

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
11 October 2001 (11.10.2001)

PCT

(10) International Publication Number
WO 01/75652 A2

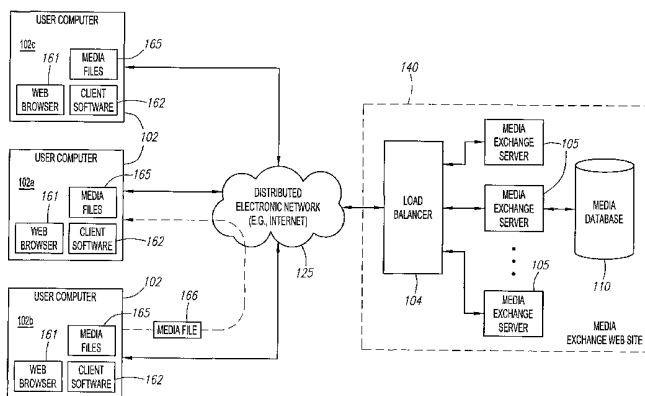
- (51) International Patent Classification⁷: G06F 17/00
- (21) International Application Number: PCT/US01/10546
- (22) International Filing Date: 30 March 2001 (30.03.2001)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
09/539,548 31 March 2000 (31.03.2000) US
- (71) Applicant: CENTERSPAN COMMUNICATIONS CORP. [US/US]; 7175 NW Evergreen Parkway, #400, Hillsboro, OR 97124 (US).
- (72) Inventors: SMILAK, Kevin; 10982 Roebling Avenue, #433, Los Angeles, CA 90024 (US). HAYKINSON, Ilya; 1127 15th Street, #D, Santa Monica, CA 90403 (US). TODD, Michael; 1245 McClellan Drive, #214, Los Angeles, CA 90025 (US).
- (74) Agents: VANDERLAAN, Christopher, A. et al.; Lyon & Lyon LLP, Suite 4700, 633 West Fifth Street, Los Angeles, CA 90071 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EE, 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, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, 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.

(54) Title: MEDIA EXCHANGE SYSTEM AND PROCESS



(57) Abstract: A media exchange system for facilitating the exchange of media files over a distributed electronic network, such as the Internet, comprises a load balancer for directing incoming requests from user computers to any of multiple media exchange servers. The media exchange servers connect to a concurrent access media database storing names and keyword information for media files of visiting users. The visiting users can simultaneously browse or conduct keyword searches of the media database. The search results are returned to the requesting user in the form of a list of media files and associated user addresses where such media files are located. Client-side software at the user computer then oversees a file transfer in a peer-to-peer transaction over the distributed electronic network, bypassing the media exchange web site. Each media exchange server may include a new record buffer for temporarily storing new records to be added to the media database when a new user logs on, and a record inserter which periodically adds the new records to the media database in a single database transaction. An exchange router may be used to facilitate media file transfers between members connected to different media exchange servers. Alternatively, a hierarchical tree architecture may include leaf node servers each having a local media database of a subset of all the logged on members for speeding up response time, and upper tier media exchange servers having more complete information that may be searched if necessary.

WO 01/75652 A2

SPECIFICATIONTITLE OF THE INVENTION5 **Media Exchange System and Process**BACKGROUND OF THE INVENTION1) Field of the Invention

10 The field of the present invention relates to systems and methods for media exchange over a distributed electronic network such as the Internet.

2) Background

15 In recent years, the exchange of media files, particularly sound files, over the Internet has become quite popular. One of the more common types of media file formats for sound files is known as the MP3 format, which is the common name for the MPEG-1 (Motion Pictures Expert Group) Audio Layer 3. The MP3 format is used to compress sound files into a much smaller group of data than, for example, a compact disc (CD), making MP3 files ideal for sending and downloading via the Internet. MP3 files are
20 commonly used to store music, which is presently by far the most popular type of content for exchange over the Internet.

A number of web sites have been developed to facilitate the exchange of music (i.e., sound) files among visitors to the site. In a typical application, a visitor to a music exchange web site is requested to register with the site so as to become a member. Client-
25 side software for the music exchange system is then typically downloaded to the member's computer, and the member is enabled to log on to the web site when convenient by invoking the client-side music exchange software application. The member designates certain directories on the member's computer for sharing with other members of the music exchange web site, and, upon log-on, the names of the music files contained in the
30 designated directories are transmitted to the music exchange web site. Those directories designated by members visiting the music exchange web site, and the music files stored in those designated directories, are temporarily made available to other members visiting the

site. At any given time, a visiting member therefore generally has access to all of the collectively shared music files by other visiting members, within the capacity of the given web site.

5 Once the member logs on to the web site, the member may specify which artist the member is interested in. The member is generally provided with the ability to enter search criteria for locating music files by artists of interest. The member's search criteria are forwarded to the server at the music exchange web site, which processes the request and conducts a search of all of the collectively shared directories available to it, within the constraints of the system. The search results are downloaded to the member's computer.

10 Typically, the member is presented with a list of music files from which one or more can subsequently be selected. Each of the music files is tagged with information indicating the source at which the music file may be found. When the member selects a music file from the list, a peer-to-peer transfer of the music file takes place. In a typical scenario, a file transfer request is made from the requesting member's computer over the Internet to the

15 member's computer at which the music file is stored, thus bypassing the central hardware at the music exchange web site. The client-side music exchange software at the other member's computer interprets the file transfer request, locates the appropriate music file, and transmits it back to the requesting member's computer, where it may be stored in any convenient location.

20 At least one web site also has been developed which allows the exchange of media files beyond only music files, and provides an on-line chat service as well.

As the popularity of music exchange web sites grows, these sites continue to attract more and more visitors who are interested in sharing music files. However, as an increasingly large number of members attempt to use a particular music exchange web

25 site, the demands on the hardware and software servicing the web site can become inordinate. For example, it can become very burdensome and inefficient for the central computer at the music exchange web site to search all of the lists of music files temporarily stored there, when the visiting members are numerous. The burdens placed on the web site by too many users can cause unacceptably slow response times.

30 One approach which attempts to alleviate the foregoing problem is to use multiple servers at the music exchange web site to increase the number of visiting users that can be supported. In one particular system, for example, when a member logs on to the web site,

the member is randomly assigned to one of the multiple servers of the web site by a redirector or other similar traffic control system component. Once assigned to a given server at the music exchange web site, the visiting member has access to the names of all of the music files of the other visiting members assigned to the same server. However, a
5 drawback with this approach is that a visiting member is limited to accessing the names of the music files of only those other visiting members who happen to be assigned to the same server. Thus, a visiting member assigned to a given server does not have access to the names of music files of other visiting members who happen to be assigned to other servers at the music exchange web site.

10 It would therefore be advantageous to provide a system and process for a music exchange which allows an increased number of visitors to share music files, each of the visitors able to access the music files of all other visitors, without causing inordinate burdens on the web site or making the web site response unacceptably slow.

15 SUMMARY OF THE INVENTION

The invention in one aspect provides a system and process for media exchange over a distributed electronic network such as the Internet.

In one embodiment, a media exchange system for facilitating the exchange of media files over a distributed electronic network, such as the Internet, comprises a load
20 balancer connected to the distributed electronic network for communicating with a plurality of user computers, and a plurality of media exchange servers connected to said load balancer. The load balancer directs incoming requests from the user computers to one of the media exchange servers. The media exchange servers are connected to a media database, which stores names and keyword information for media files of visiting users.
25 The media files are preferably stored in specially designated directories on the visiting user's computer. Visiting users can browse or conduct keyword searches of the contents of the media database. The media database is a concurrent access variety, allowing multiple media exchange servers to access it simultaneously in response to user search requests. The search results are returned to the requesting user in the form of a list of
30 media files and associated user addresses where such media files are located. Client-side software at the user computer then initiates a link to the media file owner's computer, requests a file transfer, and receives the file over the distributed electronic network in a

peer-to-peer transaction that bypasses the central hardware at the media exchange web site.

5 A preferred database schema comprises a user (or member) table, a media file table, and a keyword table. When a user logs on, the user's media file names are parsed into keywords, which are stored in the keyword table and indexed. The keyword index may later be searched by visiting users to find media files associated with certain user-specified keywords.

10 Each of the media exchange servers may include a new record buffer for temporarily storing new records to be added to the media database when a new user logs on. A record inserter (either global or local to each media exchange server) reviews the new record buffers of each of the media exchange servers periodically, and adds the new records to the media database. The new records from all of the new record buffers are preferably added by the record inserter in a single database transaction. In a preferred embodiment, the new records are added only if the new user has been logged on for a certain minimum period of time.

15 In another embodiment, a central user location register and exchange router is connected to the media exchange servers. The media exchange servers can look up logged-on members and their media server connection quickly on the central user location register, and can route media file transfer requests through the exchange router to the media exchange server connected to the media file owner, where access to the media file owner is otherwise difficult to achieve.

20 In another embodiment, a media exchange web site is configured in a hierarchical tree architecture, with a plurality of leaf node media exchange servers, a root media exchange servers, and, possibly, one or more levels of mid-tier media exchange servers. Each media exchange server preferably has its own local database having complete information as to it and its children, if any, while the root media exchange server preferably has global information about all members. Member search requests are preferably conducted on the local database first, then propagated to the next tier if not enough search results are returned. A media exchange server assignment table may assist in load balancing and keeping members together on the same media exchange server that are also topologically close, thus saving network resources.

Further embodiments as well as modifications, variations and enhancements of the invention are also described herein.

BRIEF DESCRIPTION OF THE DRAWINGS

5 FIG. 1 is a block diagram of a preferred media exchange system in accordance with one embodiment as disclosed herein.

FIG. 2 is a diagram of a partial database schema as may be used in connection with the media exchange system of FIG. 1.

10 FIG. 3 is a diagram of a process flow for sharing media files as may be accomplished, for example, utilizing the media exchange system of FIG. 1.

FIG. 4 is a diagram of a process flow for transferring a list of media file names from a member computer to a media database as may be accomplished, for example, utilizing the media exchange system of FIG. 1.

15 FIG. 5 is a diagram of an example of a process flow for searching for specified media files, in connection with a system having a database schema such as depicted in FIG. 2.

FIG. 6 is a block diagram of another embodiment of a media exchange web site, including a record inserter, as may be employed, for example, within a media exchange system such as illustrated in FIG. 1.

20 FIG. 7 is a block diagram is another embodiment of a media exchange web site, including local record inserters at some or all of the media exchange servers, as may be employed, for example, within a media exchange system such as illustrated in FIG. 1.

25 FIG. 8 is a block diagram of another embodiment of a media exchange web site, including a central user location register and exchange router, as may be employed, for example, within a media exchange system such as illustrated in FIG. 1.

FIG. 9 is a diagram illustrating an example of the operation of the media exchange web site depicted in FIG. 8 in the context of an overall media exchange system.

FIG. 10 is a process flow associated with the operation of the media exchange web site and system as illustrated in FIGs. 8 and 9.

30 FIG. 11 is a block diagram of another embodiment of a media exchange web site having a hierarchical architecture, as may be employed, for example, within a media exchange system such as illustrated in FIG. 1.

FIG. 12 is a diagram illustrating an example of the operation of the media exchange web site depicted in FIG. 11 in the context of an overall media exchange system.

FIG. 13 is an illustration of a routing table as may be used in conjunction with the media exchange web site and system as depicted in FIGs. 11 and 12.

