



- (51) International Patent Classification:  
*G06F 17/30* (2006.01)
- (21) International Application Number:  
PCT/US2013/044105
- (22) International Filing Date:  
4 June 2013 (04.06.2013)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
61/655,452 4 June 2012 (04.06.2012) US  
13/909,021 3 June 2013 (03.06.2013) US
- (71) Applicant: **GOOGLE INC.** [US/US]; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US).
- (72) Inventors: **SAITO, Yasushi**; 405 Stierlin Rd #43, Mountain View, CA 94043 (US). **GHEMAWAT, Sanjay**; 1067 Mercy St, Mountain View, CA 94041 (US). **DEAN, Jeffrey, Adgate**; 995 Matadero Avenue, Palo Alto, CA 94306 (US).
- (74) Agents: **.CRISMAN, Douglas, J.** et al.; Morgan Lewis & Bockius LLP, 2 Palo Alto Square, 3000 El Camino Real, Suite 700, Palo Alto, CA 94306 (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, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

[Continued on next page]

(54) Title: METHOD AND SYSTEM FOR DELETING OBSOLETE FILES FROM A FILE SYSTEM

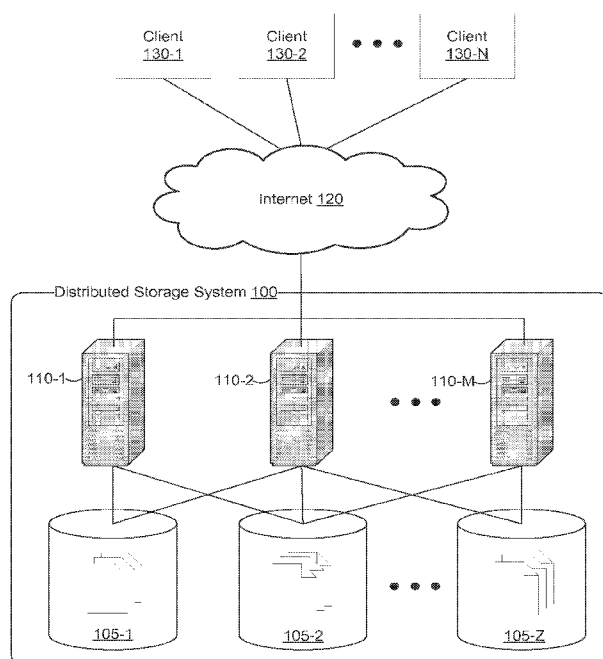


FIG. 1

(57) Abstract: A method for deleting obsolete files from a file system is provided. The method includes: receiving a request to delete a reference to a target file in a file system from a file reference data structure, wherein the file reference data structure includes target file names and reference file names; identifying a reference file name in the file reference data structure, wherein the reference file name includes a file name of the target file; deleting a reference file from the file system, wherein the reference file has the identified reference file name; checking whether the file system includes at least one reference file whose file name matches the file name of the target file; if there is no such reference file in the file system: deleting the target file from the file system; and deleting the file name of the target file from the file reference data structure.



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

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

**Published:**

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

# METHOD AND SYSTEM FOR DELETING OBSOLETE FILES FROM A FILE SYSTEM

## TECHNICAL FIELD

[0001] The disclosed implementations relate generally to the management of a file system, and in particular, to system and method for deleting obsolete files from a file system.

## BACKGROUND

[0002] A file system in a computer system is used for managing data stored in a storage device. Typically, the data is organized into many data files and each data file is stored in a particular directory of the file system. Different data files in the file system usually operate on their own such that the addition or removal of one data file in the file system has little, if any, impact on the other data files. Compared with a file system, a database system also deals with a collection of data. But the data in the database system is often organized into many inter-related data records within one or more data structures (e.g., tables). In response to a query, the database system identifies within different tables a set of data records that satisfy the query and return the set of records or a derivative thereof as a response to the query. Similarly, the deletion of a data record from one table may trigger the deletion of a data record from another table. Currently, the line separating the file system from the database system is becoming increasingly blurry because data records in the database system often include references or pointers to different data files in the file system. As a result, an update to the database system may cause an update to the file system.

## SUMMARY

[0003] In accordance with some implementations described below, a method performed by a computer server for deleting obsolete files from a file system is provided. The method includes: receiving a request to delete a reference to a target file in a file system from a file reference data structure, wherein the file reference data structure includes a plurality of target file names and a plurality of reference file names; identifying a reference file name in the file reference data structure, wherein the reference file name includes a file name of the target file in the file system; deleting a reference file from the file system, wherein the reference file has the identified reference file name; checking whether the file system includes at least one reference file whose file name matches the file name of the target file in the file system; if

there is no such reference file in the file system: deleting the target file from the file system; and deleting the file name of the target file in the file system from the file reference data structure.

**[0004]** In accordance with some implementations described below, a computer system for deleting obsolete files from a file system includes one or more processors; memory; and a plurality of program modules, wherein the plurality of program modules are stored in the memory and to be executed by the one or more processors. The plurality of program modules include instructions for: receiving a request to delete a reference to a target file in a file system from a file reference data structure, wherein the file reference data structure includes a plurality of target file names and a plurality of reference file names; identifying a reference file name in the file reference data structure, wherein the reference file name includes a file name of the target file in the file system; deleting a reference file from the file system, wherein the reference file has the identified reference file name; checking whether the file system includes at least one reference file whose file name matches the file name of the target file in the file system; if there is no such reference file in the file system: deleting the target file from the file system; and deleting the file name of the target file in the file system from the file reference data structure.

**[0005]** In accordance with some implementations described below, a non-transitory computer readable-storage medium stores one or more programs for execution by one or more processors of a computer system that includes one or more processors and memory for deleting obsolete files from a file system. The one or more program modules include instructions: receiving a request to delete a reference to a target file in a file system from a file reference data structure, wherein the file reference data structure includes a plurality of target file names and a plurality of reference file names; identifying a reference file name in the file reference data structure, wherein the reference file name includes a file name of the target file in the file system; deleting a reference file from the file system, wherein the reference file has the identified reference file name; checking whether the file system includes at least one reference file whose file name matches the file name of the target file in the file system; if there is no such reference file in the file system: deleting the target file from the file system; and deleting the file name of the target file in the file system from the file reference data structure.

## BRIEF DESCRIPTION OF DRAWINGS

[0006] The aforementioned implementation of the invention as well as additional implementations will be more clearly understood as a result of the following detailed description of the various aspects of the invention when taken in conjunction with the drawings. Like reference numerals refer to corresponding parts throughout the several views of the drawings.

[0007] FIG. 1 is a block diagram illustrating how clients interact with a distributed storage system including a plurality of servers and databases in accordance with some implementations.

[0008] FIG. 2 is a block diagram illustrating the components of a tablet in accordance with some implementations.

[0009] FIG. 3A is a flow chart illustrating how to create a new target file that belongs to a tablet in accordance with some implementations.

