

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
3 July 2008 (03.07.2008)

PCT

(10) International Publication Number
WO 2008/080143 A2

(51) International Patent Classification:
G06F 17/30 (2006.01)

[IN/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US). **PRASAD, Arun** [—/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US).

(21) International Application Number:

PCT/US2007/088736

(74) Agents: **DALEY-WATSON, Christopher, J.** et al.; Perkins Coie LLP, P.O. Box 1247, Seattle, WA 98111-1247 (US).

(22) International Filing Date:

21 December 2007 (21.12.2007)

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

60/871,735 22 December 2006 (22.12.2006) US
11/931,034 31 October 2007 (31.10.2007) US

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(71) Applicant (*for all designated States except US*): **COM-MVAULT SYSTEMS, INC.** [US/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US).

(72) Inventors; and

(75) Inventors/Applicants (*for US only*): **KOT-TOMTHARAYIL, Rajiv** [IN/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US). **MAY, Andreas** [US/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US). **BROCKWAY, Brian** [US/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US). **MULLER, Marcus, B.** [US/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US). **PRAHLAD, Anand** [US/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US). **KAVURI, Srinivas**

Published:

— *without international search report and to be republished upon receipt of that report*

(54) Title: METHOD AND SYSTEM FOR SEARCHING STORED DATA

(57) Abstract: A complete document management system is disclosed. Accordingly, systems and methods for managing data associated with a data storage component coupled to multiple computers over a network are disclosed. Systems and methods for managing data associated with a data storage component coupled to multiple computers over a network are further disclosed. Additionally, systems and methods for accessing documents available through a network, wherein the documents are stored on one or more data storage devices coupled to the network, are disclosed.



WO 2008/080143 A2

METHOD AND SYSTEM FOR SEARCHING STORED DATA

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims the benefit of U.S. Provisional Application No. 60/871,735, entitled "METHOD AND SYSTEM FOR SEARCHING STORED DATA," filed December 22, 2006, and U.S. Application No. 11/931,034, entitled "METHOD AND SYSTEM FOR SEARCHING STORED DATA," filed October 31, 2007, each of which is incorporated by reference herein.

BACKGROUND

[0002] Data protection systems contain large amounts of data. This data includes personal data, such as financial data, customer/client/patient contact data, audio/visual data, and much more. Corporate computer systems often contain word processing documents, engineering diagrams, spreadsheets, business strategy presentations, and so on. With the proliferation of computer systems and the ease of creating content, the amount of content in an organization has expanded rapidly. Even small offices often have more information stored than any single employee can know about or locate.

[0003] Some data protection applications provide functions for actively searching for files within the organization based on a previously created index of the information available in each file. A user can then search for and retrieve documents based on a topic. Typical search software operates on a single index of keywords derived from the data that has been copied for protection purposes. It is typical for an organization to maintain many secondary copies of its data and the various copies are typically stored in multiple formats in multiple devices. For example, when current copy of data is made, previous copies are often maintained so that an historical archive is created. Thus, if the most recent copy does not have the desired data for a restore operation, an older copy may be used. With the existence of multiple copies on multiple devices spanning weeks, months and even years, a search over this data can be complex and time consuming. A search over such a large amount of data can require separately searching content indices of all of the computer systems

within an organization. This can put an unexpected load on already burdened systems and can require significant time on the part of a system operator.

[0004] Typical search systems also create problems when retrieval of the desired data is attempted. First, typical systems require that retrieval of the identified data be performed as a restore operation. The typical restore operation first identifies a secondary copy of the data in question on a secondary volume and copies the identified copy of the data back onto a production server (or other primary or working volume) and overwrites the existing data files. This can be inconvenient if it is desired to maintain the production copy or if it is merely desired to inspect the contents of a secondary data store. Second, typical systems are blind to the security rights of users and database operators. Typical systems do not have an integrated data rights security control that identifies the security privileges of the operator or user for whom the data is being restored and allows or denies the restore accordingly. Additionally, typical systems do not allow a user to promote and reapply search criteria throughout the data management system.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] Figure 1 illustrates an example of a group of platforms and data types for searching.

[0006] Figure 2 is a block diagram that illustrates a hierarchical data storage system.

[0007] Figure 3 is a block diagram that illustrates components of a storage operations cell.

[0008] Figure 4 is a block diagram that illustrates interaction between a global cell and data storage cells.

[0009] Figure 5 is a block diagram that illustrates flow of data through the system.

[0010] Figure 6 is a flow diagram that illustrates processing of a content indexing component of the system.

[0011] Figure 7 is a flow diagram that illustrates processing of an index searching component of the system.

- [0012]** Figure 8 illustrates a client selection interface for searching.
- [0013]** Figure 9 illustrates a query construction interface for searching.
- [0014]** Figure 10 illustrates a search summary.
- [0015]** Figure 11 illustrates a results display in an interface for searching.
- [0016]** In the drawings, the same reference numbers and acronyms identify elements or acts with the same or similar functionality for ease of understanding and convenience.

DETAILED DESCRIPTION

[0017] The invention 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 invention. However, one skilled in the art will understand that the invention 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 invention.

[0018] 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 invention. 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.

[0019] Figure 1 illustrates a summary example of a group of platforms and data types that can be searched. As illustrated and as described in more detail herein, a search can be performed over any platform, over any data type, and for documents having been created over any period of time. As illustrated, the system described herein can operate to archive and search data files including, for example, word processing documents 101, email correspondence 102, and database files 103. These files and documents can exist as online copies 105, backup copies 110, and archive copies 115. Thus, the systems and methods described herein can be used to search for and locate virtually any document that has ever existed on an institutional system, whether it currently exists or existed at any time in the past.

These various data types and platform types can coexist in and be operated on in a hierarchical data storage system.

Suitable System

[0020] Referring to Figure 2, a block diagram illustrating a hierarchical data storage system comprises two levels: a storage operations level 210 and a global level 250. The global level 250 may contain a global operations cell 260, which may contain a global manager 261 and a database 262. The storage operations level 210 may contain storage operations cells, such as cells 220 and 230. Cells 220 and 230 may perform specified data storage operations, or may perform varied data storage operations that depend on the needs of the system.

[0021] Cell 220 contains components used in data storage operations, such as a storage manager 221, a database 222, a client 223, and a primary storage database 224. Cell 230 may contain similar components, such as storage manager 231, a database 232, a client 233, and a primary storage database 234. In this example, cell 230 also contains media agent 235 and secondary database 236. Both cells 220 and 230 communicate with global manager 260, providing information related to the data storage operations of their respective cells.

[0022] Referring to Figure 3, a block diagram illustrating components of a storage operations cell is shown. Storage operations cells (such as cells 220 or 230 of Figure 2) may contain some or all of the following components, depending on the use of the cell and the needs of the system. For example, cell 300 contains a storage manager 310, clients 320, multiple media agents 330, and multiple storage devices 340. Storage manager 310 controls media agents 330, which are responsible, at least in part, for transferring data to storage devices 340. Storage manager 310 includes a jobs agent 311, a management agent 312, a database 313, and an interface module 314. Storage manager 310 communicates with client 320. Client 320 accesses data to be stored by the system from database 322 via a data agent 321. The system uses media agents 330, which contain databases 331, to transfer and store data into storage devices 340.

[0023] Cells 300 may include software and/or hardware components and modules used in data storage operations. The cells 300 may be transfer cells that function to transfer data during data store operations. The cells 300 may perform

other storage operations in addition to operations used in data transfers. For example, cells 300 may perform creating, storing, retrieving, and/or migrating primary and secondary data copies. The data copies may include snapshot copies, secondary copies, hierarchical storage manager copies, archive copies, and so on. The cells 300 may also perform storage management functions that may push information to higher level cells, including global manager cells.

[0024] In some embodiments, the system can be configured to perform a storage operation based on one or more storage policies. A storage policy may be, for example, a data structure that includes a set of preferences or other criteria considered during storage operations. The storage policy may determine or define a storage location, a relationship between components, network pathways, accessible datapipes, retention schemes, compression or encryption requirements, preferred components, preferred storage devices or media, and so on. Storage policies may be stored in storage manager 310, 221, 231, or may be stored in global manager 261, as discussed above.

[0025] Additionally or alternatively, the system may implement or utilize schedule policies. A schedule policy may specify when to perform storage operations, how often to perform storage operations, and so on. The schedule policy may also define the use of sub-clients, where one type of data (such as email data) is stored using one sub-client, and another type of data (such as database data) is stored using another sub-client. In these cases, storage operations related to specific data types (email, database, and so on) may be distributed between cells.

[0026] Referring to Figure 4, a block diagram illustrating interaction between the global cell and data storage cells is shown. Global server 100, which may contain global load components, global filter components, and other components configured to determine actions based on received data storage information, may communicate with a database 420 and a user interface 410. Database 420 may store storage policies, schedule policies, received sample data, other storage operation information, and so on. User interface 410 may display system information to a user. Further details with respect to the user interface display are discussed below.

[0027] Global server 100 may push data to a management server 442. Server 442 communicates with a database 445 and clients 451, 452 and/or 453. Data storage servers 430 push data to the global server 100, and contain data agents 432

and can communicate with databases 435. These servers may communicate with clients 454, 455, and/or 456.

[0028] Global server 100 can be configured to perform actions (such as redistributing storage operations), and apply these actions to the data storage system via a management server. Global server 100 receives information used to determine the actions from the storage servers 430. In this example, the global server acts as a hub in the data storage system by sending information to modify data storage operations and monitoring the data storage operations to determine how to improve the operations.

