Office de la Propriété Intellectuelle du Canada
Un organisme d'Industrie Canada

Canadian Intellectual Property Office
An agency of Industry Canada

CA 2537814 A1 2005/05/19

(21) 2 537 814

(12) DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION

(13) A1

(86) Date de dépôt PCT/PCT Filing Date: 2004/07/23
(87) Date publication PCT/PCT Publication Date: 2005/05/19
(85) Entrée phase nationale/National Entry: 2006/03/03
(86) N° demande PCT/PCT Application No.: US 2004/023951
(87) N° publication PCT/PCT Publication No.: 2005/045569
(30) Priorité/Priority: 2003/10/23 (US10/691,872)

(51) Cl.Int./Int.Cl. *G06F 17/00* (2006.01), *H04L 29/02* (2006.01), *H04L 12/24* (2006.01)

(71) Demandeur/Applicant:
MICROSOFT CORPORATION, US

(72) Inventeurs/Inventors:
REASOR, STERLING, US;
BREZAK, JOHN, US;
FLO, ERIC, US;
SAMJI, MOHAMMED, US

(74) Agent: SMART & BIGGAR

(54) Titre : SYSTEME ET PROCEDE PERMETTANT DE GENERER DES VISUALISATIONS DE DONNEES AGREGEES DANS UN RESEAU INFORMATIQUE

(54) Title: SYSTEM AND METHOD FOR GENERATING AGGREGATED DATA VIEWS IN A COMPUTER NETWORK

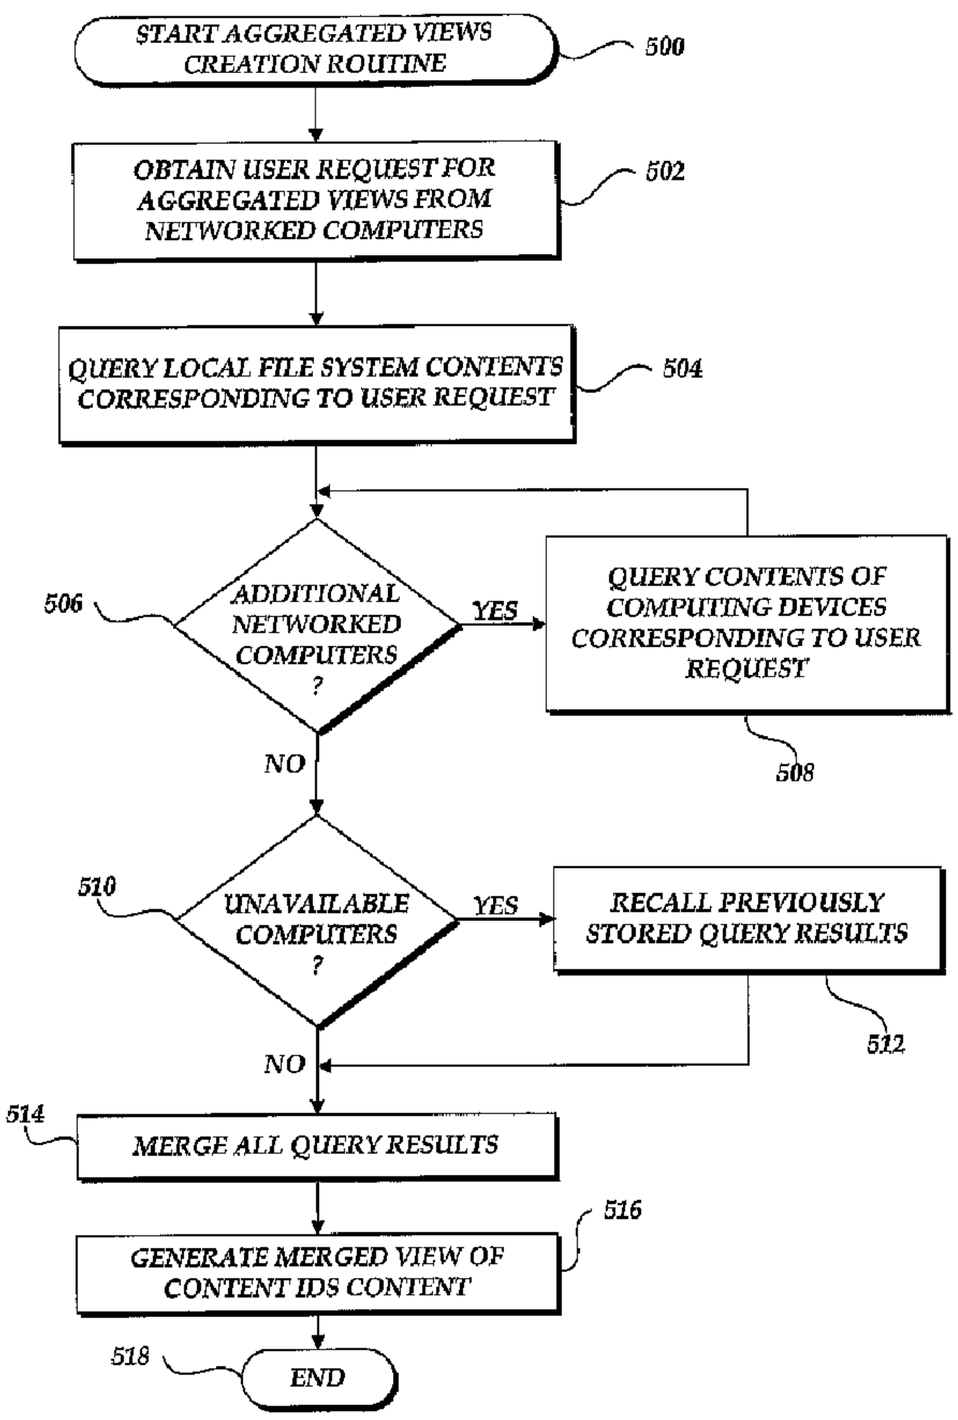

(57) Abrégé/Abstract:
A system and method for generating aggregated content views in a computing network are provided. A host computing device obtains a request for an aggregated view of content corresponding to a set of criteria. The host computing device queries itself and

CA 2537814 A1 2005/05/19

(21) 2 537 814

(13) A1

(57) Abrégé(suite)/Abstract(continued):
each computing device in a defined network for locally stored content matching the set of criteria. The query results are merged and displayed to a user at the host computing device as an aggregated list view.

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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
19 May 2005 (19.05.2005)

PCT

(10) International Publication Number
WO 2005/045569 A2

(51) International Patent Classification[7]: G06F

(21) International Application Number:
PCT/US2004/023951

(22) International Filing Date: 23 July 2004 (23.07.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/691,872 23 October 2003 (23.10.2003) US

(71) Applicant: **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, WA 98052 (US).

(72) Inventors: **REASOR, Sterling**; 6238 - 119th Place SE, Bellevue, WA 98006 (US). **BREZAK, John**; 24033 NE 188th Street, Woodinville, WA 98077 (US). **FLO, Eric**; 23410 SE 17th Place, Sammamish, WA 98075 (US). **SAMJI, Mohammed**; 1805 134th Avenue SE #31, Bellevue, WA 98005 (US).

(74) Agent: **URIBE, Mauricio, A.**; Christensen O'Connor Johnson Kindness PLLC, 1420 Fifth Avenue, Suite 2800, Seattle, WA 98101-2347 (US).

(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, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, 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),

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR GENERATING AGGREGATED DATA VIEWS IN A COMPUTER NETWORK

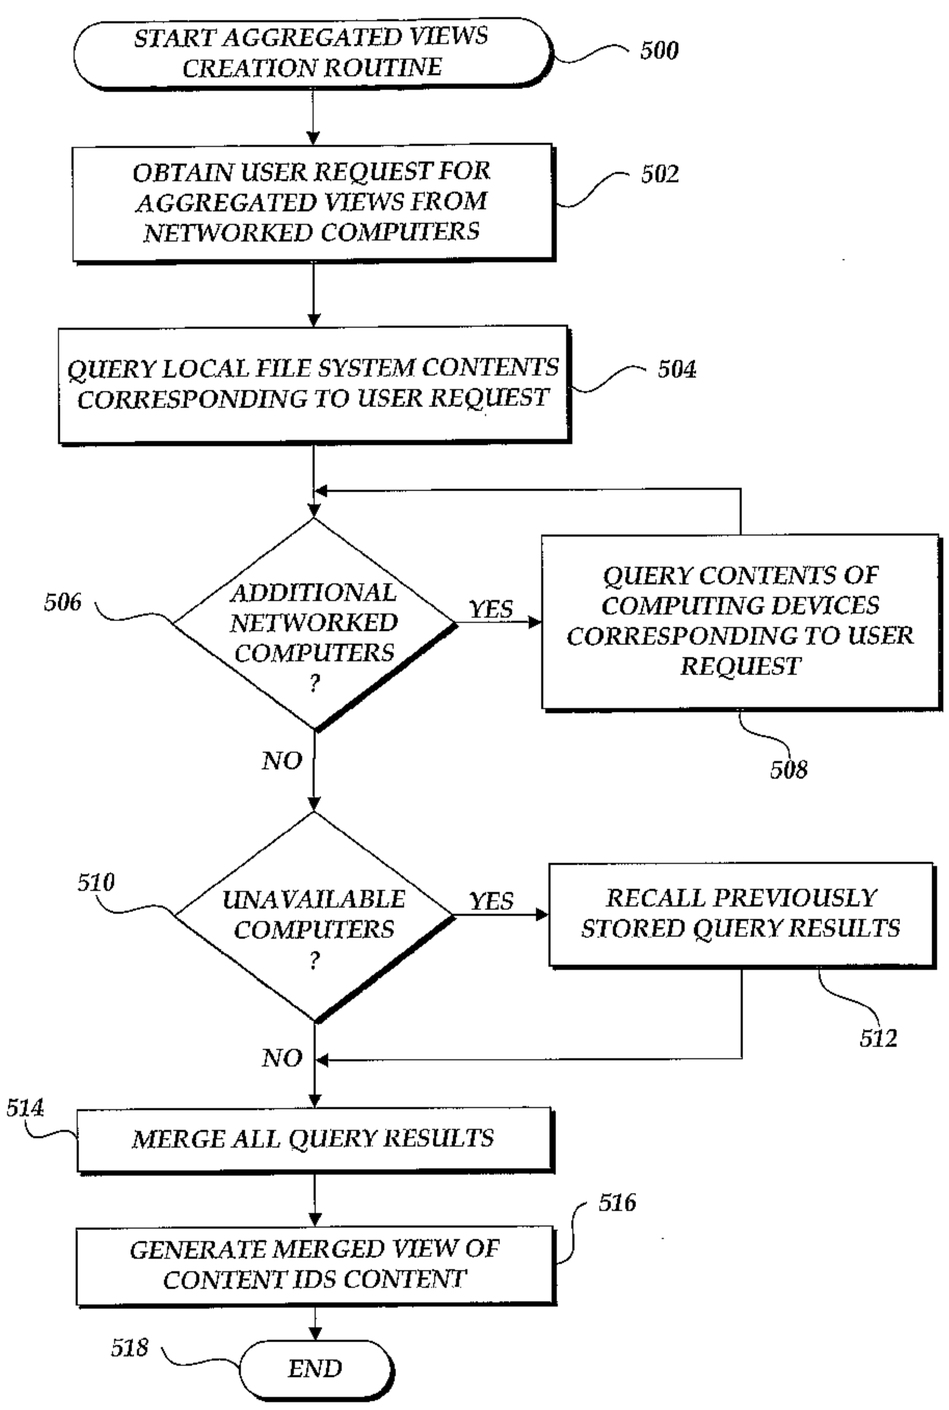

(57) Abstract: A system and method for generating aggregated content views in a computing network are provided. A host computing device obtains a request for an aggregated view of content corresponding to a set of criteria. The host computing device queries itself and each computing device in a defined network for locally stored content matching the set of criteria. The query results are merged and displayed to a user at the host computing device as an aggregated list view.

WO 2005/045569 A2

WO 2005/045569 A2

European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, 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:

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

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.

SYSTEM AND METHOD FOR GENERATING AGGREGATED DATA VIEWS IN A COMPUTER NETWORK

FIELD OF THE INVENTION

In general, the present invention relates to computer software and computer communication networks, and in particular, to a system and method for aggregating and displaying the contents of storage locations across a computer network.

BACKGROUND OF THE INVENTION

Generally described, the advancement and availability of computing devices and communication networks connecting computing devices provides a variety of operating environments for computer users. In a managed computing network environment, common to business and other large-scale computing networks, the computing network includes a centralized network authority, such as a primary domain controller, that regulates user access by maintaining passwords and permissions. The centralized network authority can also manage access to locally stored data files to authorized users of the network.

In contrast to the managed computing networks, unmanaged computing networks, common to homes and other small-scale networks, do not typically incorporate a primary domain controller for regulating user access or centralized data file storage. Instead, many unmanaged computing networks require individual user management in terms of regulating user security authorization to use the various computing devices on the network. Further, most unmanaged networks require additional user participation and knowledge to store and recall content stored on the various computing devices on the network. One attempt to provide standardized user authentication has been proposed in commonly assigned, co-pending application 10/414,354, entitled SMALL-SCALE SECURED COMPUTER NETWORK GROUP WITHOUT CENTRALIZED MANAGEMENT, and filed on April 15, 2003, in which a user's security identification is replicated to each computing device in an unmanaged network.

One attempt to provide for a more centralized storage of data in both managed and unmanaged networks corresponds to the creation of shared storage locations that can be found on one of the networked computing devices. The shared storage location can serve as a designated store for data, regardless of which of the networked computing devices a user is currently using. Although this approach attempts to centralize storage, it can

become deficient in that users are required to remember the location and name of the centralized shared storage location. Further, in the event that the shared storage location is renamed, moved or otherwise modified, users would be unable to access desired content.

Another approach for a more centralized storage of data files in managed networks corresponds to the maintenance of roaming user profiles that define where data is stored. Although a roaming user profile would assist in locating data files stored in various locations, the storage and update of the roaming profiles provides a greater burden on the processing and memory resources of the computer network. Accordingly, as the number of networked computing devices increases, the burden of a roaming user profile would also increase.

A further approach for a more centralized storage of data files in both managed and unmanaged networks corresponds to the redirection of local file system actions to a designated storage location on the network. For example, a request for the contents of a local storage location may be redirected automatically to retrieve the contents of a designated storage location on the network, which may be remote from the computing device. Similarly, a request to store data locally may be redirected automatically to store the selected content at the designated storage location on the network. Although this approach would allow all data requests to be directed towards a central storage location without requiring the user to know the location of the designated location, it prevents the storage and retrieval of content on the local machines. Further, in the event that the designated storage location is renamed, moved or otherwise modified, users would be unable to access desired content.

Thus, based upon the above-mentioned deficiencies associated with the small-scale networks there is a need for a system and method for managing content stored on two or more machines in a computer network.

SUMMARY OF THE INVENTION

A system and method for generating aggregated content views in a computing network are provided. A host computing device obtains a request for an aggregated view of content corresponding to a set of criteria. The host computing device queries itself and each computing device in a defined network for locally stored content matching the set of criteria. The query results are merged and displayed to a user at the host computing device as an aggregated list view.

In accordance with an aspect of the present invention, a method for managing data available for access on the network is provided. The method may be implemented in a computer network having two or more computing devices in communication. In accordance with the method a host computing device included as part of the computer network and associated with a user obtains a request to identify data corresponding to a set of criteria and obtains an identification of data stored on the host computing device associated with the user request and matching the set of criteria. The host computing device automatically obtains an identification of data stored on at least one computing device included in the computer network and matching the set of criteria. The host computing device merges the identification of data stored on the host computing device associated with the user request and the identification of data stored on at least one computing device included in the computer network. The host computing device then generates a result of the merging the identification of data stored on the host computing device associated with the user request and the identification of data stored on at least one computing device included in the computer network.

In accordance with another aspect of the present invention, a method for managing data available for access on the network is provided. The method may be implemented in a computer network having two or more computing devices in communication. In accordance with the method, a computing device obtains a user request to identify content stored on the two or more computing devices. The computing device automatically queries the two or more computing devices within the computer network to identify the contents of local computing device storage locations associated with a unique user identifier. The computing device merges the results of the queries. Additionally, the computing device displays the results of the merge query results.

In accordance with a further aspect of the present invention, a method for managing data available for access on the network is provided. The method may be implemented in a computer network having a computing device directly associated with a user and at least one remote computing device in communication. In accordance with the method, a computing device directly associated with a user obtains a request to identify data corresponding to a set of criteria. The computing device directly associated with a user obtains an identification of locally stored content matching the set of criteria and transmits a request to the remote computing device for an identification of content matching the set of criteria. The remote computing device obtains an identification of

locally stored content matching the set of criteria and transmits the identification of locally stored content matching the set of criteria. The computing device directly associated with the user then merges the content matching the set of criteria and generates a result of the merged content matching the set of criteria.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing aspects and many of the attendant advantages of this invention will become more readily appreciated as the same become better understood by reference to the following detailed description, when taken in conjunction with the accompanying drawings, wherein:

FIGURE 1 is a block diagram of an unmanaged computer network including three computing devices locally storing varied content in accordance with the present invention;

FIGURE 2 is a block diagram of the computer network of FIGURE 1 illustrating the initiation of a user data request and content query in accordance with the present invention;

FIGURE 3 is a block diagram of the computer network of FIGURE 1 illustrating the return and merging of matching content identifiers from the networked computing devices in accordance with the present invention;

FIGURE 4 is a block diagram of the computer network of FIGURE 1 illustrating the result of processing aggregated view requests at each computing device in the network in accordance with the present invention;

FIGURE 5 is a flow diagram illustrative of an aggregated view creation routine implemented by a networked computing device in accordance with the present invention;

FIGURES 6A-6D are block diagrams illustrative of screen displays for displaying merged content identifiers corresponding to an aggregated view in accordance with the present invention; and

FIGURE 7 is a flow diagram illustrative of an aggregated view remote document processing routine implemented by a networked computing device in accordance with the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Generally described, the present invention corresponds to a system and method for generating aggregated views of content stored in various computing devices in a computer network. More specifically, the present invention corresponds to a system and

method for automatically querying computing devices in a defined network and generating a view of the merged query responses. Although the present invention will be described with regard to a computing device network in which all the network computing devices have common security profiles for all authorized users of the network, one skilled in the relevant art will appreciate that the present invention may be implement in alternative computing networks in which each computing device may maintain separate security profiles for authorized network users. Further, one skilled in the relevant art will appreciate that some or all of the aspects of the present invention may be practiced in managed computer networks having a centralized network authority for regulating user access. Accordingly, the embodiments described with regard to the present invention are illustrative in nature and should not be construed as limiting.

FIGURE 1 is a block diagram illustrative of a computing device network 100 including three computing devices 102, 104, and 106 in accordance with the present invention. The computing devices 102, 104, 106 may be embodied as any one of a variety of devices that may be utilized to persist content to the computer network 100. Examples of computing devices include, but are not limited to, personal computing devices, hand-held computing devices, server-based computing devices, personal digital assistants, mobile telephones, stand-alone memory devices, electronic devices having some type of memory, and the like.

In an illustrative embodiment of the present invention, each computing device 102, 104 and 106 maintains a database of security information 108 that corresponds to a unique security identifier ("SID") for each authorized user. The database 108 may also maintain information associating unique user SIDs to one or more groups of users, such as administrator, users, reviewers, and a unique security identifier for the group. In accordance with the illustrative embodiment, the security information database 108 is replicated at each of the computing devices 102, 104, 106 such that a user's security profile, including a user SID, is accessible at any one of the computing devices. A more detailed explanation of the replication of security profile information across a computing network can be found in commonly assigned, co-pending application 10/414,354, entitled SMALL-SCALE SECURED COMPUTER NETWORK GROUP WITHOUT CENTRALIZED MANAGEMENT, and filed on April 15, 2003. U.S. Patent Application No. 10/414,354 is incorporated by reference herein.

As illustrated in FIGURE 1, for a particular user, a first computing device 102, "MACHINE A", maintains a first file 110, "FILE A", on a local storage location, such as a hard drive or other mass storage device, that can be access in some manner by the user. Similarly, a second computing device 104, "MACHINE B" maintains a second file 112, "File B" and a third machine 106, "MACHINE C", maintains a third file 114, "File C", both of which can be accessed in some manner by the user. As will be explained in additional detail below, in an illustrative embodiment of the present invention, each file 110, 112, 114 maintains an access control list ("ACL") that indicates which user SIDs and group SIDs that may have access to the particular file. Additionally, each file ACL may also indicate limitations to the type of access a particular user SID or group SID may have. For example, a file ACL may indicate that a particular user may access a file to read the file, but cannot modify the file in any manner. Further, each file ACL may also indicate a current owner of the file corresponding to the user SID that created the file and/or last edited the file. However, under conventional data management approaches, the user would have to remember the precise location of each piece of data to access the data remotely. Although the management of the data is described with regard to an ACL, one skilled in the relative art will appreciate that additional data management, such as security encryption, digital rights management, and the like, may also be utilized.

With reference now to FIGURES 2-4, the initiation and processing of an aggregated view request among the three computing devices 102, 104 and 106 of the computer network 100 will be described. With reference to FIGURE 2, the process is initiated with the receiving a content identification request at MACHINE A 102 by the user. Based on a unique SID associated with a particular user, or user account, MACHINE A 102 initiates a local content query corresponding the unique SID and any one of a variety of criteria that may used to select data. Concurrently with the local content query, or after the processing of the local query, MACHINE A 102 automatically initiates a distributed content request to the computing devices 104, 106 in the network 100. Because the security profiles in each computing device in the illustrative network have replicated security profile databases 108, MACHINE A 102 may issue identical content queries to each computing device. Further, although FIGURE 2 illustrates that every computing device on the network 100 receives the distributed content query, MACHINE A may also limit the content requests to a subset of computing devices within the network 100.

