and Methods for Performing Storage Operations Using Network Attached Storage, describes how individual files in a NAS device can be written to secondary storage, and are replaced in the NAS device with a stub having a pointer to the secondary storage location where the file now resides.

[0004] A NAS device may provide centralized storage to client computers on a network, but may also assist in load balancing and fault tolerance for resources such as email and/or web server systems. Additionally, NAS devices are generally smaller and easy to install to a network.

[0005] NAS device performance generally depends on traffic and the speed of the traffic on the attached network, as well as the capacity of a cache memory on the NAS device. Because a NAS device supports multiple protocols

and contains reduced processing and operating systems, its performance may suffer when many users or many operations attempt to utilize the NAS device. The contained hardware intrinsically limits a typical NAS device, because it is self-contained and self-supported. For example, the capacity of its local memory may limit a typical NAS device's ability to provide data storage to a network, among other problems.

[0006] The need exists for a system that overcomes the above problems, as well as one that provides additional benefits. Overall, the examples herein of some prior or related systems and their associated limitations are intended to be illustrative and not exclusive. Other limitations of existing or prior systems will become apparent to those of skill in the art upon reading the following Detailed Description.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 is a block diagram illustrating components of a data stream utilized by a suitable data storage system.

[0008] FIG. 2 is a block diagram illustrating an example of a data storage system.

[0009] FIG. 3 is a block diagram illustrating an example of components of a server used in data storage operations.

[0010] FIG. 4 is a block diagram illustrating a NAS device within a networked computing system.

[0011] FIG. 5 is a block diagram illustrating the components of a NAS device configured to perform data migration.

[0012] FIGS. 6A and 6B are schematic diagrams illustrating a data store before and after a block-based data migration, respectively.

[0013] FIG. 7 is a flow diagram illustrating a routine for performing block-level data migration in a NAS device.

[0014] FIG. 8 is a flow diagram illustrating a routine for performing chunk-level data migration in a NAS device.

[0015] FIG. 9 is flow diagram illustrating a routine for block-based or chunk-based data restoration and modification via a NAS device.

DETAILED DESCRIPTION

Overview

[0016] Described in detail herein is a system and method that transfers or migrates data objects within a stand-alone network storage device, such as a filer or network-attached storage (NAS) device. In some examples, a NAS device transfers segments, portions, increments, or proper subsets of data objects stored in local memory of the NAS device. The NAS device may transfer portions of files, folders, and other data objects from a cache to secondary storage based on certain criteria, such as time-based criteria, age-based criteria, and so on. A portion may be one or more blocks of a data object, or one or more chunks of a data object, or other data portions that combine to form, store, and/or contain a data object, such as a file.

[0017] In some examples, the NAS device performs block-based migration of data. A data migration component within the NAS device identifies one or more blocks of a data object stored in a cache or data storage that satisfy a certain criteria, and migrates the identified blocks. For example, the data migration component may determine that a certain number of blocks of a file have not been modified or called by a file system within a certain time period, and

migrate these blocks to secondary storage. The data migration component then maintains the other blocks of the file in primary storage. In some cases, the data migration component automatically migrates data without requiring user input. Additionally, the migration may be transparent to a user.

[0018] In some examples, the NAS device performs chunk-based migration of data. A chunk is, for example, a group or set of blocks. One or more chunks may comprise a portion of a file, folder, or other data object. The data migration component identifies one or more chunks of a data object that satisfy a certain criteria, and migrates the identified chunks. For example, the data migration component may determine that a certain number of chunks of a file have not been modified or called by a file system in a certain time period, and migrate these chunks to secondary storage. The system then maintains the other chunks of the file in the cache or data storage of the NAS device.

[0019] Network-attached storage, such as a filer or NAS device, and associated data migration components and processes, will now be described with respect to various examples. The following description provides specific details for a thorough understanding of, and enabling description for, these examples of the system. However, one skilled in the art will understand that the system may be practiced without these details. In other instances, well-known structures and functions have not been shown or described in detail to avoid unnecessarily obscuring the description of the examples of the system.

[0020] The terminology used in the description presented below is intended to be interpreted in its broadest reasonable manner, even though it is being used in conjunction with a detailed description of certain specific examples of the system. Certain terms may even be emphasized below; however, any terminology intended to be interpreted in any restricted manner will be overtly and specifically defined as such in this Detailed Description section. A suitable data storage system will first be described, followed by a description of suitable stand-alone devices. Following that, various data migration and data recovery processes will be discussed.

Suitable System

[0021] Referring to FIG. 1, a block diagram illustrating components of a data stream utilized by a suitable data storage system, such as a system that performs network attached storage, is shown. The stream **110** may include a client **111**, a media agent **112**, and a secondary storage device **113**. For example, in storage operations, the system may store, receive and/or prepare data, such as blocks or chunks, to be stored, copied or backed up at a server or client **111**. The system may then transfer the data to be stored to media agent **112**, which may then refer to storage policies, schedule policies, and/retention policies (and other policies) to choose a secondary storage device **113**, such as a NAS device that receives data and transfers data to attached secondary storage devices. The media agent **112** may include or be associated with a NAS device, to be discussed herein.