5 FIG. 14 is a conceptual diagram of a generic node as may be employed in the media exchange web site of FIG. 11.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

FIG. 1 is a block diagram of a preferred media exchange system 100 in accordance
10 with one embodiment as disclosed herein. As illustrated in FIG. 1, the media exchange system 100 comprises a media exchange web site 140 including a load balancer 104 connected to a plurality of media exchange servers 105. The media exchange servers 105 are connected to a media database 110, which stores media file names from visiting members. The media exchange web site 140 connects to a distributed electronic network
15 125, such as the Internet, and thereby allows communication with a plurality of user and/or member computers 102. It should be understood that any of the user and/or member computers 102 may be part of a local area network (LAN), and may connect to the distributed electronic network 125 through a router and/or firewall or other shared local network utilities.

20 The media files involved in the media exchange system 100 may comprise, for example, digital sound or music files such as MP3 or wave (i.e., .wav) files, graphics files such as GIF, PNG, JPEG or MPEG files, three-dimensional graphics files such as meta-stream files, multi-media files such as AVI files, or streaming digital video or audio files, to name a few. The invention is not intended to be limited to these specific types of media
25 files; rather, any type of media file may be utilized in the media exchange system 100 so long as it is transportable over a distributed electronic network 125. Further, the concepts disclosed herein are also applicable to non-media files; however, the preferred embodiments as described herein concern the exchange of media files.

The media exchange servers 105 of the media exchange system 100 preferably
30 serve a variety of functions. The media exchange servers 105 maintain connections with member computers 102 of members currently visiting the media exchange web site 140. The media exchange servers 105 communicate over one interface with client-side software

162 over the distributed electronic network 125 using, for example, a standard Internet protocol (such as hyper-text transfer protocol (HTTP)), and over another interface with the media database 110 using an appropriate database access protocol (such as SQL or CGI). The media exchange servers 105 also preferably provide system security, manage overhead and administration, and conduct log-in and log-off procedures. The media exchange servers 105 each may be embodied as a dual-processor personal computer (PC) or functional equivalent running an operating system such as, for example, Linux or UNIX. The operating system of the media exchange servers 105 may be multi-tasking in nature.

10 Operation of the media exchange system 100 in a preferred fashion may be described in more detail with occasional reference to FIG. 3, which is a diagram of a process flow 300 for exchanging media files as may be accomplished, for example, utilizing the media exchange system of FIG. 1. In a preferred embodiment, as indicated by a first step 301 in the process flow 300 depicted in FIG. 3, a visitor utilizing a user computer 102 accesses the media exchange web site 140 using a standard web browser 161 or some other similar tool for navigating the distributed electronic network 125. To reach the media exchange web site 140, the user may enter the web site's Unified Resource Locator (URL) or other standard address information, allowing the web browser 161 to connect to the media exchange web site 140 using any of a variety of conventional techniques as are well known in the art.

20 Once the media exchange web site 140 is reached, as indicated by a next step 303 in the process 300, the visitor is requested to register with the media exchange web site 140 so as to become a member. While it is not essential to the functionality of the invention that users accessing the media exchange web site 140 become actual registered members, in a preferred embodiment, users interested in utilizing the services and capabilities of the media exchange web site 140 initially register as members. Registration may entail, for example, filling out information on an electronic form (e.g., an HTML page), such as information sufficient to identify the user and means of electronically contacting the user (by URL, e-mail address, or otherwise), information relating to membership payment (if payment is required to utilize the capabilities of the media exchange web site 140), and any other personal information as may be seen as useful to collect. The contents of the electronic form may be transmitted back to the media

exchange web site 140 as part of the registration process, as indicated by step 305 in the process 300 depicted in FIG. 3. As part of the registration process, the user may be provided with a username (which may be the same as the user's e-mail address) and password for accessing the media exchange web site 140.

5 When registration is complete, in a next step 306 as indicated in FIG. 3, client-side software 162 for the media exchange system 100 is preferably downloaded to the member computer 102 of the recently registered member. The client-side software 162 may be written, for example, in JAVA® or some other platform-independent language. Once downloaded, the client-side software 162 may appear as an icon on the member computer
10 102. By invoking the client-side software 162 at the member computer 102 (such as by clicking on the media exchange system icon with a mouse or by otherwise activating the software application), the registered member is able to access and log on to the media exchange web site 140, as indicated by step 310 in FIG. 3. The client-side software 162 may, in some embodiments, invoke the web browser 161 in order to establish an initial
15 connection to the media exchange web site 140. As part of the log on process, the member may be required to enter the member's username and password.

 Before or after logging on, the member preferably designates certain directories on the member computer 102 for sharing with other members at the media exchange web site 140. Upon logging on, or at any other convenient point thereafter, the names of the media
20 files contained in the designated directories of the member computer 102 are transmitted to the media exchange web site 140 and stored in the media database 110, as reflected by step 311, and as further described herein. These media files then become temporarily available to other members visiting the media exchange web site 140. A visiting member generally has access to all of the shared media files, collectively, stored on the designated
25 directories of the member computers 102 of all the visiting members. At any given time, it is expected that only some percentage of all registered members will be visiting the media exchange web site 140. Therefore, the media file names (and, hence, the media files) available to a visiting member are typically limited to those media files of the presently visiting members (at the time a search is conducted on the media exchange web
30 site 140, as further described herein).

 Once the member logs on to the media exchange web site 140, the member may obtain information concerning the identity of media files whose names (i.e., titles) are

temporarily stored by the media database 110, and, specifically, may conduct a search of media file names or else browse through media file titles. If the member is interested in a search, then the member may specify one or more criteria for locating media files which the member is interested in, as indicated by step 312. For example, the member may be interested in a particular musical artist or a particular song title. The client-side software 162 allows the member to enter search criteria for locating media files according to whatever search criteria the member specifies. As indicated in step 313 of the process 300 in FIG. 3, the member's search criteria are forwarded from the member computer 102 to the media exchange web site 140. In a next series of steps, the media exchange web site 140 processes the search request and conducts a search of the contents of all the available media files stored in the designated directories of all of the visiting members. Accordingly, in a preferred embodiment, as indicated by step 314, the search request from the member computer 102 is directed by the load balancer 104 to one of the media exchange servers 105 at the media exchange web site 140. In step 316, the media exchange server 105 interprets the search request and formulates an appropriate database query for the media database 110.

In a preferred embodiment, the media files of a particular member generally become available to the other members when the particular member logs on to the media exchange web site 140, and no longer become available, at least from a searching standpoint, when the particular member logs off the media exchange web site 140. Therefore, the media files available to a visiting member are typically limited to those media files of the presently visiting members at the time a search is conducted on the media exchange web site 140. As described later herein, exchanges of media files may occur after one or both exchanging members have logged off; however, when a member logs off, the list of that member's media files is preferably removed from the media database 110 so that no subsequent search will return the names of media files of the member who has logged off.

At a next step 320 of the process 300 in FIG. 3, the search results of the search conducted by the media exchange web site 140 are downloaded to the member computer 102. The search results may be presented in a list format. The list may contain the names of media files along with meta-information indicating the source (e.g., member computer

102, member's URL or e-mail address, or other linking information) at which the media file may be found.

If the member selects a browse option, then, as indicated by step 332, the member's browse request is transmitted from the member computer 102 to the media exchange web site 140. The member's browse request is directed by the load balancer 104 to one of the media exchange servers 105 at the media exchange web site 140. In step 336, the media exchange server 105 interprets and processes the browse request. The media database 110 may include an index file (in a binary tree or other suitable format) which comprises an alphabetically or alphanumerically sorted list of media file names stored in the media database 110. The media file name index is preferably updated with the addition or deletion of each media file name record in the media database 110. Alternatively, but less preferably, the media exchange server 105 may conduct a database query and/or a sort routine to arrive at an ordered list of filenames starting with the appropriate letter (or other alphanumeric character selected by the browsing member). In step 337, a list of media files is transmitted to the member computer 102. Only a fraction of the media file names listed in the media database 110 are transferred to the member computer 102. As the member reviews (e.g., scrolls through) the list of media files and reaches the end of the current listing, the member may request a subsequent set of media file names from the media exchange web site 140. To facilitate ongoing browsing capability for a given member, the media exchange server 105 serving the member may use an electronic bookmark or other mechanism as a placeholder in the media file name index to provide an indication of the last place in the media database 110 from which media file name information was obtained.

Whether the member is browsing the contents of the media database 110 or else has conducted a search, the member is preferably presented with an on-screen list of media files. In step 342, the member selects, if desired, one or more media files from the list of media files that has been provided to the member computer 102, by, for example, clicking on the appropriate media file(s), or by highlighting the appropriate media file(s) and hitting an assigned key on the computer keyboard, or by any other selection means, the details of which are not critical to the functioning of the invention. When the member selects a media file, a peer-to-peer transfer of the media file takes place between the member computer at which the media file is located (for example, member computer 102b

in FIG. 1) and the requesting member computer (for example, member computer 102a in FIG. 1). Thus, in step 344, the client-side software 162 at the requesting member computer 102a preferably formats one or more media file transfer requests. In step 345, each media file transfer request is transmitted from the requesting member computer 102a over the distributed electronic network 125 to the appropriate member computer (in this example, member computer 102b) where the media file is located, thus bypassing the central hardware at the media exchange web site 140. The meta-information associated with the media file name is used by the client-side software to obtain address information for the member computer 102b at which the media file is located.

In a next step 351, the client-side software 162 at the target member computer 102b interprets the file transfer request and locates the appropriate media file. In a following step 352, the client-side software 162 downloads the media file 165 back to the requesting member computer 102a. The media file 165 may be directly downloaded to the requesting member computer 102a, or else may be sent as an attachment to an electronic message (such as an e-mail message), or may be transmitted in any other suitable way. In a next step 353, the transferred media file 165 is stored in any desired location on the requesting member computer 102a. The requesting member may, if desired, specify a default directory for the storage of transferred media files.

If the member desires to select additional media files from the last search or browsing page, then the process returns to step 342, wherein the member selects the additional media files. Otherwise, as indicated by step 361, the member may select new search criteria, if desired, by returning to step 312. Alternatively, the member may enter a browse mode by selecting the browse option, and thus proceed to step 332.

When the member is finished browsing, searching and/or selecting media files, the member preferably logs off the media exchange web site 140. As part of the log-off procedure, the media exchange web site 140 immediately deletes the file names that were added to the media database 110 as a result of the member having earlier logged on.

In a preferred embodiment, it is contemplated that a large number of members may simultaneously access the media exchange web site 140 for the purpose of searching for and exchanging media files. For example, it is possible that several tens or hundreds of thousands of members may simultaneously access the media exchange web site 140. Many of the visiting members may, for example, be conducting searches for media files at

roughly the same time. The volume of potential usage places special demands on the media exchange web site 140 to provide reasonably prompt response time for all of the requesting users. Various architectural features of the media exchange web site 140 are provided to facilitate system performance under potentially severe user demand situations and to increase system throughput.

In a preferred embodiment, when multiple users conduct searches, each search request is directed by the load balancer 104 to one of the plurality of media exchange servers 105. The load balancer 104 assesses the demands on the various media exchange servers 105 and attempts to direct each new search request to the media exchange server 105 having the least amount of work at a given time. Each media exchange server 105 is preferably capable of formatting a database query and simultaneously accessing the media database 110 for the purpose of conducting a database search according to the respective member's specified criteria. To enable simultaneous access by multiple media exchange servers 105, the media database 110 may comprise a concurrent-access database of the type commercially available from, for example, Oracle Corporation of Redwood Shores, California. The Oracle 8i database is an example of a suitable concurrent-access database. The media database 110 can be a multiple layer database, although it preferably appears as a single "virtual" database from the perspective of external applications. The media database 110 may comprise hardware components including one or more disks for storing information and one or more processors (e.g., Alpha-type processors) for executing the programming instructions for managing the media database 110.

