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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
20 November 2014 (20.11.2014)

WIPO | PCT

(10) International Publication Number
WO 2014/186444 A1

(51) International Patent Classification:
G06F 17/30 (2006.01)

(21) International Application Number:
PCT/US2014/037973

(22) International Filing Date:
14 May 2014 (14.05.2014)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/823,086 14 May 2013 (14.05.2013) US
14/109,433 17 December 2013 (17.12.2013) US

(71) Applicant: **GOOGLE INC.** [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).

(72) Inventors: **KULICK, Matthew**; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). **PREETHAM, Arcot J.**; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). **SAMBHU, Manish Mahesh**; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). **NOSHADI, Hyduke**; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).

(74) Agents: **HARRITY, John E.** et al.; Harrity & Harrity, LLP, 11350 Random Hills Road, Suite 600, Fairfax, Virginia 22030 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

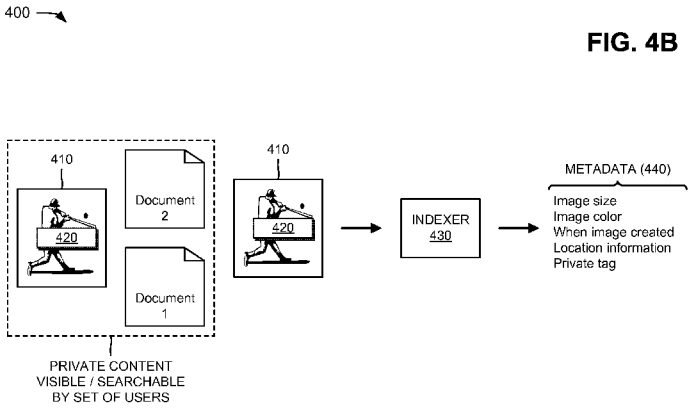
(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: INDEXING AND SEARCHING DOCUMENTS WITH RESTRICTED PORTIONS

FIG. 4B



(57) Abstract: One or more computer devices may receive a document from private content that is accessible by a set of users, and may extract metadata from the document. The metadata may provide information relating to one or more aspects of the document. The one or more computer devices may determine, based on the metadata, whether a portion of the document includes information that is only to be accessible by a particular user from the set of users, and may append a restrict indicator to the metadata when the portion of the document includes information that is only to be accessible by the particular user. The restrict indicator may identify the particular user and the portion of the document. The one or more computer devices may store the metadata and the restrict indicator in a personal search index.

WO 2014/186444 A1

INDEXING AND SEARCHING DOCUMENTS WITH RESTRICTED PORTIONS

BACKGROUND

Current personal search systems enable a set of users to search private content accessible only to the set of users. For example, the personal search systems may provide search results that include tips, photos, posts, etc. from the set of users and are personal and private to the set of users. However, in such personal search systems, if one user of the set of users has access to a document or a piece of content, all users of the set of users can search the metadata of the document in an index associated with the personal search system.

SUMMARY

Described herein are devices and techniques for appending a restrict indicator to metadata associated with a document that includes a portion that is only to be accessible by particular user(s) of a set of users. The devices and techniques may store the document metadata and the restrict indicator in a personal search index. The devices and techniques may interpret a user's search query and may add restrictions to the search query based on information associated with the user. The personal search index may return a search result that identifies the document that includes the portion that is only to be accessible by the user.

In some possible implementations, a method, performed by one or more computer devices, may include: receiving a document from private content, the private content being accessible by a set of users; extracting metadata from the document, the metadata providing information relating to one or more aspects of the document; determining, based on the metadata, whether a portion of the document includes information that is only to be accessible by a particular user from the set of users; appending a restrict indicator to the metadata when the portion of the document includes information that is only to be accessible by the particular user, the restrict indicator identifying the particular user and the portion of the document; and storing the metadata and the restrict indicator in a personal search index.

In some possible implementations, the method may further include storing only metadata in the personal search index for a second document that is accessible to the set of users.

In some possible implementations, the method may further include: receiving a search query from the particular user; adding a restriction to the search query based on information associated with the particular user; and identifying search results, based on the search query and

the restriction, from the personal search index, the search results including information identifying the document.

In some possible implementations, the method may further include: comparing the restriction and the restrict indicator; and identifying, from the personal search index, the document in the search results based on the comparing.

In some possible implementations, the method may further include: scoring the search results; and providing the scored search results to the particular user.

In some possible implementations, the method may further include: receiving a selection of a particular scored search result of the scored search results, the particular scored search result being associated with the document; and providing the document and the portion of the document to the particular user.