With reference now to FIGURE 3, each computing device receiving the content request processes the request and returns information identifying the matching content and the location of the content. For example, in the illustrative example, MACHINE B 104 would return information identifying File B 112 and its specific location within MACHINE B's storage system. Likewise, MACHINE C 106 would return information identifying File C 114 and its specific location within MACHINE C's storage system. In an illustrative embodiment of the present invention, the content located at each computing device is not transmitted to the requesting computing device. Further, as will be explained in greater detail below, previously stored content request responses or cached content request responses may also be processed.

Once the query responses are received by MACHINE A 102, the query results are merged. In an illustrative embodiment of the present invention, the merged results are embodied as single representation of all the matching content without direct reference to the location, or origin, of the content. However, MACHINE A 102 may obtain additional instructions from the user to sort, or otherwise organize, the merged results for the requesting user. For example, MACHINE A 102 may sort the merged content by size, date, origin, and the like. Still further, MACHINE A 102 may associated additional visual cues to the merged results to identify specific pieces of data, such as data stored on a computing device that is not currently available.

FIGURE 4 is a block diagram of the computer network 100 of FIGURE 1 illustrating the result of processing aggregated view requests at each of the computing devices in the network. For purposes of FIGURE 4, it is assumed that each of the computing devices 102, 104, 106 issued aggregated view requests and that the content of each computing device has not been modified. The aggregated view for MACHINE A 102 shows File A 110, File B 112, and File C 114 with File B and File C illustrated in dotted lines to show that they are stored remotely from MACHINE A. Similarly, the aggregated view for MACHINE B 104 shows the identical content with File A 110 and File C 114 illustrated in dotted lines to show that they are stored remotely. Finally, the aggregated view for MACHINE C 106 shows the identical content with File A 110 and File B 112 illustrated in dotted lines to show that they are stored remotely. Based on the aggregated views at each computing device 102, 104, 106, the user could access the same content without requiring knowledge of its actual location. Further, in the event that one or more of the computing devices are not available, the computing

device may still show the content stored on the unavailable machine with additional visual cues to indicate that it is not currently available.

FIGURE 5 is a flow diagram illustrative of an aggregated view creation routine 500 implemented by a network computing device, such as computing devices 102, 104, 106, in accordance with the present invention. At block 500, a user request for creating an aggregated view of content is obtained. In an illustrative embodiment of the present invention, the aggregated view request can correspond directly to a user's access to any type of application, or module, displaying content. For example, an aggregated view request may correspond to a request to open a file within a software application program. Similarly, an aggregated view creation request may correspond to the initiation of a file system management program. Still further, the aggregated view request may be automatically generated at the occurrence of specific events, such as at a particular time of day, after a computing device has powered up, after a computing device has connected to the network, and the like.

At block 502, the computing device queries the local computing device file system for content matching a content query. In an illustrative embodiment of the present invention, the content query is in the form of a database query, such as a query function support by Microsoft Corporation's SQL database. The database query includes a set of criteria for determining matching content. In an illustrative embodiment of the present invention, the query includes a request to match all content in which the user, identified by a user SID security token, can access. Additionally, the query can include any number of additional criteria, such as keyword searches, file types, date ranges, etc., that can be used to match content. The computing device storage system can then accept the query and identify any content that matches the query criteria. One skilled in the relevant art will appreciate that some or all portions of the query string, such as the user SID, may be entered automatically without requiring additional user input. Alternatively, some portion of the query string, such as the additional search criteria, may be specified by user at the time the search is requested or a pre-determined as part of a configuration of the computing device.

At decision block 506, a test is conducted to determine whether there are additional computing devices connected to the network 100. In an illustrative embodiment of the present invention, the test for the additional computing devices may correspond to a polling of every computing device connected to the network.

Alternatively, the test for additional computing devices may be limited to a specific subset of computing devices defined by a system administrator or the user. If there are additional computing devices within the network, at block 508, the computing device associated with the user sends content queries to the additional computing device. In an illustrative embodiment of the present invention, because the user's security profile, e.g., SID, is replicated to each computing device, the content query can be identical. Alternatively, if a user's security profile may change, each content query may be unique to accommodate for various user security identifiers. Blocks 506 and 508 will repeat until all the computing devices within the network 100, or subset of computing devices, has been queried. In accordance with an illustrative embodiment of the present invention, the results of the query are transmitted to the requesting computing device in the form of a content identifier and a path for locating the content, such as a Uniform Name Location ("UNC") path. As described above, the content does not need to be moved to the requesting computing device. Additionally, the requesting computing device does not need to make a copy of the content. Further, in an illustrative embodiment of the present invention, the computing device can maintain a copy of a previously received content request reply from the remote computer in a memory, such as RAM or cache. In addition to issuing a new content request from the remote computing device, the computing device may recall the previously received content request for the same remote computing device and begin processing the previously received content request. One the new content request is received, the computing device can update the response and store the new results for subsequent use.

Once all the networked computing devices have been queried, at decision block 510, a test is conducted to determine whether there are any previously networked computing devices that are not currently available. One skilled in the art will appreciate that any number of computing devices may be temporarily unavailable such as being powered down, out of wireless communication range or not currently docked to a network connection. If there are any previously networked computing devices that are not currently available, at block 512, the computing device attempts to recall a previously stored content request result for the unavailable computing device.

Once all of the computing devices have been queried, at block 514, the requesting computing device merges the result of the content queries as an aggregated list. In an illustrative embodiment of the present invention, the computing device may filter, or

otherwise process, the merged results. Additionally, the computing device may archive, cache or otherwise store, the merged results for use in subsequent aggregated list creation. At block 516, the merged content identifier view is generated and displayed to the user for manipulation. In an illustrative embodiment of the present invention, the results of the merging are represented to the users a flat view including content from all of the computing devices. Further, in the illustrative embodiment, the origin of the content is not initially displayed to the user. However, one skilled in the relevant art will appreciate that the requesting computing device can perform additional filtering or organization to display the matching content to the user. For example, the requesting computing device may sort the matching content by various attributes, such as data size, data type, title, date criteria, and the like. Further, the location of the origin of the content may be provided directly to the user, either as part of the initial view or as part of additional data available on the request of the user. Still further, in the event that one of the computing devices is unavailable, the merged view may be represented to the user with special visual cues, such as transparent icons, to indicate to the user the content that is likely stored on the unavailable computing device, but that can't current be accessed. In another embodiment, the computing device may utilize a previously stored record of the to compare to a current aggregated view. The computing device can then indicate which files have been modified, removed, or added. At block 518, routine 500 terminates.

FIGURES 6A-6D are block diagrams illustrative of various screen displays for displaying merged content identifiers in an aggregated list view in accordance with an aspect of the present invention. With reference to FIGURE 6A, the screen display 600 can include a first display portion 602 for allowing a user to select various file locations of interest. The screen display 600 can also include a second display portion 604 for displaying the specific content stored in a selected file space. With reference to the previous example of FIGURES 1-4, if a user were to initiate an aggregated view request by manipulating the graphical icon 606 corresponding to "My Documents", or otherwise initiating a request to view all files corresponding to the current user, the computing device would execute routine 500. The resulting three files 110, 112, 114 would be then displayed to the user in the second display portion 604 without need for the user to know the origin of the content. In contrast, under traditional file system management, the user would need to access each computing device file system independently to identify the location of each particular file of interest. For example, if files 110, 112 and 114

corresponded to icons 608, 610 and 612 respectively, the user would need to access each icon separately to view the same files created by the aggregated view.

With reference now to FIGURE 6B, one or more of the content identifiers in the second portion of the screen display 600 may include visual cues to assist the user. For example, if a computing device is not currently connected to the network and the requesting computing device has utilized a previously cached or archived content request, the content may distinguished visually as illustrated for file 112. In an illustrative embodiment of the present invention, unavailable content may be displayed with a dotted border and/or in a semi-transparent manner. In another example, content which has been determined by the computing device to have been modified since a previous aggregated view request may be highlighted to the user as illustrated for file 114. In an illustrative embodiment of the present invention, the modified piece of data may be highlighted on the screen display 600.

With reference now to FIGURES 6C and 6D, the computing device can automatically sort the merged matching content requests. Additionally, the computing device can obtain user input to sort the merged matching content requests. In one embodiment of the present invention, the merged matching content can be grouped according to the location of the content. With reference to FIGURE 6C, the screen display 600 can include additional identifiers 614, 616, 618 that indicate that the identity of the location. Additional levels of storage locations may be also be displayed on the screen display 600. In another embodiment of the present invention, the merged matching content can be grouped according to various attributes of the data. With reference to FIGURE 6D, the screen display can include a list view that displays the title of the data, the size of the data in storage and the type of data file. If the data were sorted by size as illustrated in FIGURE 6D, File A 620 would be at the top of the list, followed by File C 622, and File B 624. In this embodiment, a user could change the sorting criteria by manipulating the screen display 600. Although FIGURES 6C and 6D illustrate various sorting examples, one skilled in the relevant art will appreciate that additional sorting criteria and/or screen displays would also be included within the scope of the present invention.

With reference now to FIGURE 7, a routine 700 for manipulating content in an aggregated view will be described. At block 702, a request to manipulate content displayed in an aggregated view is obtained. One skilled in the relevant art will

appreciate that the manipulation of content can include opening a data file with a host software application, modifying content, deleting content, copying content and or moving the content. At block 704, the requesting computing device obtains a security identifier from the local security information database 108.

At decision block 706, a test is conducted to determine whether the selected content is stored locally. If the content is stored locally, at block 708, the computing device processes the request with the local security information to determine whether the user is authorized to manipulate the selected content. If the content is not stored locally, at decision block 710, a test is conducted to determine whether a copy of the requested data is maintained locally by the computing device. In an illustrative embodiment of the present invention, the computing device may maintain copies of previously accessed documents in a volatile memory, such as a computing device cache. If a copy of the file is available locally, the routine 700 proceeds to block 708 to process the request with local security information, as described above. If a copy of the requested file is not maintained locally, at block 712, the requesting computing device transmits the particular user's security information as a token along with a request to manipulate the content to the computing device actually storing the content. The receiving computing device then utilizes the security token to process the request and transmits a response. At block 714, the sending computing device receives the response from the receiving computing device. If the manipulation is authorized at block 708 or at block 712, the user is allowed to manipulate the content at block 716. The routine 700 terminates at block 718.

While illustrative embodiments of the invention have been illustrated and described, it will be appreciated that various changes can be made therein without departing from the spirit and scope of the invention.

The embodiments of the invention in which an exclusive property or privilege is claimed are defined as follows:

1. In a computer network having two or more computing devices in communication, a method for managing data available for access on the network, the method comprising:

obtaining, at a host computing device included as part of the computer network and associated with a user, a request to identify data corresponding to a set of criteria;

obtaining an identification of data stored on the host computing device associated with the user request and matching the set of criteria;

automatically obtaining an identification of data stored on at least one computing device included in the computer network and matching the set of criteria;

merging the identification of data stored on the host computing device associated with the user request and the identification of data stored on at least one computing device included in the computer network; and

generating a result of the merging the identification of data stored on the host computing device associated with the user request and the identification of data stored on at least one computing device included in the computer network.

2. The method as recited in Claim 1, wherein the request to identify data corresponding to the set of criteria corresponds to a request to view data associated with a unique identifier corresponding to the user.

3. The method as recited in Claim 2, wherein the request to identify data corresponding to the set of criteria corresponds to data in which the user is determined to be an owner of the data.

4. The method as recited in Claim 2, wherein the request to identify data corresponding to the set of criteria corresponds to data in which the user is determined to have permission to view the data.

5. The method as recited in Claim 1, wherein obtaining an identification of data stored on at least one computing device included in the computer network and matching the set of criteria includes generating a database query corresponding to a request to identify data corresponding to the set of criteria.

6. The method as recited in Claim 1, wherein obtaining an identification of data stored on at least one computing device included in the computer network and matching the set of criteria includes transmitting a security identifier corresponding to the user.

7. The method as recited in Claim 1, wherein the computer network includes three or more networked computers and wherein obtaining an identification of data stored on at least one computing device included in the computer network and matching the set of criteria includes obtaining an identification of data stored on each computer included in the computer network and matching the set of criteria.

8. The method as recited in Claim 1, wherein generating a result includes displaying the identification of the data without identifying the location of the data.

9. The method as recited in Claim 1 further comprising:

obtaining on the host computing device an indication to manipulate data corresponding to the result of the merging the identified data, wherein the data selected to be manipulated is remote from the host computing device;

transmitting a security identifier corresponding to the user;

obtaining permission to manipulate the selected data; and

manipulating the data selected to be manipulated.

10. The method as recited in Claim 1 further comprising maintaining a record of the result of the merging the identification of data stored on the host computing device associated with the user request and the identification of data stored on at least one computing device included in the computer network.

11. The method as recited in Claim 10 further comprising:

obtaining, at a host computing device, a second request to identify data corresponding to a set of criteria;

obtaining an identification of content stored on the host computing device associated with the user request and matching the set of criteria;

determining that the at least one computing device is not available to receive a content query;

recalling the record of the result of the merging the identification of data stored on the host computing device associated with the user request and the identification of data stored on at least one computing device included in the computer network;

merging the identification of data stored on the host computing device associated with the user request and the record of the result of the merging the identification of data stored on the host computing device associated with the user request and the identification of data stored on at least one computing device included in the computer network; and

generating a result of the merging the identification of data stored on the host computing device associated with the user request and the record of the result of the merging the identification of data stored on the host computing device associated with the user request and the identification of data stored on at least one computing device included in the computer network.

12. The method as recited in Claim 11, wherein generating a result of the merging the identification of data stored on the host computing device associated with the user request and the record of the result of the merging the identification of data stored on the host computing device associated with the user request and the identification of data stored on at least one computing device included in the computer network includes generating visual cues corresponding to the data not currently available to the user.

13. A computer-readable medium having computer-executable instructions for performing the method recited in Claim 1.

14. A computer system having a processor, a memory and an operating system, the computer system operable to perform the method recited in Claim 1.

15. In a computer network having two or more computing devices in communication, a method for managing data available for access on the network, the method comprising:

obtaining a user request to identify content stored on the two or more computing devices;

automatically querying the two or more computing devices within the computer network to identify the contents of local computing device storage locations associated with a unique user identifier;

merging the results of the queries; and

displaying the results of the merge query results.

16. The method as recited in Claim 15, wherein the request to identify content corresponding to the unique user identifier corresponds to content in which the user is determined to be an owner of the content .

17. The method as recited in Claim 15, wherein the request to identify content corresponding to the unique user identifier corresponds to content in which the user is determined to have permission to view the content.

18. The method as recited in Claim 15, wherein automatically querying the two or more computing devices within the computer network to identify the contents of local computing device storage locations corresponding to the unique user identifier includes generating a database query corresponding to a request to identify content corresponding to the unique user identifier.

19. The method as recited in Claim 15, wherein automatically querying the two or more computing devices within the computer network to identify the contents of local computing device storage locations corresponding to the unique user identifier includes processing a security identifier associated with the unique user identifier.

20. The method as recited in Claim 15, wherein the computer network includes three or more networked computers and wherein automatically querying the two or more computing devices within the computer network to identify the contents of local computing device storage locations corresponding to the unique user identifier includes querying each computer included in the computer network to identify the contents of local computing device storage locations corresponding to the unique user identifier.

21. The method as recited in Claim 15 further comprising:

obtaining a request to manipulate one or more pieces of content included in the merged results, wherein the content selected to be manipulated is remote from a computing obtaining the request;

transmitting a security identifier corresponding to the unique user identifier;

obtaining permission to manipulate the selected content; and

manipulating the content selected to be manipulated.

22. The method as recited in Claim 15 further comprising maintaining a record of the result of the merging the results of the queries.

23. The method as recited in Claim 22 further comprising:

obtaining a second user request to identify content stored on the two or more computing devices;

obtaining an identification of content stored on a computing device associated with the unique user identifier and matching the unique user identifier;

determining that the at least one computing device is not available to receive a content query;

recalling the record of the result of the merging of the results of the queries;

merging the identification of data stored on the computing device associated with the unique user identifier and the record of the result of the merging of the results of the queries; and

generating a result of the merging the identification of data stored on the computing device associated with the unique user identifier and the record of the result of the merging of the results of the queries.

24. The method as recited in Claim 23, wherein displaying the results of the merge query results includes generating visual cues corresponding to the data not currently available to the user.

25. A computer-readable medium having computer-executable instructions for performing the method recited in Claim 15.

26. A computer system having a processor, a memory and an operating system, the computer system operable to perform the method recited in Claim 15.

27. In a computer network having a computing device directly associated with a user and at least one remote computing device in communication, a method for managing data available for access on the network, the method comprising:

obtaining, by the computing device directly associated with a user, a request to identify data corresponding to a set of criteria;

obtaining, by the computing device directly associated with a user, an identification of locally stored content matching the set of criteria;

transmitting, by the computing device directly associated with a user, a request to the remote computing device for an identification of content matching the set of criteria;

obtaining, by the remote computing device, an identification of locally stored content matching the set of criteria;

transmitting, by the remote computing device, the identification of locally stored content matching the set of criteria

merging, by the computing device directly associated with the user, content matching the set of criteria; and

generating, by the computing device directly associated with the user, a result of the merged content matching the set of criteria.

28. The method as recited in Claim 27, wherein the request to identify data corresponding to the set of criteria corresponds to a request to view data associated with a unique identifier correspond to the user.

29. The method as recited in Claim 28, wherein the request to identify data corresponding to the set of criteria corresponds to data in which the user is determined to be an owner of the data.

30. The method as recited in Claim 28, wherein the request to identify data corresponding to the set of criteria corresponds to data in which the user is determined to have permission to view the data.

31. The method as recited in Claim 27 further comprising transmitting, by the computing device directly associated with the user, a security identifier corresponding to the user.

32. The method as recited in Claim 27, wherein the computer network includes a second networked computer remote from the user, the method further comprising:

transmitting, by the computing device directly associated with a user, a request to the second remote computing device for an identification of content matching the set of criteria;

obtaining, by the second remote computing device, an identification of locally stored content matching the set of criteria;

transmitting, by the second remote computing device, the identification of locally stored content matching the set of criteria; and

merging, by the computing device directly associated with the user, content matching the set of criteria.

33. The method as recited in Claim 27, wherein generating the merged content matching the set of criteria includes displaying the identification of the content without identifying the location of the content.

34. The method as recited in Claim 27 further comprising:

obtaining, by the computing device directly associated with the user, an indication to manipulate data corresponding to the result of the merged content, wherein the data selected to be manipulated is stored on the remote computing device;

transmitting, by the computing device directly associated with the user, the request to manipulate the selected data and a security identifier to the remote computing device;

obtaining, by the remote computing device, the request to manipulate the selected data and the security identifier corresponding to the user;

processing, by the remote computing device, the request to manipulate the selected data and the security identifier to determine whether the user is authorized to manipulate the selected data;

transmitting, by the remote computing device, permission to manipulate the selected data ;

obtaining, by the computing device directly associated with the user, the permission; and

manipulating, by the computing device directly associated with the user, the selected data.

35. The method as recited in Claim 27 further comprising maintaining a record of the result of the merging of content matching the set of criteria.

36. The method as recited in Claim 35 further comprising:

obtaining, by the computing device directly associated with the user, a second request to identify data corresponding to a set of criteria;

obtaining, by the computing device directly associated with the user, an identification of locally stored content matching the set of criteria;

determining, by the computing device directly associated with the user, that the remote computing device is not available to receive a content query;

recalling, by the computing device directly associated with the user, the record of the result of the merging of content matching the set of criteria;

merging, by the computing device directly associated with the user, the identification of locally stored content matching the set of criteria and the record of the result of the merging of content matching the set of criteria; and

generating a result of the merging of the identification of locally stored content matching the set of criteria and the record of the result of the merging of content matching the set of criteria.

37. The method as recited in Claim 36, wherein generating a result of the merging of the identification of locally stored content matching the set of criteria and the record of the result of the merging of content matching the set of criteria includes generating visual cues corresponding to the data not currently available to the user.

38. A computer-readable medium having computer-executable instructions for performing the method recited in Claim 27.

39. A computer system having a processor, a memory and an operating system, the computer system operable to perform the method recited in Claim 27.

