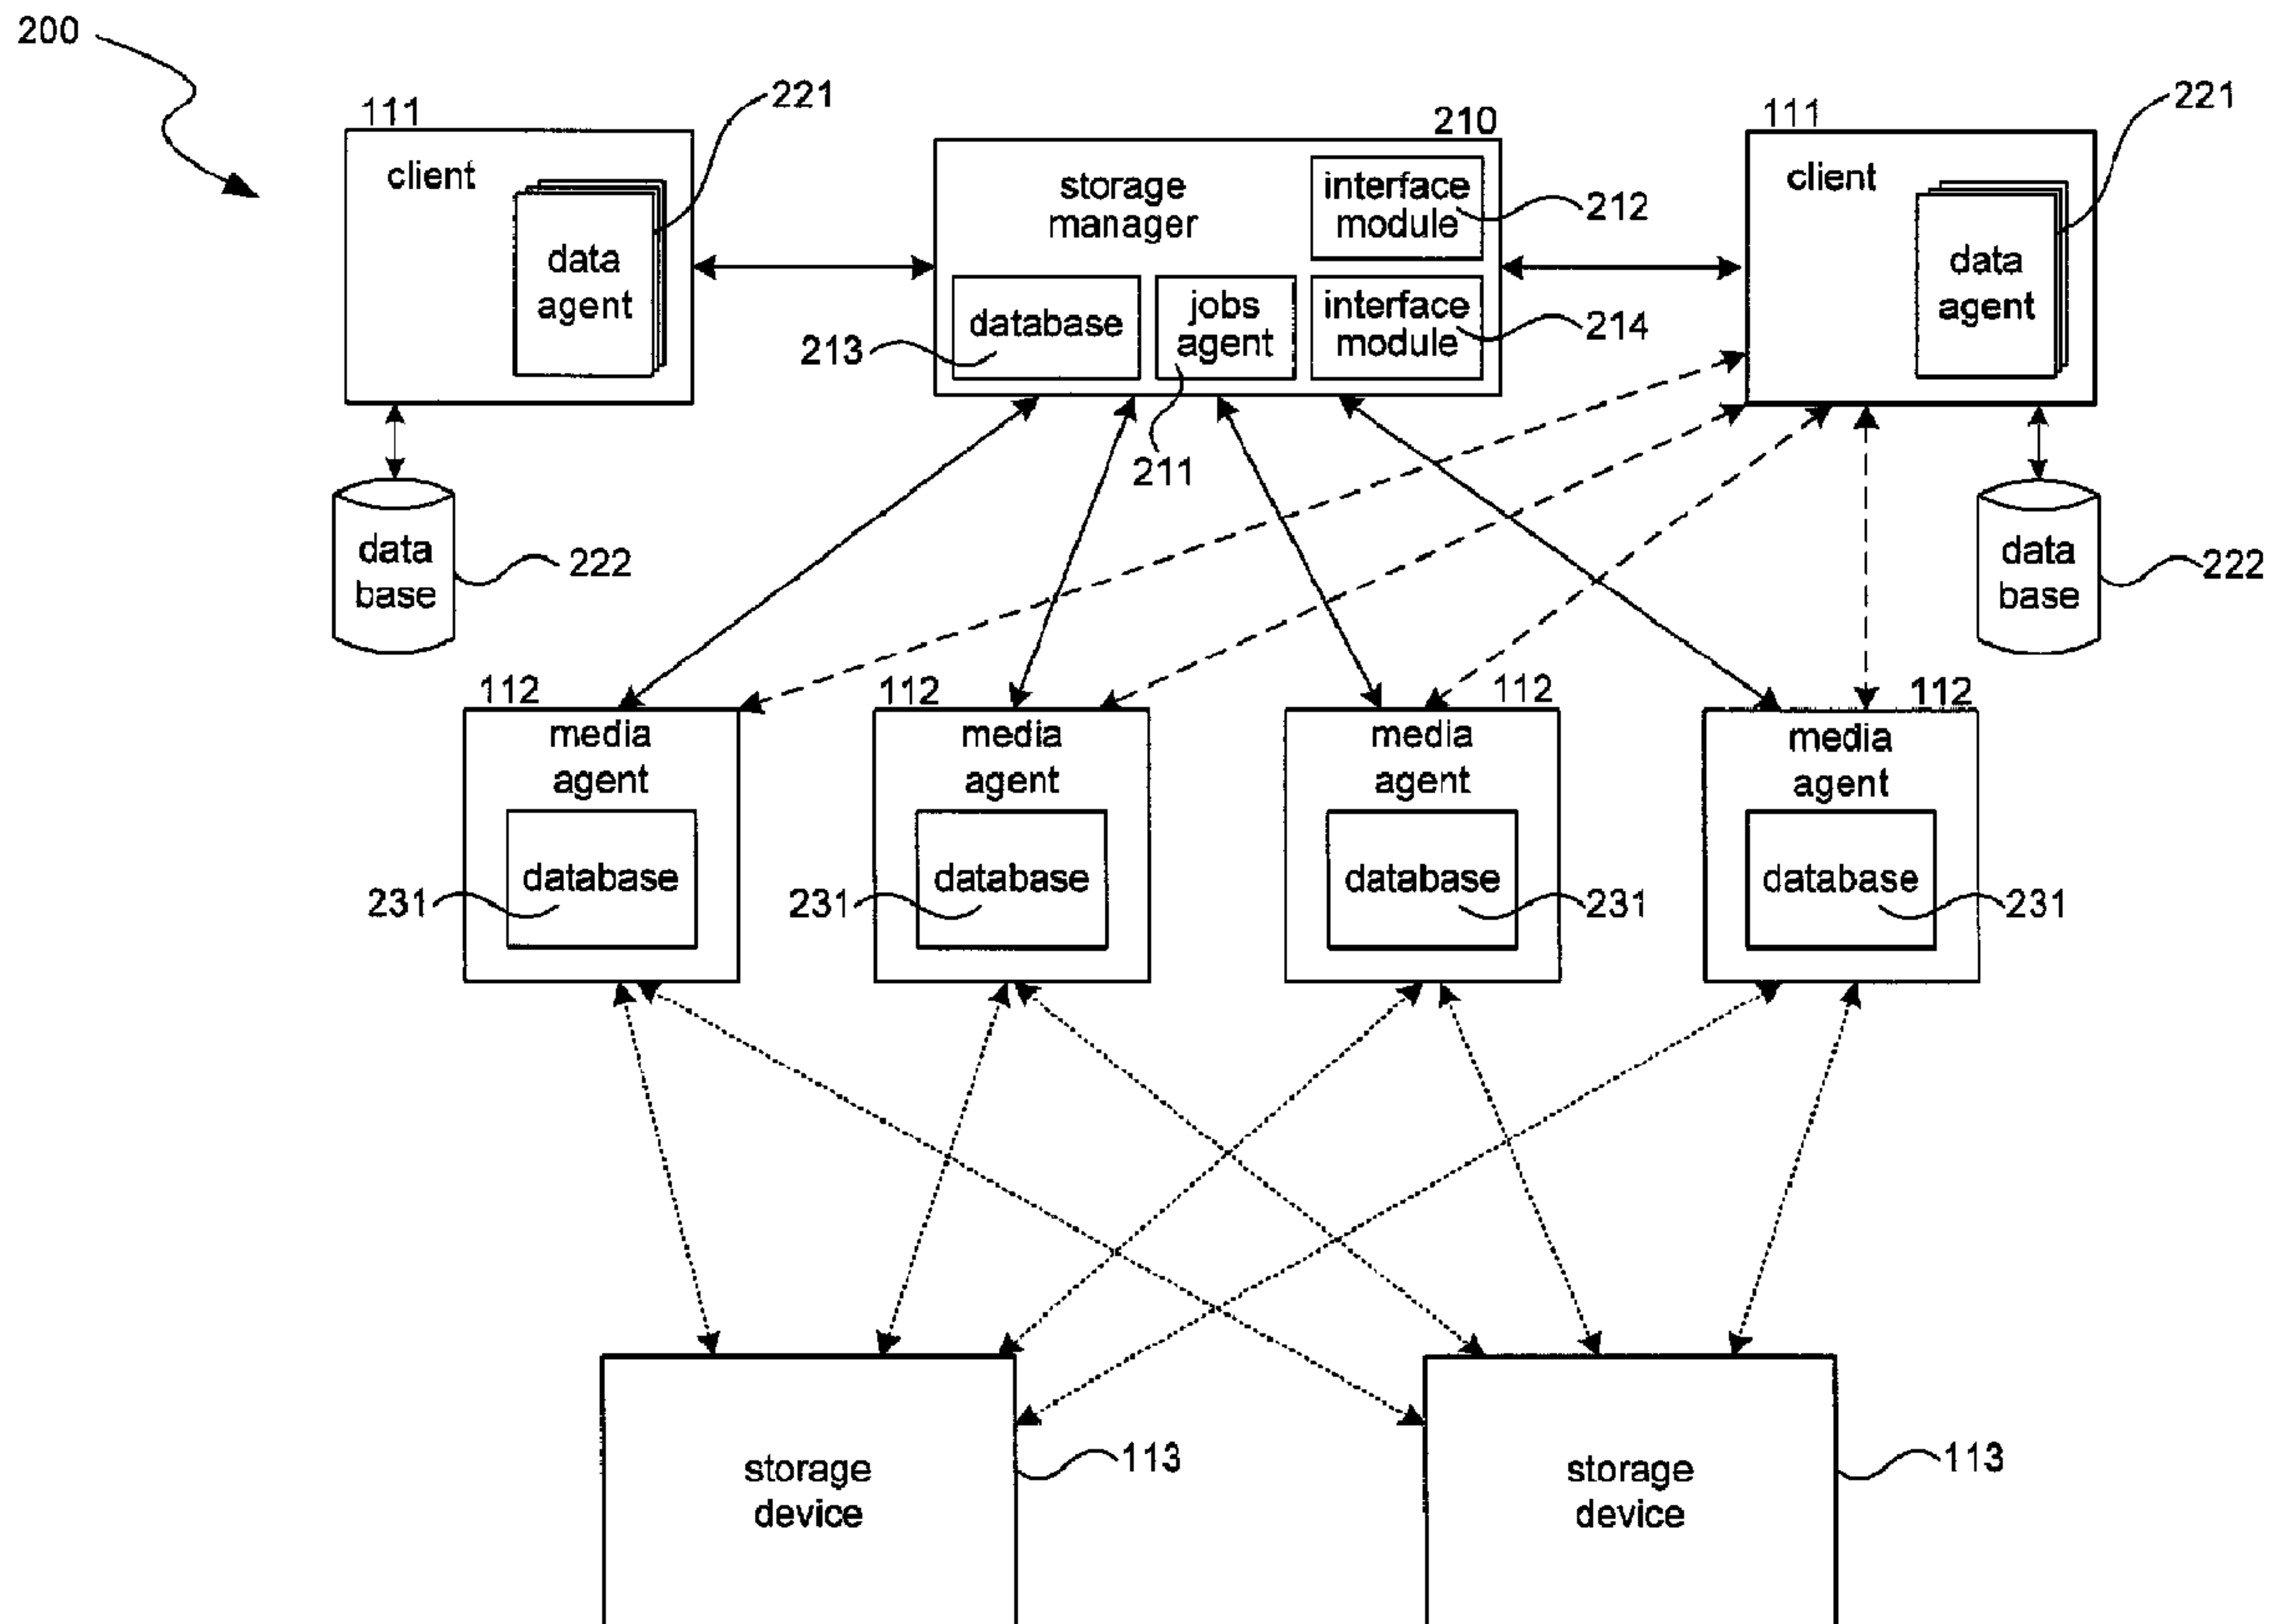




(86) **Date de dépôt PCT/PCT Filing Date:** 2009/01/29  
 (87) **Date publication PCT/PCT Publication Date:** 2009/12/23  
 (45) **Date de délivrance/Issue Date:** 2015/08/04  
 (85) **Entrée phase nationale/National Entry:** 2010/11/12  
 (86) **N° demande PCT/PCT Application No.:** US 2009/032325  
 (87) **N° publication PCT/PCT Publication No.:** 2009/154806  
 (30) **Priorité/Priority:** 2008/06/18 (US12/141,776)

(51) **Cl.Int./Int.Cl. G06F 17/00** (2006.01),  
**G06F 12/16** (2006.01)  
 (72) **Inventeur/Inventor:**  
KOTTOMTHARAYIL, RAJIV, US  
 (73) **Propriétaire/Owner:**  
COMMVAULT SYSTEMS, INC., US  
 (74) **Agent:** OYEN WIGGS GREEN & MUTALA LLP

(54) **Titre : PLANIFICATION DE PROTECTION DE DONNEES, TELLE QUE LA FOURNITURE D'UNE FENETRE DE SAUVEGARDE FLEXIBLE DANS UN SYSTEME DE PROTECTION DE DONNEES**  
 (54) **Title: DATA PROTECTION SCHEDULING, SUCH AS PROVIDING A FLEXIBLE BACKUP WINDOW IN A DATA PROTECTION SYSTEM**



(57) **Abrégé/Abstract:**

A data protection scheduling system provides a flexible or rolling data protection window that analyzes various criteria to determine an optimal or near optimal time for performing data protection or secondary copy operations. While prior systems may have

**(57) Abrégé(suite)/Abstract(continued):**

scheduled backups at an exact time (e.g., 2:00 a.m.), the system described herein dynamically determines when to perform the backups and other data protection storage operations, such as based on network load, CPU load, expected duration of the storage operation, rate of change of user activities, frequency of use of affected computer systems, trends, and so on.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau(43) International Publication Date  
23 December 2009 (23.12.2009)(10) International Publication Number  
**WO 2009/154806 A1**(51) International Patent Classification:  
*G06F 19/00* (2006.01)(21) International Application Number:  
PCT/US2009/032325(22) International Filing Date:  
29 January 2009 (29.01.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
12/141,776 18 June 2008 (18.06.2008) US(71) Applicant (for all designated States except US): **COMM-VAULT SYSTEMS, INC.** [US/US]; 21 Crescent Place, Oceanport, NJ 07757-0090 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **KOT-TOMTHARAYIL, Rajiv** [US/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US).(74) Agents: **ANANIAN, Michael** et al.; Perkins Coie LLP, P.O.. Box 1247, Seattle, WA 98111-1247 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, 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, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, 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, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: DATA PROTECTION SCHEDULING, SUCH AS PROVIDING A FLEXIBLE BACKUP WINDOW IN A DATA PROTECTION SYSTEM

