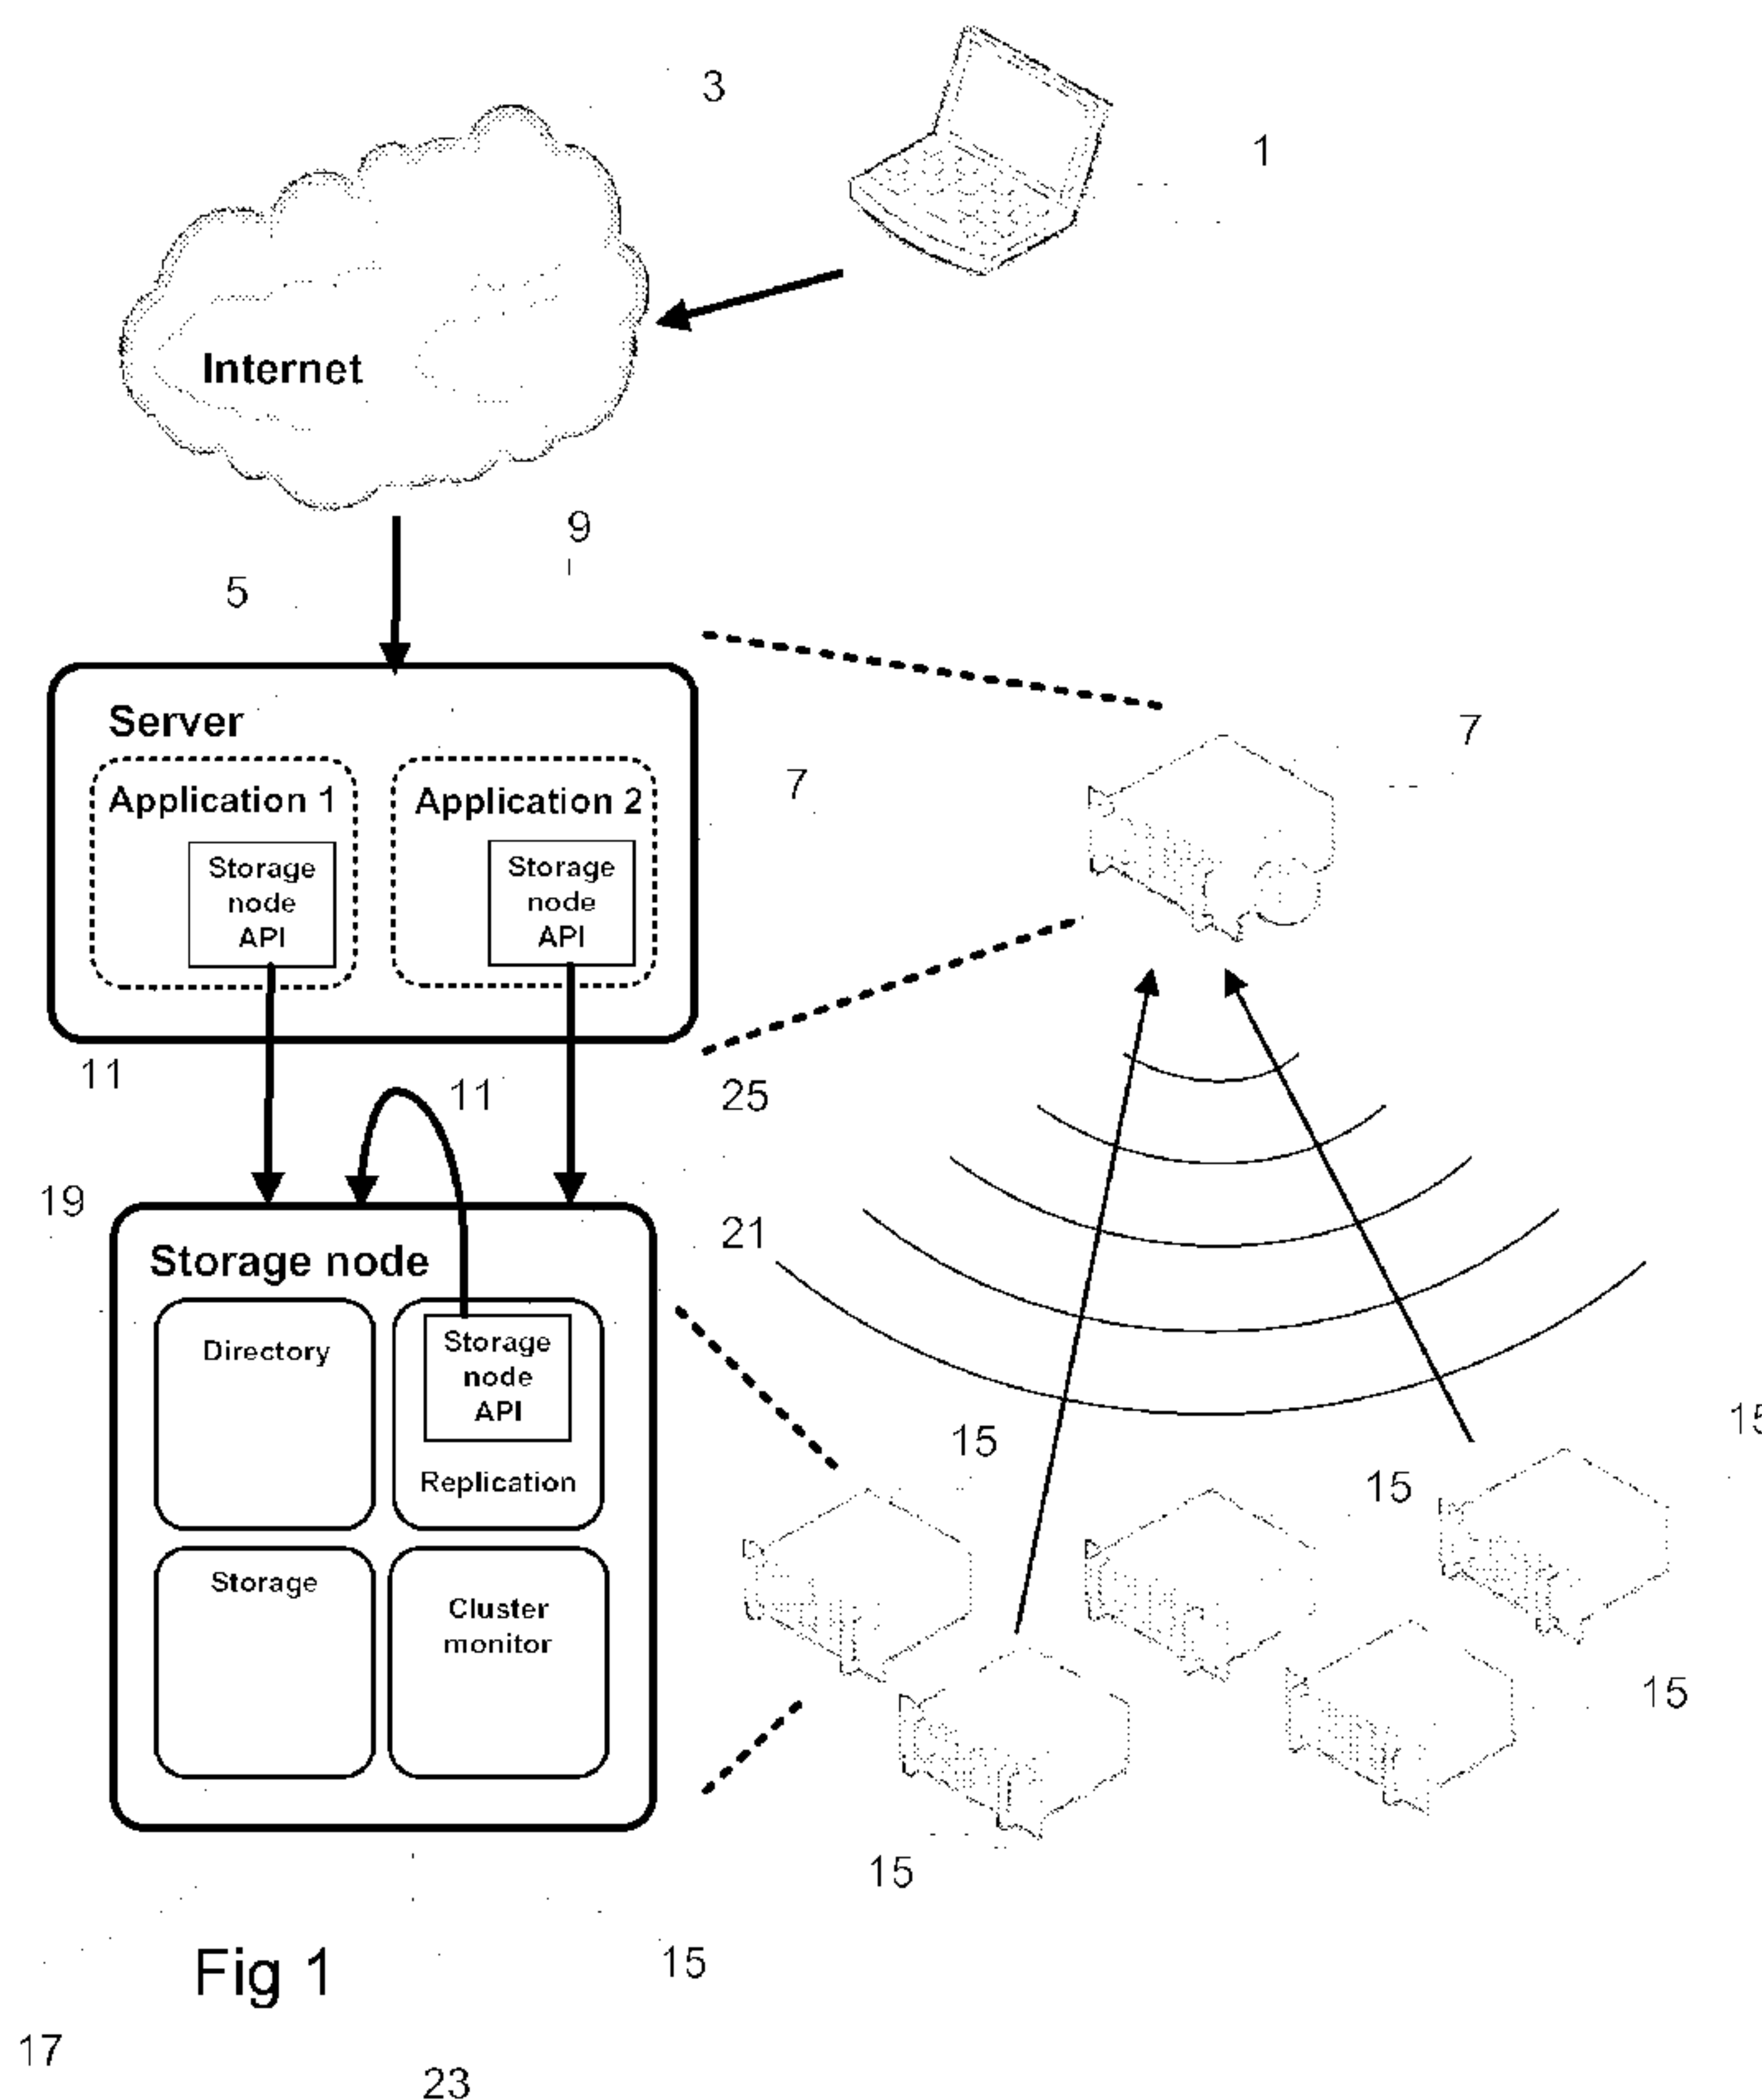




(86) Date de dépôt PCT/PCT Filing Date: 2009/10/21
 (87) Date publication PCT/PCT Publication Date: 2010/04/29
 (85) Entrée phase nationale/National Entry: 2011/04/21
 (86) N° demande PCT/PCT Application No.: EP 2009/063796
 (87) N° publication PCT/PCT Publication No.: 2010/046393
 (30) Priorité/Priority: 2008/10/24 (SE0802277-4)

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01)
 (71) Demandeur/Applicant:
 ILT PRODUCTIONS AB, SE
 (72) Inventeurs/Inventors:
 MELANDER, CHRISTIAN, SE;
 BERNBO, STEFAN, SE;
 PETERSSON, GUSTAV, SE;
 PERSSON, ROGER, SE
 (74) Agent: BORDEN LADNER GERVAIS LLP

(54) Titre : STOCKAGE DE DONNEES DISTRIBUE
 (54) Title: DISTRIBUTED DATA STORAGE



(57) Abrégé/Abstract:

The present invention relates to a distributed data storage system comprising a plurality of storage nodes. Using unicast and multicast transmission, a server application may read and write data in the storage system. Each storage node may monitor reading

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

and writing operations on the system as well as the status of other storage nodes. In this way, the storage nodes may detect a need for replication of files on the system, and may carry out a replication process that serves to maintain a storage of a sufficient number of copies of files with correct versions at geographically diverse locations.

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

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
29 April 2010 (29.04.2010)(10) International Publication Number
WO 2010/046393 A2(51) International Patent Classification:
G06F 17/30 (2006.01)na (SE). **PERSSON, Roger** [SE/SE]; Nyhemsvägen 1 C,
S-SE-371 43 Karlskrona (SE).(21) International Application Number:
PCT/EP2009/063796(74) Agent: **HJALMARSSON, Magnus**; Awapatent AB,
P.O. Box 99, S-351 04 Växjö (SE).(22) International Filing Date:
21 October 2009 (21.10.2009)(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ,
CA, CH, 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, KM, KN, KP,
KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD,
ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI,
NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD,
SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT,
TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(25) Filing Language: English

(26) Publication Language: English


(30) Priority Data:
0802277-4 24 October 2008 (24.10.2008) SE(71) Applicant (for all designated States except US): **ILT
Productions AB** [SE/SE]; P.O. Box 177, S-SE-371 22
Karlskrona (SE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **MELANDER,
Christian** [SE/SE]; Selenvägen 3, S-SE-370 30 Rödeby
(SE). **BERNBO, Stefan** [SE/SE]; Arklimästaregatan 46
A, S-SE-371 36 Karlskrona (SE). **PETERSSON, Gustav**
[SE/SE]; Östra Köpmansgatan 29, S-SE-371 32 Karlskro-(84) Designated States (unless otherwise indicated, for every
kind of regional protection available): ARIPO (BW, GH,
GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM,
ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ,
TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE,
ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM,

[Continued on next page]

(54) Title: DISTRIBUTED DATA STORAGE

(57) Abstract: The present invention relates to a distributed
data storage system comprising a plurality of storage nodes.
Using unicast and multicast transmission, a server applica-
tion may read and write data in the storage system. Each
storage node may monitor reading and writing operations
on the system as well as the status of other storage nodes.
In this way, the storage nodes may detect a need for replica-
tion of files on the system, and may carry out a replica-
tion process that serves to maintain a storage of a sufficient
number of copies of files with correct versions at geograph-
ically diverse locations.

WO 2010/046393 A2 

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

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

DISTRIBUTED DATA STORAGE

Technical field

The present disclosure relates to methods for writing and maintaining data in a data storage system comprising a plurality of data storage nodes, the methods being employed in a server and in a storage node in the data storage system. The disclosure further relates to storage nodes or servers capable of carrying out such methods.

Background

Such a method is disclosed e.g. in US, 2005/0246393, A1. This method is disclosed for a system that uses a plurality of storage centres at geographically disparate locations. Distributed object storage managers are included to maintain information regarding stored data.

One problem associated with such a system is how to accomplish simple and yet robust and reliable writing as well as maintenance of data.

Summary of the invention

One object of the present disclosure is therefore to realise robust writing or maintenance of data in a distributed storage system, without the use of centralised maintenance servers, which may themselves be a weak link in a system. This object is achieved by means a method of the initially mentioned kind which is accomplished in a storage node and comprises: monitoring the status of other storage nodes in the system as well as writing operations carried out in the data storage system, detecting, based on the monitoring, conditions in the data storage system that imply the need for replication of data between the nodes in the data storage system, and initiating a replication process in case such a condition is detected. The replication process includes sending a multicast message, to a plurality of storage nodes, the message enquiring which of those storage nodes store specific data.

By means of such a method, each storage node can be active with maintaining data of the entire system. In case a storage node fails, its data

can be recovered by other nodes in the system, which system may therefore be considered to be self-healing.

The monitoring may include listening to heartbeat signals from other storage nodes in the system. A condition that implies need for replication is then a malfunctioning storage node.

The data includes files and a condition implying need for replications may then be one of a file deletion or a file inconsistency.

A replication list, including files that need replication, may be maintained and may include priorities.