To facilitate rapid, parallel searching of large numbers of records stored in the media database 110, a preferred database schema such as that depicted in FIG. 2 is preferably utilized. As shown in the example of FIG. 2, three primary tables are utilized in connection with a preferred database schema: a keyword table 201, a file table 221, and a user table 241. Records in the keyword table 201 each preferably comprise a word field 202 and an ID field 203. Records in the file table 221 each preferably comprise an ID field 223, a name field 224, a size field 225, a date field 226, a bit rate field 227, and a user ID field 228. Records in the user table 241 each preferably comprise an e-mail address field 242, a login information field 243, a state field 244, and a user ID field 245. A record exists in the user table 241 for each member of the media exchange system 100. A record exists in the file table 221 for each media file designated for sharing by a given

member. Each media file record in the file table 221 is associated to a given member record in the user table 241 through the user ID fields 228 and 245. A record exists in the keyword table 201 for each word (with the possible exception of certain articles such as “a” and “the” and/or certain other very common words providing little or no meaningful searchable content) in the name of each media file in the file table 221. Each keyword file in the keyword table 201 is associated to a given media file record in the file table 221 through the file ID fields 203 and 223.

With the database scheme illustrated in FIG. 2, for each member of the media exchange system 100, all of the media files currently available through that member are readily ascertainable by locating records in the media file table 221 having the user ID matching that of the member record in the user table 241. Conversely, the owner of each media file record in the file table 221 can be readily ascertained by locating the appropriate record in the user table 241. For each media file record in the file table 221, all of the constituent words in the name of the media file (with the possible exception of certain very common words) which are useful for searching are readily ascertainable by locating records in the keyword table 201. Conversely, the media file record in the file table 221 for any given keyword stored in the keyword table 201 can be readily ascertained by locating the appropriate record in the file table 221.

When a member is logged off the media exchange web site 140, the state field 244 is set to “off-line”. When a member initially accesses the media exchange web site 140, the member enters log-in information which is cross-checked against the contents of the login information field 243 for the particular member. Preferably, the load balancer 204 directs the login request to one of the media exchange servers 105 for handling the login procedure. If login is successfully completed, the state field 244 for the appropriate member record in the user table 241 is changed from “off-line” to “on-line”.

Once login has been successfully completed, the names of the media files designated in the shared directories of the member computer 102 are transmitted to the media exchange web site 140. FIG. 4 is a diagram illustrating one example of a process flow for transferring media file names from a member computer 102 to the media database 110. As indicated in FIG. 4, in a first step 402, the client-side software 162 reads the names of the media files 165 stored on the member computer 102 in the designated directories. In a next step 404, the client-side software 162 transmits the list of media file

names to the media exchange web site 140. The media exchange web site 140 receives the list of file names in step 405. Then, in a next step 406, the load balancer 104 at the media exchange web site 140 directs the list of file names to the media exchange server 105 handling the login procedure. The media exchange server 105, in a next step 408, parses each of the media file names into constituent words. In a next step 409, the media exchange server 105 generates a file record compatible with file table 221 for each media file name received from the member computer 102. Then, in step 410, the media exchange server 105 generates a keyword record compatible with keyword table 201 for each constituent word in each media file name. The media exchange server 105 then delivers the file records and the keyword records to the media database 110 for storage. A database server of the media database 110 adds the file records and keyword records to the media database 110, and takes care of indexing as necessary.

Preferably, as previously indicated herein, the media database 110 includes a media file name index which comprises an alphabetical or alphanumerically sorted listing of media file names stored in the file table 221. If such is the case, then, in step 411, the media file name index is updated to include the new media file names of the records added to the file table 221 (specifically using the contents of the name field 224 of the new records added to the file table 221). In addition, the media database 110 also preferably includes a keyword index which comprises an alphabetical or alphanumerically sorted listing of keywords stored in the keyword table 201. Accordingly, in step 413, the keyword index is updated to include the new keyword records added to the keyword table 201 (specifically using the contents of the word field 202 in the new records added to the keyword table 201).

FIG. 5 is a diagram of an example of a process flow for searching for given media files and returning the search results, in connection with a system having a database schema such as depicted in FIG. 2. Thus, when a search for a particular keyword or set of keywords is conducted, as indicated by step 502 in the process 500, a query is preferably run against the keyword index stored and maintained in the media database 110. The keyword index facilitates rapid searching of the keyword table 201. In step 503, a determination is made as to whether a multiple keyword search is being conducted. If so, then in step 504, the results from each keyword search are merged according to the parameters specified in the search request (i.e., a union of the results is generated for

search terms connected by a boolean OR relationship, and an intersection of the results is generated for search terms connected by a boolean AND relationship). For example, if a member desires to search for the name of a media file "Star Spangled Banner" and enters the search terms "Star" and "Spangled" with an AND relationship, the responsible media server 105 at the media exchange web site 140 will conduct a database query under each of the two search terms, and the search results of "Star" will be crossed with the search results of "Spangled", yielding a final list of file names having both the words "Star" and "Spangled" in their titles. If the same search were conducted using an OR relationship, the final list of file names would include all names having either the word "Star" or the word "Spangled" in the title.

In a preferred embodiment, the search of the keyword records in the keyword table is conducted as an "identity" search, in the sense that an exact match of the search term and the keyword in the keyword table needs to occur to return a hit. An advantage of an identity search is that it returns results relatively quickly, as compared to, for example, a substring search, which would generally require a compare operation on each constituent substring in all of the keywords stored in the keyword table.

In a next step 505, for each keyword record located in the keyword table 201 matching the search criteria (i.e., the search results), the corresponding media file record in the file table 221 is identified, using the link between the file ID fields 203 and 223 in the keyword table 201 and file table 221, respectively. For each media file record in the file table 221 retrieved, the contents of the name field 224 are read out and collected by the media exchange server 105, as indicated by step 506. The media exchange server 105 then, in step 507, locates the corresponding owner member information (i.e., the member record stored in user table 241) for each media file record by using the linking user ID fields 228 and 245. In step 508, the media exchange server 105 reads out the contents of the e-mail address field 242 from the member record in the user table 241 corresponding to each media file record that was returned in the search. Then, in step 510, the media exchange server 105 generates a list of media file names and corresponding user e-mail addresses for the owner members of each media file. In a next step 515, the media exchange server transmits the list of media file names and corresponding user e-mail addresses to the requesting member computer 102. At the member computer 102, in step 517, the client-side software 162 displays the list of media files on-screen. The

corresponding e-mail addresses are preferably not displayed, but are securely and confidentially maintained by the client-side software 162.

If the requesting member desires to obtain a media file, the requesting member selects the media file from those listed on the screen. In response, the client-side software 162 generates a link to the member owner's computer using the member's e-mail address. To establish such a link, the client-side software 162 may utilize the web browser 161 resident at the member computer 102. When the link is established, the client-side software 162 at the requesting member computer 102 communicates the requested media file name to the client-side software 162 at the owner member computer 102. The client-side software 162 at the owner member computer 102 retrieves the media file from the shared directory at the owner member computer 102, and transmits back over the link to the requesting member computer 102. The transfer of the media file may be conducted according to a hyper-text transmission protocol (HTTP), or by a non-HTTP packet-switched or TCP/IP transmission protocol. In the case of streaming files, the transfer of the media file may also alternatively be conducted using a circuit-switched connection protocol. The client-side software 162 at the requesting member computer 102 then may store the obtained media file in a special designated directory, and may also display a message to the member that a media file transfer was successfully completed. If the file comprises streaming data, then, rather than, or in addition to, being stored at the requesting member computer 102, it may be displayed on-screen as the streaming data is received.

As an optimization for facilitating rapid searching of the media database 110, in one embodiment the media database 110 is physically partitioned into different areas (or different disks or spindles) with each area corresponding to a specific alphanumeric range. For example, the media database 110 may be segmented into twenty-seven different areas (or different disks or spindles), each of the first twenty-six areas corresponding to a different letter of the alphabet and the twenty-seventh area corresponding to numerals 0 through 9. Each media file is then stored in the area or on the disk or spindle associated with the first letter or number (i.e., the first alphanumeric character) in its title. To browse through the media files stored in the media database 110, the media exchange server 105 may read from the "top" the database segment associated with the letter 'a' (or other alphanumeric character selected by the browsing member).

When a member logs off the media exchange web site 140, the contents of the state field 244 in the member record stored in the user table 241 is changed from “on-line” back to “off-line”. As immediately as practicable, all of the media file name records and keyword records associated with the logged-off user are deleted from the media database 110. Accordingly, as part of the log-off procedure, the media exchange server 105 handling the visiting member searches the media database 110 and requests deletion of all media file records in the file table 221 corresponding to the member, and all keyword records in the keyword table 201 corresponding to the member’s media file records. The deletion of records pertaining to members logging off is preferably given very high priority, so that future searches by other members do not return the records of members who are no longer logged onto the media exchange web site 140. However, even if the file names of a logged off member are held in the media database 110 for some transient period of time after log-off, little harm is expected to arise, since it is likely that the shared media files will remain on the logged off member’s computer 102 for some time after log off. Thus, a file transfer request directed to a member that has logged off recently is still likely to result in an effective file transfer, so long as the member computer 102 of the media file owner is still connected to the distributed electronic network 125.

FIG. 6 is a diagram of an alternative embodiment of a media exchange system 600 having additional features for reducing overhead associated with the media database. In the media exchange system 600 illustrated in FIG. 6, a load balancer 604 serves a similar purpose as load balancer 104 shown in FIG. 1, and connects to a distributed electronic network and thereby to various user or member computers, not explicitly shown in FIG. 6, in a manner similar to that described with respect to FIG. 1. For purposes of clarity, only features of the media exchange web site 640 are illustrated in FIG. 6. As further shown in FIG. 6, the media exchange web site 640 comprises a plurality of media exchange servers 605, which again serve similar general functions as the media exchange servers 105 depicted in FIG. 1. However, each media exchange server 605 in the embodiment illustrated in FIG. 6 comprises a new record buffer 606 for temporarily storing new media file records to be added to the file table 221 and new keyword records to be added to the keyword table 201. At least some, but preferably all, of the media exchange servers 605 are connected to a record inserter 607 at the media exchange web site 640.

In operation, as each new member logs on, the media exchange server 605 handling the log-on procedure stores new media file records to be added to the file table 221 and new keyword records to be added to the keyword table 201 in its new record buffer 606. The record inserter 607 is inactive for most of the time, but “wakes up” or activates periodically (e.g., every 5 minutes or other suitable time period) to check all of the new record buffers 606 for each of the media exchange servers 605 and determine whether any new records need to be added to the media database 610. For each new record buffer 606 that is not empty, the record inserter 607 opens up a connection with the media database 610, writes all of the records in the new record buffer 606 to the media database 610, then closes the connection with the media database 610.