[0022] The secondary storage device **113** receives the data from the media agent **112** and stores the data as a secondary copy, such as a backup copy. Secondary storage devices may be magnetic tapes, optical disks, USB and other similar

media, disk and tape drives, and so on. Of course, the system may employ other configurations of stream components not shown in the Figure.

[0023] Referring to FIG. 2, a block diagram illustrating an example of a data storage system **200** is shown. Data storage systems may contain some or all of the following components, depending on the needs of the system. FIG. 2 and the following discussion provide a brief, general description of a suitable computing environment in which the system can be implemented. Although not required, aspects of the system are described in the general context of computer-executable instructions, such as routines executed by a general-purpose computer, e.g., a server computer, wireless device or personal computer. Those skilled in the relevant art will appreciate that the system can be practiced with other communications, data processing, or computer system configurations, including: Internet appliances, network PCs, mini-computers, mainframe computers, and the like. Indeed, the terms “computer,” “host,” and “host computer” are generally used interchangeably herein, and refer to any of the above devices and systems, as well as any data processor.

[0024] Aspects of the system can be embodied in a special purpose computer or data processor that is specifically programmed, configured, or constructed to perform one or more of the computer-executable instructions explained in detail herein. Aspects of the system can also be practiced in distributed computing environments where tasks or modules are performed by remote processing devices, which are linked through a communications network, such as a Local Area Network (LAN), Wide Area Network (WAN), Storage Area Network (SAN), Fibre Channel, or the Internet. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0025] Aspects of the system may be stored or distributed on computer-readable media, including tangible storage media, such as magnetically or optically readable computer discs, hard-wired or preprogrammed chips (e.g., EEPROM semiconductor chips), nanotechnology memory, biological memory, or other data storage media. Alternatively, computer implemented instructions, data structures, screen displays, and other data under aspects of the system may be distributed over the Internet or over other networks (including wireless networks), on a propagated signal on a propagation medium (e.g., an electromagnetic wave(s), a sound wave, etc.) over a period of time, or they may be provided on any analog or digital network (packet switched, circuit switched, or other scheme). Those skilled in the relevant art will recognize that portions of the system reside on a server computer, while corresponding portions reside on a client computer, and thus, while certain hardware platforms are described herein, aspects of the system are equally applicable to nodes on a network.

[0026] For example, the data storage system **200** contains a storage manager **210**, one or more clients **111**, one or more media agents **112**, and one or more storage devices **113**. Storage manager **210** controls media agents **112**, which may be responsible for transferring data to storage devices **113**. Storage manager **210** includes a jobs agent **211**, a management agent **212**, a database **213**, and/or an interface module **214**. Storage manager **210** communicates with client(s) **111**. One or more clients **111** may access data to be stored by the system from database **222** via a data agent **221**. The system uses media agents **112**, which contain databases **231**, to transfer and store data into storage devices **113**. The storage

devices **113** may include network attached storage, such as the NAS devices described herein. Client databases **222** may contain data files and other information, while media agent databases may contain indices and other data structures that assist and implement the storage of data into secondary storage devices, for example.

[0027] The data storage and recovery system may include software and/or hardware components and modules used in data storage operations. The components may be storage resources that function to copy data during storage operations. The components may perform other storage operations (or storage management operations) other than operations used in data stores. For example, some resources may create, store, retrieve, and/or migrate primary or secondary data copies of data. Additionally, some resources may create indices and other tables relied upon by the data storage system and other data recovery systems. The secondary copies may include snapshot copies and associated indices, but may also include other backup copies such as HSM copies, archive copies, auxiliary copies, and so on. The resources may also perform storage management functions that may communicate information to higher level components, such as global management resources.

[0028] In some examples, the system performs storage operations based on storage policies, as mentioned above. For example, a storage policy includes a set of preferences or other criteria to be considered during storage operations. The storage policy may determine or define a storage location and/or set of preferences about how the system transfers data to the location and what processes the system performs on the data before, during, or after the data transfer. In some cases, a storage policy may define a logical bucket in which to transfer, store or copy data from a source to a data store, such as storage media. Storage policies may be stored in storage manager **210**, or may be stored in other resources, such as a global manager, a media agent, and so on. Further details regarding storage management and resources for storage management will now be discussed.

[0029] Referring to FIG. 3, a block diagram illustrating an example of components of a server used in data storage operations is shown. A server, such as storage manager **210**, may communicate with clients **111** to determine data to be copied to storage media. As described above, the storage manager **210** may contain a jobs agent **211**, a management agent **212**, a database **213**, and/or an interface module. Jobs agent **211** may manage and control the scheduling of jobs (such as copying data files) from clients **111** to media agents **112**. Management agent **212** may control the overall functionality and processes of the data storage system, or may communicate with global managers. Database **213** or another data structure may store storage policies, schedule policies, retention policies, or other information, such as historical storage statistics, storage trend statistics, and so on. Interface module **215** may interact with a user interface, enabling the system to present information to administrators and receive feedback or other input from the administrators or with other components of the system (such as via APIs).

