

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
22 February 2007 (22.02.2007)

PCT

(10) International Publication Number
WO 2007/020121 A1

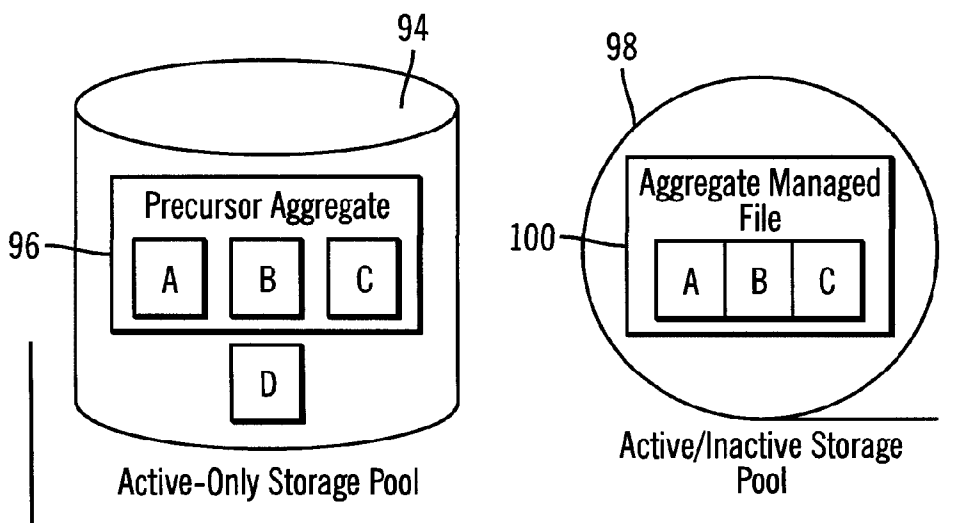
- (51) International Patent Classification:
G06F 17/30 (2006.01)
- (21) International Application Number:
PCT/EP2006/063455
- (22) International Filing Date: 22 June 2006 (22.06.2006)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/206,496 17 August 2005 (17.08.2005) US
- (71) Applicant (for all designated States except US): **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).
- (71) Applicant (for MG only): **IBM UNITED KINGDOM LIMITED** [GB/GB]; PO Box 41, Portsmouth Hampshire PO6 3AU (GB).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **MARTIN, Howard, Newton** [US/US]; 10538 South Peregrine Ridge Court, Vail, Arizona 85641 (US). **CANNON, David, Maxwell** [US/US]; 9791 East Kleindale Road, Tucson, Arizona 85749 (US).

- (74) Agent: **LING, Christopher, John**; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester Hampshire SO21 2JN (GB).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, 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, HU, IE, IS, IT, LT, LU, LV, MC, 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).

Published:
— with international search report

[Continued on next page]

(54) Title: MAINTAINING AN AGGREGATE INCLUDING ACTIVE FILES IN A STORAGE POOL IN A RANDOM ACCESS MEDIUM



(57) Abstract: Provided are a method, system, and program for maintaining an aggregate including active files in a storage pool in a random access medium. An active-only storage pool is configured in a random access device. The active-only storage pool is intended to include only active files and not inactive versions of files. A plurality of files in the active-only storage pool are associated with a precursor aggregate in the active-only storage pool. One file associated with the precursor aggregate in the active only storage pool is deactivated to produce an inactive version of the deactivated file. The precursor aggregate is associated with the inactive version of the file in response to the deactivation. A determination is made of the inactive version of at least one file in the precursor aggregate. The determined inactive version of the at least one file is deleted from the active only storage pool.

WO 2007/020121 A1



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

**MAINTAINING AN AGGREGATE INCLUDING
ACTIVE FILES IN A STORAGE POOL IN A RANDOM ACCESS MEDIUM**

Field of the Invention

5

The present invention relates to a method, system, and program for maintaining an aggregate including active files in a storage pool in a random access device.

10

Background of the Invention

Storage management software maintains a repository for data by backing-up, archiving or migrating data from client nodes in a computer network. A storage management server stores data objects or files in one or more storage pools maintained in storage devices. The storage management server may use a database to track information about the stored files, including their location, status, policy information on the files, etc. The storage management software may also be used to restore data. The storage management software may maintain a hierarchy of storage devices, where a first level of the hierarchy includes the faster access devices, such as disk drives, storing more frequently used and accessed files. Lower levels in the hierarchy provide slower access storage devices, such as lower performance disk drives and tape drives, to store less frequently accessed or inactive data. One example of storage management software is the Tivoli® Storage Manager product marketed by International Business Machines Corporation (IBM®). (Tivoli and IBM are registered trademarks of IBM).

To improve data transfer performance, the storage management software may aggregate individual files in a storage pool, so that moving and copying operations may be performed with respect to an aggregate of files instead of single files. U.S. Pat. No. 6,098,074 describes an aggregation technique in which objects being stored are aggregated into a "managed file." The objects may thereafter be tracked and moved as a single managed file within the storage hierarchy. When appropriate, individual objects can be processed individually such as for deletion or retrieval operations. PCT patent application no. WO 2005/071565 describes further techniques for managing files in aggregates.

As the amount of data users store in storage pools continues to increase, files and aggregates of files are increasingly migrated to lower

5 levels in the storage hierarchy. Thus, data that may need to be restored may have been migrated to a slower access device, such as a tape storage media, which has a low restore performance. Further, to restore data from tape, the data may first be staged from tape to a hard disk drive pool and then restored from the hard disk drive.

Further improvements in data storage may be useful in a variety of applications.

10 DISCLOSURE OF THE INVENTION

15 Provided are a method, system, and program for maintaining an aggregate including active files in a storage pool in a random access medium. An active-only storage pool is configured in a random access device. The active-only storage pool is intended to include only active files and not inactive versions of files. A plurality of files in the active-only storage pool are associated with a precursor aggregate in the active-only storage pool. One file associated with the precursor aggregate in the active only storage pool is deactivated to produce an inactive version of the deactivated file. The precursor aggregate is associated with the inactive version of the file in response to the deactivation. A determination is made of the inactive version of at least one file in the precursor aggregate. The determined inactive version of the at least one file is deleted from the active only storage pool.

25 In a further embodiment, the file is deactivated as a result of an update to the file. Updating the file produces an active version of the file in addition to the inactive version of the file. Indication of the association of the deleted inactive version of the at least one file with the precursor aggregate is removed. The precursor aggregate is associated with only the active version of files in response to removing the indication.

35 In a further embodiment, the files associated with the precursor aggregate are copied to an aggregate managed file in an additional storage pool before deleting the determined inactive version of the at least one file. The aggregate managed file maintains the inactive version of the at least one file deleted from the active only storage pool.

In a further embodiment, the files are written sequentially to the aggregate managed file and the aggregate managed file comprises a sequential file.

5 In a further embodiment, the additional storage pool is implemented in a sequential access device.

10 In a further embodiment, the precursor aggregate has a first identifier and the aggregate managed file has a second identifier. A data structure is provided having information on defined aggregates including an aggregate identifier and storage pool including the aggregate for each indicated aggregate. Indication is made in the information for the aggregate identified by one aggregate identifier whether the aggregate is a precursor aggregate comprising an association of files or an aggregate managed file in which files are written.

15 In a further embodiment, indicating in the data structure information that the aggregate is the precursor aggregate comprises indicating that the precursor aggregate has an aggregate size of zero. 20 Indicating in the information that the aggregate is the aggregate managed file comprises indicating that the aggregate managed file has an aggregate size that is a cumulative size of the files written in the aggregate managed file.

25 In a further embodiment, the information for each aggregate in the data structure includes a logical size and an actual size. Indicating that the precursor aggregate has the aggregate size of zero comprises indicating that the actual size of the precursor aggregate is zero and indicating that the aggregate managed file has the aggregate size that is the cumulative size comprises indicating that the actual size of the 30 aggregate managed file is the cumulative size. Indication is made in the information for the precursor aggregate and the aggregate managed file that the logical size is the cumulative size of the files associated with the precursor aggregate and written to the aggregate managed file, 35 respectively.

In a further embodiment, the file is deactivated as a result of an update to the file. Updating the file produces an active version of the file in addition to the inactive version of the file. Information is 40 added to the data structure for the updated file having a third identifier and indicating that the updated file is in the active-only storage pool.

In a further embodiment, the data structure comprises a first data structure. A second data structure has information on each file grouped in one aggregate including the aggregate identifier of the aggregate including the file. The aggregate identifier for the files associated with the precursor aggregate comprises the first identifier and the aggregate identifier for the files written to the aggregate managed file comprises the second identifier.

In a further embodiment, the information is deleted in the second data structure for the inactive version of the at least one file having the first identifier that was deleted from the active storage pool. The second data structure includes information on the inactive version of the at least one file having the second identifier maintained in the aggregate managed file.

In a further embodiment, the information for the files in the second data structure indicates an offset and file size of the file. Files associated with one precursor aggregate have an offset of zero and files included in one aggregate managed file have an offset at which the file is stored in the aggregate managed file.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

FIG. 1 illustrates an embodiment of a network computing environment;

FIGS. 2, 3, and 4 illustrate embodiments of information on files and aggregates of files maintained in storage pools;

FIG. 5 illustrates an embodiment of an active-only storage pool;

FIG. 6 illustrates an embodiment of operations to create a precursor aggregate in an active-only storage pool;

FIG. 7 illustrates an embodiment of operations to copy files in a precursor aggregate in the active-only storage pool;

FIG. 8 illustrates an embodiment of operations to deactivate a file in a precursor aggregate in an active-only storage pool;

FIG. 9 illustrates an embodiment of operations to expire inactive files in the active-only storage pool; and

5 FIGs. 10 and 11 illustrate examples of an active-only storage pool, an active/inactive storage pool, and tables of information on the files and aggregates in the active-only and active/inactive storage pools.