Index Searching

[0029] The hierarchical storage system described herein can be used for searching multiple indices of content, retrieving the identified data in accordance with integrated data security policies, and applying the search criteria as a data management policy. Some or all of these functions can be performed via a simple interface accessed, e.g., from a web browser.

[0030] The content indices searched can be created by a content indexing system. Indices of this data can be created using any known technique including those described in the assignee's co-pending application serial no. 11/694,869 entitled "Method and System for Offline Indexing of Content and Classifying Stored Data" (Attorney Docket No. 60692-8046), the contents of which are herein incorporated by reference.

[0031] The content indexing system can create an index of an organization's content by examining files generated from routine secondary copy operations performed by the organization. The content indexing system can index content from current secondary copies of the system as well as older copies that contain data that may no longer be available on the organization's network. For example, the organization may have secondary copies dating back several years that contain older data that is no longer available, but may still be relevant to the organization. The content indexing system may associate additional properties with data that are not part of traditional indexing of content, such as the time the content was last available or user attributes associated with the content. For example, user attributes such as a project name with which a data file is associated may be stored.

[0032] Members of the organization can search the created index to locate content on a secondary storage device that is no longer online. For example, a user can search for content related to a project that was cancelled a year ago. In this way, content indexing is not affected by the availability of the system that is the original source of the content and users can find additional organization data that is not available in traditional content indexing systems.

[0033] In some embodiments, members of the organization can search for content within the organization independent of the content's source through a single, unified user interface, which may be available through a web browser. For example, members may search for content that originated on a variety of computer systems within the organization. Members may also search through any copy of the content including any primary, secondary, and/or tertiary or auxiliary copies of the content.

[0034] In some embodiments, the content indexing system searches for content based on availability information related to the content. For example, a user may search for content available during a specified time period, such as email received during a particular month. A user may also search specifically for content that is no longer available, such as searching for files deleted from the user's primary computer system. The user may perform a search based on the attributes described above, such as a search based on the time an item was deleted or based on a project with which the item was associated. A user may also search based on keywords associated with user attributes, such as searching for files that only an executive of the organization would have access to, or searching for files tagged as confidential.

[0035] Figure 5 is a block diagram that illustrates the procedural flow of data, in one embodiment. Content is initially stored on a data server 505 that may be a user computer, data warehouse server, or other information store accessible via a network. The data is accessed by a secondary copy manager 510 to perform a regular copy of the data. Secondary copies of data are stored in a secondary copy data store 515 such as a network attached storage device or secondary copy server. The secondary copy data store 515 provides the data to the content indexing system 520 to perform the functions described above. As illustrated in the diagram, because the content indexing system 520 works with a copy of the data, the original data server 505 is not negatively impacted by the operations of the content indexing system 520. Search system 525 can operate on the data in the content indexing

system 520 to provide search functionality for the data having been stored in the secondary copy data store 515.

[0036] Figures 6-7 are representative flow diagrams that depict processes used in some embodiments. These flow diagrams do not show all functions or exchanges of data, but instead they provide an understanding of commands and data exchanged under the system. Those skilled in the relevant art will recognize that some functions or exchange of commands and data may be repeated, varied, omitted, or supplemented, and other (less important) aspects not shown may be readily implemented.

[0037] Figure 6 is a flow diagram that illustrates the processing of a content indexing component for later searching, according to one embodiment. The component is invoked when new content is available or additional content is ready to be added to the content index. In step 610, the component selects a copy of the data to be indexed. For example, the copy may be a secondary copy of the data or a data snapshot. In step 620, the component identifies content within the copy of the data. For example, the component may identify data files such as word processing documents, spreadsheets, and presentation slides within the secondary data store. In step 630, the component updates an index of content to make the content available for searching. For example, the component may add information such as the location of the content, keywords found within the content, and other supplemental information about the content that may be helpful for locating the content during a search. After step 630, these steps conclude.

[0038] Figure 7 is a flow diagram that illustrates the processing of an index searching component of the system, in one embodiment. In step 710, the component receives a search request specifying criteria for finding matching target content. For example, the search request may specify one or more keywords that will be found in matching documents. The search request may also specify boolean operators, regular expressions, and other common search parameters to identify relationships and precedence between terms within the search query. The search request may also specify data stores to be searched. The request may specify that the search is to include one or more of an original copy, a primary secondary copy, and secondary or auxiliary copies of the content. As described in more detail below, in some embodiments, a user may be provided with an interface by which to select

one or more classes of data stores for search. In some embodiments, an interface may be provided by which a user can specify a security clearance and corresponding operators. For example, a user could form a search query for all documents on a certain class of data store having medium security or higher clearance.

[0039] In step 720, the component searches the content index to identify matching content items that are added to a set of search results. For example, the component may identify documents containing specified keywords or other criteria and add these to a list of search results. In step 730, the component selects a first or next search result. In decision step 740, if the search results indicate that the identified content is offline, then the component continues at step 750, else the component continues at step 760. For example, the content may be offline because it is on a tape that has been sent to an offsite storage location. In step 470, the component retrieves the archived content. Additionally or alternatively, the component may provide an estimate of the time required to retrieve the archived content and add this information to the selected search result. In step 760 the component provides the search results in response to the search query. For example, the user may receive the search results through a web browser interface that lists the search results or the search results may be provided to another component for additional processing through an application programming interface (API). After step 760, these steps conclude.

Federated Search

[0040] The search described herein can include indices of data, where the data is a snapshot, primary copy, secondary copy, auxiliary copy, and so on. An organization may have several copies of data available on different types of media. Data may be available on, for example, a tape, on a secondary copy server, or through network attached storage.

[0041] The search capability can be extended to handle an end-user based search via a web interface, a user-based search (*e.g.*, all files that can belong to "Bob" or that can be viewed by "Bob"), search results across several application types (*e.g.*, file copies, Microsoft Exchange mailbox copies, Microsoft Exchange data agents, Microsoft Exchange public folders, etc.) and search results across multiple computers.

[0042] Using a graphical user interface, search criteria can be provided to specify data that is stored on any number and type of volumes and any type of data. An interface such as the interface 800 illustrated in Figure 8 can be used to specify a search term 801 and one or more clients or volumes to search. As illustrated in Figure 8, a list of available clients 805 can be presented. A set of controls 810 can be used to select one or more of the available clients. Selected clients can be displayed in region 815. Variations on this embodiment of the interface can be used to allow a user to select various volumes for the search. For example, the interface can allow a user to specify that the search is to be over the original copy, a primary secondary copy, and secondary or auxiliary copies of the content. The interface can also be configured to allow the user to specify that the search is to include file contents. An exemplary interface for allowing this option and receiving additional related parameters from a user can include an enabling check box 820 for searching in files, a search by field 825, a file name field 830, and a folder path 835 field.

[0043] The search criteria can also specify that the data be from any of multiple applications or of any type. An example of an interface for receiving additional search parameters is shown in Figure 9. The search interface 900 can include fields for a search term 905, file name 906, file size 907, folder 908, modification date 909, email subject 910, email sender 911, email recipient 912, folder 913, date of receipt 914, and various advanced options such as client 915, iDA 916, owner 917, accessibility 918, sample 919, indexing time 920, and time zone 921.

[0044] Through the same interface or a separate interface, the user can also select the various types of application data to be searched. The graphical interface for performing the search can provide an efficient means for a user to enter search terms and perform that search over multiple volumes and data types. For example, the interface can provide check boxes or other population routines for identifying hardware or resources and display the list whereby a user can select specific volumes by name or address or whereby a user can select volumes by type or classification. Similarly, a user may be prompted to specify data types or classes.

[0045] In some embodiments, the search performed over multiple secondary copies and physical devices will be made with reference to metadata stored in one or more metabases or other forms of databases. A data collection agent may traverse a network file system and obtain certain characteristics and other attributes of data in the system. In some embodiments, such a database may be a collection of metadata and/or other information regarding the network data and may be referred to herein as a metabase. Generally, metadata refers to data or information about data, and may include, for example, data relating to storage operations or storage management, such as data locations, storage management components associated with data, storage devices used in performing storage operations, index data, data application type, or other data. Operations can be performed on this data using any known technique including those described in the assignee's co-pending application serial no. 11/564,119 entitled "Systems and Methods for Classifying and Transferring Information in a Storage Network" (Attorney Docket No. 60692-8029) the contents of which are herein incorporated by reference.

[0046] With this arrangement, when a search over multiple secondary copies is to be performed, a system administrator or system process may simply consult the metabase for such information rather than iteratively access and analyze each data item in the network. This approach significantly reduces the amount of time required to obtain data object information by substantially reducing or eliminating the need to obtain information from the source data, and furthermore reduces or minimizes the involvement of network resources in this process, thereby reducing the processing burden on the host system.

[0047] In some embodiments, a query may be received by the system for certain information. This request may be processed and analyzed by a manager module or

other system process that determines or otherwise identifies which metabase or metabases within the system likely include at least some of the requested information. For example, the query itself may suggest which metabases to search and/or the management module may consult an index that contains information regarding metabase content within the system. The identification process may include searching and identifying multiple computing devices within an enterprise or network that may contain information satisfying search criteria.