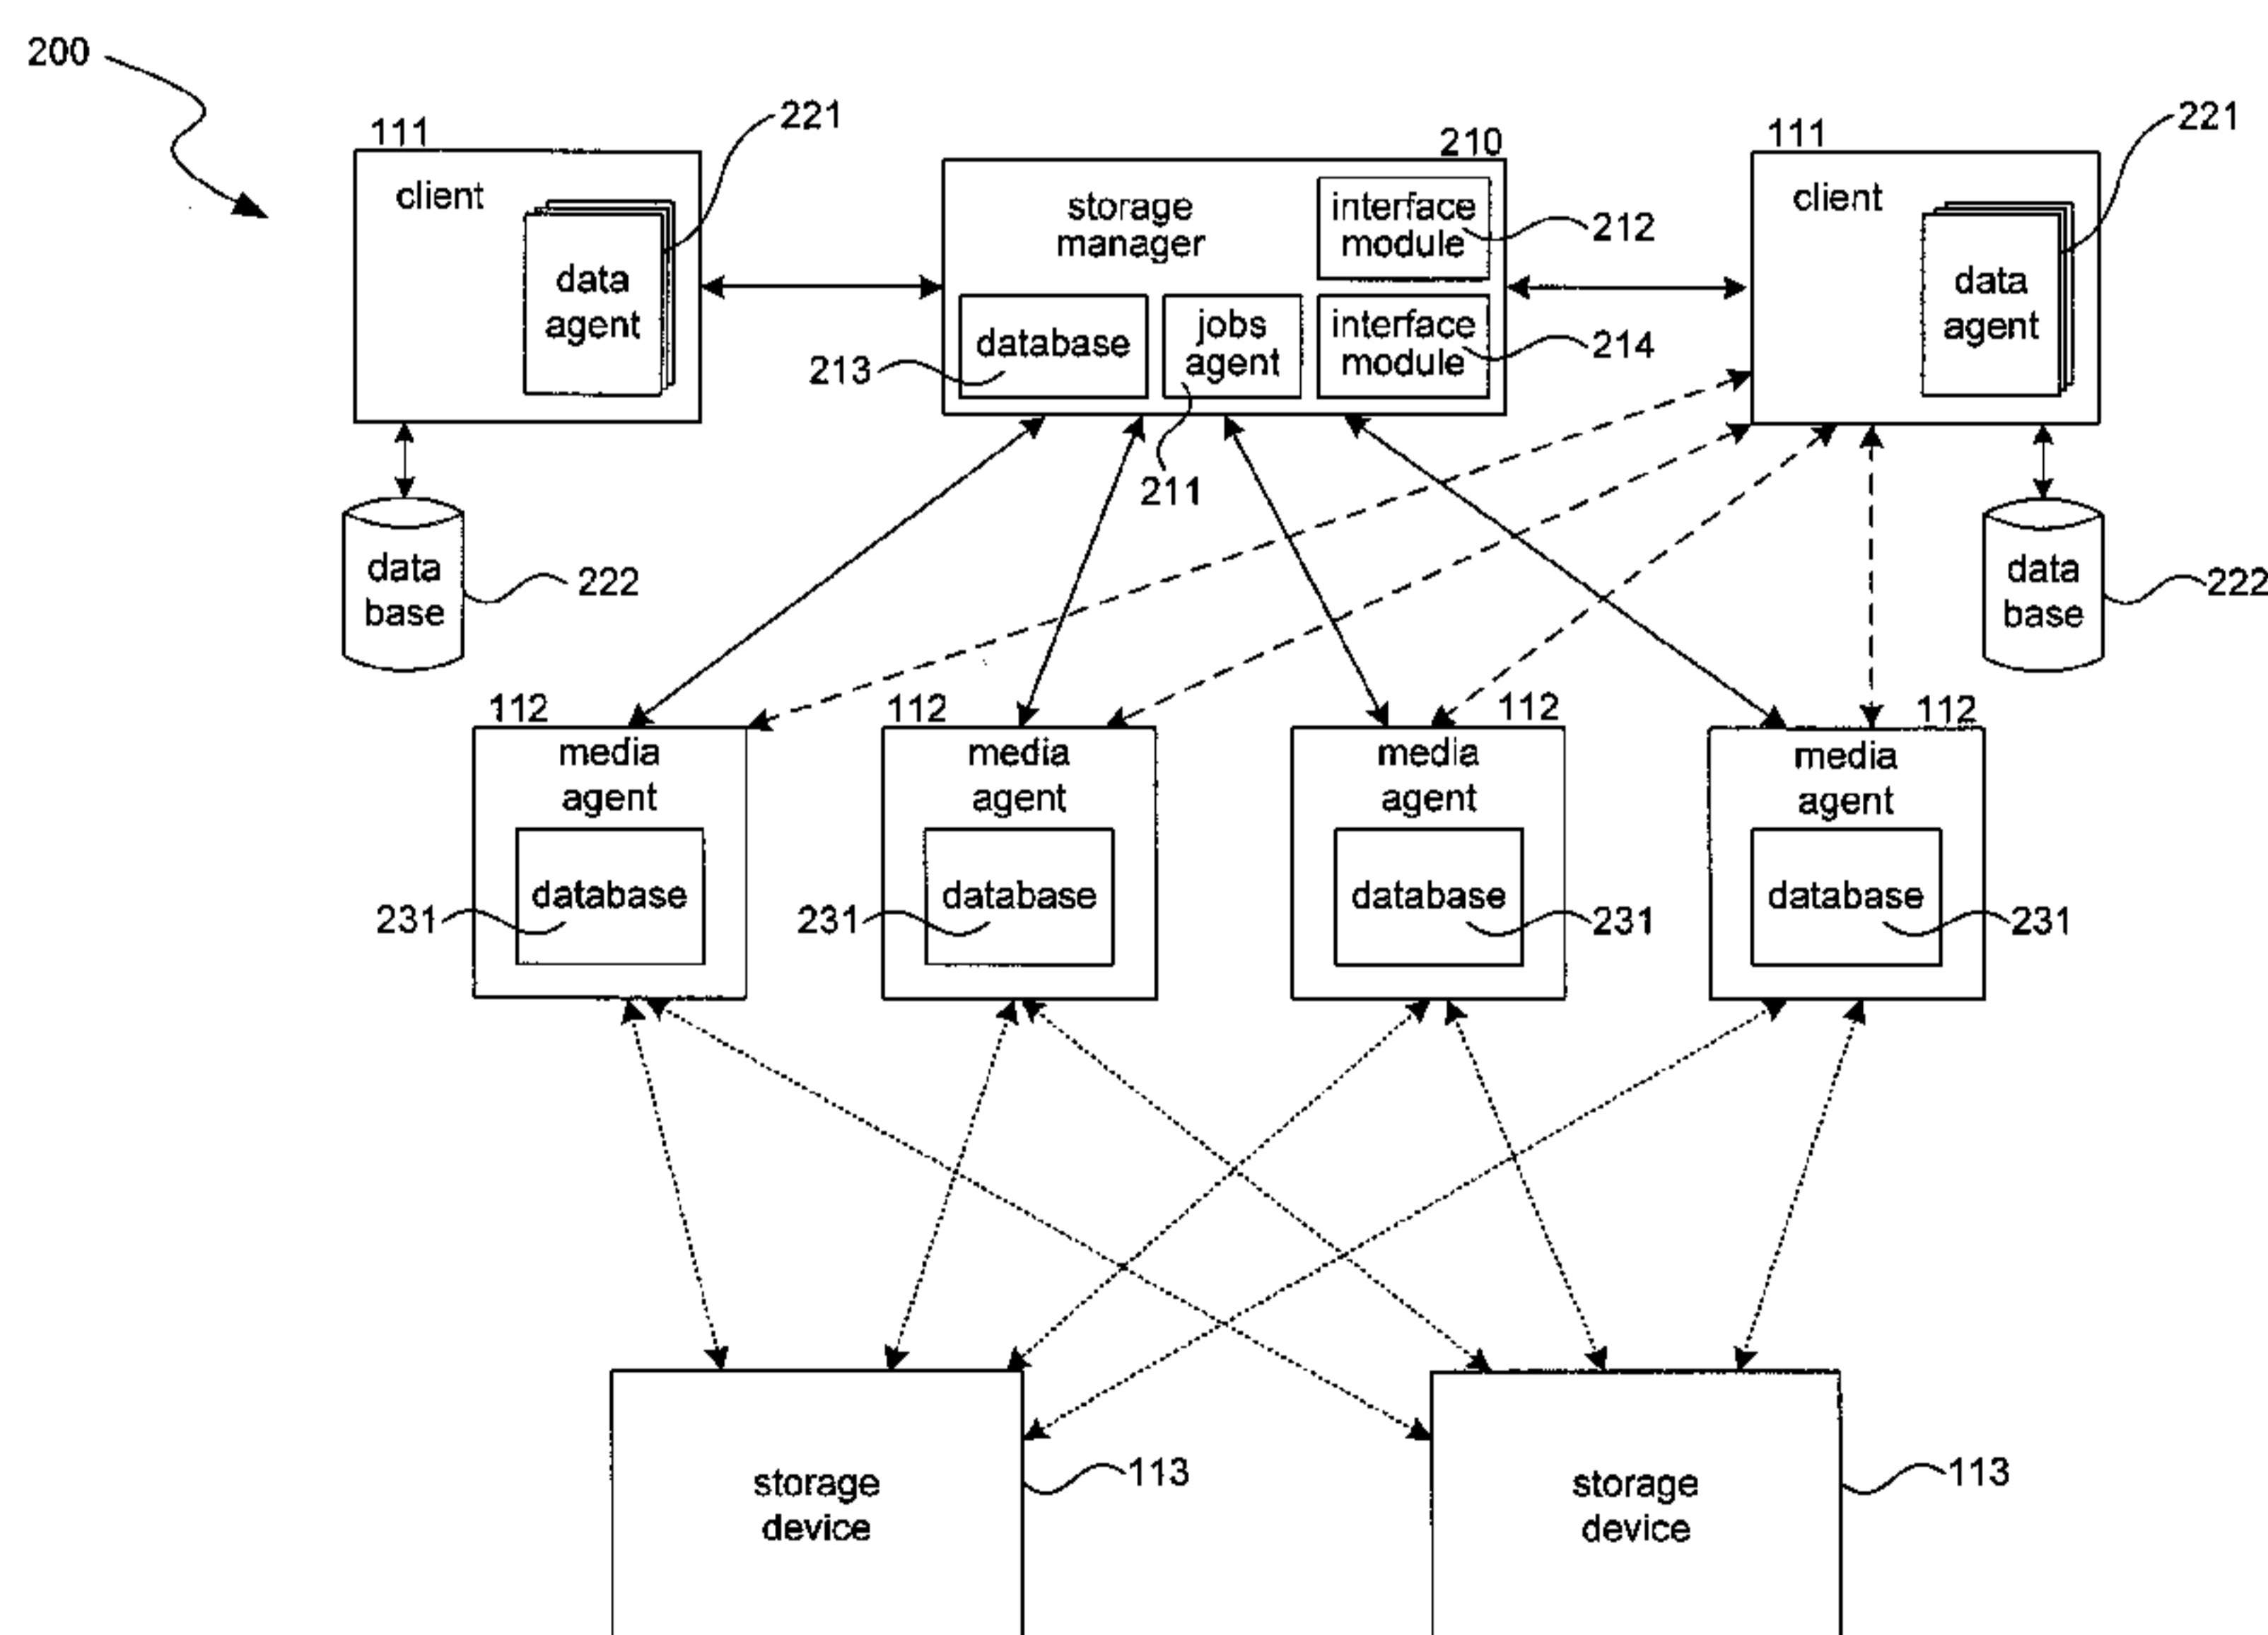


FIG. 2

(57) Abstract: A data protection scheduling system provides a flexible or rolling data protection window that analyzes various criteria to determine an optimal or near optimal time for performing data protection or secondary copy operations. While prior systems may have scheduled backups at an exact time (e.g., 2:00 a.m.), the system described herein dynamically determines when to perform the backups and other data protection storage operations, such as based on network load, CPU load, expected duration of the storage operation, rate of change of user activities, frequency of use of affected computer systems, trends, and so on.

DATA PROTECTION SCHEDULING, SUCH AS PROVIDING A  
FLEXIBLE BACKUP WINDOW IN A DATA PROTECTION SYSTEM

BACKGROUND

**[0001]** Computer systems contain large amounts of information. This information includes personal information, such as financial information, customer/client/patient contact information, business information, audio/visual information, and much more. This information also includes information related to the correct operation of the computer system, such as operating system files, application files, user settings, and so on. With the increased reliance on computer systems to store critical information, the importance of protecting information has grown. Traditional storage systems receive an identification of a file to protect, and then create one or more secondary copies, such as backup files, containing the contents of the file. These secondary copies can then later be used to restore the original data should anything happen to the original data.

**[0002]** In corporate environments, protecting information is generally part of a routine process that information technologists perform for many computer systems within an organization. For example, a company might back up critical computing systems related to e-commerce such as databases, file servers, web servers, and so on as part of a daily, weekly, or monthly maintenance schedule. The company may also protect computing systems used by each of its employees, such as those used by an accounting department, marketing department, engineering department, and so forth.

**[0003]** Often, these systems are required to store large amounts of data (e.g. all of a company's data files) during a time period known as a "storage window." The storage window defines a duration and actual time period when the system may perform storage operations. For example, a storage window may be for twelve hours, between 6 PM and 6 AM (that is, twelve non-business hours).

**[0004]** Often, storage windows are rigid and unable to be modified. Therefore, when data storage systems attempt to store increasing data loads, they may need to do so without increasing the time in which they operate. Additionally, many systems

perform daily stores, which may add further reliance on completing storage operations during allotted storage windows.

**[0005]** Additionally, or alternatively, current systems may attempt to store a large number of distinct jobs, or groups of data, chunks of data, and so on. The system may look at each job as a separate storage operation, which often leads to fragmentation on secondary storage devices (tapes, magnetic disks, and so on) that receive data stores as the storage devices develop small gaps of unused space between spaces containing data. In these cases, the system may inefficiently restore stored data because of the fragmentation that occurs during the data storage process.

**[0006]** The foregoing examples of some existing limitations are intended to be illustrative and not exclusive. Other limitations will become apparent to those of skill in the art upon a reading of the Detailed Description below. These and other problems exist with respect to data storage management systems.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0007]** FIG. 1A is a block diagram illustrating an example of components used in data storage operations.

**[0008]** FIG. 1B is a block diagram illustrating an alternative example of components used in data storage operations.

**[0009]** FIG. 1C is a block diagram illustrating an alternative example of components used in data storage operations.

**[0010]** FIG. 2 is a block diagram illustrating an example of a data storage system.

**[0011]** FIG. 3 is a block diagram illustrating an example of components of a server used in data storage operations.

**[0012]** FIG. 4 is a flow diagram illustrating an example of a routine for selecting storage resources in a data storage operation.

**[0013]** FIG. 5 is a flow diagram illustrating an example of a routine for performing a selected storage operation.

**[0014]** FIG. 6 is a flow diagram illustrating the overall process performed in scheduling data storage operations.

**[0015]** FIG. 7 is a table illustrating criteria used in scheduling a data storage operation.

## DETAILED DESCRIPTION

### Overview

**[0016]** Described in detail below is a system for dynamically providing a flexible or "rolling" data protection window that analyzes various criteria to determine an (optimal or near optimal) time for performing data protection or secondary copy operations within certain parameters. While prior systems may have scheduled backups at an exact time (e.g., 2:00 a.m.), the system described herein dynamically determines when to perform the backups and other data storage operations, such as based on network load, CPU load, expected duration of the storage operation, rate of change of user activities, frequency of use of affected computer systems, trends, data priority, compliance factors, and so on.

**[0017]** In some examples, the system first receives a request to perform a storage operation. For example, a data storage system may receive a request to protect all of the data from a particular computer system. The request may include, among other things, a deadline to complete data protection of the data, an identification of storage processes to be used in protecting the data, and/or other criteria used to guide the storage operation.

**[0018]** The system may then review the criteria included in the request as well as performance and other metrics tied to the storage operation to determine how and when to perform the storage operation. For example, the metrics may indicate the use capacity of storage operation components, may predict when a storage network is under (or over) a threshold, may indicate when a media library to which data will be transferred is not in use, and so on.

**[0019]** Then, the system may schedule the data storage operation based on the reviewed metrics. In some cases, the system changes or modifies a previously de-  
termined schedule based on other scheduled storage operations,

new information about network load, or other factors that may affect the successful completion or timing of the storage operation. Also, by dynamically scheduling storage operations, the system is able to avoid at least some conflicts between computer systems over resources needed to perform the storage operations and reduces the impact on primary storage resources. Thus, the system can provide a dynamically determined schedule for data storage operations, such as a schedule that is more likely to be able to accommodate most or all of the desired operations within a storage window.

**[0020]** Therefore, the system dynamically determines a time to perform the storage operation and the components used to facilitate the storage operation based on criteria included in a request to protect data and based on (predicted) characteristics of the storage operation used to carry out the request. That is, the system may review a dynamic storage policy, a policy that provides storage operation instructions to the system based on the needs of a storage request and the characteristics of the system.

**[0021]** In some cases, the dynamic storage policy may be part of a flexible priority system for performing storage operations, or jobs. The priority system may instruct a storage system to perform certain categories of jobs in a predetermined order, while allowing the system to dynamically determine the order of other categories of jobs. For example, the priority system may instruct a storage system to perform all hard coded jobs first (jobs that must be completed within a time window), to perform all soon to expire jobs second (that is, any jobs having a deadline of completion less than a threshold deadline, or previously postponed jobs), and to perform all other jobs based on a flexible and dynamically determined schedule of operations.

**[0022]** Various examples of the system will now be described. The following description provides specific details for a thorough understanding and enabling description of these examples. One skilled in the art will understand, however, that the system may be practiced without many of these details. Additionally, some well-known structures or functions may not be shown or described in detail, so as to avoid unnecessarily obscuring the relevant description of the various examples.

**[0023]** 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.

### Suitable System

**[0024]** Referring to Figure 1A, a block diagram illustrating components of a data storage system with which the data protection scheduling system can be used is shown. The data storage system 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 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 for storage of the data. Secondary storage devices may be magnetic tapes, optical disks, USB, SSD and other similar media, disk and tape drives, and so on. The combination of all of the components (or a device containing all the components) needed to perform a storage operation (e.g., a client, media agent, and secondary storage device) may be referred to as a data stream.

**[0025]** Referring to Figure 1B, a block diagram illustrating components of multiple selectable data streams is shown. Client 111 and any one of multiple media agents 112 may form a stream 110. For example, one stream may contain client 111, media agent 121, and storage device 131, while a second stream (or multiple streams) may use media agent 125, storage device 133, and the same client 111. Additionally, media agents may contain additional subpaths 123, 124 that may increase the number of possible streams for client 111. Examples of subpaths 123, 124 include host bus adapter (HBA) cards, Fibre Channel cards, SCSI cards, wireless paths, and so on. Thus, the system is able to stream data from client 111 to multiple secondary storage devices 113 via multiple media agents 112 using multiple streams.



**[0026]** Referring to Figure 1C, a block diagram illustrating components of alternative multiple selectable data streams is shown. In this example, the system may transfer data from multiple media agents 151, 152 to the same storage device 113. For example, one stream may be from client 141, to media agent 151, to secondary storage device 113, and a second stream may be from client 142, to media agent 152, to secondary storage device 113. Thus, the system is able to copy data to one secondary storage device 113 using multiple streams 110. Additionally, the system may stream data from one client to two media agents and to one storage device. Of course, the system may employ other configurations of stream components not shown in the Figures.

**[0027]** Referring to Figure 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. 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. 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.

**[0028]** The data storage 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. A primary copy of data (for example, the copy of data created by a file) and a secondary copy may be any copy of the primary copy. Example

secondary copies may include snapshot copies, backup copies, HSM copies, archive 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.

**[0029]** 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 that instruct or guide storage operations. The storage policy may define, identify, or indicate 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. For example, the storage policy may be a dynamically changing policy based on factors, metrics and other characteristics of storage operations and requests to perform storage operations. These characteristics may include a job priority (such as a priority provided by the flexible priority system or by an administrator), a scheduling priority (such as a priority based on deadlines of completion of some or all scheduled jobs), the type of data (e.g., exchange and SQL data may have different priorities), and so on. In some cases, a (dynamic) 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, based on a current or predicted status of the storage operation. 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.

**[0030]** Referring to Figure 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 primary or secondary storage. 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, and other data, such as historical storage statistics, storage trend

WO 2009/154806

PCT/US2009/032325

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).

**[0031]** Further examples of suitable systems and methods may be found in U.S. Patent Application Publication No. 2008/0288947, published 20 November 2008, entitled SYSTEMS AND METHODS OF DATA STORAGE MANAGEMENT, SUCH AS DYNAMIC DATA STREAM ALLOCATION.