[0010] FIG. 3B is a flow chart illustrating how to create a new reference to an existing target file that belongs to a tablet in accordance with some implementations.

[0011] FIG. 3C is a flow chart illustrating how to remove a reference to an existing target file that belongs to a tablet in accordance with some implementations.

[0012] FIG. 4 is a block diagram illustrating the components of the distributed storage system in accordance with some implementations.

## DETAILED DESCRIPTION

[0013] FIG. 1 is a block diagram illustrating how clients (130-1, 130-2, 130-N) interact with a distributed storage system 100 including a plurality of servers (110-1, 110-2, 110-M) and storage devices (105-1, 105-2, 105-Z) in accordance with some implementations. Both the clients 130 and the distributed storage system 100 are connected to the Internet 120. In some implementations, a client is a desktop/laptop computer through which an end user accesses services on the Internet, e.g., checking emails or playing on-line games, etc. In some other implementations, the client is a computer server itself that offers on-line services. For example, the client is a web server that provides on-line video streaming service to end

users and it uses the distributed storage system 100 for hosting a large amount of video streams. On behalf of end users, the web server may send requests to the distributed storage system 100 for the user-requested video streams and return them to the requesting end users.

**[0014]** As a distributed storage system, the system 100 may include multiple data centers located at different geographic regions (e.g., one on the East Coast of the United States and another one on the West Coast of the United States) or even across the world. A data access request submitted to the distributed storage system 100 is received and processed by one of the servers 110 (usually the one that is geographically close to the source of the request). Within the distributed storage system 100, the server identifies one or more storage devices 105 that have the requested data and returns the requested data to the requesting client. Note that the client-requested data may be stored in the data storage devices at the same data center as the server or in the data storage devices at another data center remote from the server. A data update request (e.g., adding a video stream file to the distributed storage system or deleting a video stream file from the distributed storage system) submitted to the distributed storage system 100 may cause delete operations at one or more storage devices that store the video stream file.

**[0015]** In some implementations, data within a storage device 105 is organized into tablets to combine the benefit of a file system and that of a database system. For example, the storage device 105 includes a table-like data structure including multiple rows, each row including data associated with a user account of an on-line service and one or more pointers or references to files managed by a file system associated with the storage device 105. A tablet is a data unit including one or more rows (or partial rows) that associated a file in the file system with a particular entry in the tablet.

**[0016]** FIG. 2 is a block diagram illustrating the components of a tablet 200 in accordance with some implementations. For simplicity, this exemplary tablet 200 focuses on the relationship between a tablet descriptor file 210 and a plurality of files referenced by the tablet descriptor file 210. In some implementations, the tablet descriptor file 210 is also a file managed by the file system. In some other implementations, the tablet descriptor file 210 is a specific data structure designed for hosting the file referencing data.

**[0017]** In some implementations, files associated with the tablet descriptor file 210 are divided into two groups: a group of target files such as files 240-1 and 240-2 and a group of

reference files 250-1, 250-2, and 250-M. A target file corresponds to a file that has some data (e.g., an email message or a segment of a video stream) managed by a storage device in connection with an on-line service. A reference file corresponds to a file in the file system that represents a reference to a target file, which may be generated in response to a data access request from a particular client to the distributed storage system 100. Accordingly, the tablet descriptor file 210 includes a list of target file names 220 and a list of references 230. Each file name (220-1, 220-2) in the list of target file names 220 identified a target file (240-1, 240-2) in the file system. In some implementations, the file name includes a directory path to the target file in the file system.

**[0018]** A reference 230-1 in the list of references 230 further includes a status indicator 231-1, a target file name 233-1, and a reference file name 235-1. The status indicator may have one of multiple attributes indicating that the current status of a reference file associated with the reference. For example, the status indicator of a reference file that has been generated in the file system but whose corresponding reference is still being generated in the tablet descriptor file 210 is “TENTATIVE.” After the corresponding reference is generated in the tablet descriptor file 210, the status indicator changes from “TENTATIVE” to “LIVE,” indicating the completion of generating the reference. In response to a request for deleting the reference and before the deletion of the reference from the tablet descriptor file 210, the status indicator changes from “LIVE” to “TO-DELETE.” The target file name 233-1 identifies a target file in the file system and it is the same as the corresponding target file name 220-1 in the list of target file names 220 (as indicated by the dashed-line between the target file name 233-1 and the target file 240-1). The reference file name 235-1 identifies a reference file 250-1 in the file system. In some implementations, there is a predefined naming convention for the reference file name 235-1. As shown in FIG. 2, the file name of the reference file 250-1 includes three components: the file name of the corresponding target file 240-1, a text string “.ref.” and a unique identifier <guid-1>. In some implementations, the unique identifier further includes one or more of a current timestamp of the distributed storage system 100, an IP address or the like associated with a storage device 105 or a server 110 or both.

**[0019]** In some implementations, a reference 230-1 is generated in the list of references 230 in connection with the creation of a target file 240-1 in the file system because the creation of the target file 240-1 is usually triggered by a client access request for the data in

the target file 240-1. In this case, the reference 230-1 is typically the first reference in the tablet descriptor file 210 for the target file 240-1. In addition to the reference 230-1, a reference file 250-1 is generated in the file system. In some implementations, the reference file 250-1 is an empty file because its existence is to indicate that there is at least one reference to the target file 240-1. In response to another client access request for the data in the target file 240-1, a new reference 230-M is generated in the list of references 230 and a new reference file 250-M is generated in the file system, indicating that one target file may be referenced for multiple times by different clients and is therefore associated with multiple reference files in the file system.

**[0020]** Of course, it is also possible that a target file 240-2 may be referenced by only one client and therefore has only one reference file 250-2. It should be noted that the reference entry 230-2 in the list of references 230 may or may not be the entry corresponding to the creation of the target file 240-2. For example, a target file may, at some point, be referenced by multiple clients and has therefore multiple references in the tablet descriptor file. But as time goes on, many references (including the one associated with the creation of the target file 240-2) may be deleted such that there is only one reference 230-2 in the list of references 230 for the target file 240-2. As will be described below, a deletion of the reference 230-2 will trigger the deletion of the target file 240-2 because it has no active reference in the tablet descriptor file 210, which indicates that the target file 240-2 and the reference file 250-2 should be cleaned up from the file system so that the file system can reclaim the space occupied by the two files for other uses.

**[0021]** FIG. 3A is a flow chart illustrating how to create a new target file that belongs to a tablet T in accordance with some implementations. Assuming that the tablet T has been created before, the first step is for a server to get (301) a tablet descriptor file D(T) associated with the tablet T. For the new target file X, the server generates (303) a global unique identifier “guid.” In some implementations, the global unique identifier “guid” may include the server’s IP address, a current timestamp or the like generated by the server such that the global unique identifier “guid” is unique within the distributed storage system 100. The server then creates (305) a new reference R for the target file X, which may include the following attributes {Status=TENTATIVE, Target\_Name=X, Ref\_Name=X.ref.<guid>}. As noted above, the reference R’s status indicator is set to “TENTATIVE” because neither the target file X and the corresponding reference file has been created in the file system yet.