The replication process may include: sending a multicast message to a plurality of storage nodes request enquiring which of those storage nodes store specific data, receiving responses from those storage nodes that contain said specific data, determining whether said specific data is stored on a sufficient number of storage nodes, and, if not, selecting at least one additional storage node and transmitting said specific data to that storage node. Further, the specific data on storage nodes containing obsolete versions thereof may be updated.

Additionally, the replication process may begin with the storage node attempting to attain mastership for the file, to be replicated, among all storage nodes in the system.

The monitoring may further include monitoring of reading operations carried out in the data storage system.

The present disclosure further relates to a data storage node, for carrying out maintenance of data, corresponding to the method. The storage node then generally comprises means for carrying out the actions of the method.

The object is also achieved by means of a method for writing data to a data storage system of the initially mentioned kind, which is accomplished in a server running an application which accesses data in the data storage system. The method comprises: sending a multicast storage query to a plurality of storage nodes, receiving a plurality of responses from a subset of said storage nodes, the responses including geographic data relating to the geographic position of each server, selecting at least two storage nodes in the

subset, based on said responses, and sending data and a data identifier, corresponding to the data, to the selected storage nodes.

This method accomplishes robust writing of data in the sense that a geographic diversity is realised in an efficient way.

The geographic position may include latitude and longitude of the storage node in question, and the responses may further include system load and/or system age for the storage node in question.

The multicast storage query may include a data identifier, identifying the data to be stored.

Typically, at least three nodes may be selected for storage, and a list of storage nodes, successfully storing the data, may be sent to the selected storage nodes.

The present disclosure further relates to a server, for carrying out writing of data, corresponding to the method. The server then generally comprises means for carrying out the actions of the method.

Brief description of the drawings

Fig 1 illustrates a distributed data storage system.

Figs 2A –2C, and fig 3 illustrate a data reading process.

Figs 4A –4C, and fig 5 illustrate a data writing process.

Fig 6 illustrates schematically a situation where a number of files are stored among a number of data storage nodes.

Fig 7 illustrates the transmission of heartbeat signals.

Fig 8 is an overview of a data maintenance process.

Detailed description

The present disclosure is related to a distributed data storage system comprising a plurality of storage nodes. The structure of the system and the context in which it is used is outlined in fig 1.

A user computer 1 accesses, via the Internet 3, an application 5 running on a server 7. The user context, as illustrated here, is therefore a regular client-server configuration, which is well known per se. However, it should be noted that the data storage system to be disclosed may be useful

also in other configurations.

In the illustrated case, two applications 5, 9 run on the server 7. Of course however, this number of applications may be different. Each application has an API (Application Programming Interface) 11 which provides an interface in relation to the distributed data storage system 13 and supports requests, typically write and read requests, from the applications running on the server. From an application's point of view, reading or writing information from/to the data storage system 13 need not appear different from using any other type of storage solution, for instance a file server or simply a hard drive.

Each API 11 communicates with storage nodes 15 in the data storage system 13, and the storage nodes communicate with each other. These communications are based on TCP (Transmission Control Protocol) and UDP (User Datagram Protocol). These concepts are well known to the skilled person, and are not explained further.

It should be noted that different APIs 11 on the same server 7 may access different sets of storage nodes 15. It should further be noted that there may exist more than one server 7 which accesses each storage node 15. This, however does not to any greater extent affect the way in which the storage nodes operate, as will be described later.

The components of the distributed data storage system are the storage nodes 15 and the APIs 11, in the server 7 which access the storage nodes 15. The present disclosure therefore relates to methods carried out in the server 7 and in the storage nodes 15. Those methods will primarily be embodied as software implementations which are run on the server and the storage nodes, respectively, and are together determining for the operation and the properties of the overall distributed data storage system.

The storage node 15 may typically be embodied by a file server which is provided with a number of functional blocks. The storage node may thus comprise a storage medium 17, which typically comprises of a number of hard drives, optionally configured as a RAID (Redundant Array of Independent Disk) system. Other types of storage media are however conceivable as well.

The storage node 15 may further include a directory 19, which comprises lists of data entity/storage node relations as a host list, as will be

discussed later.

In addition to the host list, each storage node further contains a node list including the IP addresses of all storage nodes in its set or group of storage nodes. The number of storage nodes in a group may vary from a few to hundreds of storage nodes. The node list may further have a version number.

Additionally, the storage node 15 may include a replication block 21 and a cluster monitor block 23. The replication block 21 includes a storage node API 25, and is configured to execute functions for identifying the need for and carrying out a replication process, as will be described in detail later. The storage node API 25 of the replication block 21 may contain code that to a great extent corresponds to the code of the server's 7 storage node API 11, as the replication process comprises actions that correspond to a great extent to the actions carried out by the server 7 during reading and writing operations to be described. For instance, the writing operation carried out during replication corresponds to a great extent to the writing operation carried out by the server 7. The cluster monitor block 23 is configured to carry out monitoring of other storage nodes in the data storage system 13, as will be described in more detail later.

The storage nodes 15 of the distributed data storage system can be considered to exist in the same hierarchical level. There is no need to appoint any master storage node that is responsible for maintaining a directory of stored data entities and monitoring data consistency, etc. Instead, all storage nodes 15 can be considered equal, and may, at times, carry out data management operations vis-à-vis other storage nodes in the system. This equality ensures that the system is robust. In case of a storage node malfunction other nodes in the system will cover up the malfunctioning node and ensure reliable data storage.

The operation of the system will be described in the following order: reading of data, writing of data, and data maintenance. Even though these methods work very well together, it should be noted that they may in principle also be carried out independently of each other. That is, for instance the data reading method may provide excellent properties even if the data writing

method of the present disclosure is not used, and vice versa.

The reading method is now described with reference to figs 2A-2C and 3, the latter being a flowchart illustrating the method.

The reading, as well as other functions in the system, utilise multicast communication to communicate simultaneously with a plurality of storage nodes. By a multicast or IP multicast is here meant a point-to-multipoint communication which is accomplished by sending a message to an IP address which is reserved for multicast applications.

For example, a message, typically a request, is sent to such an IP address (e.g. 244.0.0.1), and a number of recipient servers are registered as subscribers to that IP address. Each of the recipient servers has its own IP address. When a switch in the network receives the message directed to 244.0.0.1, the switch forwards the message to the IP addresses of each server registered as a subscriber.

In principle, only one server may be registered as a subscriber to a multicast address, in which case a point-to-point, communication is achieved. However, in the context of this disclosure, such a communication is nevertheless considered a multicast communication since a multicast scheme is employed.

Unicast communication is also employed referring to a communication with a single recipient.

With reference to fig 2A and fig 3, the method for retrieving data from a data storage system comprises the sending 31 of a multicast query to a plurality of storage nodes 15. In the illustrated case there are five storage nodes each having an IP (Internet Protocol) address 192.168.1.1, 192.168.1.2, etc. The number of storage nodes is, needless to say, just an example. The query contains a data identifier "2B9B4A97-76E5-499E-A21A6D7932DD7927", which may for instance be a Universally Unique Identifier, UUID, which is well known per se.

The storage nodes scan themselves for data corresponding to the identifier. If such data is found, a storage node sends a response, which is received 33 by the server 7, cf. fig 2B. As illustrated, the response may optionally contain further information in addition to an indication that the

storage node has a copy of the relevant data. Specifically, the response may contain information from the storage node directory about other storage nodes containing the data, information regarding which version of the data is contained in the storage node, and information regarding which load the storage node at present is exposed to.

Based on the responses, the server selects 35 one or more storage nodes from which data is to be retrieved, and sends 37 a unicast request for data to that/those storage nodes, cf. fig 2C.

In response to the request for data, the storage node/nodes send the relevant data by unicast to the server which receives 39 the data. In the illustrated case, only one storage node is selected. While this is sufficient, it is possible to select more than one storage node in order to receive two sets of data which makes a consistency check possible. If the transfer of data fails, the server may select another storage node for retrieval.

The selection of storage nodes may be based on an algorithm that take several factors into account in order to achieve a good overall system performance. Typically, the storage node having the latest data version and the lowest load will be selected although other concepts are fully conceivable.

Optionally, the operation may be concluded by server sending a list to all storage nodes involved, indicating which nodes contains the data and with which version. Based on this information, the storage nodes may themselves maintain the data properly by the replication process to be described.

Figs 4A –4C, and fig 5 illustrate a data writing process for the distributed data storage system.

With reference to fig 4A and fig 5 the method comprises a server sending 41 a multicast storage query to a plurality of storage nodes. The storage query comprises a data identifier and basically consists of a question whether the receiving storage nodes can store this file. Optionally, the storage nodes may check with their internal directories whether they already have a file with this name, and may notify the server 7 in the unlikely event that this is the case, such that the server may rename the file.

In any case, at least a subset of the storage nodes will provide responses by unicast transmission to the server 7. Typically, storage nodes

having a predetermined minimum free disk space will answer to the query. The server 7 receives 43 the responses which include geographic data relating to the geographic position of each server. For instance, as indicated in fig 4B, the geographic data may include the latitude, the longitude and the altitude of each server. Other types of geographic data may however also be conceivable, such as a ZIP code or the like.

In addition to the geographic data, further information may be provided that serves as an input to a storage node selection process. In the illustrated example, the amount of free space in each storage node is provided together with an indication of the storage node's system age and an indication of the load that the storage node currently experiences.

Based on the received responses, the server selects 45 at least two, in a typical embodiment three, storage nodes in the subset for storing the data. The selection of storage nodes is carried out by means of an algorithm that take different data into account. The selection is carried out in order to achieve some kind of geographical diversity. At least it should preferably be avoided that only file servers in the same rack are selected as storage nodes. Typically, a great geographical diversity may be achieved, even selecting storage nodes on different continents. In addition to the geographical diversity, other parameters may be included in the selection algorithm. As long as a minimum geographic diversity is achieved, free space, system age and current load may also be taken into account.

When the storage nodes have been selected, the data to be stored and a corresponding data identifier is sent to each selected node, typically using a unicast transmission.

Optionally, the operation may be concluded by each storage node, which has successfully carried out the writing operation, sending an acknowledgement to the server. The server then sends a list to all storage nodes involved indicating which nodes have successfully written the data and which have not. Based on this information, the storage nodes may themselves maintain the data properly by the replication process to be described. For instance if one storage node's writing failed, there exists a need to replicate the file to one more storage node in order to achieve the desired number of

storing storage nodes for that file.

The data writing method in itself allows an API in a server 7 to store data in a very robust way, as excellent geographic diversity may be provided.

In addition to the writing and reading operations, the API in the server 7 may carry out operations that delete files and update files. These processes will be described in connection with the data maintenance process below.

The aim of the data maintenance process is to make sure that a reasonable number of non-malfunctioning storage nodes each store the latest version of each file. Additionally, it may provide the function that no deleted files are stored at any storage node. The maintenance is carried out by the storage nodes themselves. There is thus no need for a dedicated "master" that takes responsibility for the maintenance of the data storage. This ensures improved reliability as the "master" would otherwise be a weak spot in the system.

Fig 6 illustrates schematically a situation where a number of files are stored among a number of data storage nodes. In the illustrated case, twelve nodes, having IP addresses consecutively numbered from 192.168.1.1 to 192.168.1.12, are depicted for illustration purposes. Needless to say however, the IP address numbers need not be in the same range at all. The nodes are placed in a circular order only to simplify the description, i.e. the nodes need not have any particular order. Each node store one or two files identified, for the purpose of simplicity, by the letters A-F.

With reference to fig 8, the method for maintaining data comprises the detecting 51 conditions in the data storage system that imply the need for replication of data between the nodes in the data storage system, and a replication process 53. The result of the detection process 51 is a list 55 of files for which the need for replication has been identified. The list may further include data regarding the priority of the different needs for replication. Based on this list the replication process 53 is carried out.

The robustness of the distributed storage relies on that a reasonable number of copies of each file, correct versions, are stored in the system. In the illustrated case, three copies of each file is stored. However, should for instance the storage node with the address 192.168.1.5 fail, the desired

number of stored copies for files “B” and “C” will be fallen short of.

One event that results in a need for replication is therefore the malfunctioning of a storage node in the system.

Each storage node in the system may monitor the status of other storage nodes in the system. This may be carried out by letting each storage node emit a so-called heartbeat signal at regular intervals, as illustrated in fig 7. In the illustrated case, the storage node with address 192.168.1.7 emits a multicast signal 57 to the other storage nodes in the system, indicating that it is working correctly. This signal may be received by all other functioning storage nodes in the system carrying out heartbeat monitoring 59 (cf. fig 8), or a subset thereof. In the case with the storage node with address 192.168.1.5 however, this node is malfunctioning and does not emit any heartbeat signal. Therefore, the other storage nodes will notice that no heartbeat signal has been emitted by this node in a long time which indicates that the storage node in question is down.

The heartbeat signal may, in addition to the storage node’s address, include its node list version number. Another storage node, listening to the heartbeat signal and finding out that the transmitting storage node has a later version node list, may then request that transmitting storage node to transfer its node list. This means that addition and removal of storage nodes can be obtained simply by adding or removing a storage node and sending a new node list version to one single storage node. This node list will then spread to all other storage nodes in the system.

Again with reference to fig 8, each storage node searches 61 its internal directory for files that are stored by the malfunctioning storage node. Storage nodes which themselves store files “B” and “C” will find the malfunctioning storage node and can therefore add the corresponding file on their lists 55.

The detection process may however also reveal other conditions that imply the need for replicating a file. Typically such conditions may be inconsistencies, i.e. that one or more storage nodes has an obsolete version of the file. A delete operation also implies a replication process as this process may carry out the actual physical deletion of the file. The server’s delete operation

then only need make sure that the storage nodes set a deletion flag for the file in question. Each node may therefore monitor reading and writing operations carried out in the data storage system. Information provided by the server 7 at the conclusion of reading and writing operations, respectively, may indicate that one storage node contains an obsolete version of a file (in the case of a reading operation) or that a storage node did not successfully carry out a writing operation. In both cases there exists a need for maintaining data by replication such that the overall objects of the maintenance process are fulfilled.

In addition to the basic reading and writing operations 63, 65, at least two additional processes may provide indications that a need for replication exists, namely the deleting 67 and updating 69 processes that are now given a brief explanation.

The deleting process is initiated by the server 7 (cf. fig 1). Similar to the reading process, the server sends a query by multicasting to all storage nodes, in order to find out which storage nodes has data with a specific data identifier. The storage nodes scan themselves for data with the relevant identifier, and respond by a unicast transmission if they have the data in question. The response may include a list, from the storage node directory, of other storage nodes containing the data. The server 7 then sends a unicast request, to the storage nodes that are considered to store the file, that the file be deleted. Each storage node sets a flag relating to the file and indicating that it should be deleted. The file is then added to the replication list, and an acknowledgement is sent to the server. The replication process then physically deletes the file as will be described.

The updating process has a search function, similar to the one of the deleting process, and a writing function, similar to the one carried out in the writing process. The server sends a query by multicasting to all storage nodes, in order to find out which storage nodes has data with a specific data identifier. The storage nodes scan themselves for data with the relevant identifier, and respond by a unicast transmission if they have the data in question. The response may include a list, from the storage node directory, of other storage nodes containing the data. The server 7 then sends a unicast

request, telling the storage nodes to update the data. The request of course contains the updated data. The storage nodes updating the data sends an acknowledgement to the server, which responds by sending a unicast transmission containing a list with the storage nodes that successfully updated the data, and the storage nodes which did not. Again, this list can be used by the maintenance process.

Again with reference to fig 8 the read 63, write 65, delete 67, and update 69 operations may all indicate that a need for replication exists. The same applies for the heartbeat monitoring 59. The overall detection process 51 thus generates data regarding which files need be replicated. For instance, a reading or updating operation may reveal that a specific storage node contains an obsolete version of a file. A deletion process may set a deletion flag for a specific file. The heartbeat monitoring may reveal that a number of files, stored on a malfunctioning storage node need be replicated to a new storage node.

Each storage nodes monitors the need for replication for all the files it stores and maintains a replication list 55. The replication list 55 thus contains a number of files that need be replicated. The files may be ordered in correspondence with the priority for each replication. Typically, there may be three different priority levels. The highest level is reserved for files which the storage node holds the last online copy of. Such a file need be quickly replicated to other storage nodes such that a reasonable level of redundancy may be achieved. A medium level of priority may relate to files where the versions are inconsistent among the storage nodes. A lower level of priority may relate to files which are stored on a storage node that is malfunctioning.

The storage node deals with the files on the replication list 55 in accordance with their level of priority. The replication process is now described for a storage node which is here called the operating storage node, although all storage nodes may operate in this way.