DETAILED DESCRIPTION

10 FIG. 1 illustrates a computing environment in which embodiments are implemented. A plurality of clients 2a, 2b...2n, storage devices 4a, 4b, and a storage management server 6 are in communication over a network 8. The storage management server 6 may access storage devices 4a, 4b over the network 8. Alternatively, a storage device may be attached directly to
15 the storage management server 6 and accessed over a bus. The clients 2a, 2b...2n include client storage agents 10a, 10b...10n and the storage management server 6 includes storage management software 12. The client storage agents 10a, 10b...10n and the storage management software 12 interact to manage the storage of files in storage pools 14a, 14b
20 implemented in the storage devices 4a, 4b. The storage management software 12 may maintain a storage hierarchy system in which storage pools 14a, 14b are defined to have level information, such that files at a higher level may be migrated to a lower level according to a hierarchical storage policy, e.g., to move less frequently accessed files to storage
25 pools at a lower level in the hierarchy having slower access devices. The storage in the higher level may comprise faster access devices, such as hard disk drives, whereas the storage at lower levels may comprise slower access devices, such as slower hard disk drives, tape drives, etc. For instance, storage pools high in the hierarchy include more frequently or
30 recently used data and are implemented in relatively faster access storage devices, whereas storage pools lower in the hierarchy may provide more long term storage of data and be implemented in a tape storage medium. In certain embodiments, the client storage agents 10a, 10b...10n may access, backup, archive and restore data through the storage management software
35 12, which manages access to files in the storage pools 14a, 14b. Further, the storage management software 12 may backup and archive data from the clients 2a, 2b, 2n.

40 The storage management software 12 maintains information on the files in the storage pools, including file information 16, aggregate information 18, and storage pool information 20. The information 16, 18,

and 20 may be implemented in one or more database tables of a relational database or other suitable data structures known in the art. The file information 16 may comprise an inventory table having information on every file in the storage pools 14a, 14b, including client and policy information. The aggregate information 18 comprises information on aggregates defined in the storage pools 14a, 14b. A precursor aggregate comprises an association of files in the storage pools 14a, 14b, an aggregate managed file comprises a file or object in one storage pool in which one or more files are written and associated. Files grouped by a precursor aggregate are not stored in an aggregate managed file, but remain stored as separate files in the file system. An aggregate simplifies file movement operations (e.g., storage pool backup, restore, reclamation or movement to another pool which includes migration to another location within the hierarchy) because the storage management software 12 need only specify a data transfer operation with respect to an aggregate (precursor or aggregate managed file), and the storage management software 12 will then perform the requested operation with respect to the files grouped by the aggregate.

The storage pool information 20 contains information about where each managed file is stored in the storage hierarchy implemented in the storage pools 14a, 14b. The storage table contains an entry for each managed file.

A storage pool 14a, 14b may be implemented in a random access device or as a sequential access storage pool. A sequential access storage pool may be implemented in a sequential access media, e.g., tape, or in a sequential file volume in a random access media. When data is stored sequentially in a sequential file volume in a random access device, the volume is a file in the underlying file system. Space from a deleted file in a sequential-access disk pool is recovered by consolidating valid data on a new file volume during a reclamation operation. In a random-access disk pool, space is allocated in random blocks, such that once a file is deleted from a random-access storage pool, the space is immediately available for use.

In one embodiment, a storage pool may be defined as an active only storage pool 22, such that only active files are maintained in that storage pool 22, not inactive files. An inactive file is any file that has been deactivated according to some criteria. For instance, a file may be considered inactive if it is updated, deleted or is a file whose age

exceeds an aging policy. A file may be updated or deleted on a client system 2a, 2b...2n, and then when that update is supplied to the storage management server 6, the deactivated file may be marked as inactive. A file may also be deactivated directly by the storage management software 12. An update to a file results in both an active version of the file having the update and an inactive version of the pre-updated file.

The clients 2a, 2b...2n may comprise a suitable computational device known in the art, such as a workstation, desktop computer, server, mainframe, hand held computer, telephony device, etc. The storage management server 6 may comprise a suitable server class machine. The network 8 may comprise a Local Area Network (LAN), Storage Area Network (SAN), Wide Area Network (WAN), wireless network, etc. The storage devices 4a, 4b may be implemented in storage devices known in the art, such as one hard disk drive, a plurality of interconnected hard disk drives configured as Direct Access Storage Device (DASD), Redundant Array of Independent Disks (RAID), Just a Bunch of Disks (JBOD), etc., a tape device, an optical disk device, a non-volatile electronic memory device (e.g., Flash Disk), a tape library, etc. The storage devices 4a, 4b may comprise random access devices, such as hard disk drives, electronic memory or storage, etc., where data may be written randomly to blocks in the storage device or sequential access devices where data must be written sequentially to the storage media, such as tape media.

FIG. 2 illustrates an embodiment of file/aggregate information 50 maintained for every aggregate managed file and file outside of an aggregate managed file within a storage pool. The file/aggregate information 50 may be part of the aggregate information 18. Each file/aggregate information 50 instance includes an identifier (ID) 52 of the file or aggregate managed file; a size 54 (e.g., byte length) of the file or aggregate, where an aggregate size is the size of all files included in the aggregate managed file; and a storage pool 56 identifying the storage pool 14a, 14b, 22 including the file or aggregate managed file. A precursor aggregate may have a size 54 of zero, indicating it is a logical entity and an aggregate managed file may have a size 54 including its actual size, which includes the cumulative size of all files written to the aggregate managed file. For entries 50 for files, as opposed to aggregates, an active status field 58 indicates whether the file is active or inactive.

FIG. 3 illustrates an embodiment of aggregate information 70 maintained for every aggregate. The aggregate information 70 may be part of the aggregate information 18. Each aggregate information 70 instance includes: an identifier 72 of the aggregate; an actual size 74 of the aggregate managed file or a value indicating that the aggregate is a precursor aggregate, such as zero (a precursor aggregate may not have an actual size because it comprises an association of files and is not an object storing files); a logical size 76 including the size of all files associated with that aggregate; and a number of files 78 grouped by the aggregate (precursor or managed file). An aggregate managed file may have a logical size 76 equal to the cumulative size of the files written to the aggregate managed file and a precursor aggregate has a logical size of the files associated with the precursor aggregate.

FIG. 4 illustrates an embodiment of aggregate file information 80 maintained for every file associated with an aggregate, i.e., a precursor aggregate or included in one aggregate managed file. The aggregate file information 80 may be part of the file information 16 or the aggregate information 18. Each aggregate file information 80 instance includes: a file identifier 82, such as the file name and location; an aggregate identifier 84 indicating the aggregate (precursor or managed file) in which the file is grouped; an offset 86 indicating the byte offset from the beginning of the aggregate managed file at which the file starts; a length 88 of the file; and an active status 90 indicating whether the file is active or inactive. A file associated with a precursor aggregate may have an offset 86 of zero or some other value indicating that the file is associated with a precursor aggregate. A file associated with a precursor aggregate may have a length 88 equal to the byte length of the file.

FIG. 5 illustrates an active-only storage pool 94 implemented as a random access device in one of the storage devices 4a, 4b including one precursor aggregate 96 having three files A, B, C. The files A, B, C may be stored at non-sequential or sequential locations in the storage pool 94, and are shown as stored at non-sequential random locations. The active-only storage pool 94 may include additional precursor aggregates or aggregate managed files as well as files not included in an aggregate, such as large files that do not get aggregated. A file D is not associated with the precursor aggregate 96.

FIG. 5 further illustrates an active/inactive storage pool 98 in which data may be written sequentially and an aggregate managed file 100

in which files A, B, C are written sequentially. An active/inactive storage pool comprises either a copy storage pool or a primary storage pool. A copy storage pool is a secondary storage pool for receiving a copy of data from an active-only storage pool or maintaining a backup copy for disaster recovery. A primary storage pool contains both active and inactive files, and inactive files may be restored from the primary storage pool. The active/inactive storage pool 98 may be implemented in a sequential access media or in a sequential file volume in a random access media in which data is written sequentially. A sequential file volume comprises a file on a random access media, such as a hard disk drive, that is managed as a tape device where data is written sequentially. A storage pool can have multiple sequential file volumes including aggregate managed files. When the sequential file volume reaches a predefined size, then the sequential file volume is closed and a new sequential file volume is added to the storage pool to fill-up.

FIG. 6 illustrates operations performed by the storage management software 12 to create a precursor aggregate 96 in the active-only storage pool 94. The invocation of the storage management software 12 (at block 100) to create a precursor aggregate may be initiated by one client storage agent 10a, 10b...10n or an administrator at the storage management server 6. A file/aggregate information 50 (FIG. 2) instance is added (at block 102) to the file information 16 identifying the precursor aggregate 52 being created, a size 54 indicating a precursor aggregate (e.g., 0), and an identifier in field 56 of the active-only storage pool 94 in which the precursor aggregate is created. For each file to associate with a precursor aggregate, an aggregate file information instance 80 (FIG. 4) is added (at block 104) to the file 16 or aggregate 18 information identifying the file 82 associated with the precursor aggregate, the precursor aggregate ID 84, an offset 86 of zero, a length 88 of the file, and a status 90 of the file being associated. The storage management software 12 further adds (at block 106) an aggregate information instance 70 (FIG. 3) identifying the created precursor aggregate ID 72, actual size 74 indicating a precursor aggregate (e.g., 0), a logical size 76 of the size of all files associated with the precursor aggregate and a number of files 78 associated with the precursor aggregate. When associating a new file with the precursor aggregate, the logical size 76 and number of files 78 would be updated to reflect the added file.

FIG. 7 illustrates an embodiment of operations performed by the storage management software 12 to copy a precursor aggregate 96 from the

active-only storage pool 94 to a target aggregate managed file 100 in an active/inactive storage pool 98, which may comprise a sequential access device or sequential file. The operations of FIG. 7 may further apply to the copying of files not included in an aggregate, such as file D (FIG. 5). Upon initiating (at block 150) the copy operation, the storage management software 12 creates (at block 152) an aggregate managed file 100 in the active/inactive storage pool 98 and creates an aggregate information instance 70 (FIG. 3) for the aggregate managed file, initializing the aggregate ID 72 and other values, and creates a file/aggregate information instance 50 (FIG. 2) identifying the created aggregate managed file 52 and the active/inactive storage pool 56. The files associated with the precursor aggregate 96 are copied (at block 154) to the created aggregate managed file 100 in the active/inactive storage pool 98 in which the files are written sequentially. The storage management software 12 adds (at block 156) an aggregate file information instance 80 (FIG. 4) for each file written to the aggregate 100 object including the aggregate managed file identifier 84, the offset 86 at which the file was written, the length of the file 88 and the status 90 as inactive or active. Further, the size field 54 in the file aggregate information 50 and the actual size 74, logical size 78, and number of files 78 fields in the aggregate information 70 are updated to reflect the size of the aggregate managed file 100 including the number of files written thereto.