**[0022]** Next, the server adds (307) the new reference R to the tablet descriptor file D(T) and creates an empty reference file in the file system using the reference file name Ref\_Name. The server then adds (309) the new target file name X to the list of target file names of the tablet descriptor file D(T) and creates the target file X in the file system. After the creation of the target file X and the reference file Ref\_Name in the file system, the server updates (311) the reference R to change its status indicator from “TENTATIVE” to “LIVE.” In some implementations, the status indicator of the reference R serves like a lock to the reference R in the tablet descriptor file D(T) to deny a client request for a target file through a particular reference R if the reference’s status indicator is not “LIVE.”

**[0023]** FIG. 3B is a flow chart illustrating how to create a new reference to an existing target file that belongs to a tablet in accordance with some implementations. Using the same example, the server gets (321) the tablet descriptor file D(T) of the tablet T and then generates (323) a new global unique identifier “guid” for the new reference to the target file X. The server creates (325) a new reference R for the target file X, which may include the following attributes {Status=TENTATIVE, Target\_Name=X, Ref\_Name=X.ref.<guid>}. As noted above, the naming convention of the global unique identifier <guid> guarantees that the new reference is different from the previous reference created in connection with the creation of the target file X in the file system. Next, the server adds (327) the new reference R to the tablet descriptor file D(T) and creates an empty reference file in the file system using the reference file name Ref\_Name. Finally, the server updates (329) the reference R to change its status indicator from “TENTATIVE” to “LIVE,” making the reference R available to the requesting client.

**[0024]** FIG. 3C is a flow chart illustrating how to remove a reference to an existing target file that belongs to a tablet in accordance with some implementations. Using the same example, the server gets (341) the tablet descriptor file D(T) of the tablet T and then identifies (343) a reference R for the target file X in the tablet descriptor file D(T). In some implementations, the server request a request for removing the reference from the tablet descriptor file D(T) and the request includes an identifier associated with the reference to be removed. As an initial step of removing the reference, the server updates (345) the status indicator of the reference R from “LIVE” to “TO-DELETE.” As noted above, the status indicator serves like a lock to the reference R such that any subsequent access to the reference R may be denied as a result of the change of status. After identifying the reference file name

in the reference R, the server deletes (347) the reference file corresponding to the reference R from the file system.

**[0025]** Next, the server checks (349) whether there is any other existing reference to the target file X. In some implementations, the server performs this operation by checking whether there is any reference file in the file system including the file name of the target file X. If so (349-Yes), the server stops taking further actions. Otherwise (349-No), the server determines that all the references to the target file X have been deleted, which renders the target file X obsolete. As such, the removal of the target file X from the file system is assumed to have no impact on the performance of the distributed storage system 100. Accordingly, the server deletes (351) the target file X from the file system. The server also deletes (353) the file name “X” and all the references to the file name “X” (whose status indicators are presumably “TO-DELETE”) from the tablet descriptor file D(T).

**[0026]** FIG. 4 is a block diagram illustrating the components of the distributed storage system 100 in accordance with some implementations. The distributed storage system 100 includes one or more processing units (CPU’s) 402 for executing modules, programs and/or instructions stored in the memory 414 and thereby performing processing operations; one or more network or other communications interfaces 404; memory 414; and one or more communication buses 412 for interconnecting these components. In some implementations, the distributed storage system 100 includes an optional user interface 406 comprising a display device 408 and one or more input devices 410 (e.g., keyboard or mouse or remote control). In some implementations, the memory 414 includes high-speed random access memory, such as DRAM, SRAM, DDR RAM or other random access solid state memory devices. In some implementations, memory 414 includes non-volatile memory, such as one or more magnetic disk storage devices, optical disk storage devices, flash memory devices, or other non-volatile solid state storage devices. In some implementations, memory 414 includes one or more storage devices remotely located from the CPU(s) 402. Memory 414, or alternately the non-volatile memory device(s) within memory 414, comprises a non-transitory computer readable storage medium. In some implementations, memory 414 or the non-transitory computer readable storage medium of memory 414 stores the following elements, or a subset of these elements, and may also include additional elements:

- an operating system 416 that includes procedures for handling various basic system services and for performing hardware dependent tasks;

- a network communications module 418 that is used for connecting the distributed storage system 100 to other devices (e.g., the clients 130) via the communication network interfaces 404 and one or more communication networks (wired or wireless), such as the Internet 120, other wide area networks, local area networks, metropolitan area networks, and so on;
- one or more server applications 420, each server application further including a file creation module 422 as described above in connection with FIG. 3A, a reference creation module 424 as described above in connection with FIG. 3B, and a reference deletion module 425 as described above in connection with FIG. 3C; and
- one or more storage devices 426, each storage device including a plurality of tablets 428 and each tablet further including a tablet descriptor file 430, one or more target files 432 and one or more reference files 434.

**[0027]** Reference has been made in detail to implementations, examples of which are illustrated in the accompanying drawings. While particular implementations are described, it will be understood it is not intended to limit the invention to these particular implementations. On the contrary, the invention includes alternatives, modifications and equivalents that are within the spirit and scope of the appended claims. Numerous specific details are set forth in order to provide a thorough understanding of the subject matter presented herein. But it will be apparent to one of ordinary skill in the art that the subject matter may be practiced without these specific details. In other instances, well-known methods, procedures, components, and circuits have not been described in detail so as not to unnecessarily obscure aspects of the implementations.

**[0028]** Although the terms first, second, etc. may be used herein to describe various elements, these elements should not be limited by these terms. These terms are only used to distinguish one element from another. For example, first ranking criteria could be termed second ranking criteria, and, similarly, second ranking criteria could be termed first ranking criteria, without departing from the scope of the present invention. First ranking criteria and second ranking criteria are both ranking criteria, but they are not the same ranking criteria.

**[0029]** The terminology used in the description of the invention herein is for the purpose of describing particular implementations only and is not intended to be limiting of the invention.

As used in the description of the invention and the appended claims, the singular forms “a,” “an,” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will also be understood that the term “and/or” as used herein refers to and encompasses any and all possible combinations of one or more of the associated listed items. It will be further understood that the terms “includes,” “including,” “comprises,” and/or “comprising,” when used in this specification, specify the presence of stated features, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, operations, elements, components, and/or groups thereof.

**[0030]** As used herein, the term “if” may be construed to mean “when” or “upon” or “in response to determining” or “in accordance with a determination” or “in response to detecting,” that a stated condition precedent is true, depending on the context. Similarly, the phrase “if it is determined [that a stated condition precedent is true]” or “if [a stated condition precedent is true]” or “when [a stated condition precedent is true]” may be construed to mean “upon determining” or “in response to determining” or “in accordance with a determination” or “upon detecting” or “in response to detecting” that the stated condition precedent is true, depending on the context.