#### Using a Data Storage Window to Affect Storage Operations

**[0032]** In some cases, the system may modify or change storage operations based on a review of a storage window. In some cases the system considers the data storage window to be dynamically determined period of time when the system may perform data storage operations. In other cases, the storage window is rigid for some or all storage operations and the system should complete such any associated data transfers within the window. Therefore, a dynamic review of the storage window during or prior to data storage operations may assist storage systems in completing storage tasks within an allotted window of time.

**[0033]** Referring to Figure 4, a flow diagram illustrating a routine 400 as an example of selecting storage resources in a data storage operation begins in step 410, where the system may compare the storage window with an estimated time remaining to complete data storage operations. For example, the system may calculate an estimate of the time required to complete all pending job transfers, and compare the estimated time with the time allotted to run data transfers. In step 420, if the time allotted is larger than the time estimate, routine 400 ends and the system performs scheduled operations within the window, else routine 400 proceeds to step 430. In step 430, the system performs corrective operations. Examples of corrective operations may include the dynamic stream management discussed above, using more resources, selecting a subset of the remaining jobs to store, sending remaining jobs to an alternative or "standby" data storage system, moving or modifying the window for certain jobs, obtaining guidance from the flexible priority system, moving certain jobs outside the window (to be discussed herein), and so on.

After performing corrective actions, routine 400 proceeds back to step 420, and compares the new estimated time against the time allotment.

**[0034]** In some cases, the system may review, monitor, or track default pathways (such as streams) and modify storage operations if there is not enough time in the storage window to complete all data transfers using the default pathways. For example, the system may select high-speed pathways instead of default pathways for data of a certain type and nature (such as high priority or unprotected data).

**[0035]** The system may perform routine 400 as infrequently or as often as necessary, depending on the needs of the system or the progress of data storage operations. For example, the system may track the performed corrective measures to determine their effectiveness, and determine a frequency in which to perform the routine within the window. When a corrective measure causes the estimated time of completion to fall within the storage window, the system may perform the routine at a lesser frequency than when the estimated time of completion is not within the storage window. Also, the system may perform routine 400 to obtain information about data storage operations, to be used in performing corrections at a later time. The system may determine patterns, statistics, metrics, criteria, characteristics and/or historical information about the storage operations and associated resources from routine 400. For example, in a 12 hour time allotted storage window, the system may run routine 400 twelve times, once per hour. Comparing the twelve iterations, the system may determine a pattern of high resource use, low resource use, and so on, and modify future data storage operations accordingly.

**[0036]** In some cases, the system may be able to delay the transfer of some types of data in order to transfer other types of data within the storage window. Referring to Figure 5, a flow diagram illustrating an example of performing a selected storage operation is shown. In step 510, the system may compare the storage window with an estimated time remaining to complete data storage operations. For example, the system may estimate the time required to complete all pending job transfers, and compare the estimated time with the time allotted to run transfer operations. In step 520, if the time allotted is larger than the time estimate, routine 500 proceeds to step 530, the system may select certain jobs to store

based on information received from the flexible priority system, and delay other jobs. For example, the system may be able to store some types of data outside of the storage window, to be discussed herein. The system may then select these jobs and move them out of the job queue, to a delayed jobs queue.

**[0037]** After selecting "priority" jobs, routine 500 proceeds back to step 520, and compares the new estimated time against the time allotment. The system transfers all "priority" jobs, and only goes to the delayed job queue after the main job queue is empty of priority jobs. The system may then transfer the delayed jobs during the remaining time of the storage window, may transfer the jobs outside of the job window, or may be able to send the jobs to the next scheduled storage operation or data transfer, and transfer the remaining jobs during that operation.

**[0038]** Assigning some jobs a priority may be arbitrary or contingent on the needs of the system. The system may assign priorities to types of files (e.g., metadata, such as user-defined content) or jobs within a storage policy 210. The system may enable users to determine what types of jobs are priority jobs. The system may maintain some jobs as always being priority, or may change these preferences on a case-by-case basis. For example, a user may set a policy to flag all financial data as "priority," and set a policy to never flag email data (or email from certain user groups) as "priority." However, in some cases, the reverse may be more desirable. The system may also assign a higher priority to a job that was delayed or missed in a previous window. In some cases, the system may update or modify metadata, data classification or other preferences, and may assign priorities to characteristics of data as well as to data.

**[0039]** As discussed herein, a flexible priority system may instruct a data storage system to perform certain jobs before other jobs. The flexible priority system may identify jobs having the highest priority (such as hard coded jobs), jobs having a next highest priority (such as jobs set to expire, jobs required to be completed within a storage window), and jobs having a flexible priority that changes (based on the factors described herein) in order to optimize the performance of the data storage system. Further details will now be discussed.

### Flexible Storage Window

**[0040]** The system looks at various criteria to determine an (optimal or near optimal) period for performing a storage operation. The criteria can include one or more of the following: job priority, types of data within the job, network traffic or load, a disk load, CPU usage, expected backup duration, and so on. For example, the system may gather relevant metrics or other parameters during a predetermined period, such as over a 24- to 48-hour period, which may identify within a one to two day period a time in which to perform a data storage operation for a job that was not completed during a storage window.

**[0041]** Figure 6 is a flow diagram that illustrates the overall process performed by the data protection scheduling system in one embodiment. In step 610, the system receives a request to perform a storage operation. The storage operation may include one or more request criteria that indicate parameters of the request, such as a deadline for completing the request, a desired frequency (e.g., once per hour) that the request should be performed, and so forth. In step 620, the system applies one or more criteria to determine a suitable time to perform the request. For example, the system may determine a priority for the request, the computer systems that will be used for performing the request, the network loads on or between the computer systems, the availability of resources on the computing systems (e.g., processor time, memory, and storage space), the availability of a media library or storage media, and so on.

**[0042]** The criteria may also take into account any request criteria. For example, if the network load is at its lowest point in two hours, and its second to lowest point in 45 minutes, then the system may select the second lowest point to perform the storage operation if a request criteria requests that the storage operation be completed within the next hour. In step 630, the system schedules the data storage operation based on the performance criteria, among other factors. The system may revisit the schedule periodically based on changes to conditions that affect the scheduling of requests. For example, a new media agent may be added to the network to which some of the load can be distributed. The system dynamically schedules performing storage operations.

**[0043]** The system may also look at historical data or other periods outside a one to two day period. For example, by analyzing one or more months of data, the system can identify trends or historical data. Thus, the system can determine a rate of change, frequency of use, particularly with respect to particular data, data sources, departments, etc. Thus, the system may determine that when a particular user goes home, network activity drops. They can notice trends, such as trends based on vacations, school schedules, etc., and perform actions based on those events. Thus, users, processes, historical trends, pattern recognition, rate of change and amount of access may all be analyzed by the system to determine an optimal or new optimal window for backup processes.

**[0044]** In some examples, a priority system determines when the system is to perform required and/or requested storage operations. In some cases, the priority system instructs a storage system to perform all hard coded jobs first (jobs that must be completed within a time window). The hard coded jobs may be certain daily backups or other data required to be copied within a time window. Once these jobs are completed, the storage system performs all soon to expire jobs second (that is, any jobs having a deadline of completion less than a threshold deadline, or previously postponed jobs), and performs all other jobs based on the flexible and dynamically determined schedule of operations described herein. Such a flexible priority system ensures, among other things, that a data storage system can complete all data storage operations when they are required to be completed, even when certain operations are not completed when originally scheduled.

#### Dynamically Determining when to Perform Storage Operations

**[0045]** The system may review historical data or information about previously performed jobs or usage of storage operation resources, and predict when to start a job, even if the start time falls outside a storage window. The data may identify a time when the usage resources are predicted to be at a "low" usage level (e.g., below a threshold level of a load on the resources), and start the job at that time. For example, the system may determine that the usage level for certain resources drops during a typical lunch hour (e.g., 12:00 to 1:00 P.M.), and select a job to complete during that time that is expected to be completed within an hour.

**[0046]** Once a job starts, the system may also monitor the usage of the resources and make decisions based on the monitored usage. For example, the system may start a job and after a certain time the job (and other jobs or operations) may cause the load on the resources to exceed a threshold value. Such an occurrence may cause the system to stop the job until the load drops below the threshold value (or, until the predicted load drops below the threshold value). At this time, the system may restart the job using the resources. Additionally, the system may assign a higher priority to partially completed jobs, which may affect the threshold value.

**[0047]** As discussed with respect to Figure 6, the system identifies and applies criteria to determine a time in which to start or perform a storage operation, such as a job. In some cases, the system applies multiple criteria. For example, the system may look at (1) criteria associated with the job or jobs, (2) criteria associated with the data storage system, and/or (3) criteria associated with historical information with respect to the job or jobs, the system and its resources, and so on. The system may then use the criteria as weighting factors within rules or algorithms that determine threshold requirements for starting jobs.

**[0048]** Referring to Figure 7, a table 700 including criteria used in scheduling a job is shown. The criteria include a jobs criterion 710, system criterion 720, and historical data criterion 730. Each of the criterion includes various metrics 715 associated with that criterion. For example, for a given job, the job criterion 710 may include a metric attributed to the time window in which to complete the job, a metric attributed to an assigned priority for the job, and so on. Similarly, the system criterion 720 may be affected by the current usage of system resources, a scheduled usage, a predicted usage, and so on. The historical criterion 730 may consider historical trends or patterns, such as those discussed herein.

**[0049]** Thus, the system may look to table 700 when scheduling a job. For example, after a storage window ends, a number of jobs (e.g., two) remain to be performed. Job A is a high priority job but historical data determines that the job causes a high load to system resources (e.g. CPU load). Job B is a lower priority job but causes a lower load to the system resources. Using the table, the system may de



Job A: Start job whenever system resources usage is less than 40% threshold from a maximum system load (may be a calculation of current usage plus predicted usage due to job)

Job B: Start job whenever system resources usage is less than 70% threshold.

Also preempt Job B with Job A when Job A threshold is obtained.

**[0050]** Thus, when a usage factor of system resources moves below the 70% threshold or is predicted to do so by historical data (such as the lunch hour example described herein), a data storage system begins performing Job B. Should the usage factor move below the 40% threshold, Job B is stopped and Job A begins. Once the usage factor moves above the 40% threshold (but less than 70%), Job A is stopped and Job B is restarted. Should the usage factor move above the 70% threshold, both jobs are stopped.

**[0051]** As discussed herein, the rules assigned to the jobs may dynamically change during storage operations. For example, if Job B is more than half way completed, the system may assign the job a higher priority than the priority assigned to Job A. In this example, Job B would continue to perform even when the usage factor moves below the 40% threshold.

**[0052]** In some examples, the system looks at a large period of time (such as a months' worth of data) to identify lulls in the system resources, and make predictions about the usage of the system resources. The system then determines when to start a job based on the predictions, and schedules the job.

**[0053]** In some examples, the flexible priority system exempts certain jobs from the dynamically determined scheduling. The system may determine priorities for particular jobs, departments, users, and so forth. These jobs/departments/users/etc. may be exempt from flexible scheduling, and instead have a set data storage window, such as a window specified by an administrator. For example, hard coded and expiring jobs may not be flexibly scheduled. Thus, the flexible priority system may increase the priority of certain jobs/departments/users/etc. based on historical data, current performance, and so on. As one example, if a client was last in the queue, the client may be moved closer to the top of the queue for the current scheduling on various factors.

**[0054]** In some examples, the system employs a scoring metric. For example, the system may add some weighting to encourage a data storage operation to be scheduled sooner if a particular job did not get processed on a given date, is close to expiration, was started but not completed, and so on. As another example, a job may get a negative weighting for the next data storage window if the job was placed high in the job queue and completed for a current data storage window.

**[0055]** In some examples, the system may employ an alternating approach to flexibly assigning jobs to available storage windows. For example, the flexible priority system may identify jobs having relatively equal priority and schedule the jobs to be performed one-by-one, or in parallel where possible based on available resources.

**[0056]** In some examples, the system looks at request criteria to dynamically determine the scheduling of data storage operations. For example, a request may specify that a data storage operation must occur at least once per hour, such as for protecting financial transaction data or other frequently changing data. The data protection scheduling system then looks at past scheduling of the job to determine the current schedule of the job so that the system schedules the job in a way that satisfies the request criteria. Such request criteria may identify hard coded or other "highest" priority jobs, or may identify jobs always available as being flexible to the system.

### Conclusion

**[0057]** The system may perform some or all of the above examples in combination with one another. For example, the system may use aspects of dynamic stream management to choose a stream to transfer a data store job, and may transfer that job within pre-allocated disk space for multiple jobs. The system may trigger dynamic stream management processes based on a review of the storage window.

**[0058]** Systems and modules described herein may comprise software, firmware, hardware, or any combination(s) of software, firmware, or hardware suitable for the purposes described herein. Software and other modules may reside on personal computers, computerized tablets, PDAs, and other devices for the purposes described herein. In other words, the

software and other modules described herein may be 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 aspects of 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," "server," "host," "host system," and the like are generally used interchangeably herein, and refer to any of the above devices and systems, as well as any data processor. Furthermore, 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.

**[0059]** Software and other modules may be accessible via local memory, via a network, via a browser or other application in an ASP context, or via other means suitable for the purposes described herein. Examples of the technology 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. Data structures described herein may comprise computer files, variables, programming arrays, programming structures, or any electronic information storage schemes or methods, or any combinations thereof, suitable for the purposes described herein. User interface elements described herein may comprise elements from graphical user interfaces, command line interfaces, and other interfaces suitable for the purposes described herein. Screenshots presented and described herein can be displayed differently as known in the art to input, access, change, manipulate, modify, alter, and work with information.

**[0060]** Examples of the technology may be stored or distributed on computer-readable media including magnetically or optically readable computer discs, hard-wire 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).

**[0061]** 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." As used herein, the terms "connected," "coupled," or any variant thereof, means any connection or coupling, either direct or indirect, between two or more elements; the coupling of connection between the elements can be physical, logical, or a combination thereof. 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, 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.

**[0062]** While certain aspects of the technology are presented below in certain claim forms, the inventors contemplate the various aspects of the technology in any number of claim forms. For example, while only one aspect of the technology is recited as a means-plus-function claim under 35 U.S.C. sec. 112, other aspects may likewise be embodied as a means-plus-function claim. 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 technology.

**[0063]** The above detailed description of examples of the technology is not intended to be exhaustive or to limit the invention to the precise form disclosed above. For example, although certain types of storage operations have been described, a scheduling system can dynamically schedule many specific embodiments of, and examples for, the invention

WO 2009/154806

PCT/US2009/032325

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 to provide alternative or subcombinations. 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.

**[0064]** The teachings of the technology 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 examples.

Aspects of the invention can be modified, if necessary, to employ the systems, functions, and concepts of the various references described above to provide yet further examples of the technology.

**[0065]** These and other changes can be made to the invention in light of the above Detailed Description. While the above description describes 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 and method for classifying and transferring information may vary considerably in its 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 technology under the claims. While certain aspects of the technology are presented below in certain claim forms, the inventors contemplate the various aspects of the technology in any number of claim forms. For example, while only one aspect of the technology is recited as stored in a computer memory, other aspects may likewise be stored in a computer memory. 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 technology.

## WHAT IS CLAIMED IS:

1. A method for dynamically determining a schedule for performing a data storage operation, the method comprising:  
5 receiving a request to perform a storage operation;  
determining whether the storage operation is exempt from dynamically determined scheduling, and  
when the storage operation is not exempt from dynamically determined scheduling, then:  
10 applying one or more performance criteria to dynamically determine a time to perform the storage operation,  
wherein the applying one or more performance criteria includes adding weighting to the storage operation to change a previously determined schedule of  
15 operations, if the storage operation did not get processed previously;  
scheduling the data storage operation based on the one or more applied performance criteria, and  
monitoring usage of resources after starting the storage operation for dynamically altering the schedule for performing the storage operation;  
20 wherein one or more of the applied performance criteria are unassociated with an amount of data to store under the storage operations, and  
25 wherein at least one of the one or more applied performance criteria relate to a predicted usage of computing resources in performing the storage operation;

wherein scheduling the data storage operation comprises:

- 5 determining a storage window at least partly based on a future network load;
- 10 determining a schedule for completing multiple requested data storage operations within the storage window,
- 15 comparing, during execution of the schedule, the storage window with an estimated time remaining to complete data storage operations;
- delaying, based on the comparison, storage of some types of data in order to store other types of data within the storage window;
- and
- sending some data storage jobs to other data storage systems;

when the storage operation is exempt from dynamically determined scheduling, then:

- 20 performing the storage operation in an order in which hard coded storage operations to be completed within a time window are performed first, followed by
- performing storage operations having a deadline of completion less than a threshold deadline.

25

2. The method of claim 1 wherein the request includes a deadline for performing the storage operation based on past trends of storage operations.



3. The method of claim 1 wherein the storage operation includes one of storing, receiving and preparing data to be stored, copied or backed up, and the request includes a desired frequency for performing the storage operation.  
5
4. The method of claim 1 wherein the request includes a request criterion and wherein applying at least one performance criterion takes into account the request criterion.
- 10 5. The method of claim 1 wherein at least one performance criterion of the one or more performance criteria relates to a time when resources of the computer systems affected by the storage operation are unlikely to be in use, a time when the network will have a relatively lower load, or a time when a media library to which data  
15 will be transferred by the storage operation is not in use.
6. The method of claim 1 including reviewing the determined schedule before the storage operation starts and changing the schedule based on other scheduled storage operations.  
20
7. The method of claim 1 including reviewing the determined schedule before the storage operation starts and changing the schedule based on information received after the storage operation was last scheduled.  
25
8. The method of claim 1 wherein scheduling the data storage operation avoids at least some conflicts between computer systems over resources needed to perform the storage operation.