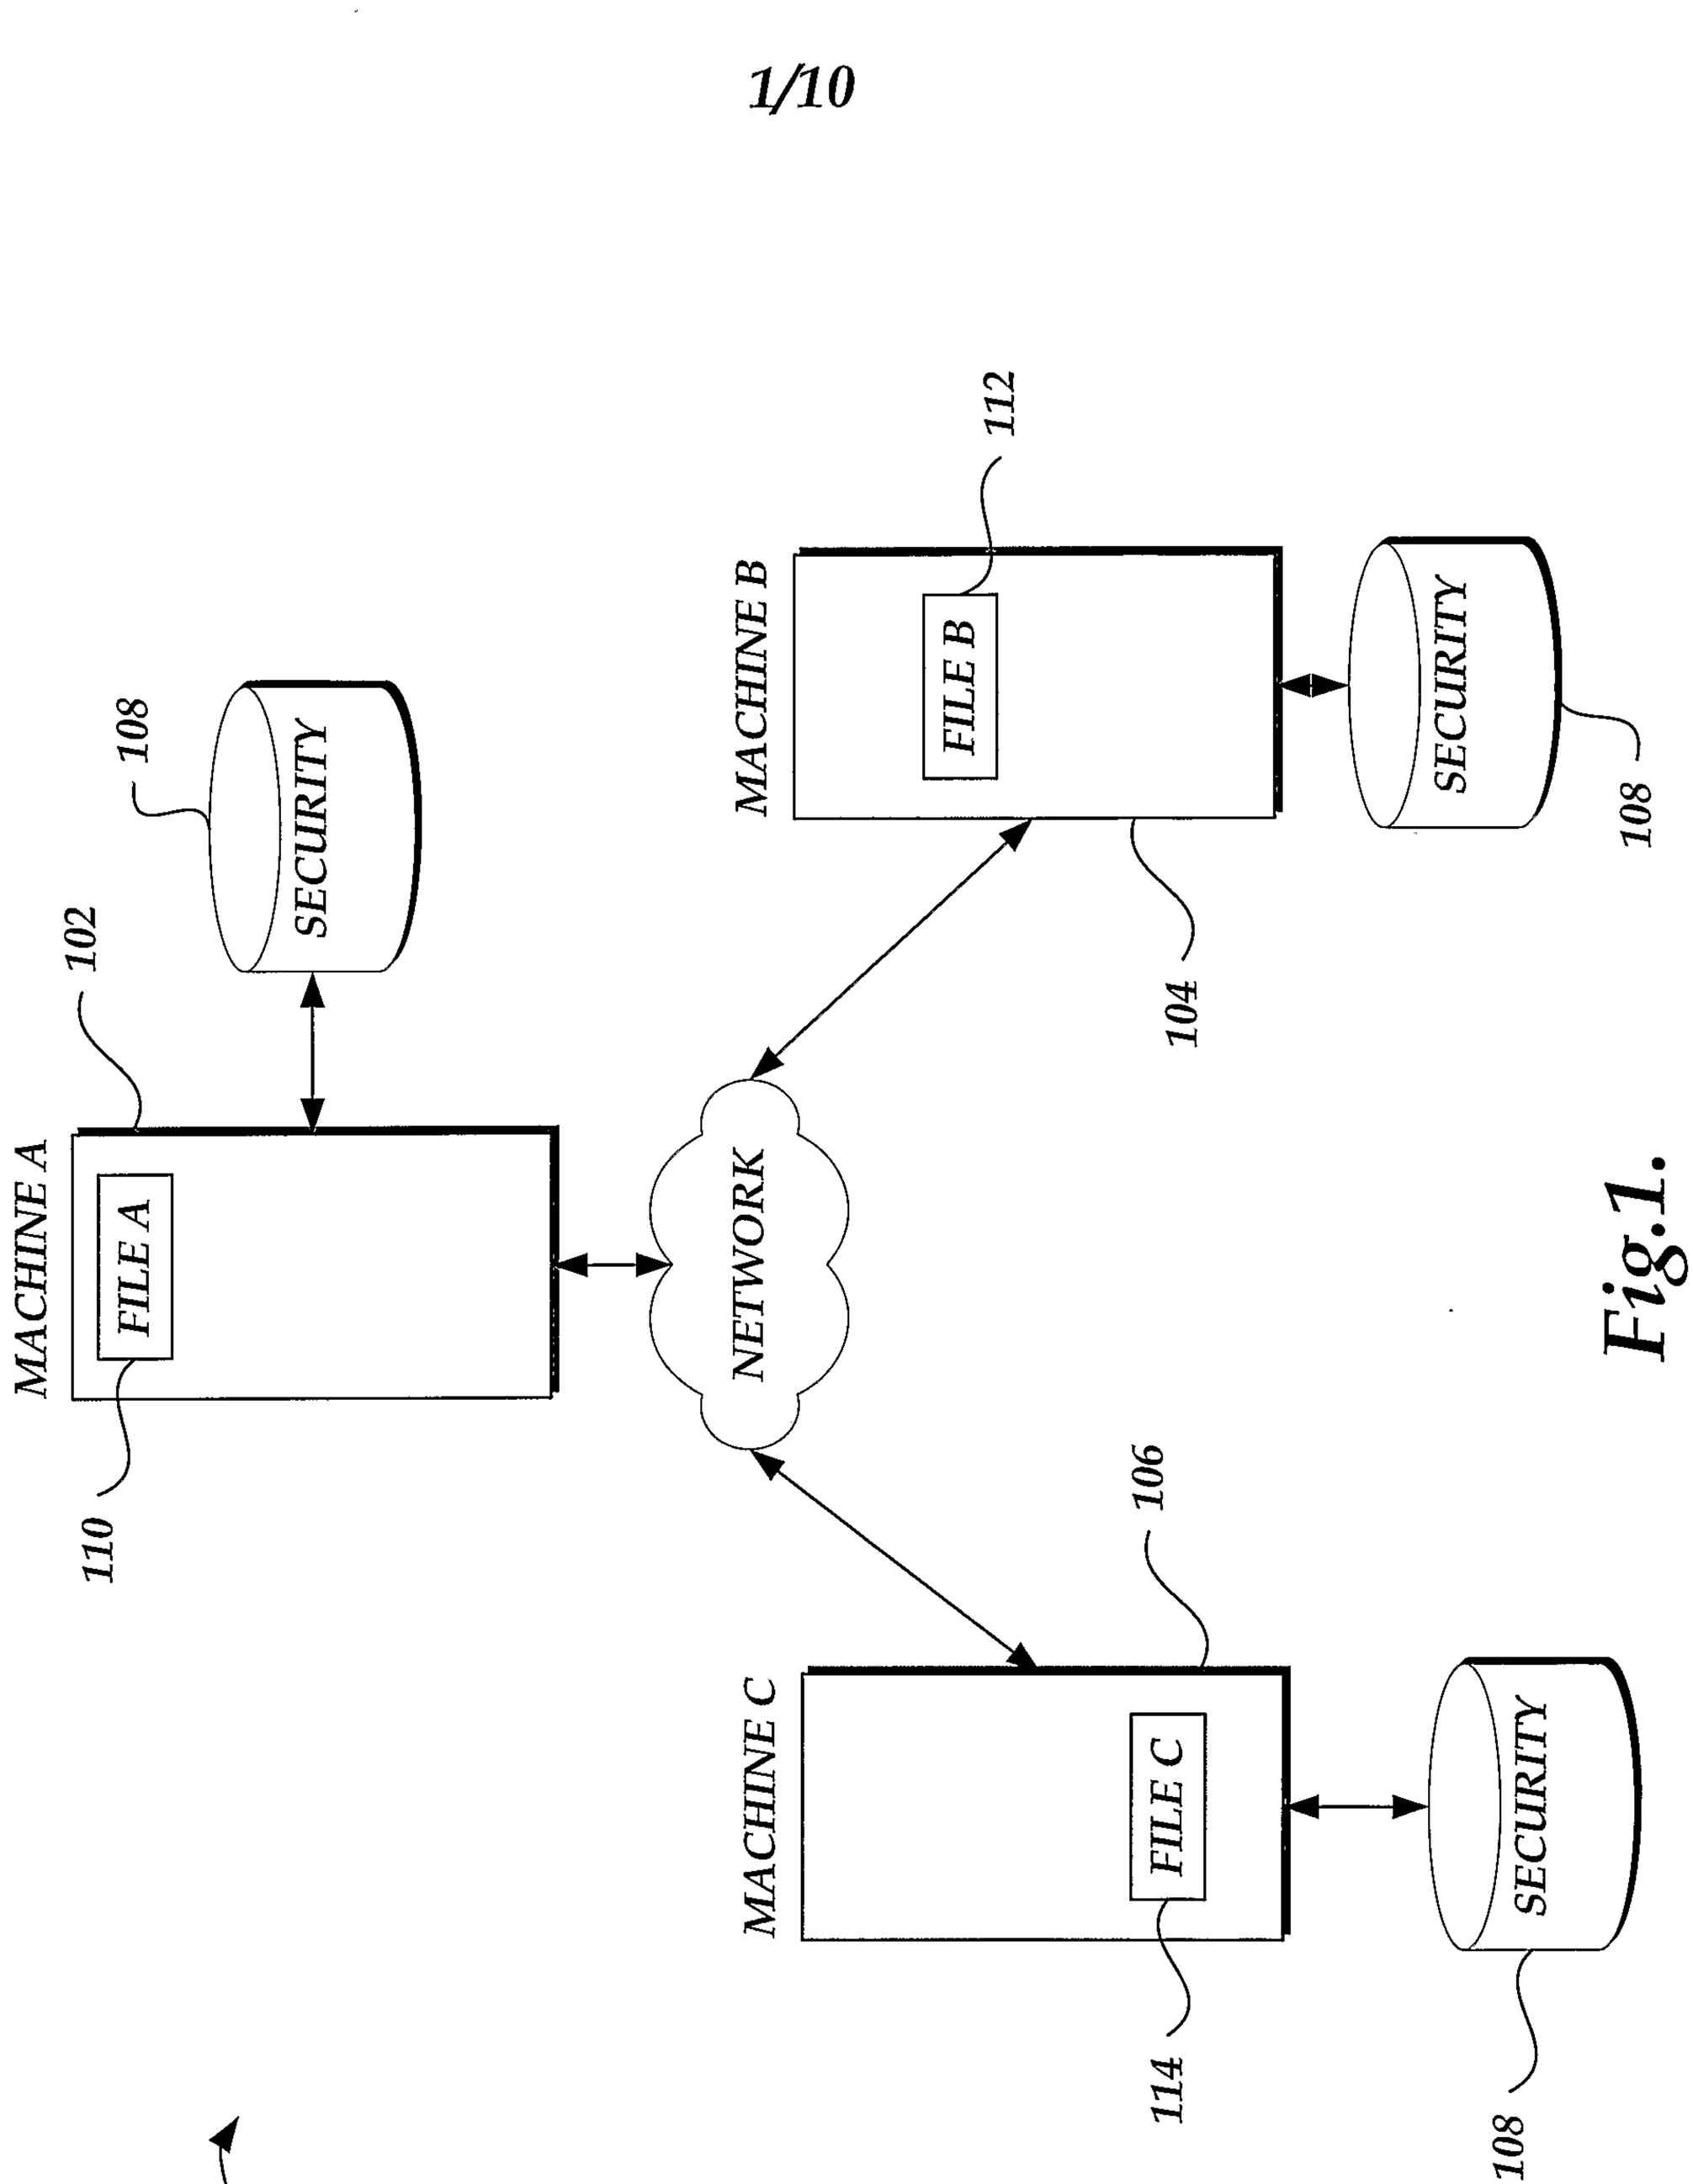
100
MACHINE A
102
FILE A
110
SECURITY
108
NETWORK
MACHINE B
112
FILE B
104
SECURITY
108
MACHINE C
106
FILE C
114
SECURITY
108

Fig.1.

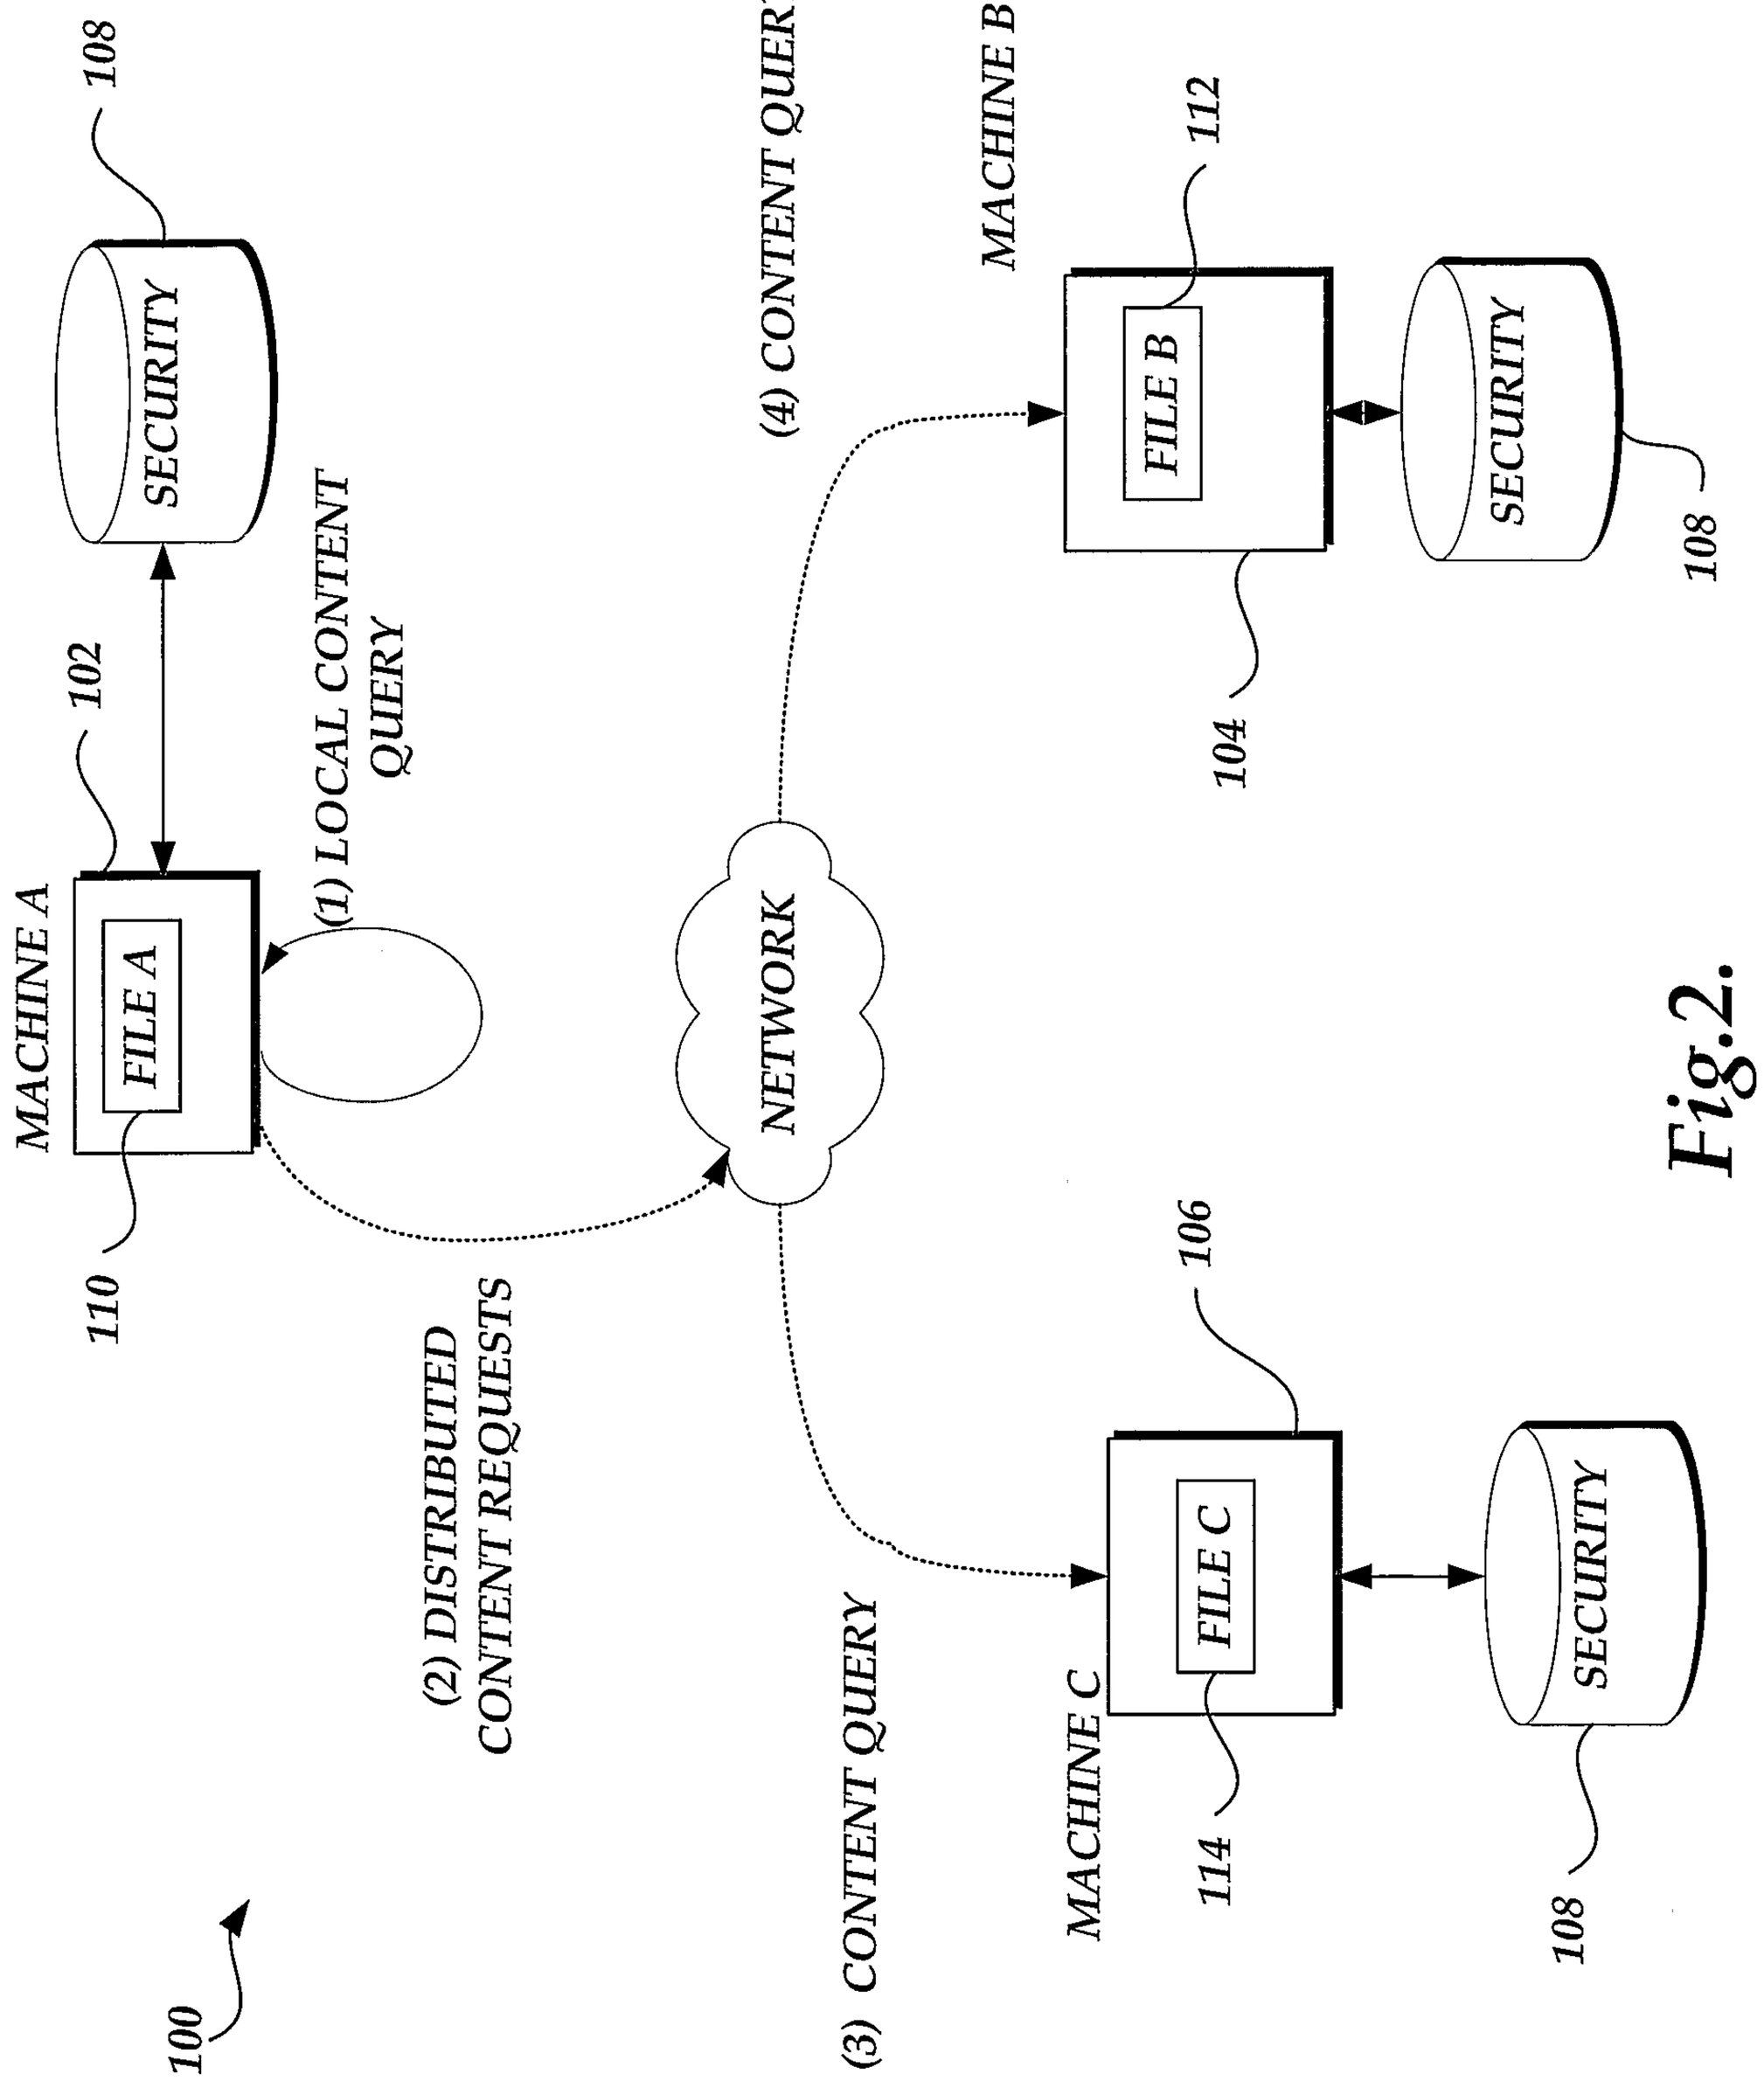
100
MACHINE A
102
FILE A
SECURITY
108
(1) LOCAL CONTENT QUERY
110
(2) DISTRIBUTED CONTENT REQUESTS
NETWORK
(3) CONTENT QUERY
(4) CONTENT QUERY
MACHINE B
112
FILE B
SECURITY
108
104
MACHINE C
106
FILE C
SECURITY
108
114

Fig.2.