**[0031]** Although some of the various drawings illustrate a number of logical stages in a particular order, stages that are not order dependent may be reordered and other stages may be combined or broken out. While some reordering or other groupings are specifically mentioned, others will be obvious to those of ordinary skill in the art and so do not present an exhaustive list of alternatives. For example, it is possible for the set top box to send raw audio signal to the TV content recognition server, which is then responsible for converting the audio signal into audio fingerprints. Moreover, it should be recognized that the stages could be implemented in hardware, firmware, software or any combination thereof.

**[0032]** The foregoing description, for purpose of explanation, has been described with reference to specific implementations. However, the illustrative discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The implementations were chosen and described in order to best explain principles of the invention and its practical applications, to thereby enable others skilled in the art to best utilize the invention and various implementations with various modifications as are suited to the particular use contemplated. Implementations include alternatives, modifications and

equivalents that are within the spirit and scope of the appended claims. Numerous specific details are set forth in order to provide a thorough understanding of the subject matter presented herein. But it will be apparent to one of ordinary skill in the art that the subject matter may be practiced without these specific details. In other instances, well-known methods, procedures, components, and circuits have not been described in detail so as not to unnecessarily obscure aspects of the implementations.

## WHAT IS CLAIMED IS:

1. A method of deleting obsolete files from a file system, comprising:
  - at a computer server including one or more processors and memory:
    - receiving a request to delete a reference to a target file in a file system from a file reference data structure, wherein the file reference data structure includes a plurality of target file names and a plurality of reference file names;
    - identifying a reference file name in the file reference data structure, wherein the reference file name includes a file name of the target file in the file system;
    - deleting a reference file from the file system, wherein the reference file has the identified reference file name;
    - checking whether the file system includes at least one reference file whose file name matches the file name of the target file in the file system;
    - if there is no such reference file in the file system:
      - deleting the target file from the file system; and
      - deleting the file name of the target file in the file system from the file reference data structure.
2. The method of claim 1, wherein the file reference data structure is in form of a file stored in the file system.
3. The method of claim 1, wherein the file reference data structure is in form of a data record stored in a database.
4. The method of claim 1, wherein the file reference data structure includes a plurality of reference records, each record including a status indicator of a reference file in the file system, a file name of a target file in the file system referenced by the reference file, and a file name of the reference file in the file system.
5. The method of claim 4, wherein the status indicator of the reference file is one selected from the group consisting of “TENTATIVE”, “LIVE”, and “TO-DELETE”.
6. The method of claim 4, wherein the file name of the reference file includes the file name of the target file and a unique identifier assigned to the reference file.
7. The method of claim 6, wherein the unique indicator assigned to the reference file includes at least one of an IP address, a current timestamp or the like of the computer server.

8. The method of claim 1, wherein a reference file in the file system is an empty file.
9. The method of claim 1, further comprising:
  - receiving a request to create a new target file in the file system, the request including a new target file name and a piece of content;
  - generating a new reference file name using the new target file name;
  - adding the new target file name and the new reference file name to the plurality of target file names and the plurality of reference file names in the file reference data structure, respectively; and
  - creating a new target file and a new reference file in the file system using the new target file name and the new reference file name, respectively, wherein the new target file includes the piece of content.
10. The method of claim 9, wherein the new reference file name includes a unique identifier assigned to the new reference file.
11. The method of claim 10, wherein the unique indicator assigned to the reference file includes at least one of an IP address, a current timestamp or the like of the computer server.
12. The method of claim 9, further comprising:
  - generating a new reference record, the record including a status indicator of the new reference file, the new target file name, and the new reference file name; and
  - adding the new reference record to the file reference data structure.
13. The method of claim 9, wherein the new reference file is an empty file in the file system.
14. The method of claim 1, further comprising:
  - receiving a request to create a new reference to the target file in the file system, the request including the target file name;
  - generating a new reference file name using the target file name;
  - adding the new reference file name to the plurality of reference file names in the file reference data structure; and
  - creating a new reference file in the file system using the new reference file name.
15. The method of claim 14, wherein the new reference file name includes a unique identifier assigned to the new reference file.

16. The method of claim 15, wherein the unique indicator assigned to the new reference file includes at least one of an IP address, a current timestamp or the like of the computer server.
17. The method of claim 14, further comprising:  
generating a new reference record, the record including a status indicator of the new reference file, the target file name, and the new reference file name; and  
adding the new reference record to the file reference data structure.
18. The method of claim 14, wherein the new reference file is an empty file in the file system.
19. A computer server system for deleting obsolete files from a file system, comprising:  
one or more processors;  
memory; and  
a plurality of program modules, wherein the plurality of program modules are stored in the memory and to be executed by the one or more processors, and the plurality of program modules include instructions for:  
receiving a request to delete a reference to a target file in a file system from a file reference data structure, wherein the file reference data structure includes a plurality of target file names and a plurality of reference file names;  
identifying a reference file name in the file reference data structure, wherein the reference file name includes a file name of the target file in the file system;  
deleting a reference file from the file system, wherein the reference file has the identified reference file name;  
checking whether the file system includes at least one reference file whose file name matches the file name of the target file in the file system;  
if there is no such reference file in the file system:  
deleting the target file from the file system; and  
deleting the file name of the target file in the file system from the file reference data structure.
20. A non-transitory computer readable storage medium storing one or more program modules configured for execution by a computer server system that includes one or more



processors and memory for deleting obsolete files from a file system, the one or more program modules comprising instructions for:

- receiving a request to delete a reference to a target file in a file system from a file reference data structure, wherein the file reference data structure includes a plurality of target file names and a plurality of reference file names;

- identifying a reference file name in the file reference data structure, wherein the reference file name includes a file name of the target file in the file system;