An advantage of the architecture of the media exchange web site 640 illustrated in FIG. 6 is that each member logging on does not result in the media exchange server 605 opening up a connection with the media database 610, writing the member’s media file records and keyword records to the media database 610, and closing the connection with the media database 610, all of which consumes substantial overhead, particularly when such activity is being done on a continuous and immediate basis to service all of the logging-on members from all of the media exchange servers 605. Moreover, members that log-on and log-off rapidly (i.e., within the 5 minute or other period between activity of the record inserter 607) could otherwise result in a large number of records being written to the media database 610 and then being deleted very soon afterwards. By temporarily buffering the new records for each logging-on member for a certain period of time, using the new record buffer 606, the overhead associated with a member that quickly logs on and off is avoided. As an added means for avoiding this type of overhead, the record inserter 607 may be configured to process only those new records in the new record buffer 606 associated with members that have been logged on for a certain period of time (such as, for example, 5 minutes), which may be kept track of at the media exchange server 605 using a special log-on time variable (not shown) associated with each member record in the user table 241 or with member data temporarily stored in the media exchange server 605 for the purpose of managing and administering visiting member activity.

In an alternative embodiment, a second or backup record inserter (not shown) identical in functionality to record inserter 607 is provided at the media exchange web site 640, and is connected to all of the media exchange servers 605. The record inserter 607

and the backup record inserter may “heartbeat” one another to provide redundant functionality of the record inserter 607.

In a preferred embodiment, each media exchange server 105 (or 605) comprises a cache (not shown), and as much of the media file and member information as possible is stored in the cache so as to minimize access to the media database 110 (or 610). The
5 media database 110 (or 610) ordinarily is preferably accessed only during database insert and query operations (to add new records and search existing records, respectively), with other operations being conducted, to the extent possible, on information stored in the cache. Each cache may be flushed periodically according to, for example, a least-
10 frequently used algorithm.

FIG. 7 is a block diagram is another embodiment of a media exchange web site 740 having similar functionality to the media exchange web site 640 shown in FIG. 6, but including local record inserters at some or all of the media exchange servers as opposed to a separate record inserter 607 as used in the embodiment shown in FIG. 6. Thus, similar
15 to the media exchange web site 640 shown in FIG. 6, the media exchange web site 740 of FIG. 7 includes a load balancer 704 and media database 710 which serve the same purpose and have the same functionality as the counterpart components illustrated in FIG. 6. For example, the load balancer 704 connects to a distributed electronic network and thereby to various user or member computers, not explicitly shown in FIG. 7, in a manner similar
20 to that described with respect to FIGs. 1 and 6. As further shown in FIG. 7, the media exchange web site 740 comprises a plurality of media exchange servers 705, which have the same general functionality as the media exchange servers 105 or 605 depicted in FIGs. 1 and 6, respectively. Further, each media exchange server 705 in the embodiment illustrated in FIG. 7 comprises a new record buffer 706 for temporarily storing new media
25 file records to be added to the file table 221 and new keyword records to be added to the keyword table 201.

Each of the media exchange servers 705 also includes a local record inserter 707 which serves the same general purpose as the record inserter 607 illustrated in FIG. 6, except only serving the particular media exchange server 705. In operation, as each new
30 member logs on, the media exchange server 705 handling the log-on procedure stores new media file records to be added to the file table 221 and new keyword records to be added to the keyword table 201 in its new record buffer 706. The record inserter 707 is inactive

for most of the time, but “wakes up” or activates periodically (e.g., every 5 minutes or other suitable time period) to check the new record buffer 706 for its media exchange server 705 and to determine whether any new records need to be added to the media database 710. If the new record buffer 706 is not empty, the record inserter 707 opens up
5 a connection with the media database 710, writes all of the records in the new record buffer 706 to the media database 710, then closes the connection with the media database 710.

Similar to the media exchange web site 640 shown in FIG. 6, an advantage of the architecture of the media exchange web site 740 illustrated in FIG. 7 is that each member
10 logging on does not result in the media exchange server 705 opening up a connection with the media database 710, writing the member’s media file records and keyword records to the media database 710, and closing the connection with the media database 710, all of which consumes substantial overhead, particularly when such activity is being done on a
15 continuous and immediate basis to service all of the logging-on members from all of the media exchange servers 705. Moreover, members that log-on and log-off rapidly do not increase the overhead of accessing the media database 710 significantly, due to the temporary buffering of new records for each member logging on.. To reduce database access overhead further, the record inserter 707 may be configured to process only those
20 new records in the new record buffer 706 associated with members that have been logged on for a certain period of time (such as, for example, 5 minutes), which may be kept track of at the media exchange server 705 using a special log-on time variable associated with each member record in the user table 241 or with member data temporarily stored in the media exchange server 705.

FIG. 8 is a block diagram of another embodiment of a media exchange web site
25 840, including a central user location register and exchange router 850 which preferably serves a variety of useful information and routing functions. As illustrated in FIG. 8, the media exchange web site 840 comprises load balancer 804 connected to multiple media exchange servers 805, each of which serve the same general purpose and have similar functionality as the counterpart components illustrated in FIG. 1. For example, the load
30 balancer 804 connects to a distributed electronic network and thereby to various user or member computers, not explicitly shown in FIG. 8, in a manner similar to that described with respect to FIG. 1. As further illustrated in FIG. 8, the media exchange web site 840

also includes a media database 810 which is connected to the media exchange servers 805, and which serves the same general purpose of the media database 110 described with respect to FIG. 1. The central user location register and exchange router 850 also connects to the media exchange servers 805. The central user location register and exchange router 850 stores, in a centralized location, information relating to each member's current connection with the media exchange web site 840. In particular, the central user location register and exchange router 850 stores in table form a listing of all currently connected members along with the media exchange server 805 to which each member is connected. The table stored at the central user location register and exchange router 850 may also include information relating to other member attributes, such as the number of media files available from each member. The central user location register and exchange router 850 may keep a running tally of the number of members currently connected to the media exchange web site 840 as well as the total number of media files available at a given time.

In the embodiment depicted in FIG. 8, when a member first logs on, the media exchange server 805 handling the connection preferably retrieves information from the central user location register and exchange router 850. For example, the media exchange server 805 may retrieve the tally of the total number of members currently connected and/or the total number of media files available, and pass that information to the user computer for display to the member. The media exchange server 805 handling the member connection may also retrieve information as to the on-line status of other members of particular interest to the logging-on member. In particular, the client-side software 162 of the user computer 102 (see FIG. 1) may store a list (hereafter referred to as a "hotlist") of members of particular interest to the logging-on member. The hotlist may be created by the member by on-screen selection of members (i.e., media file owners) appearing in returned listings of media files and their associated owners. For example, if the member identifies a particular media file owner that appears to have a number of media files of particular interest to the member, then the member can add that media file owner to the member's hotlist.

When a member first logs on, the member's hotlist is transmitted from the user computer 102 to the media exchange server 805 handling the connection. The media exchange server 805 then issues a query to the central user location register and exchange router 850 for each of the members on the hotlist. The media exchange server 805 keeps

track of each query returning a hit (i.e., the member is currently logged on), and conveys this information to the user computer 102 for display. Thus, when the member first logs on, the member will be presented with a listing of other members appearing on the logging-on member's hotlist, who are also currently logged on, in addition to, if desired, the total number of members currently logged on and/or the total number of media files currently available.

The central user location register and exchange router 850 may also serve an additional function within the media exchange web site 840, for facilitating the transfer of media files when a member owning a desired media file is inaccessible to other members by virtue of, for example, being protected by a firewall. In such a case, the central user location register and exchange router 850 may serve a routing function between the appropriate media exchange servers 805.

This functionality of the central user location register and exchange router 850 may be explained with respect to FIGs. 9 and 10. FIG. 9 is a diagram illustrating an example of the operation of the media exchange web site depicted in FIG. 8 in the context of an overall media exchange system, while FIG. 10 is a decision process flow associated with an example of the operation of the media exchange web site and system as illustrated in FIGs. 8 and 9. While four media exchange servers 805a-d are depicted in FIG. 9 for simplicity, the concepts described are applicable to an arbitrary number of media exchange servers. The process 900 set forth in FIG. 9 starts with the assumption that a member has selected a media file from a listing of available media files, and wants to retrieve it from the media file owner. Thus, in a first step 901, a determination is made as to whether the media file owner is accessible or not. If the media file owner is directly accessible, then the requesting member sends a media file request to the media file owner (step 905) as previously described herein. If, on the other hand, the media file owner is behind a firewall 803, as illustrated in FIG. 9 with respect to user computer 802a, then attempts to reach the media file owner by the requesting member's user computer 802b may fail. In such a case, as represented by step 910, a request for the media file is sent from the user computer 802b to the media exchange server 805b (see FIG. 9) handling the requesting member.

If the media file owner's user computer 802a is connected to the same media exchange server 805b (a situation not illustrated in the example of FIG. 9), then the media

exchange server 805b would forward the media file request to the media file owner's user computer 802a., as indicated by step 915. If, on the other hand, the media file owner's user computer 802a is connected to a different media exchange server (in the example of FIG. 9, to media exchange server 802a), then the central user location register and exchange router 850 comes into play. Specifically, the media exchange server 802b handling the requesting member forwards the media file request to the central user location register and exchange router 850. In response, the central user location register and exchange router 850 refers to its table of currently logged on members and determines which media exchange server 805a-d the media file owner is connected to. The central user location register and exchange router 850 then routes the media file request to the appropriate media exchange server 805a. The media exchange server 805a receiving the routed media file request then forwards it to the media file owner's user computer 102a., whereupon the request is fulfilled. Because the media file owner's user computer 102a is already connected to the media exchange web site 840, the firewall 803 does not pose a hindrance to the media file request.

The path of the media file request from the originating user computer 802b through the load balancer 804 to the requestor's media exchange server 805b, the central user location register and exchange router 850, the media file owner's media exchange server 805b, and to the target user computer 802a through the load balancer 804 again is represented by the message arrows 870 through 875 in FIG. 9.

It should be noted that the features described with respect to the central user location register and exchange router 850 of media exchange web site 840 can be used in conjunction with other principles and concepts as described herein, such as, for example, the record inserter(s) as depicted in embodiments of FIGs. 6 or 7.

FIG. 11 is a block diagram of another embodiment of a media exchange web site 1140, as may be employed, for example, within a media exchange system such as illustrated in FIG. 1. In one aspect, the media exchange web site 1140 illustrated in FIG. 11 includes an architecture of a hierarchical nature, preferably structured so as to facilitate rapid response times to member requests, to minimize hardware needs, and to provide ease of scalability, among other things. As illustrated in FIG. 11, the media exchange web site 1140 preferably comprises load balancer 1104 connected to multiple media exchange servers 1105, each of which serve the same general purpose and have similar functionality

as the counterpart components illustrated in FIG. 1. For example, the load balancer 804 connects to a distributed electronic network and thereby to various user or member computers, not explicitly shown in FIG. 11, in a manner similar to that described with respect to FIG. 1. As further illustrated in FIG. 11, the media exchange servers 1105 constitute leaf nodes in a hierarchical or tree structure, and directly service incoming user requests transferred from the load balancer 1104. Each group of media exchange servers 1105 is preferably connected to a mid-tier media exchange server 1130, and all of the mid-tier media exchange servers 1130 are, in turn, preferably connected to a root media exchange server 1150.

10 The media exchange web site 1140 may also include a media exchange server assignment table 1112 connected to the load balancer 1104, the function of which is described later herein.

Each of the leaf node media exchange servers 1105 preferably is configured with its own local media database 1120. Similarly, each mid-tier media exchange server 1130 has its own local media database 1135, and the root media exchange server 1160 is connected to a global media database 1160. The local media database 1120 at each media exchange server 1105 preferably stores the user and media file information for each member logged on to that particular media exchange server 1105. Such information would include, for example, a keyword table 201, file table 221 and user table 241 such as described earlier herein with respect to FIG. 2, but only for the members logged on to the particular media exchange server 1105.