In some possible implementations, the method may further include: receiving a selection of a particular scored search result of the scored search results, the particular scored search result being associated with a second document that includes a restricted portion that is not accessible by the particular user; and providing the second document, without the restricted portion, to the particular user.

In some possible implementations, a device may include one or more processors to: receive a document from private content that is accessible by a set of users; extract metadata from the document, the metadata providing information relating to one or more aspects of the document; determine, based on the metadata, whether a portion of the document includes information that is only to be accessible by a particular user from the set of users; append a restrict indicator to the metadata when the portion of the document includes information that is only to be accessible by the particular user, the restrict indicator identifying the particular user and the portion of the document; and store the metadata and the restrict indicator in a personal search index.

In some possible implementations, the one or more processors may be further to: store only metadata in the personal search index for a second document that is accessible to the set of users.

In some possible implementations, the one or more processors may be further to: receive a search query from the particular user; add a restriction to the search query based on information associated with the particular user; and identify search results, based on the search query and the

restriction, from the personal search index, the search results including information identifying the document.

In some possible implementations, the one or more processors may be further to: compare the restriction and the restrict indicator; and identify, from the personal search index, the document in the search results based on the comparing.

In some possible implementations, the one or more processors may be further to: score the search results; and provide the scored search results to the particular user.

In some possible implementations, the one or more processors may be further to: receive a selection of a particular scored search result of the scored search results, the particular scored search result being associated with the document; and provide the document and the portion of the document to the particular user.

In some possible implementations, the one or more processors may be further to: receive a selection of a particular scored search result of the scored search results, the particular scored search result being associated with a second document that includes a restricted portion that is not accessible by the particular user; and provide the second document, without the restricted portion, to the particular user.

In some possible implementations, a computer-readable medium may store one or more instructions that, when executed by a processor of a device, cause the processor to: receive a document from private content that is accessible by a set of users; extract metadata from the document, the metadata providing information relating to one or more aspects of the document; determine, based on the metadata, whether a portion of the document includes information that is only to be accessible by a particular user from the set of users; append a restrict indicator to the metadata when the portion of the document includes information that is only to be accessible by the particular user, the restrict indicator identifying the particular user and the portion of the document; and store the metadata and the restrict indicator in a personal search index.

In some possible implementations, the computer-readable medium may further include one or more instructions that, when executed by the processor, cause the processor to: store only metadata in the personal search index for a second document that is accessible to the set of users.

In some possible implementations, the computer-readable medium may further include one or more instructions that, when executed by the processor, cause the processor to: receive a search query from the particular user; add a restriction to the search query based on information

associated with the particular user; and identify search results, based on the search query and the restriction, from the personal search index, the search results including information identifying the document.

In some possible implementations, the computer-readable medium may further include one or more instructions that, when executed by the processor, cause the processor to: compare the restriction and the restrict indicator; and identify, from the personal search index, the document in the search results based on the comparing.

In some possible implementations, the computer-readable medium may further include one or more instructions that, when executed by the processor, cause the processor to: score the search results; and provide the scored search results to the particular user.

In some possible implementations, the computer-readable medium may further include one or more instructions that, when executed by the processor, cause the processor to: receive a selection of a particular scored search result of the scored search results, the particular scored search result being associated with the document; and provide the document and the portion of the document to the particular user.

In some possible implementations, a system may include means for receiving a document from private content that is accessible by a set of users; means for extracting metadata from the document, the metadata providing information relating to one or more aspects of the document; means for determining, based on the metadata, whether a portion of the document includes information that is only to be accessible by a particular user from the set of users; means for appending a restrict indicator to the metadata when the portion of the document includes information that is only to be accessible by the particular user, the restrict indicator identifying the particular user and the portion of the document; and means for storing the metadata and the restrict indicator in a personal search index.

The above discussion mentions examples in which some implementations may be implemented via one or more methods performed by one or more processors of one or more devices. In some implementations, one or more systems and/or one or more devices may be configured to perform one or more of the acts mentioned above. In some implementations, a computer-readable medium may include computer-executable instructions which, when executed by one or more processors, cause the one or more processors to perform one or more of the acts mentioned above.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate some implementations described herein and, together with the description, explain these implementations. In the drawings:

Figs. 1A and 1B are diagrams illustrating an overview of some implementations described herein;

Fig. 2 is a diagram of an example environment in which systems and/or methods described herein may be implemented;

Fig. 3 is a flowchart of an example process for indexing documents with restricted portions;

Figs. 4A-4C are diagrams of an example of the process of Fig. 3;

Fig. 5 is a flowchart of an example process for searching documents with restricted portions;

Figs. 6A-6D are diagrams of an example of the process of Fig. 5; and

Fig. 7 is an example of a generic computer device and a generic mobile computer device according to some implementations.