Suitable Storage Devices

[0030] Referring to FIG. 4, a block diagram illustrating components of a networked data storage device, such as a filer or NAS device **440**, configured to perform data migration within a networked computing system is shown. (While the examples below discuss a NAS device, any architecture

or networked data storage device employing the following principles may be used, including a proxy computer coupled to the NAS device). The computing system **400** includes a data storage system **410**, such as the tiered data storage system **200**. Client computers **420**, including computers **422** and **424**, are associated with users that generate data to be stored in secondary storage. The client computers **422** and **424** communicate with the data storage system **410** over a network **430**, such as a private network such as an Intranet, a public network such as the Internet, and so on. The networked computing system **400** includes network attached storage, such as NAS device **440**. The NAS device **440** includes NAS-based storage or memory, such as a cache **444**, for storing data received from the network, such as data from client computers **422** and **424**. (The term “cache” is used generically herein for any type of storage, and thus the cache **444** can include any type of storage for storing of data files within the NAS device, such as magnetic disk, optical disk, semiconductor memory, or other known types of storage such as magnetic tape or types of storage hereafter developed.) The cache **444** may include an index or other data structure in order to track where data is eventually stored or the index may be stored elsewhere, such as on the proxy computer. The index may include information associating the data with information identifying a secondary storage device that stored the data, or other information. For example, as described in detail below, the index may include both an indication of which blocks have been written to secondary storage (and where they are stored in secondary storage), and a look up table that maps blocks to individual files stored within the NAS and NAS device **440**.

[0031] The NAS device **440** also includes a data migration component **442** that performs data migration on data stored in the cache **444**. While shown in FIG. 4 as being within the NAS device **440**, the data migration component **442** may be on a proxy computer coupled to the NAS device. In some cases, the data migration component **442** is a device driver or agent that performs block-level data migration of data stored in the cache. In some cases, the data migration component **442** performs chunk-based data migration of data stored in the cache. Additionally, in some cases the data migration component **442** may perform file-based data migration, or a combination of two or more types of data migration, depending on the needs of the system. During data migration, the NAS device transfers data from the cache of the device to one or more secondary storage devices **450** located on the network **430**, such as magnetic tapes **452**, optical disks **454**, or other secondary storage **456**. The NAS device may include various data storage components when identifying and transferring data from the cache **444** to the secondary storage devices **450**. These components will now be discussed.

[0032] Referring to FIG. 5, a block diagram illustrating the components of a NAS device **440** configured to perform data migration is shown. In addition to a data migration component **442** and cache **444**, the NAS device **440** may include an input component **510**, a data reception component **520**, a file system **530**, and an operating system **540**. The input component **510** may receive various inputs, such as via an iSCSI protocol. That is, the NAS device may receive commands or control data from a data storage system **410** over IP channels. For example, the data storage system **410** may send commands to a NAS device's IP address in order to provide instructions to the NAS device. The data reception

component **520** may receive data to be stored over multiple protocols, such as NFS, CIFS, and so on. For example, a UNIX based system may send data to be stored to the NAS device over a NFS communication channel, while a Windows based system may send data to be stored to the NAS device over a CIFS communication channel.

[0033] Additionally, the NAS device **440** may include a number of data storage resources, such as a data storage engine **560** to direct reads from writes to the data store **444**, and one or more media agents **570**. The media agents **570** may be similar to the media agents **112** described herein. In some cases, the NAS device **440** may include two or more media agents **570**, such as multiple media agents **570** externally attached to the NAS device **440**. The NAS device **440** may expand its data storage capabilities by adding media agents **570**, as well as other components.

[0034] As discussed herein, the NAS device **440** includes a data migration component capable of transferring some or all of the data stored in the cache **442**. In some examples, the data migration component **442** requests and/or receives information from a callback layer **550**, or other intermediate component, within the NAS device **440**. Briefly, the callback layer **550** intercepts calls for data between the file system **530** and the cache **444**, and tracks these calls to provide information to the data migration component **442** regarding when data is changed, updated, and/or accessed by the file system **530**. Further details regarding the callback layer **550** and other intermediate components will now be discussed.

[0035] In some examples, the NAS device monitors the transfer of data from the file system **530** to the cache **444** via the callback layer **550**. The callback layer **550** not only facilitates the migration of data portions from data storage on the NAS device to secondary storage, but also facilitates read back or callback of that data from the secondary storage back to the NAS device. While described at times herein as a device driver or agent, the callback layer **550** may be a layer, or additional file system, that resides on top of the file system **530**. The callback layer **550** may intercept data requests from the file system **530**, in order to identify, track and/or monitor data requested by the file system **530** and store information associated with these requests in a data structure, such as a bitmap similar to the one shown in Table 1. Thus, the callback layer stores information identifying when a data portion is accessed by tracking calls from the file system **530** to the cache **530**. For example, Table 1 provides entry information that tracks calls to a data store:

TABLE 1

Chunk of File1	Access Time
File1.1	Sep. 5, 2008 @12:00
File1.2	Sep. 5, 2008 @12:30
File1.3	Sep. 5, 2008 @13:30
File1.4	Jun. 4, 2008 @ 12:30

[0036] In this example, the file system **530** creates a data object named “File1,” using a chunking component (described herein) to divide the file into four chunks: “File1.1,” “File1.2,” “File1.3,” and “File1.4.” The file system **530** stores the four chunks to the cache **444** on 06.04.2008. According to the table, the file system can determine that it has not accessed chunk File1.4 since its creation, and most recently accessed the other chunks on Sep. 5, 2008. Of course, Table 1 may include additional, other or different

information, such as information identifying a location of the chunks, information identifying the type of media storing the chunks, information identifying the blocks within the chunk, and/or other information or metadata.