The local media database 1135 for each mid-tier media exchange server 1130 stores the user and media file information for all of the members logged on to any of the leaf node media exchange servers 1105 connected to the mid-tier media exchange server 1130. Thus, the mid-tier exchange server 1130 comprises all of the information collectively stored in local media databases 1120 connected to the leaf node media exchange servers 1105. Similarly, the global media database 1160 connected to the root media exchange server 1150 stores user and media file information for all of the members currently logged on. Thus, the global media database 1160 comprises all of the information collectively stored in all of the local media databases 1120 connected to the leaf node media exchange servers 1105 (or, equivalently, all of the information collectively stored in all of the local media databases 1135 of the mid-tier media exchange

5 servers 1130). In one aspect, each of the media databases 1120, 1135 and 1160 has a different knowledge set. Each of the local media databases 1120 and 1135 has only partial knowledge of the complete set of user and media information at the media exchange web site 1140, but each of the local media databases 1120 has complete knowledge with respect to members connected to its associated media exchange server 1105, and each of the mid-tier local media databases 1135 has complete knowledge with respect to members connected to any of its "children" (i.e., the leaf node media exchange servers 1105 connected to it).

10 In operation, members log on to the media exchange web site 1140 much in the same manner as previously described with respect to the media exchange web site 140 depicted in FIG. 1. Upon log-on, the members media files are transferred (as well as parsed into keywords and other associated information, as described previously with respect to FIG. 2) to the local media database 1120 of the media exchange server 1105 to which the member is connected. Once the member's information has been parsed and stored in an appropriate format, the media files and associated data are propagated up the chain to the mid-tier media exchange server 1130, whereupon the information is stored in the mid-tier local media database 1135, and to the root media exchange server 1150, whereupon it is stored in the global media database 1160. In this manner, the media databases 1120, 1135 and 1160 are provided with their respective content.

20 When a member conducts a search for media files, the search request is sent to the media exchange server 1105 handling the particular member connection. In a preferred embodiment, the media exchange server 1105 first performs a query on its own local media database 1120 to determine if matches are found. Preferably, along with each search request, the member specifies a target number of hits desired for the particular search (the target number may be limited to a certain maximum number either by the media exchange server 1105 or the client-side software, or both). The media exchange server 1105 tallies all of the hits from its local media database 1120 and, if the number of hits are equal to or greater than the specified target number of hits, provides a response to the member. In a variation of this approach, the media exchange server 1105 and/or local media database 1120 is configured to stop searching once the specified target number of hits is reached. If the search is not so constrained, however, a selection is preferably made of the returned hits, so as to meet the target number of hits specified by the user. The

selection may be based on any of a variety of factors, such as the first hits returned, the newest media files that were hit, the oldest media files that were hit, or else a random selection of hits, to name a few possibilities. The list of media files is then transmitted from the media exchange server 1105 to the member's computer, as previously described with respect to the media exchange system 100 depicted in FIG. 1. If there are more than the target number, a message may be transmitted to the user along with the search results that there may be more media files available than those provided in the transmitted list.

If, however, not enough hits are returned by the search of the local media database 1120 connected to the media exchange server 1105, then whatever search results were obtained by the media exchange server 1105 are preferably discarded, and the search request is propagated to the mid-tier media exchange server 1130. The mid-tier media exchange server 1130 then performs a query on the mid-tier local media database 1135 according to the same search criteria provided by the member. The mid-tier media exchange server 1130 tallies all of the hits from its local media database 1135 and, if the number of hits are equal to or greater than the member-specified target number of hits, provides a response to the media exchange server 1105, which in turn conveys the response to the member. As before, if the search returns more hits than the member-specified target number, a selection of the media files that came up in the search can be accomplished in any of a wide variety of ways.

Optionally, if the leaf node media exchange server 1105 does not get the target number of search results, it may nevertheless deliver those that did come up to the member, then push the request to the mid-tier media exchange media server 1130 with a request to do the remaining search asynchronously, but block out any search results pertaining to the requesting leaf node media exchange server 1105. Preferably, in such an embodiment, the leaf node media exchange server 1105 subtracts the number of hits at the leaf node level from the member-specified target when passing the request for further searching along to the mid-tier media exchange server 1130, so as not to waste system resources.

If the search by the mid-tier media exchange server 1130 does not result in enough hits to meet the user-specified target number, then whatever search results were obtained by the mid-tier media exchange server 1130 are preferably discarded, and the search request is propagated again, this time to the root media exchange server 1150. The root

media exchange server 1150 performs a query on the global media database 1160 according to the same search criteria provided by the member. The root media exchange server 1150 tallies all of the hits from the global media database 1160 and, if the number of hits are equal to or greater than the member-specified target number of hits, provides a response to the mid-tier media exchange server 1130, which in turn conveys the response to the member via the connecting leaf node media exchange server 1105. As before, if the search returns more hits than the member-specified target number, a selection of the media files that came up in the search can be accomplished in any of a wide variety of ways. However, if the search does not return enough hits to meet the member-specified target number, then either the search results can be entirely thrown out and the user informed that not enough results were found, or else the search results can be forwarded to the member with a message explaining that the search resulting in less than the target number of specified media files being achieved.

In a separate aspect of the media exchange web site 1140 illustrated in FIG. 11, a media exchange server assignment table 1112 is utilized by the load balancer 1104 in determining which media exchange server 1105 to connect members to. FIG. 13 is an illustration of a preferred media exchange server assignment table 1130 as may be used, for example, in connection with the media exchange web site 1140 of FIG. 11. As shown in FIG. 13, the media exchange server assignment table 1130 correlates user Internet protocol (IP) addresses with the different media exchange servers 1105. When a member logs on, the member's IP address is applied to the media exchange server assignment table 1130 to determine which media exchange server 1105 to assign the member to. The media exchange server assignment table 1130 divides the various possible IP addresses into ranges (as shown in column 1131 of the table), and with each range associates one of the media exchange servers 1105 (as shown in column 1132 of the table). As a result, members with IP addresses in near proximity will generally be assigned to the same media exchange server 1105.

IP addresses are organized numerically in a hierarchical fashion, and are generally placed into one of three classes: A.x.x.x, B.B.x.x., or C.C.C.x. In this notation, the top tier is at the farthest left, and descend towards the right. The "." separates different hierarchical levels, and "x" indicates all of the IP addresses owned or controlled within the class. Each tier can have up to 256 different numerical addresses. Thus, an owner of a

class A block of IP addresses would own 256^3 different IP addresses (and, hence, could connect and individually identify up to that many different machines); an owner of class B block of IP addresses would own 256^2 different IP addresses; and an owner of a class C block of IP addresses would own 256 different IP addresses. Because all of the IP addresses within a block are commonly owned by the same entity, they are close topologically, and may even share hardware or be within the same local area network. The media exchange server assignment table 1130 therefore keeps members who are likely to be topologically close together connected to the same media exchange server 1105.

The contents of the media exchange server assignment table 1130 may be fixed, or else may be varied dynamically depending upon the nature of the members currently logged on. While having a fixed media exchange server assignment table 1130 is simpler in some respects from an implementation standpoint, varying the contents of the media exchange server assignment table 1130 dynamically may help in load balancing. To vary the contents of the media exchange server assignment table 1130 dynamically, the media exchange server assignment table 1130 may be manipulated so that each range of IP addresses has roughly the same number of addresses. To ensure that the members connecting to the same media exchange server 1105 are close from a topological standpoint (i.e., have relatively close IP addresses), if any adjustments need to be made to the load carried at a given media exchange server 1105, those members having IP addresses near the border of a neighboring IP address range are the most likely to be transferred into an adjacent group.

FIG. 12 is a diagram illustrating an example of the operation of the media exchange web site 1140 including the media exchange server assignment table 1112. As illustrated in FIG. 12, user computers 1102a and 1102b at the same local area network 1180a are connected, by virtue of the nearness of their IP addresses, to the same media exchange server 1105. When members at either of the user computers 1102a and 1102b log on, the member's IP address is applied to the media exchange server assignment table 1112, which directs the load balancer 1104 to route the user computer 1102a or 1102b to the appropriate media exchange server 1105. Likewise, user computers 1102c, 1102d and 1102e are connected within the same local area network 1180b and, by virtue of the nearness of their IP addresses, will be directed by the load balancer to the same media

exchange server 1105 (but a different media exchange server 1105 potentially than that connecting to user computers 1102a and 1102b).

When a member makes a search request, the search results are returned first, if possible from the media exchange server 1105, assuming enough media files are found to meet the member's specified target. Therefore, when the member attempts to send a media file transfer request to the media file owner and effectuate a transaction thereby, the member is most likely to connect to another member who is topologically close. As a result, the demands on the network infrastructure are reduced, since media file transfers can often be made by connections internal to a local area network 1180a or 1180b, or the same Internet service provider (ISP), rather than having to connect externally to other systems. This is an advantage both to the members involved in the media file transfer as well as the particular network in which the members are located.

A variety of advantages, in addition to those already mentioned, can be experienced when utilizing the media exchange web site architecture illustrated in FIG. 11. First, most searches of media files are expected to return faster because the local media database 1120 at the media exchange server 1105 can be relatively small. Since it is expected that many media file searches will be for the most popular media files, it would be expected that, despite the fact only a subset of the members are connected to a given media exchange server 1105, the relevant media files are likely to be found on the same media exchange server 1105 receiving the searching member's search request. Even a search at a mid-tier level (i.e., a search of mid-tier local database 1135) is expected to be relatively quick, since again it only contains information from a subset (albeit a larger one) of all of the logged on members.

Since many of the searches will be able to be served by the leaf node media exchange servers 1105 or mid-tier media exchange servers 1130, the global media database 1160 will receive fewer search requests. Consequently, the global media database 1160 does not have as strict speed requirements and can be less costly. The leaf-node media exchange servers 1105 can be embodied as relatively simple, and hence inexpensive, machines such as dual-processor personal computers (PCs) or functional equivalents. Since the local media databases 1120 can be relatively small, they too can be relatively inexpensive.

Another advantage of the media exchange web site 1140 is that the architecture can be scaled rather easily. In some embodiments, for example, the mid-tier media exchange servers 1130 may be eliminated, so that only two tiers in the tree structure are implemented. Similarly, additional tiers may be added, so as to result in a tree structure having an arbitrary number of tiers. To help illustrate this concept, FIG. 14 is a conceptual diagram of a generic node 1400 as may be employed in the media exchange web site of FIG. 11 or variations thereof. The node 1400 includes a media exchange server 1405 which connects to a local database 1420, and may also have a local table or cache 1421 for increasing response time. The media exchange server 1400 connects to N children (i.e., other nodes), and generally has a single connection to a parent (although connections to multiple parents are also possible). Expanding the media exchange web site 1140 through the addition of further tiers or nodes is relatively quick and inexpensive, given the consistency of the node-based architecture throughout.

Yet another advantage of the media exchange web site architecture illustrated in FIG. 11 is that it is robust. For example, if a mid-tier media exchange server 1130 or the root media exchange server 1150 becomes inoperable for any reason, the sub-network (i.e., media exchange servers 1105 or sub-network media exchange servers 1125, as the case may be) connected to the failed upstream node can still operate and respond to search requests on their own. Thus, system functionality is maintained despite a node failure. As an alternative, the sub-network can be temporarily switched to a different parent upon such a failure. For example, if a first mid-tier media exchange server 1130 failed, the leaf node media exchange servers 1105 connected to it could be temporarily switched over to a different mid-tier media exchange server 1130, or even to the root media exchange server 1150.