DETAILED DESCRIPTION

The following detailed description refers to the accompanying drawings. The same reference numbers in different drawings may identify the same or similar elements.

Systems and/or methods, as described herein, may provide an indexing mechanism that appends a restrict indicator to metadata associated with a document that includes a portion that is only to be accessible by particular user(s) (e.g., an owner or creator of the document) of a set of users. The indexing mechanism may store the document metadata and the restrict indicator in a personal search index. The systems and/or methods may provide a query mechanism that interprets a user's search query and adds restrictions to the search query based on information associated with the user (e.g., user identification). For example, the query mechanism may add a restriction that enables the user to search, in the personal search index, for a document that includes a portion that is only to be accessible by the user. The personal search index may return a search result that identifies the document that includes the portion that is only to be accessible by the user. Whereas another user issuing the same search query, and who does not have access to the portion of the document, may not receive the search result that identifies the document.

The term document, as used herein, is to be broadly interpreted to include any machine-readable and machine-storable work product. A document may include, for example, a website, an e-mail, a business listing, a file, a combination of files, one or more files with embedded links to other files, a news group posting, a blog, etc. In the context of the Internet, a common document is a web page. Documents often include textual information and may include embedded information, such as meta information, images, hyperlinks, etc., and/or embedded instructions, such as Javascript, etc. In some implementations, metadata (e.g., face tags, computer vision labels, and other types of metadata) may relate to, be derived, and/or be associated with a document.

User interfaces, as described herein, may include graphical user interfaces (GUIs) and/or non-graphical user interfaces, such as text-based interfaces. The user interfaces may provide information to users via customized interfaces (e.g., proprietary interfaces) and/or other types of interfaces (e.g., browser-based interfaces, etc.). The user interfaces may receive user inputs via one or more input devices, may be user-configurable (e.g., a user may change the sizes of the user interfaces, information displayed in the user interfaces, color schemes used by the user interfaces, positions of text, images, icons, windows, etc., in the user interfaces, etc.), and/or may not be user-configurable. Information associated with the user interfaces may be selected and/or manipulated by a user of a user device (e.g., via a touch screen display, a mouse, a keyboard, a keypad, voice commands, etc.).

Figs. 1A and 1B are diagrams illustrating an overview of some implementations described herein. For the overview and as shown in Fig. 1A, assume that a personal search system includes private content (e.g., one or more documents) that may be visible and/or searchable (i.e., accessible) by a set of users. Further assume that the private content includes a particular document (e.g., Document 3) with a portion that is to be visible and/or searchable only by particular user(s) of the set of users. As further shown in Fig. 1A, the personal search system may include an indexer component that receives the documents from the private content. The indexer component may extract metadata associated with each document, and may store the document metadata in a personal search index. If the document is an image, the metadata may include, for example, information that describes how large the image is, a color depth of the image, a resolution of the image, when the image was created, etc. If the document is a text document, the metadata may include, for example, information about a length of the document,

an author of the document, when the document was written, a summary of the document, etc. For example, as shown in Fig. 1A, the document metadata may include metadata associated with a first document (e.g., Document 1), a second document (e.g., Document 2), etc.

The indexer component may determine, based on the document metadata, whether a portion of a document includes information that is only to be accessible by a particular user(s) from the set of users. For example, the indexer component may determine that a portion of the particular document (e.g., Document 3) includes information that is only to be accessible by the particular user(s). The indexer component may append a restrict indicator to the metadata of the particular document. The restrict indicator may identify the particular user and the portion of the particular document. In some implementations, the indexer component may determine that a portion of a document is accessible to a certain class of users, and may assign a restrict indicator to the certain class of users (e.g., as opposed to enumerating each member of the certain class). As further shown in Fig. 1A, the indexer component may store the particular document metadata and the restrict indicator in the personal search index.

With reference to Fig. 1B, assume that personal search system includes a search engine component that receives a search query from a user. Assume further that the search query relates to the particular document (e.g., Document 3) and/or the restricted portion of the particular document. The search engine component may add a restriction(s) to the search query based on information associated with the user (e.g., user identification). The restriction(s) may include information identifying document portions that are accessible by the user. For example, the restriction(s) may specify that the portion of the particular document (e.g., Document 3) is accessible by the user.