FIG. 8 illustrates an embodiment of operations performed by the storage management software 12 to process (at block 170) a deactivation of a file associated with one precursor aggregate 96 in the active-only storage pool 94. As mentioned a file may be deactivated if the file is updated, producing both an active version having the update and an inactive version. Certain deactivations may produce only an inactive version of the file, such as if the file is deleted or the file's age exceeds a policy criteria. In response to the deactivation, the storage management software 12 indicates (at block 172) the active status fields 58, 90 in the file/aggregate information 50 and aggregate file information 80 of the deactivated file associated with the precursor aggregate 96 as inactive. If (at block 174) the file was deactivated as a result of an update operation, then the active version is added (at block 176) to the active-only storage pool by updating the file/aggregate information 50 to identify the file or aggregate and the active-only storage pool. This operation may involve adding a file/aggregate information instance 50 (FIG. 2) for the updated file and assigning a new ID 52 and indicating the

size 54 and storage pool 56 including the updated file. At this point, the active-only storage pool 94 includes inactive and active versions of the updated file. Further, files not part of an aggregate may have their status field 58 in their file/aggregate information 50 updated to indicate
5 inactive when they are updated. Yet further, the updated active file may be stored in an aggregate with other files and/or copied to other storage pools, such as the active/inactive storage pool 98. If (at block 174) the file was not deactivated in a manner that results in both an active and inactive version, then control ends.

10
FIG. 9 illustrates an embodiment of operations implemented by the storage management software 12 to expire inactive files in the active-only storage pool 94 in a random access device. In certain embodiments, a copy of the precursor aggregate in the active/inactive storage pool (made
15 according to the operations of FIG. 7) occurs before inactive files are expired according to the operations of FIG. 9. At block 200, an operation is invoked to expire inactive files on one precursor aggregate 96 in the active-only storage pool 94. This operation may be invoked periodically or in response to an event, such as an update to files in the active-only
20 aggregate. For each file in the active-only storage pool 94, a loop is performed at blocks 202 through 216. Files in the active only storage pool 94 may be determined as those files whose file/aggregate information instances 50 (FIG. 2) indicate in field 56 the active-only storage pool. If (at block 204) the file is active or inactive (which may be determined
25 from status field 58 in the file/aggregate information 50 (FIG. 2) for the file being considered) and not migrated to the active/inactive storage pool 98, then control proceeds (at block 216) to consider a next file in the active-only storage pool 94. If (at block 204) the file is both inactive and migrated to the active/inactive storage pool 98, then the
30 file is deleted (at block 206) from the active-only storage pool 94 and the file/aggregate information entry 50 (FIG. 2) for the deleted file identified in field 56 as in the active-only storage pool 96 is also removed (at block 208). The aggregate file information entry 80 for the deleted file that is also maintained in the aggregate managed file 100 in
35 the active/inactive storage pool 98 may remain because the copy of the inactive file removed from the active-only storage pool 96 is maintained in the active/inactive storage pool 98.

40
If (at block 210) the deleted file is not in a precursor aggregate, then control proceeds to block 216 to consider a next file in the active-only storage pool 94. The storage management software 12 may

determine that a removed file is in a precursor aggregate by determining whether there is an aggregate file information entry 80 for the removed file identifying the removed file as included in a precursor aggregate, which has an actual size 74 (FIG. 3) of zero. Otherwise, if (at block 5 210) the removed file is in a precursor aggregate, then the aggregate file information entry 80 having the ID 82 of the file deleted from the active only storage pool that is associated with another aggregate (precursor or managed file), as indicated in field 84, is updated (at block 214) to indicate the status 90 to inactive. In this way, only the entry 80 for 10 the file removed from the active-only storage pool 94 is removed, and not entries 80 for the file in other storage pools. Thus, during expiration, files having an inactive status are deleted from precursor aggregates in the active-only storage pool 94, yet may remain as inactive files in other storage pools, such as the active/inactive storage pool 98.

15
FIG. 10 illustrates an example of storage pools and tables having information on the aggregates in the storage pools. An active only storage pool 300, having an ID of "1", is implemented in a random access media including three files A, B, C having IDs of 100, 200, 300, 20 respectively. These files A, B, C are grouped in a precursor aggregate 302 having an ID of 400. Before updating files in the active-only storage pool 300, a copy is made of the files associated with the precursor aggregate 302 to an aggregate managed file 304, having an ID of 500, in the active/inactive storage pool 306, having an ID of "-1". The 25 active/inactive storage pool 306 may comprise a sequential access device or store files and aggregate managed files in a sequential file volume. File B 600 comprises an updated version of file B 200 associated with the precursor aggregate 302. Any changes to files associated with the precursor aggregate 302 does not effect the copy of the files in the 30 aggregate managed file 304 in the active/inactive storage pool 306, other than changing the status of the copy of the files from active to inactive. File B may also be copied to the active/inactive storage pool 306.

35
In one embodiment, the storage management software 12 maintains a file/aggregate information table 310, an aggregate information table 312, and an aggregate file information table 314 having information on the files and aggregates in the storage pools 300 and 306. The file/aggregate information table 310 includes entries, such as entry 50 (FIG. 2), for each aggregate, including entries for precursor aggregate 400 and 40 aggregate managed file 500 (having reference numbers 302 and 304, respectively) that are in the storage pools identified as 1 and "-1"

(having reference numbers 300 and 306, respectively) and the updated file B having identifier 600.

5 The aggregate file information table 314 includes entries, such as entry 80 (FIG. 4), for each file in the aggregates 302 and 304. The files in the precursor aggregate 302 have offsets of zero and the files in the aggregate managed file 304 have offsets of their location in the aggregate managed file. In FIG. 10, the aggregate information 312 has information on the precursor aggregate 302 and the aggregate managed file 304, including an actual size field of zero for precursor aggregate 400
10 indicating that the aggregate is a precursor. The table 314 further indicates whether the file is active, "Y" or active and "N" for inactive.

15 FIG. 11 illustrates an example of storage pools and tables having information on the aggregates in the storage pools after expiration occurs with respect to the storage pool 302 in FIG. 10. FIG. 11 shows storage pools 300 and 306 as in FIG. 10 and precursor aggregate 302 in storage pool 300 and aggregate managed file 304 in storage pool 306. However, in the active-only storage pool 300, the updated file B was deleted from the precursor aggregate 302 by updating the information defining the precursor
20 aggregate according to the operations of FIG. 9. However, the inactive file B 200 remains in the aggregate managed file 304 (having ID 500) in the active/inactive storage pool 306. The updated file B 600 is also in the active-only storage pool 300.

25 The updated file/aggregate information table 322 for FIG. 11 includes the entries from FIG. 10, with the entry for the inactive file B 200 removed from the active-only storage pool as a result of expiration, leaving only an entry for the updated file B 600. The entry in the updated aggregate information table 324 for the precursor aggregate 400 is
30 updated to include the new logical size of 20 to reflect that a file was removed. The entry in the aggregate file information table 326 for file B 200 in the precursor aggregate 400 is updated to indicate the inactive status. However, the length is set to 0 to indicate that that file B 200 is no longer included in the precursor aggregate 96. In this embodiment,
35 all inactive files are removed from the precursor aggregate 96 in the active-only storage pool 94 and the inactive file is maintained in the aggregate managed file 304 in the active/inactive storage pool 306.

40 Described embodiments provide an active only storage pool implemented in a random access media, such that any files in a precursor

aggregate in the active-only storage pool that become inactive as a result of an update or other deactivation are removed from the active only storage pool and precursor aggregate during an expiration operation. Further, by maintaining only active only files in a storage pool, the client may restore the active data faster from either the active-only storage pool or a backup copy of the precursor aggregate having active data, which may be on disk or tape. In this way, the active files may be streamed from the aggregate managed file on the tape or disk media. This allows the client to restore the active files faster, which the client is more likely to need than inactive files.

Additional Embodiment Details

The described operations may be implemented as a method, apparatus or article of manufacture using standard programming and/or engineering techniques to produce software, firmware, hardware, or any combination thereof. The described operations may be implemented as code maintained in a "computer readable medium", where a processor may read and execute the code from the computer readable medium. A computer readable medium may comprise media such as magnetic storage medium (e.g., hard disk drives, floppy disks, tape, etc.), optical storage (CD-ROMs, DVDs, optical disks, etc.), volatile and non-volatile memory devices (e.g., EEPROMs, ROMs, PROMs, RAMs, DRAMs, SRAMs, Flash Memory, firmware, programmable logic, etc.), etc. The code implementing the described operations may further be implemented in hardware logic (e.g., an integrated circuit chip, Programmable Gate Array (PGA), Application Specific Integrated Circuit (ASIC), etc.). Still further, the code implementing the described operations may be implemented in "transmission signals", where transmission signals may propagate through space or through a transmission media, such as an optical fiber, copper wire, etc. The transmission signals in which the code or logic is encoded may further comprise a wireless signal, satellite transmission, radio waves, infrared signals, Bluetooth, etc. The transmission signals in which the code or logic is encoded is capable of being transmitted by a transmitting station and received by a receiving station, where the code or logic encoded in the transmission signal may be decoded and stored in hardware or a computer readable medium at the receiving and transmitting stations or devices. An "article of manufacture" comprises computer readable medium, hardware logic, and/or transmission signals in which code may be implemented. Of course, those skilled in the art will recognize that many modifications may be made to this configuration without departing from the scope of the

present invention, and that the article of manufacture may comprise suitable information bearing medium known in the art.