[0048] A processor can be configured to search metabases or other indices corresponding to multiple volumes and data stores to identify an appropriate data set that may potentially have information related to the query. This may involve performing iterative searches that examine results generated by previous searches and subsequently searching additional, previously unidentified metabases to find responsive information that may not have been found during the initial search. Thus, the initial metabase search may serve as a starting point for searching tasks that may be expanded based on returned or collected results. The returned results may be optionally analyzed for relevance, arranged, and placed in a format suitable for subsequent use (e.g., with another application), or suitable for viewing by a user and reported.

[0049] Once a search has been performed and at least one document or other discrete data item identified, a list of the identified documents or data items can be provided. An example interface 1000 for displaying the results of an email search is illustrated in Figure 10. The interface 1000 can include a summary area 1005 with summary information as well as a search results section 1010.

[0050] In some further embodiments, the one or more identified documents can be retrieved without performing a restore of the data back to the production volume. Such a transfer may involve copying data objects and metadata from one data store and metabase to another, or in some embodiments, may involve migrating the data from its original location to a second location and leaving a pointer or other reference to the second location so the moved information may be quickly located from information present at the original location.

[0051] In some embodiments, a preview pane can be provided so that a user can view at least a portion of the contents of the identified file. One such preview pane 1100 is illustrated in Figure 11. This preview can be provided before any

restore or retrieve operation is executed. In some embodiments, the preview can be generated by reading the identified file from the original data store and displaying the contents on the screen. In other embodiments, the identified file can be copied to a local disk and the preview generated based on file as it resides on a local disk. In some embodiments, the interface can display a portion of content 1105 from the data file returned by the search query and, in some further embodiments, prompt a user to refine the search. Data retrieval can also be performed using any known technique including those described in the assignee's co-pending application serial no. 11/694,890 entitled "System and Method for Data Retrieval, Including Secondary Copy Precedence Optimizations" (Attorney Docket No. 60692-8039), the contents of which are herein incorporated by reference.

Data Management Policy Integration

[0052] In some embodiments, the search criteria provided by a user as part of a search can later be applied as a data management policy. For example, a user could develop search terms that identify a certain set of data files. These search terms can then be stored as a data management policy which can then be applied at any other point in the data storage system. A data management policy created in this manner can be a data structure or other information source that includes a set of preferences and other storage criteria associated with performing a storage operation. The data management policy created based on a user-supplied search criteria can also be used as part of a schedule policy.

[0053] A schedule policy may specify when to perform storage operations and how often, and may also specify performing certain storage operations on sub-clients of data and how to treat those sub-clients. A sub-client may represent static or dynamic associations of portions of data of a volume and are typically mutually exclusive. Thus, a portion of data may be given a label and the association is stored as a static entity in an index, database or other storage location used by the system. Sub-clients may also be used as an effective administrative scheme of organizing data according to data type, department within the enterprise, storage preferences, etc. The search criteria provided by a user can be used as a file selector in connection with any schedule policy.

[0054] In some embodiments, the data management policy can include various storage preferences, for example, those expressed by a user preference or storage

policy. As non-limiting examples, the data management policy can specify a storage location, relationships between system components, network pathway to utilize, retention policies, data characteristics, compression or encryption requirements, preferred system components to utilize in a storage operation, and other criteria relating to a storage operation. Thus, a storage policy may indicate that certain data is to be stored in a specific storage device, retained for a specified period of time before being aged to another tier of secondary storage, copied to secondary storage using a specified number of streams, etc. A storage policy and/or a schedule policy may be stored in a storage manager database or in other locations or components of the system.

Integrated Data Rights Security Control

[0055] Some organizations may have multiple levels of security according to which some users can access certain files while others cannot. For example, a high security user group can be defined and this group can be granted access to all documents created by the organization; a medium security group can be granted access to only certain classes of documents; a low security group can be granted access only to certain predefined documents.

[0056] The search interface described herein can be configured to be accessible by any type of user including a secondary copy administrator, an end user who does not have any administrative privileges, or a user of any security clearance. Additionally, the data files stored in the data management system can be tagged with security information. This information tag can be stored in a metabase or any other form of content index and can be used to leverage existing security schema. In embodiments in which a search is performed on one or more context indices, corresponding security tag information can be stored therein. Security information can include identification of various classes of users who are granted rights to access the document as well as identification of classes of users who are denied access rights.

[0057] In some embodiments, security information can be stored in the form of user tags. User tags are further described in the assignee's co-pending application serial no. 11/694,784 entitled "System and Method Regarding Security And Permissions" (Attorney Docket No. 60692.8042), the contents of which are herein incorporated by reference.

[0058] In some further embodiments, the search results can be filtered based on the user's security clearance or access privileges. After a user enters search parameters, data files matching those parameters may be identified, and a list of the identified files displayed to the user. If the user does not have the required security clearance or access privileges, the interface can be configured not to display the file.

[0059] It is possible that a secondary copy administrator may not have sufficient security clearance to inspect a file that is being restored or retrieved. In such a circumstance, the administrator will not be allowed to preview the file or otherwise inspect the contents of it during the search process. The interface providing results may be configured to not display a preview of such a file. If a secondary copy administrator had sufficient security clearance, then a preview may be provided or the administrator may be allowed to make a local copy of the file.

[0060] If the secondary copy administrator does not have sufficient security clearance for a specific file or group or class of files, an interface may be provided through which the administrator may initiate a copy of that file directly from the secondary copy device to a directory or disk associated with a user who has sufficient security clearance. In some instances, the user associated with the file may be the owner of the file. If the secondary copy administrator or other user executing a search query has sufficient security clearance to inspect the contents of the one or more files identified in the search, a preview of the data file may be displayed.

System Embodiments

[0061] The following discussion provides a brief, general description of a suitable computing environment in which the invention can be implemented. Although not required, aspects of the invention 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 invention can be practiced with other communications, data processing, or computer system configurations, including: Internet appliances, hand-held devices (including personal digital assistants (PDAs)), wearable computers, all manner of cellular or mobile phones, multi-processor systems, microprocessor-based or programmable consumer electronics, set-top boxes, 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.

[0062] Aspects of the invention 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 invention 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), or the Internet. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0063] Aspects of the invention may be stored or distributed on computer-readable media, including 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. Indeed, computer implemented instructions, data structures, screen displays, and other data under aspects of the invention 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 invention reside on a server computer, while corresponding portions reside on a client computer such as a mobile or portable device, and thus, while certain hardware platforms are described herein, aspects of the invention are equally applicable to nodes on a network.

Conclusion

[0064] From the foregoing, it will be appreciated that specific embodiments of the 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 invention. For example, although files have been described, other types of content such as user settings, application data, emails, and other data objects can all be indexed by the system. Accordingly, the invention is not limited except as by the appended claims.

[0065] 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.

[0066] The above detailed description of embodiments of the invention is not intended to be exhaustive or to limit the invention to the precise form disclosed above. While specific embodiments of, and examples for, the invention are described above for illustrative purposes, various equivalent modifications are possible within the scope of the invention, 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.

[0067] The teachings of the invention 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.

[0068] These and other changes can be made to the invention in light of the above Detailed Description. While the above description details certain embodiments of the invention and describes the best mode contemplated, no matter how detailed the above appears in text, the invention can be practiced in many ways.

Details of the system may vary considerably in implementation details, while still being encompassed by the invention disclosed herein. As noted above, particular terminology used when describing certain features or aspects of the invention 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 invention with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the invention to the specific embodiments disclosed in the specification, unless the above Detailed Description section explicitly defines such terms. Accordingly, the actual scope of the invention encompasses not only the disclosed embodiments, but also all equivalent ways of practicing or implementing the invention under the claims.

[0069] While certain aspects of the invention are presented below in certain claim forms, the inventors contemplate the various aspects of the invention in any number of claim forms. For example, while only one aspect of the invention is recited as embodied in a computer-readable medium, other aspects may likewise be embodied in a computer-readable medium. Accordingly, the inventors reserve the right to add additional claims after filing the application to pursue such additional claim forms for other aspects of the invention.

CLAIMS

We claim:

1. A system for managing data associated with one or more data storage devices, wherein the data storage devices are coupled to a network, the system comprising:

a data storage management component for managing production copies of data over the network stored within the data storage devices and managing non-production copies of data stored within the data storage devices, wherein the non-production copies are copies made from at least one production copy;

a content indexing component for creating at least one index of data stored within the production and non-production copies managed by the data storage management component; and

a search component accessible through a web browser that searches the data stored within the production and non-production copies in response to a single query to create search results, wherein the production and non-production copies contain data in multiple software application formats.

2. The system of claim 1, further comprising a database associated with the data storage management component, the database storing metadata referring to the stored data managed by the data storage management component

3. The system of claim 1, wherein the search component is configured to search for a user-specified parameter in the at least one index.

4. The system of claim 1, wherein the multiple software application formats include at least two of a word processing document, a spreadsheet, an audiovisual media file, an email message, and a database.

5. The system of claim 1, wherein the data storage management component performs storage operations based on a set of preferences specified by a user of the system.

6. The system of claim 1, further comprising a data security component, the data security component further configured to determine a user's level of access to the data included in the search, and in response to determining that the user does not have access to a particular search result, preventing the search result from being displayed to the user.

7. The system of claim 1, further comprising a data security component, the data security component further configured to determine a user's level of access to the data included in the search, and in response to determining that the user does not have access to a particular search result, displaying the search result with an indication that the user does not have access to retrieve data associated with the result.