The replication part 53 of the maintaining process starts with the operating storage node attempting 71 to become the master for the file it intends to replicate. The operating storage nodes sends a unicast request to become master to other storage nodes that are known store the file in

question. The directory 19 (cf. fig 1) provides a host list comprising information regarding which storage nodes to ask. In the event, for instance in case of a colliding request, that one of the storage nodes does not respond affirmatively, the file is moved back to the list for the time being, and an attempt is instead made with the next file on the list. Otherwise the operating storage node is considered to be the master of this file and the other storage nodes set a flag indicating that the operating storage node is master for the file in question.

The next step is to find 73 all copies of the file in question in the distributed storage system. This may be carried out by the operating storage node sending a multicast query to all storage nodes, asking which ones of them have the file. The storage nodes having the file submit responses to the query, containing the version of the file they keep as well as their host lists, i.e. the list of storage nodes containing the relevant file that is kept in the directory of each storage node. These host lists are then merged 75 by the operating storage node, such that a master host list is formed corresponding to the union of all retrieved host lists. If additional storage nodes are found, which were not asked when the operating storage node attempted to become master, that step may now be repeated for the additional storage nodes. The master host list contains information regarding which versions of the file the different storage nodes keep and illustrate the status of the file within the entire storage system.

Should the operating storage node not have the latest version of the file in question, this file is then retrieved 77 from one of the storage nodes that do have the latest version.

The operating storage node then decides 79 whether the host list need to be changed, typically if additional storage nodes should be added. If so, the operating storage node may carry out a process very similar to the writing process as carried out by the server and as described in connection with figs 4A-4C, and 5. The result of this process is that the file is written to a new storage node.

In case of version inconsistencies, the operating storage node may update 81 copies of the file that are stored on other storage nodes, such that

all files stored have the correct version.

Superfluous copies of the stored file may be deleted 83. If the replication process is initiated by a delete operation, the process may jump directly to this step. Then, as soon as all storage nodes have accepted the deletion of the file, the operating storage node simply requests, using unicast, all storage nodes to physically delete the file in question. The storage nodes acknowledge that the file is deleted.

Further the status, i.e. the master host list of the file is updated. It is then optionally possible to repeat steps 73-83 to make sure that the need for replication no longer exists. This repetition should result in a consistent master host list that need not be updated in step 85.

Thereafter, the replication process for that file is concluded, and the operating storage node may release 87 the status as master of the file by sending a corresponding message to all other storage nodes on the host list.

This system where each storage node takes responsibility for maintaining all the files it stores throughout the set of storage nodes provides a self-repairing (in case of a storage node malfunction) self-cleaning (in case of file inconsistencies or files to be deleted) system with excellent reliability. It is easily scalable and can store files for a great number of different applications simultaneously.

The invention is not restricted to the specific disclosed examples and may be varied and altered in different ways within the scope of the appended claims.

CLAIMS

1. A method for maintaining data in a data storage system comprising a plurality of data storage nodes, the method being employed in a storage node in the data storage system and comprising:

-monitoring the status (59) of other storage nodes in the system as well as writing operations (65, 67, 69) carried out in the data storage system, the storage node having access to a host list, including storage nodes storing a data entity;

-detecting (51), based on the monitoring, conditions in the data storage system that imply the need for replication of data between the nodes in the data storage system; and

-initiating a replication process (53) in case such a condition is detected, wherein the replication process includes sending a multicast message, to a plurality of storage nodes, the message enquiring which of those storage nodes store specific data.

2. A method according to claim 1, wherein the monitoring includes listening (59) to heartbeat signals from other storage nodes in the system, and wherein a condition implying need for replication is a malfunctioning storage node.

3. A method according to claim 1 or 2, wherein the data includes files and a condition is one of a file deletion or a file inconsistency.

4. A method according to any of the preceding claims, wherein a replication list, including files that need replication, is maintained.

5. A method according to claim 4, wherein the replication list includes priorities.

6. A method according to any of the preceding claims, wherein the replication process further includes:

-receiving responses from those storage nodes that contain said specific data;

-determining whether said specific data is stored on a sufficient number of storage nodes; and

-if not, selecting at least one additional storage node and transmitting said specific data to that storage node.

7. A method according to claim 6, further comprising updating said specific data on storage nodes containing obsolete versions thereof.

8. A method according to claim 6 or 7, wherein the replication process begins with the storage node attempting to attain mastership for the file to be replicated among all storage nodes in the system.

9. A method according to any of the preceding claims, wherein the monitoring further includes monitoring of reading operations (63) carried out in the data storage system.

10. A data storage node for maintaining data in a data storage system comprising a plurality of data storage nodes, the data storage node comprising:

-means for monitoring the status of other storage nodes in the system as well as writing operations carried out in the data storage system, wherein the monitoring is based on a host list, including storage nodes storing a data entity;

- means for detecting, based on the monitoring, conditions in the data storage system that imply the need for replication of data between the nodes in the data storage system; and

-means for initiating a replication process in case such a condition is detected, wherein the replication process includes sending a multicast message, to a plurality of storage nodes, the message enquiring which of those storage nodes store specific data.

11. A method for writing data to a data storage system comprising a plurality of data storage nodes, the method being employed in a server running an application which accesses data in the data storage system, and comprising:

-sending (41) a multicast storage query to a plurality of said storage nodes;

-receiving (43) a plurality of responses from a subset of said storage nodes, the responses including geographic data relating to the geographic position of each storage node;

-selecting (45) at least two storage nodes in the subset, based on said responses; and

-sending (47) data and a data identifier, corresponding to the data, and a list of storage nodes successfully storing the data to the selected storage nodes.

12. A method according to claim 11, wherein the geographic position includes latitude and longitude of the storage node in question.

13. A method according to claim 12, wherein the responses further includes system age for the storage node in question.

14. A method according to claim 12 or 13, wherein the responses further includes system load for the storage node in question.

15. A method according to any of claim 12-14, wherein the multicast storage query includes a data identifier, identifying the data to be stored.

16. A method according to any of claim 12-15, wherein at least three nodes are selected.

~~17. A method according to any of claim 12-16, wherein a list of storage nodes successfully storing the data is sent to the selected storage nodes.~~

~~17.~~ 17. A server adapted for writing data to a data storage system comprising a plurality of data storage nodes, the server comprising:

-means for sending a multicast storage query to a plurality of said storage nodes;

-means for receiving a plurality of responses from a subset of said storage nodes, the responses including geographic data relating to the geographic position of each storage node;

-means for selecting at least two storage nodes in the subset, based on said responses; and

-means for sending data and a data identifier, corresponding to the data, and a list of storage nodes successfully storing the data to the selected storage node.