[0037] Thus, providing data migration to the NAS device enables the device to facilitate inexpensive, transparent storage to a networked computing system, to free up storage space by migrating or archiving stale data to other locations, among other benefits. Of course, non-networked computing systems may also store data to the NAS devices described herein. Because the NAS devices described herein can be easily and quickly installed on networks, they provide users, such as network administrators, with a quick and efficient way to expand their storage capacity without incurring the typical costs associated with typical NAS devices that do not perform data migration.

[0038] For example, adding a NAS device described herein to an existing networked computing system can provide the computing system with expanded storage capabilities, but can also provide the computing system with other data storage functionality. In some examples, the NAS device described herein includes a data storage engine (e.g., a common technology engine, or CTE, provided by Commvault Systems, Inc. of Oceanport, N.J.), the NAS device may act as a backup server. For example, such a device may perform various data storage functions normally provided by a backup server, such as single instancing, data classification, mirroring, content indexing, data backup, encryption, compression, and so on. Thus, in some examples, the NAS device described herein acts as a fully functional and independent device an administrator can attach to a network to perform virtually any data storage function.

[0039] Also, in some cases, the NAS device described herein may act to perform fault tolerance in a data storage system. For example, the clustering of NAS devices on a system may provide a higher level of security, because processes on one device can be replicated on another. Thus, attaching two or more of the NAS devices described herein may provide an administrator with the redundancy or security required in some data storage systems.

Data Migration in Storage Devices

[0040] As described herein, in some examples, the NAS device leverages block-level or chunk-based data migration in order to provide expanded storage capabilities to a networked computing system.

[0041] Block-level migration, or block-based data migration, involves migrating disk blocks from the data store or cache **444** to secondary media, such as storage devices **550**. Using block-level migration, the NAS device **440** transfers blocks from the cache that have not been recently accessed to secondary storage, freeing up space on the cache.

[0042] As described above, the system can transfer or migrate certain blocks of a data object from one data store to another, such as from a cache in a NAS device to secondary storage. Referring to FIGS. 6A-6B, a schematic diagram illustrating contents of two data stores before and after a block-based data migration is shown. In FIG. 6A, a first data store **610** contains primary copies (i.e., production copies) of two data objects, a first data object **620** and a second data object **630**. The first data object comprises blocks A and A¹, where blocks A are blocks that satisfy or meet certain storage criteria (such as blocks that have not been modified since creation or not been modified within a

certain period of time) and blocks A' are blocks that do not meet the criteria (such as blocks that have been modified within the certain time period). The second data object comprises blocks B and B', where blocks B satisfy the criteria and blocks B' do not meet the criteria.

[0043] FIG. 6B depicts the first data store 610 after a block-based data migration of the two data objects 620 and 630. In this example, the system only transfers the data from blocks that satisfy a criteria (blocks A and B) from the first data store 610 to a second data store 640, such as secondary storage 642, 644. The secondary storage may include one or more magnetic tapes, one or more optical disks, and so on. The system maintains data in the remaining blocks (blocks A' and B') within the first data store 610.

[0044] The system can perform file system data migration at a block level, unlike previous systems that only migrate data at the file level (that is, they have a file-level granularity). By tracking migrated blocks, the system can also restore data at the block level, which may avoid cost and time problems associated with restoring data at the file level.

[0045] Referring to FIG. 7, a flow diagram illustrating a routine 700 for performing block-level data migration in a NAS device is shown. In step 710, the NAS device, via the data migration component 442, identifies data blocks within a cache that satisfy a certain criteria. The data migration component 442 may compare some or all of the blocks (or, information associated with the blocks) in the cache with predetermined criteria. The predetermined criteria may be time-based criteria within a storage policy or data retention policy.

[0046] In some examples, the data migration component 442 identifies blocks set to be "aged off" from the cache. That is, the data migration component 442 identifies blocks created, changed, or last modified before a certain date and time. For example, the system may review a cache for all data blocks that satisfy a criterion or criteria. The data store may be an electronic mailbox or personal folders (.pst) file for a Microsoft Exchange user, and the criterion may define, for example, all blocks or emails last modified or changed thirty days ago or earlier. The component 442 compares information associated with the blocks, such as metadata associated with the blocks, to the criteria, and identifies all blocks that satisfy the criteria. For example, the component 442 identifies all blocks in the .pst file not modified within the past thirty days. The identified blocks may include all the blocks for some emails and/or a portion of the blocks for other emails. That is, for a given email (or data object), a first portion of the blocks that include the email may satisfy the criteria, while a second portion of the blocks that include the same email may not satisfy the criteria. In other words, a file or a data object can be divided into parts or portions, and only some of the parts or portions change.