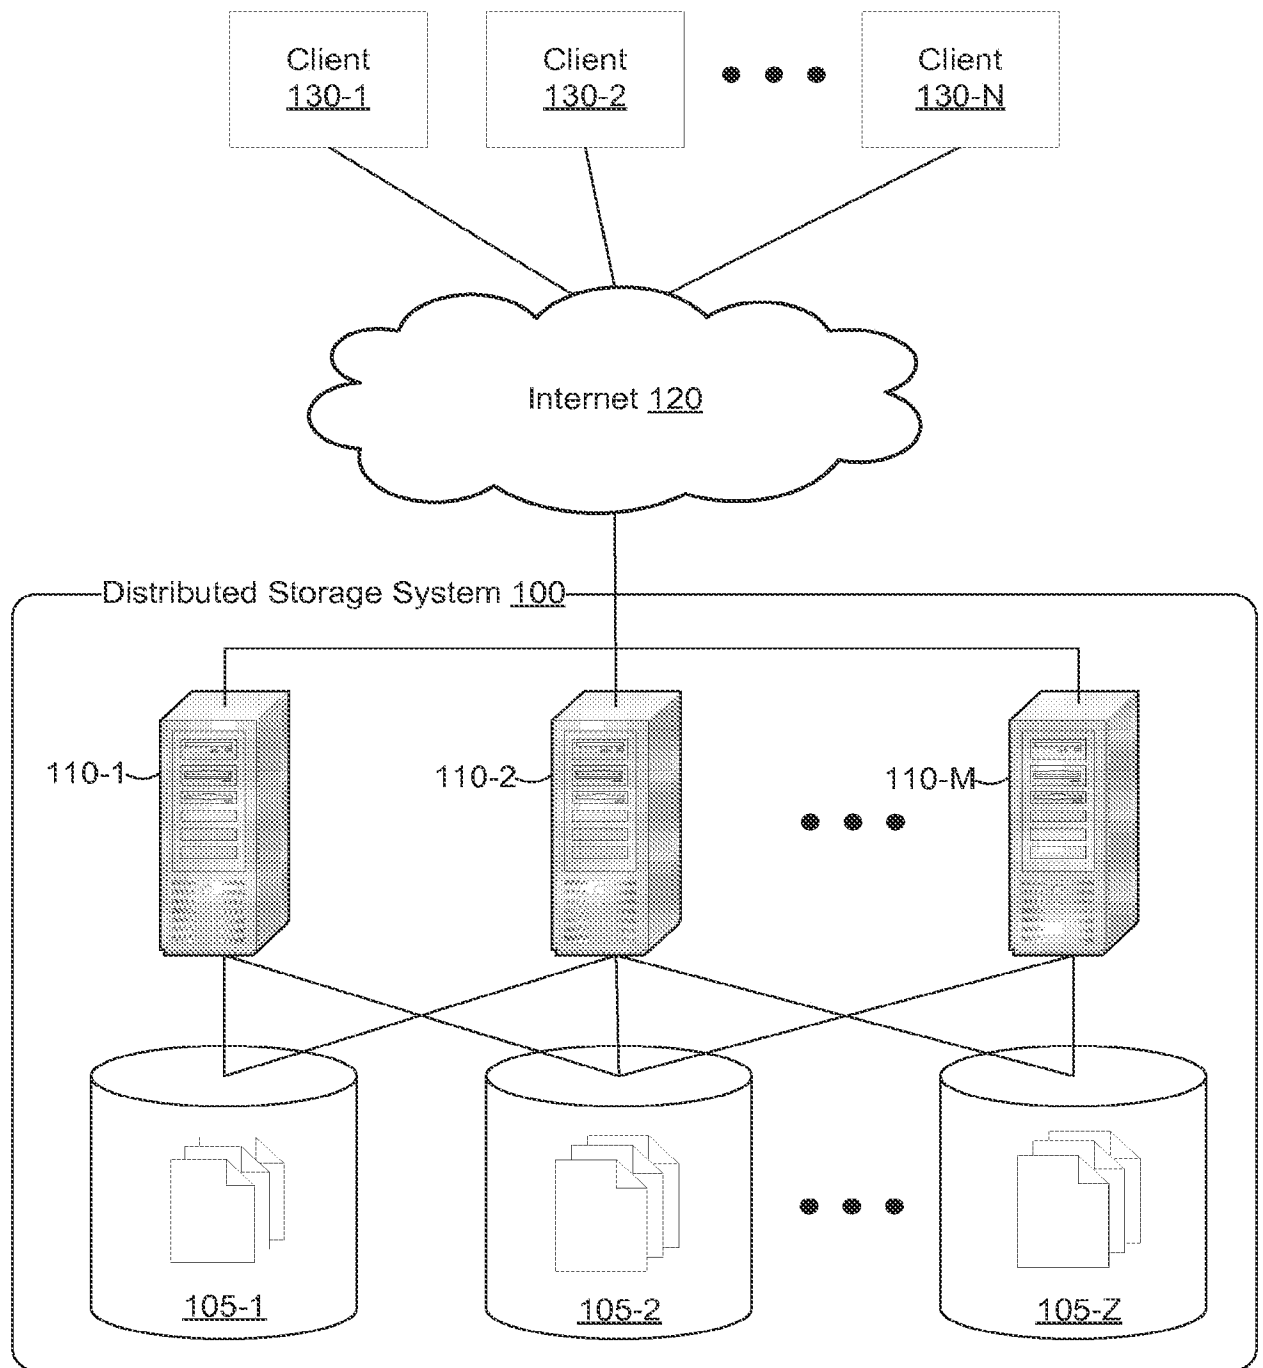
- deleting a reference file from the file system, wherein the reference file has the identified reference file name;

- checking whether the file system includes at least one reference file whose file name matches the file name of the target file in the file system;

- if there is no such reference file in the file system:

- deleting the target file from the file system; and

- deleting the file name of the target file in the file system from the file reference data structure.