9. A non-transitory, computer-readable medium storing instructions, which, when executed by at least one computing device, performs a method for dynamically determining a schedule for performing a data storage operation, the method comprising:
- 5 receiving a request to perform a storage operation;  
determining if the storage operation is exempt from dynamically determined scheduling, and  
when the storage operation is not exempt from dynamically determined scheduling, then:
- 10 applying one or more performance criteria to dynamically determine a time to perform the storage operation, wherein the applying one or more performance criteria includes adding weighting to the storage operation to change a previously determined schedule of operations, if the storage operation did not get processed previously; and
- 15 scheduling the data storage operation based on the one or more applied performance criteria, and  
monitoring usage of resources after starting the storage operation for dynamically altering the schedule for performing the storage operation;
- 20 wherein one or more of the applied performance criteria are unassociated with an amount of data to store under the storage operations, and
- 25 wherein one or more applied performance criteria relate to a predicted usage of computing resources in performing the storage operation;  
wherein scheduling the data storage operation comprises:

5 determining a storage window at least partly  
 based on a future network load;  
 determining a schedule that is sufficient to com-  
 10 plete multiple requested data storage opera-  
 tions within the storage window,  
 comparing, during execution of the schedule, the  
 storage window with an estimated time re-  
 maining to complete data storage opera-  
 tions;  
 15 delaying, based on the comparison, storage of  
 some types of data in order to store other  
 types of data within the storage window;  
 and  
 sending some data storage jobs to other data stor-  
 20 age systems;  
 when the storage operation is exempt from dynamically determined  
 scheduling, then:  
 performing the storage operation in an order in which hard  
 coded storage operations to be completed within a time  
 25 window are performed first, followed by  
 performing storage operations having a deadline of  
 completion less than a threshold deadline.

10. The non-transitory, computer-readable medium of claim 9 wherein  
 25 the request includes a deadline for performing the storage operation  
 based on past trends of storage operations.

11. The non-transitory, computer-readable medium of claim 9 wherein the request includes a desired frequency for performing the storage operation.
- 5 12. The non-transitory, computer-readable medium of claim 9 wherein the request includes a request criterion and wherein applying at least one performance criterion takes into account the request criterion.
- 10 13. The non-transitory, computer-readable medium of claim 9 wherein at least one performance criterion of the one or more performance criteria relates to a time when resources of the computer systems affected by the storage operation are unlikely to be in use.
- 15 14. The non-transitory, computer-readable medium of claim 9 wherein at least one performance criterion of the one or more performance criteria relates to a time when the network will have a relatively lower load.
- 20 15. The non-transitory, computer-readable medium of claim 9 wherein at least one performance criterion of the one or more performance criteria relates to a time when a media library to which data will be transferred by the storage operation is not in use.
- 25 16. The non-transitory, computer-readable medium of claim 9 including reviewing the determined schedule before the storage operation starts and changing the schedule based on other scheduled storage operations.

17. The non-transitory, computer-readable medium of claim 9 including reviewing the determined schedule before the storage operation starts and changing the schedule based on information received after the storage operation was last scheduled.
- 5
18. The non-transitory, computer-readable medium of claim 9 wherein scheduling the data storage operation avoids at least some conflicts between computer systems over resources needed to perform the storage operation.
- 10
19. A system to dynamically determine a schedule for performing a data storage operation, the system comprising:  
a processor; and  
a memory for storing computer-executable instructions, said instructions, when executed by the processor, causing the processor  
15 to implement a method comprising:  
receiving a request to perform a storage operation;  
determining if the storage operation is exempt from dynamically  
determined scheduling, and  
20 when the storage operation is not exempt from dynamically determined scheduling, then:  
applying one or more performance criteria to dynamically  
determine a time to perform the storage operation,  
wherein the applying one or more performance criteria in-  
25 cludes adding weighting to the storage operation to  
change a previously determined schedule of operations,  
if the storage operation did not get processed previously;

scheduling the data storage operation based on the one or  
more applied performance criteria, and  
monitoring usage of resources after starting the storage operation for dynamically altering the schedule for performing the storage operation;  
5 wherein one or more applied performance criteria are unassociated with an amount of data to store under the storage operations, and  
wherein one or more applied performance criteria relate  
10 to a predicted usage of computing resources in performing the storage operation;  
wherein the scheduling the data storage operation comprises:  
determining a storage window at least partly  
15 based on a future network load;  
determining a schedule for completing multiple requested data storage operations within the storage window,  
comparing, during execution of the schedule, the  
20 storage window with an estimated time remaining to complete data storage operations;  
delaying, based on the comparison, storage of  
25 some types of data in order to store other types of data within the storage window;  
and  
sending some data storage jobs to other data storage systems;

when the storage operation is exempt from dynamically determined scheduling, then:

performing the storage operation in an order in which hard coded storage operations to be completed within a time window are performed first, followed by performing storage operations having a deadline of completion less than a threshold deadline.

5

20. The system of claim 19 wherein the request includes a deadline for performing the storage operation based on past trends of storage operations.

10

21. The system of claim 19 wherein the request includes a desired frequency for performing the storage operation.

15

22. The system of claim 19 wherein the request includes a request criterion and wherein the means for applying at least one performance criterion takes into account the request criterion.

23. The system of claim 19 wherein at least one performance criterion of the one or more performance criteria relates to a time when resources of the computer systems affected by the storage operation are unlikely to be in use.

20

24. The system of claim 19 wherein at least one performance criterion of the one or more performance criteria relates to a time when the network will have a relatively lower load.

25

25. The system of claim 19 wherein at least one performance criterion of the one or more performance criteria relates to a time when a media library to which data will be transferred by the storage operation is not in use.
- 5
26. The system of claim 19, wherein the method further comprises reviewing the determined schedule before the storage operation starts and means for changing the schedule based on other scheduled storage operations.
- 10
27. The system of claim 19, wherein the method further comprises reviewing the determined schedule before the storage operation starts and changing the schedule based on information received after the storage operation was last scheduled.
- 15
28. The system of in claim 19 wherein the scheduling the data storage operation avoids at least some conflicts between computer systems over resources needed to perform the storage operation.
- 20
29. The method of claim 1, further including:  
looking at a period of time to identify a lull in system resources;  
making predictions about usage of the system resources; and  
starting a new storage operation based on the predictions.
- 25
30. The method of claim 1, further including adding weighting to an expiring job and a job that was started but not completed.
31. The method of claim 1 further including:



dynamically determining the schedule based on the resources required by each request and a predicted availability of resources at various times within the storage window.

- 5 32. The system of claim 19 wherein the processor controls a database agent to store historical data about an availability of resources for performing storage operations.
- 10 33. The system of claim 19, wherein the processor implements adding weighting to an expiring job and a job that was started but not completed.
- 15 34. The system of claim 19, wherein the processor further:  
looks at a period of time to identify a lull in system resources;  
makes predictions about usage of the system resources; and  
starts a new storage operation based on the predictions.