5 In described embodiments, the aggregate managed files in which files are written were contained in a sequential access media, such as a sequential disk or tape. In an alternative embodiment, the described operations may apply to an aggregate managed file in a random access media.

10 The terms "an embodiment", "embodiment", "embodiments", "the embodiment", "the embodiments", "one or more embodiments", "some embodiments", and "one embodiment" mean "one or more (but not all) embodiments of the present invention(s)" unless expressly specified otherwise.

15 The terms "including", "comprising", "having" and variations thereof mean "including but not limited to", unless expressly specified otherwise.

20 The enumerated listing of items does not imply that any or all of the items are mutually exclusive, unless expressly specified otherwise.

The terms "a", "an" and "the" mean "one or more", unless expressly specified otherwise.

25 Devices that are in communication with each other need not be in continuous communication with each other, unless expressly specified otherwise. In addition, devices that are in communication with each other may communicate directly or indirectly through one or more intermediaries.

30 A description of an embodiment with several components in communication with each other does not imply that all such components are required. On the contrary a variety of optional components are described to illustrate the wide variety of possible embodiments of the present invention.

35 Further, although process steps, method steps, algorithms or the like may be described in a sequential order, such processes, methods and algorithms may be configured to work in alternate orders. In other words, any sequence or order of steps that may be described does not necessarily indicate a requirement that the steps be performed in that order. The

40

steps of processes described herein may be performed in any order practical. Further, some steps may be performed simultaneously.

5 When a single device or article is described herein, it will be readily apparent that more than one device/article (whether or not they cooperate) may be used in place of a single device/article. Similarly, where more than one device or article is described herein (whether or not they cooperate), it will be readily apparent that a single device/article may be used in place of the more than one device or article or a different
10 number of devices/articles may be used instead of the shown number of devices or programs. The functionality and/or the features of a device may be alternatively embodied by one or more other devices which are not explicitly described as having such functionality/features. Thus, other embodiments of the present invention need not include the device itself.

15 Further, when a reference letter, such as "a", "b", or "n" is used to denote a certain number of items, the reference "a", "b" or "n" used with different elements may indicate the same or different number of such elements.

20 FIGs. 2, 3, 4, 10, and 11 show information maintained in a certain format. In alternative embodiments, the information shown in FIGs. 2, 3, 4, 10, and 11 may be maintained in alternative data structures and formats, and in different combinations.

25 The illustrated operations of FIGs. 6, 7, 8, and 9 show certain events occurring in a certain order. In alternative embodiments, certain operations may be performed in a different order, modified or removed. Moreover, steps may be added to the above described logic and still
30 conform to the described embodiments. Further, operations described herein may occur sequentially or certain operations may be processed in parallel. Yet further, operations may be performed by a single processing unit or by distributed processing units.

35 The foregoing description of various embodiments of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It is intended that the scope of the invention be
40 limited not by this detailed description, but rather by the claims appended hereto. The above specification, examples and data provide a

complete description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made without departing from the spirit and scope of the invention, the invention resides in the claims hereinafter appended.

CLAIMS

1. A method, comprising:

5 providing an active-only storage pool in a random access device,
wherein the active-only storage pool is intended to include only active
files and not inactive versions of files;

10 associating a plurality of files in the active-only storage pool
with a precursor aggregate in the active-only storage pool;

15 deactivating one file associated with the precursor aggregate in the
active only storage pool to produce an inactive version of the deactivated
file, wherein the precursor aggregate is associated with the inactive
version of the file in response to the deactivation;

determining the inactive version of at least one file in the
precursor aggregate; and

20 deleting the determined inactive version of the at least one file
from the active only storage pool.

25 The method of claim 1, wherein the file is deactivated as a result
of an update to the file, wherein updating the file produces an active
version of the file in addition to the inactive version of the file,
further comprising:

30 removing an indication of the association of the deleted inactive
version of the at least one file with the precursor aggregate, wherein the
precursor aggregate is associated with only the active version of files in
response to removing the indication.

3. The method of claim 1, further comprising:

35 copying the files associated with the precursor aggregate to an
aggregate managed file in an additional storage pool before deleting the
determined inactive version of the at least one file, wherein the
aggregate managed file maintains the inactive version of the at least one
file deleted from the active only storage pool.

4. The method of claim 3, wherein the files are written sequentially to the aggregate managed file and wherein the aggregate managed file comprises a sequential file.

5. The method of claim 4, wherein the additional storage pool is implemented in a sequential access device.

6. The method of claim 1, wherein the precursor aggregate has a first identifier and the aggregate managed file has a second identifier, further comprising:

providing a data structure having information on defined aggregates including an aggregate identifier and storage pool including the aggregate for each indicated aggregate; and

indicating in the information for the aggregate identified by one aggregate identifier whether the aggregate is a precursor aggregate comprising an association of files or an aggregate managed file in which files are written.

7. The method of claim 6, wherein indicating in the data structure information that the aggregate is the precursor aggregate comprises indicating that the precursor aggregate has an aggregate size of zero and wherein indicating in the information that the aggregate is the aggregate managed file comprises indicating that the aggregate managed file has an aggregate size that is a cumulative size of the files written in the aggregate managed file.

8. The method of claim 7, wherein the information for each aggregate in the data structure includes a logical size and an actual size, wherein indicating that the precursor aggregate has the aggregate size of zero comprises indicating that the actual size of the precursor aggregate is zero and wherein indicating that the aggregate managed file has the aggregate size that is the cumulative size comprises indicating that the actual size of the aggregate managed file is the cumulative size, further comprising:

indicating in the information for the precursor aggregate and the aggregate managed file that the logical size is the cumulative size of the files associated with the precursor aggregate and written to the aggregate managed file, respectively.

9. The method of claim 6, wherein the file is deactivated as a result of an update to the file, wherein updating the file produces an active version of the file in addition to the inactive version of the file, further comprising:

5

adding information to the data structure for the updated file having a third identifier and indicating that the updated file is in the active-only storage pool.

10. The method of claim 6, wherein the data structure comprises a first data structure, further comprising:

15 providing a second data structure having information on each file grouped in one aggregate including the aggregate identifier of the aggregate including the file, wherein the aggregate identifier for the files associated with the precursor aggregate comprises the first identifier and wherein the aggregate identifier for the files written to the aggregate managed file comprises the second identifier.

20 11. The method of claim 10, further comprising:

deleting the information in the second data structure for the inactive version of the at least one file having the first identifier that was deleted from the active storage pool, wherein the second data structure includes information on the inactive version of the at least one file having the second identifier maintained in the aggregate managed file.

25 12. The method of claim 10, wherein the information for the files in the second data structure indicates an offset and file size of the file, wherein files associated with one precursor aggregate have an offset of zero and wherein files included in one aggregate managed file have an offset at which the file is stored in the aggregate managed file.

30 13. A system in communication with a random access device, comprising:

a processor;

35 a computer readable medium including a program executed by the processor to perform operations, the operations comprising:

40

configuring an active-only storage pool in the random access device, wherein the active-only storage pool is intended to include only active files and not inactive versions of files;

5 associating a plurality of files in the active-only storage pool with a precursor aggregate in the active-only storage pool;

10 deactivating one file associated with the precursor aggregate in the active only storage pool to produce an inactive version of the deactivated file, wherein the precursor aggregate is associated with the inactive version of the file in response to the deactivation;

15 determining the inactive version of at least one file in the precursor aggregate; and

deleting the determined inactive version of the at least one file from the active only storage pool.

20 14. The system of claim 13, wherein the file is deactivated as a result of an update to the file, wherein updating the file produces an active version of the file in addition to the inactive version of the file, and wherein the operations further comprise:

25 removing an indication of the association of the deleted inactive version of the at least one file with the precursor aggregate, wherein the precursor aggregate is associated with only the active version of files in response to removing the indication.

30 15. The system of claim 13, wherein the operations further comprise:

35 copying the files associated with the precursor aggregate to an aggregate managed file in an additional storage pool before deleting the determined inactive version of the at least one file, wherein the aggregate managed file maintains the inactive version of the at least one file deleted from the active only storage pool.

40 16. The system of claim 15, wherein the files are written sequentially to the aggregate managed file and wherein the aggregate managed file comprises a sequential file.

17. The system of claim 16, wherein the additional storage pool is implemented in a sequential access device.

18. The system of claim 13, wherein the precursor aggregate has a first identifier and the aggregate managed file has a second identifier, further comprising:

a memory accessible to the processor;

wherein the operations further comprise:

providing a data structure in the memory having information on defined aggregates including an aggregate identifier and storage pool including the aggregate for each indicated aggregate; and

indicating in the information for the aggregate identified by one aggregate identifier whether the aggregate is a precursor aggregate comprising an association of files or an aggregate managed file in which files are written.

19. The system of claim 18, wherein indicating in the data structure information that the aggregate is the precursor aggregate comprises indicating that the precursor aggregate has an aggregate size of zero and wherein indicating in the information that the aggregate is the aggregate managed file comprises indicating that the aggregate managed file has an aggregate size that is a cumulative size of the files written in the aggregate managed file.

20. The system of claim 18, wherein the information for each aggregate in the data structure includes a logical size and an actual size, wherein indicating that the precursor aggregate has the aggregate size of zero comprises indicating that the actual size of the precursor aggregate is zero and wherein indicating that the aggregate managed file has the aggregate size that is the cumulative size comprises indicating that the actual size of the aggregate managed file is the cumulative size, and wherein the operations further comprise:

indicating in the information for the precursor aggregate and the aggregate managed file that the logical size is the cumulative size of the files associated with the precursor aggregate and written to the aggregate managed file, respectively.

21. The system of claim 18, wherein the file is deactivated as a result of an update to the file, wherein updating the file produces an active version of the file in addition to the inactive version of the file, and wherein the operations further comprise:

5

adding information to the data structure for the updated file having a third identifier and indicating that the updated file is in the active-only storage pool.

10 22. The system of claim 18, wherein the data structure comprises a first data structure, and wherein the operations further comprise:

15 providing a second data structure in the memory having information on each file grouped in one aggregate including the aggregate identifier of the aggregate including the file, wherein the aggregate identifier for the files associated with the precursor aggregate comprises the first identifier and wherein the aggregate identifier for the files written to the aggregate managed file comprises the second identifier.