2/7

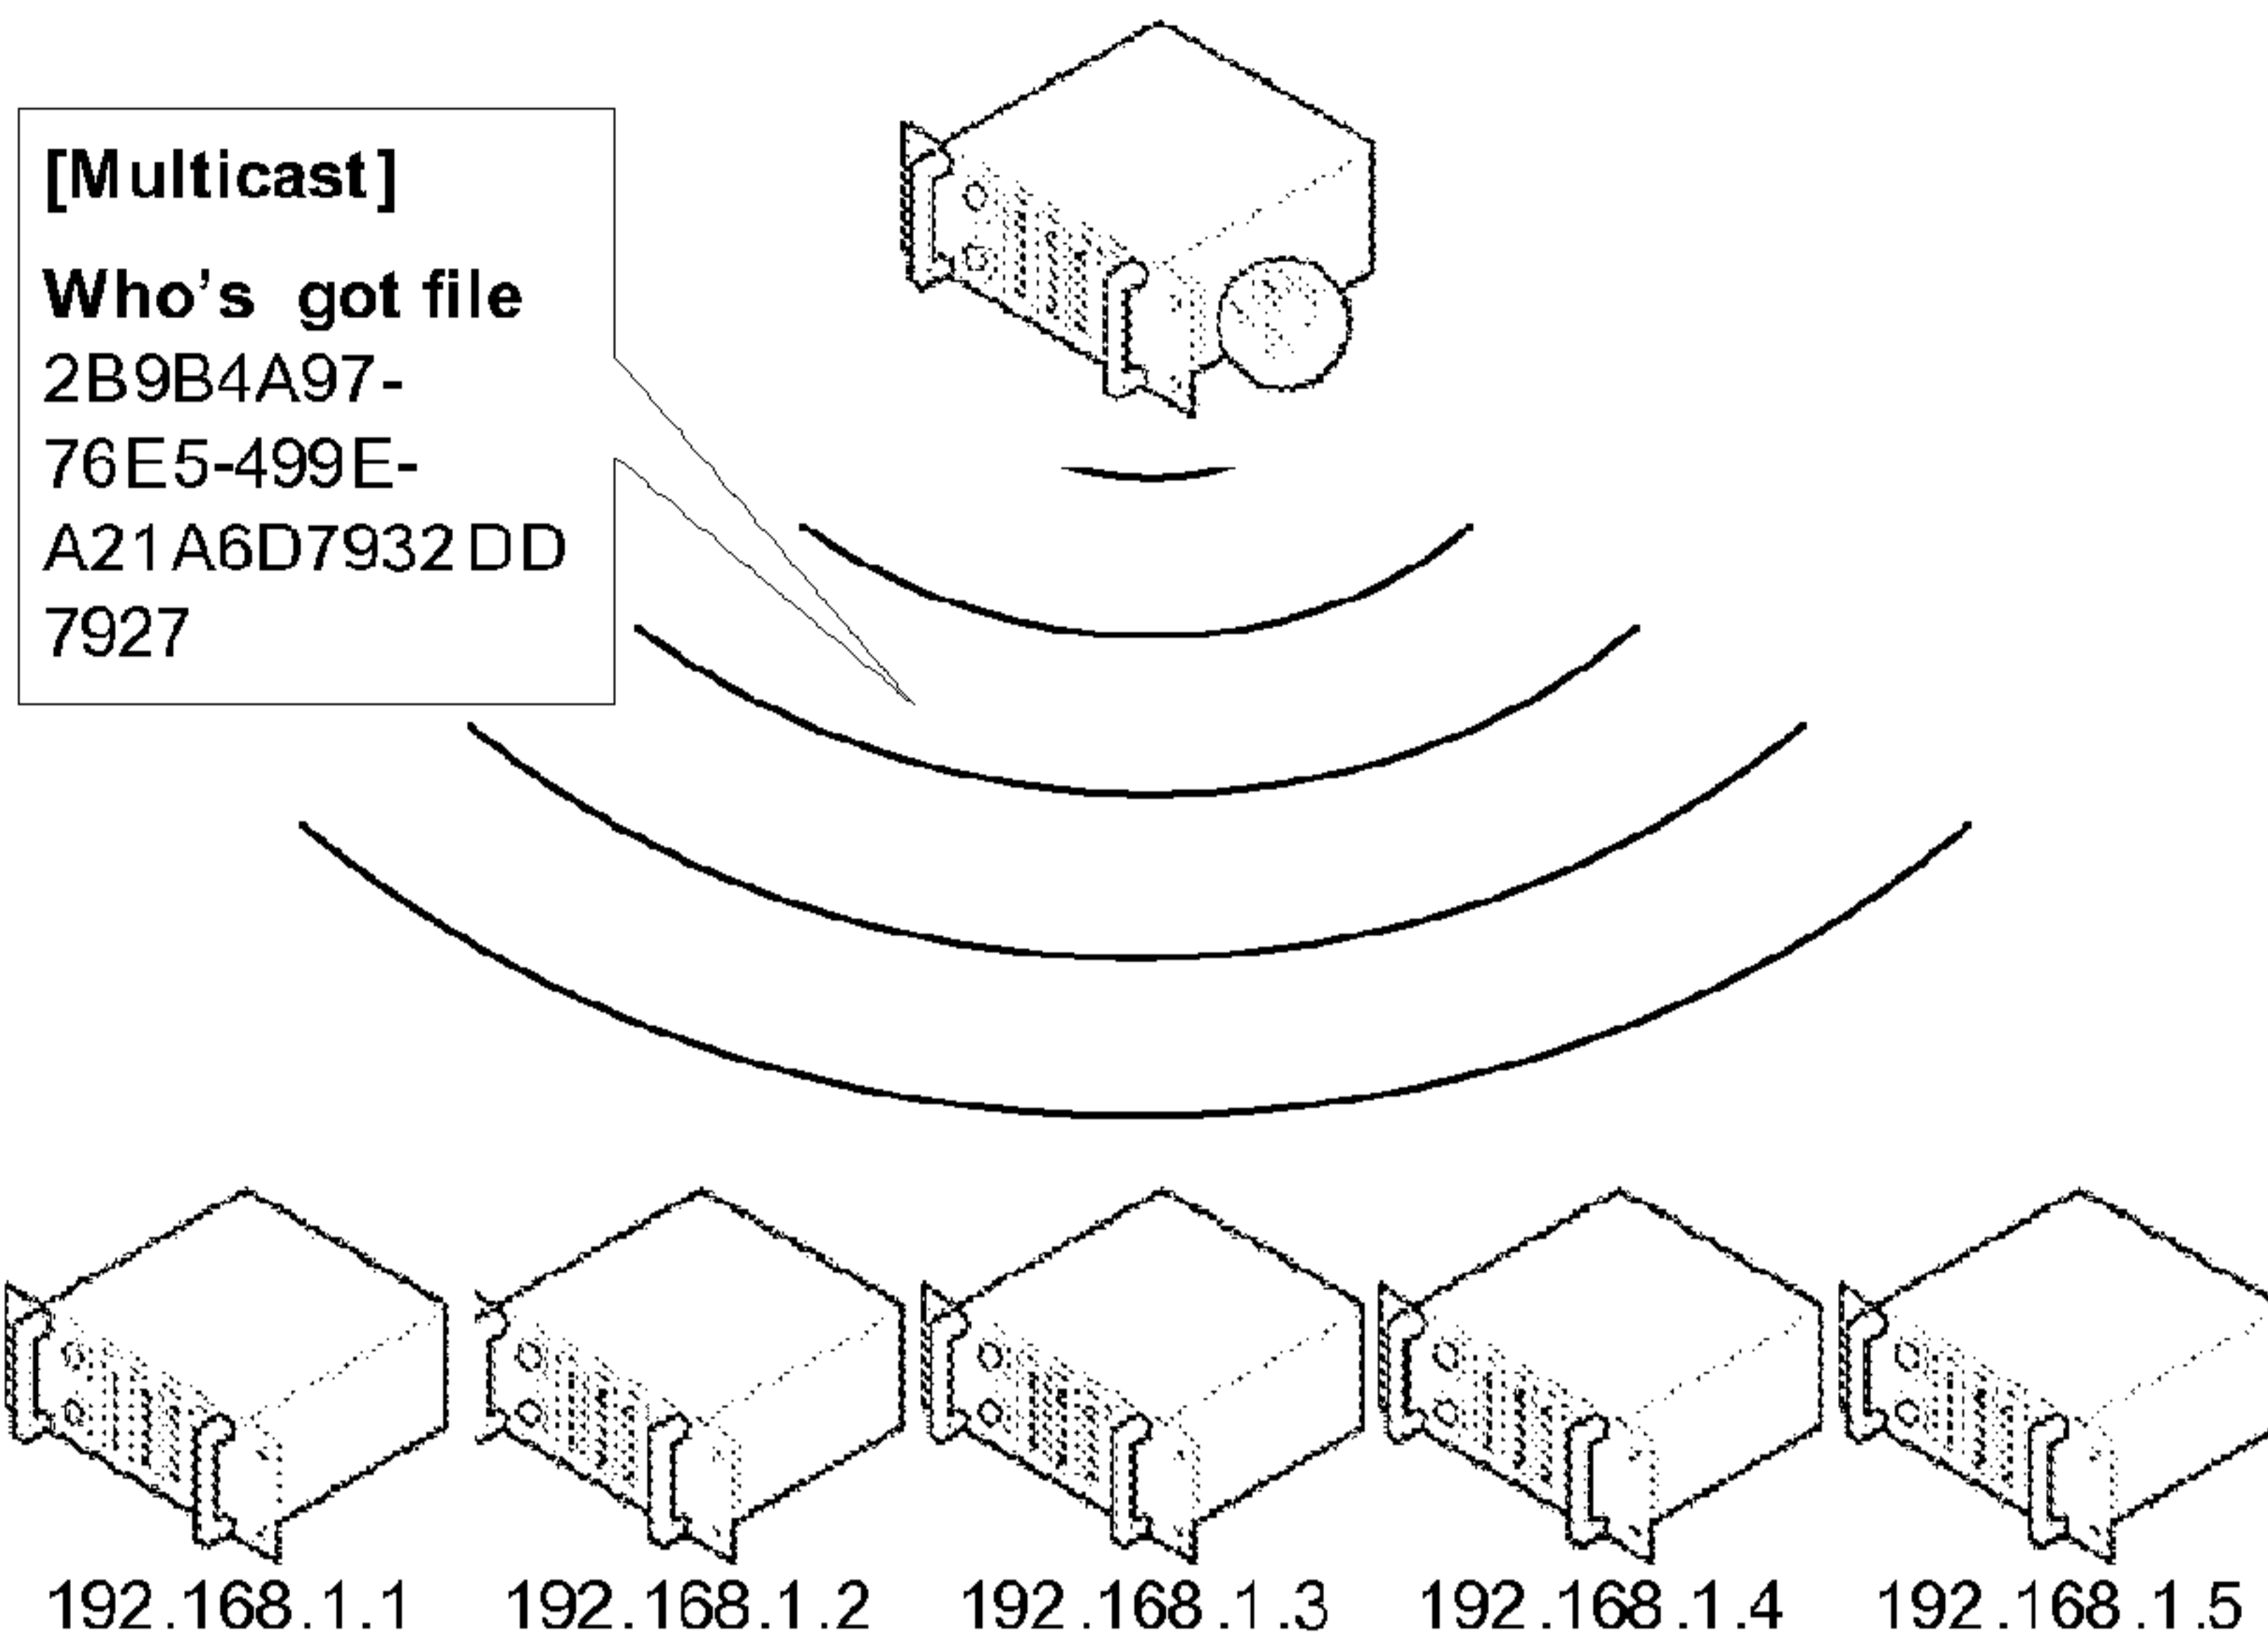


Fig 2A




Fig 2B

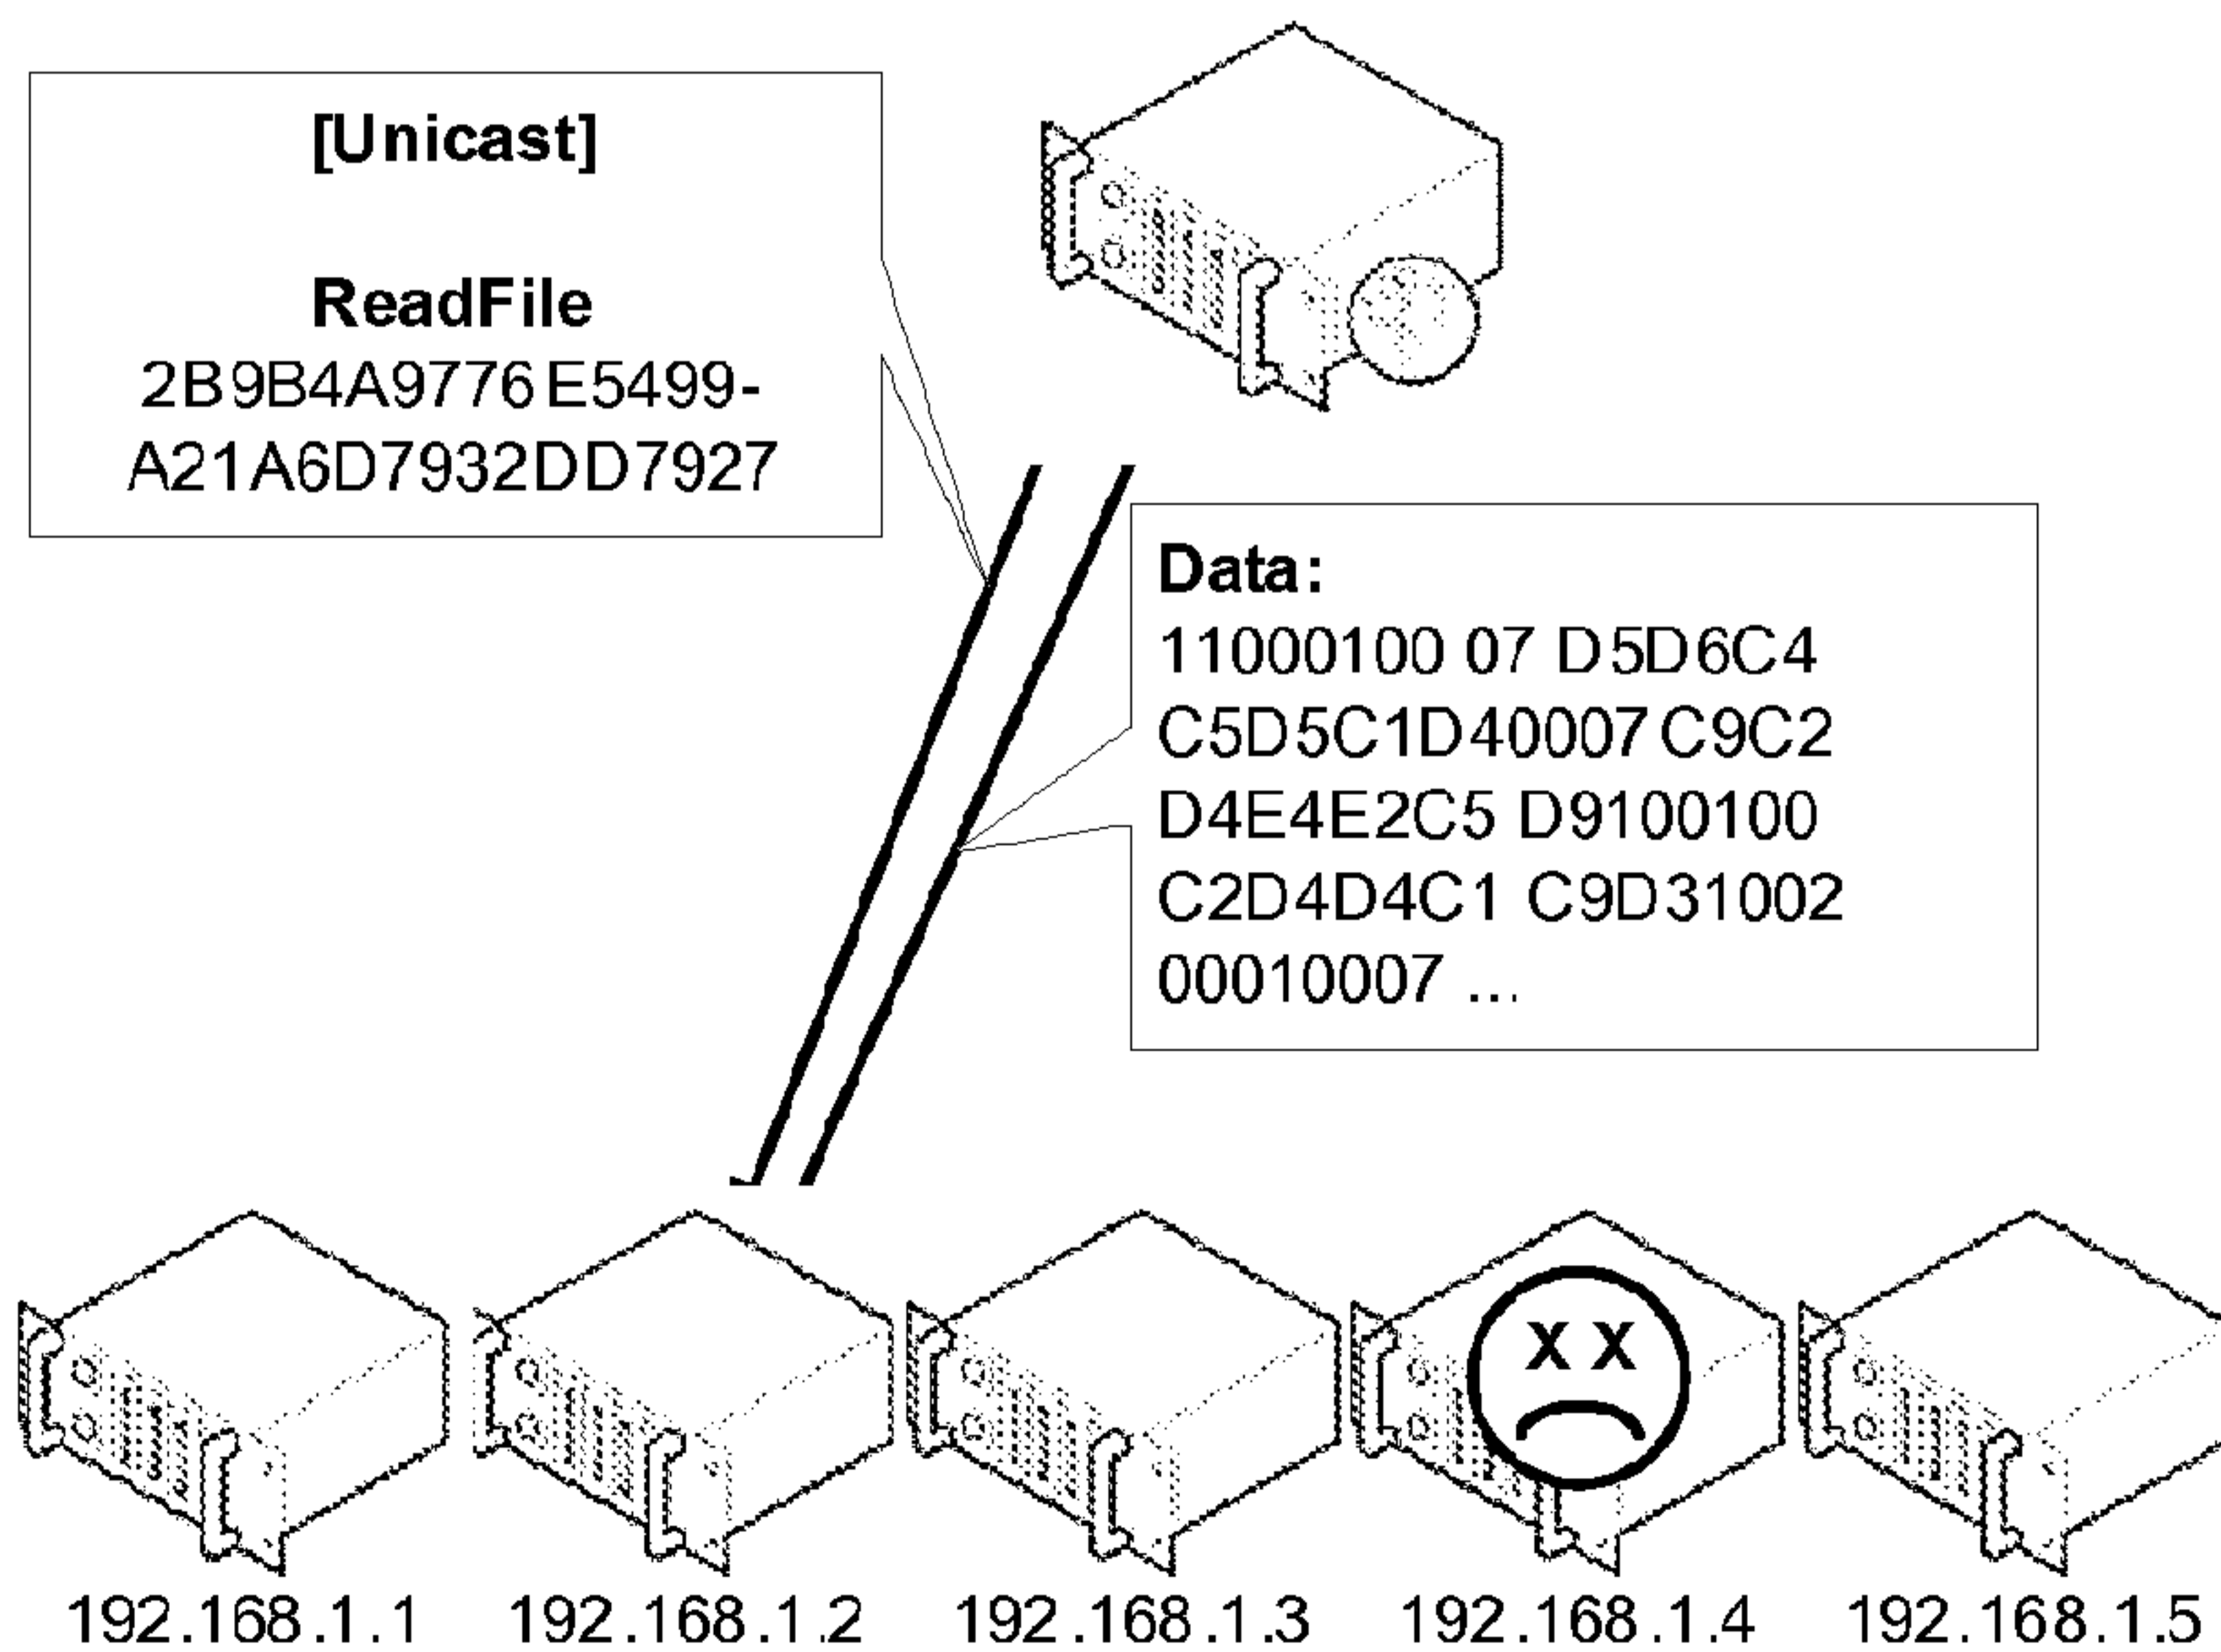