1/7

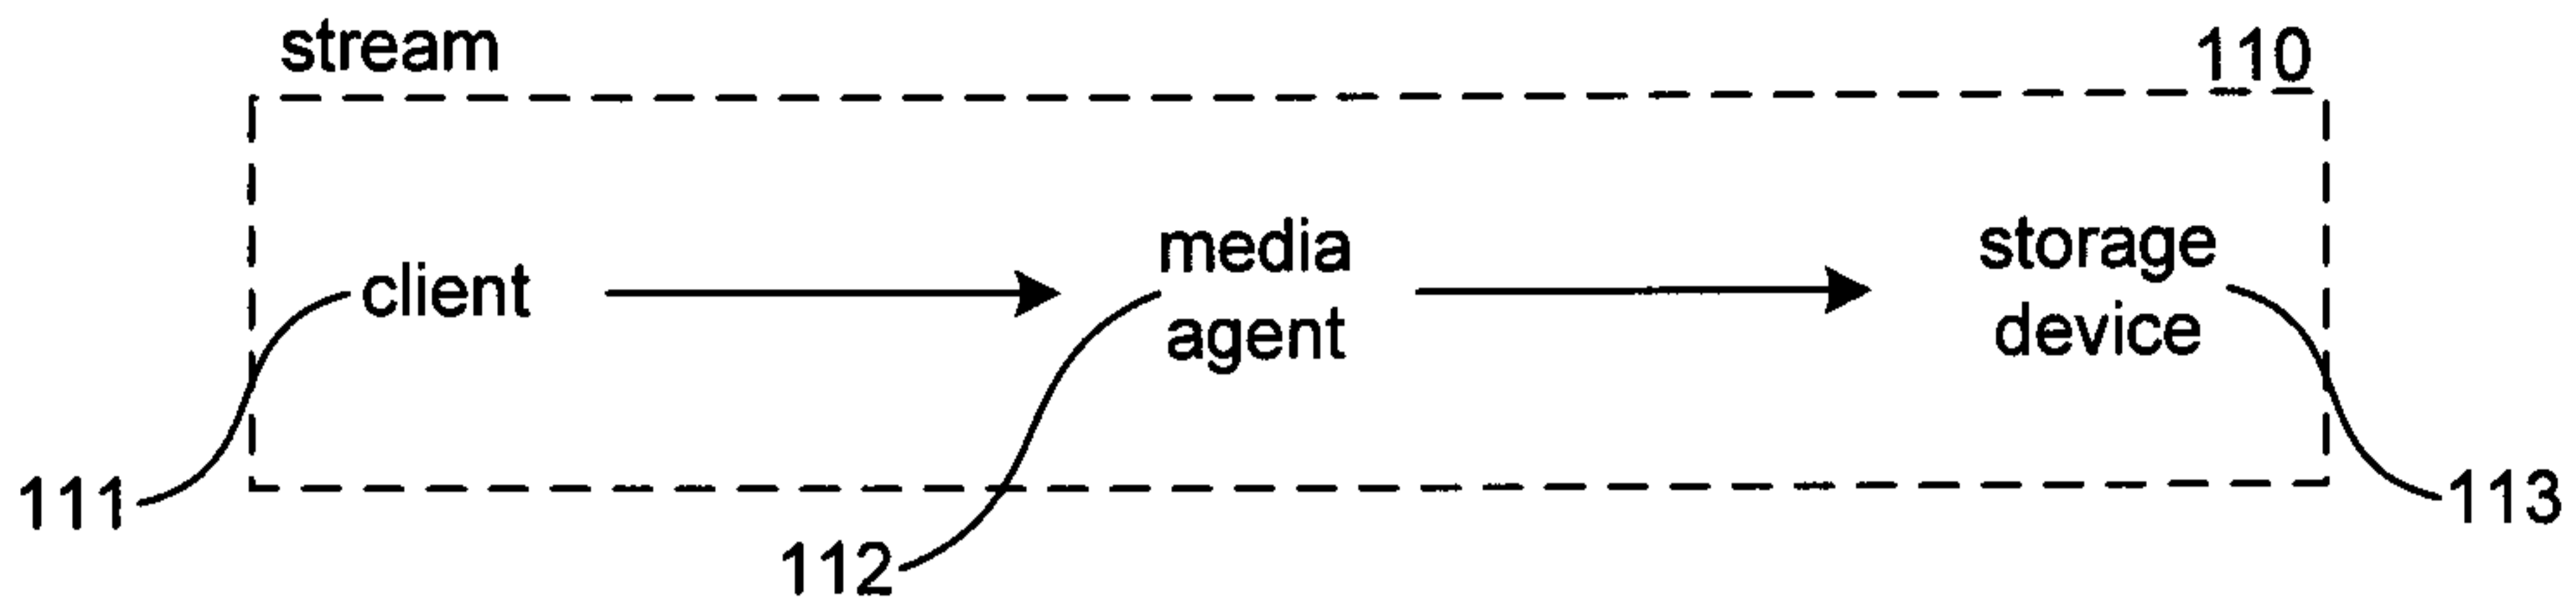


FIG. 1A

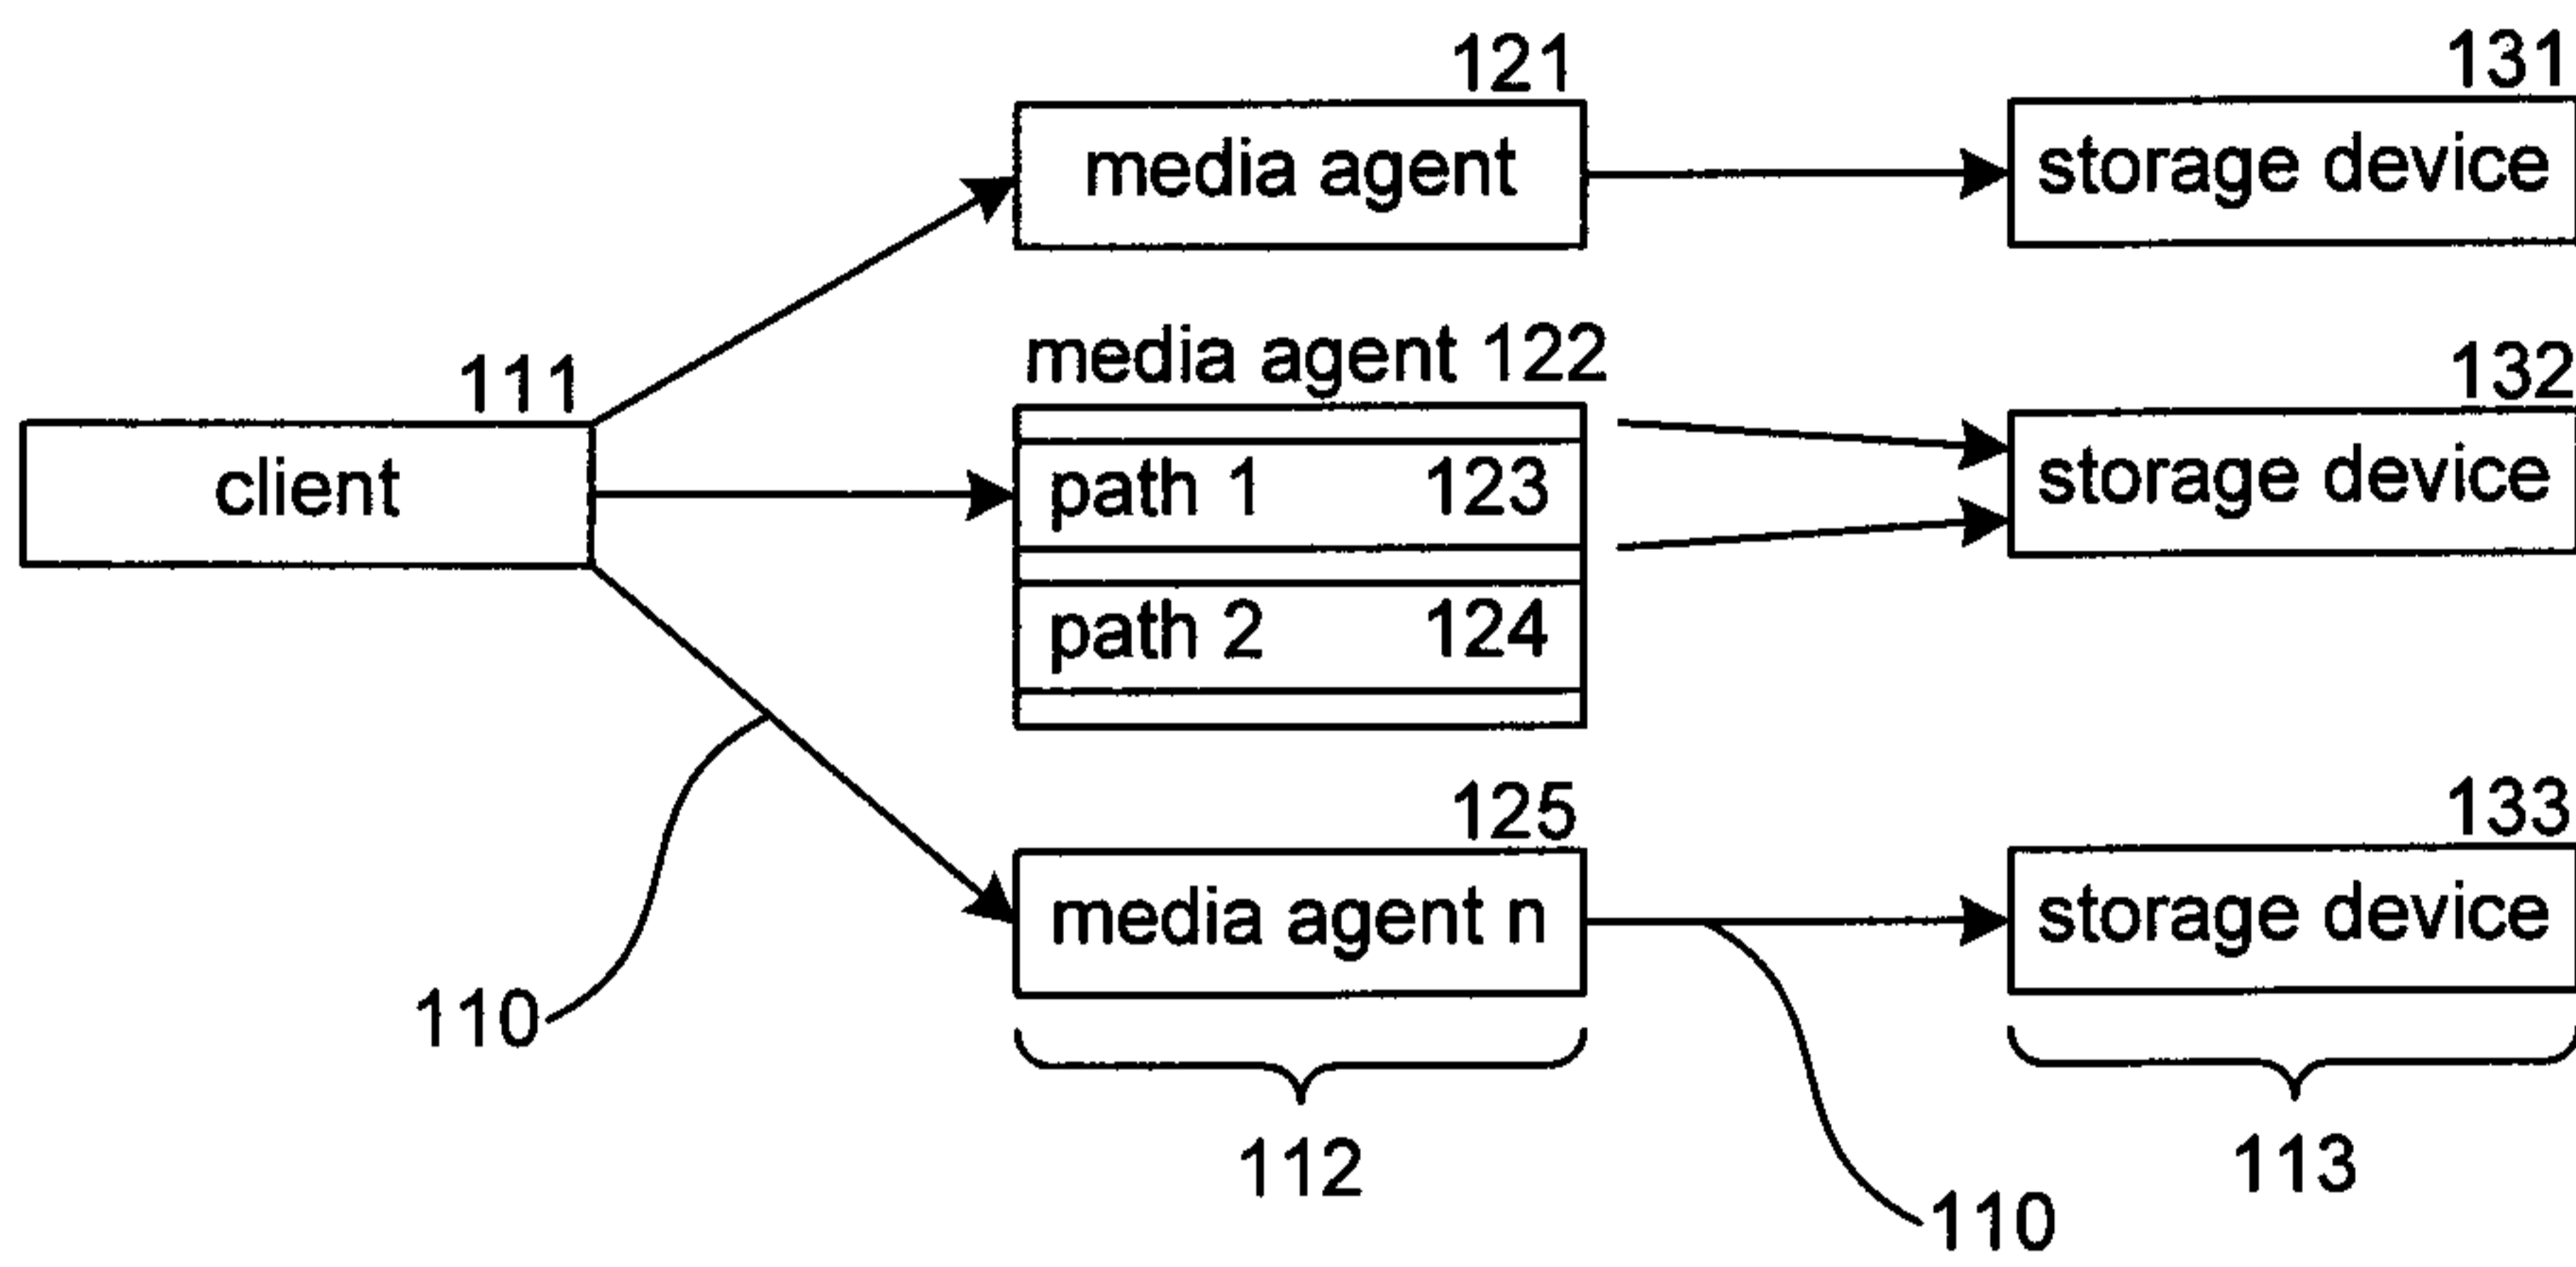


FIG. 1B