It should be noted that the features of the media exchange web site 1140 illustrated in FIG. 11 can be used in conjunction with other principles and concepts as described herein, such as, for example, the record inserter(s) as depicted in embodiments of FIGs. 6 or 7.

In various alternative embodiments, less functionality, or additional functionality, may be provided at the media exchange web site. For example, the media exchange web site need not support a media file search capability, or need not support a browse capability. Furthermore, the media exchange web site in any of the various embodiments

described herein may provide features such as electronic bulletin (message posting) boards, on-line chat, and other standard features.

Logs may be maintained for each transaction carried out at the media exchange web site, in any of the embodiments described herein. For example, a log may be maintained of the search criteria submitted by each visiting member. A log may also be maintained of what files were requested by each member in response to receiving the search results. The file transfer information may be maintained by the client-side software at the user computer, and downloaded to the media exchange web site each time the member logs on.

The media exchange server(s) utilized in any of the various embodiments described herein may have the capability of blocking out members on a semi-permanent or permanent basis, should it become necessary and desirable to do so. To facilitate this functionality, a special parameter may be stored in the user state field 244 of the member record in the user table 241 to indicate that a member is to be denied the services of the media exchange web site.

Although certain embodiments herein have been described with respect to user computers connecting to a distributed electronic network, it will be appreciated that other types of electronic tools (such as personal digital assistants, telephones, or web-televisions) which may or may not strictly be termed "computers" are just as capable of serving the same functionality as that described with respect to the user computers herein. Therefore, the scope of the invention is not intended to be limited by the nature of the user systems or tools used to access the media exchange web site.

While preferred embodiments of the invention have been described herein, many variations are possible which remain within the concept and scope of the invention. Such variations would become clear to one of ordinary skill in the art after inspection of the specification and the drawings. The invention therefore is not to be restricted except within the spirit and scope of any appended claims.

CLAIMS

What is claimed is:

- 5 1. A system for facilitating media exchange over a distributed electronic network, comprising:
- a load balancer connected to the distributed electronic network for communicating with a plurality of user systems;
- a plurality of media exchange servers connected to said load balancer; and
- 10 a media database for storing names of media files from a plurality of visiting users in a plurality of media file name records, and for storing keyword records corresponding to individual words in the names of said media files;
- wherein said plurality of media exchange servers are each capable of accessing said media database to read media file records or keyword records in said media database.
- 15 2. The system of claim 1, wherein at least two of said media exchange servers are capable of simultaneously accessing said media database
3. The system of claim 1, wherein at least two of said media exchange servers
- 20 are capable of simultaneously performing queries on keyword records stored in said media database
4. The system of claim 1, wherein a selected list of the media file names stored in said media database is transmitted to a first user system in response to a request
- 25 from said first user system.
5. The system of claim 4, wherein said first user system comprises client-side software for accessing said load balancer over said distributed electronic network.
- 30 6. The system of claim 5, wherein said first user system and a second user system are collectively capable of conducting a peer-to-peer of a media file transfer.

7. The system of claim 6, wherein said peer-to-peer media file transfer is conducted in response to a user selection of a media file name from said list of media file names.

5 8. The system of claim 7, wherein said client-side software displays said list of media file names on a screen at said user system, wherein said user selection of said media file name is conveyed to said client-side software, and wherein said client-side software initiates said peer-to-peer media file transfer.

10 9. The system of claim 1, wherein said media database comprises a file table comprising said media file name records, a keyword table comprising said keyword records, and a user table comprising a plurality of user records.

15 10. The system of claim 9, wherein said media file name records are linked to a user record associated with a visiting user whose user system stores the media files corresponding to the media file name records.

20 11. The system of claim 1, wherein at least one media exchange server is configured to carry out an identity search on said keyword records in response to search criteria received from one of said user systems.

12. The system of claim 1, wherein at least one of said media exchange servers comprises a cache for storing data from said media database.

25 13. The system of claim 1, wherein each of said media exchange servers comprises a new record buffer for temporarily storing new media file name records and new keyword records.

30 14. The system of claim 13, further comprising a new record inserter connected to each of said media exchange servers, said new record inserter periodically transferring the contents of each new record buffer to the media database in a single collective disk access operation.

15. The system of claim 14, wherein said new record inserter remains dormant for a majority of time, periodically activating to transfer the contents of each new record buffer to the media database.

5

16. The system of claim 1, wherein the media files name records for a visiting user are erased from said media database immediately upon log-off of said visiting user.

17. The system of claim 1, wherein said user systems comprise computers.

10

18. The system of claim 1, wherein said distributed electronic network comprises the Internet.

19. The system of claim 1, wherein two or more of said media exchange servers are connected to an exchange router, said exchange router capable of routing media file transfer requests from one media exchange server to another media exchange server.

15

20. A system for facilitating the exchange of digital files over a distributed electronic network, comprising:

20

a load balancer connected to the distributed electronic network for communicating with a plurality of user computers;

a plurality of servers connected to said load balancer; and

a database for storing names of digital files from a plurality of visiting users in a plurality of file name records, and for storing keyword records corresponding to individual words in the names of said files;

25

wherein said load balancer, servers and database collectively comprise a website accessible to said user computers over said distributed electronic network; and

wherein said plurality of servers are each capable of accessing said database to read file records or keyword records in said database, at least two of said servers capable of simultaneously accessing said database.

30

21. The website of claim 20, wherein at least two of said servers are capable of simultaneously performing queries on keyword records stored in said database in response to requests received from said user computers over said distributed electronic network.

5 22. The system of claim 20, wherein selected lists of file names stored in said database are transmitted to user computers in response to search requests received from said user computers over said distributed electronic network.

10 23. The system of claim 22, wherein each of said user computers comprises client-side software for accessing said website over said distributed electronic network.

15 24. The system of claim 23, wherein said user systems are capable of conducting peer-to-peer of file transfers, said peer-to-peer file transfers managed by the client-side software at said user systems.

20 25. The system of claim 24, wherein said selected lists of file names include, for each file name in a list, linking information to the user owning the file associated with the file name.

25 26. The system of claim 25, wherein said peer-to-peer file transfers are conducted in response to user selections of file names from said selected lists of file names transmitted to said user computers.

30 27. The system of claim 25, wherein the client-side software at a user computer displays a selected list of file names on a screen at the user computer on which said client-side software is run, and wherein said client-side software initiates a peer-to-peer file transfer using the linking information from a file name selected at said user computer.

35 28. The system of claim 20, wherein said database comprises a file table comprising said file name records, a keyword table comprising said keyword records, and a user table comprising a plurality of user records.

29. The system of claim 28, wherein said file name records are linked to a user record associated with a visiting user whose user computer stores the files corresponding to the file name records.

5 30. The system of claim 20, wherein each of said servers comprises a new record buffer for temporarily storing new file name records and new keyword records.

31. The system of claim 30, further comprising a new record inserter connected to each of said servers, said new record inserter periodically transferring the contents of a plurality of the new record buffers to the database in a single disk access operation.
10

32. The system of claim 31, wherein said new record inserter remains dormant for a majority of time, periodically activating to transfer the contents of the new record buffers to the database.
15

33. The system of claim 32, wherein the files name records for a visiting user are erased from said database immediately upon log-off of said visiting user.

34. The system of claim 20, wherein said distributed electronic network comprises the Internet.
20

35. A media exchange system, comprising:

a plurality of media exchange servers arranged in a hierarchy;

a plurality of media databases;

25 a load balancer connected to said media exchange servers and to a distributed electronic network;

means for conducting a search of media file names stored in said plurality of media databases.

30 36. A method for obtaining digital media files over a distributed electronic network, comprising the steps of:

connecting a user computer, over the distributed electronic network, to a media file exchange computer system;

entering media file search criteria at the user computer;

5 transmitting from the user computer, over the distributed electronic network, said media file search criteria to a first server of the media file exchange computer system; and

10 receiving at said user computer, from the first server of the media file exchange system and over the distributed electronic network, a list of files meeting the media file search criteria, said list of files originating from a plurality of user systems connected to a plurality of servers, including said first server and at least one other server, of the media file exchange computer system.

15 37. The method of claim 36, further comprising the step of transferring from the user computer, over the distributed electronic network, identifying information for one or more media files stored on said user computer to said first server of the media file exchange computer system.

38. The method of claim 37, wherein said one or more media files are stored in a designated directory on said user computer.

20 39. The method of claim 36, wherein the distributed electronic network comprises the Internet.

25 40. The method of claim 36, wherein said media files include one or more of digital audio files, video files, multi-media files, text files, data files, and graphics files.

41. The method of claim 36, further comprising the steps of:
transmitting from said user computer to said media file exchange computer system, over the distributed electronic network, a pre-designated list of users; and
receiving from the media file exchange computer system, over the distributed
30 electronic network, an indication of which of the users is currently engaged in an active session with the media file exchange computer system.

42. The method of claim 41, wherein said acts further comprise:

receiving at the user computer from the media file exchange computer system, over the distributed electronic network, a list of user systems associated with the list of files received from the media file exchange computer system;

5 displaying the list of user systems at the user computer;

receiving a user selection, input at said user computer, indicating a selected user system; and

adding said user system to the pre-designated list of users.

10 43. The method of claim 36, wherein said list of files includes, for each file in the list, linking information to the user system on which the file is presently stored.

44. The method of claim 43, further comprising the step of displaying, at said user computer, said list of files.

15

45. The method of claim 44, further comprising the steps of:

receiving, at said user computer, a user selection of a file from said list of files;

transmitting a file transfer request from said user computer, over the distributed electronic network, to a remote user computer on which the selected file is presently stored, said remote user computer identified by said linking information; and

20

receiving at said user computer, in response to said file transfer request and over the distributed electronic network, a media file from said remote user computer.

46. The method of claim 45, wherein said step of connecting said user computer to the media file exchange computer system comprises the step of logging on said user computer to the media file exchange computer system, said method further comprising the step of logging off said user computer prior to said step of transmitting said file request to the remote user computer.

25

30

47. The method of claim 36, further comprising the step of:

receiving at the user computer a user selection of a specified target number of desired hits for the media file search criteria;

transmitting, from the user computer and over the distributed electronic network, the specified target number of desired hits to the media file exchange computer system along with the media file search criteria.

5 48. A computer readable medium having stored therein one or more sequences of instructions for obtaining digital media files over a distributed electronic network, said one or more sequences of instructions causing one or more processors to perform a plurality of acts, said acts comprising:

10 connecting a user computer, over the distributed electronic network, to a media file exchange computer system;

 receiving media file search criteria input at the user computer;

 transmitting from the user computer, over the distributed electronic network, said media file search criteria to a first server of the media file exchange computer system; and

15 receiving at said user computer, from the first server of the media file exchange system and over the distributed electronic network, a list of files meeting the media file search criteria, said list of files originating from a plurality of user systems connected to a plurality of servers, including said first server and at least one other server, of the media file exchange computer system.

20 49. The computer readable medium of claim 48, wherein said acts further comprise transferring from the user computer, over the distributed electronic network, identifying information for one or more media files stored on said user computer to said first server of the media file exchange computer system.

25 50. The computer readable medium of claim 49, wherein said one or more media files are stored in a designated directory on said user computer.

 51. The computer readable medium of claim 48, wherein the distributed electronic network comprises the Internet.

52. The computer readable medium of claim 48, wherein said media files include one or more of digital audio files, video files, multi-media files, text files, data files, and graphics files.