Fig 2C

3/7

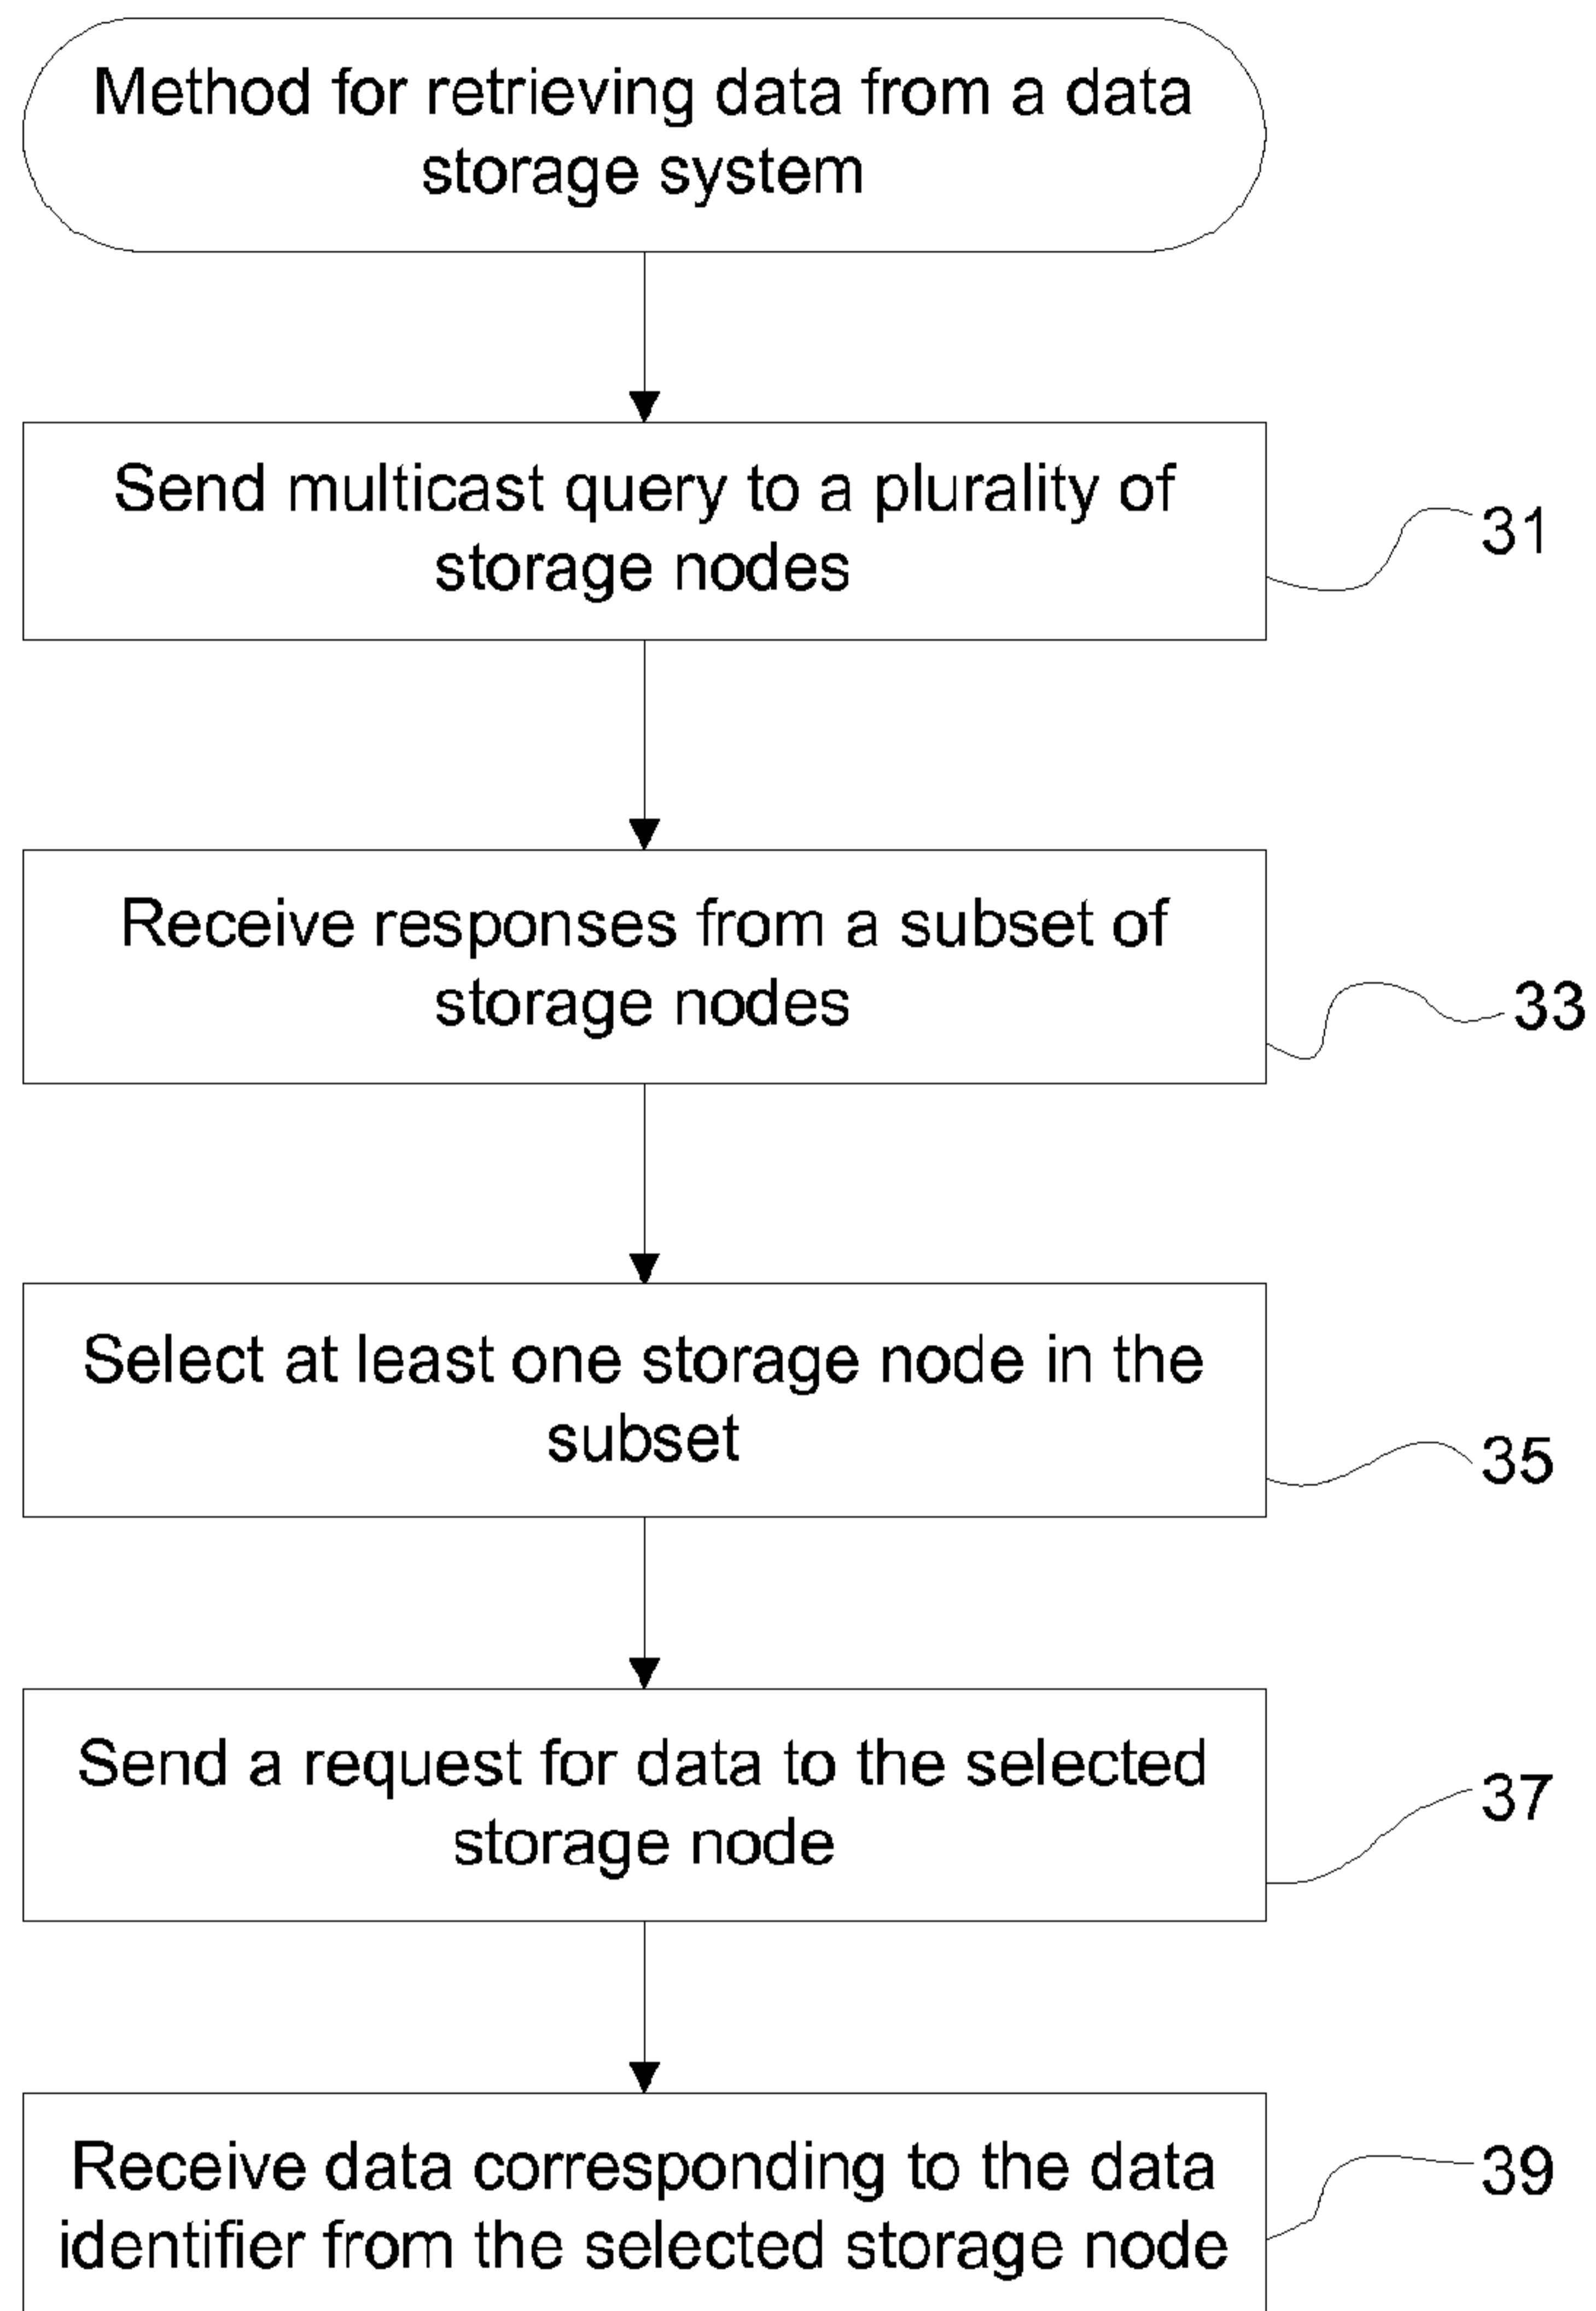


Fig 3

4/7