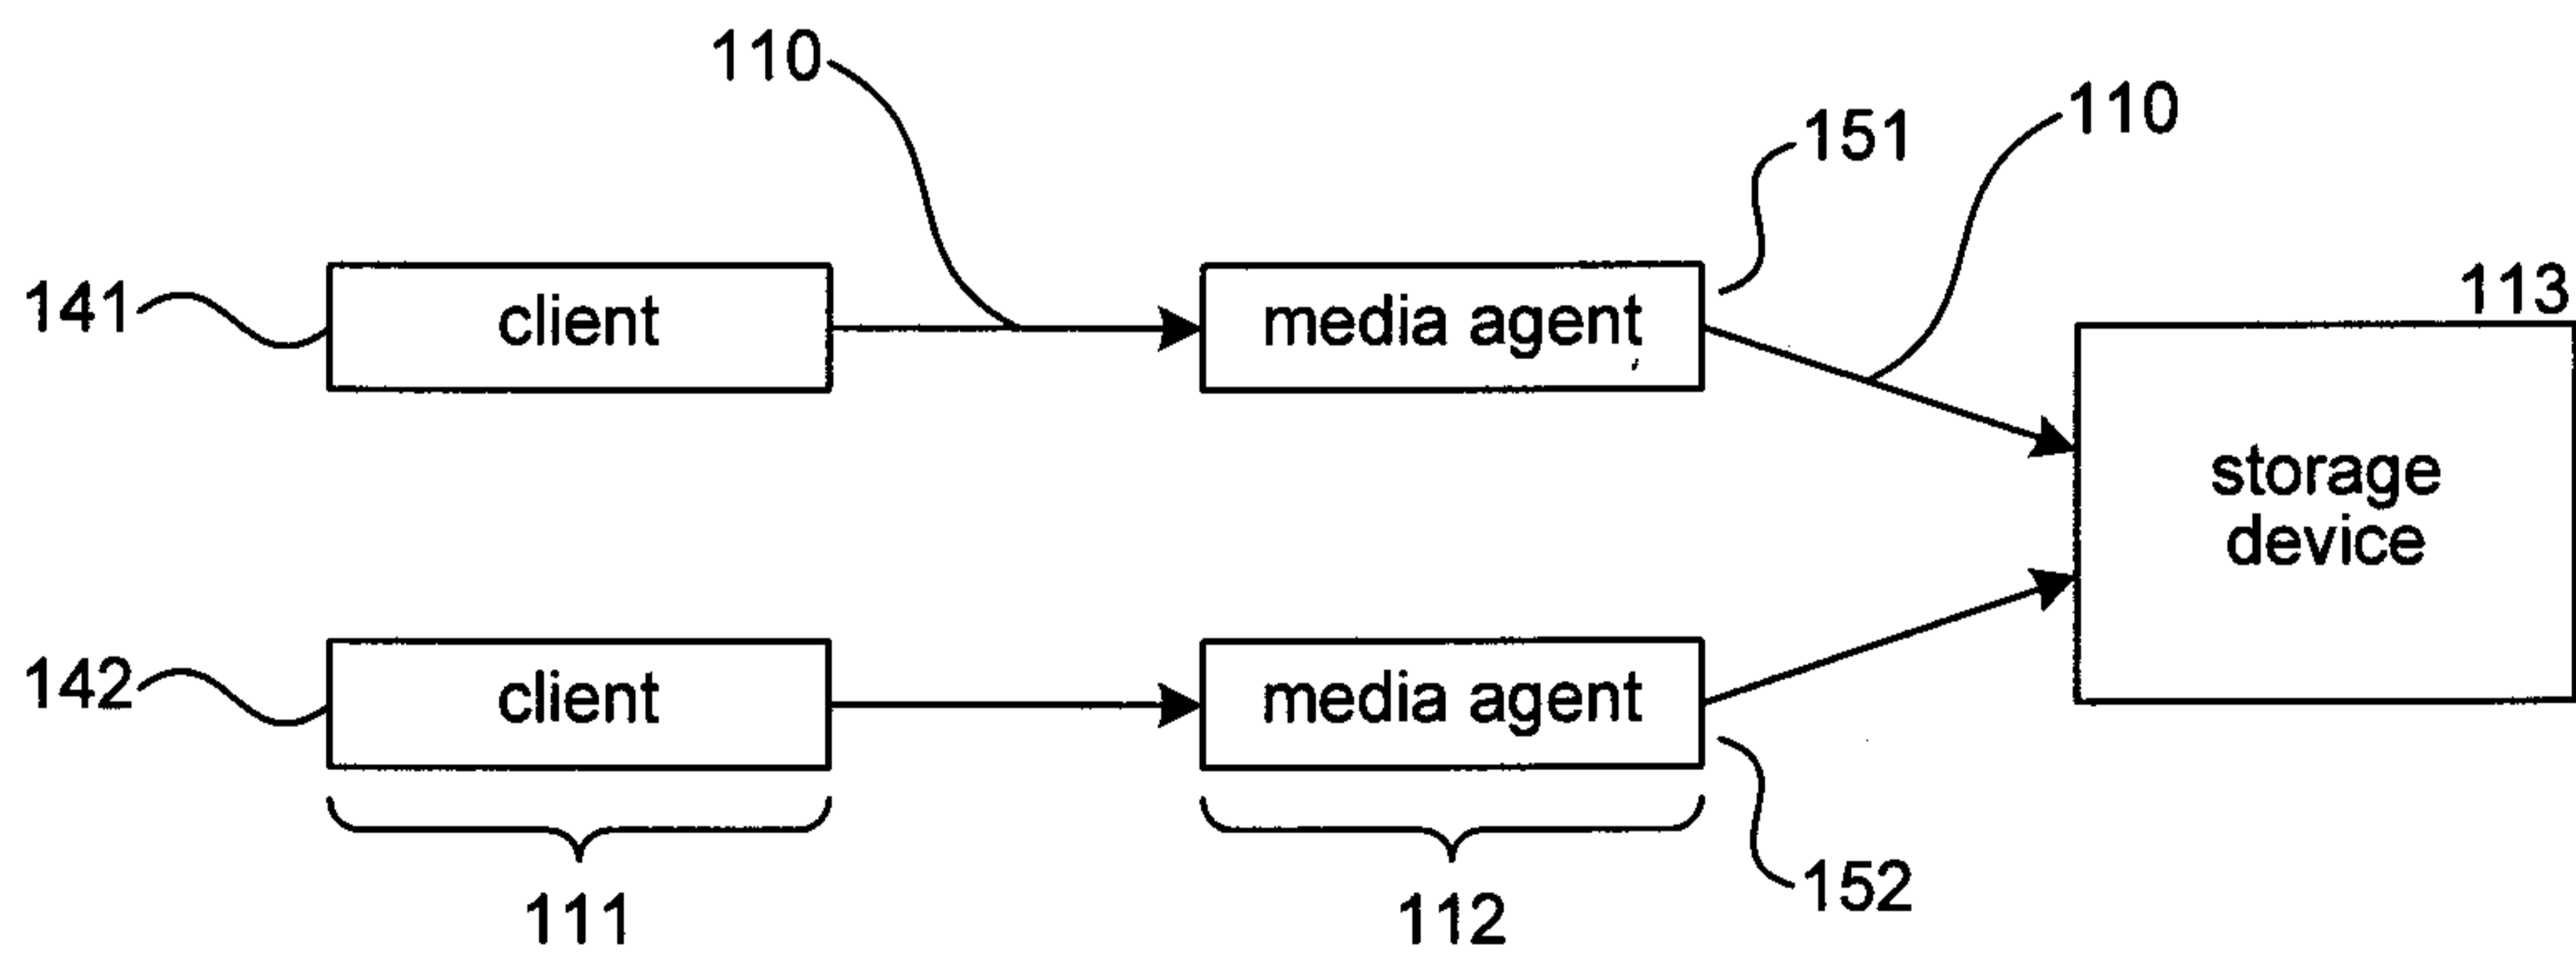


FIG. 1C

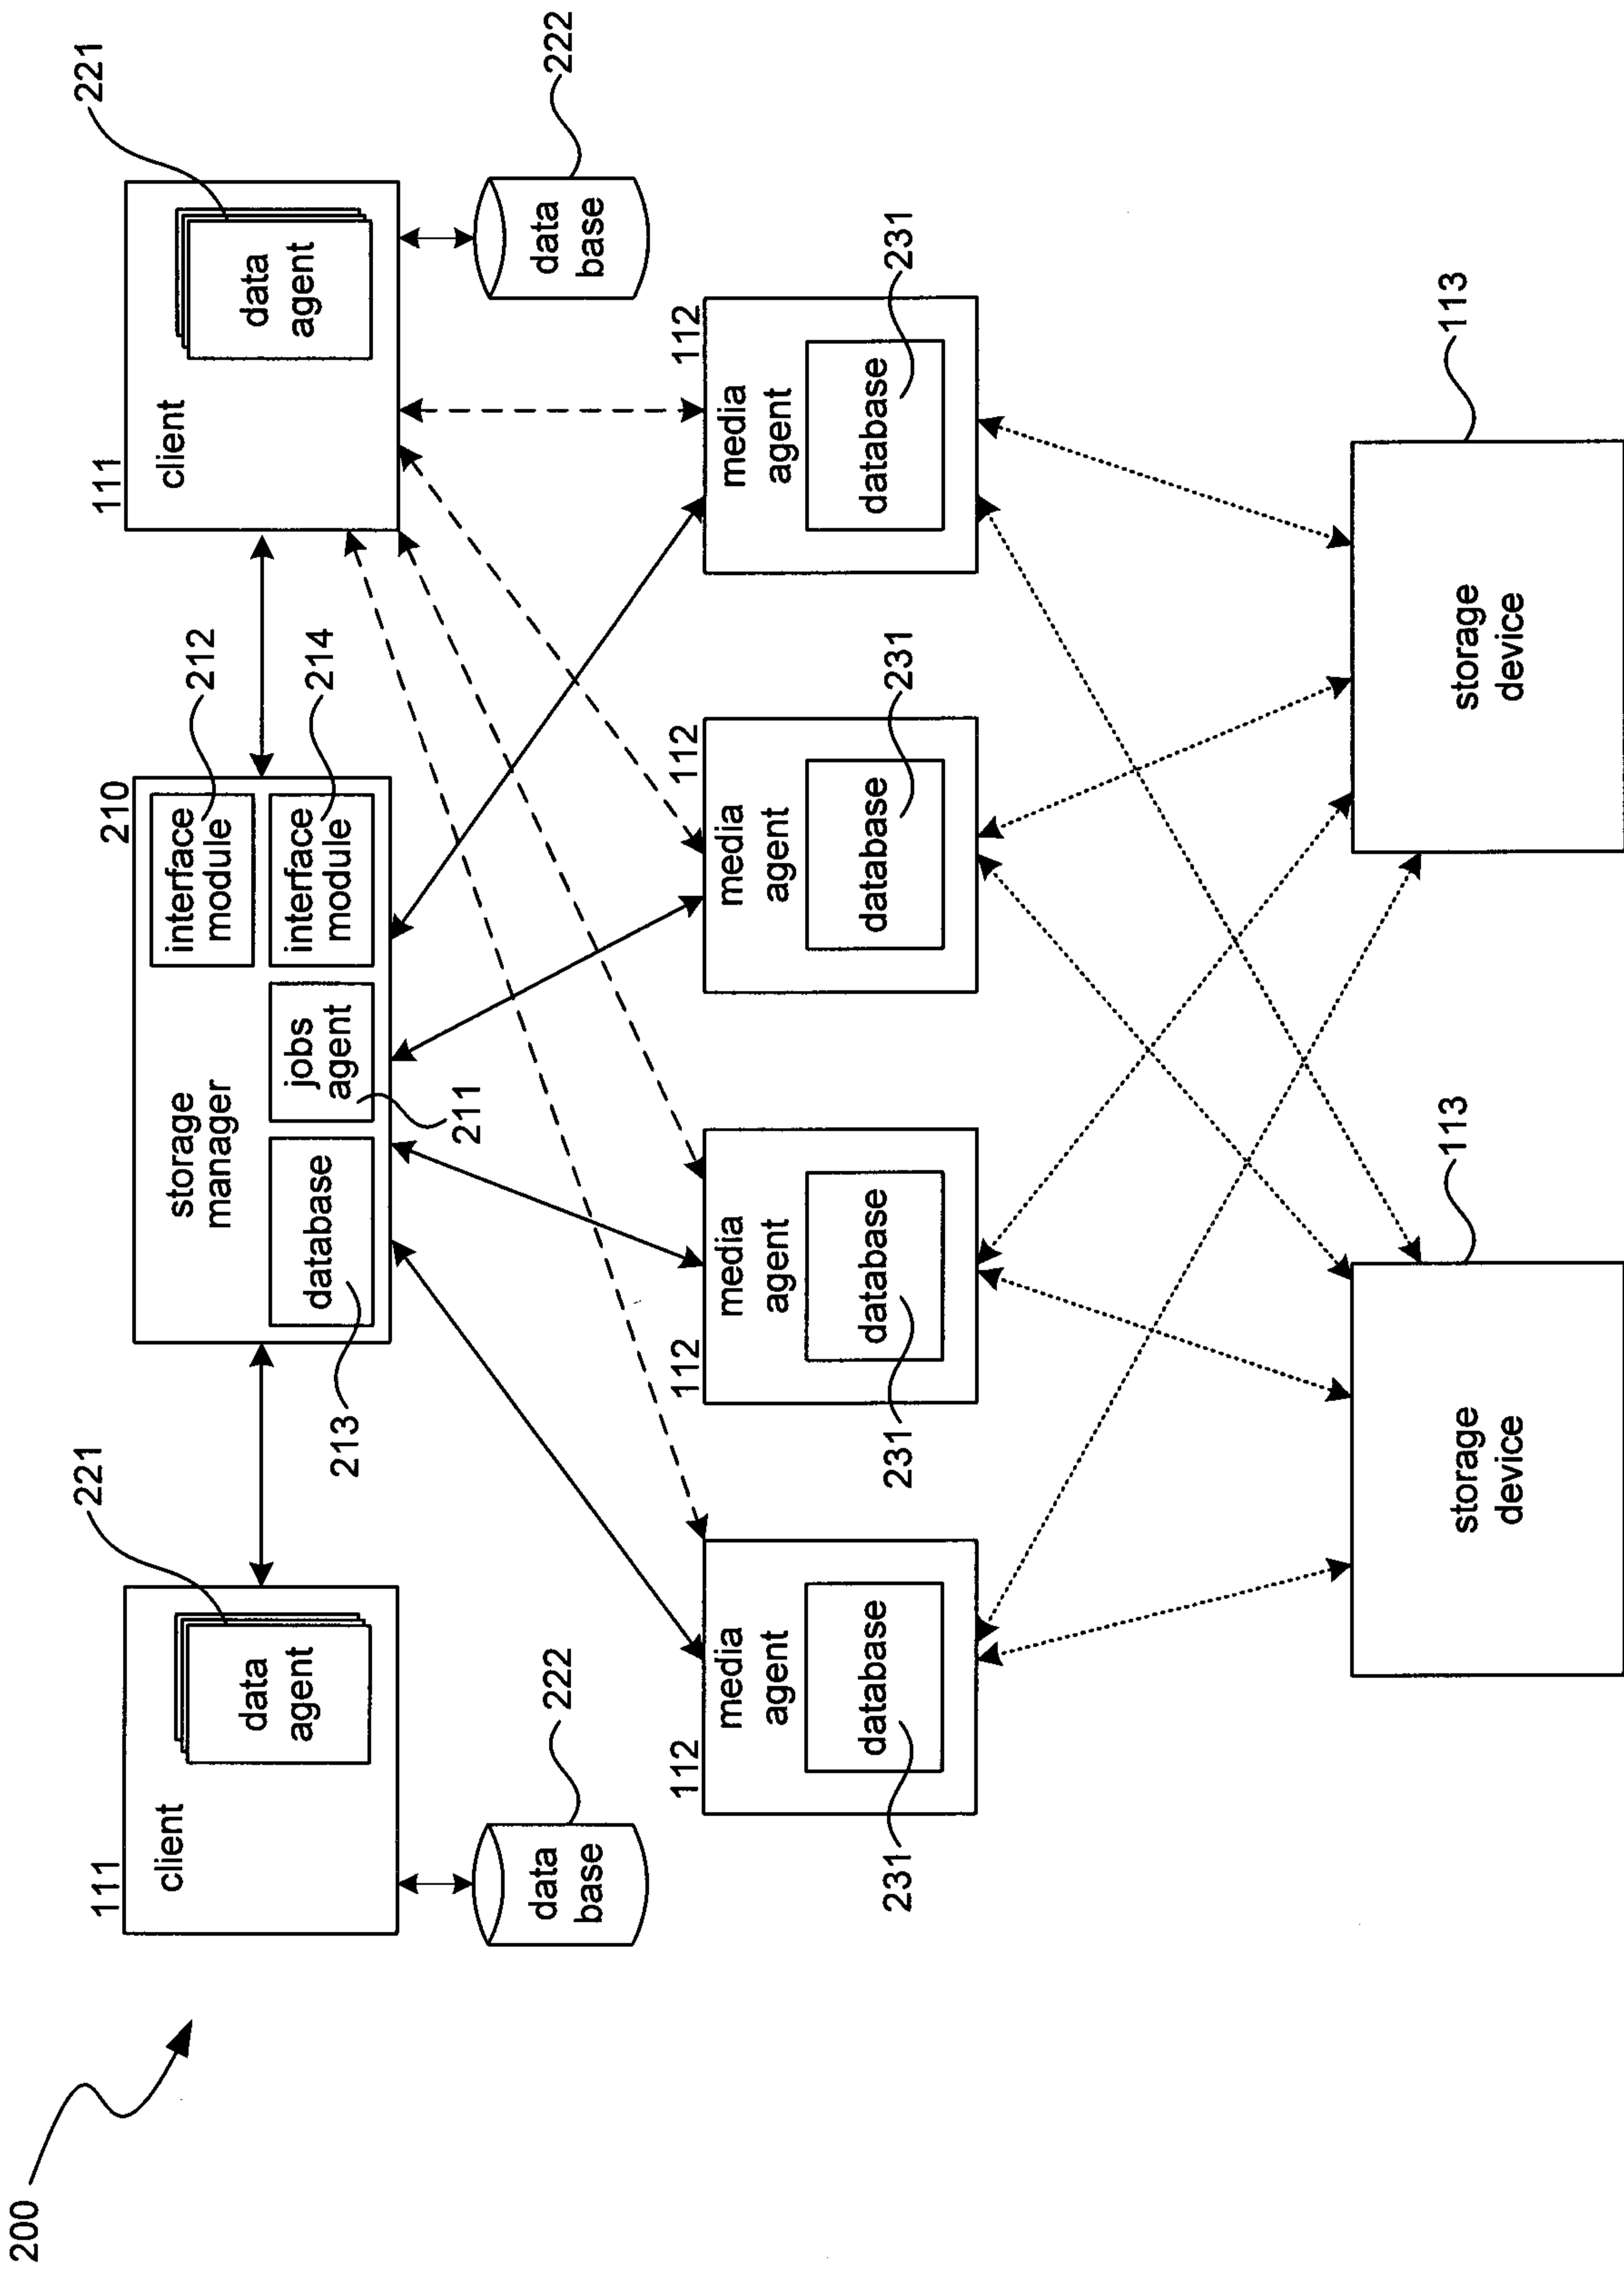


FIG. 2

3/7