5 53. The computer readable medium of claim 48, wherein said acts further comprise:

transmitting from said user computer to said media file exchange computer system, over the distributed electronic network, a pre-designated list of users; and

10 receiving from the media file exchange computer system, over the distributed electronic network, an indication of which of the users is currently engaged in an active session with the media file exchange computer system.

54. The computer readable medium of claim 53, wherein said acts further comprise:

15 receiving at the user computer from the media file exchange computer system, over the distributed electronic network, a list of user systems associated with the list of files received from the media file exchange computer system;

displaying the list of user systems at the user computer;

20 receiving a user selection, input at said user computer, indicating a selected user system; and

adding said user system to the pre-designated list of users.

55. The computer readable medium of claim 48, wherein said list of files includes, for each file in the list, linking information to the user system on which the file is
25 presently stored.

56. The computer readable medium of claim 55, wherein said acts further comprise displaying, at said user computer, said list of files.

30 57. The computer readable medium of claim 56, wherein said acts further comprise:

receiving a user selection, input at said user computer, indicating a file from said list of files;

transmitting a file transfer request from said user computer, over the distributed electronic network, to a remote user computer on which the selected file is presently stored, said remote user computer identified by said linking information; and

receiving at said user computer, in response to said file transfer request and over the distributed electronic network, a media file from said remote user computer.

58. The computer readable medium of claim 57, wherein the act of connecting said user computer to the media file exchange computer system comprises the act of logging on said user computer to the media file exchange computer system, and wherein said one or more sequences of instructions further causing the one or more processors to perform the act of logging off said user computer prior to transmitting said file request to the remote user computer.

15

59. The computer readable medium of claim 48, wherein said acts further comprise:

receiving at the user computer a user selection of a specified target number of desired hits for the media file search criteria;

20

transmitting, from the user computer and over the distributed electronic network, the specified target number of desired hits to the media file exchange computer system along with the media file search criteria.

60. A system for facilitating media exchange over a distributed electronic network, comprising:

25

a plurality of front-end media exchange servers for receiving information from and sending information to remote user systems, each of said front-end media exchange servers connected to a local database for storing names of data files from said remote user systems;

30

a global database for storing names of data files from said remote user systems; and

a root media exchange server electronically accessible to said plurality of front-end media exchange servers and connected to said global database;

5 wherein each front-end media exchange server responds to incoming file search requests from a remote user system communicating therewith by initially checking its local database for hits meeting the file search request and, if the number of hits are less than a specified minimum number, relaying the file search request upstream to another server.

10 61. The system of claim 60, further comprising a load balancer connected to the distributed electronic network for communicating with said remote user systems and distributing newly connected remote user systems to said plurality of front-end media exchange servers.

15 62. The system of claim 60, wherein said names of data files are stored in said global database in a plurality of file name records, and wherein said global database stores keyword records corresponding to individual words in the names of said data files.

20 63. The system of claim 60, wherein, if the number of hits are less than said specified minimum number, the file search request is relayed from the front-end media server to said root media exchange server.

64. The system of claim 60, wherein said specified minimum number is transmitted from a remote user system along with the file search request.

25 65. The system of claim 60, further comprising a plurality of mid-tier media exchange servers interposed between said plurality of front-end media exchange servers and said root media exchange server, each of said mid-tier media exchange servers connected to a unique group of said front-end media exchange servers, each of said mid-tier media exchange servers connected to a mid-tier local database for storing the names of
30 the data files collectively from the local databases of the front-end media exchange servers to which the mid-tier media exchange server is connected.

66. The system of claim 65, wherein, if the number of hits are less than said specified minimum number, the file search request is relayed from the front-end media server to the mid-tier media exchange server, which responds to the file search request by initially checking its mid-tier local database for hits meeting the file search request and, if
5 the number of hits are less than a specified minimum number, relaying the file search request upstream to said root media exchange server.