[0047] To determine which blocks have changed, and when, the NAS device can monitor the activity of a NAS device's file system 530 via the callback layer 550. The NAS device may store a data structure, such as a bitmap, table, log, and so on within the cache 444 or other memory in the NAS device or elsewhere, and update the data structure whenever the file system calls the cache 444 to access and update or change data blocks within the cache 444. The callback layer 550 traps commands to the cache 444, where that command identifies certain blocks on a disk for access or modifications, and writes to the data structure the changed blocks and the time of the change. The data structure may

include information such as an identification of changed blocks and a date and a time the blocks were changed. The data structure, which may be a table, bitmap, or group of pointers, such as a snapshot, may also include other information, such as information that maps file names to blocks, information that maps chunks to blocks and/or file names, and so on, and identify when accesses/changes were made. Table 2 provides entry information for tracking the activity of a file system with the "/users" directory:

TABLE 2

Blocks	Date and Time Modified
/users/blocks1-100	Sep. 8, 2008 @14:30
/users/blocks101-105	Sep. 4, 2008 @12:23
/users2/blocks106-110	Sep. 4, 2008 @11:34
/users3/blocks110-1000	Aug. 5, 2008 @10:34

[0048] Thus, if a storage policy identified the time 08.30.2008 @ 12:00 as a threshold time criteria, where data modified after the time is to be retained, the system would identify, in step 710, blocks 110-1000 as having satisfied the criteria. Thus, the system, via the intermediate component 420, can monitor what blocks are requested by a file system, and act accordingly, as described herein.

[0049] In step 720, the NAS device transfers data within the identified blocks from the cache to a media agent 570, to be stored in a different data store. The system may perform some or all of the processes described with respect to FIGS. 1-3 when transferring the data to the media agent. For example, before transferring data, the system may review a storage policy as described herein to select a media agent, such as media agent 112, based on instructions within the storage policy. In step 725, the system optionally updates an allocation table, such as a file allocation table (FAT) for the file system 530 associated with the NAS device, to indicate the data blocks that no longer contain data and are now free to receive and store data from the file system.

[0050] In step 730, via the media agent 570, the NAS device 440 stores data from the blocks to a different data store. In some cases, the NAS device, via the media agent 570, stores the data from the blocks to a secondary storage device, such as a magnetic tape 452 or optical disk 454. For example, the NAS device may store the data from the blocks in secondary copies of the data store, such as a backup copy, an archival copy, and so on.

[0051] The NAS device may create, generate, update, and/or include an allocation table, (such as a table for the data store) that tracks the transferred data and the data that was not transferred. The table may include information identifying the original data blocks for the data, the name of the data object (e.g., file name), the location of any transferred data blocks (including, e.g., offset information), and so on. For example, Table 3 provides entry information for an example .pst file:

TABLE 3

Name of Data Object	Location of data
Email1	C:/users/blocks1-100
Email2.1 (body of email)	C:/users/blocks101-120
Email2.2 (attachment)	X:/remov1/blocks1-250
Email3	X:/remov2/blocks300-500

[0052] In the above example, the data for “Email2” is stored in two locations, the cache (C:/) and an off-site data store (X:/). The system maintains the body of the email, recently modified or accessed, at a location within a data store associated with a file system, “C:/users/blocks101-120.” The system stores the attachment, not recently modified or accessed, in a separate data store, “X:/remov1/blocks1-250.” Of course, the table may include other information, fields, or entries not shown. For example, when the system stored data to tape, the table may include tape identification information, tape offset information, and so on.

[0053] Chunked file migration, or chunk-based data migration, involves splitting a data object into two or more portions of the data object, creating an index that tracks the portions, and storing the data object to secondary storage via the two or more portions. Among other things, the chunk-based migration provides for fast and efficient storage of a data object. Additionally, chunk-based migration facilitates fast and efficient recall of a data object, such as the large files described herein. For example, if a user modifies a migrated file, chunk-based migration enables a data restore component to only retrieve and migrate back to secondary storage the chunk containing the modified portion of the file, and not the entire file.

[0054] As described above, in some examples the NAS device migrates chunks of data (sets of blocks) that comprise a data object from the cache 444 to another. A data object, such as a file, may comprise two or more chunks. A chunk may be a logical division of a data object. For example, a .pst file may include two or more chunks: a first chunk that stores associated with an index of a user’s mailbox, and one or more chunks that stores email, attachments, and so on within the user’s mailbox. A chunk is a proper subset of all the blocks that contain a file. That is, for a file contained or defined by n blocks, the largest chunk of the file contains at most n-1 blocks.

[0055] In some cases, the data migration component 442 may include a chunking component that divides data objects into chunks. The chunking component may receive files to be stored in the cache 444, divide the files into two or more chunks, and store the files as two or more chunks in the cache. The chunking component may update an index that associated information associated with files with the chunks of the file, the data blocks of the chunks, and so on.

[0056] The chunking component may perform different processes when determining how to divide a data object. For example, the chunking component may include indexing, header, and other identifying information or metadata in a first chunk, and include the payload in other chunks. The chunking component may identify and/or retrieve file format or schema information from an index, FAT, NFS, or other allocation table in the file system to determine where certain chunks of a data object reside (such as the first or last chunk of a large file). The chunking component may follow a rules-based process when dividing a data object. The rules may define a minimum or maximum data size for a chunk, a time of creation for data within a chunk, a type of data within a chunk, and so on.