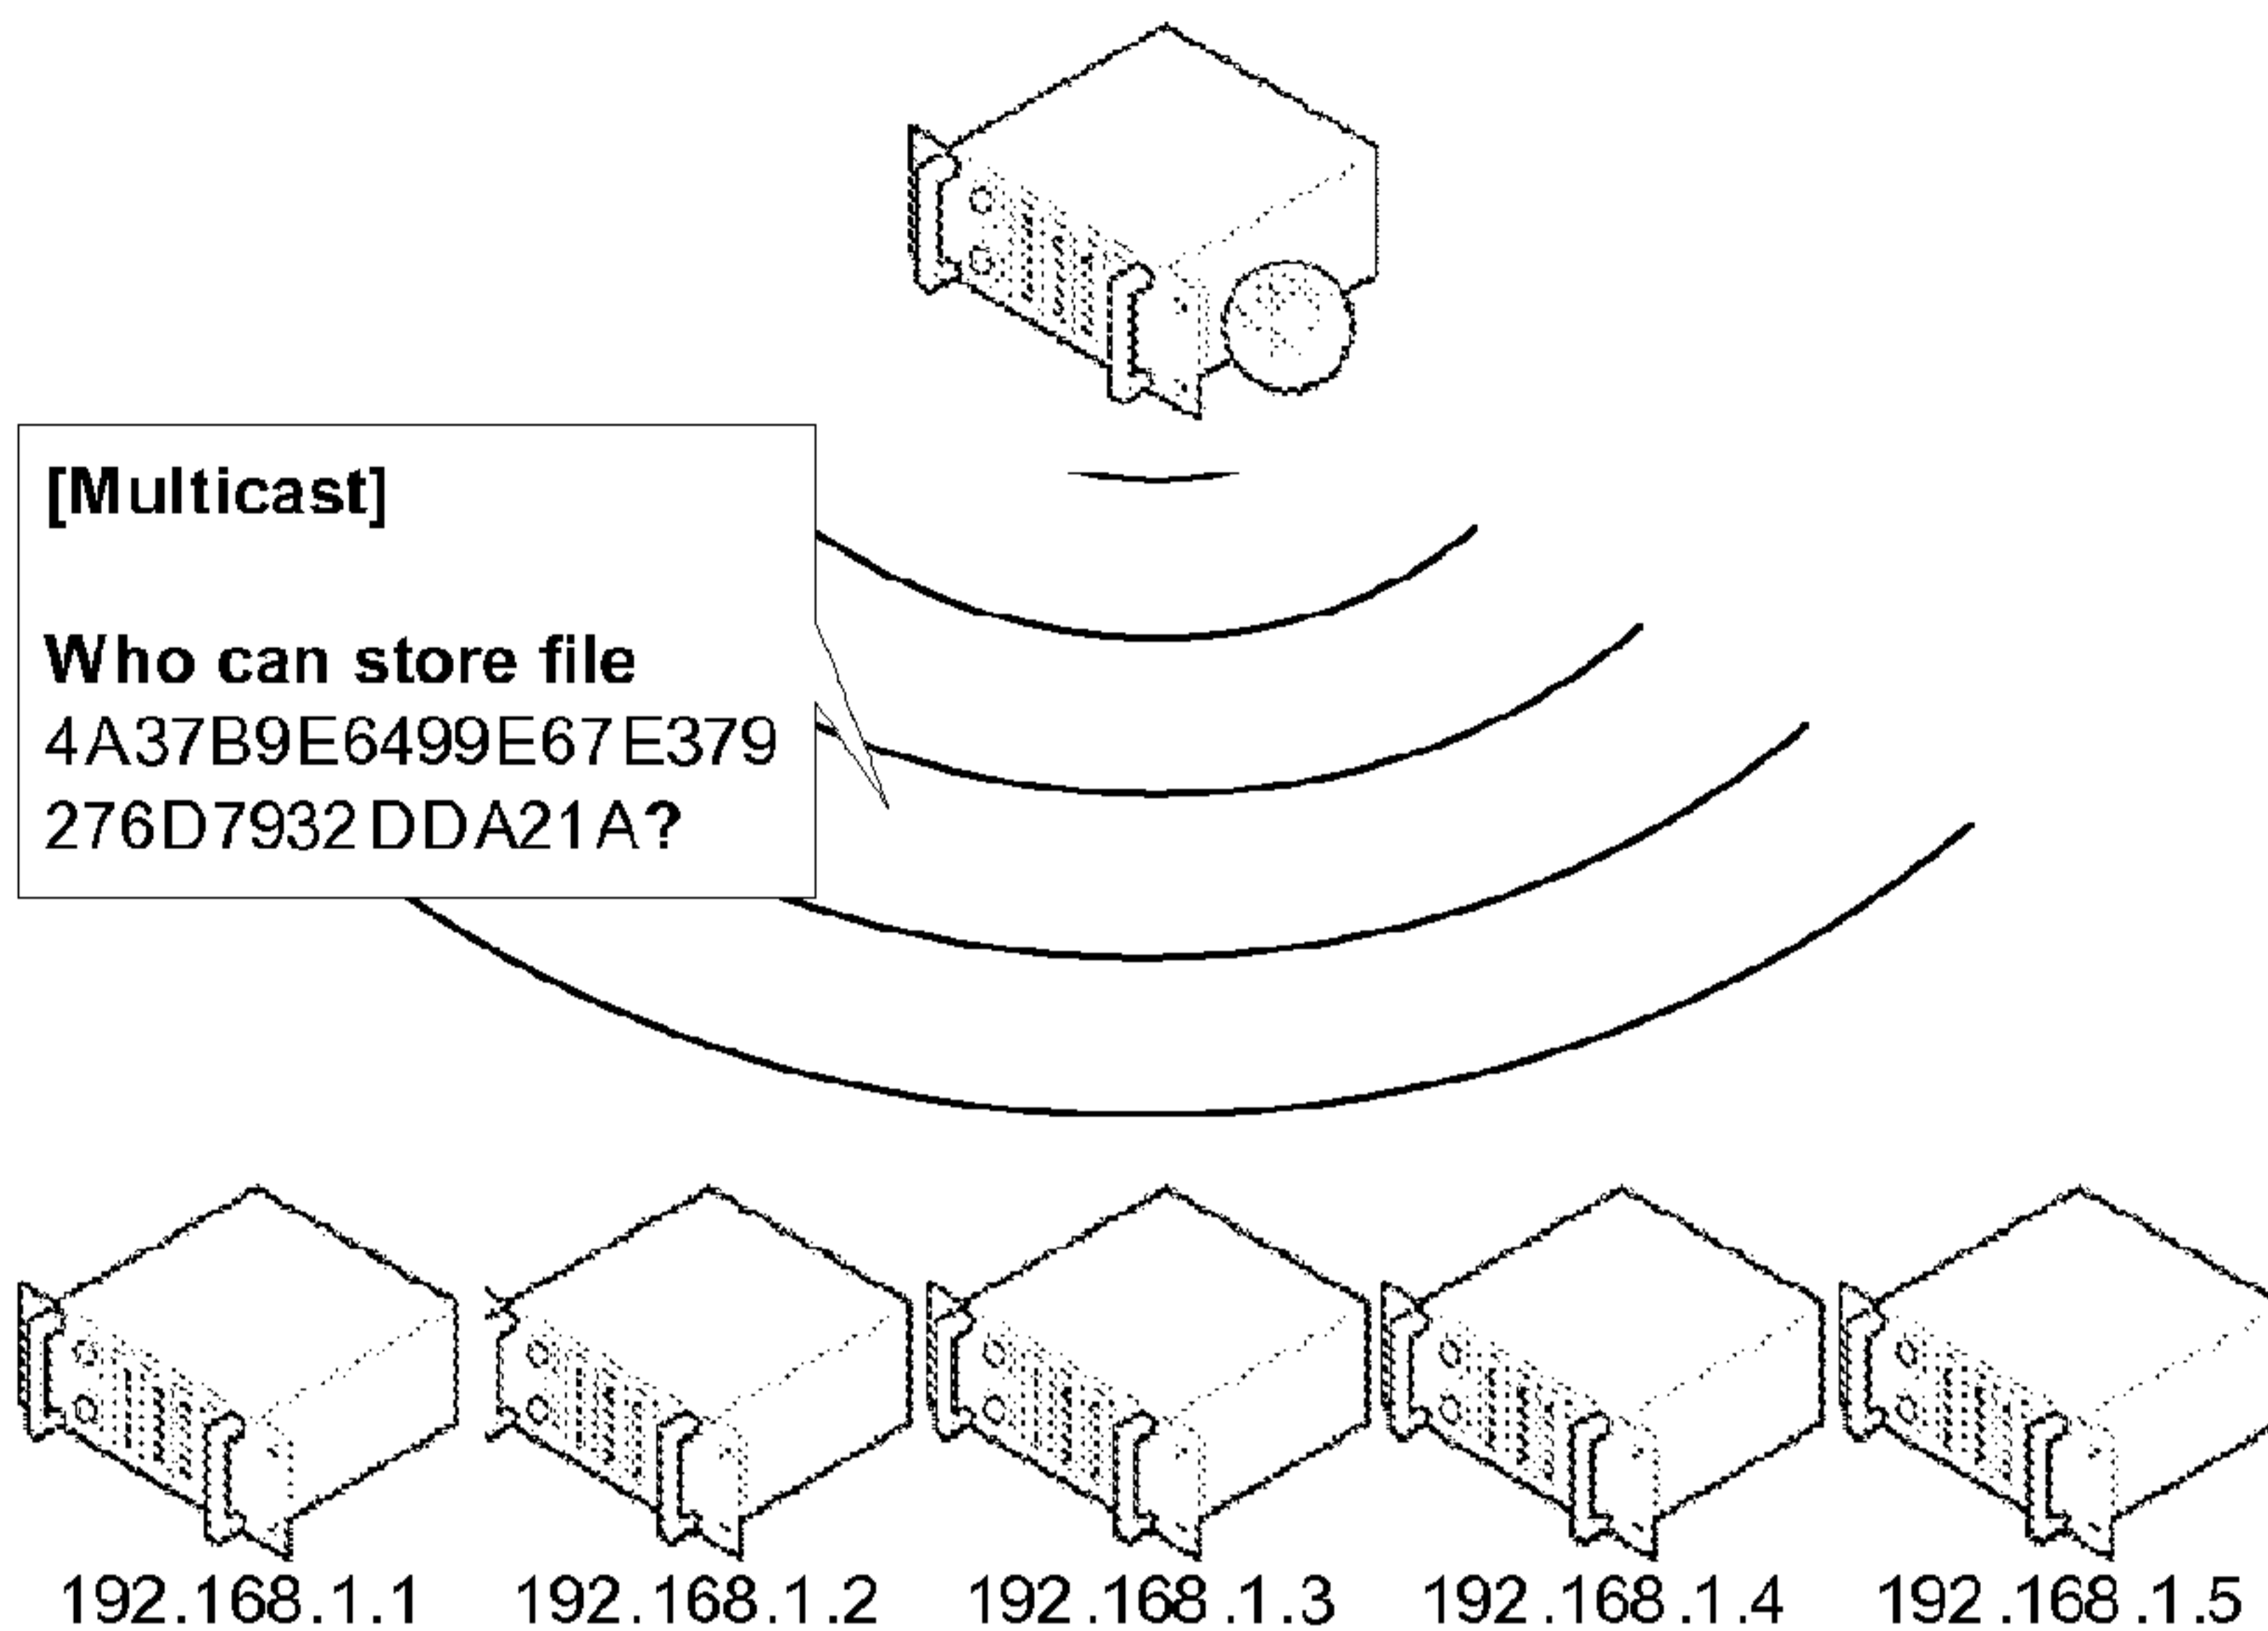


Fig 4A




Fig 4B

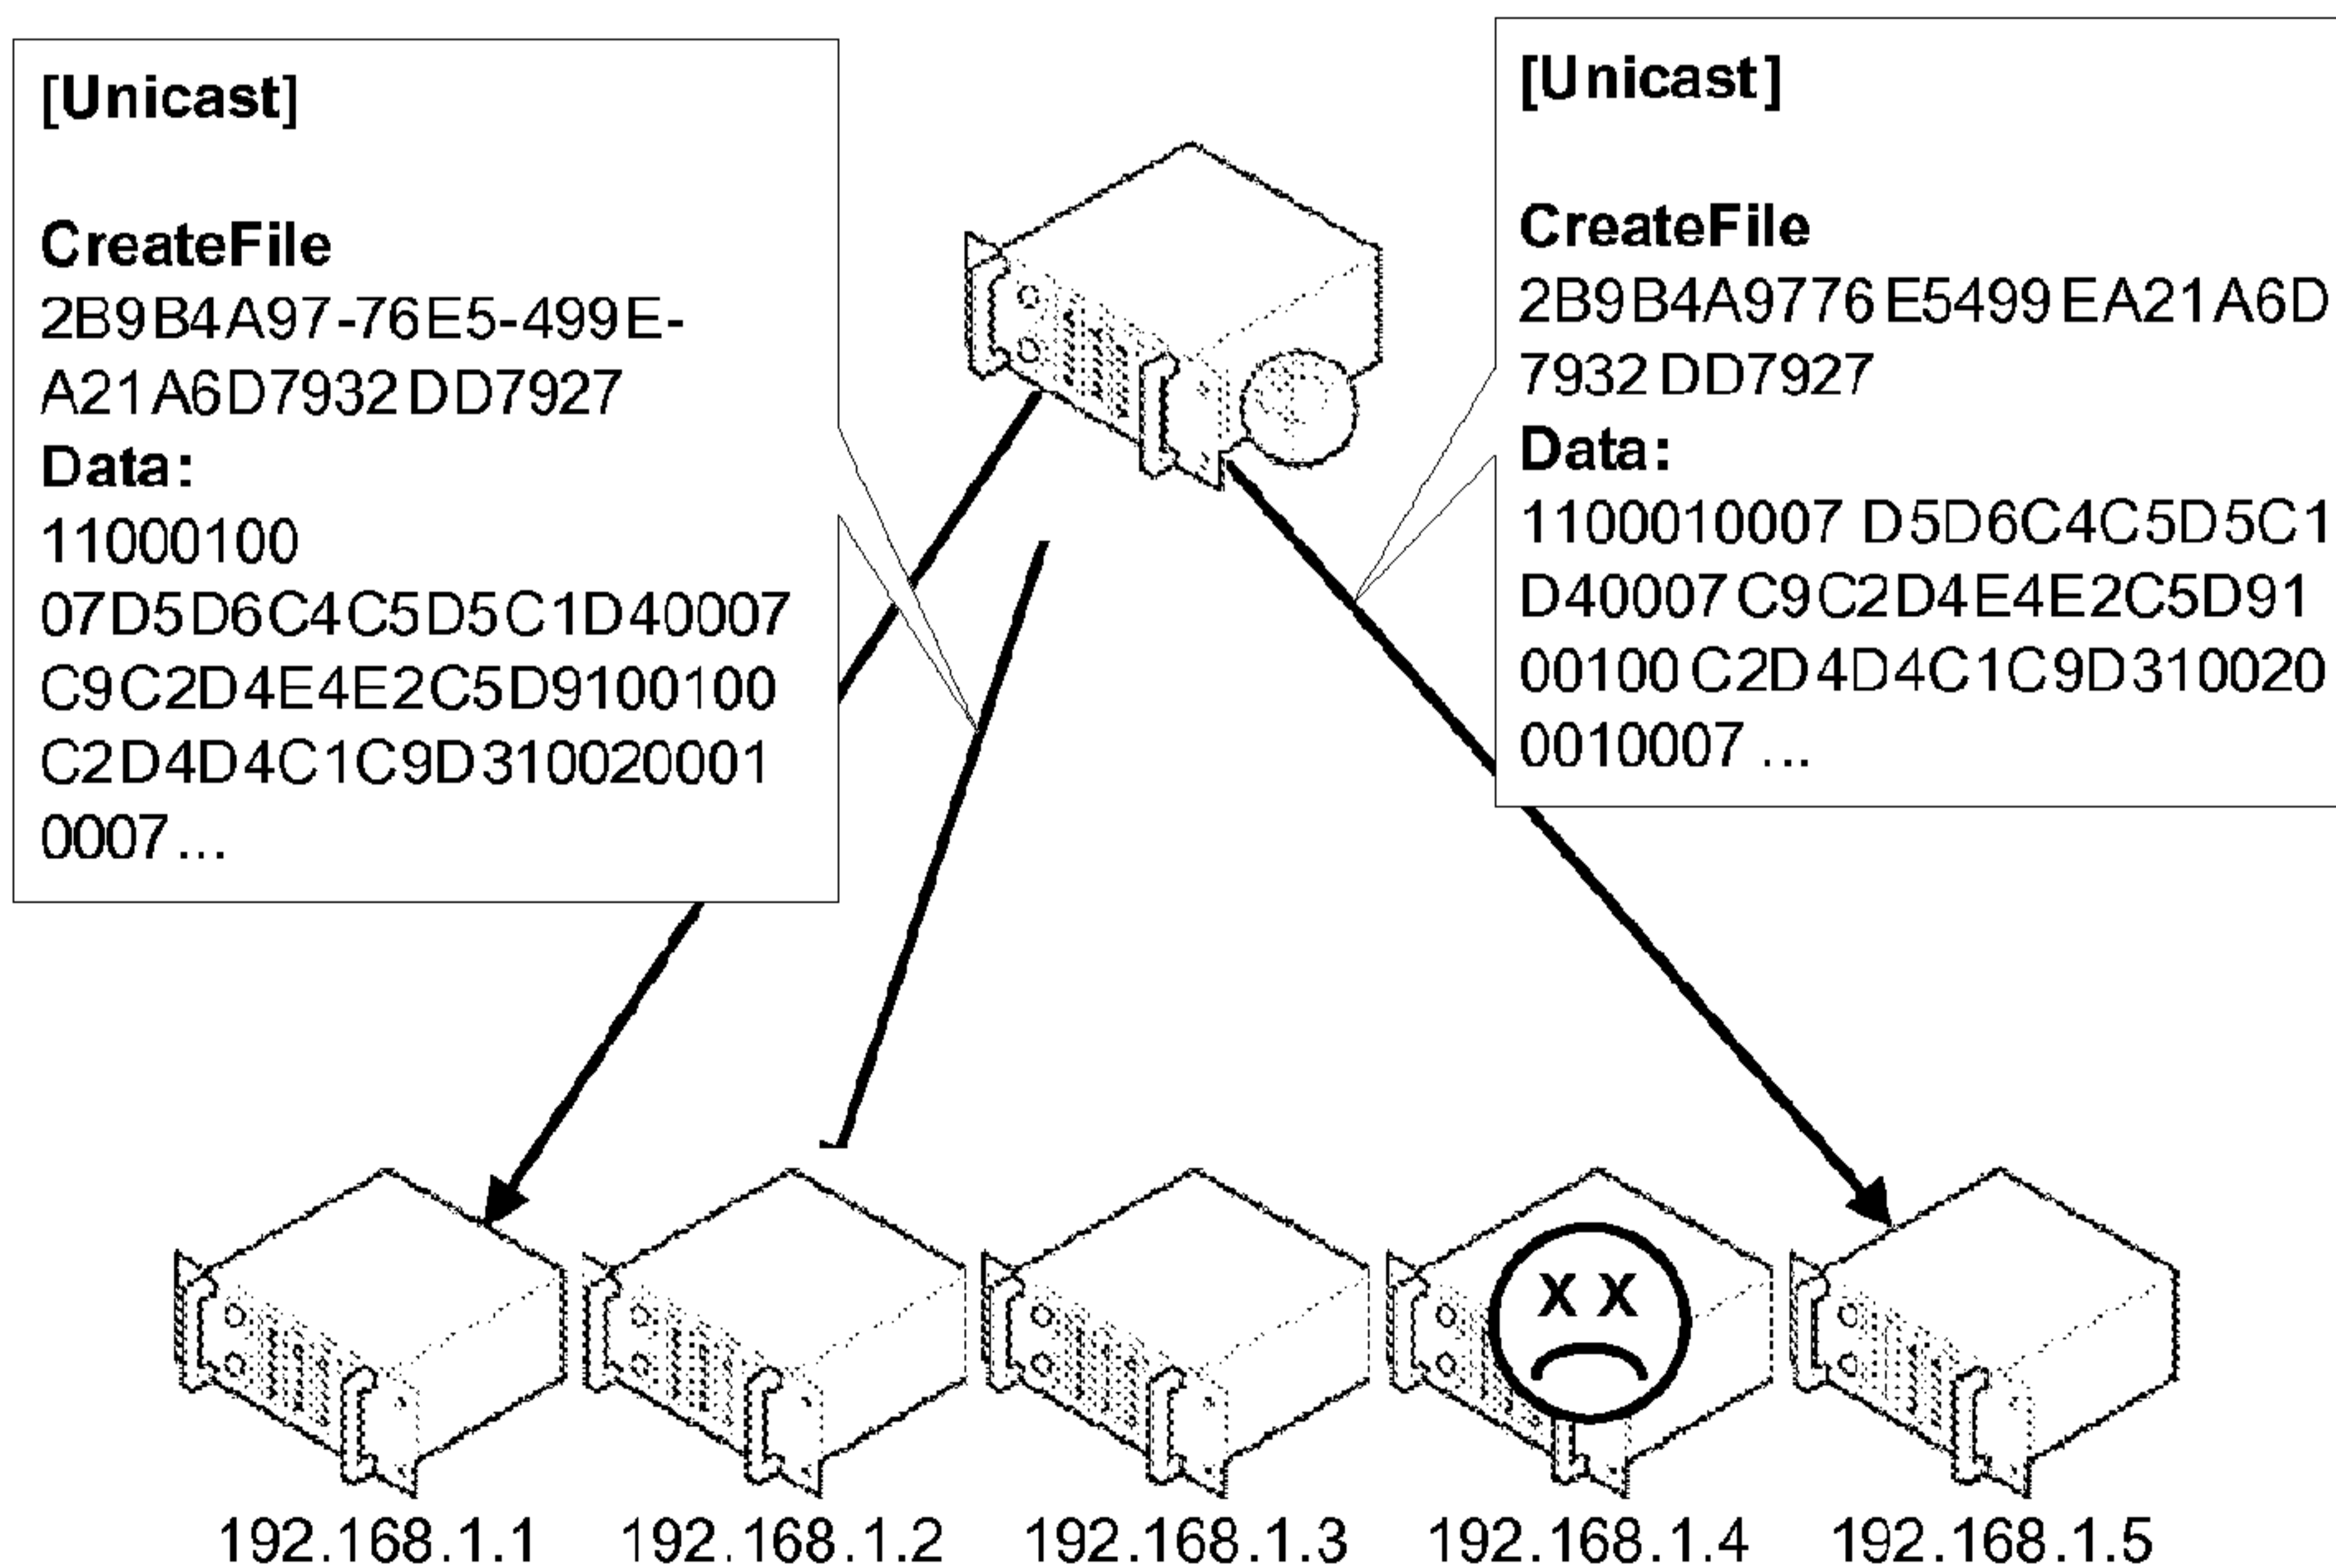


Fig 4C