**FIG. 1**

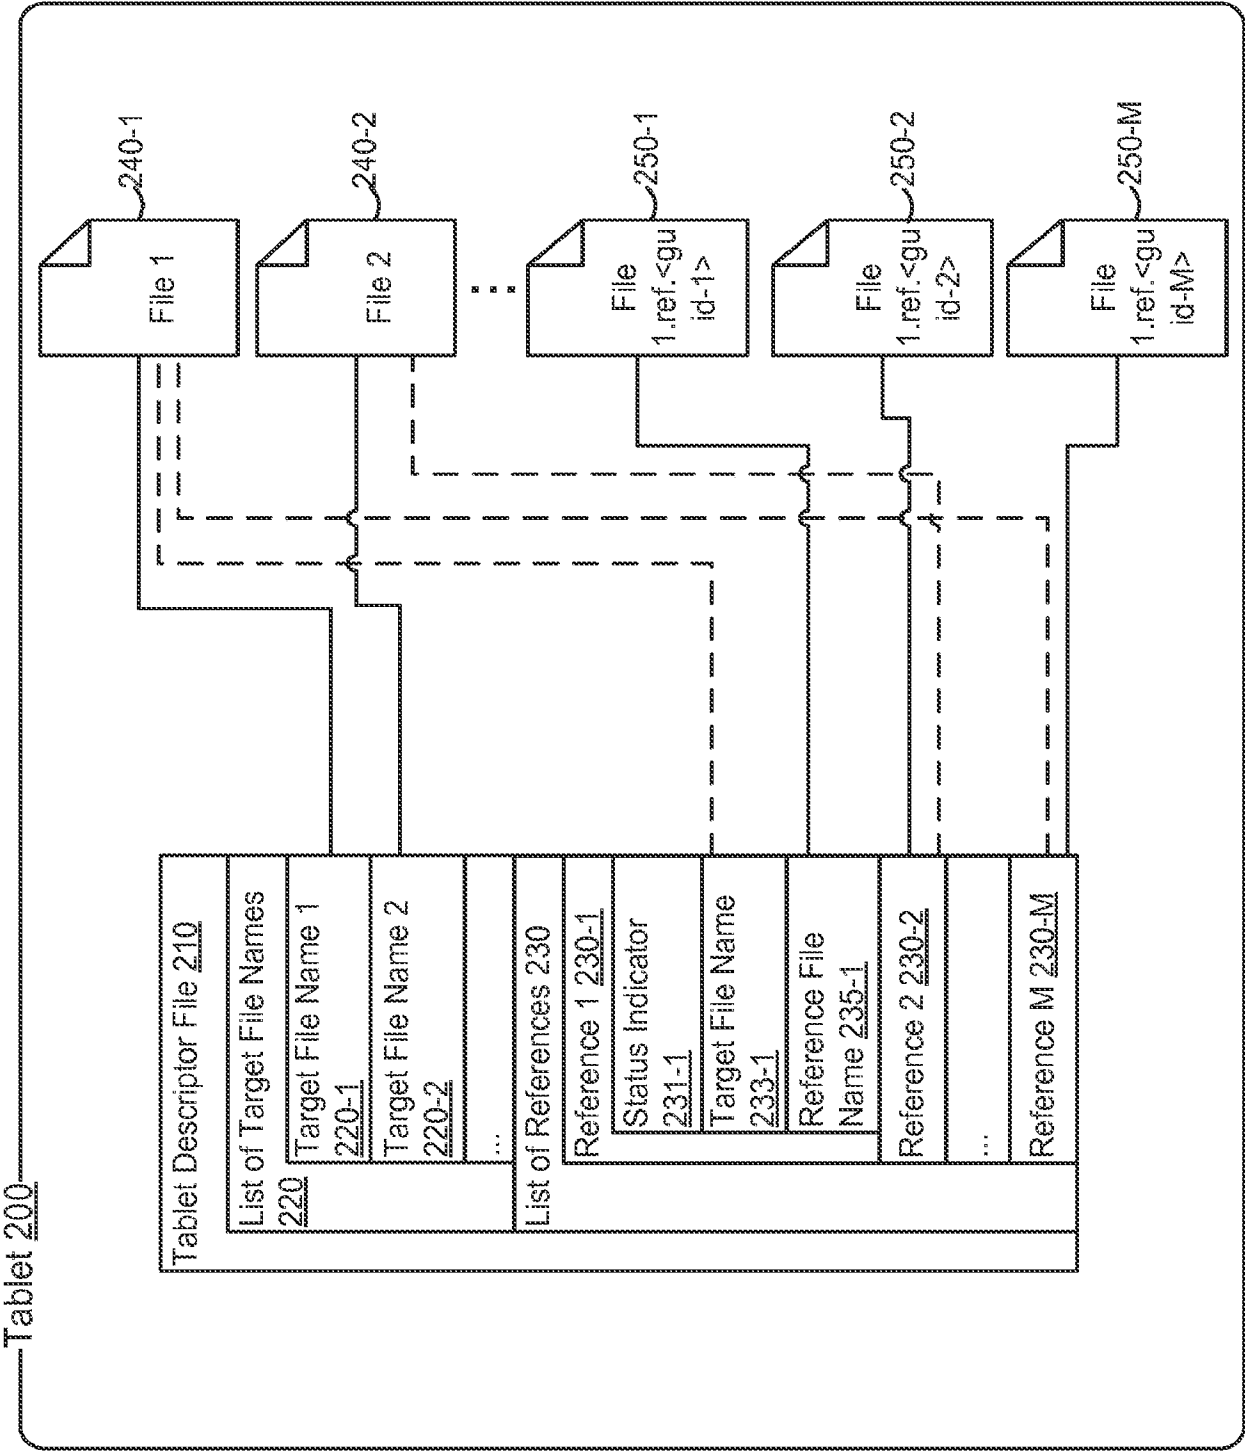
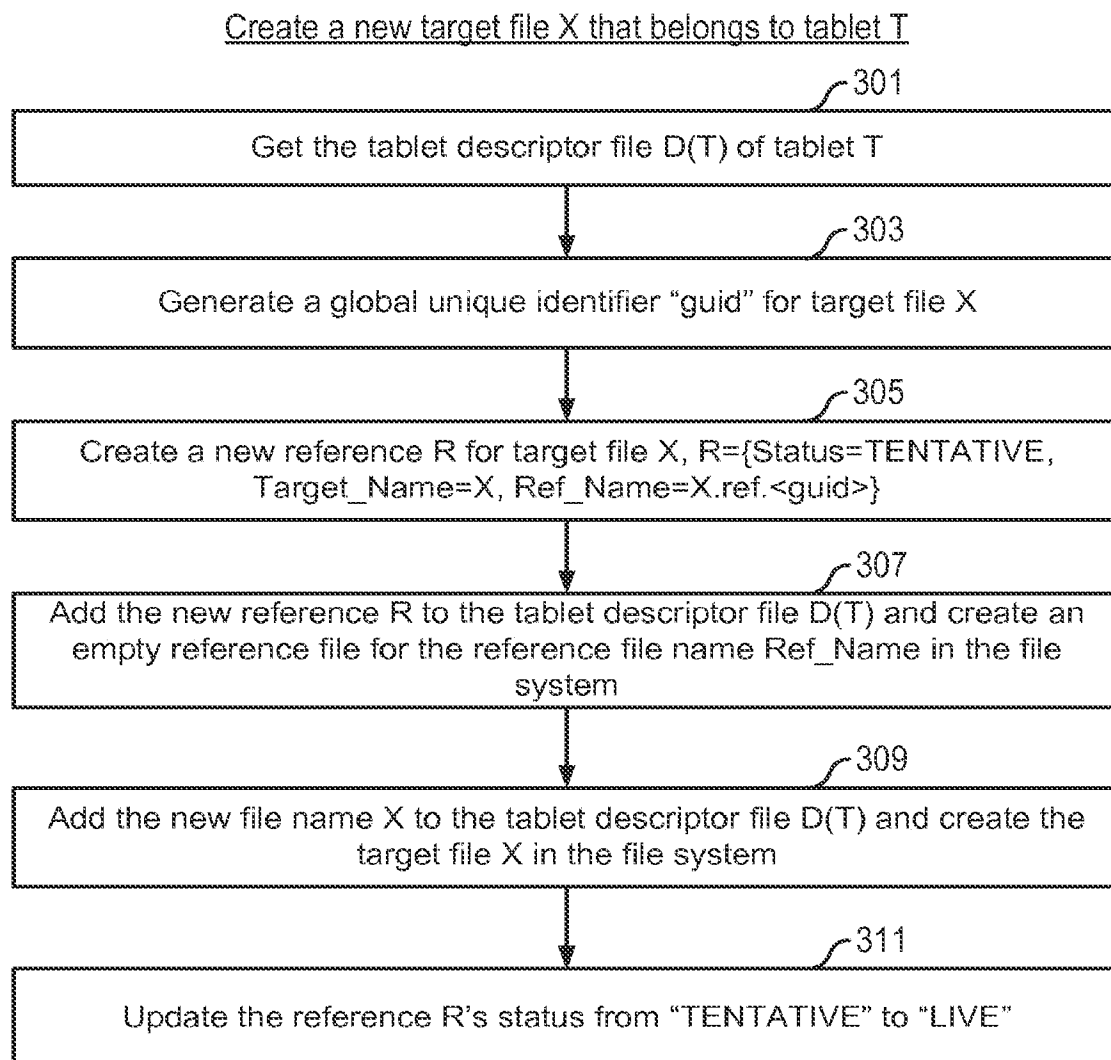