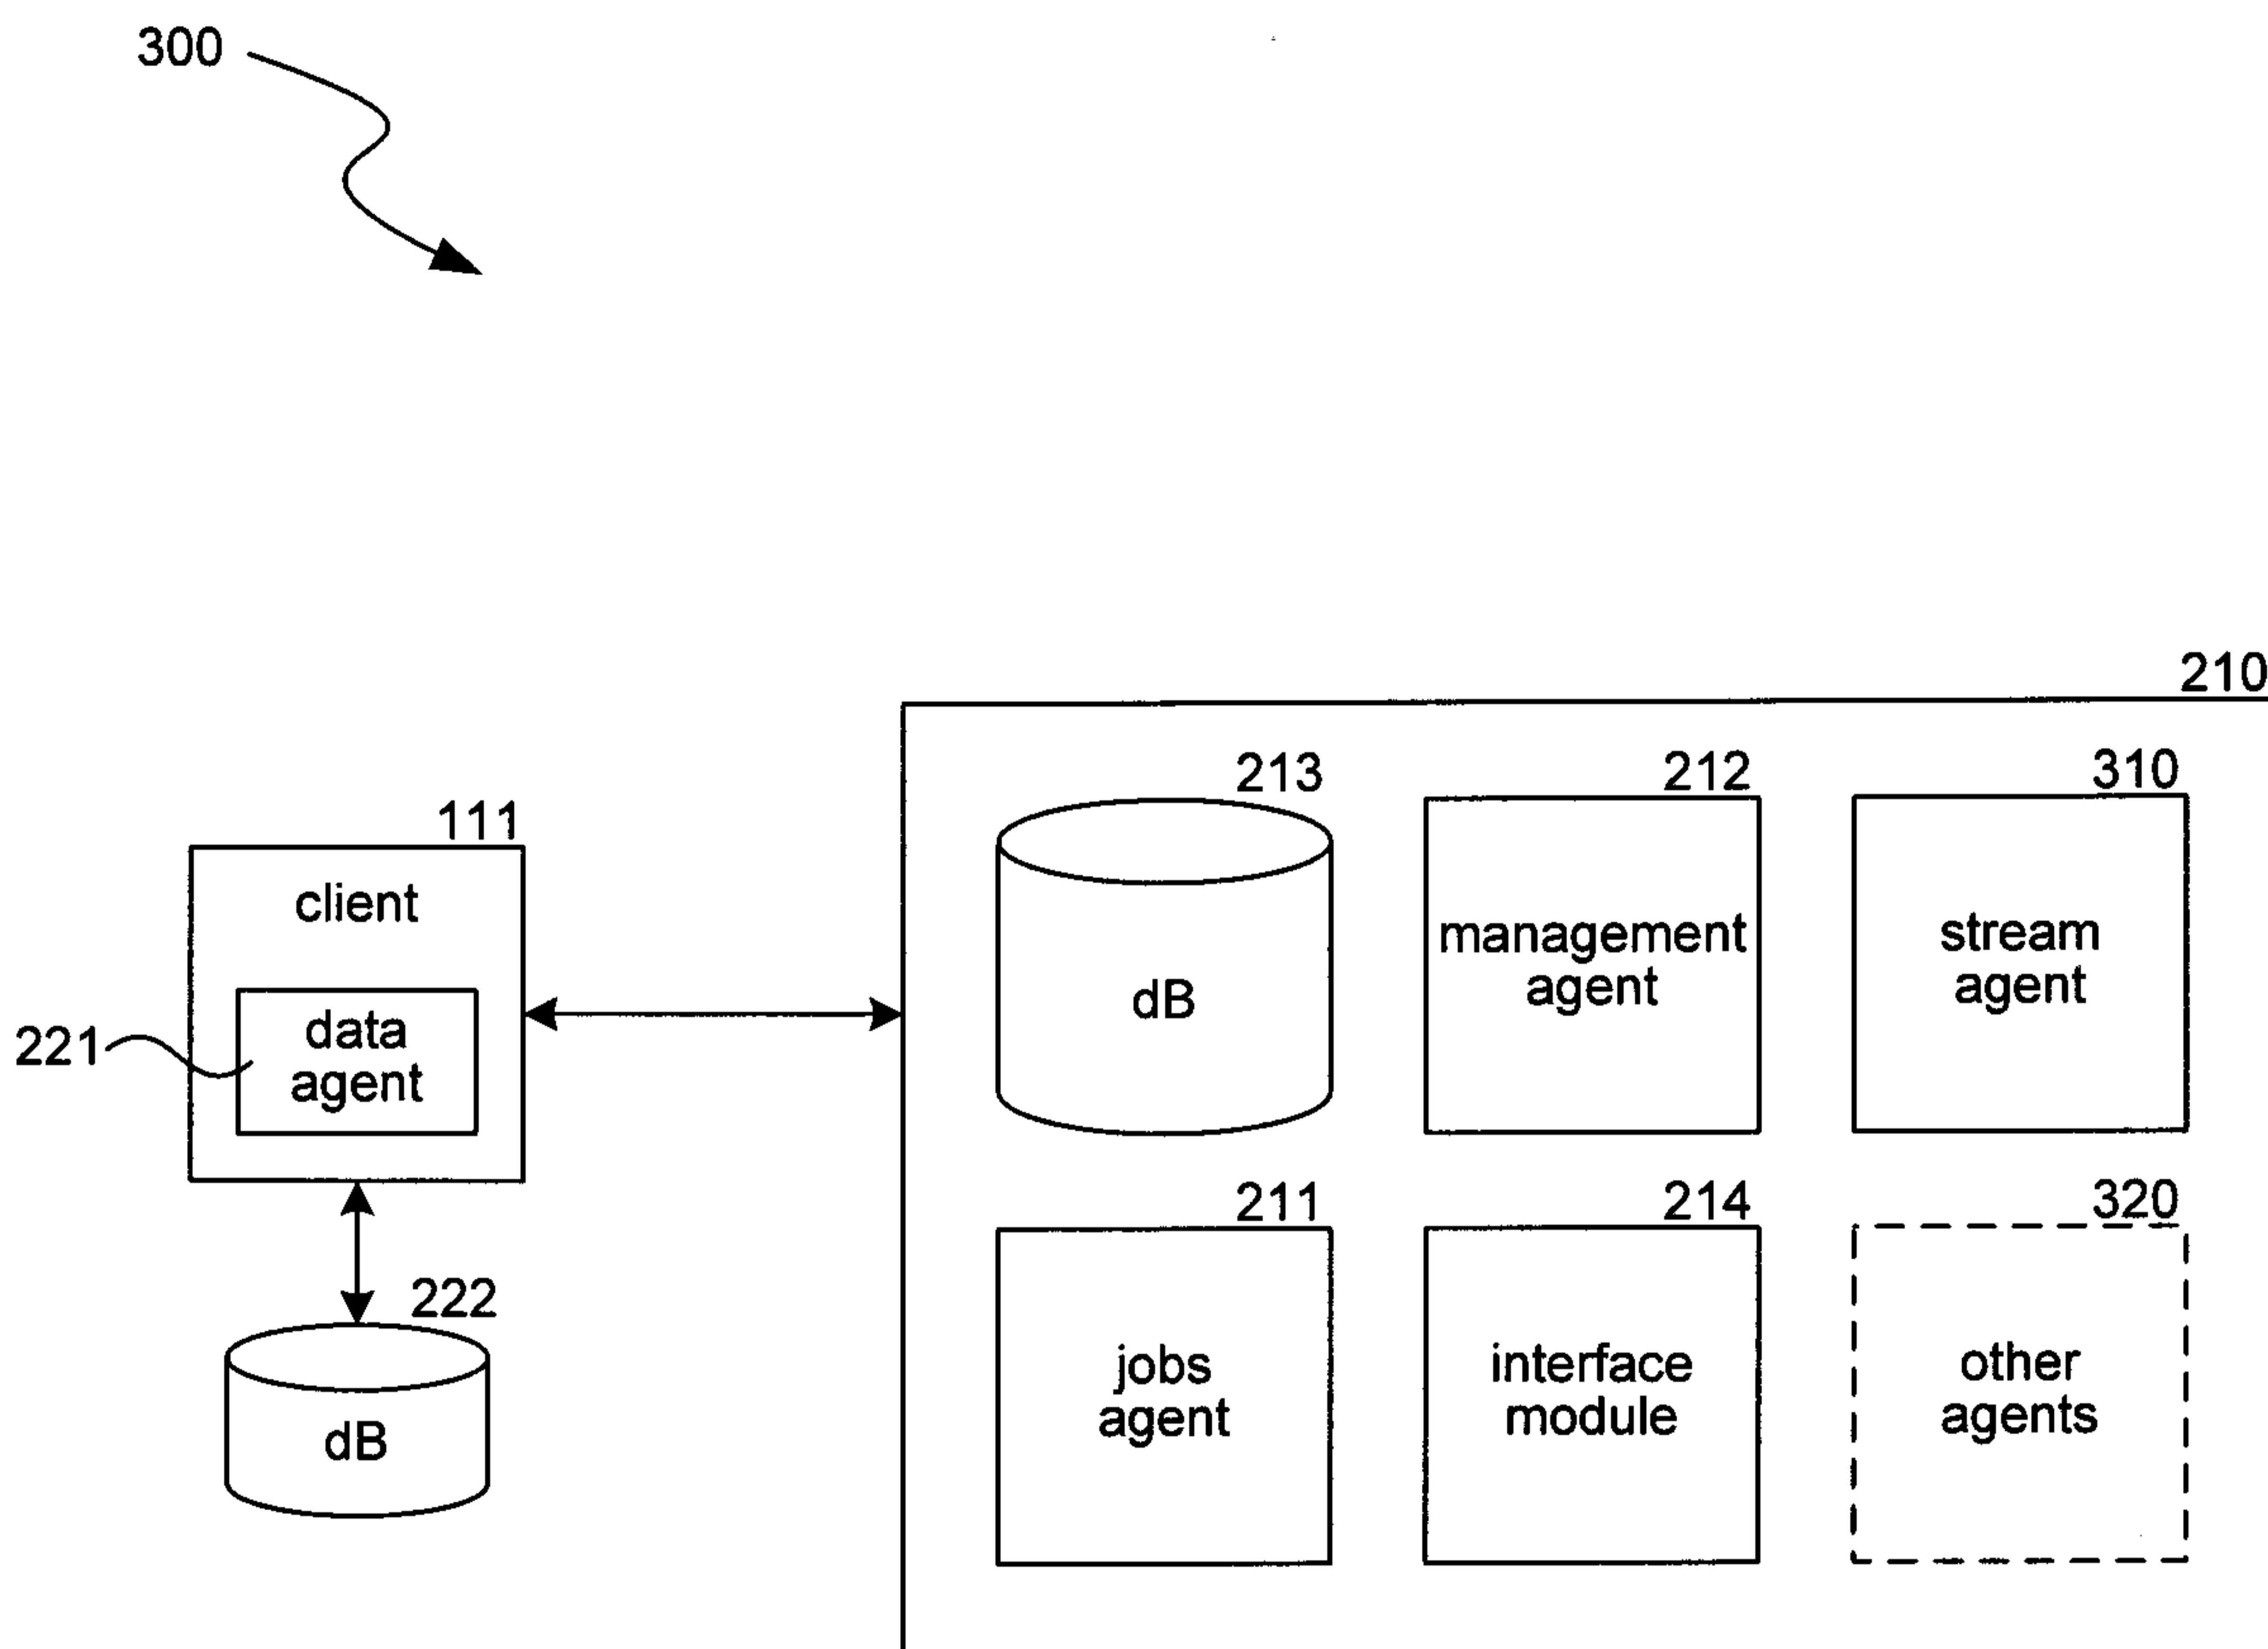
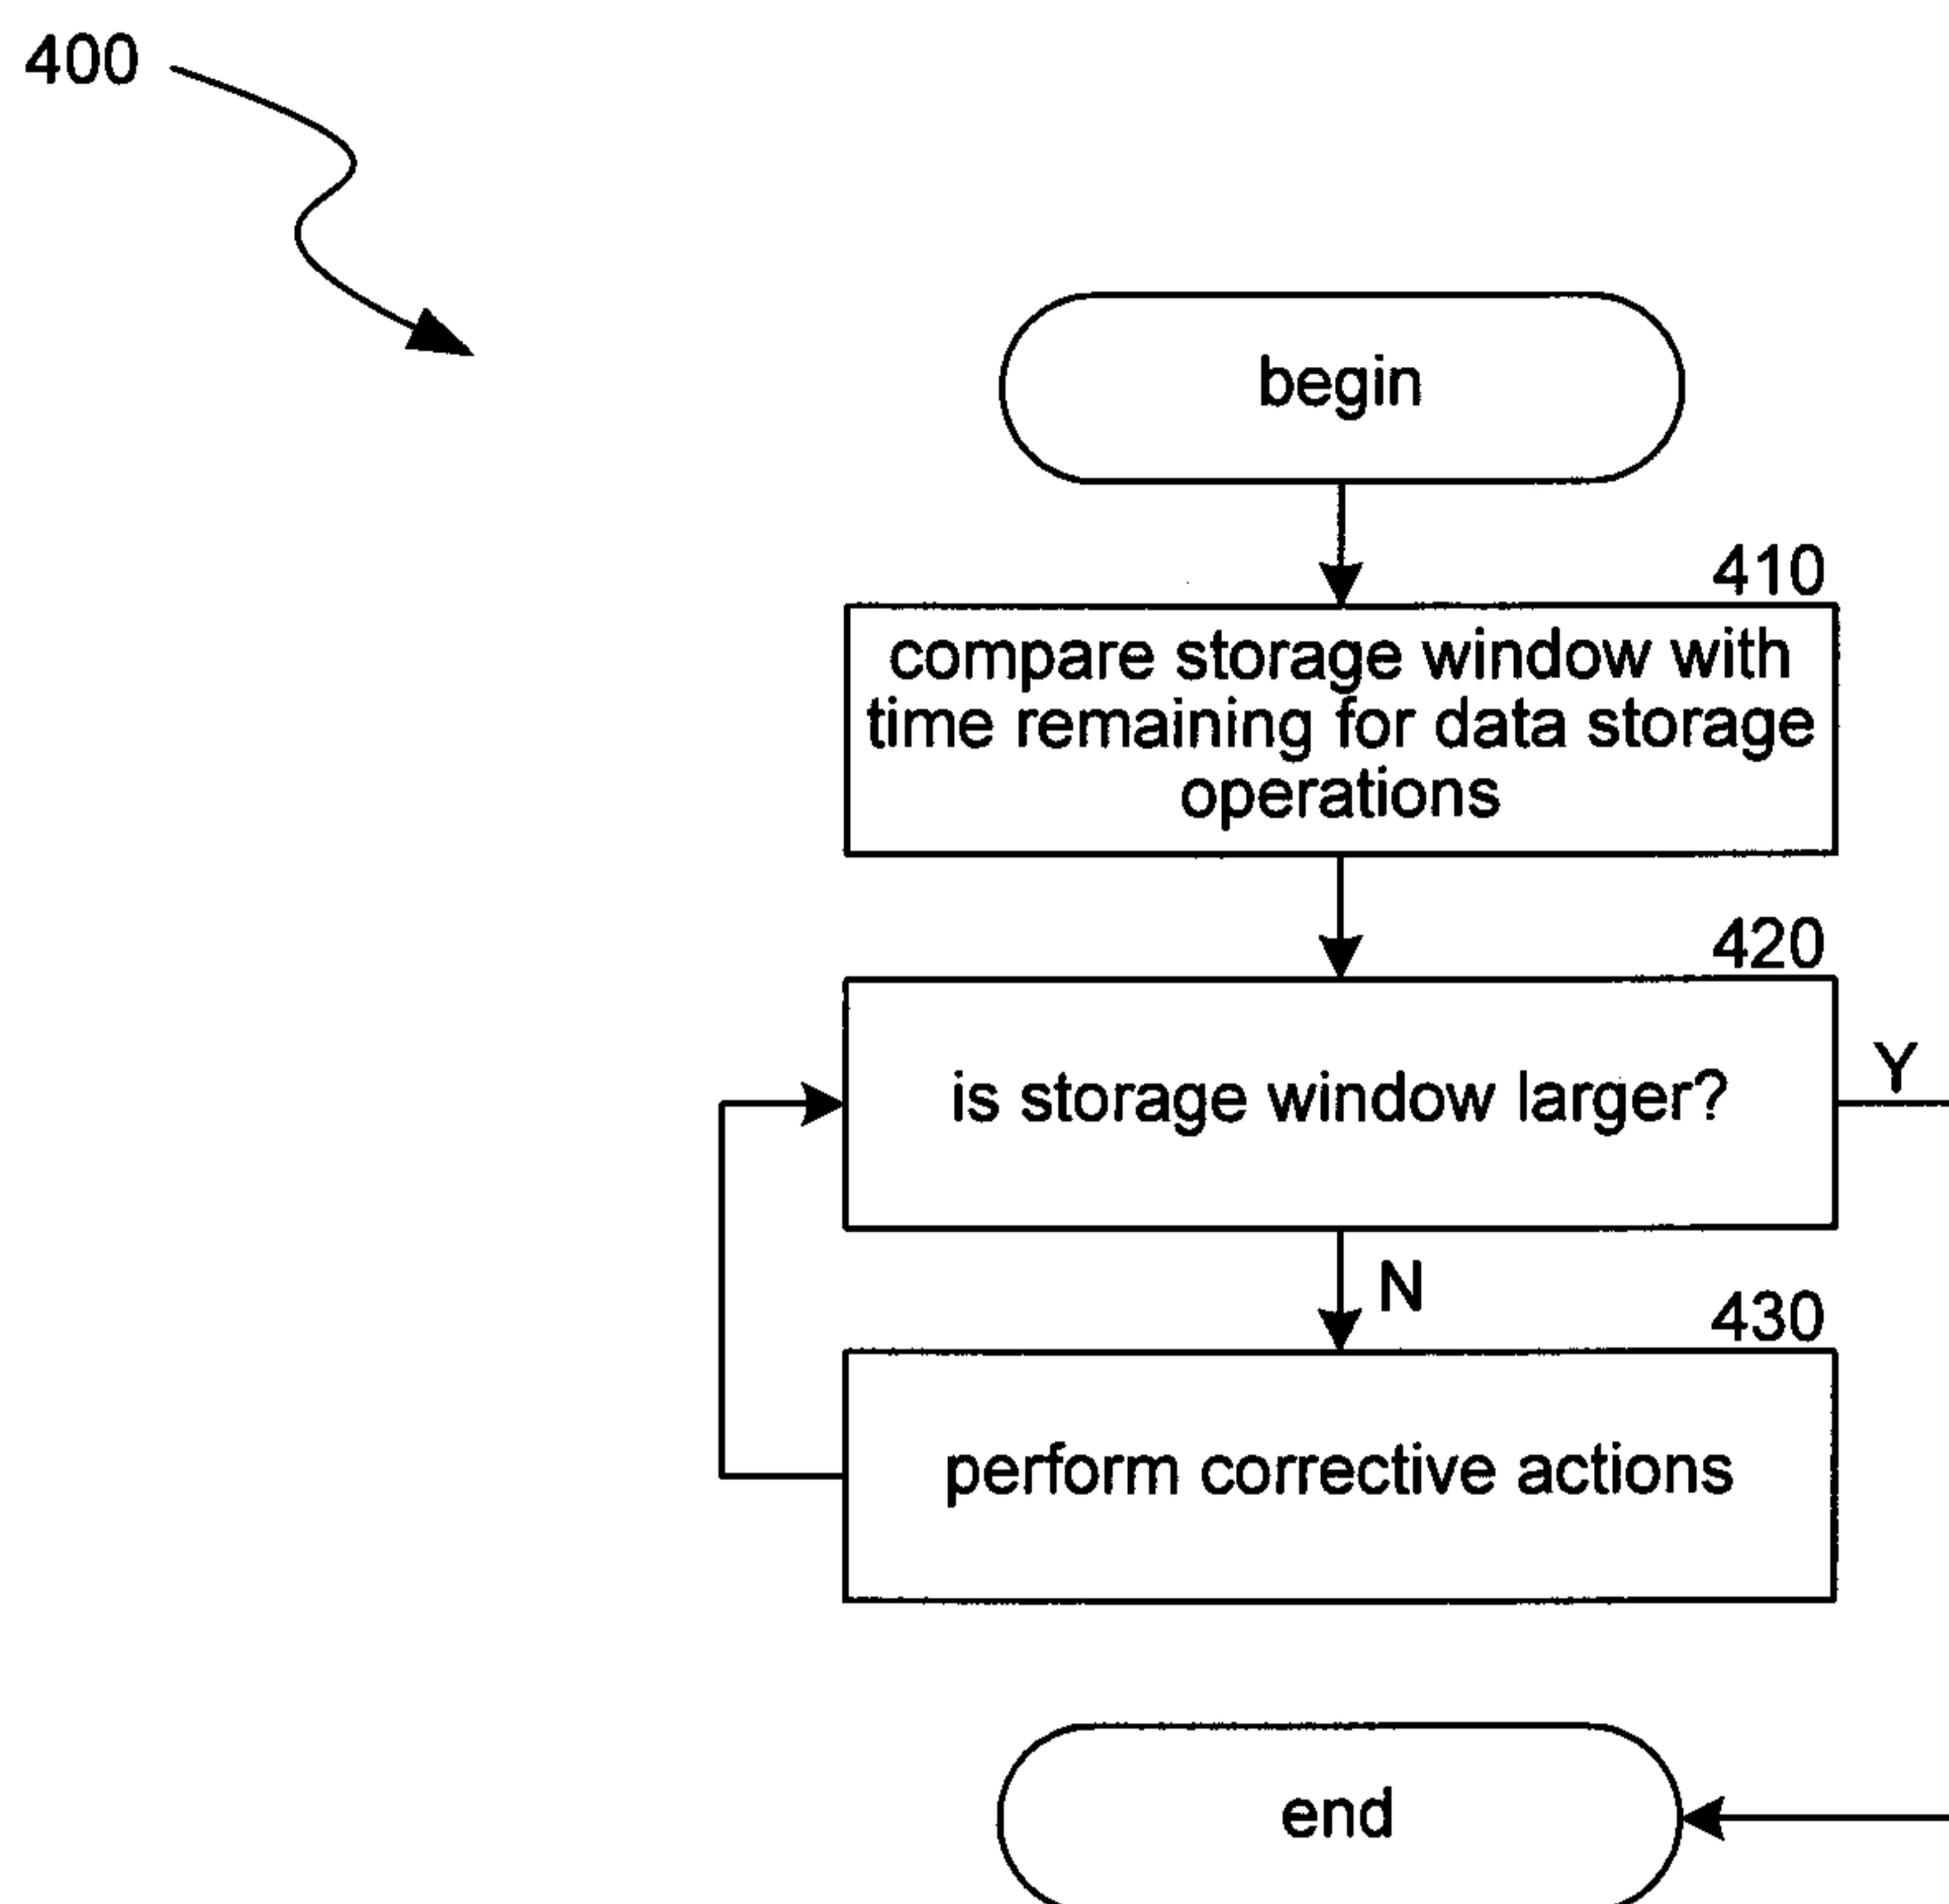
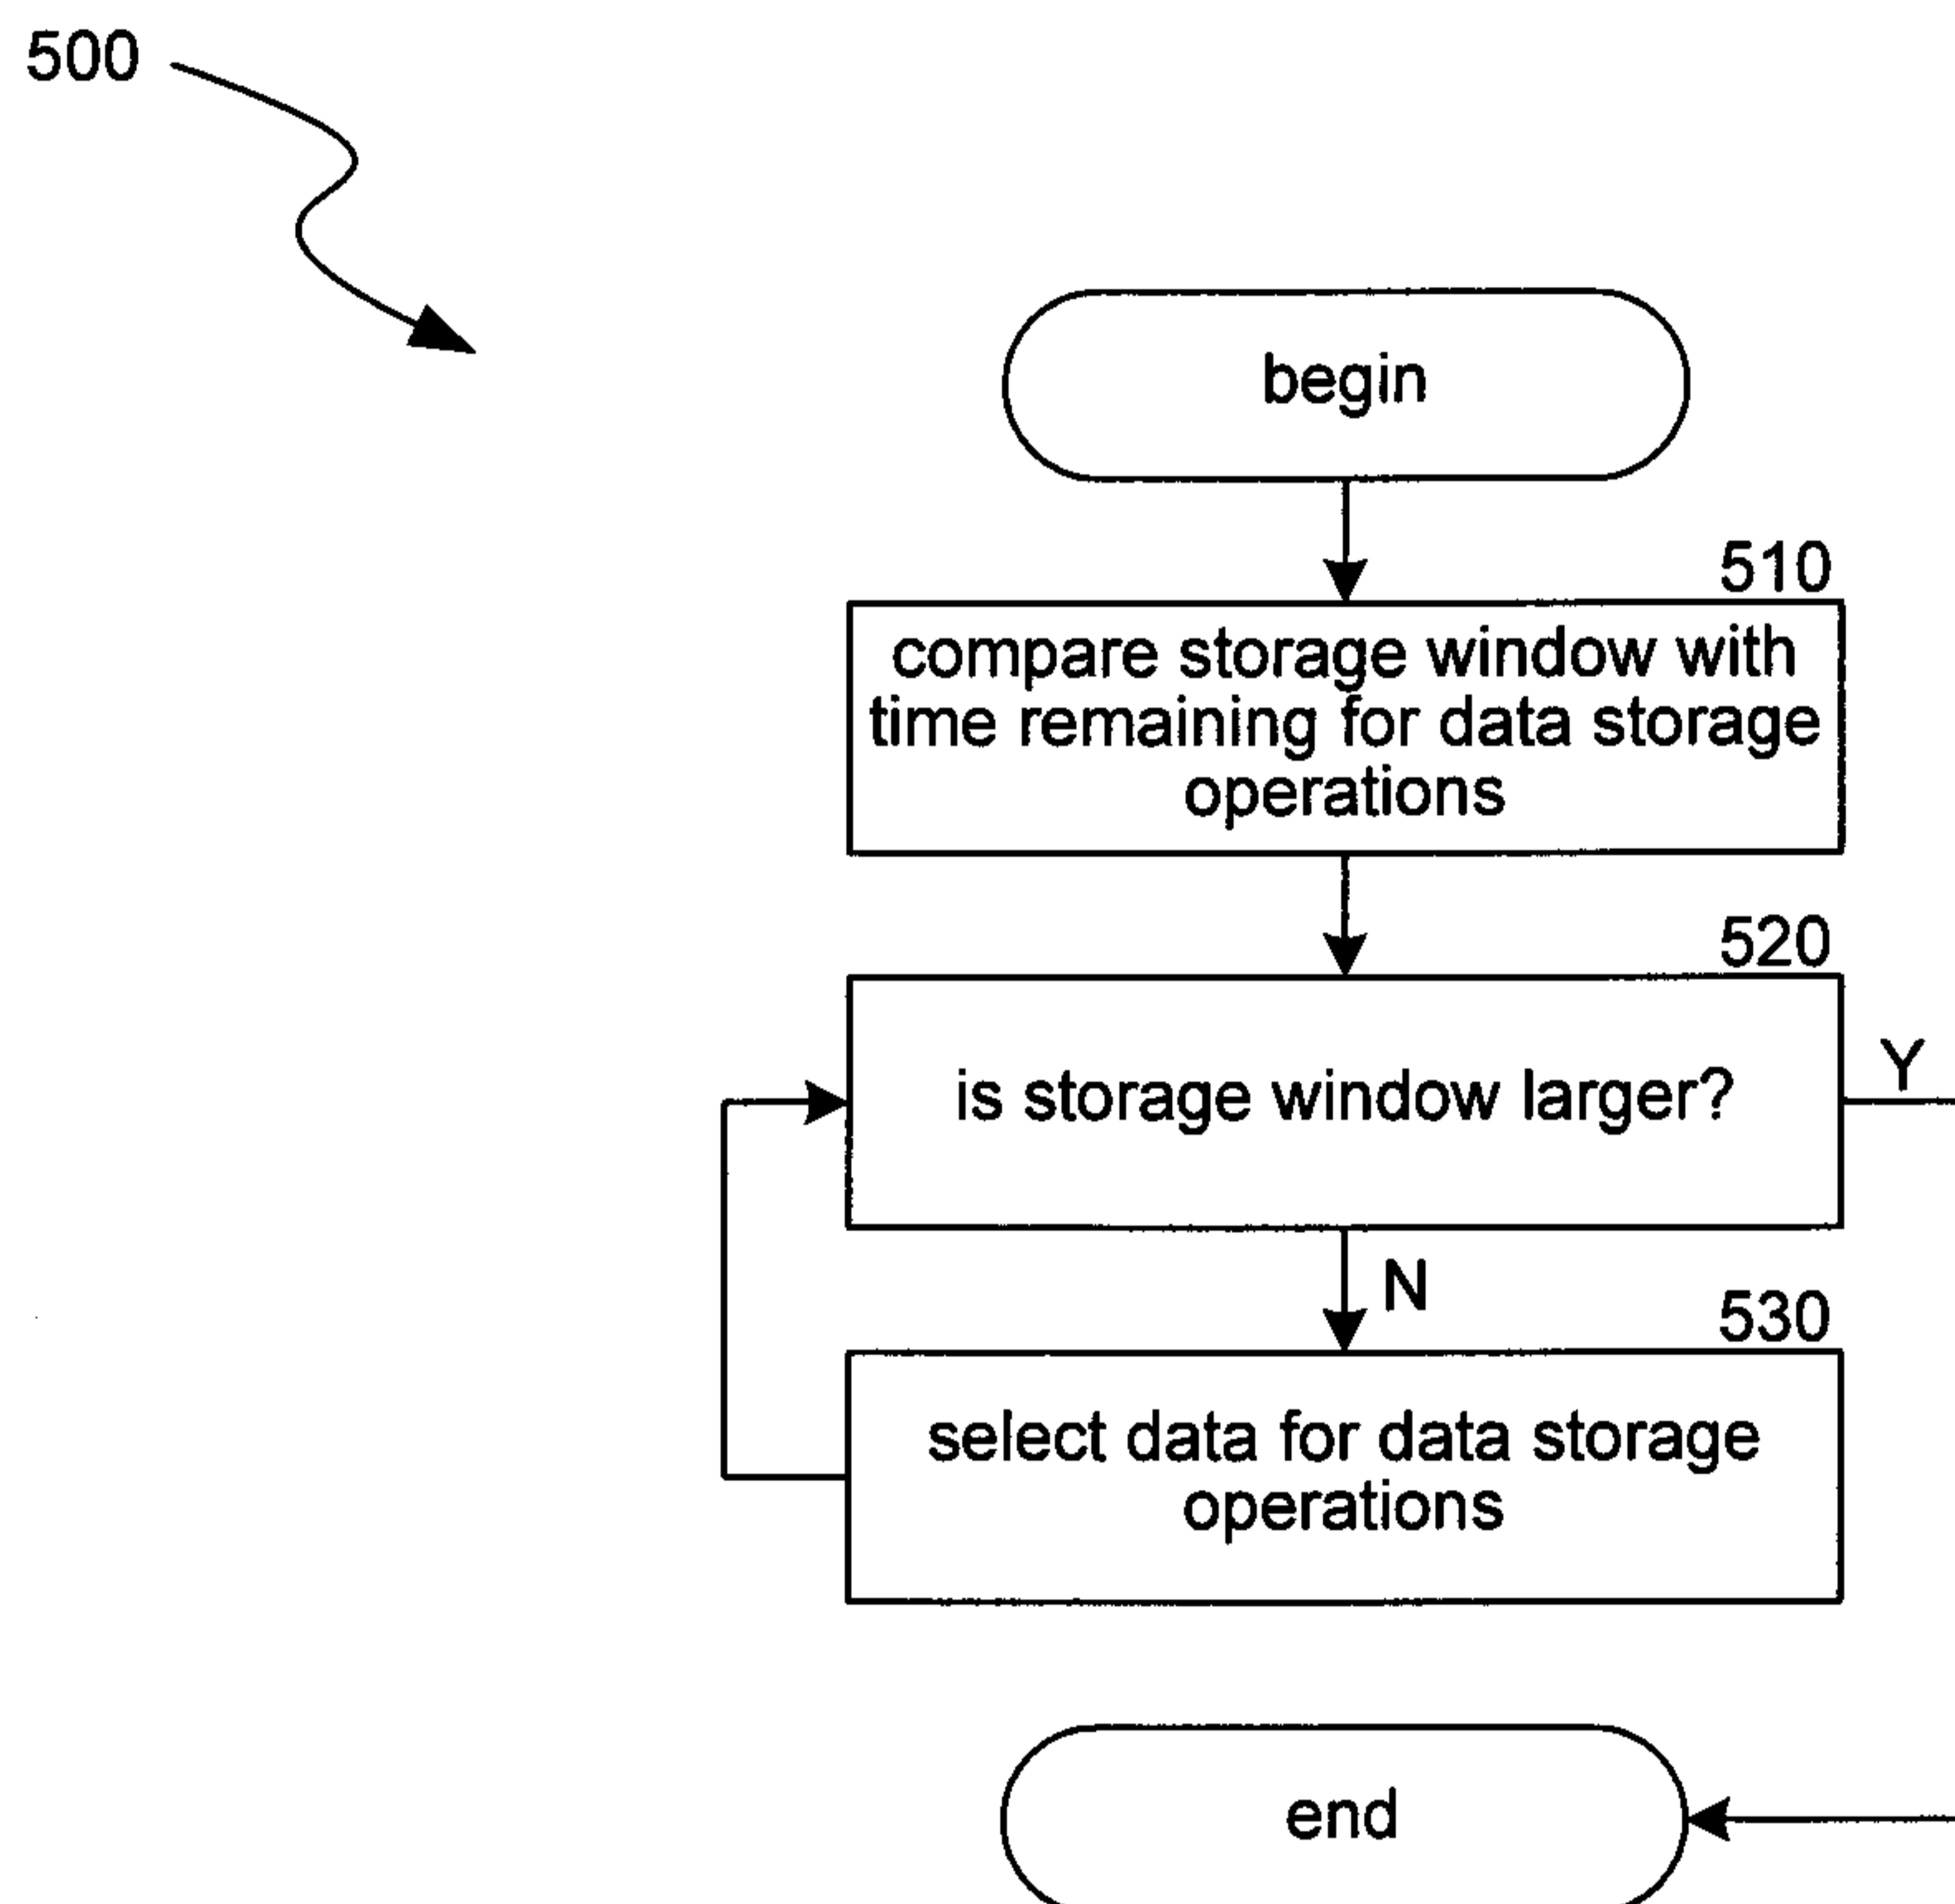