5/7

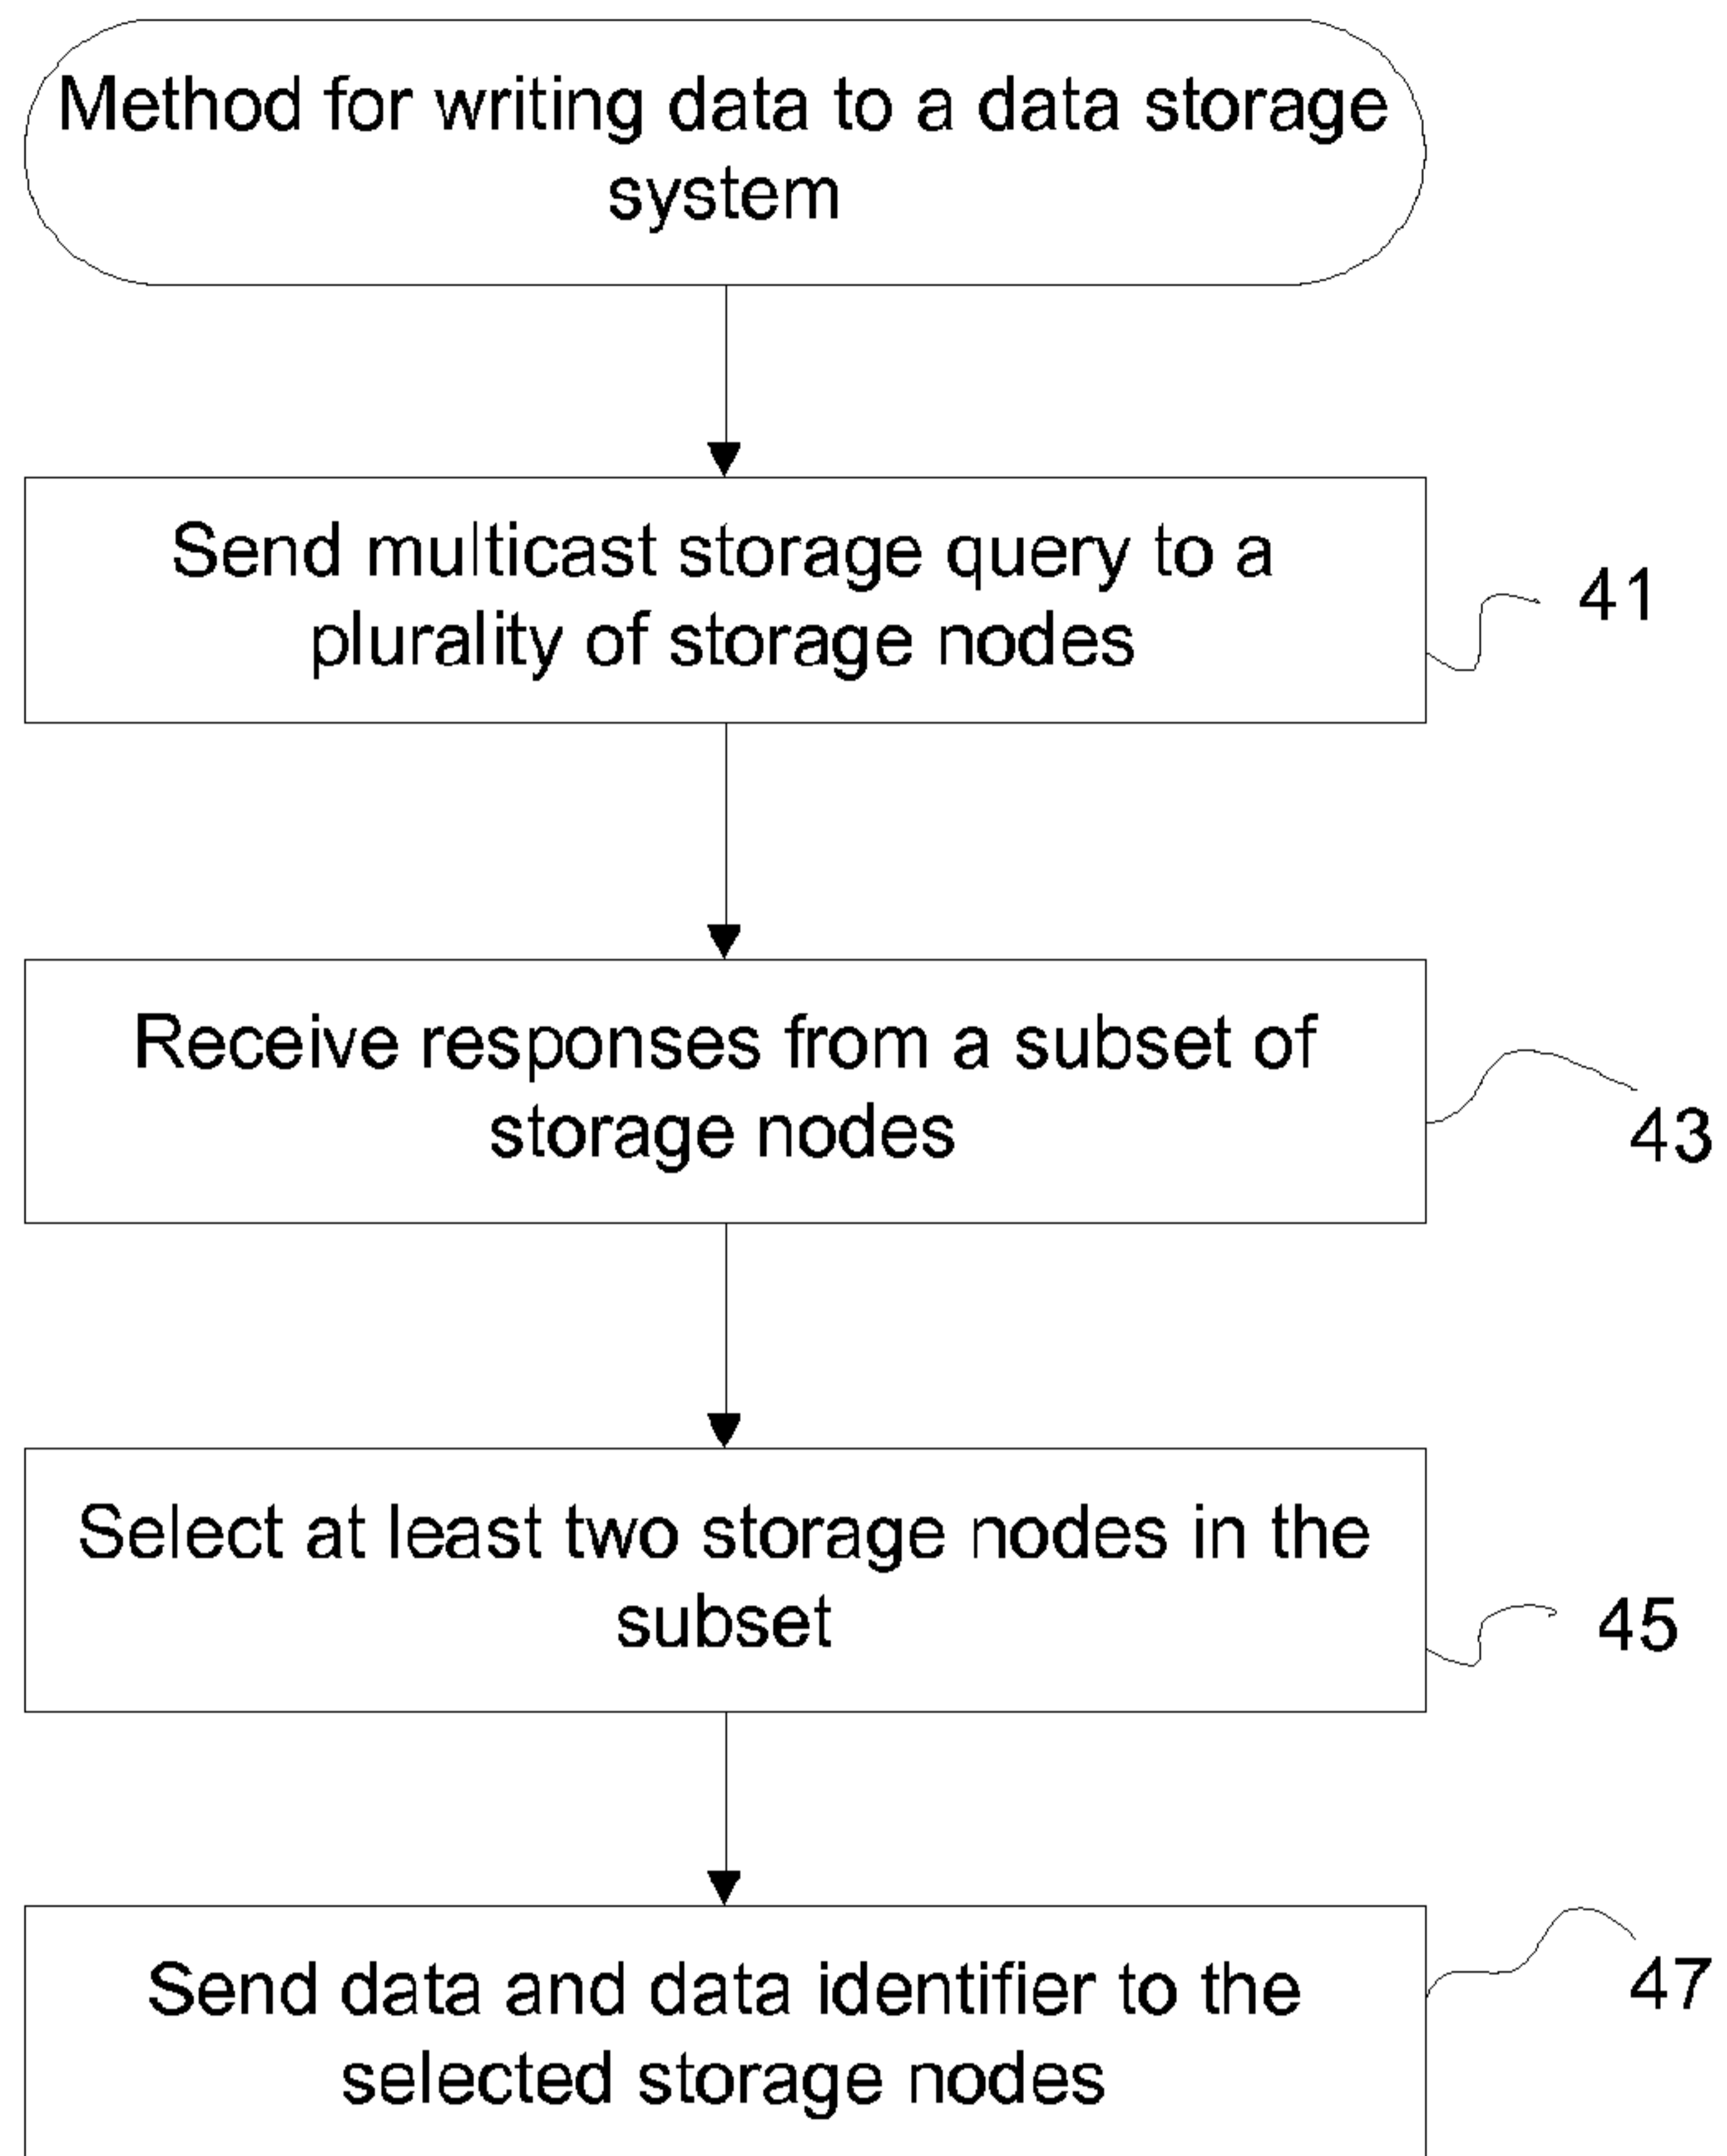


Fig 5




Fig 6




Fig 7

7/7