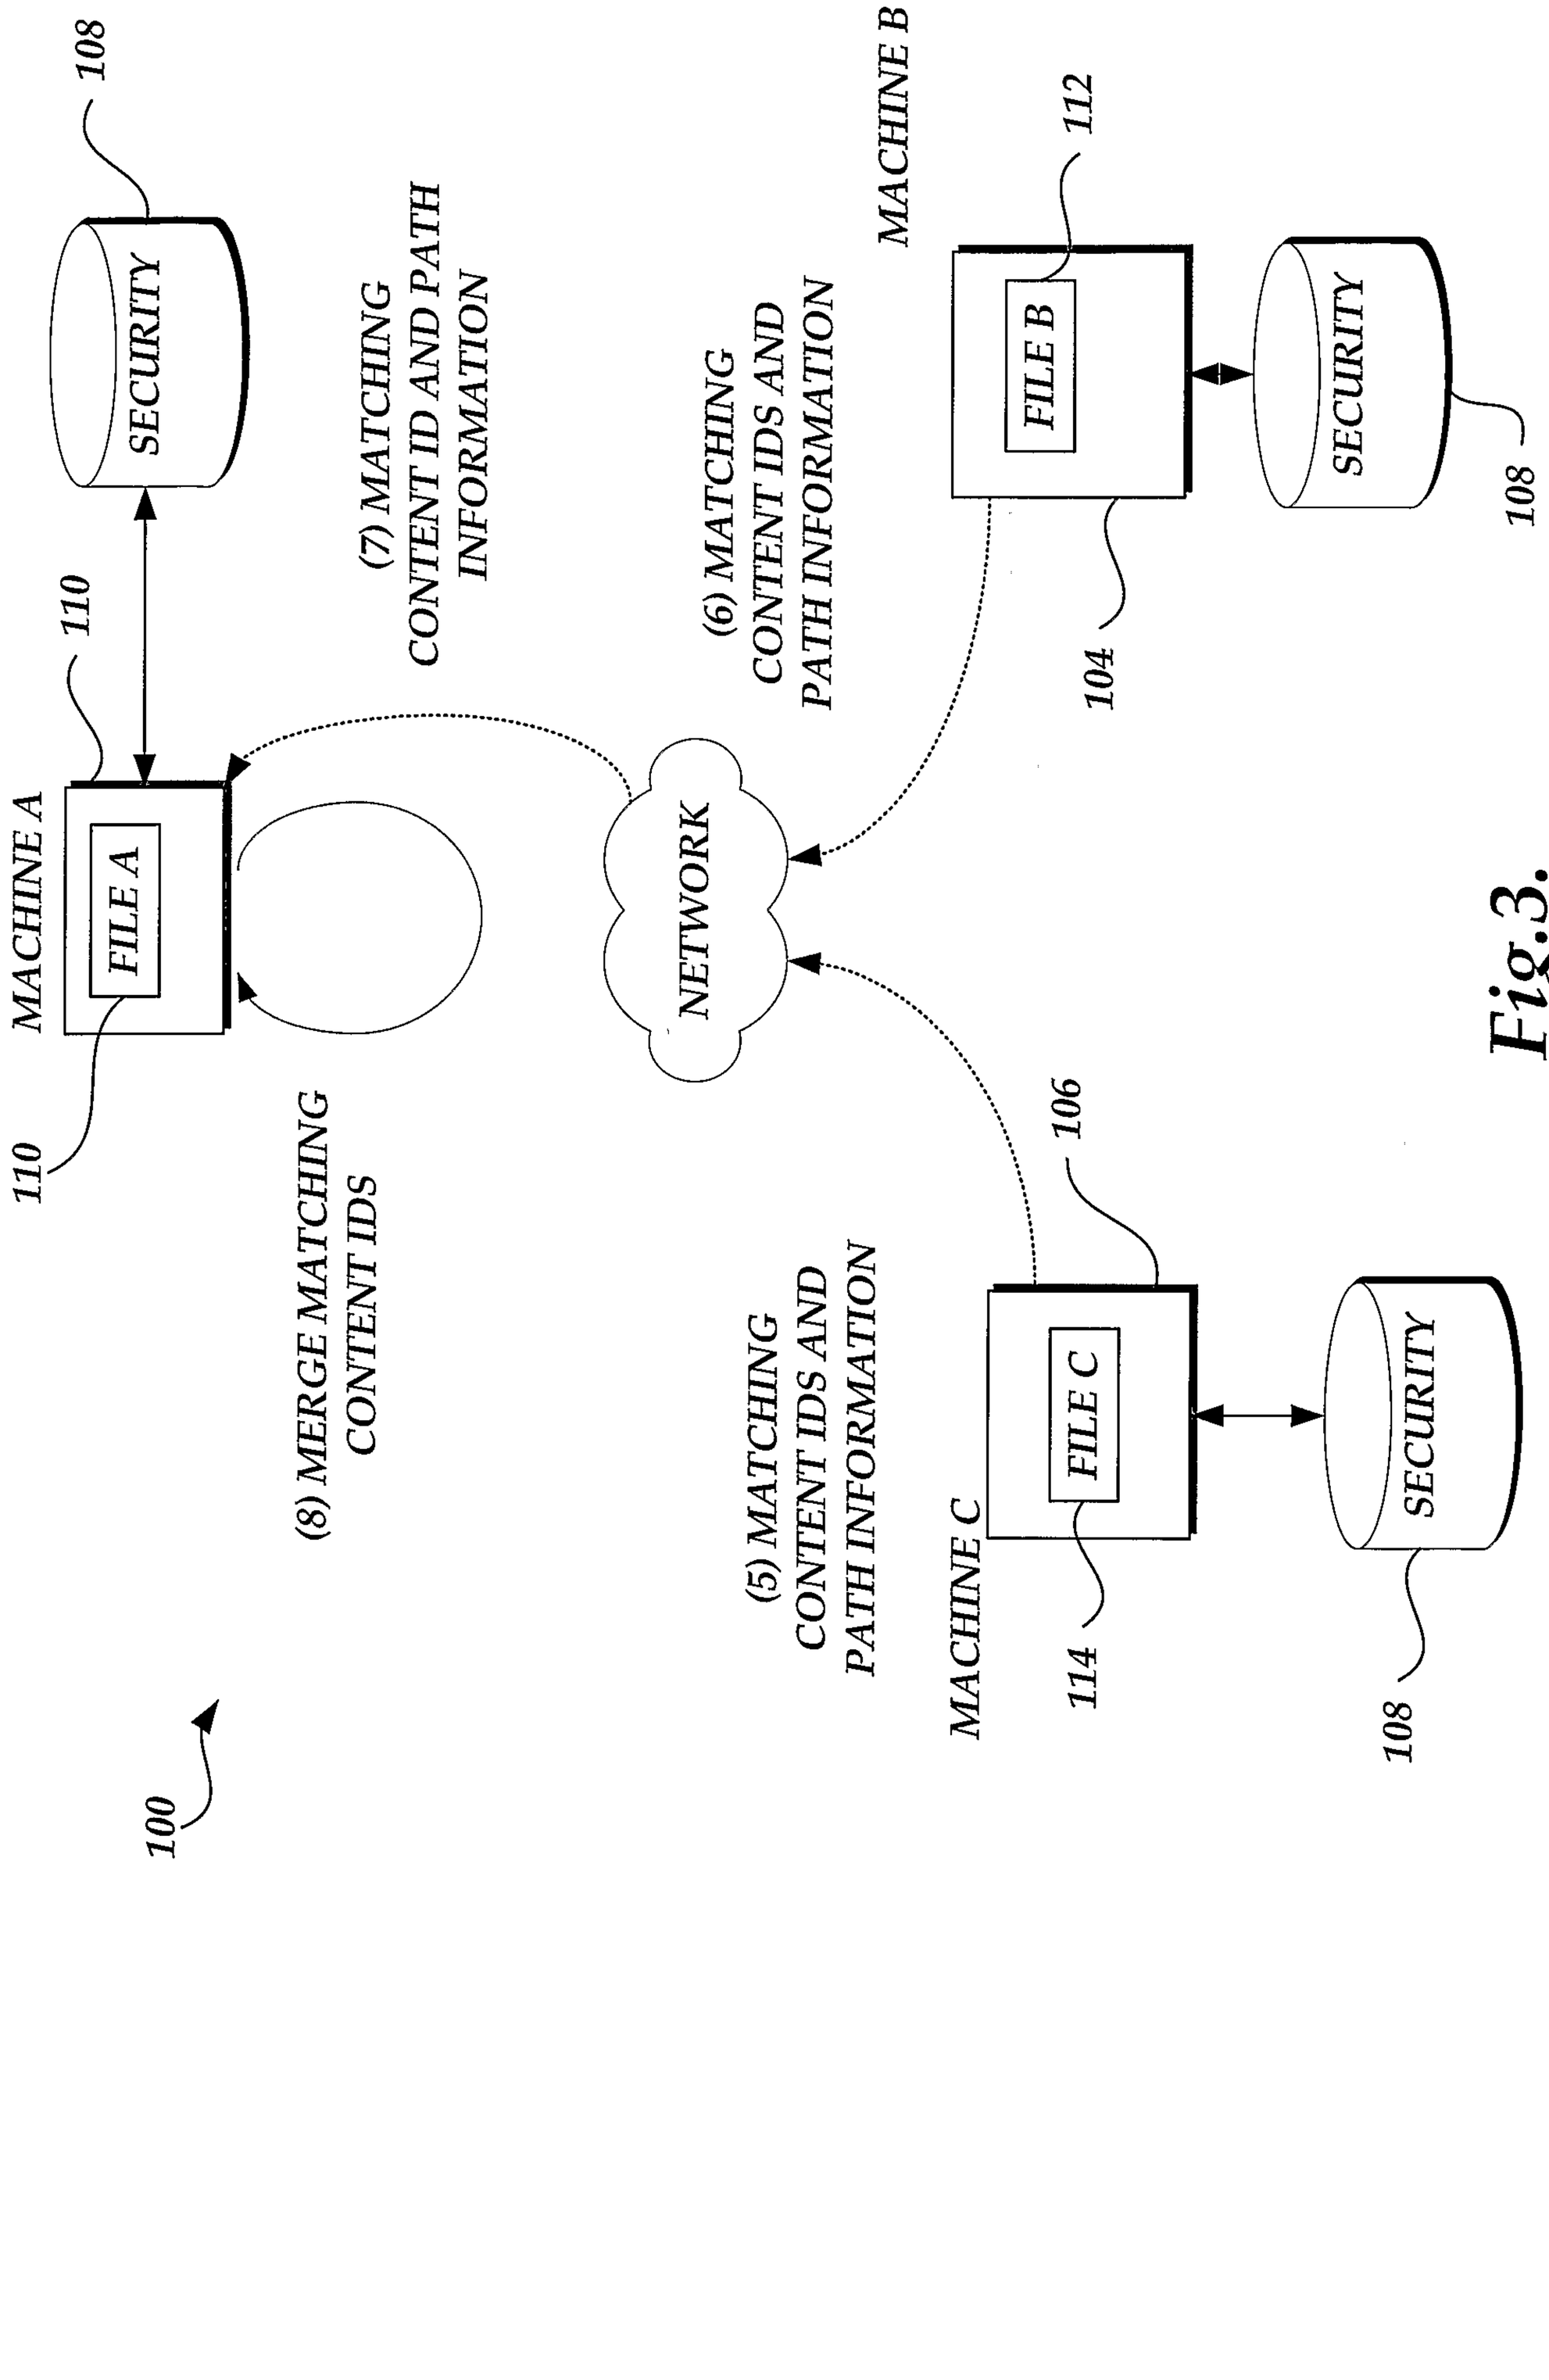
100
MACHINE A
110
FILE A
110
SECURITY
(8) MERGE MATCHING CONTENT IDS
(7) MATCHING CONTENT ID AND PATH INFORMATION
(6) MATCHING CONTENT IDS AND PATH INFORMATION
(5) MATCHING CONTENT IDS AND PATH INFORMATION
NETWORK
MACHINE B
112
FILE B
104
SECURITY
108
MACHINE C
114
FILE C
106
SECURITY
108

Fig.3.

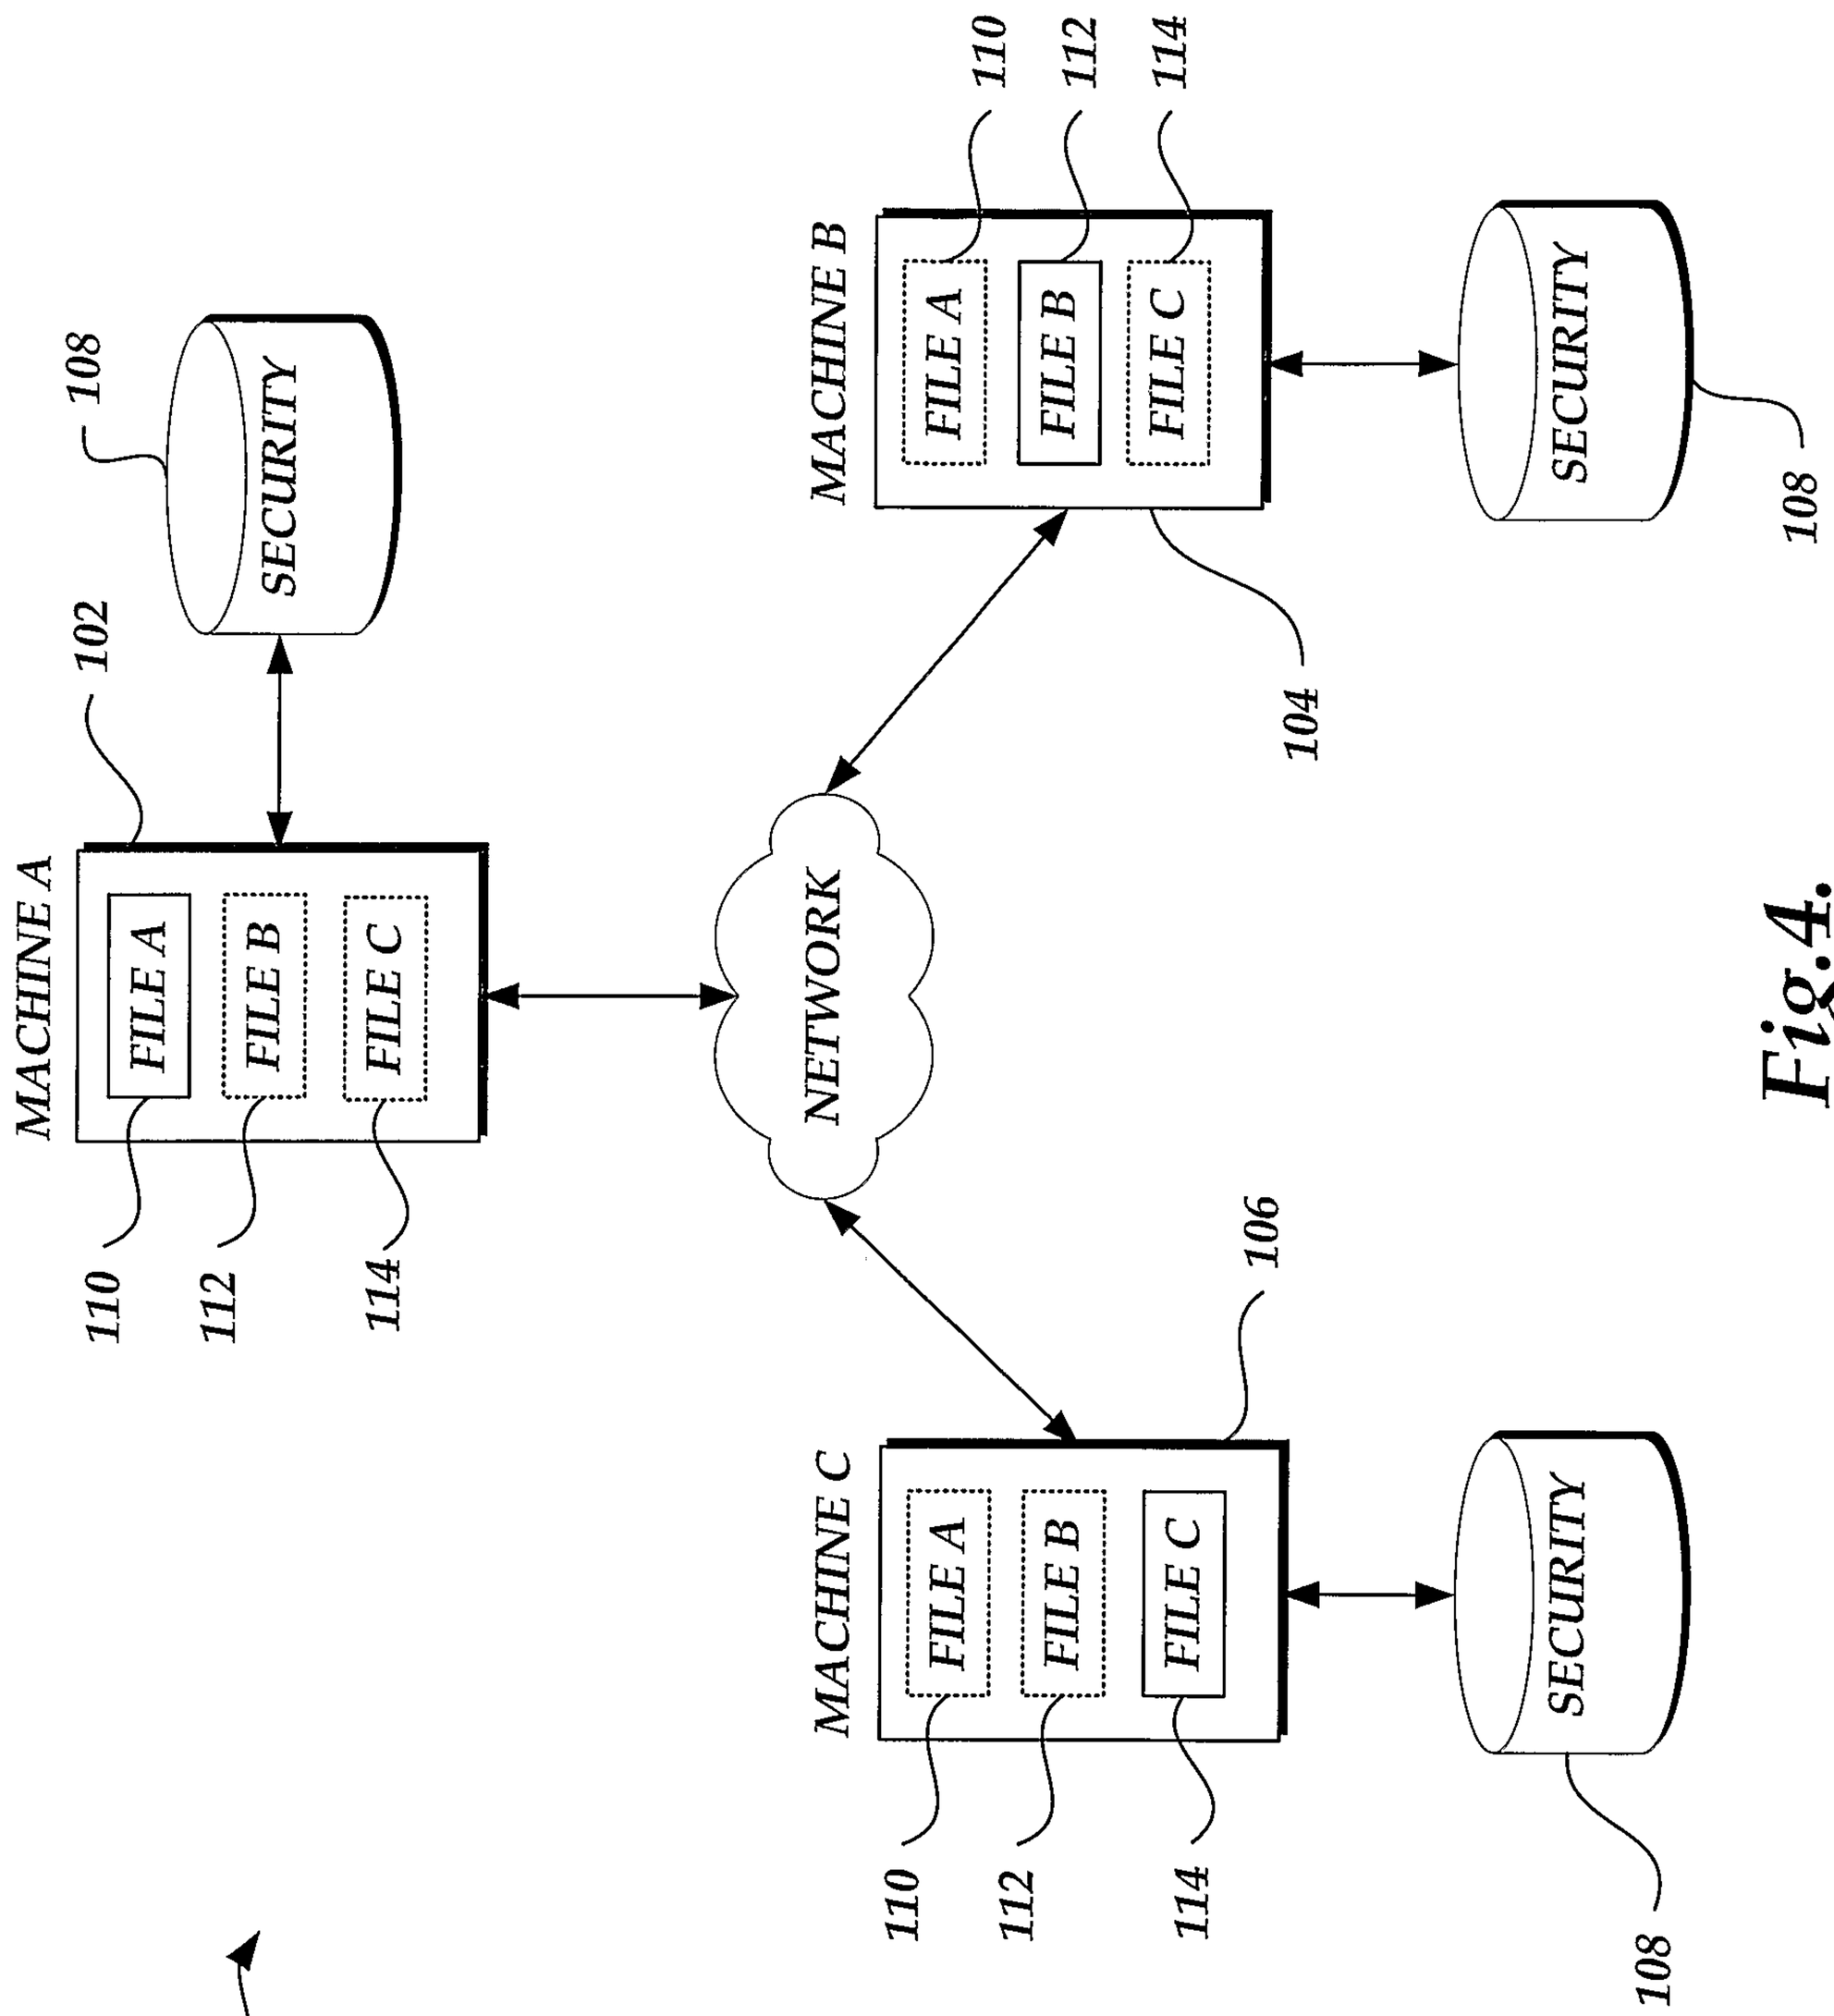
100
MACHINE A
FILE A
FILE B
FILE C
110
112
114
102
SECURITY
108
NETWORK
MACHINE B
FILE A
FILE B
FILE C
104
SECURITY
108
MACHINE C
FILE A
FILE B
FILE C
106
SECURITY
108

Fig.4.