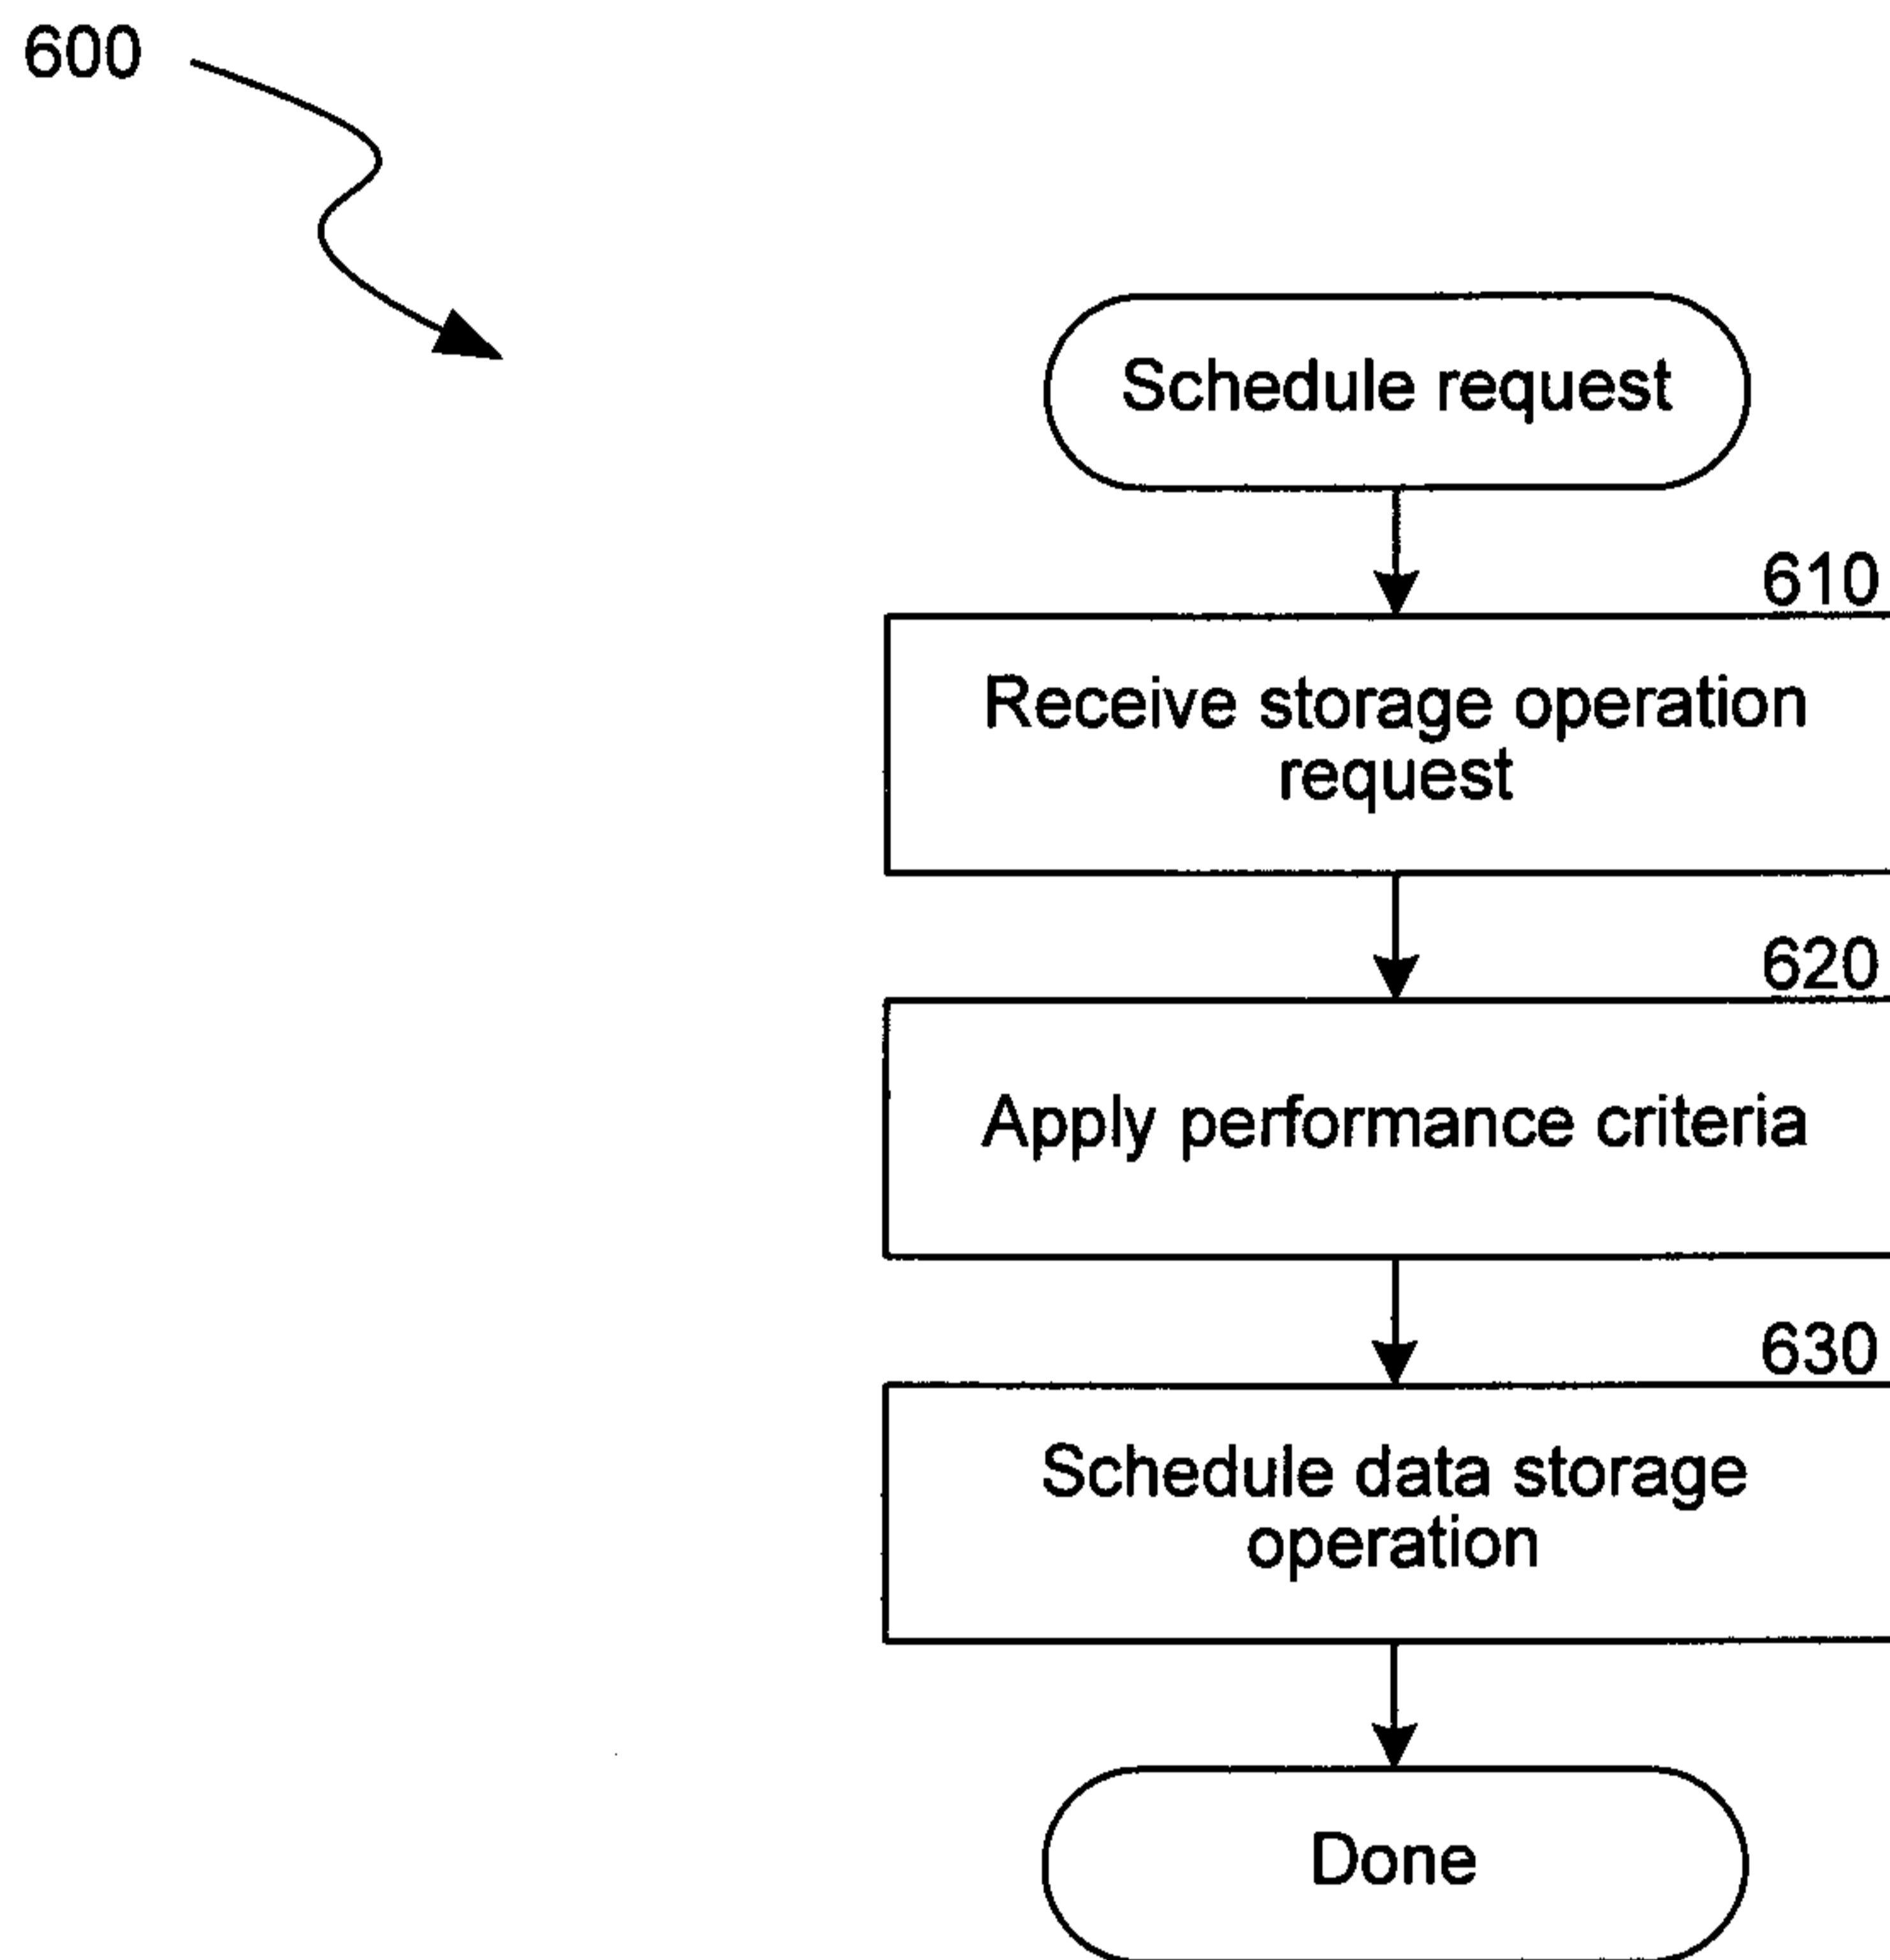
FIG. 3

4/7

**FIG. 4**

**5/7****FIG. 5**

6/7

**FIG. 6**

7/7




FIG. 7



200