8. The system of claim 1, wherein the web-based search component provides an interface for receiving a search parameter from a user, and wherein the search parameter specifies at least one client computer or volume for searching on the client computer.

9. The system of claim 1, wherein the web-based search component provides an interface for receiving a search parameter from a user, and wherein the search parameter specifies at least one client computer or volume for searching across multiple copies of data associated with the client computer or volume.

10. The system of claim 1, wherein the web-based search component is configured to confirm a specified security clearance level from a user.

11. The system of claim 1, wherein the web-based search component is configured to display a preview of search results.

12. The system of claim 1, wherein the non-production copies comprise copies in a snapshot format.

13. A method for managing data associated with a data storage component, wherein the data storage component is coupled to multiple computers over a network, the method comprising:

managing data stored within the data storage component and managing secondary copies of data stored within the data storage component, and wherein the secondary copies are copies made from at least one primary copy;

creating at least one index of data stored within the data storage component and managed by the data storage management component; and

searching stored data associated with multiple software applications, and associated with the multiple secondary copies, based on a single query.

14. The method of claim 13, further comprising creating a metabase associated with the data storage management component, the metabase storing metadata referring to the stored data managed by the data storage management component.

15. The method of claim 13, wherein the search component is configured to search for a user-specified parameter in the at least one index.

16. The method of claim 13, where the multiple applications include at least two of a word processing application, email application, and database application.

17. The method of claim 13, further comprising performing storage operations on the data stored within the data storage component based on a set of preferences or other criteria.

18. The method of claim 13, further comprising permitting access only to files satisfying a predefined security level.

19. The method of claim 13, wherein the secondary copies comprise backup copies and archive copies.

20. A computer-readable storage medium having computer-executable instructions that, when executed by a computer perform a method of managing data associated with a data storage component, wherein the data storage component is coupled to multiple computers over a network, the method comprising:

managing data stored within the data storage component and managing secondary copies of data stored within the data storage component, and wherein the secondary copies are copies made from at least one primary copy;

creating at least one index of the stored data managed by the data storage management component; and

searching stored data associated with multiple applications, and associated with the multiple secondary copies, based on a single query.

21. A system for managing data associated with at least one data storage component, wherein the data storage component is coupled to multiple computers over a network, the system comprising:

a data storage management component for managing data stored within the data storage component, wherein data stored within the data storage component includes multiple secondary copies of primary data;

a metabase associated with the data storage management component that stores descriptive data referring to the stored data managed by the data storage management component; and

a search component for searching stored data associated with multiple applications and associated with the multiple secondary copies based on a single query,

wherein the stored data includes data types relating to at least two or more different applications.

22. The system of claim 21, further comprising a HTML-based graphical user interface for receiving user input identifying a type of data to be used by the search component in a search.

23. The system of claim 22, wherein the user interface is further configured to display a preview of search results that provides a summary of the search results.

24. The system of claim 21, wherein the metabase comprises multiple content indices, and wherein each of the content indices is associated with a logical volume that comprises one or more physical volumes.

25. The system of claim 21, further comprising a security component for storing multiple security policies relating to multiple users and further defining data access rights for those users including which search results can be displayed to the user in response to a search.

26. The system of claim 21, wherein the search component is configured to receive search criteria input via a web browser.

27. The system of claim 26, wherein the storage management component is further configured to receive from a user an indication that the search results satisfy the criteria for which the user is searching and to apply the search criteria as a data management policy for future searches by the user to receive similar results.

28. The system of claim 21, wherein the data types relate to different types of applications that generate different formats of data.

29. The system of claim 21, wherein the secondary copies comprise backup copies and archive copies.

30. A method for managing data associated with one or more data storage devices, wherein the data storage devices are coupled to multiple computers over a network, the method comprising:

receiving data through a storage operation specified by a data storage policy for managing data stored within the data storage device, wherein the storage operation creates at least one secondary copy of the data;

maintaining an index of the received data and the secondary copy of the data; and

searching stored data associated with multiple different secondary copies, wherein the search is based on a single query provided by a user, and wherein the stored data is in multiple application formats.

31. The method of claim 30, further comprising receiving user input through a graphical user interface identifying a type of data to be used for searching.

32. The method of claim 30, wherein the index comprises multiple content indices, and wherein each of the content indices is associated with a logical volume that represents one or more physical volumes.

33. The method of claim 30, further comprising storing multiple security policies relating to multiple users, wherein the policies further define data access rights for those users that specify the data that can be accessed by each user.

34. The method of claim 30, further comprising receiving search criteria from a user and an indication that the search criteria produce the type of results that a user wants to find on a regular basis, and storing the search criteria as a data management policy that the user can apply to receive new search results in the future that match the search criteria.

35. The method of claim 30, wherein the secondary copies comprise copies of secondary copies on which additional data storage operations have been performed.

36. The method of claim 30, wherein two or more of the secondary copies differ in type of document and time of creation.

37. A computer-readable storage medium having computer-executable instructions that, when executed by a computer perform a method for managing data associated with a data storage component, wherein the data storage component is coupled to multiple computers over a network, the method comprising:

receiving data from a first computer over the network in a first software application format;

receiving data from a second computer over the network in a second software application format;

receiving a data storage policy that specifies one or more storage operations to be performed on the data received from the first and second computer;

maintaining an index comprising the data received from the first and second computer, wherein the index parses the first and second application formats to identify indexable content within the received data;

receiving a query from a user to identify content stored across multiple computers and in multiple software application formats that matches a search criteria; and

in response to the query, searching the indexed data received from the first and second computer to identify data that matches the search criteria.

38. The method of claim 37, wherein the index comprises multiple content indices, and wherein each of the content indices is associated with a different type of content.

39. The method of claim 37, further comprising storing multiple security policies relating to multiple users, wherein the policies define the data for which a user can receive search results.

40. The method of claim 37, further comprising receiving an indication from the user that the searching identifies a type of search results that the user would like to receive on a regular basis, and storing the search criteria as a data management policy that the user can apply in future queries to receive similar search results based on new data received by the system.

41. A method for managing data associated with a data storage device, wherein the data storage device is coupled to multiple computers over a network, the method comprising:

storing application data received from multiple computers over the network in the data storage device, the application data comprising data in multiple software application formats;

receiving a data storage policy that specifies one or more data storage operations to be performed on the stored application data ;

storing in an index summary data describing the stored application data so that the application data can be identified based on one or more received search parameters;

receiving a search request from a user comprising at least one search parameter;

searching the index for application data based on the at least one received search parameter, wherein the search accesses the summary data and identifies matching application data from multiple computers and in multiple application formats.

42. The method of claim 41, further comprising performing a storage operation based on the storage policy, wherein the storage policy defines performing a backup of the search results.

43. The method of claim 41, further comprising storing the at least one received search parameter as a new data storage policy, wherein the new data

storage policy comprises a set of preferences or storage criteria associated with performing a storage operation to retrieve similar search results in a subsequent query.

44. The method of claim 41, wherein the storage policy is associated a scheduling component and defines when the scheduling component should perform a storage operation and how often to perform the storage operation.

45. The method of claim 41, wherein the data storage policy is stored in a storage manager, the storage manager further comprising a jobs agent configured to carry out data storage operations, a management agent configured to monitor the status of the data storage operations, and an interface module configured to display the status of the data storage operations to a user.

46. The method of claim 41, further comprising receiving the search request from the user across the network via a web browser, wherein the search identifies data stored across multiple data storage tiers in the network satisfying the search request.

47. The method of claim 41, further comprising identifying one or more documents satisfying the at least one search parameter and retrieving one or more identified documents for access by a user without overwriting a current copy of the one or more identified documents on the data storage component.

48. A computer-readable storage medium having computer-executable instructions that, when executed by a computer perform a method for managing data associated with a data storage component, wherein the data storage component is coupled to multiple computers over a network, the method comprising:
storing application data received from multiple computers over the network in the data storage device, the application data comprising data in multiple software application formats;
receiving a data storage policy that specifies one or more data storage operations to be performed on the stored application data ;

storing in an index summary data describing the stored application data so that the application data can be identified based on one or more received search parameters;

receiving a search request from a user comprising at least one search parameter;

searching the index for application data based on the at least one received search parameter, wherein the search accesses the summary data and identifies matching application data from multiple computers and in multiple application formats.

49. The computer-readable storage medium of claim 48, further comprising performing a storage operation based on the storage policy, wherein the storage policy defines performing a backup of the search results.

50. The computer-readable storage medium of claim 48, further comprising storing the at least one received search parameter as a new data storage policy, wherein the new data storage policy comprises a set of preferences or storage criteria associated with performing a storage operation to retrieve similar search results in a subsequent query.

51. A method for accessing documents available through a network, wherein the documents are stored on one or more data storage devices coupled to the network, the method comprising:

receiving a request to access one or more documents, wherein the request is associated with a selected security clearance level;

searching a metabase to locate metadata describing one or more documents satisfying the request based on a security clearance level stored with the metadata that describes characteristics of the documents;

resolving references in the located metadata to identify one of the one or more data storage devices containing the data, wherein the resolving at least assists in determining which of the one or more data storage devices to access the data from, and wherein the

access is at least in part based on a data access criteria that specifies the results that a user has access to view; and providing a signal to access, via the network and in response to the request, the one or more documents at the one or more identified data storage devices.

52. The method of claim 51, further comprising receiving one or more search parameters relating to documents stored on at least one of the data storage devices.