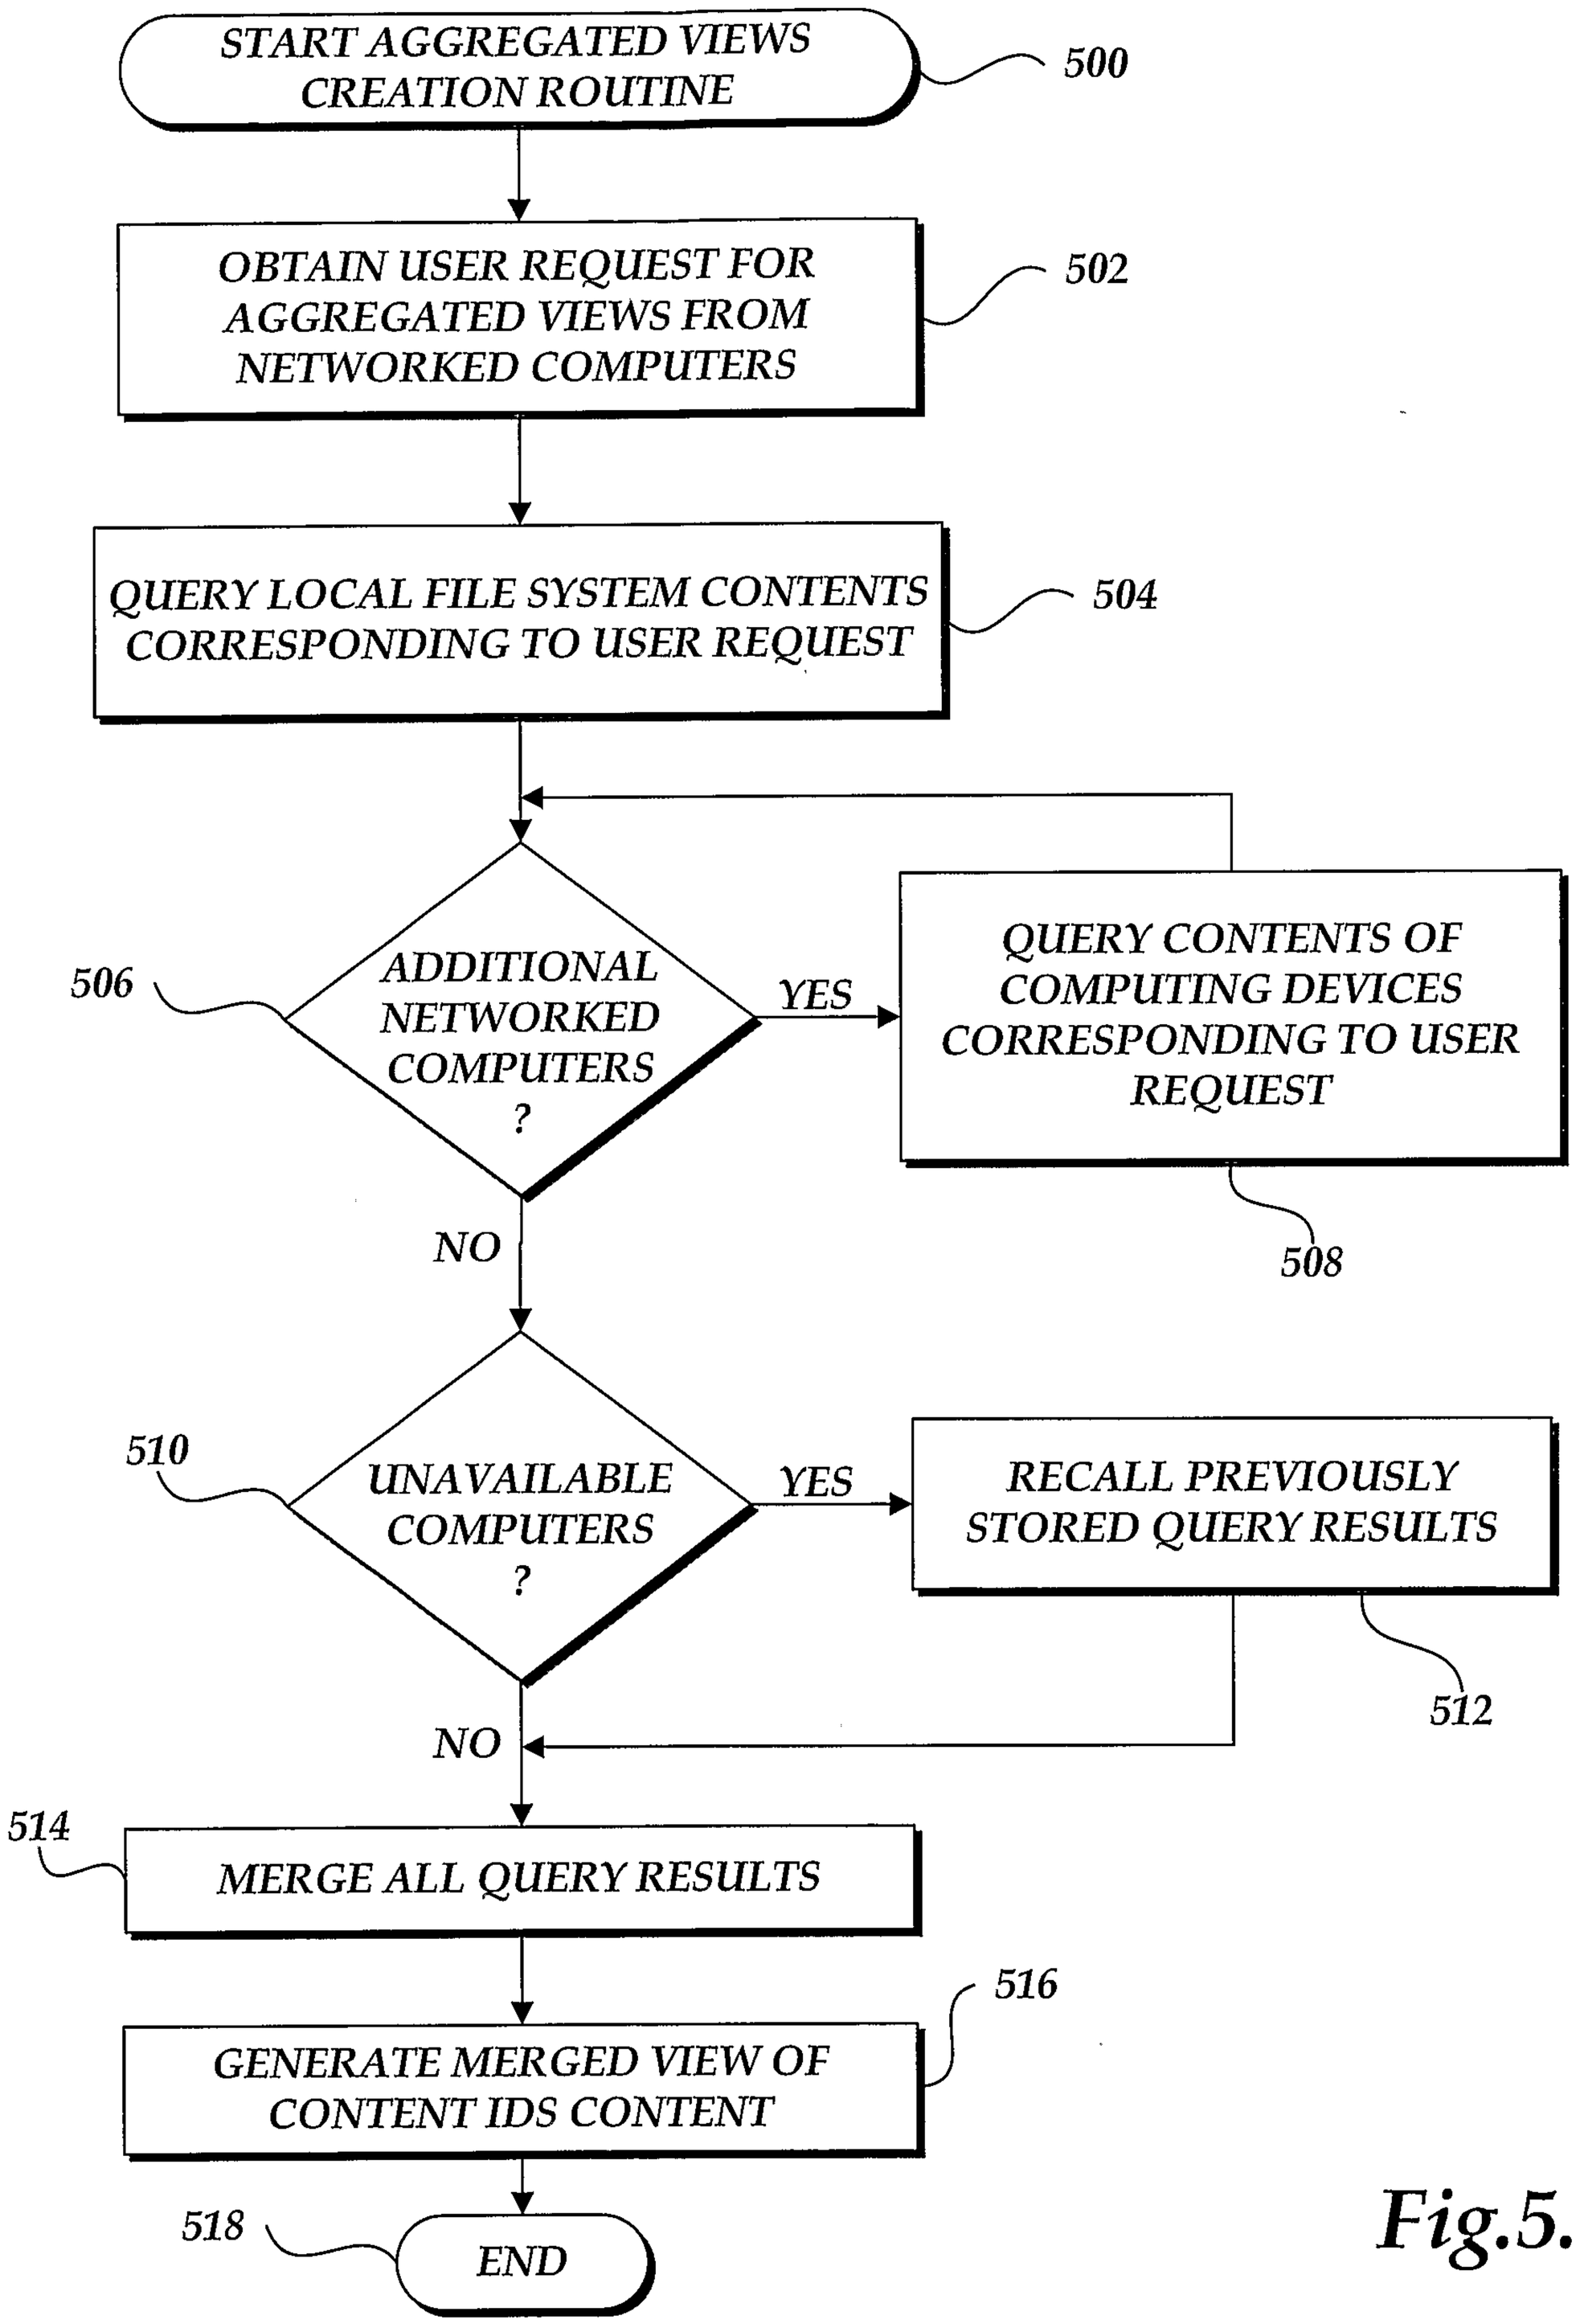
START AGGREGATED VIEWS CREATION ROUTINE
500
OBTAIN USER REQUEST FOR AGGREGATED VIEWS FROM NETWORKED COMPUTERS
502
QUERY LOCAL FILE SYSTEM CONTENTS CORRESPONDING TO USER REQUEST
504
506
ADDITIONAL NETWORKED COMPUTERS ?
YES
QUERY CONTENTS OF COMPUTING DEVICES CORRESPONDING TO USER REQUEST
508
NO
510
UNAVAILABLE COMPUTERS ?
YES
RECALL PREVIOUSLY STORED QUERY RESULTS
512
NO
514
MERGE ALL QUERY RESULTS
516
GENERATE MERGED VIEW OF CONTENT IDS CONTENT
518
END

Fig.5.

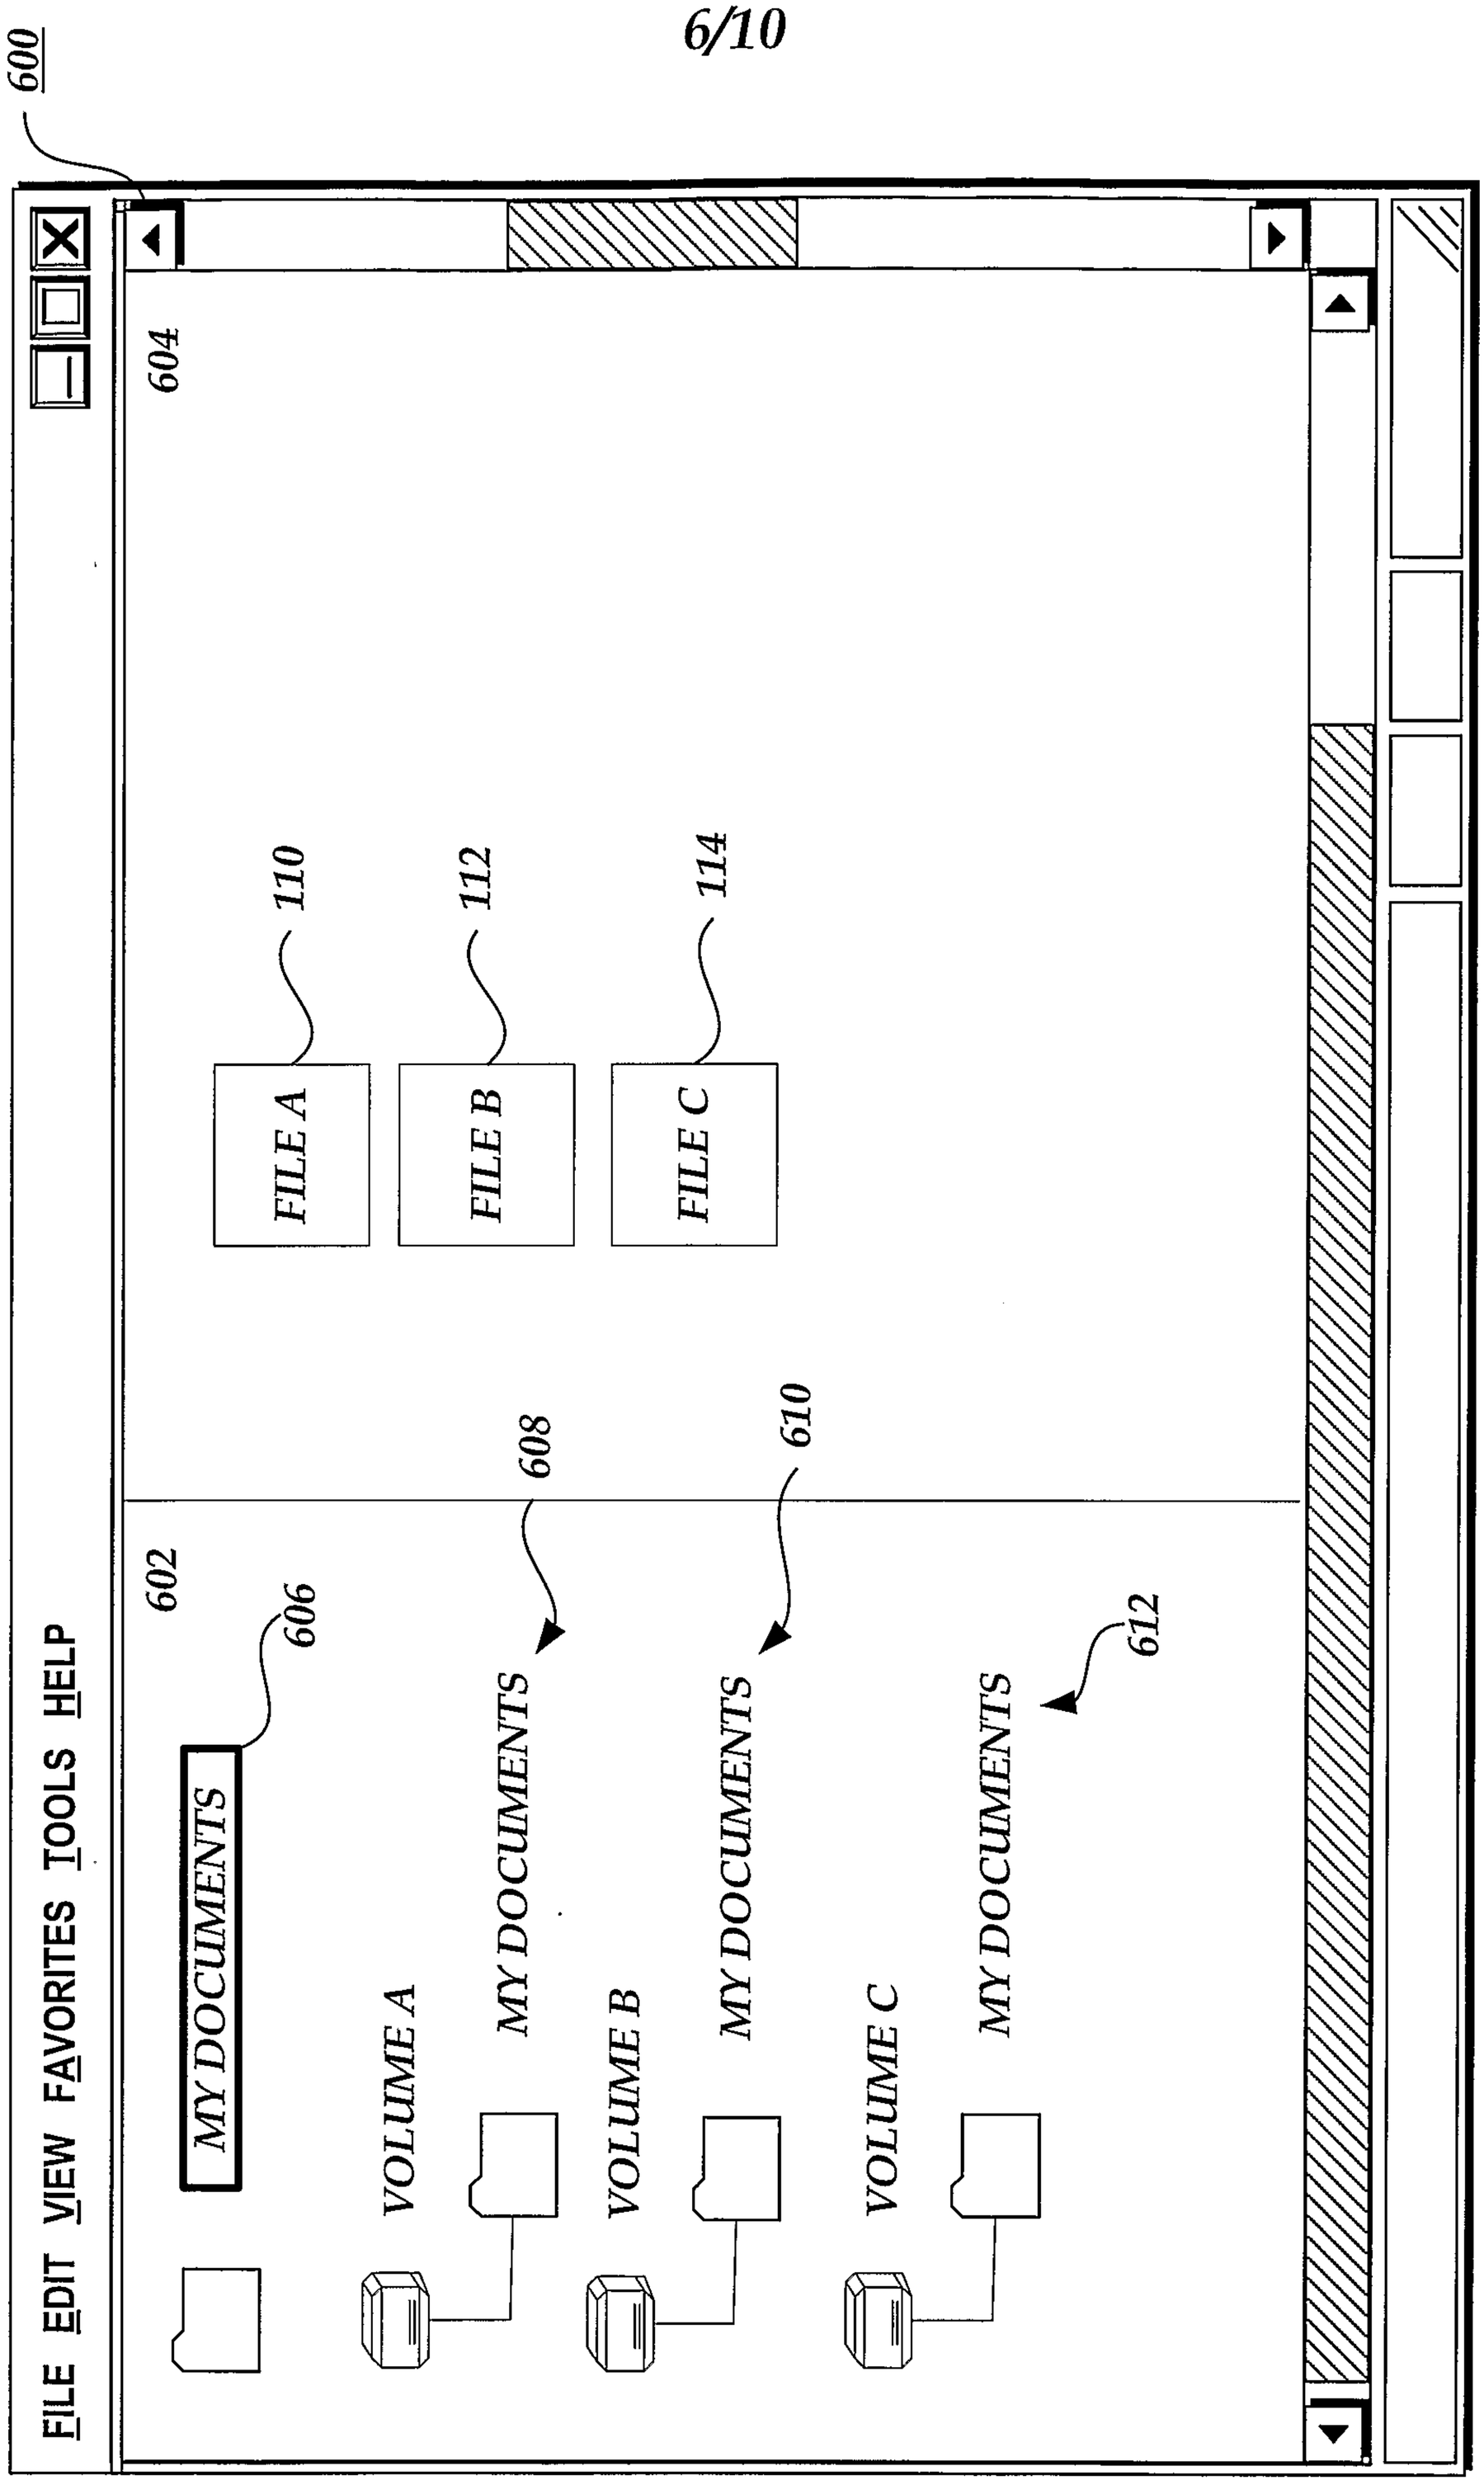
600
604
110
112
114
FILE A
FILE B
FILE C
602
606
608
610
612
FILE EDIT VIEW FAVORITES TOOLS HELP
MY DOCUMENTS
MY DOCUMENTS
MY DOCUMENTS
MY DOCUMENTS
VOLUME A
VOLUME B
VOLUME C