20 23. The system of claim 22, wherein the operations further comprise:

25 deleting the information in the second data structure for the inactive version of the at least one file having the first identifier that was deleted from the active storage pool, wherein the second data structure includes information on the inactive version of the at least one file having the second identifier maintained in the aggregate managed file.

30 24. The system of claim 22, wherein the information for the files in the second data structure indicates an offset and file size of the file, wherein files associated with one precursor aggregate have an offset of zero and wherein files included in one aggregate managed file have an offset at which the file is stored in the aggregate managed file.

35 25. An article of manufacture including code capable of communicating with a random access device and causing operations to be performed, the operations comprising:

40 configuring an active-only storage pool in the random access device, wherein the active-only storage pool is intended to include only active files and not inactive versions of files;

associating a plurality of files in the active-only storage pool with a precursor aggregate in the active-only storage pool;

5 deactivating one file associated with the precursor aggregate in the active only storage pool to produce an inactive version of the deactivated file, wherein the precursor aggregate is associated with the inactive version of the file in response to the deactivation;

10 determining the inactive version of at least one file in the precursor aggregate; and

 deleting the determined inactive version of the at least one file from the active only storage pool.

15 26. The article of manufacture of claim 25, wherein the file is deactivated as a result of an update to the file, wherein updating the file produces an active version of the file in addition to the inactive version of the file, further comprising:

20 removing an indication of the association of the deleted inactive version of the at least one file with the precursor aggregate, wherein the precursor aggregate is associated with only the active version of files in response to removing the indication.

25 27. The article of manufacture of claim 25, wherein the operations further comprise:

30 copying the files associated with the precursor aggregate to an aggregate managed file in an additional storage pool before deleting the determined inactive version of the at least one file, wherein the aggregate managed file maintains the inactive version of the at least one file deleted from the active only storage pool.

35 28. The article of manufacture of claim 27, wherein the files are written sequentially to the aggregate managed file and wherein the aggregate managed file comprises a sequential file.

40 29. The article of manufacture of claim 28, wherein the additional storage pool is implemented in a sequential access device.

30. The article of manufacture of claim 25, wherein the code is further in communication with a memory, wherein the precursor aggregate has a first identifier and the aggregate managed file has a second identifier, and wherein the operations further comprise:

5

providing a data structure in the memory having information on defined aggregates including an aggregate identifier and storage pool including the aggregate for each indicated aggregate; and

10

indicating in the information for the aggregate identified by one aggregate identifier whether the aggregate is a precursor aggregate comprising an association of files or an aggregate managed file in which files are written.

15

31. The article of manufacture of claim 30, wherein indicating in the data structure information that the aggregate is the precursor aggregate comprises indicating that the precursor aggregate has an aggregate size of zero and wherein indicating in the information that the aggregate is the aggregate managed file comprises indicating that the aggregate managed file has an aggregate size that is a cumulative size of the files written in the aggregate managed file.

20

32. The article of manufacture of claim 31, wherein the information for each aggregate in the data structure includes a logical size and an actual size, wherein indicating that the precursor aggregate has the aggregate size of zero comprises indicating that the actual size of the precursor aggregate is zero and wherein indicating that the aggregate managed file has the aggregate size that is the cumulative size comprises indicating that the actual size of the aggregate managed file is the cumulative size, and wherein the operations further comprise:

30

indicating in the information for the precursor aggregate and the aggregate managed file that the logical size is the cumulative size of the files associated with the precursor aggregate and written to the aggregate managed file, respectively.

35

33. The article of manufacture of claim 30, wherein the file is deactivated as a result of an update to the file, wherein updating the file produces an active version of the file in addition to the inactive version of the file, wherein the operations further comprise:

5

adding information to the data structure for the updated file having a third identifier and indicating that the updated file is in the active-only storage pool.

10 34. The article of manufacture of claim 30, wherein the data structure comprises a first data structure, wherein the operations further comprise:

15 providing a second data structure in the memory having information on each file grouped in one aggregate including the aggregate identifier of the aggregate including the file, wherein the aggregate identifier for the files associated with the precursor aggregate comprises the first identifier and wherein the aggregate identifier for the files written to the aggregate managed file comprises the second identifier.

20 35. The article of manufacture of claim 34, wherein the operations further comprise:

25 deleting the information in the second data structure for the inactive version of the at least one file having the first identifier that was deleted from the active storage pool, wherein the second data structure includes information on the inactive version of the at least one file having the second identifier maintained in the aggregate managed file.

30 36. The article of manufacture of claim 34, wherein the information for the files in the second data structure indicates an offset and file size of the file, wherein files associated with one precursor aggregate have an offset of zero and wherein files included in one aggregate managed file have an offset at which the file is stored in the aggregate managed file.

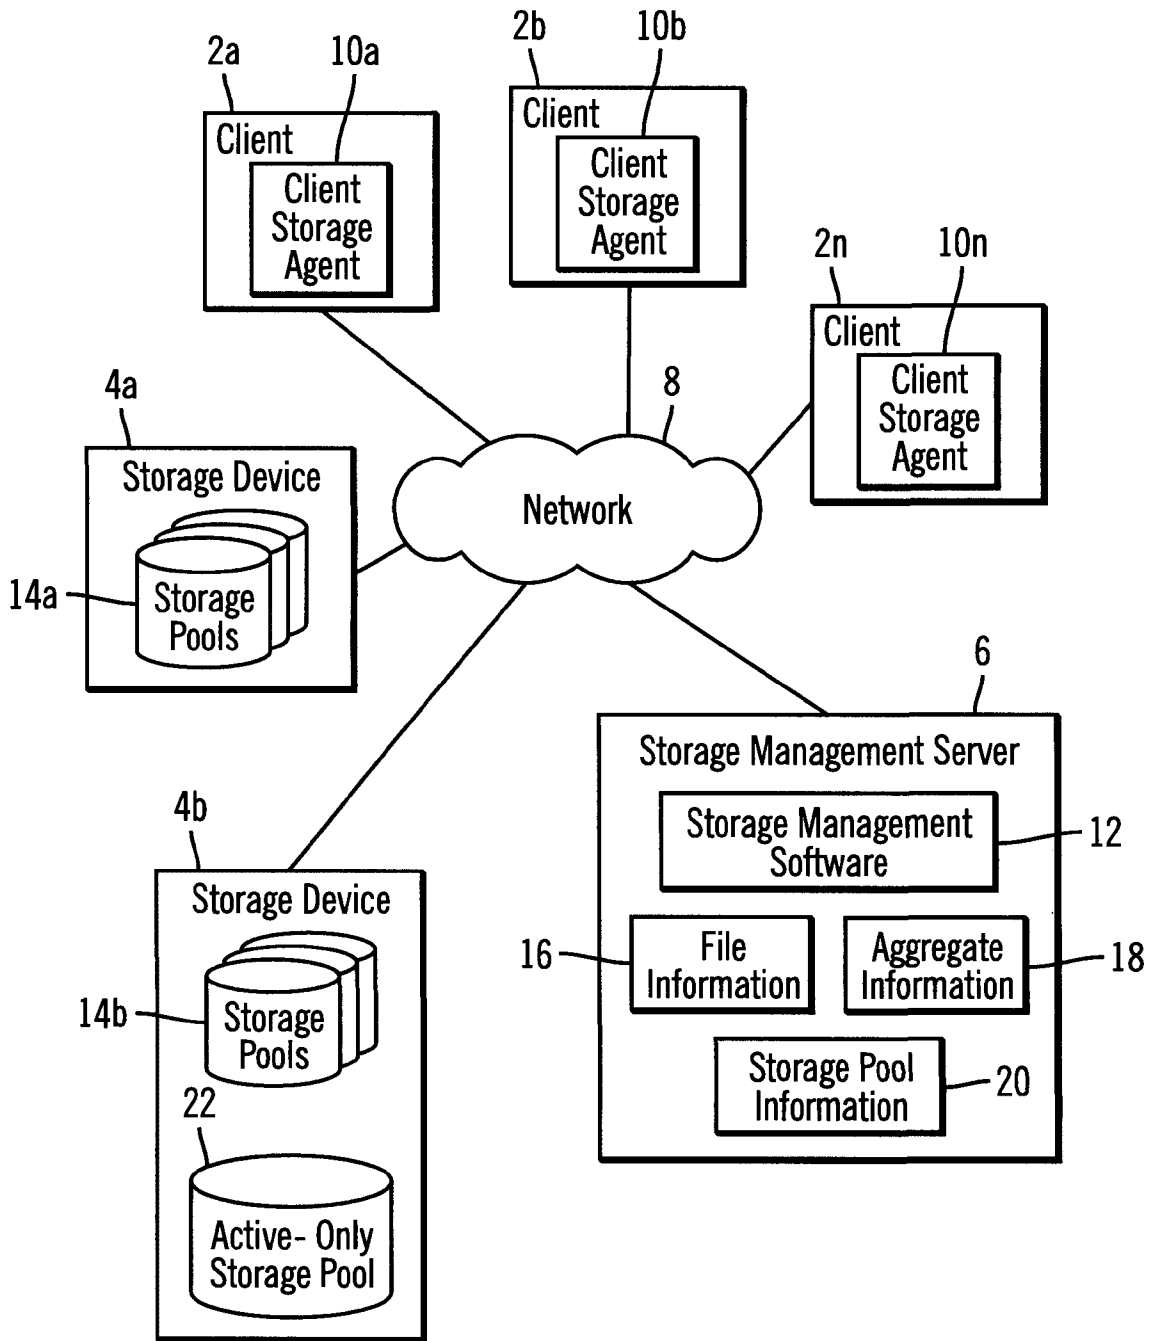
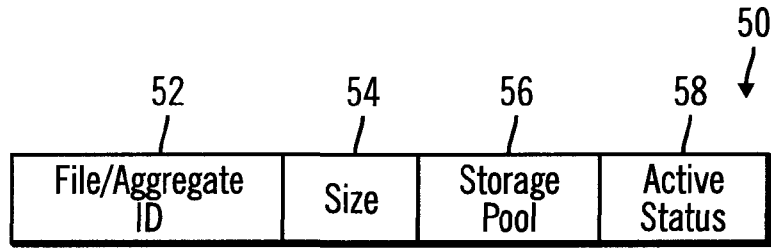