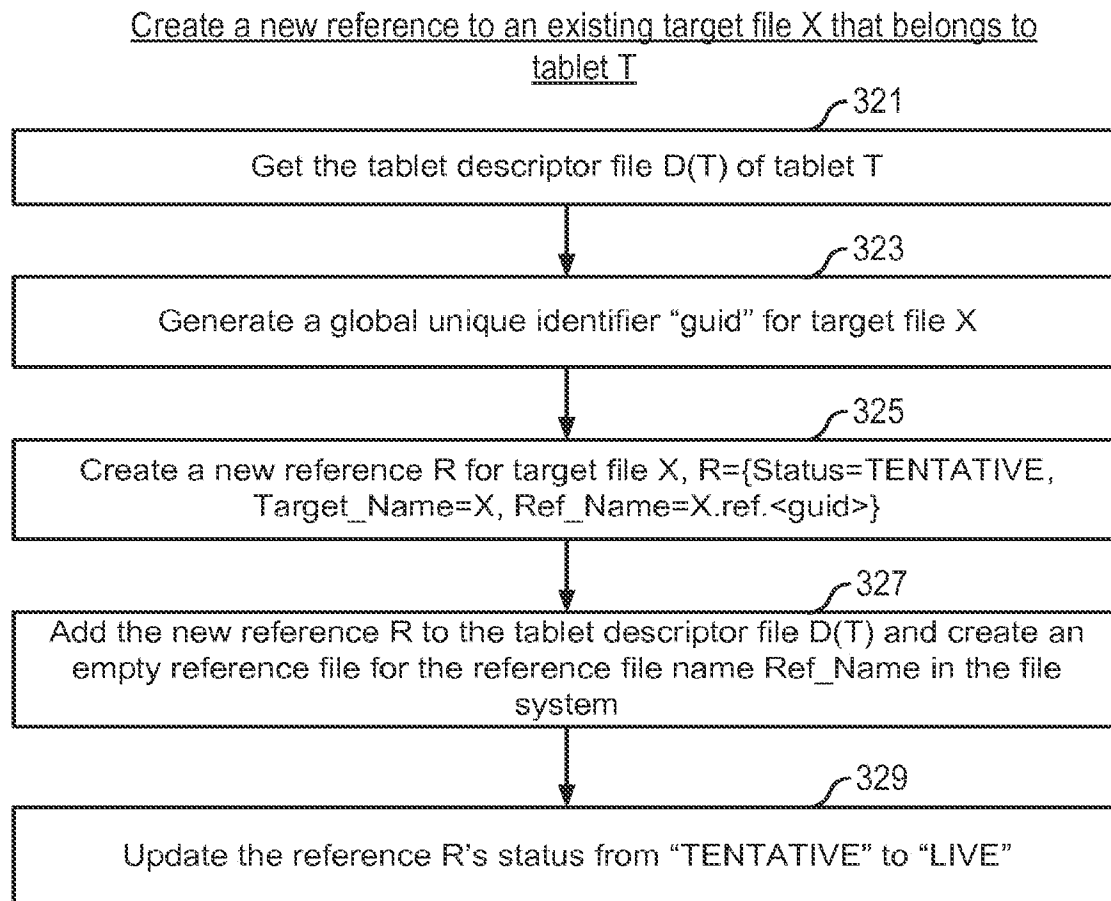
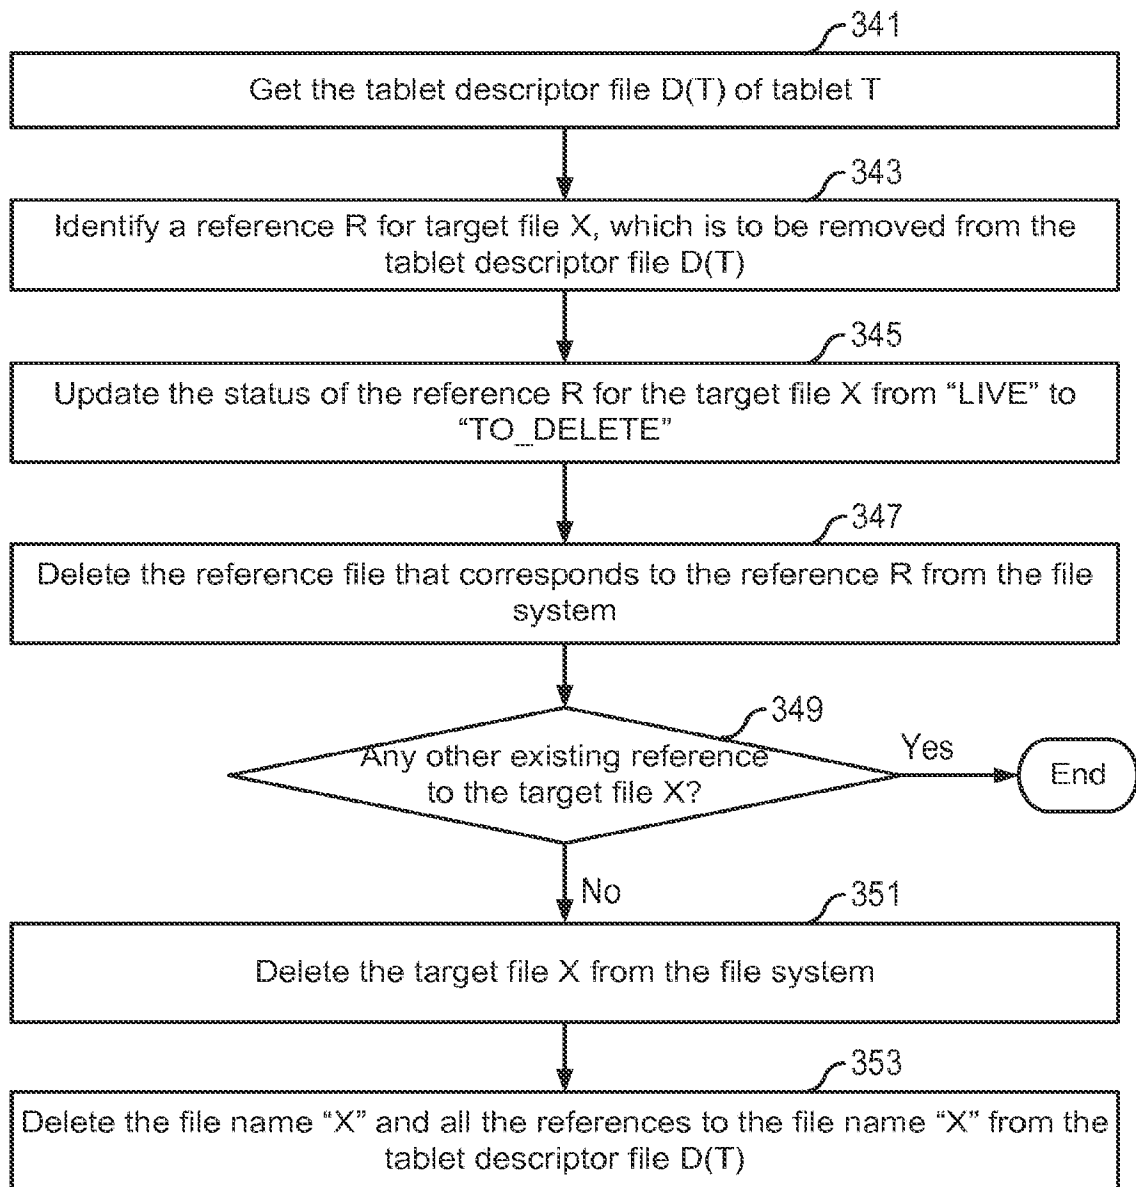