Fig.6A.

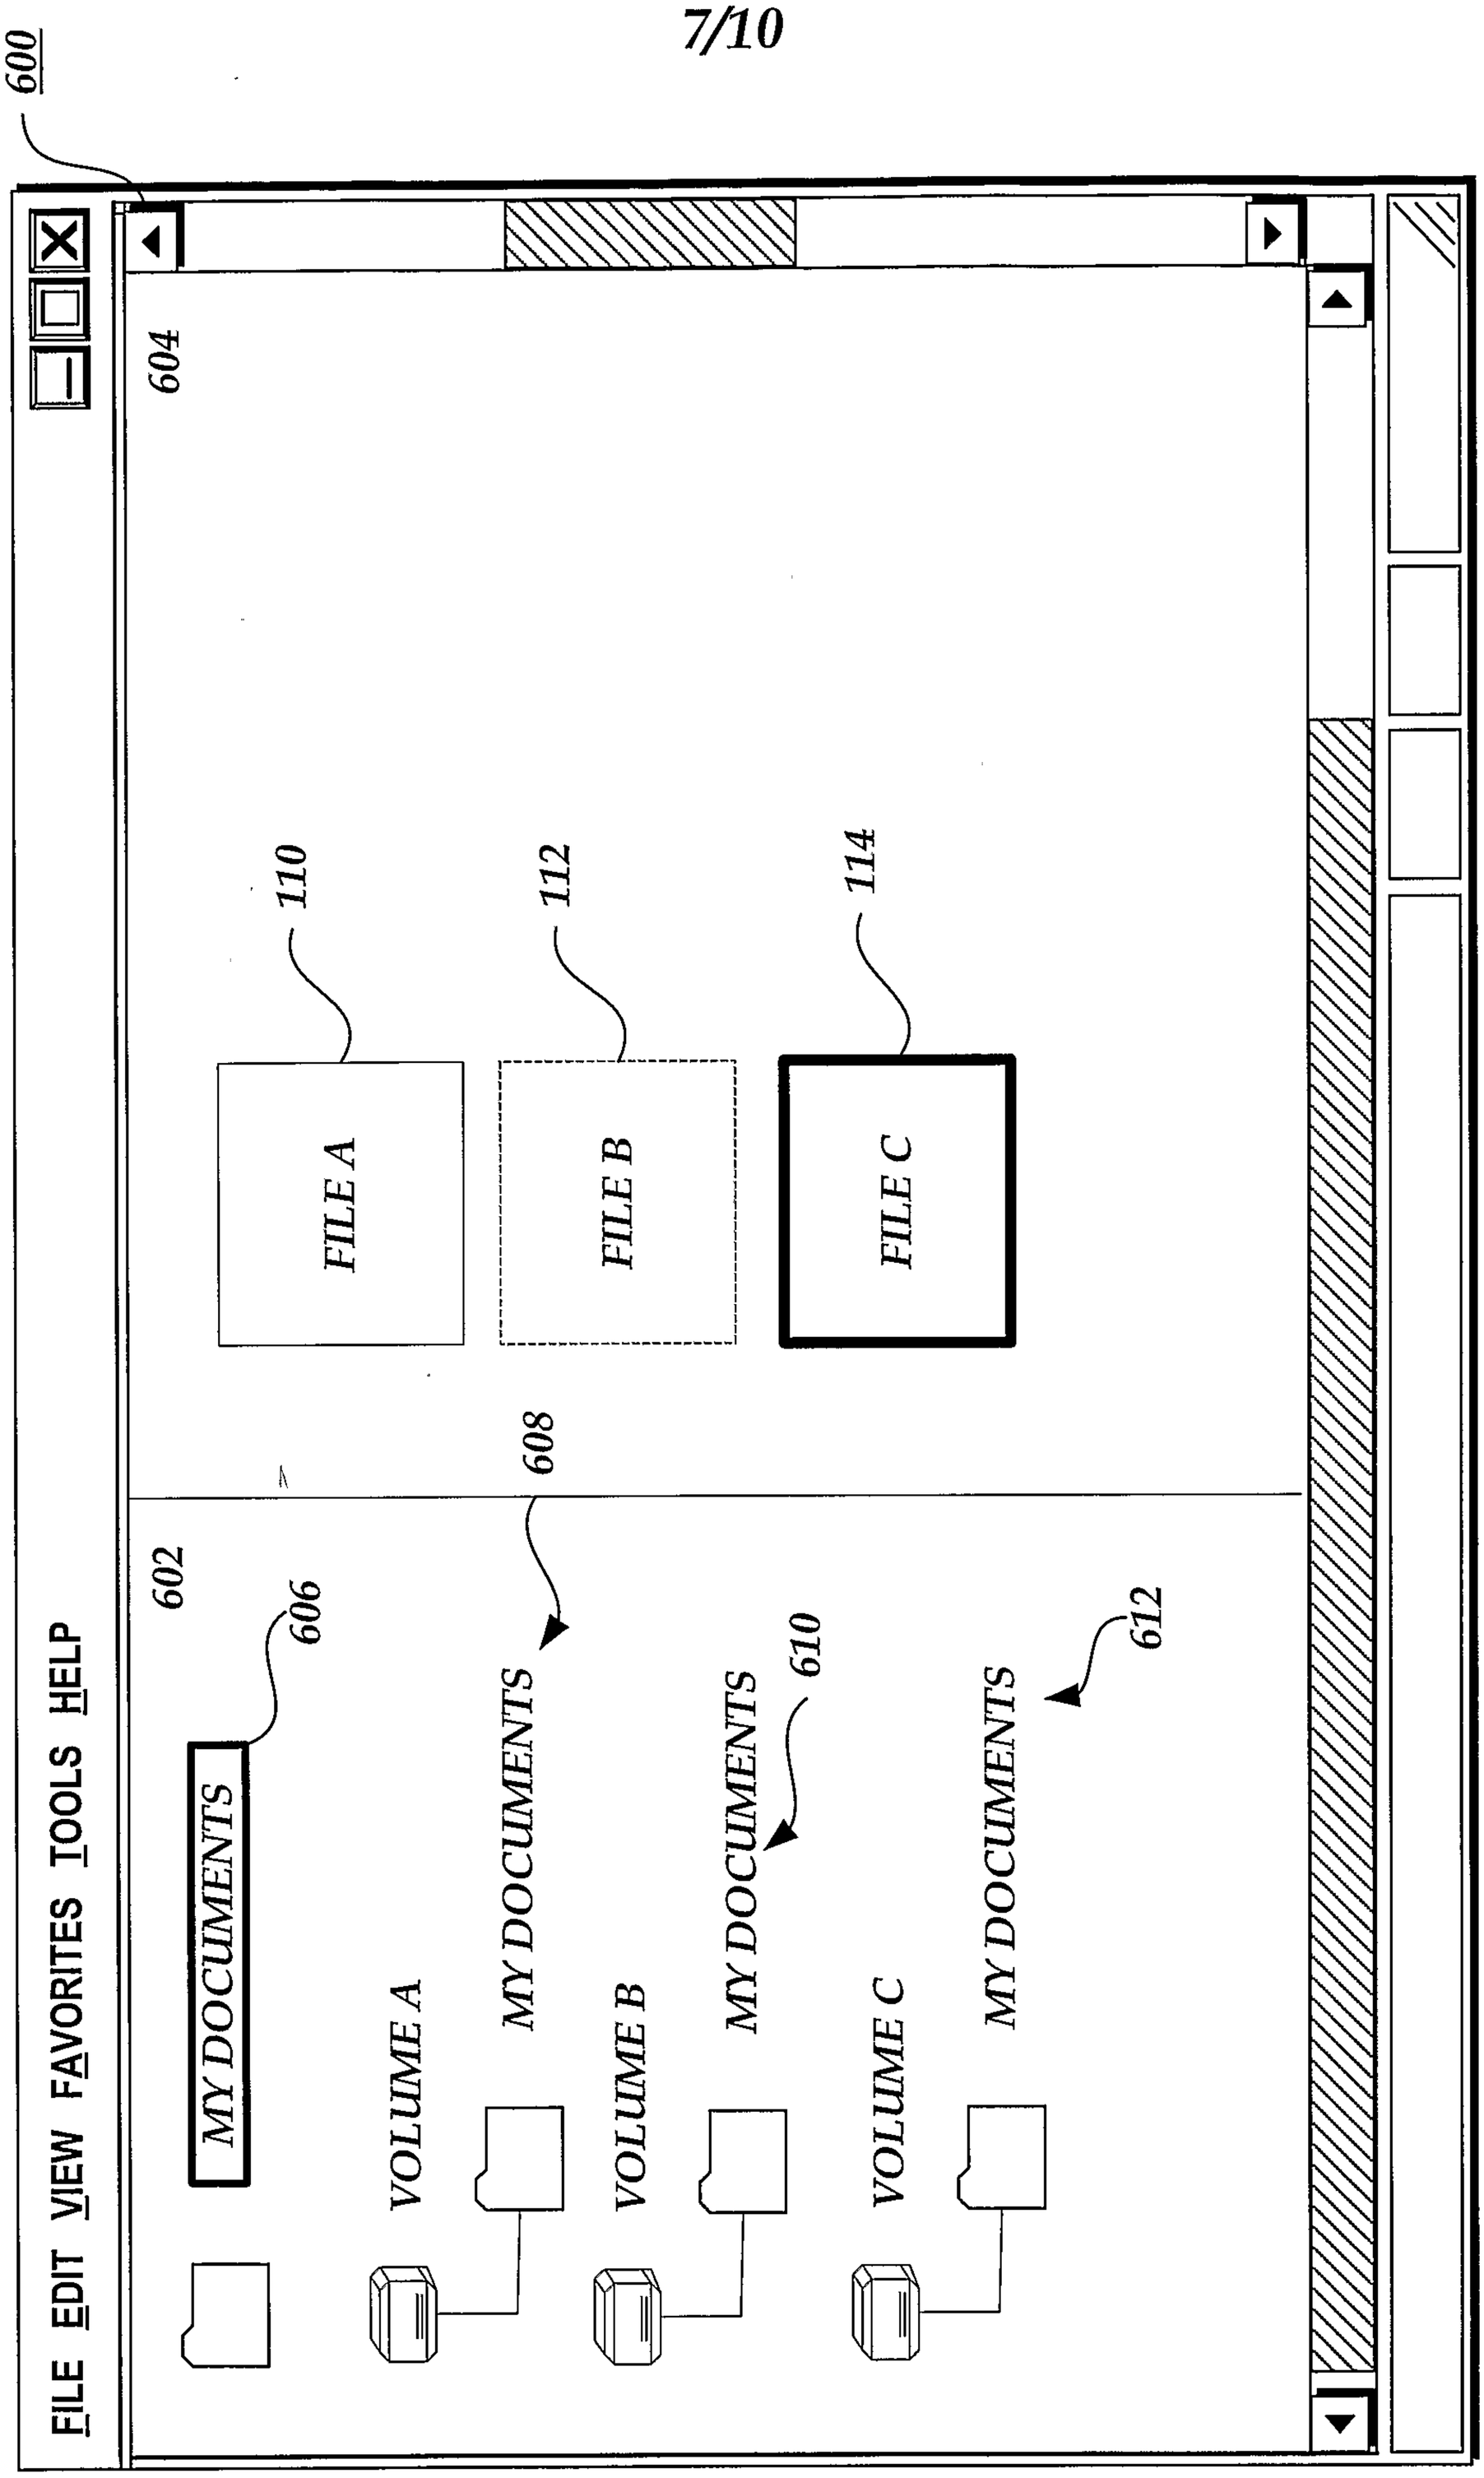
600
604
110
112
114
FILE A
FILE B
FILE C
602
606
608
610
612
MY DOCUMENTS
MY DOCUMENTS
MY DOCUMENTS
MY DOCUMENTS
VOLUME A
VOLUME B
VOLUME C
FILE EDIT VIEW FAVORITES TOOLS HELP

Fig.6B.