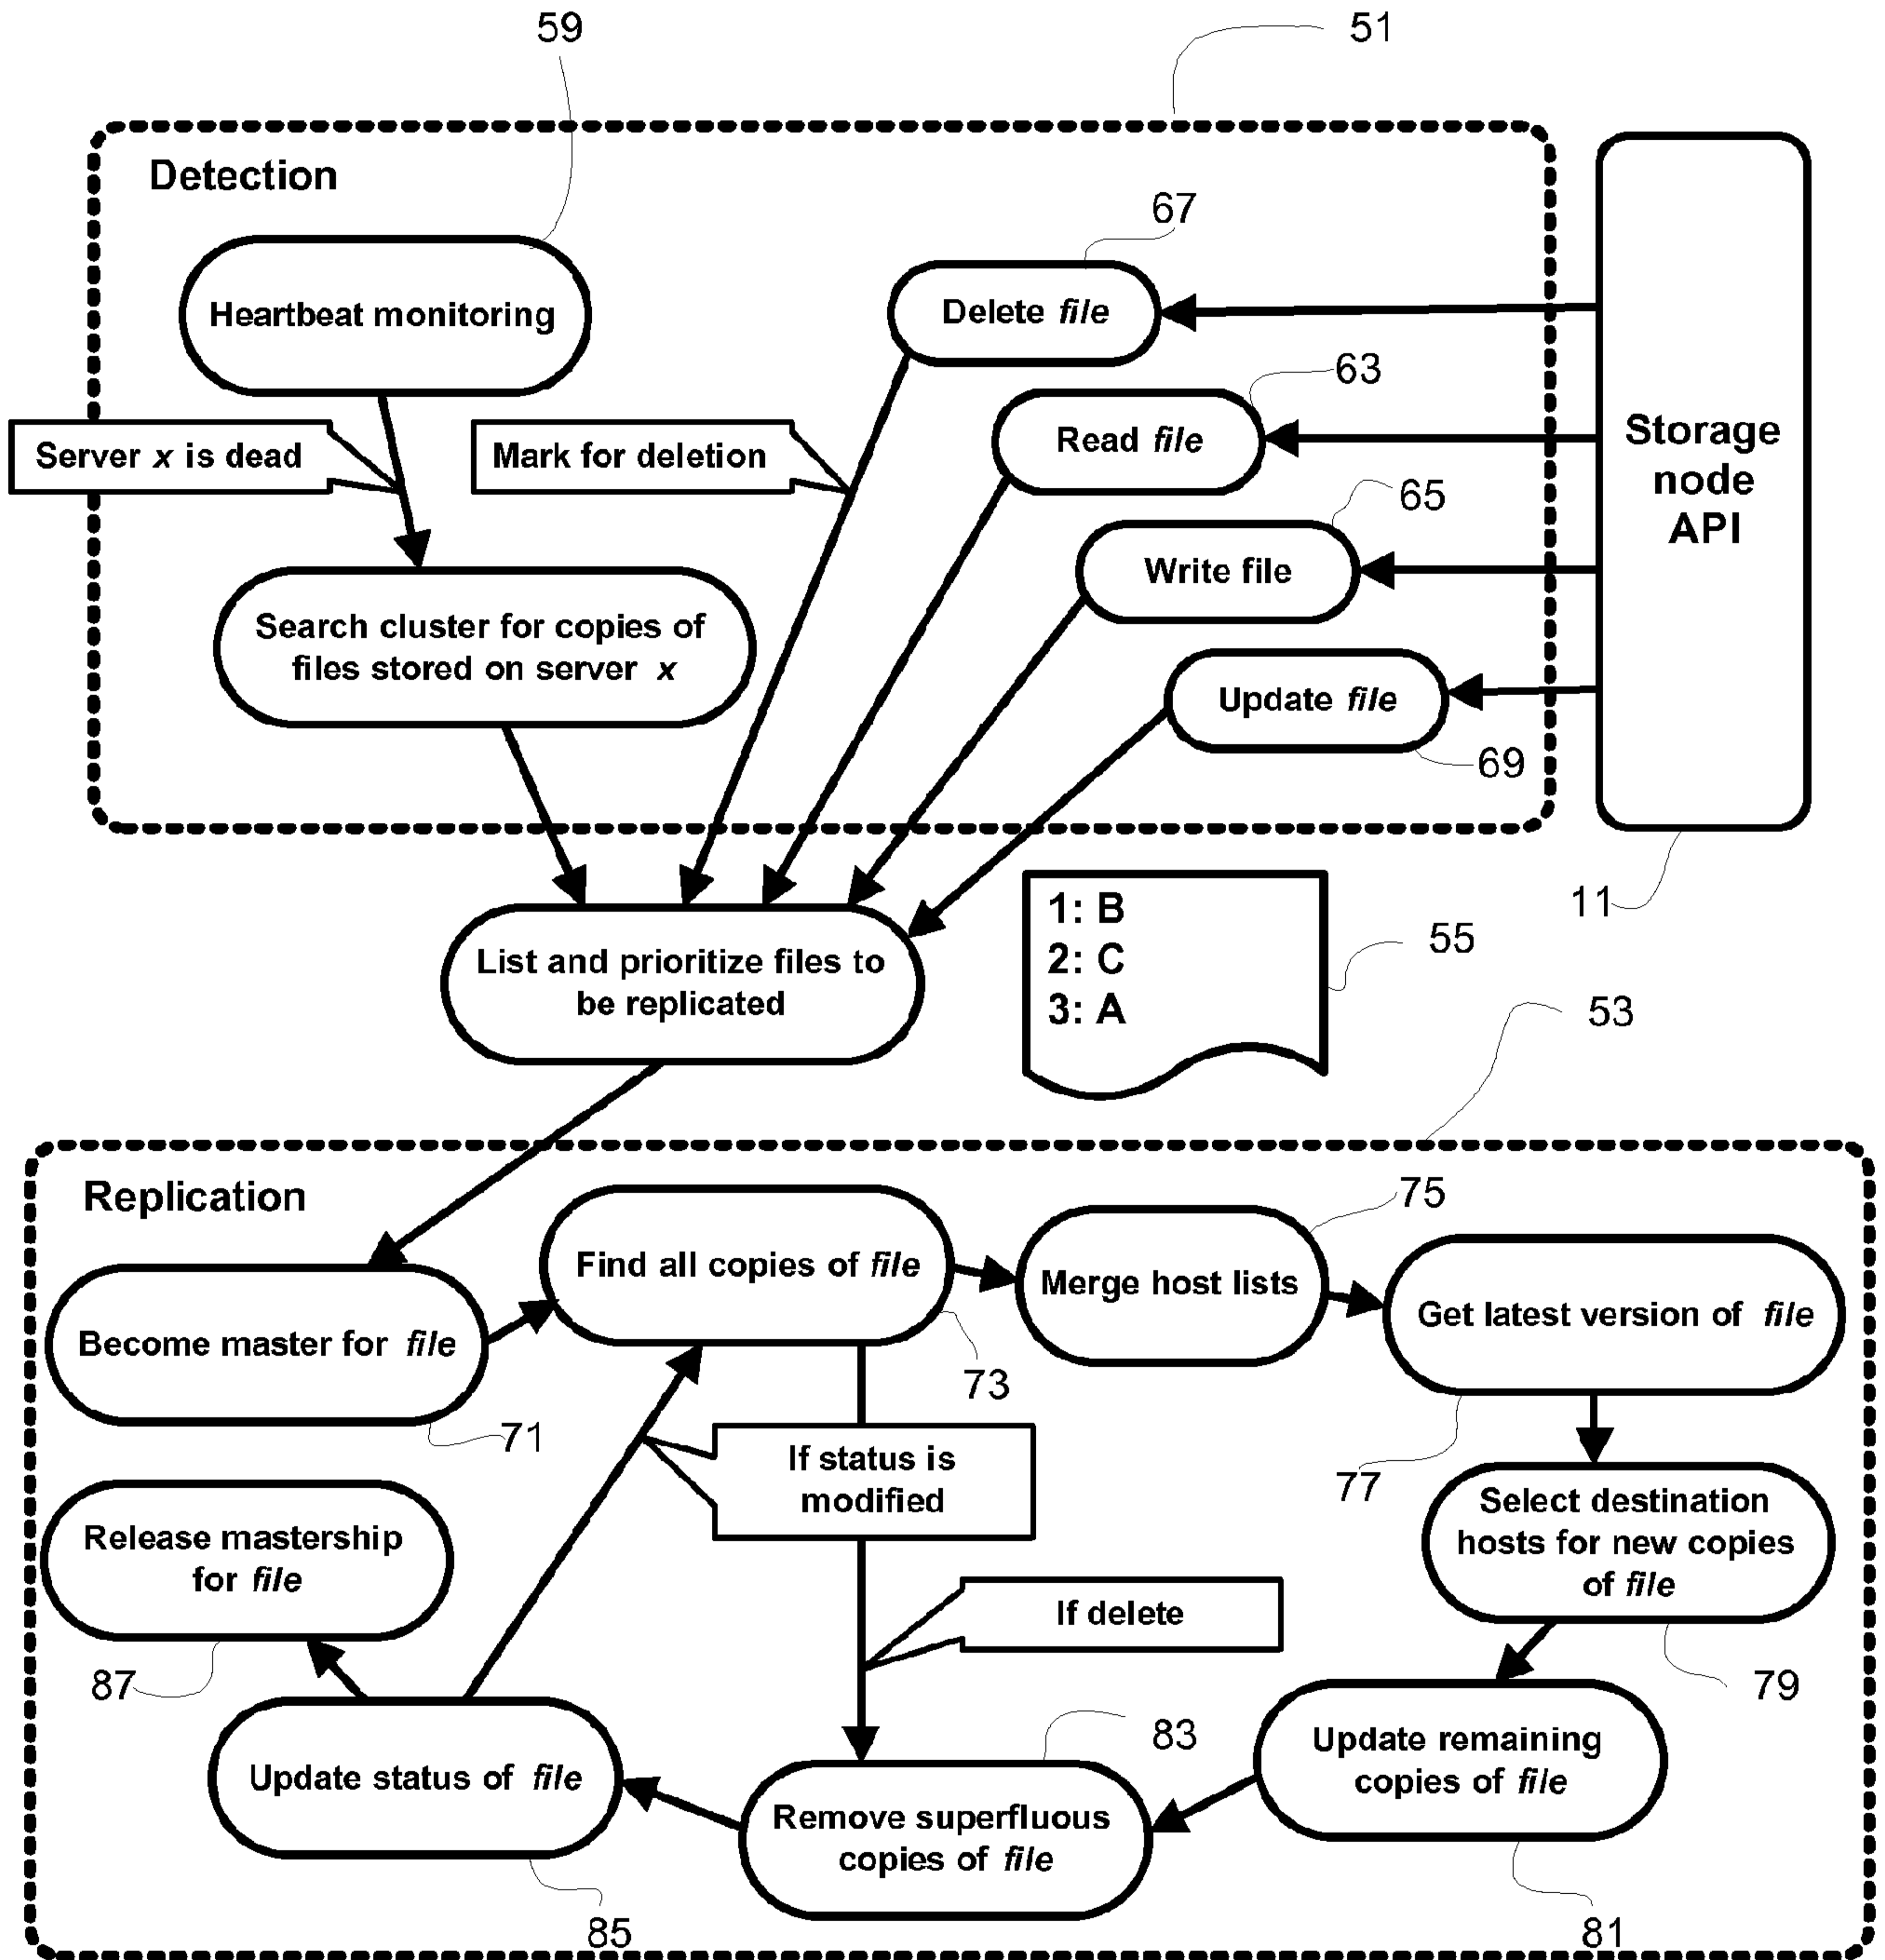


Fig 8

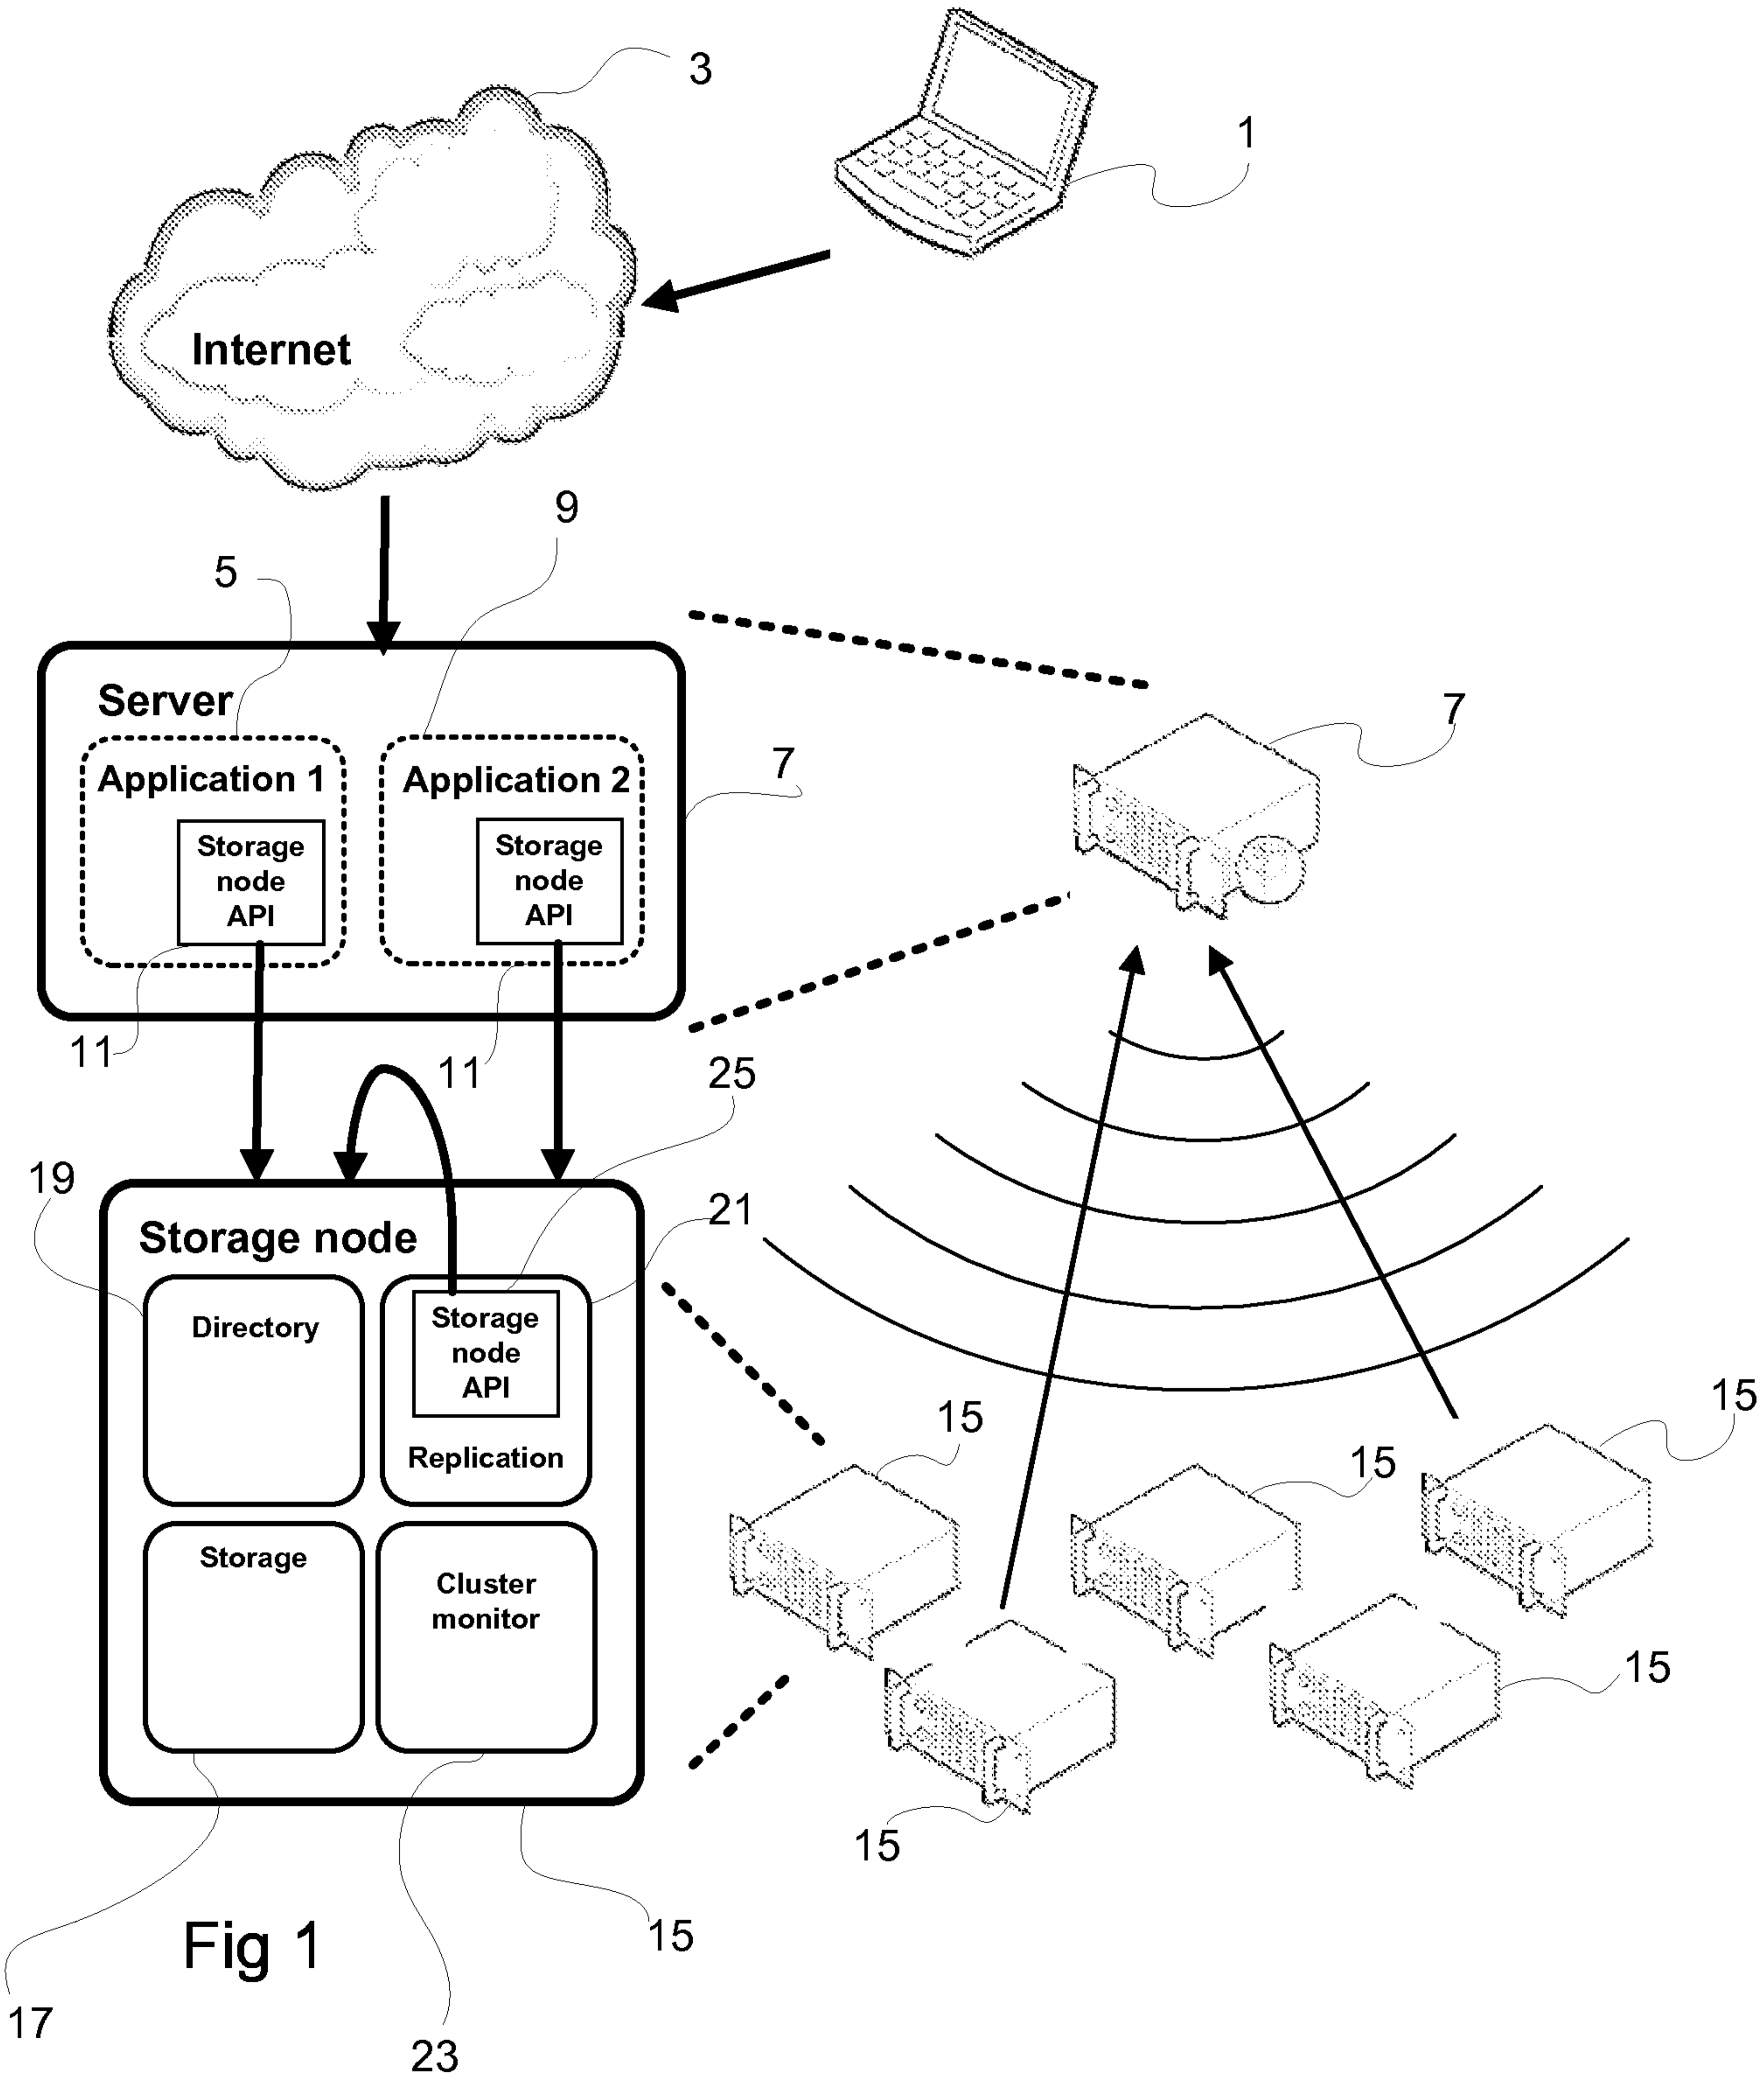


Fig 1

17

23

15