[0057] For example, the chunking component may divide a user mailbox (such as a .pst file) into a number of chunks, based on various rules that assign emails within the mailbox to chunks based on the metadata associated with the emails. The chunking component may place an index of the mailbox

in a first chunk and the emails in other chunks. The chunking component may then divide the other chunks based on dates of creation, deletion or reception of the emails, size of the emails, sender of the emails, type of emails, and so on. Thus, as an example, the chunking component may divide a mailbox as follows:

User1/Chunk1	Index
User1/Chunk2	Sent emails
User1/Chunk3	Received emails
User1/Chunk4	Deleted emails
User1/Chunk5	All Attachments.

Of course, other divisions are possible. Chunks may not necessarily fall within logical divisions. For example, the chunking component may divide a data object based on information or instructions not associated with the data object, such as information about data storage resources, information about a target secondary storage device, historical information about previous divisions, and so on.

[0058] Referring to FIG. 8, a flow diagram illustrating a routine 800 for performing chunk-level data migration in a NAS device is shown. In step 810, the system identifies chunks of data blocks within a data store that satisfy one or more criteria. The data store may store large files (>50 MB), such as databases associated with a file system, SQL databases, Microsoft Exchange mailboxes, virtual machine files, and so on. The system may compare some or all of the chunks (or, information associated with the chunks) of the data store with predetermined and/or dynamic criteria. The predetermined criteria may be time-based criteria within a storage policy or data retention policy. The system may review an index with the chunking component 815 when comparing the chunks with applicable criteria.

[0059] In step 820, the NAS device transfers data within the identified chunks from the data store to a media agent, to be stored in a different data store. The NAS device may perform some or all of the processes described with respect to FIGS. 1-3 when transferring the data to the media agent. For example, the NAS device may review a storage policy assigned to the data store and select a media agent based on instructions within the storage policy. In step 825, the system optionally updates an allocation table, such as a file allocation table (FAT) for a file system associated with the NAS device, to indicate the data blocks that no longer contain data and are now free to receive and store data from the file system.

[0060] In step 930, via one or more media agents 570, the NAS device 440 stores the data from the chunks to a different data store. In some cases, the system, via the media agent, stores the data to a secondary storage device, such as a magnetic tape or optical disk. For example, the system may store the data in secondary copies of the data store, such as a backup copy, and archive copy, and so on.

Data Recovery in Storage Devices

[0061] A data storage system, using a NAS device leveraging the block-based or chunk-based data migration processes described herein, is able to restore portions of files instead of entire files, such as individual blocks or chunks that comprise portions of the files. Referring to FIG. 9, a flow diagram illustrating a routine 900 for block-based or chunk-based data restoration and modification is shown. In

step **910**, the system, via a restore or data recovery component, receives a request to modify a file located in a cache of a NAS device or in secondary storage in communication with a NAS device. For example, a user submits a request to a file system to provide an old copy of a large PowerPoint presentation so the user can modify a picture located on slide **5** of **300** of the presentation.

[0062] In step **920**, the system identifies one or more blocks or one or more chunks associated with the request. For example, the callback layer **550** of the system looks to a table similar to Table 3, identifies blocks associated with page **5** of the presentation and blocks associated with a table of contents of the presentation, and contacts a NAS device that stored or migrated the blocks on secondary storage.

[0063] In step **930**, the system, via the NAS device, retrieves the identified blocks or chunks from the secondary storage and presents them to the user. For example, the system only retrieves page **5** and table of contents of the presentation and presents the pages to the user.

[0064] In step **940**, the system receives input from a user to modify the retrieved blocks or chunks. For example, the user updates the PowerPoint presentation to include a different picture. In step **950**, the system transfers data associated with the modified blocks or chunks back to the NAS device, where it remains in a cache or is transferred to secondary storage. For example, the system transfers the modified page **5** to the data store. The system may also update a table that tracks access to the data store, such as Table 1 or Table 3.

[0065] Thus, the system, leveraging block-based or chunk-based data migration in a NAS device, restores only portions of data objects required by a file system. Such restoration can be, among other benefits, advantageous over systems that perform file-based restoration, because those systems restore entire files, which can be expensive, time consuming, and so on. Some files, such as .pst files, may contain large amounts of data. File-based restoration can therefore be inconvenient and cumbersome, among other things, especially when a user only requires a small portion of a large file.

[0066] For example, a user submits a request to the system to retrieve an old email stored in a secondary copy on removable media via a NAS device. The system identifies a portion of a .pst file associated with the user that contains a list of old emails in the cache of the NAS device, and retrieves the list. That is, the system has knowledge of the chunk that includes the list (e.g., a chunking component may always include the list in a first chunk of a data object), accesses the chunk, and retrieves the list. The other portions (e.g., all the emails with the .pst file), were transferred from the NAS device secondary storage. The user selects the desired email from the list. The NAS device, via an index in the cache that associates chunks with data or files (such as an index similar to Table 3), identifies the chunk that contains the email, and retrieves the chunk from associated secondary storage for presentation to the user. Thus, the NAS device is able to restore the email without restoring the entire mailbox (.pst file) associated with the user.