53. The method of claim 52, further comprising identifying one or more documents satisfying the one or more received search parameters.

54. The method of claim 53, further comprising determining access control information of the one or more documents determined to satisfy the received search parameters and comparing the access control information to a security clearance level of a user associated with the search.

55. The method of claim 54, further comprising displaying a preview of each document determined to satisfy the received search parameters and for which the security level is sufficient for viewing the documents based on the access control information.

56. A computer-readable storage medium having computer-executable instructions that, when executed by a computer perform a method for accessing documents available through a network, wherein the documents are stored on one or more data storage devices coupled to the network, the method comprising:
receiving a request to access one or more documents, wherein the request is associated with a selected security clearance level;
searching a metabase to locate metadata describing one or more documents satisfying the request based on a security clearance level stored with the metadata that describes characteristics of the documents;

resolving references in the located metadata to identify one of the one or more data storage devices containing the data, wherein the resolving at least assists in determining which of the one or more data storage devices to access the data from, and wherein the access is at least in part based on a data access criteria; and providing a signal to access, via the network and in response to the request, the one or more documents at the one or more identified data storage devices.

57. The method of claim 56, further comprising receiving one or more search parameters relating to documents stored on at least one of the data storage devices.

58. The method of claim 57, further comprising identifying one or more documents satisfying the one or more received search parameters.

59. The method of claim 58, further comprising determining a security clearance level of the one or more documents determined to satisfy the received search parameters.

60. The method of claim 59, further comprising displaying a preview of the one or more documents determined to satisfy the received search parameters if the security level is greater than a predetermined level for the one or more documents.