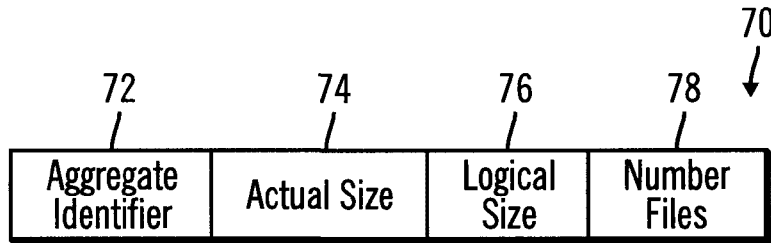
FIG. 1

2/7



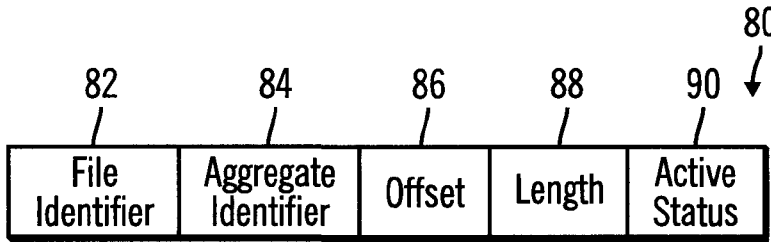
File/Aggregate Information

FIG. 2



Aggregate Information

FIG. 3



Aggregate File Information

FIG. 4

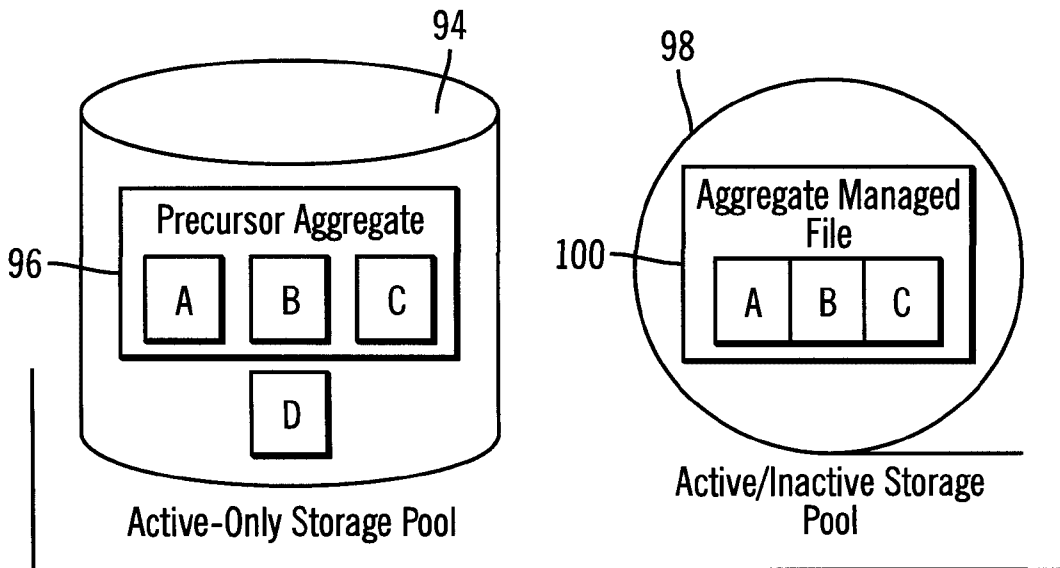


FIG. 5

3/7

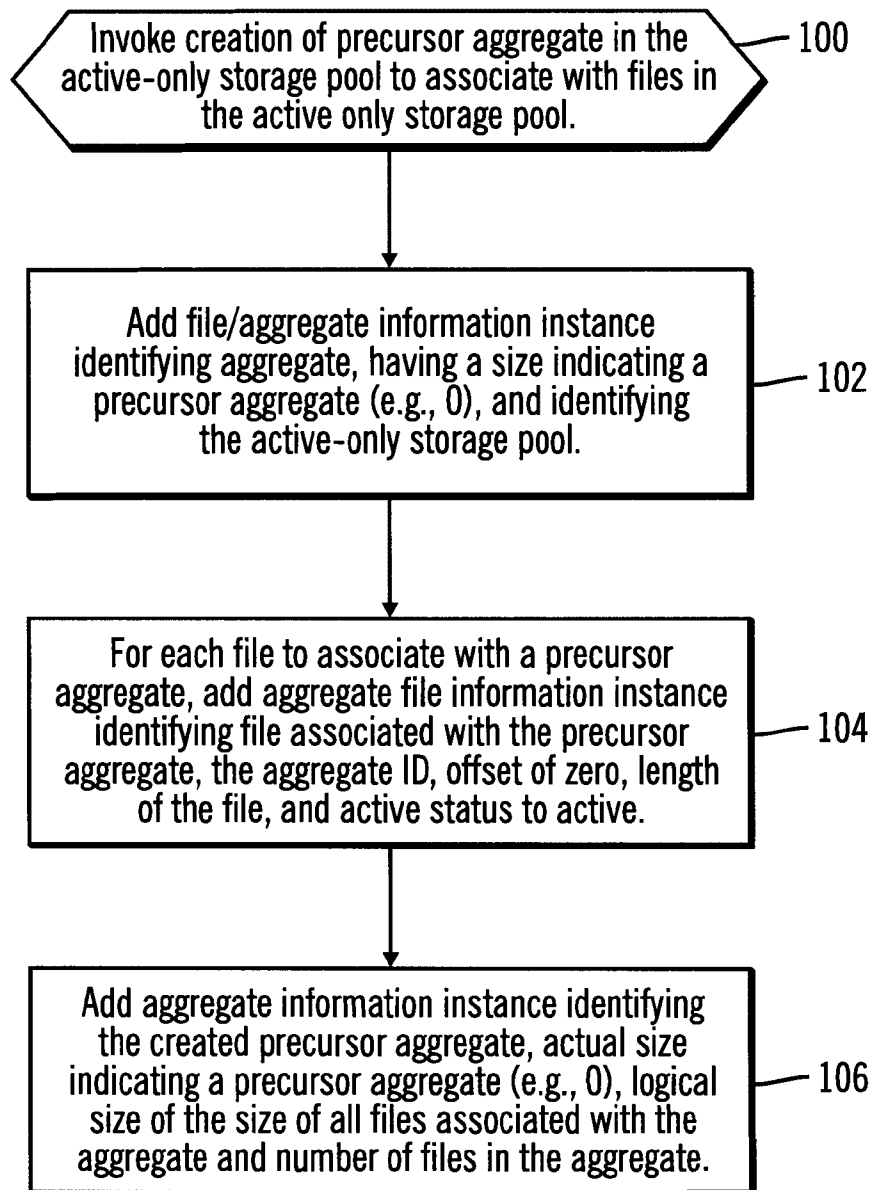


FIG. 6

4/7

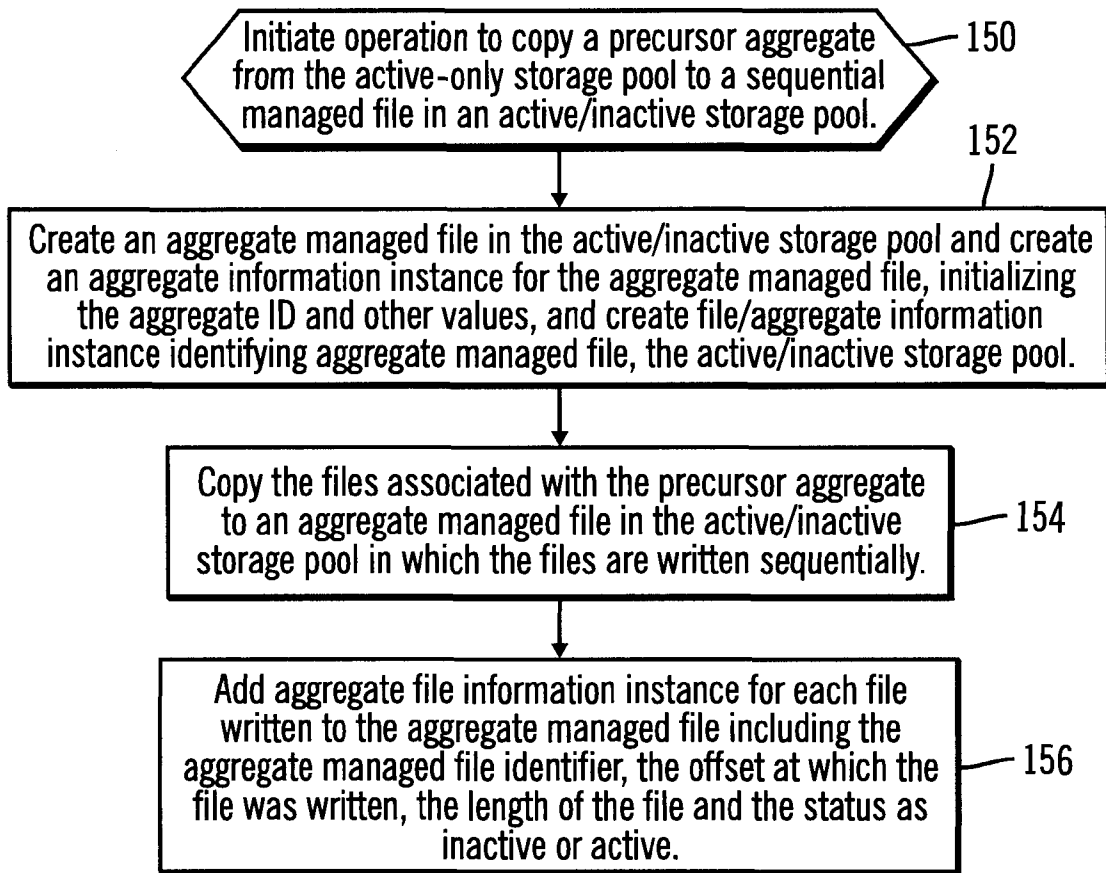


FIG. 7

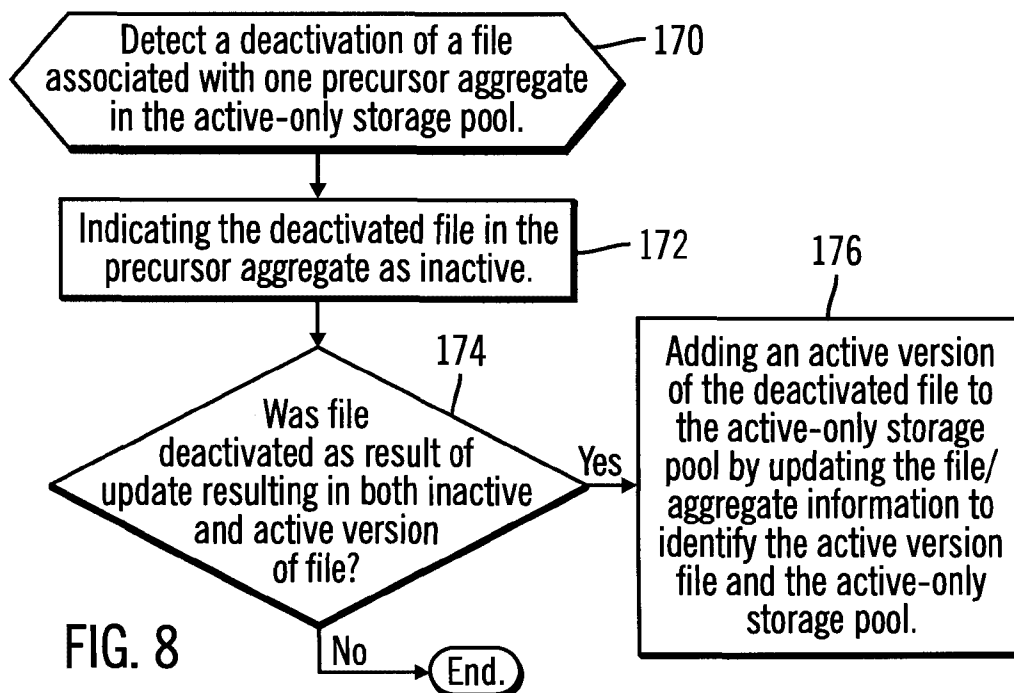


FIG. 8

5/7

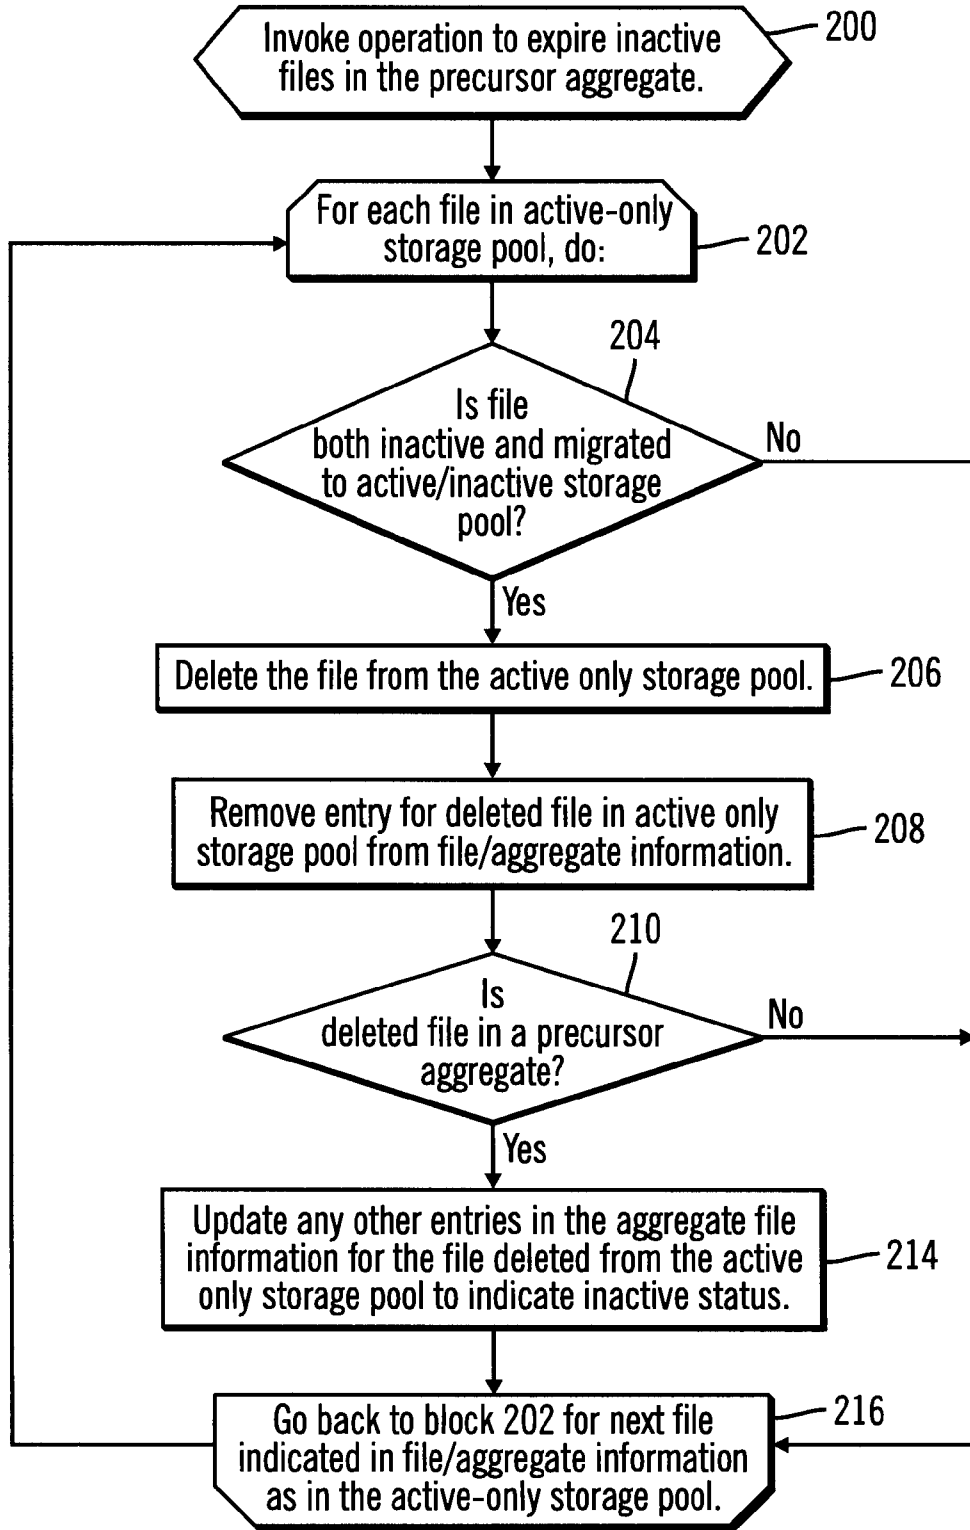


FIG. 9

6/7

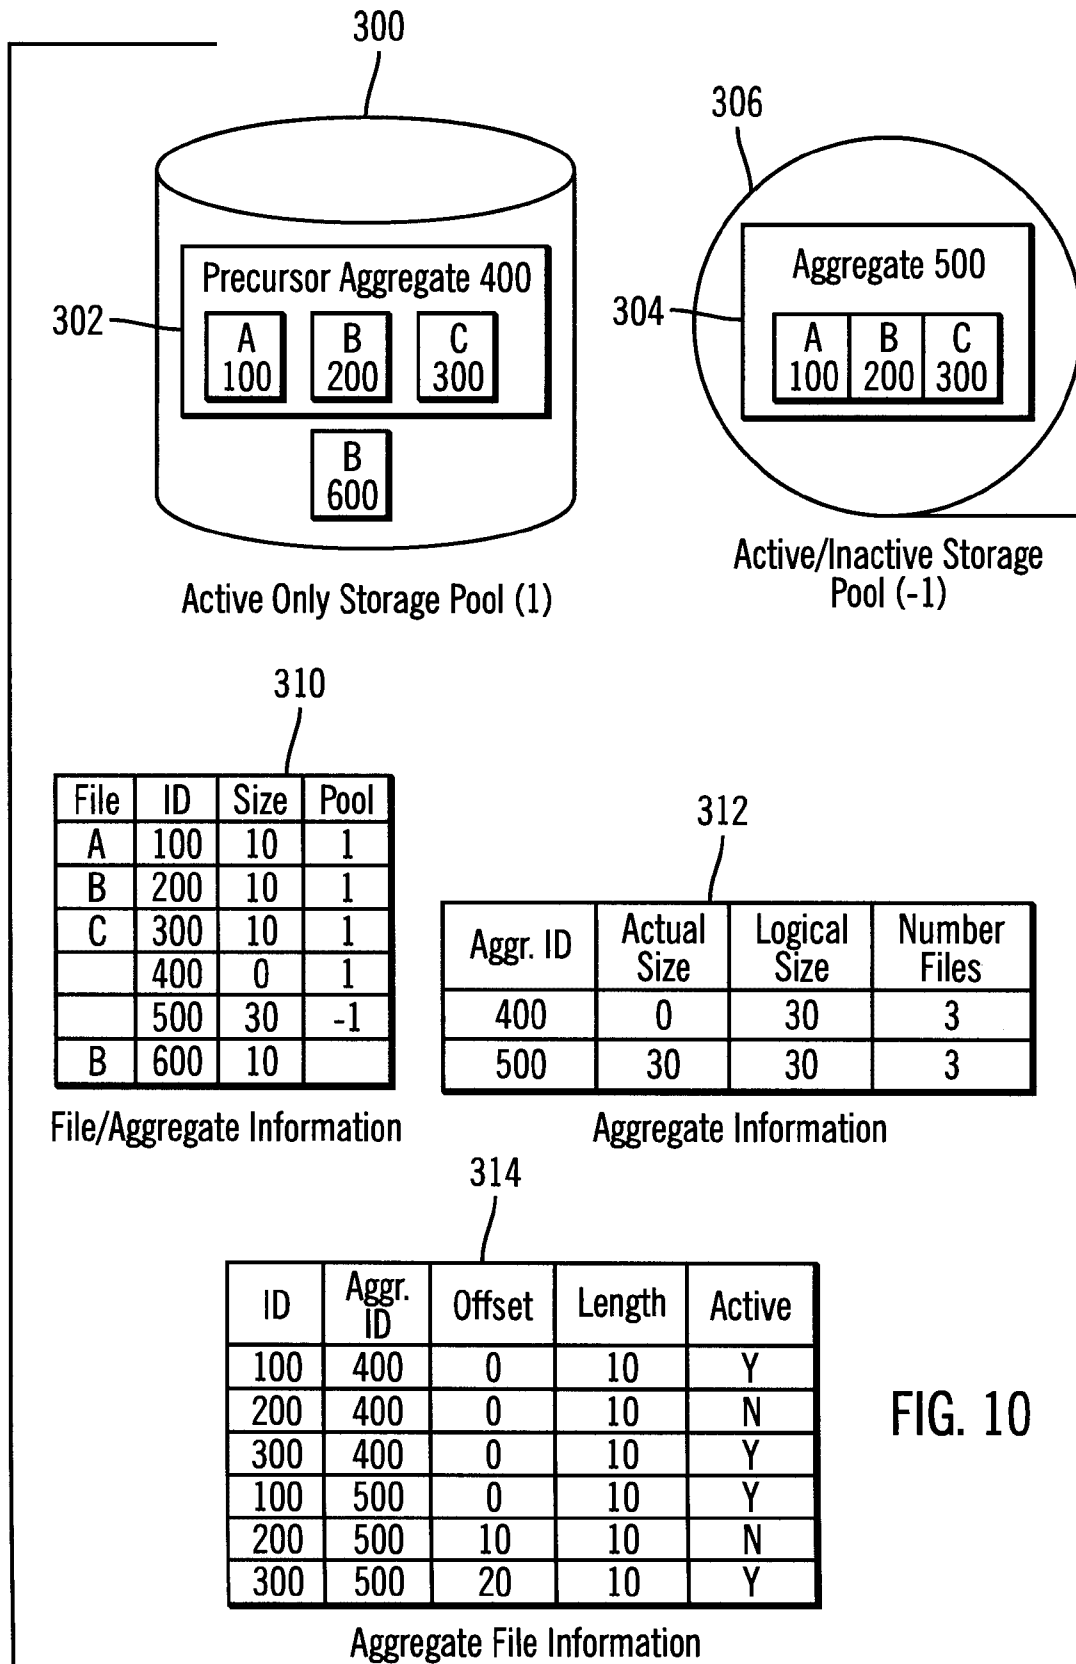


FIG. 10

7/7

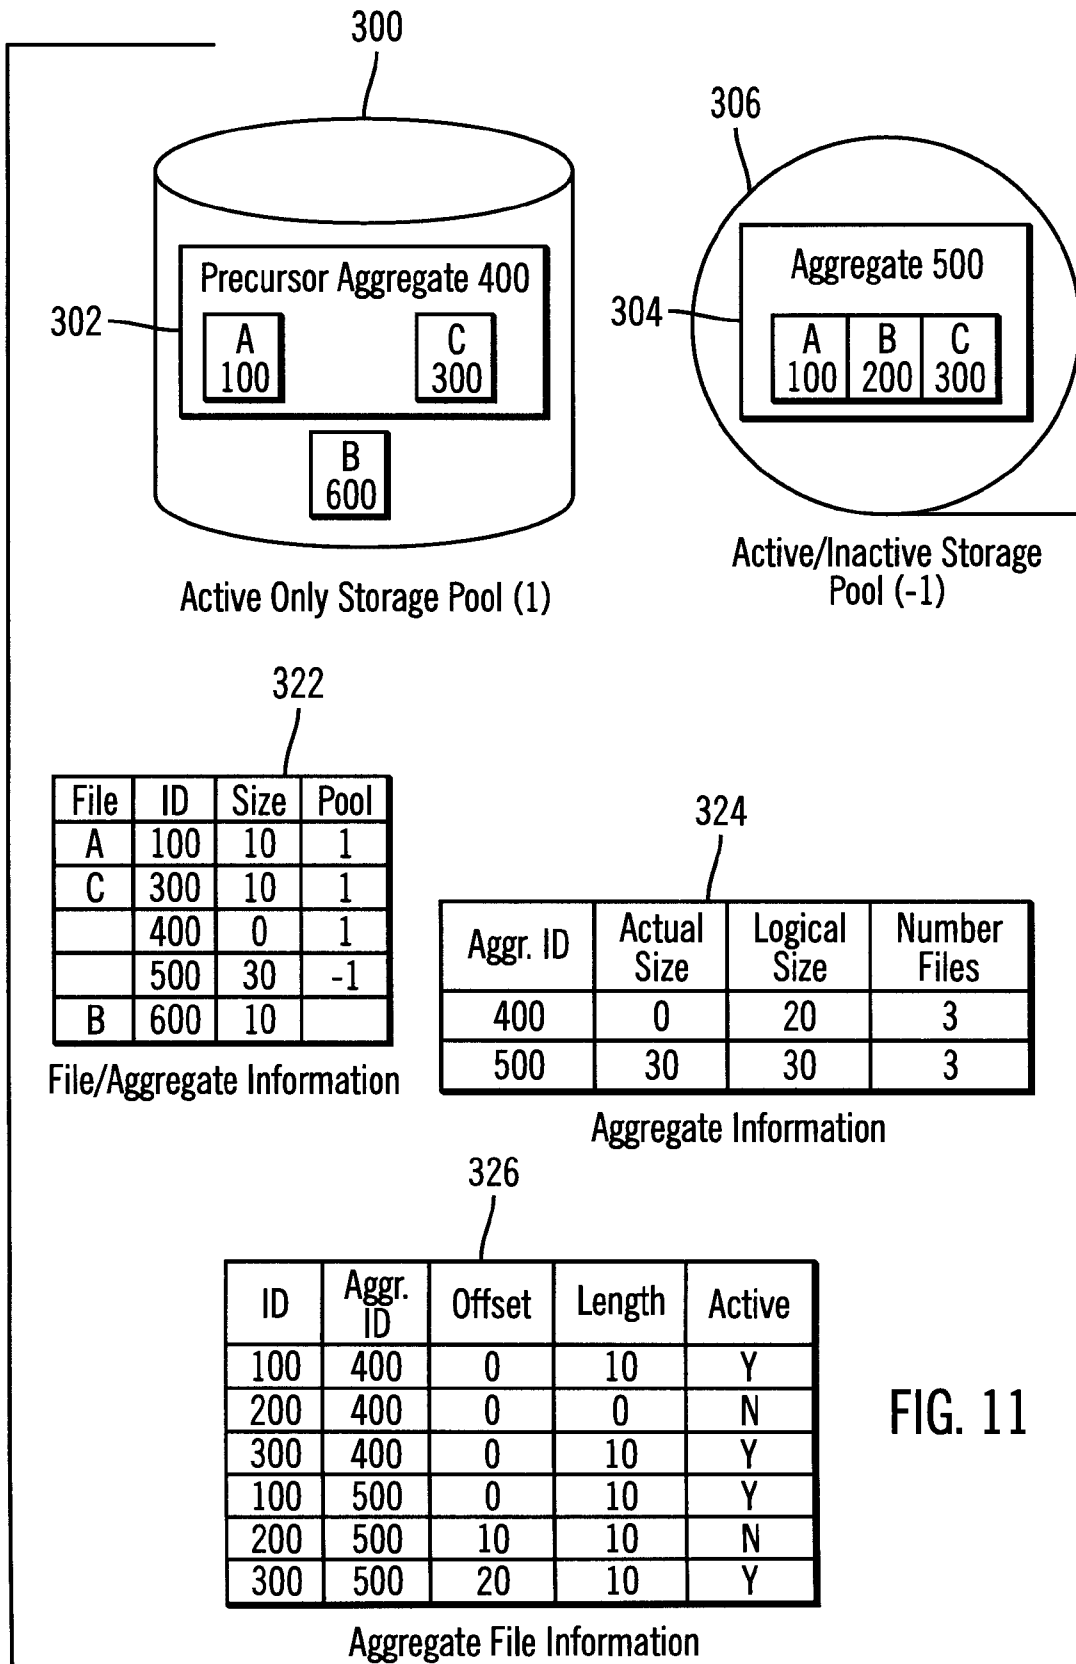


FIG. 11

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2006/063455

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, INSPEC, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2005/165722 A1 (CANNON DAVID M [US] ET AL) 28 July 2005 (2005-07-28) cited in the application abstract paragraph [0010] - paragraph [0014] paragraph [0035] - paragraph [0084]	1-36
A	US 5 983 239 A (CANNON DAVID MAXWELL [US]) 9 November 1999 (1999-11-09) abstract column 2, line 14 - column 4, line 29	1-36
A	US 6 330 572 B1 (SITKA LARRY [US]) 11 December 2001 (2001-12-11) abstract column 1, line 64 - column 7, line 49	1-36
	-/--	

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

6 November 2006

Date of mailing of the international search report

13/11/2006

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Wienold, Norbert

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2006/063455

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2004/199521 A1 (ANGLIN MATTHEW JOSEPH [US] ET AL) 7 October 2004 (2004-10-07) abstract paragraph [0001] - paragraph [0011] -----	1-36

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2006/063455

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2005165722 A1	28-07-2005	WO 2005071565 A2	04-08-2005
US 5983239 A	09-11-1999	US 6041334 A	21-03-2000
US 6330572 B1	11-12-2001	NONE	
US 2004199521 A1	07-10-2004	NONE	