[0067] As noted above, the callback layer **550** maintains a data structure that not only tracks where a block or chunk resides on secondary storage, but also which file was affected based on the migration of that block or chunk. Portions of large files may be written to secondary storage to free up space in the data store **444** of the NAS device **440**.

Thus, to the network, the total data storage of the NAS device is much greater than that actually available within the data store **444**. For example, while the data store **444** may have only a 100 gigabyte capacity, its capacity may actually appear as 300 gigabytes, with over 200 gigabytes migrated to secondary storage.

[0068] To help ensure sufficient space to write back data from secondary storage to the data store **444** of the NAS device **440**, the data store may be partitioned to provide a callback or read-back cache. For example, a disk cache may be established in the data store **444** of the NAS device **440** for the NAS device to write back data read from secondary storage. The amount of the partition is configurable, and may be, for example, between 5 and 20 percent of the total capacity of the data store **440**. In the above example, with a 100 gigabyte data store **444**, 10 gigabytes may be reserved (10 percent) for data called back from secondary storage to the NAS device **440**. This disk partition or callback cache can be managed in known ways, such that data called back to this disk partition can have the oldest data overwritten when room is needed to write new data.

CONCLUSION

[0069] From the foregoing, it will be appreciated that specific examples of the data recovery system have been described herein for purposes of illustration, but that various modifications may be made without deviating from the spirit and scope of the system. For example, although files have been described, other types of content such as user settings, application data, emails, and other data objects can be imaged by snapshots. Accordingly, the system is not limited except as by the appended claims.

[0070] Unless the context clearly requires otherwise, throughout the description and the claims, the words “comprise,” “comprising,” and the like are to be construed in an inclusive sense, as opposed to an exclusive or exhaustive sense; that is to say, in the sense of “including, but not limited to.” The word “coupled”, as generally used herein, refers to two or more elements that may be either directly connected, or connected by way of one or more intermediate elements. Additionally, the words “herein,” “above,” “below,” and words of similar import, when used in this application, shall refer to this application as a whole and not to any particular portions of this application. Where the context permits, words in the above Detailed Description using the singular or plural number may also include the plural or singular number respectively. The word “or” in reference to a list of two or more items, that word covers all of the following interpretations of the word: any of the items in the list, all of the items in the list, and any combination of the items in the list.

[0071] The above detailed description of embodiments of the system is not intended to be exhaustive or to limit the system to the precise form disclosed above. While specific embodiments of, and examples for, the system are described above for illustrative purposes, various equivalent modifications are possible within the scope of the system, as those skilled in the relevant art will recognize. For example, while processes or blocks are presented in a given order, alternative embodiments may perform routines having steps, or employ systems having blocks, in a different order, and some processes or blocks may be deleted, moved, added, subdivided, combined, and/or modified. Each of these processes or blocks may be implemented in a variety of

different ways. Also, while processes or blocks are at times shown as being performed in series, these processes or blocks may instead be performed in parallel, or may be performed at different times.

[0072] The teachings of the system provided herein can be applied to other systems, not necessarily the system described above. The elements and acts of the various embodiments described above can be combined to provide further embodiments.

[0073] These and other changes can be made to the system in light of the above Detailed Description. While the above description details certain embodiments of the system and describes the best mode contemplated, no matter how detailed the above appears in text, the system can be practiced in many ways. Details of the system may vary considerably in implementation details, while still being encompassed by the system disclosed herein. As noted above, particular terminology used when describing certain features or aspects of the system should not be taken to imply that the terminology is being redefined herein to be restricted to any specific characteristics, features, or aspects of the system with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the system to the specific embodiments disclosed in the specification, unless the above Detailed Description section explicitly defines such terms. Accordingly, the actual scope of the system encompasses not only the disclosed embodiments, but also all equivalent ways of practicing or implementing the system under the claims.

[0074] While certain aspects of the system are presented below in certain claim forms, the applicant contemplates the various aspects of the system in any number of claim forms. For example, while only one aspect of the system is recited as a means-plus-function claim under 35 U.S.C. sec. 112, sixth paragraph, other aspects may likewise be embodied as a means-plus-function claim, or in other forms, such as being embodied in a computer-readable medium. (Any claims intended to be treated under 35 U.S.C. § 112, will begin with the words “means for”.) Accordingly, the applicant reserves the right to add additional claims after filing the application to pursue such additional claim forms for other aspects of the system.

1-20. (canceled)

21. A network attached storage device comprising:

a first storage device having computer-executable instructions stored thereon; and

one or more processors that, having executed the computer-executable instructions, configure the network attached storage device to:

store a first data file in the first storage device;

update a first data structure with a plurality of entries based on the first data file, wherein:

each entry of the plurality of entries identifies a portion of the first data file and a modification time of the portion; and

each portion is less than all of the first data file;

identify a first one or more portions of the first data file for migrating to a second storage device based on a comparison of a modification time criterion with each modification time for the identified one or more portions;

migrate the identified one or more portions from the first storage device to a second storage device outside of the network attached storage device; and

identify space in the first storage device previously used by the migrated one or more portions as free for storing a second data file.