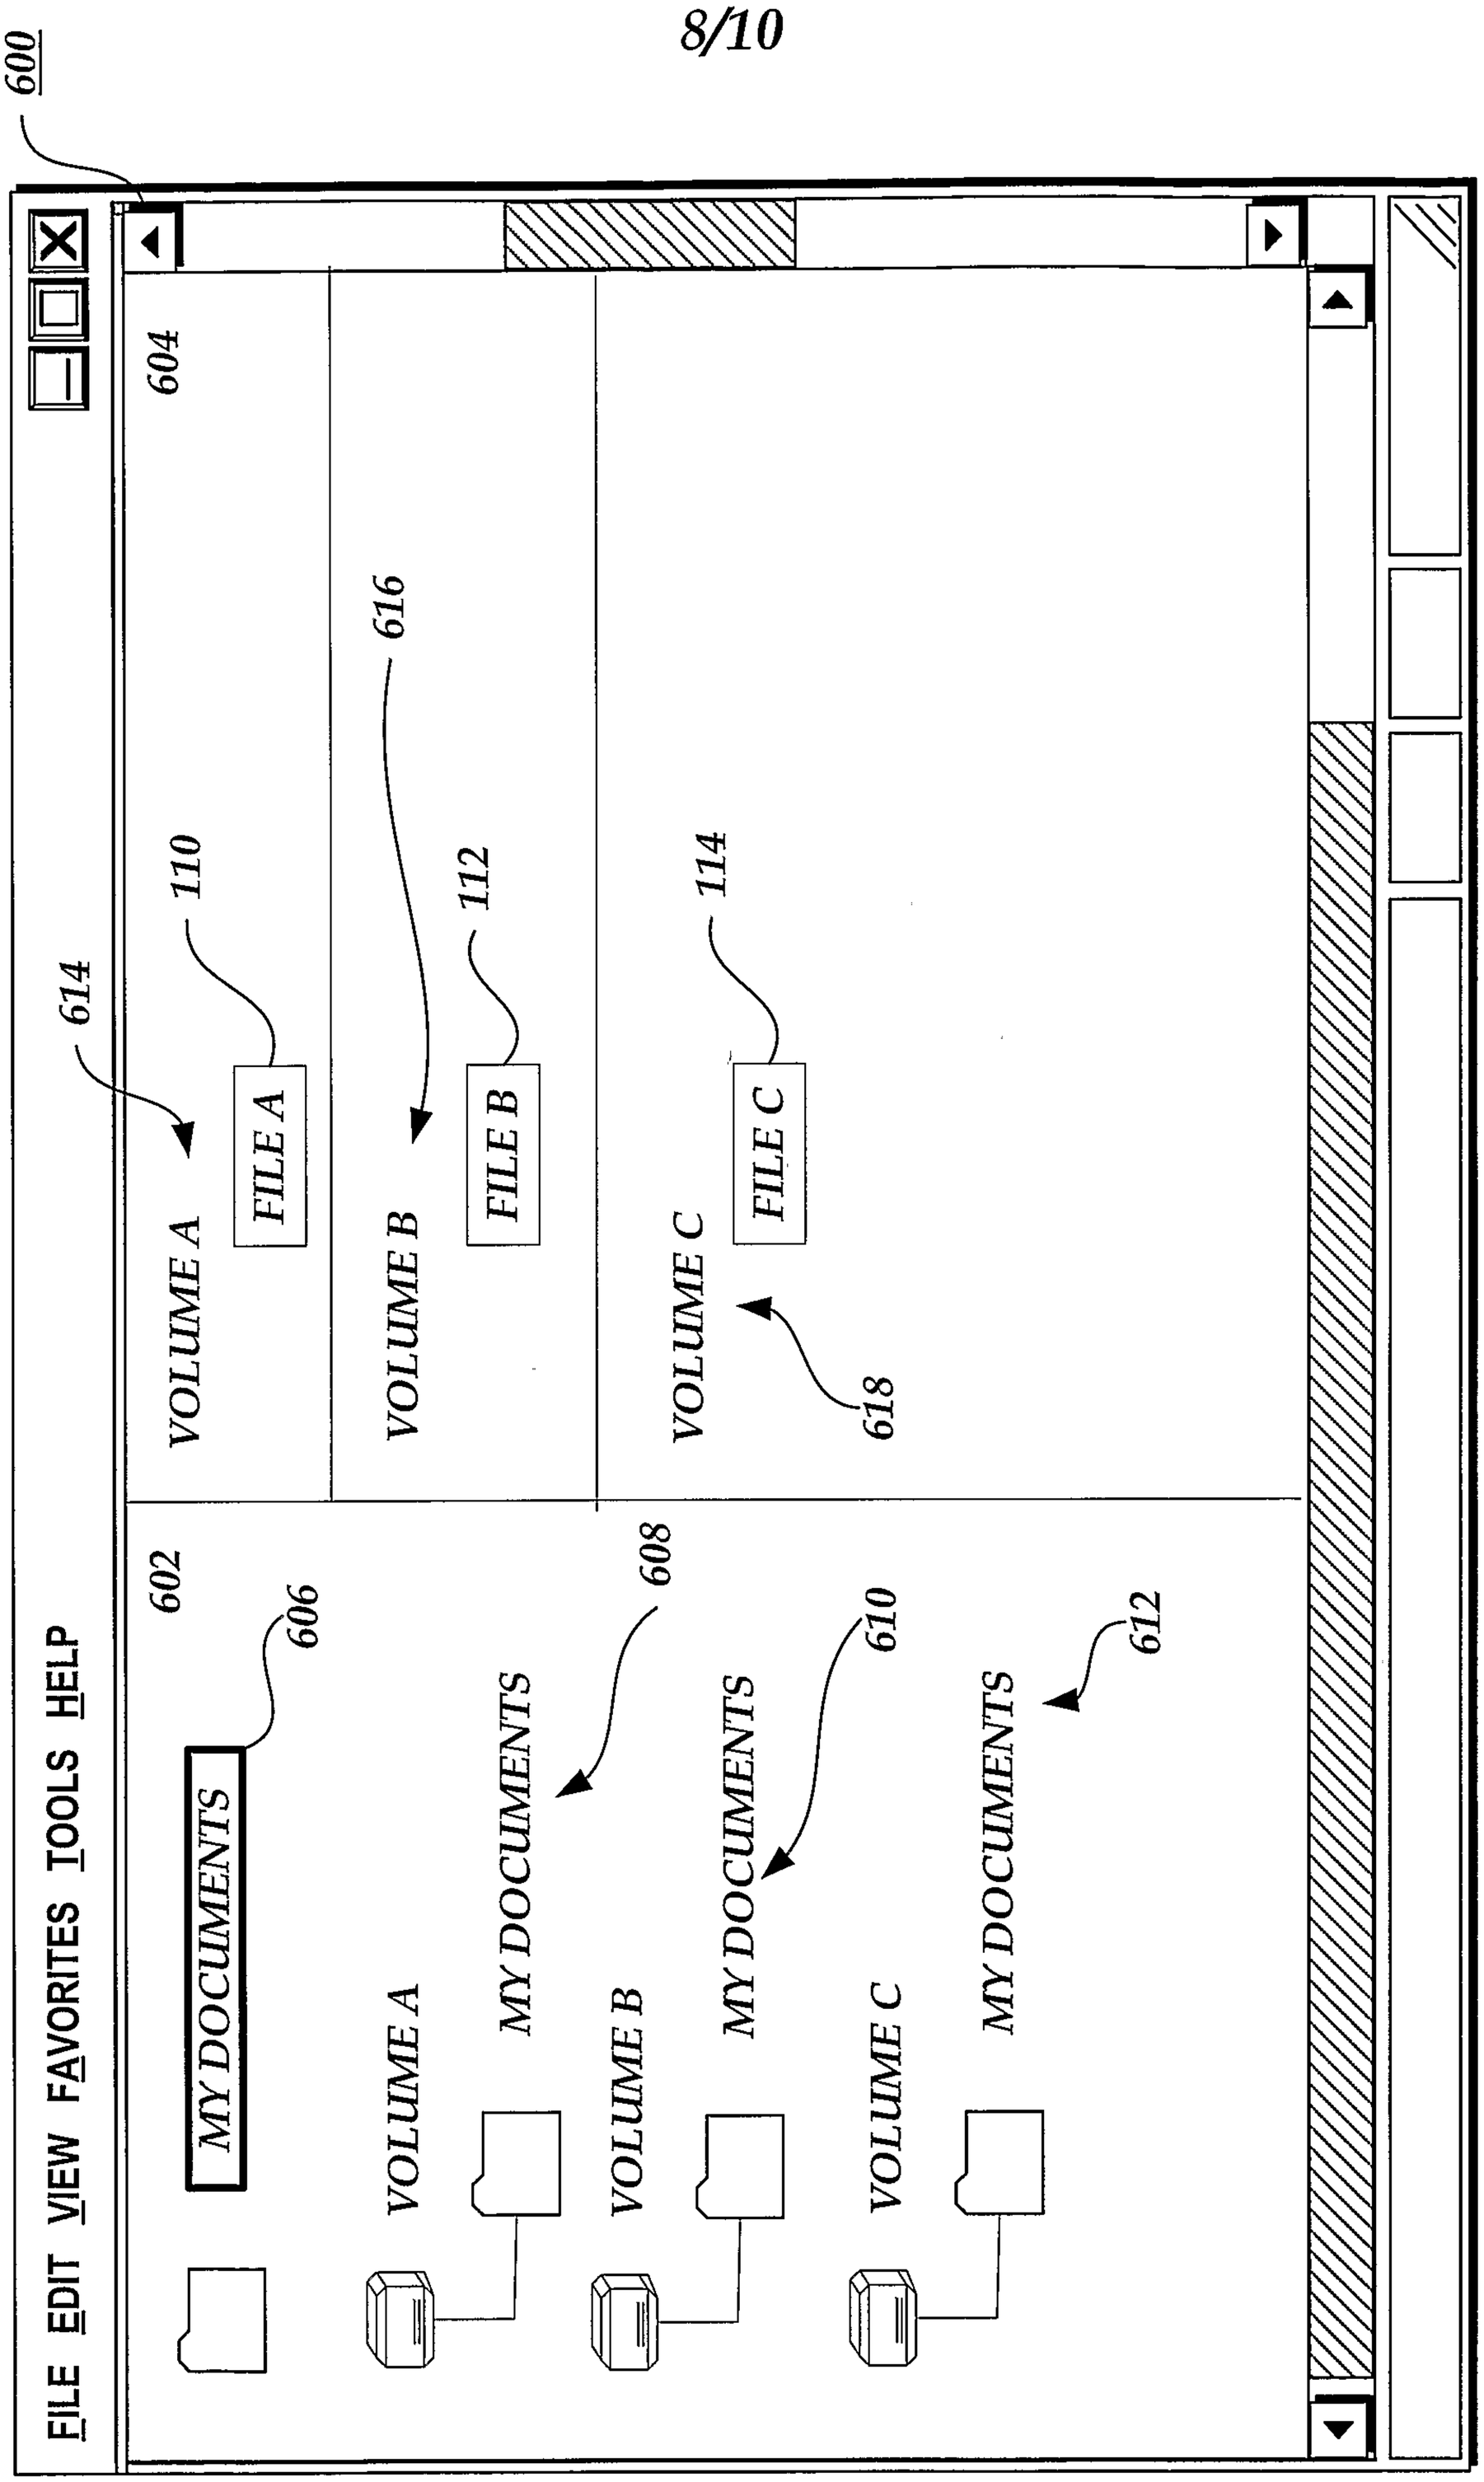
600
604
614
110
FILE A
616
112
FILE B
114
FILE C
FILE EDIT VIEW FAVORITES TOOLS HELP
602
606
MY DOCUMENTS
VOLUME A
608
MY DOCUMENTS
VOLUME B
610
MY DOCUMENTS
VOLUME C
612
MY DOCUMENTS
VOLUME A
VOLUME B
VOLUME C
618

Fig.6C.

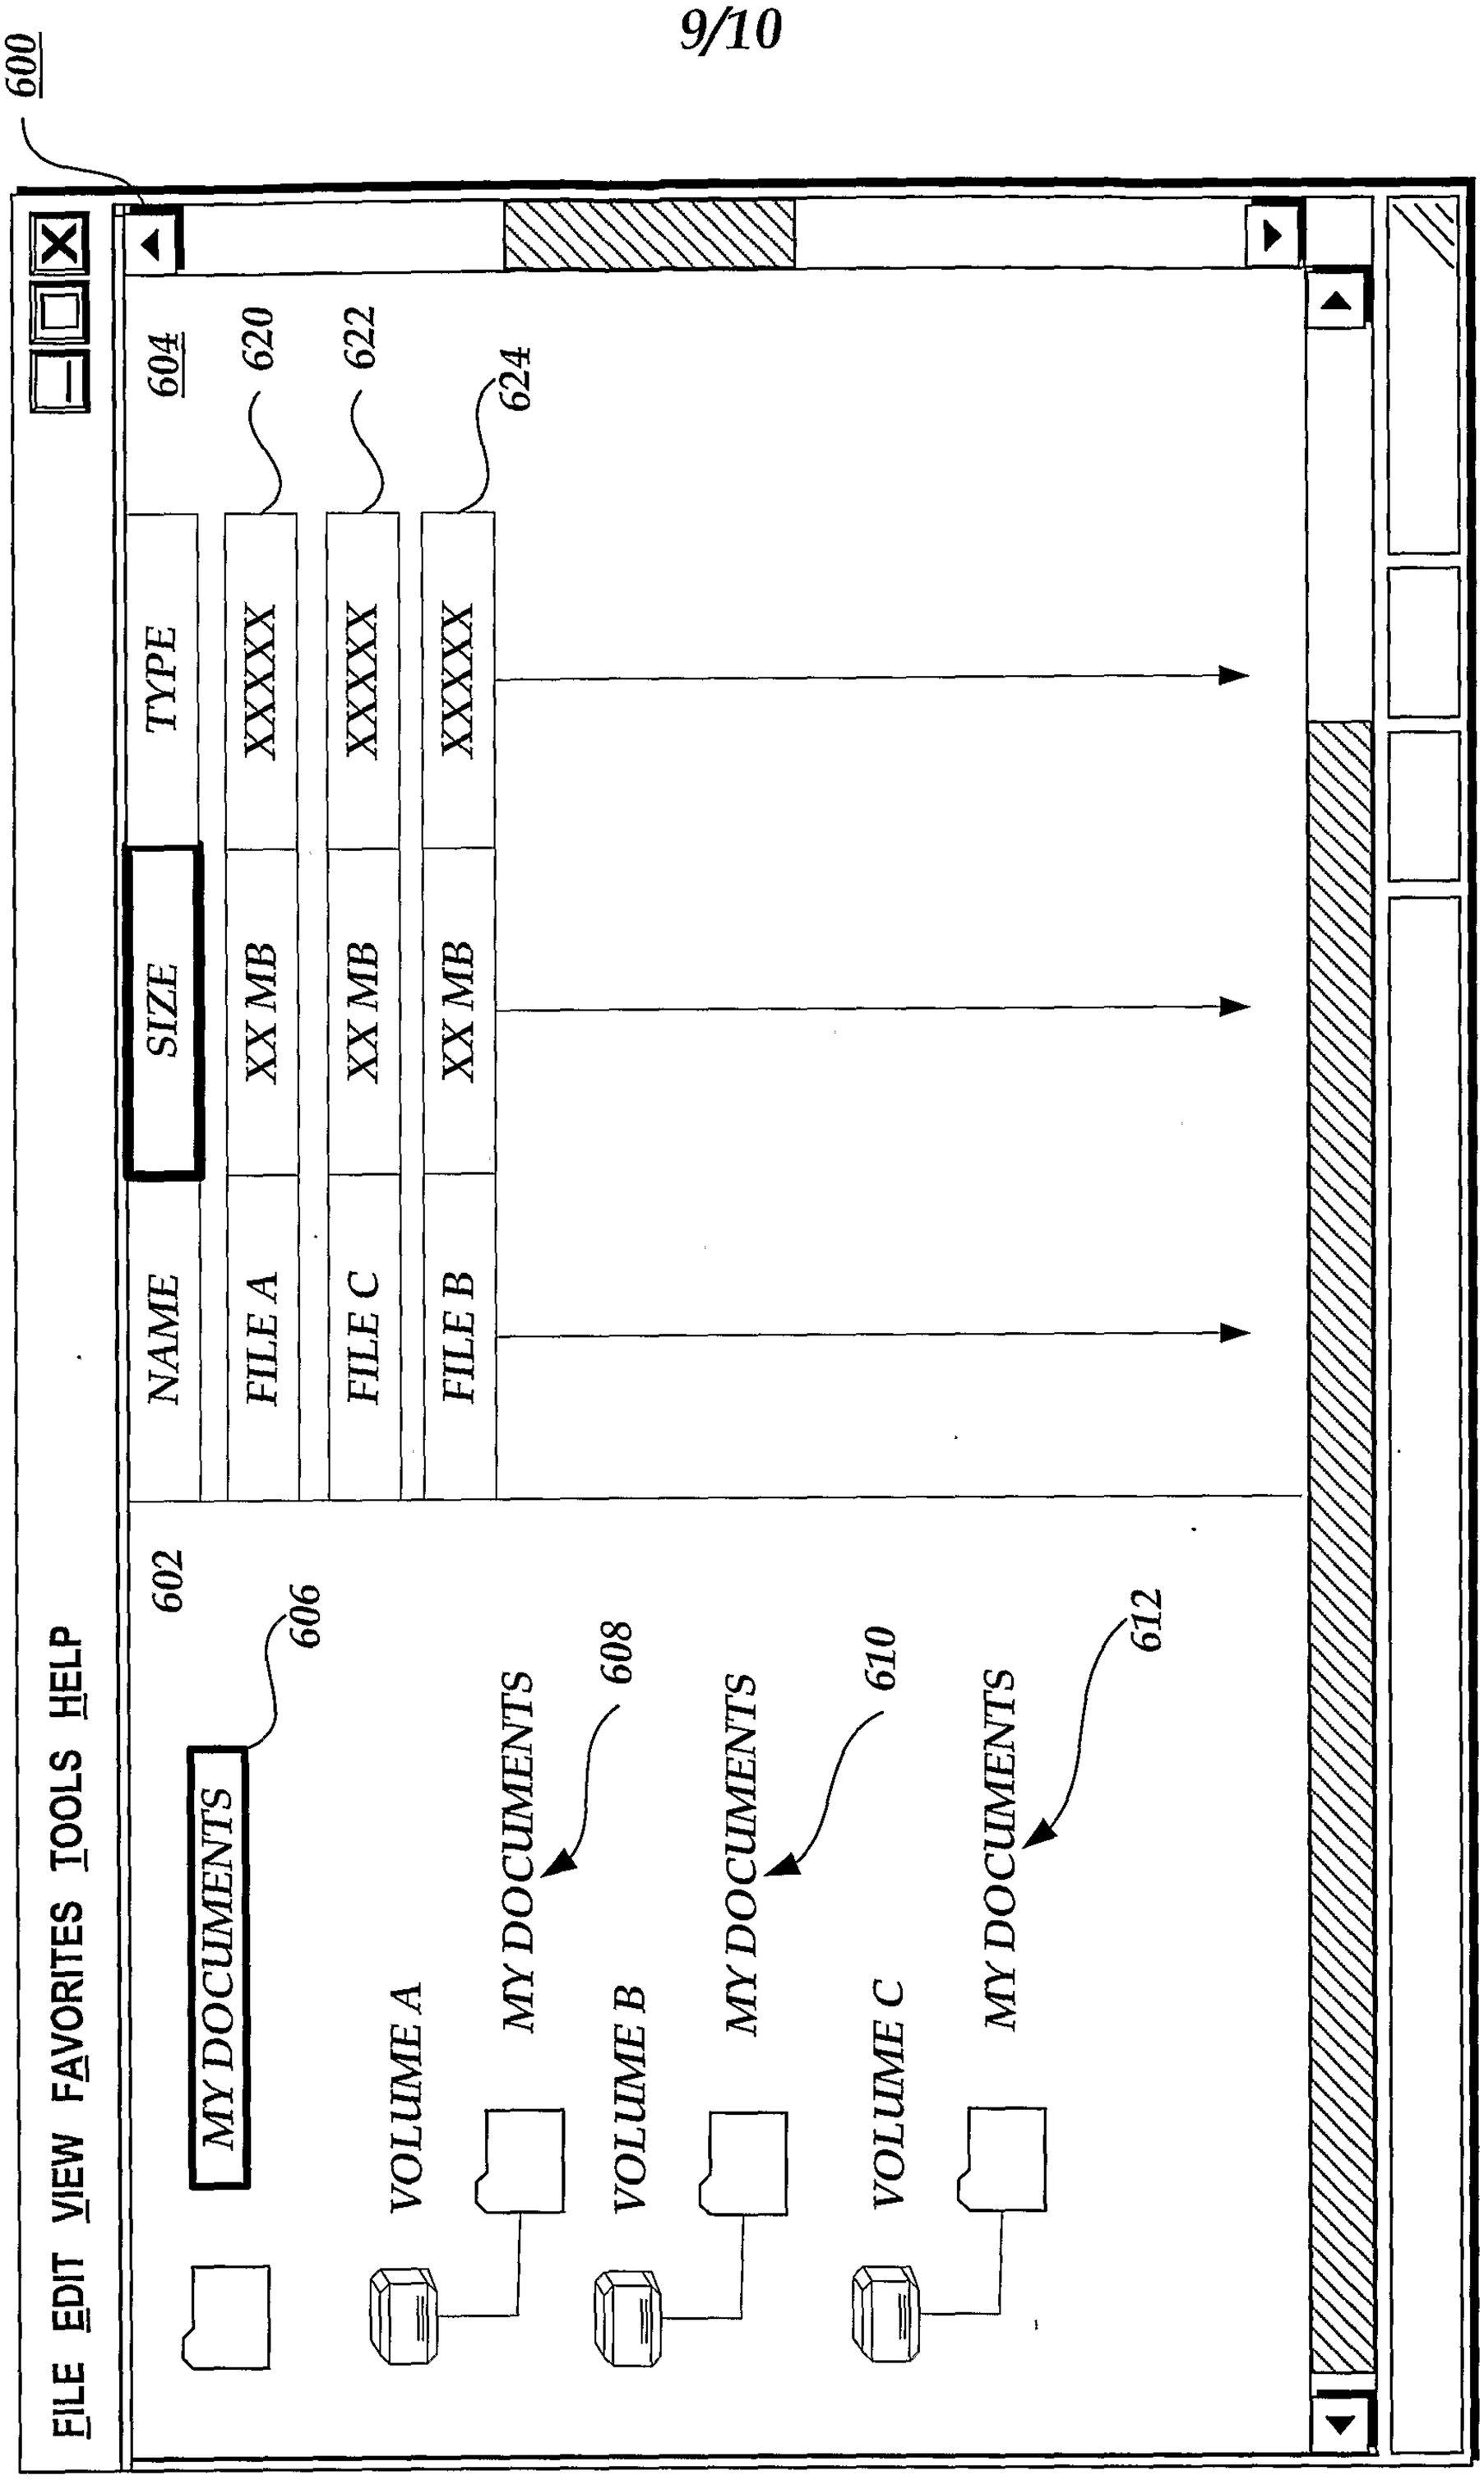
600
604
620
622
624
TYPE
XXXXX
XXXXX
XXXXX
SIZE
XX MB
XX MB
XX MB
FILE A
FILE C
FILE B
NAME
602
606
608
610
612
FILE EDIT VIEW FAVORITES TOOLS HELP
MY DOCUMENTS
MY DOCUMENTS
MY DOCUMENTS
MY DOCUMENTS
VOLUME A
VOLUME B
VOLUME C

Fig.6D.