FIG. 2

**FIG. 3A**

**FIG. 3B**

Remove a reference to an existing target file X that belongs to tablet T



**FIG. 3C**

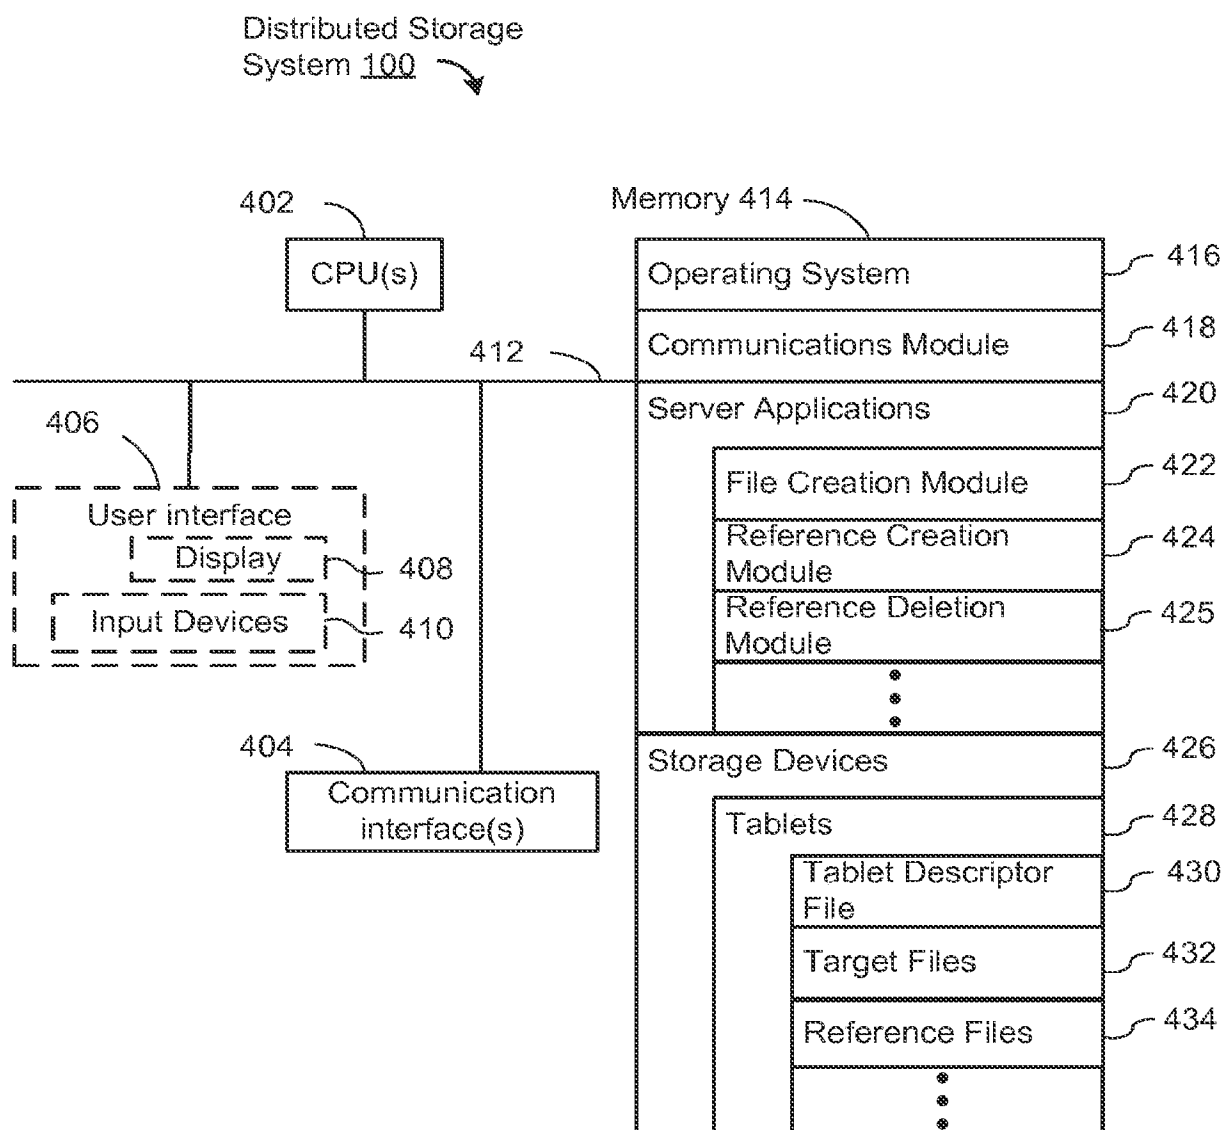


FIG. 4

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/US2013/044105

A. CLASSIFICATION OF SUBJECT MATTER  
INV. G06F17/30  
ADD.

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 practicable, search terms used)

EPO-Internal, WPI Data

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6 477 544 B1 (BOLOSKY WILLIAM J [US] ET AL) 5 November 2002 (2002-11-05) page 1 - page 14	1-20
X	US 2008/243879 A1 (GOKHALE PARAG [US] ET AL) 2 October 2008 (2008-10-02) page 1 - page 5	1-20



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 application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

7 November 2013

Date of mailing of the international search report

14/11/2013

Name and mailing address of the ISA/

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

Authorized officer

Korkuzas, Valdas



# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2013/044105

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6477544	B1	05-11-2002	NONE
-----			
US 2008243879	A1	02-10-2008	US 2008243879 A1 02-10-2008
			US 2008243914 A1 02-10-2008
			US 2008243957 A1 02-10-2008
			US 2008243958 A1 02-10-2008
			US 2012023070 A1 26-01-2012
			US 2013006946 A1 03-01-2013
-----			