22. The network attached storage device of claim **21**, wherein the one or more processors are further configured to:

update a second data structure with a logical location for each portion of the first data file.

23. The network attached storage device of claim **22**, wherein a first logical location for a first portion of the first data file is on the network attached storage device and a second logical location for a second portion of the first data file is on the second storage device.

24. The network attached storage device of claim **21**, wherein the one or more processors are further configured to:

receive a request to modify the first data file;

identify the one or more portions of the first data file that have been migrated to the second storage device based on the received request;

retrieve the identified one or more portions of the first data file that have been migrated to the second storage device; and

present content of the retrieved one or more portions for modification.

25. The network attached storage device of claim **21**, wherein one or more portions are represented as one or more file system blocks based on a file system implemented on the network attached storage device.

26. The network attached storage device of claim **21**, wherein one or more portions are represented as one or more chunks, wherein a size of each chunk is determined from a rule-based process of dividing the first data file into the one or more portions.

27. The network attached storage device of claim **21**, wherein the one or more processors are further configured to:

partition the first data storage device to obtain a read-back cache;

retrieve one or more of the portions of the first data file that have been migrated to the second storage device; and

store the one or more of the retrieved portions of the first data file to the read-back cache, wherein the contents of the retrieved portions are presented to a user after the one or more of the retrieved portions of the first data file have been stored in the read-back cache.

28. A method comprising:

storing a first data file in a first storage device of a network attached storage device;

updating a first data structure with a plurality of entries based on the first data file, wherein:

each entry of the plurality of entries identifies a portion of the first data file and a modification time of the portion; and

each portion is less than all of the first data file;

identifying a first one or more portions of the first data file for migrating to a second storage device based on a comparison of a modification time criterion with each modification time for the identified one or more portions;

migrating the identified one or more portions from the first storage device to a second storage device outside of the network attached storage device; and

identifying space in the first storage device previously used by the migrated one or more portions as free for storing a second data file.

29. The method of claim **28**, further comprising: updating a second data structure with a logical location for each portion of the first data file.

30. The method of claim **30**, wherein a first logical location for a first portion of the first data file is on the network attached storage device and a second logical location for a second portion of the first data file is on the second storage device.

31. The method of claim **28**, further comprising: receiving a request to modify the first data file; identifying the one or more portions of the first data file that have been migrated to the second storage device based on the received request; retrieving the identified one or more portions of the first data file that have been migrated to the second storage device; and presenting content of the retrieved one or more portions for modification.

32. The method of claim **28**, wherein one or more portions are represented as one or more file system blocks based on a file system implemented on the network attached storage device.

33. The method of claim **28**, wherein one or more portions are represented as one or more chunks, wherein a size of each chunk is determined from a rule-based process of dividing the first data file into the one or more portions.

34. The method of claim **28**, further comprising: partitioning the first data storage device to obtain a read-back cache; retrieving one or more of the portions of the first data file that have been migrated to the second storage device; and storing the one or more of the retrieved portions of the first data file to the read-back cache, wherein the contents of the retrieved portions are presented to a user after the one or more of the retrieved portions of the first data file have been stored in the read-back cache.

35. A non-transitory, computer-readable media having computer-executable instructions stored thereon that, when executed by one or more processors, configures a network attached storage device to perform a plurality of operations comprising:

storing a first data file in a first storage device of a network attached storage device;
 updating a first data structure with a plurality of entries based on the first data file, wherein:
 each entry of the plurality of entries identifies a portion of the first data file and a modification time of the portion; and
 each portion is less than all of the first data file;

identifying a first one or more portions of the first data file for migrating to a second storage device based on a comparison of a modification time criterion with each modification time for the identified one or more portions;

migrating the identified one or more portions from the first storage device to a second storage device outside of the network attached storage device; and

identifying space in the first storage device previously used by the migrated one or more portions as free for storing a second data file.

36. The non-transitory, computer-readable media of claim **35**, wherein the plurality of operations further comprises: updating a second data structure with a logical location for each portion of the first data file.

37. The non-transitory, computer-readable media of claim **36**, wherein a first logical location for a first portion of the first data file is on the network attached storage device and a second logical location for a second portion of the first data file is on the second storage device.

38. The non-transitory, computer-readable media of claim **35**, wherein the plurality of operations further comprises: receiving a request to modify the first data file;

identifying the one or more portions of the first data file that have been migrated to the second storage device based on the received request;

retrieving the identified one or more portions of the first data file that have been migrated to the second storage device; and

presenting content of the retrieved one or more portions for modification.

39. The non-transitory, computer-readable media of claim **35**, wherein one or more portions are represented as one or more file system blocks based on a file system implemented on the network attached storage device.

40. The non-transitory, computer-readable media of claim **35**, wherein the plurality of operations further comprises:

partitioning the first data storage device to obtain a read-back cache;

retrieving one or more of the portions of the first data file that have been migrated to the second storage device; and

storing the one or more of the retrieved portions of the first data file to the read-back cache, wherein the contents of the retrieved portions are presented to a user after the one or more of the retrieved portions of the first data file have been stored in the read-back cache.

* * * * *