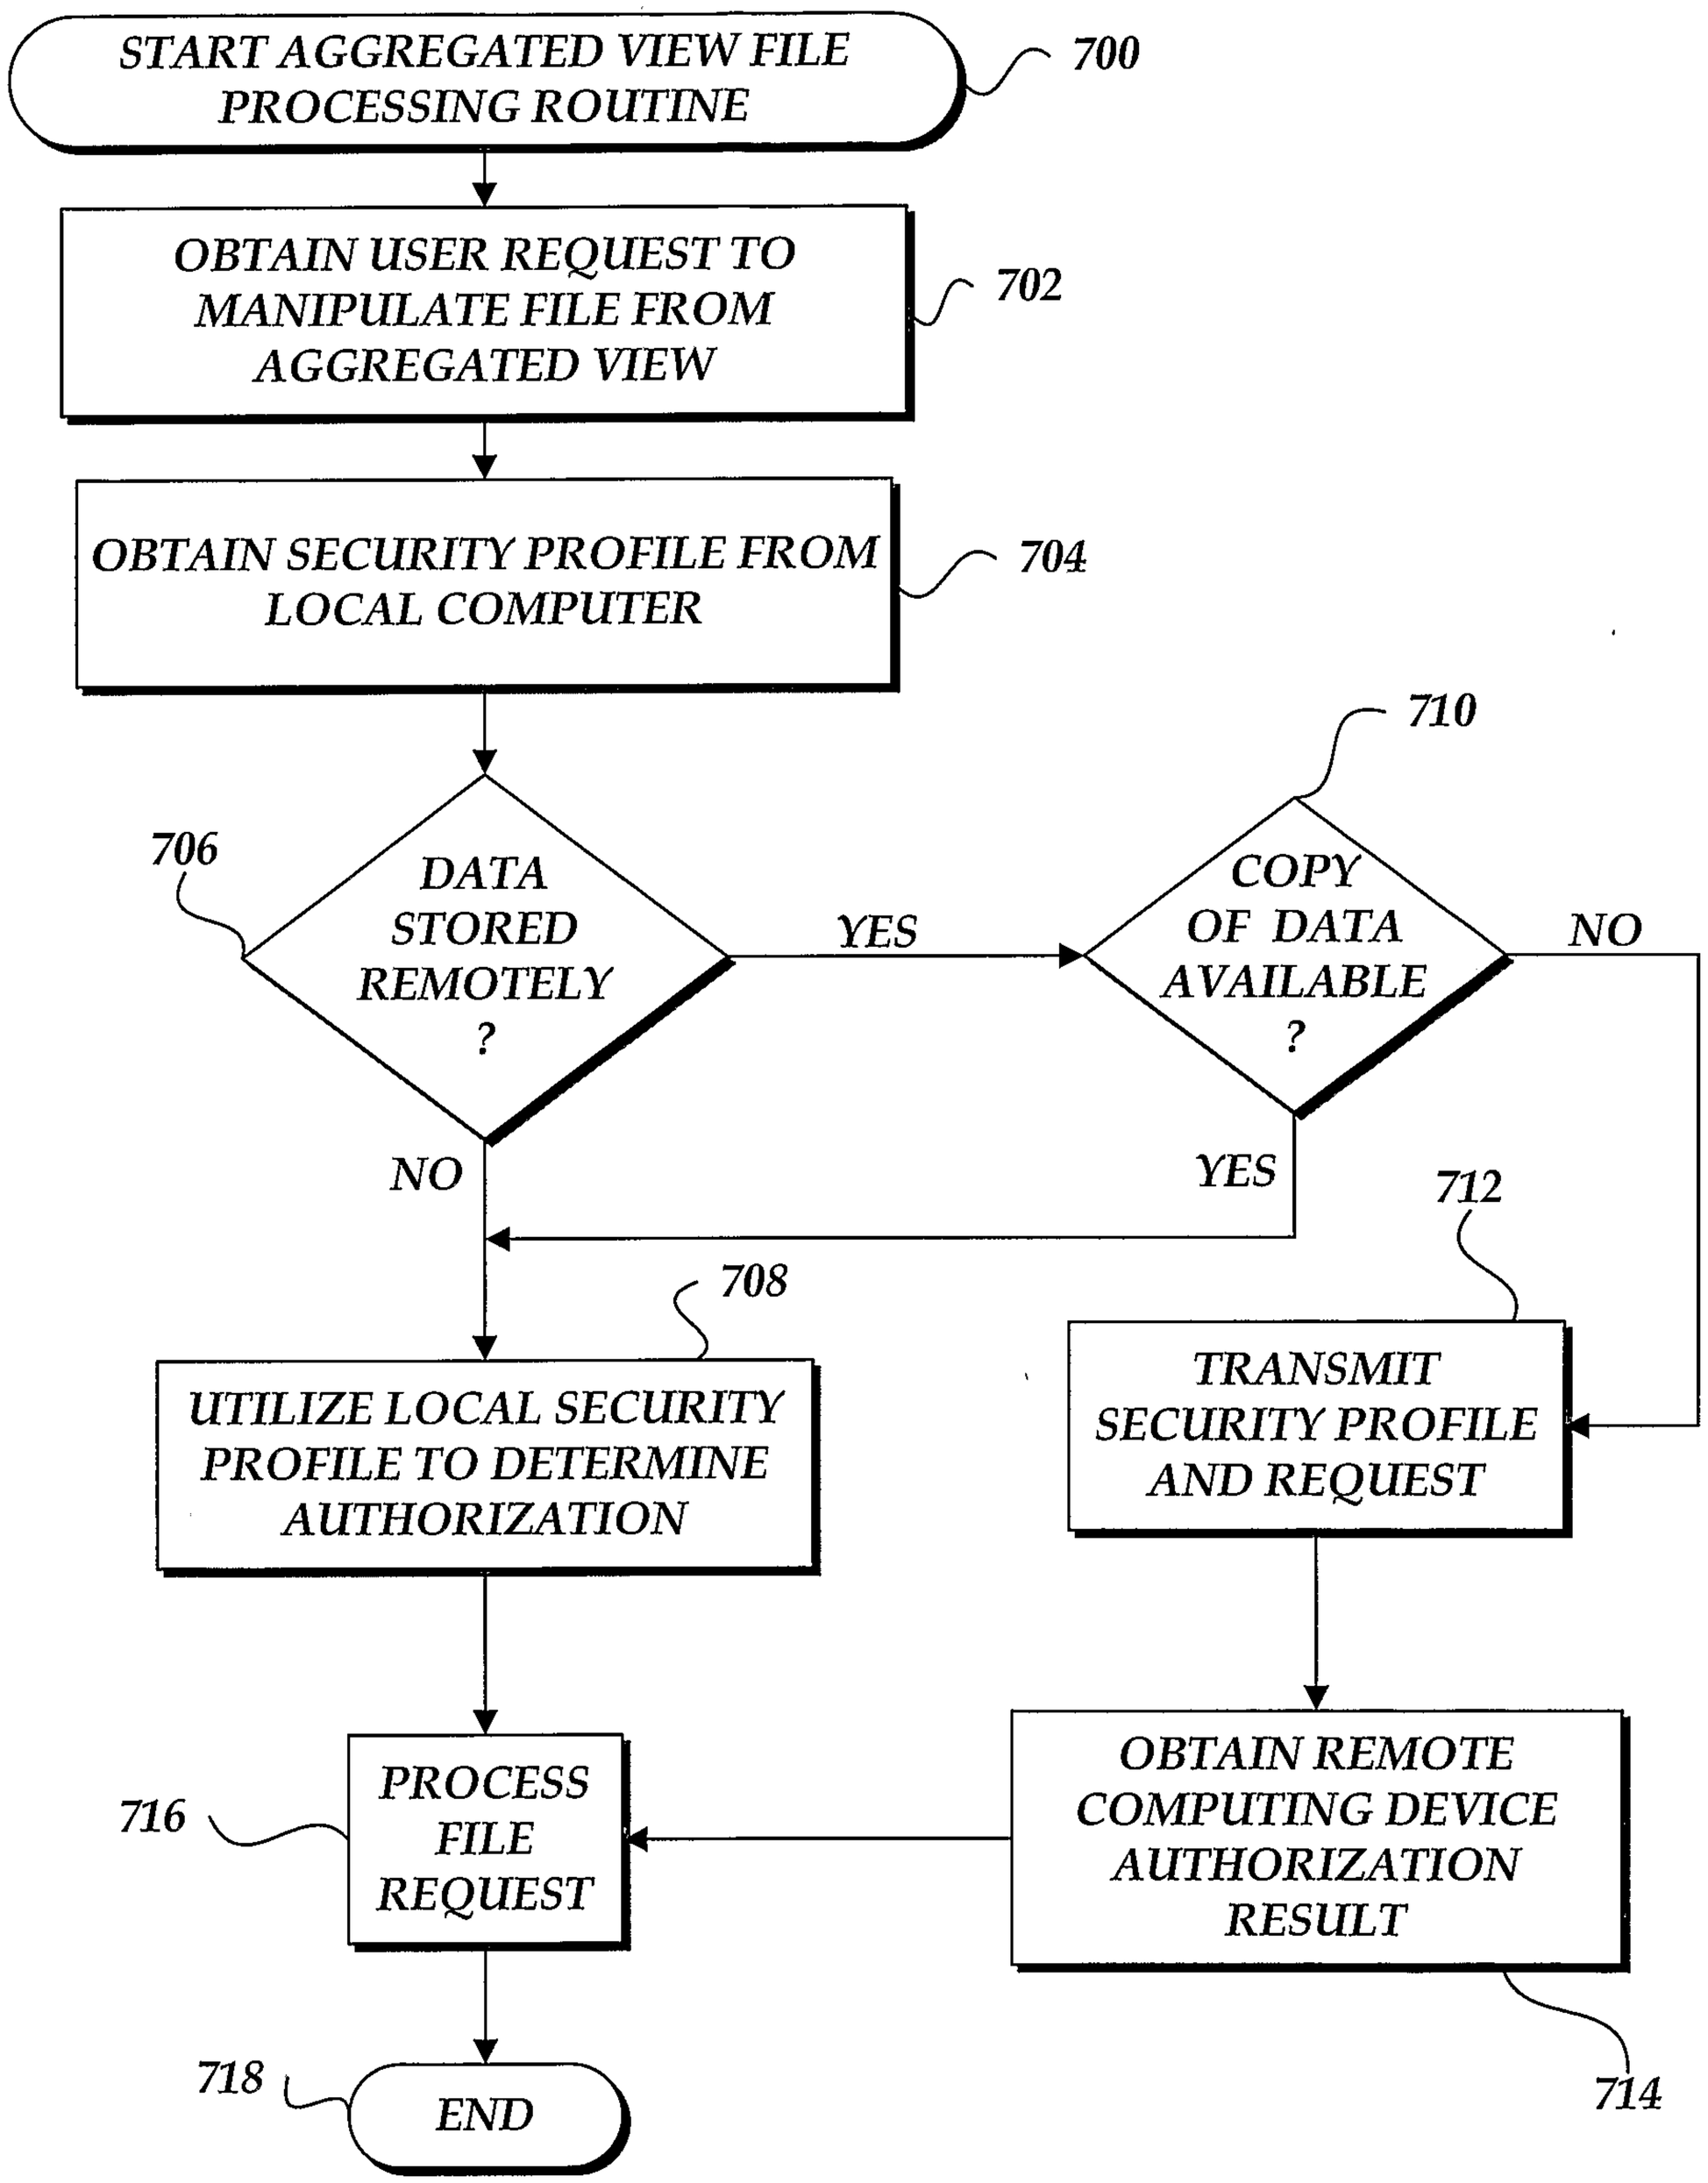
START AGGREGATED VIEW FILE PROCESSING ROUTINE
700
OBTAIN USER REQUEST TO MANIPULATE FILE FROM AGGREGATED VIEW
702
OBTAIN SECURITY PROFILE FROM LOCAL COMPUTER
704
706
DATA STORED REMOTELY ?
YES
NO
710
COPY OF DATA AVAILABLE ?
NO
YES
712
708
UTILIZE LOCAL SECURITY PROFILE TO DETERMINE AUTHORIZATION
TRANSMIT SECURITY PROFILE AND REQUEST
716
PROCESS FILE REQUEST
OBTAIN REMOTE COMPUTING DEVICE AUTHORIZATION RESULT
714
718
END

Fig.7.

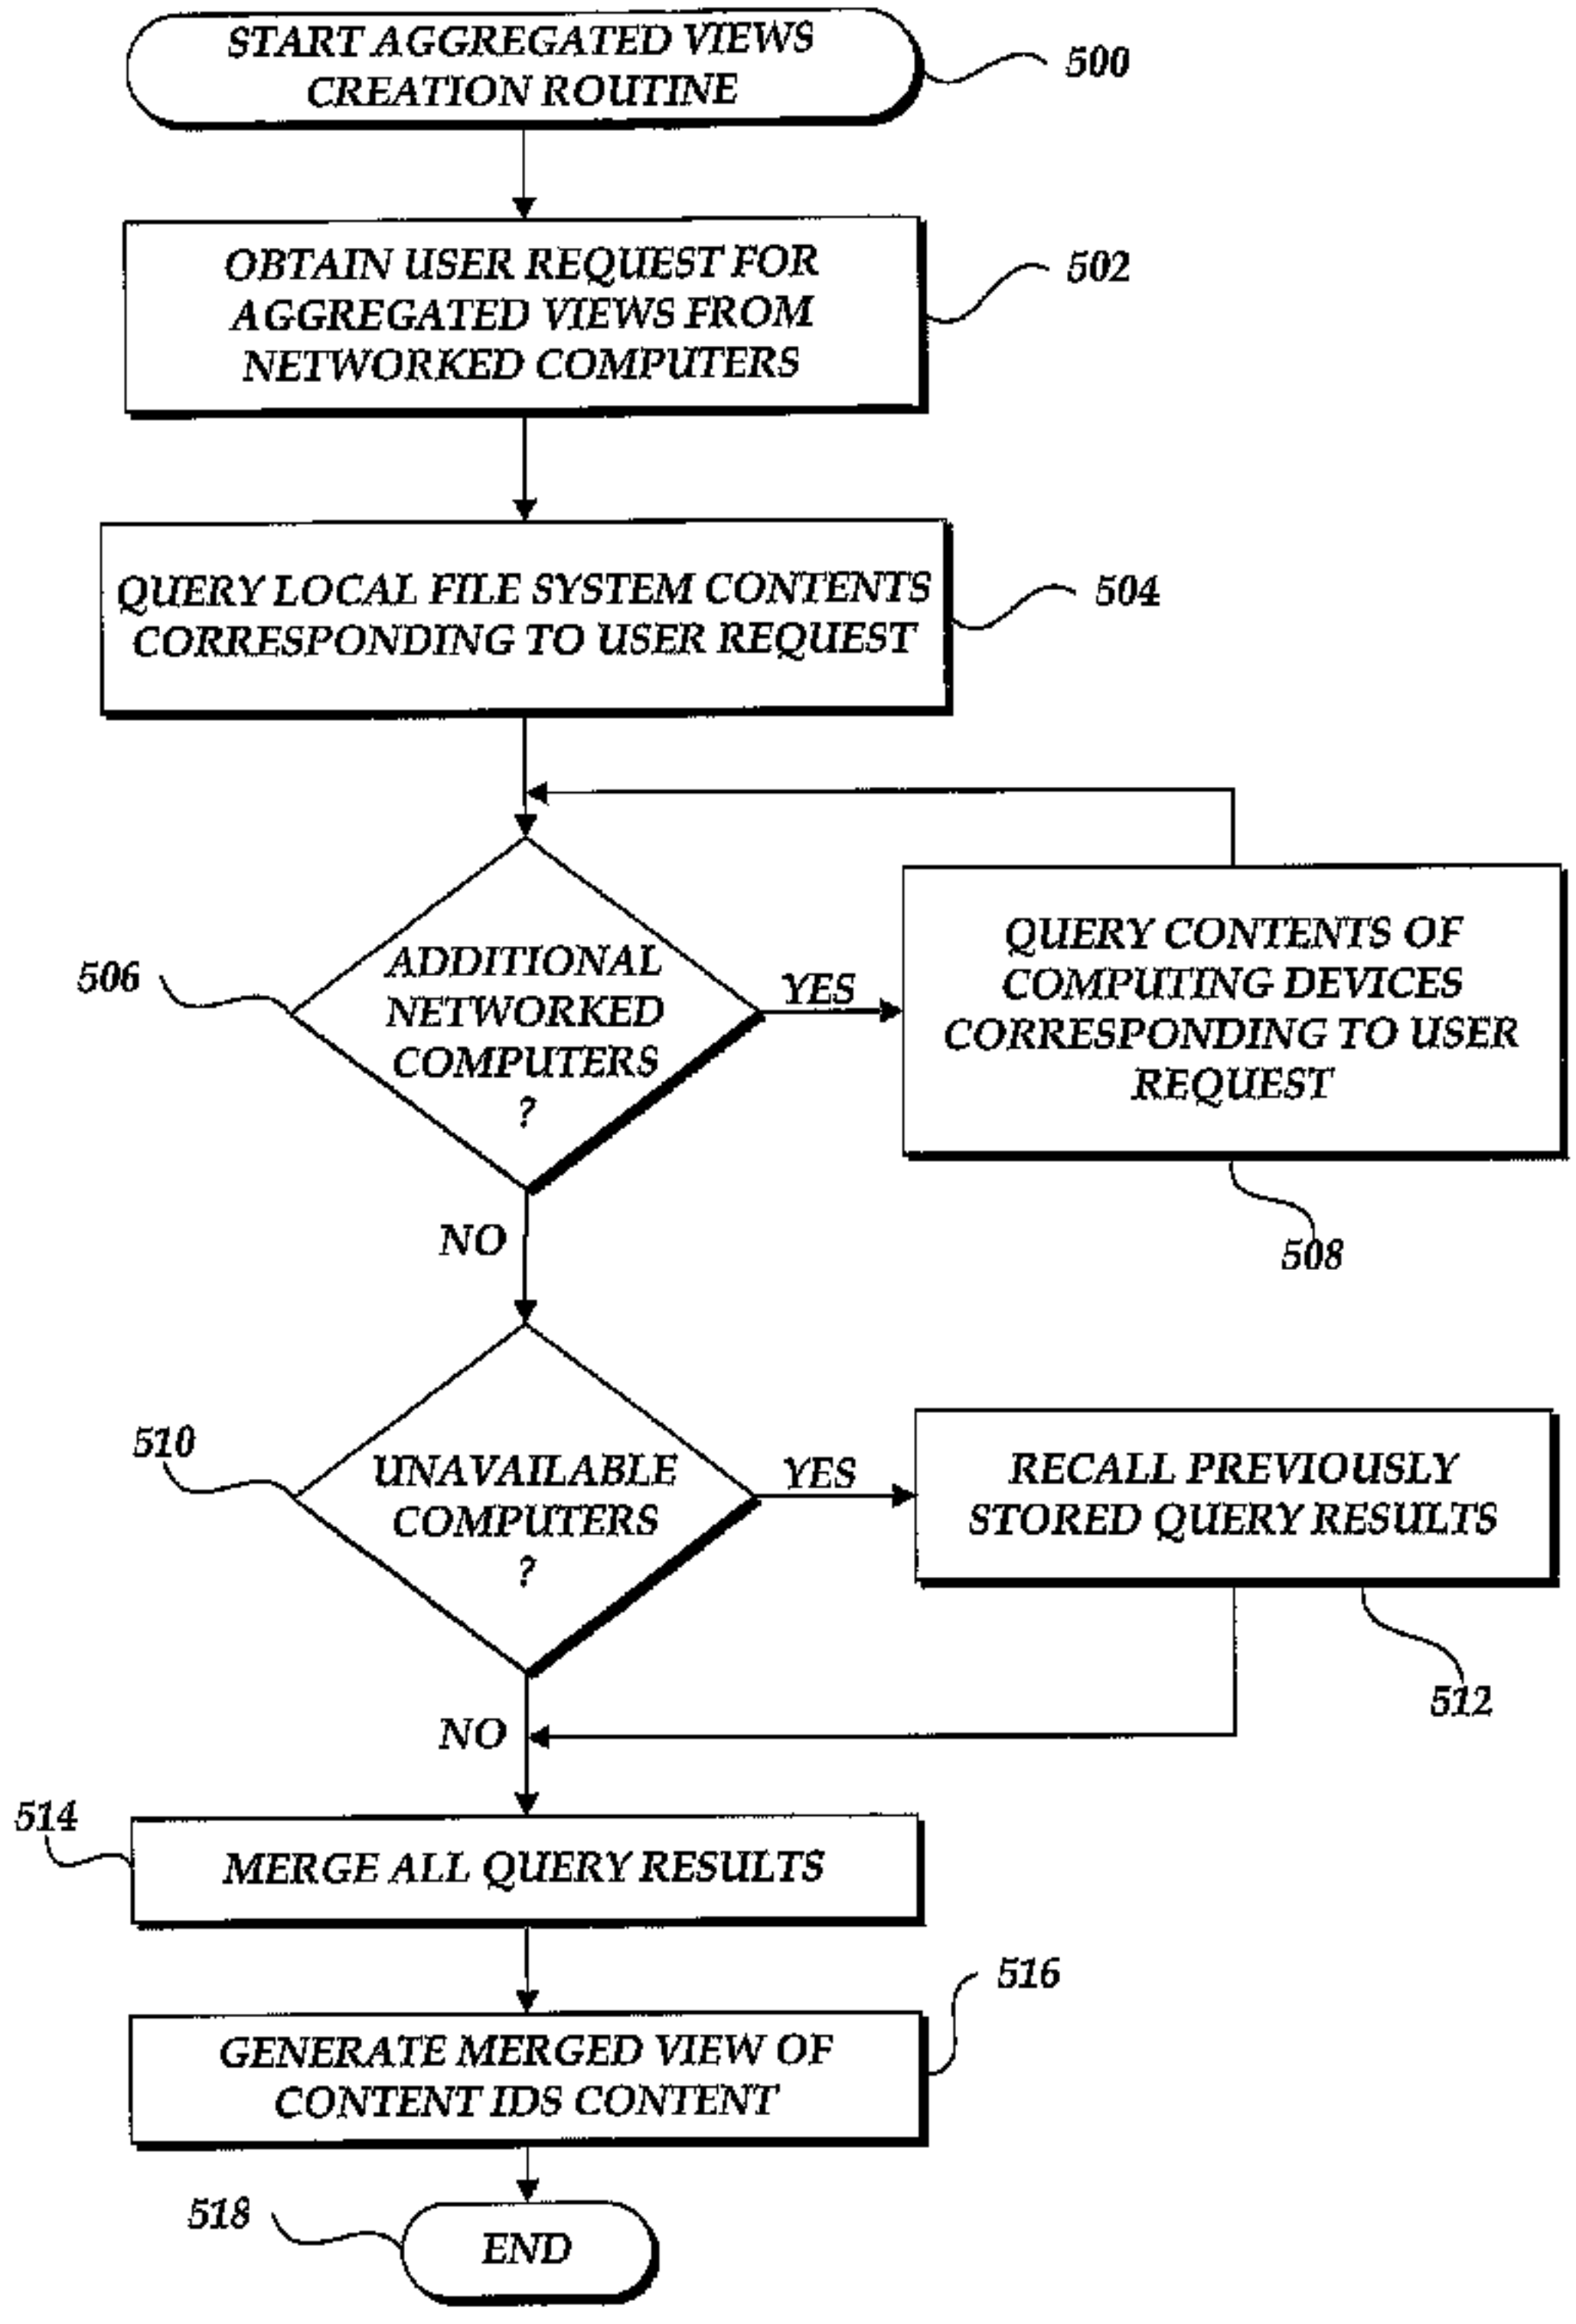
START AGGREGATED VIEWS CREATION ROUTINE
500
OBTAIN USER REQUEST FOR AGGREGATED VIEWS FROM NETWORKED COMPUTERS
502
QUERY LOCAL FILE SYSTEM CONTENTS CORRESPONDING TO USER REQUEST
504
506
ADDITIONAL NETWORKED COMPUTERS ?
YES
QUERY CONTENTS OF COMPUTING DEVICES CORRESPONDING TO USER REQUEST
508
NO
510
UNAVAILABLE COMPUTERS ?
YES
RECALL PREVIOUSLY STORED QUERY RESULTS
512
NO
514
MERGE ALL QUERY RESULTS
516
GENERATE MERGED VIEW OF CONTENT IDS CONTENT
518
END