FIG. 1

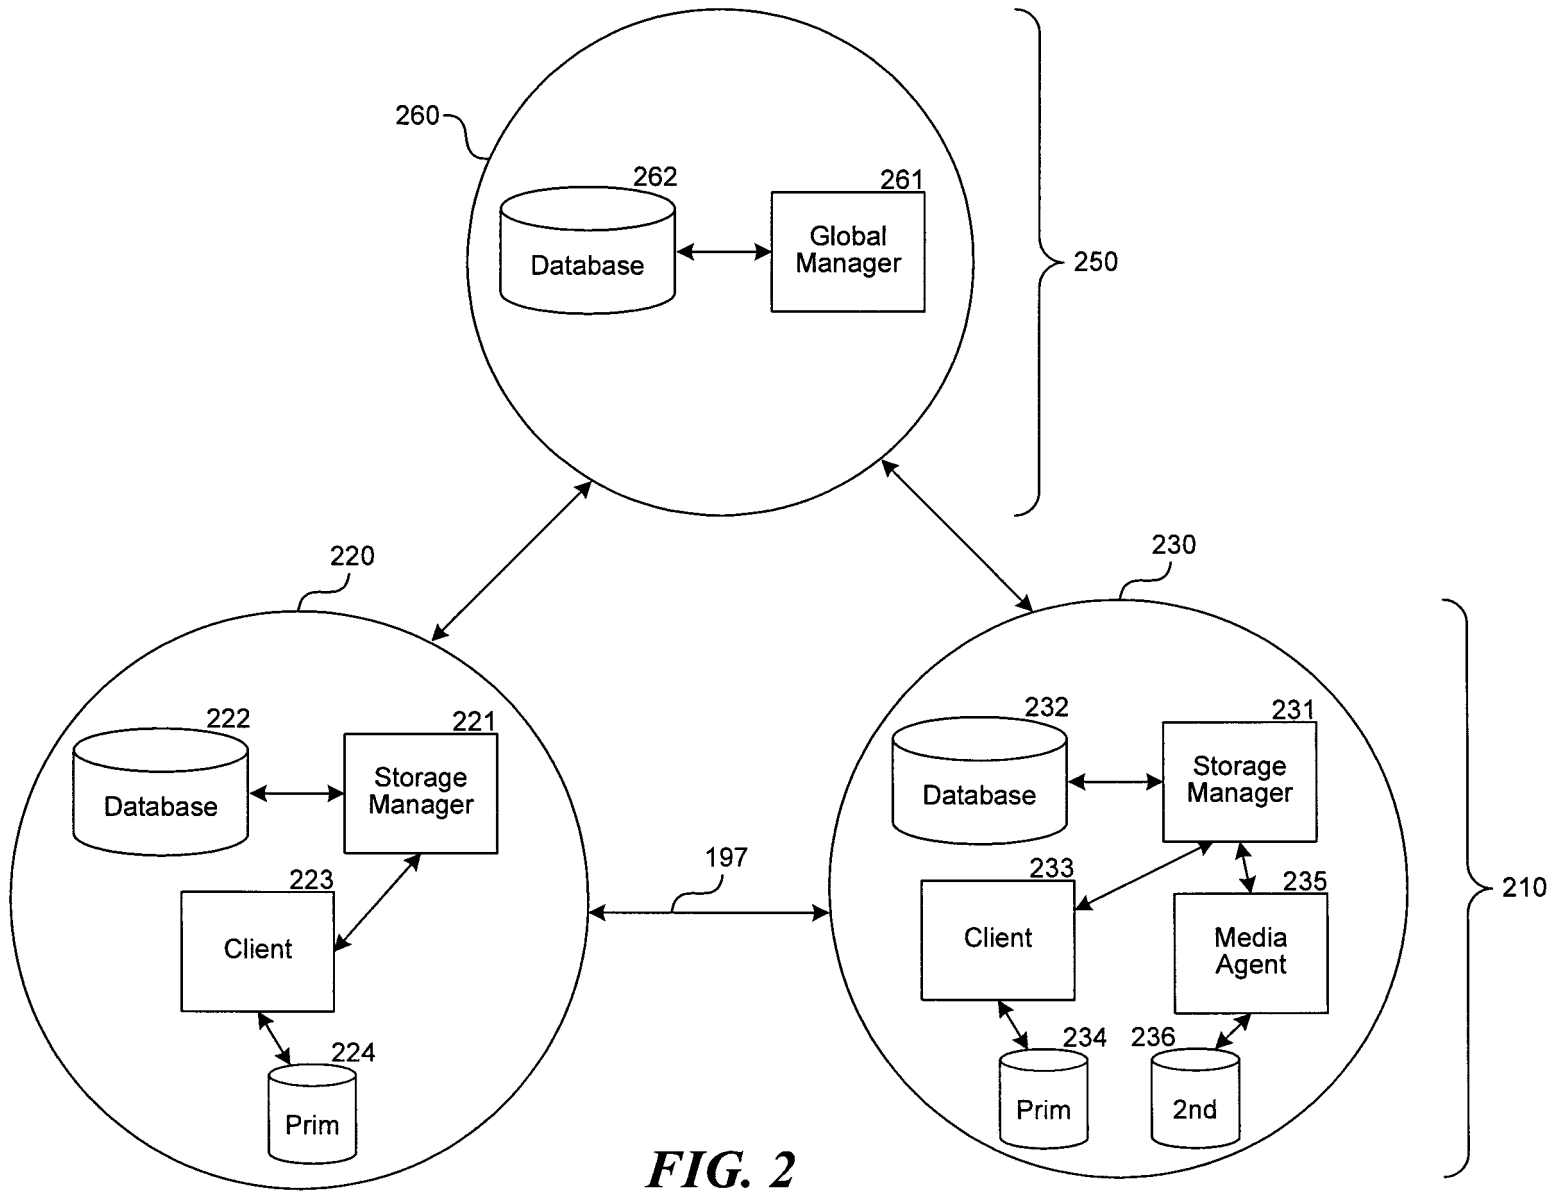


FIG. 2

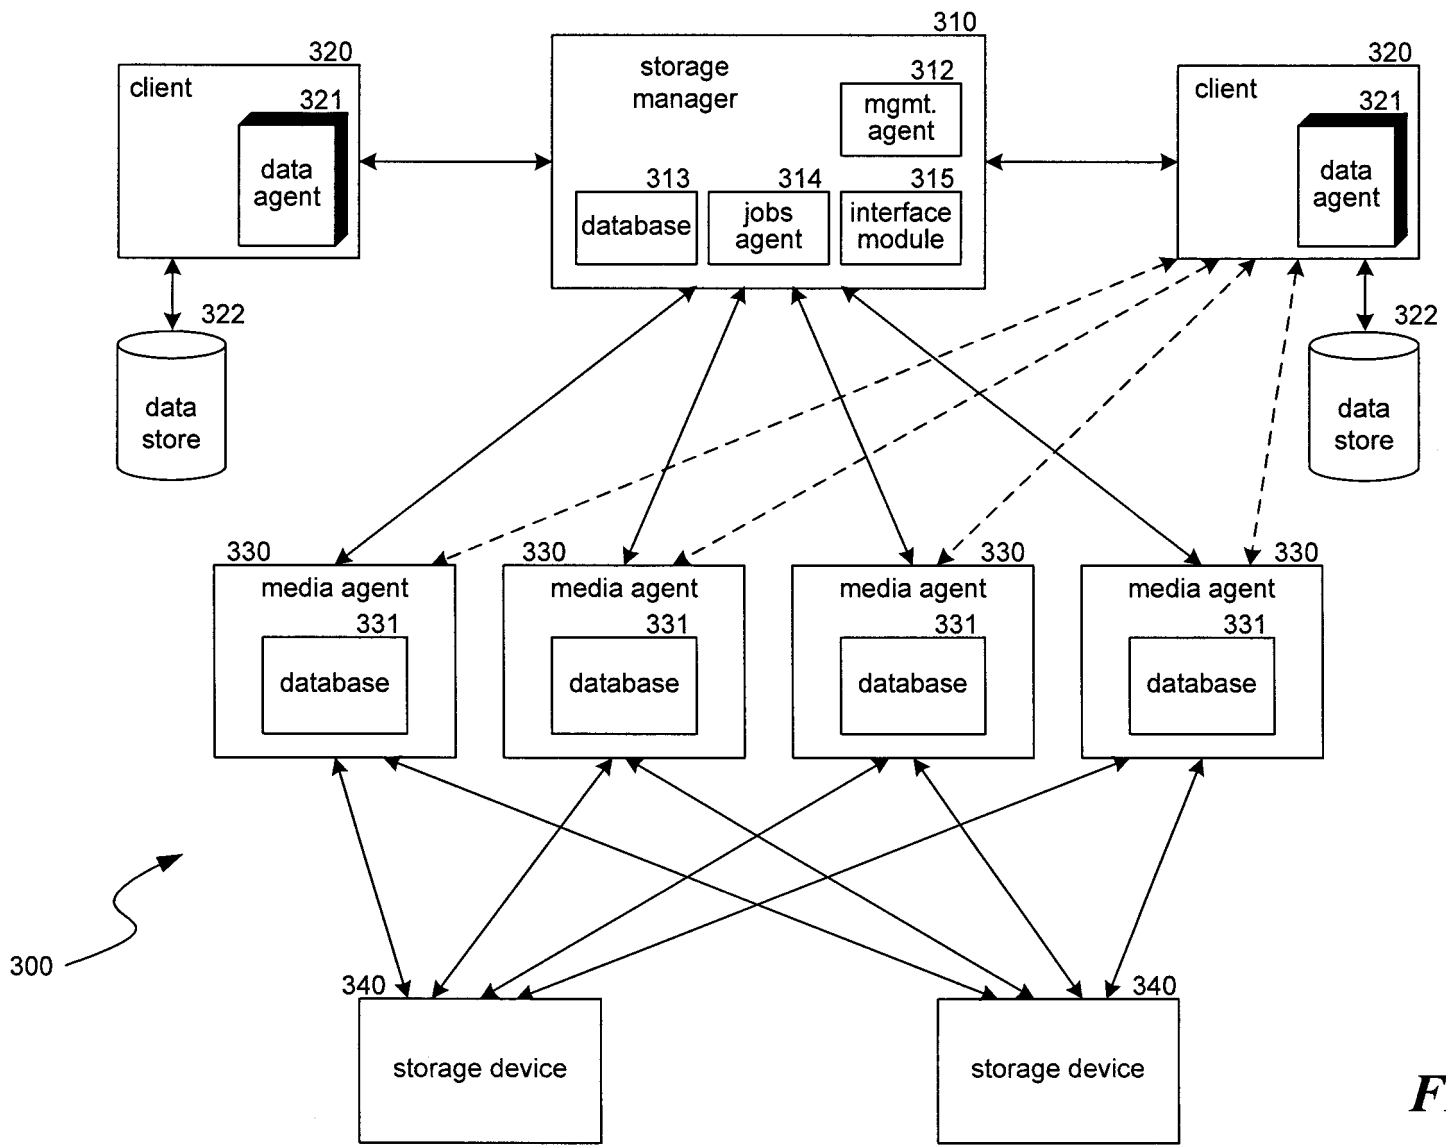


FIG. 3

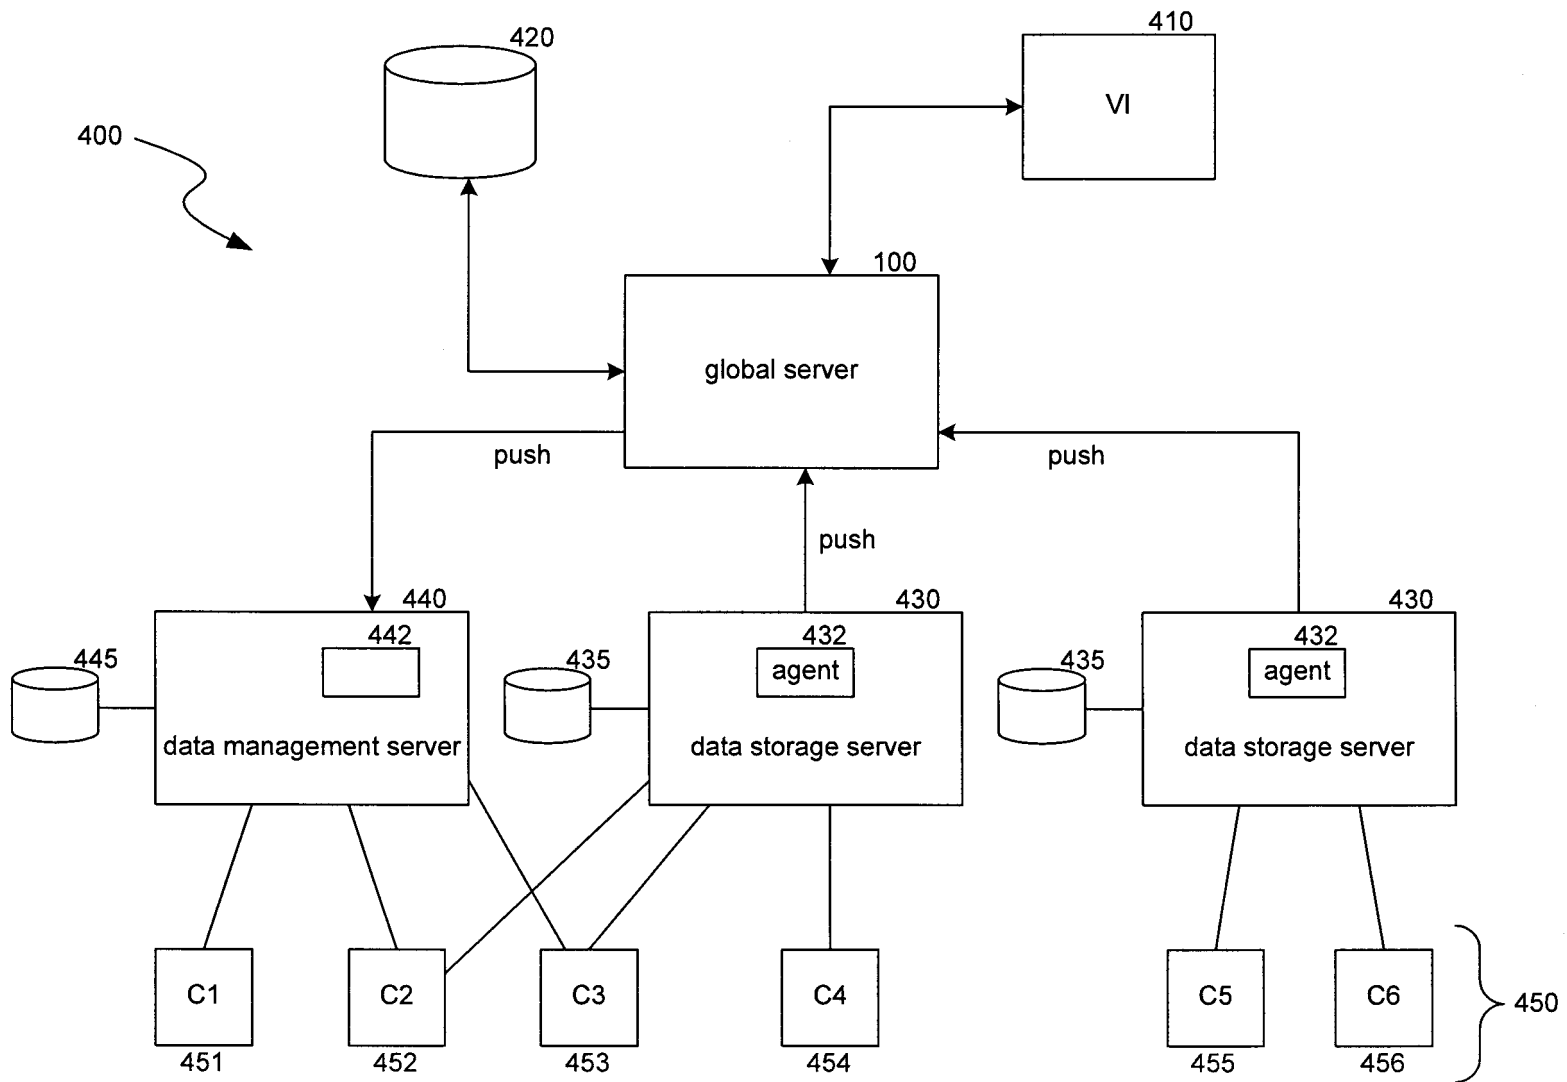


FIG. 4




FIG. 5




FIG. 6

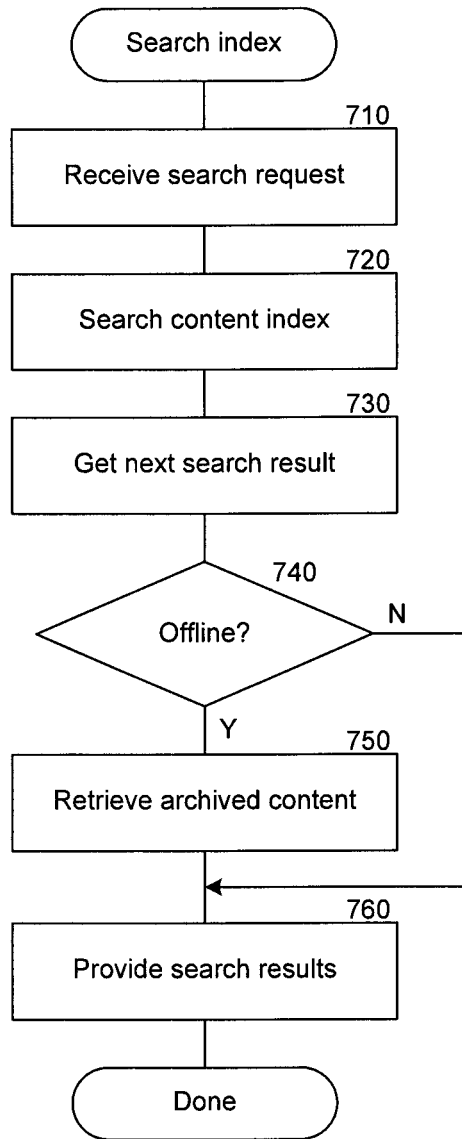


FIG. 7

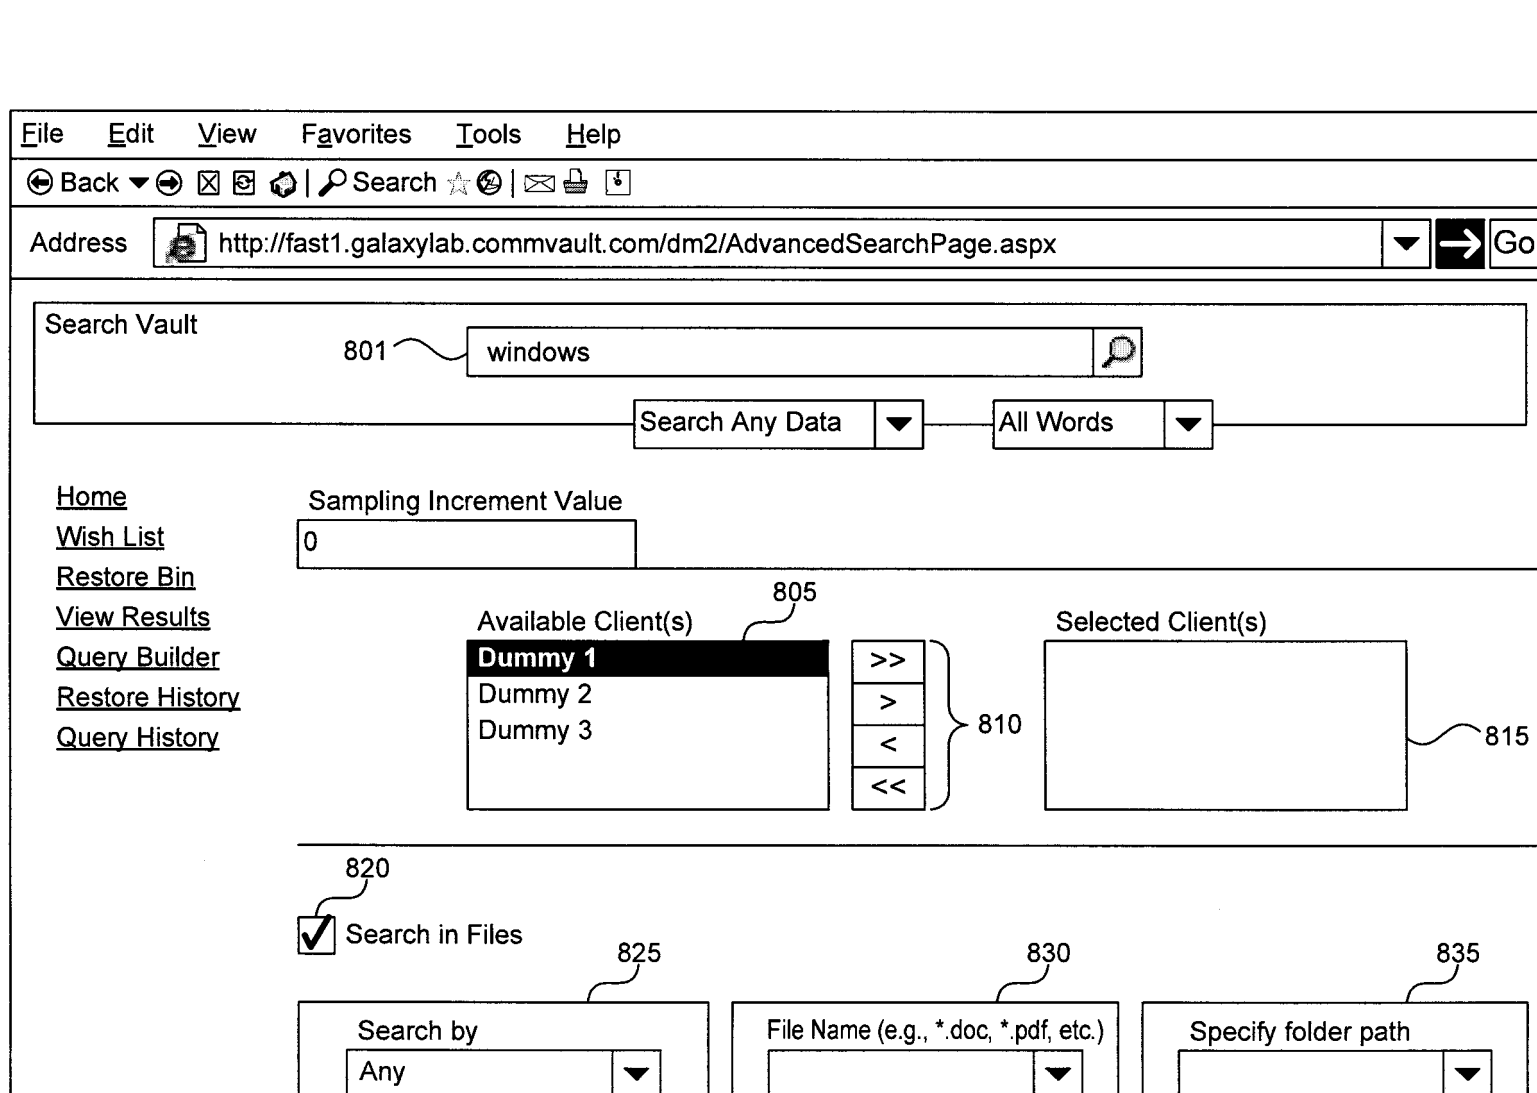


FIG. 8

800

Super Search