FIG. 1




FIG. 2

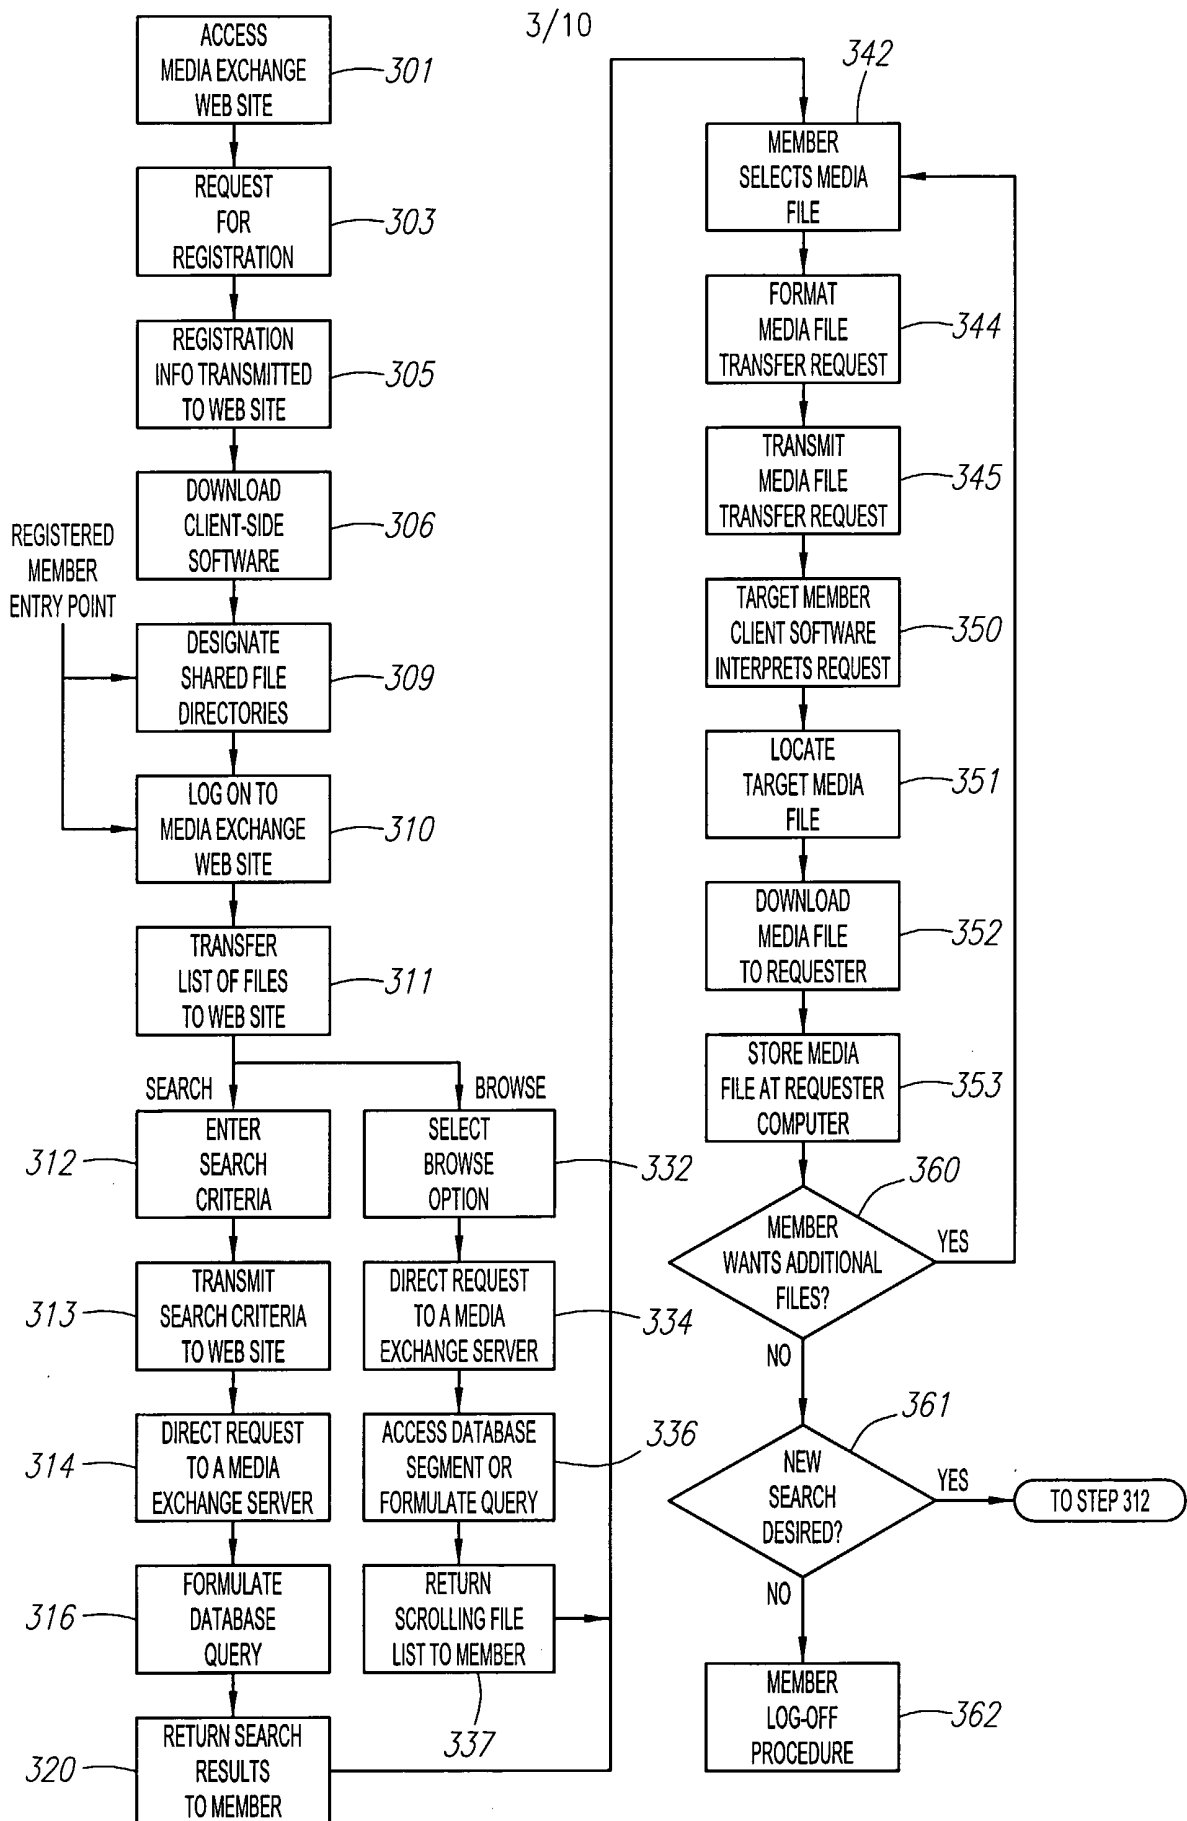


FIG. 3

4/10

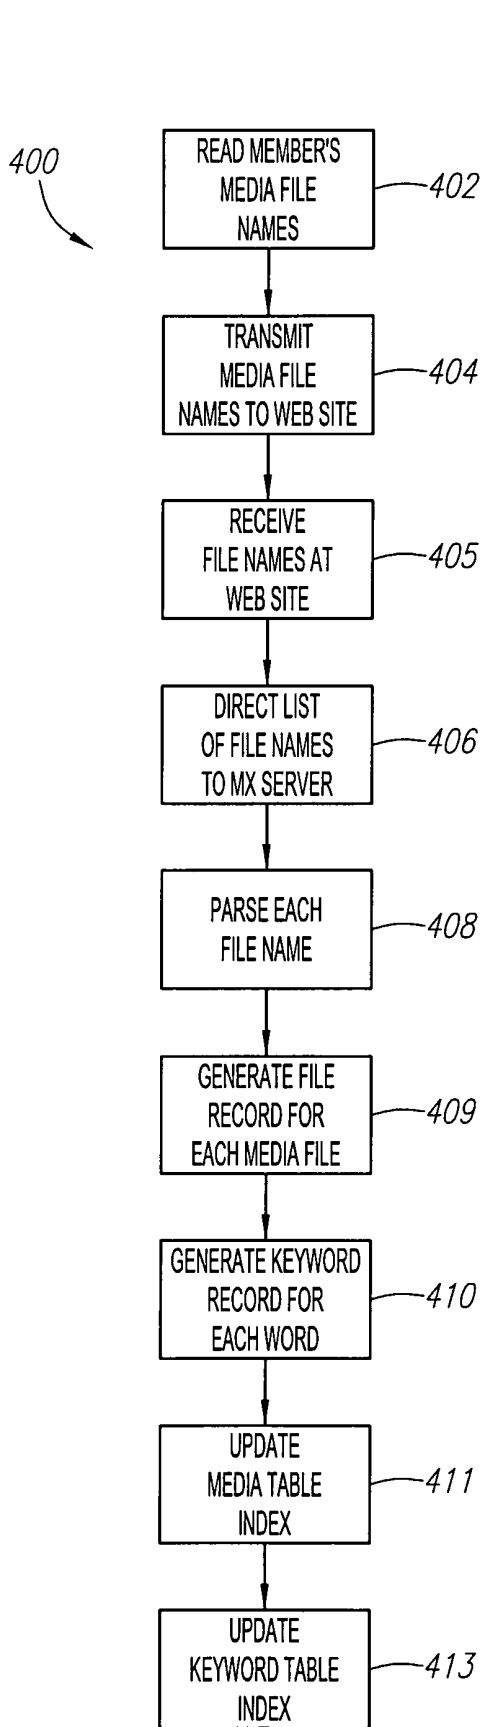


FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9

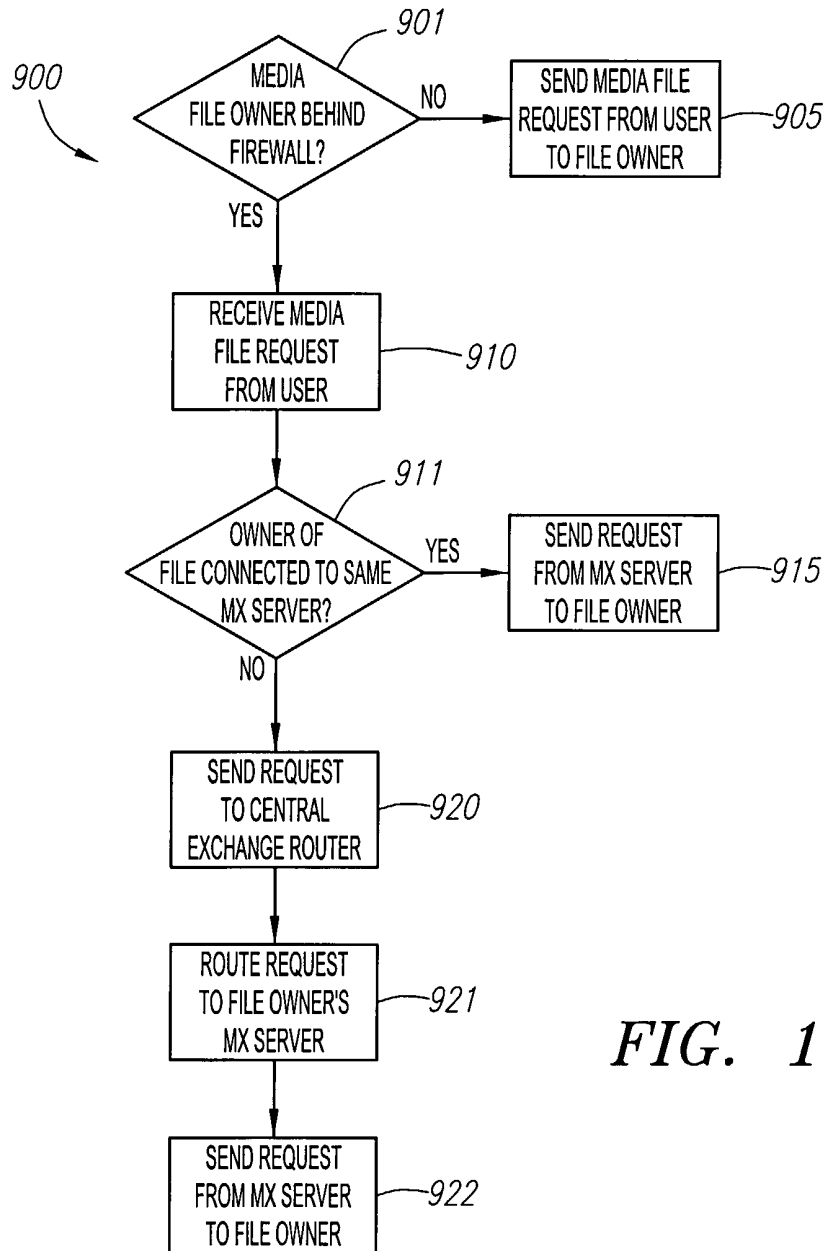


FIG. 10

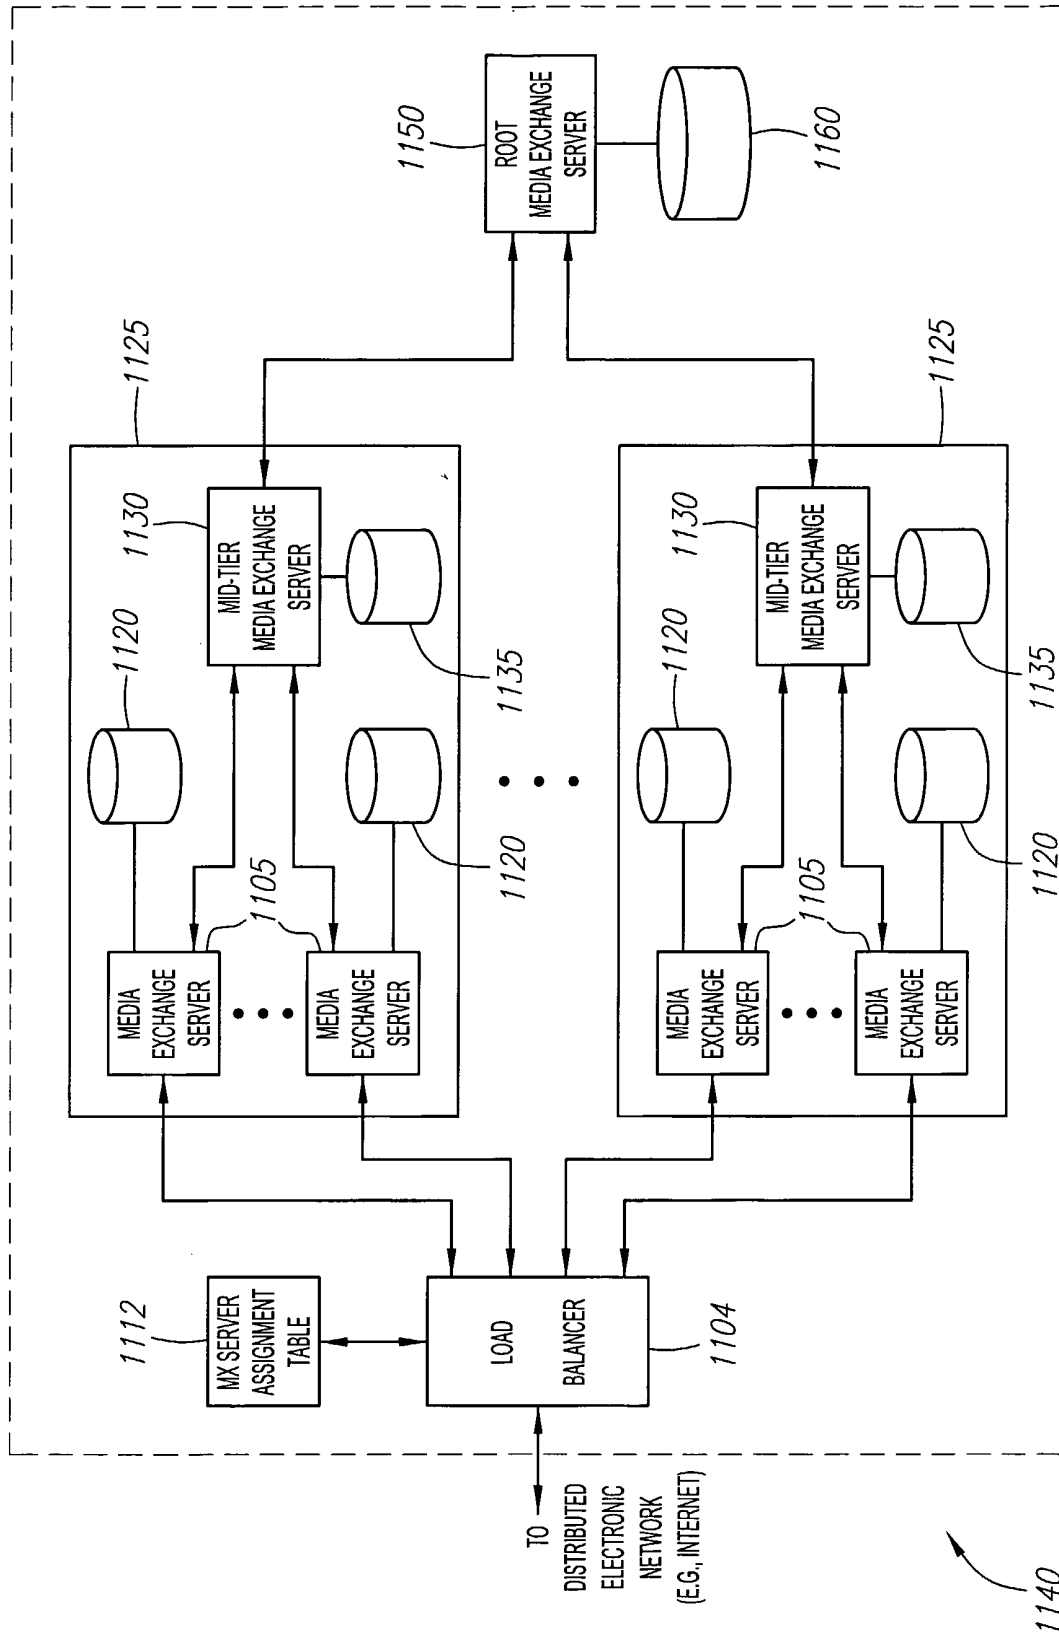


FIG. 11




FIG. 12

1301

1302

1300

IP ADDRESS	MEDIA EXCHANGE SERVER ASSIGNMENT
RANGE 1	MX SERVER #1
RANGE 2	MX SERVER #2
RANGE 3	MX SERVER #3
⋮	
RANGE N	MX SERVER #N

MEDIA EXCHANGE SERVER ASSIGNMENT TABLE

FIG. 13




FIG. 14