Search Options

905 Search For

#Files#

906 #Name/Pattern#

907 #Size# Any

908 #Folder#

909 #Modified Data#

Emails

910 Subject

911 From

912 To

913 Folder

914 Receive Date

#Advanced Options#

#Search# #Clear# Schedule a Search Operation

#Advanced Options#

Clients
Itouch; 915

DAs
Backedup Files; 916

Search for data owned by 917

Search for data accessible by 918

#Sample Every #
 messages # 919

#Indexing Time#
 #Any# 920

#Time Zone#
(GMT-05:00) Eastern Time (US & Canada) 921

900

FIG. 9

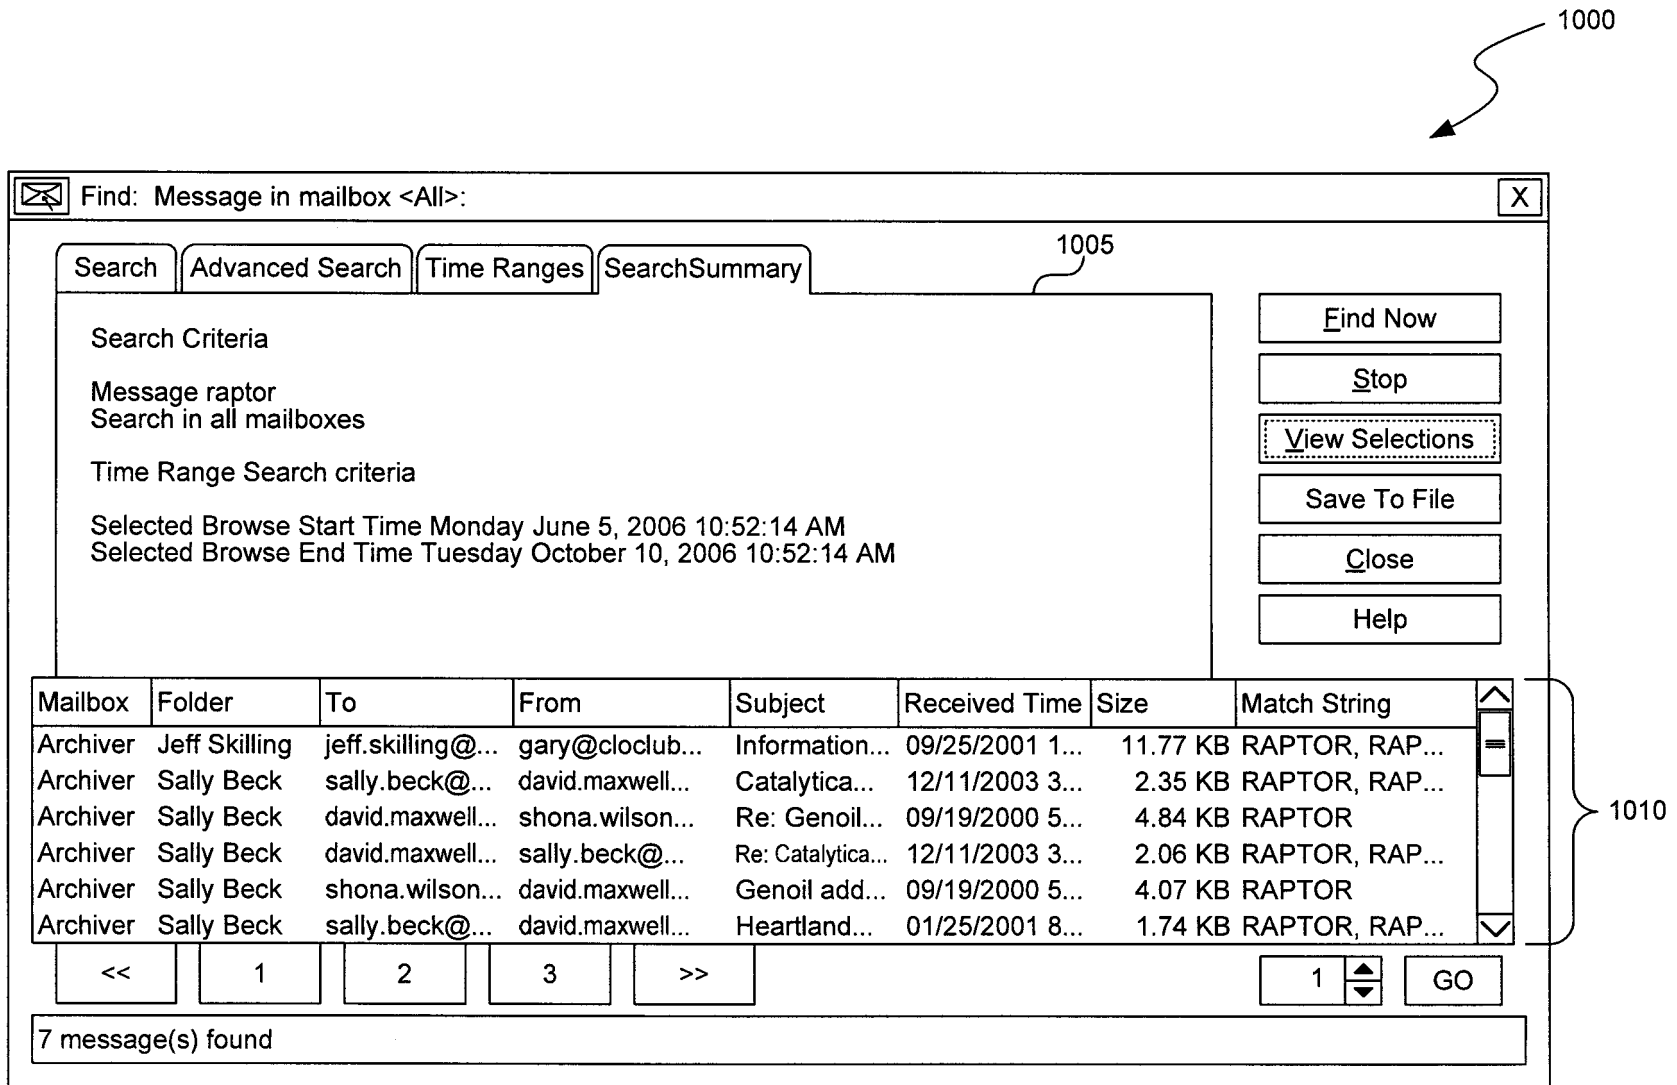


FIG. 10

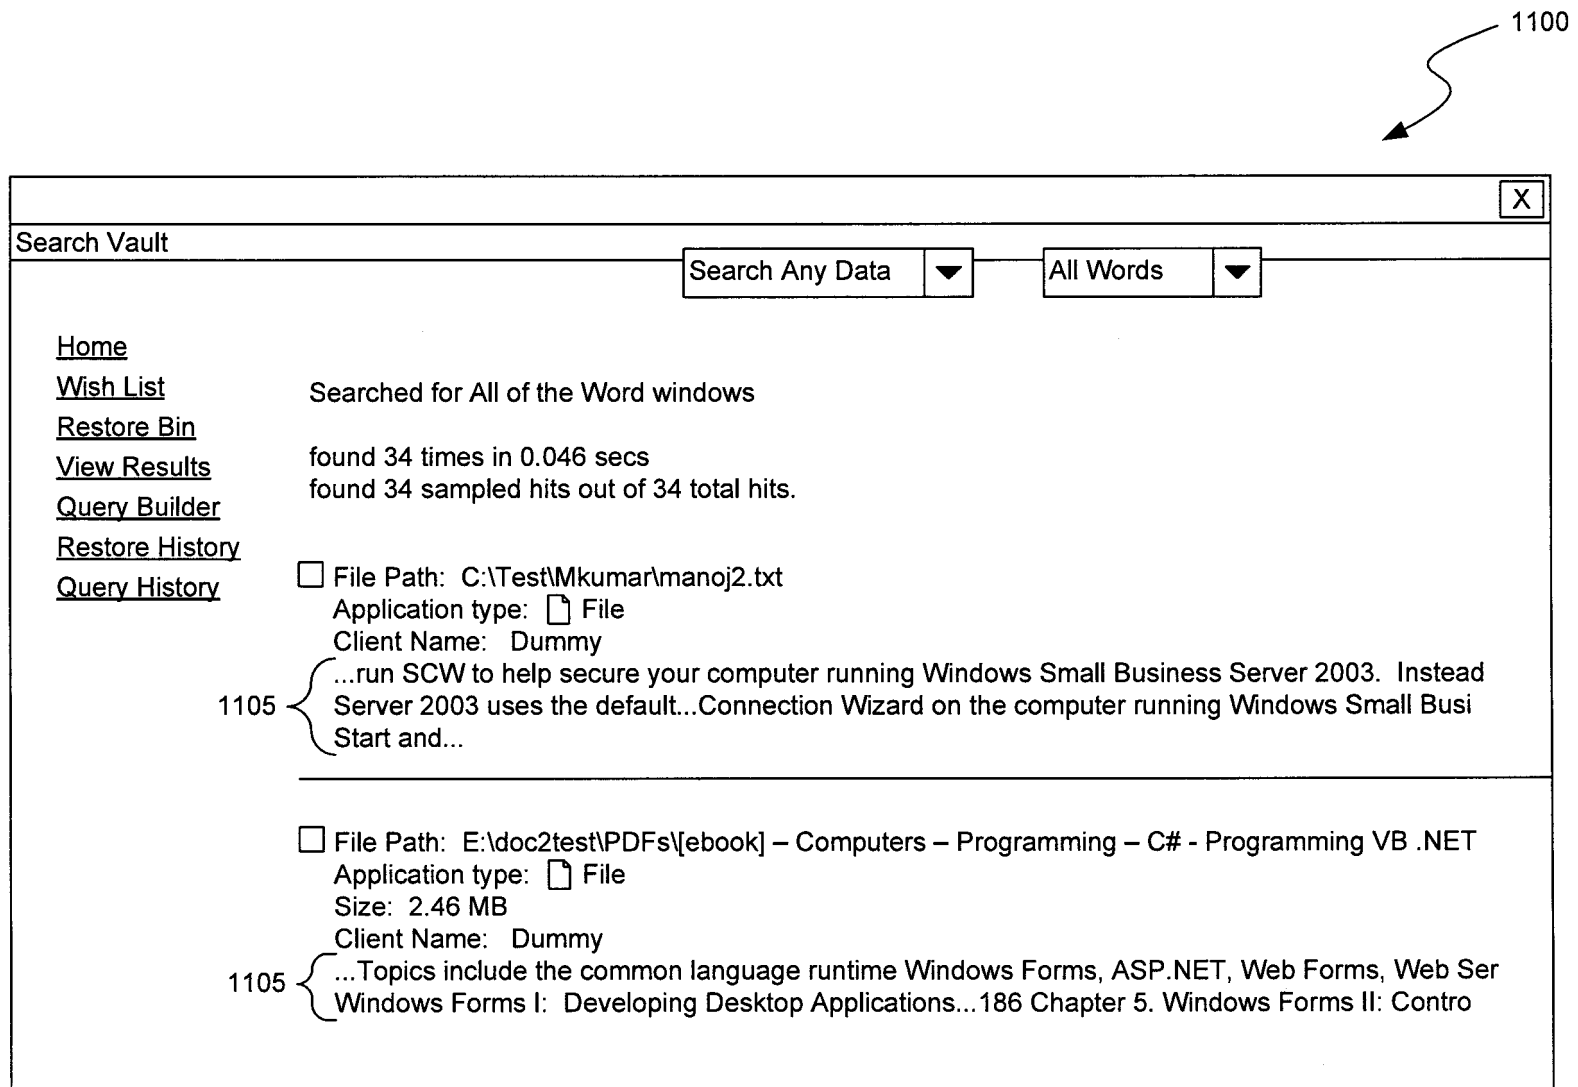


FIG. 11