The search engine component may compare the search query and the restriction(s) with the personal search index, and may identify search result(s) based on the search query and the restriction(s). For example, the search engine component may identify the particular document in the search result(s) based on the restriction(s) matching the restrict indicator. The search engine component may provide the search result(s) to a scoring component of the personal search system. The scoring component may score the search result(s), and may provide the scored search result(s) to the user, via the search engine component. At least one of the scored search result(s) may include, for example, information associated with the particular document (e.g., Document 3), the restricted portion of the particular document, etc.

Fig. 2 is a diagram of an example environment 200 in which systems and/or methods described herein may be implemented. Environment 200 may include a user device 210 connected to multiple server devices 220-240 via a network 250.

User device 210 may include a device capable of interacting with server device 220 to perform a search. Examples of user device 210 may include a personal computer, a laptop computer, a tablet computer, a personal digital assistant, a mobile telephone, such as a smart phone, etc. User device 210 may include user interfaces presented through one or more browsers, e.g., web browsers.

Server devices 220-240 may include server devices that gather, process, search, and/or implement functions in a manner described herein. Each server device 220-240 may be implemented as a single server device or a collection of server devices that may be co-located or remotely located. In some implementations, any two or more of server devices 220-240 may be implemented within a single, common server device or a single, common collection of server devices. While server devices 220-240 are shown as separate components, it may be possible for one or more of server devices 220-240 to perform one or more of the functions of another one or more of server devices 220-240.

As shown in Fig. 2, server device 220 may implement a search system 225 that receives search queries from user device 210, and that provides search results that are responsive to the search queries. In some implementations, search system 225 may include a personal search system that enables a set of users to search private content accessible only to the set of users. Server device 220 may crawl a corpus of documents, e.g., web pages, index the documents, and store information associated with the documents. Server devices 230 and 240 may store or maintain documents that may be crawled or analyzed by server device 220.

Network 250 may include any type of network, such as, for example, a local area network, also referred to as a "LAN," a wide area network, also referred to as a "WAN," a telephone network, such as the Public Switched Telephone Network (PSTN) or a cellular network, an intranet, the Internet, a combination of networks, etc. User device 210 and server devices 220-240 may connect to network 250 via wired and/or wireless connections. In other words, any one of user device 210 and server devices 220-240 may connect to network 250 via a wired connection, a wireless connection, or a combination of a wired connection and a wireless connection.

Although Fig. 2 shows example components of environment 200, in some implementations, environment 200 may include additional components, fewer components, different components, and/or differently arranged components than those depicted in Fig. 2. Additionally, or alternatively, one or more components of environment 200 may perform one or more tasks described as being performed by one or more other components of environment 200.

Fig. 3 is a flowchart of an example process 300 for indexing documents with restricted portions. In some implementations, process 300 may be performed by server device 220, such as by search system 225 of server device 220. In some implementations, process 300 may be performed by one or more other devices instead of, or possibly in conjunction with, server device 220.

Process 300 may include receiving a document from private content that includes a group of documents and is accessible by a set of users (block 310). For example, server device 220 and/or servers 230/240 may include a repository of private content. The private content may include one or more documents that may only be visible and/or searchable (i.e., accessible) by a set of users. In some implementations, the private content may include blog posts, images, videos, text documents, etc. stored by the set of users in the repository. Server device 220 (e.g., search system 225) may receive a document from the private content in a number of ways. For example, search system 225 may request the document from the repository, and the repository may provide the document to search system 225 based on the request. In some implementations, search system 225 may crawl the private content, and may receive the document based on crawling the private content.

In some implementations, one or more documents of the private content may include portion(s) that are only to be accessible to a subset (e.g., one or more particular users) of the set of users. For example, a user of the set of users may provide photos in the private content and may identify who is in each photo by tagging their faces with names. Different users of the set of users may tag a same photo as the user. The set of users may have access to the user's photos, but the user may mark the tags as private and restricted to particular users.

Process 300 may further include extracting metadata from the document (block 320). For example, server device 220, e.g., an indexer component of search system 225, may extract metadata from the document. If the document is an image, the metadata may include, for example, information that describes how large the image is, a color depth of the image, a

resolution of the image, when the image was created, etc. If the document is a text document, the metadata may include, for example, information about a length of the document, an author of the document, when the document was written, a summary of the document, etc.

In some implementations, the indexer component may extract information from text and/or hypertext markup language (HTML) code of the document, and may associate the extracted information, as keywords, with information identifying the document. In one example, the indexer component may create, from the HTML code of the document, keywords relating to visual features of the document. For example, the indexer component may create, from the HTML code of the document, a keyword corresponding to the color of the background of the document, a keyword corresponding to the color of the text in the document, a keyword corresponding to the size of the text in the document, a keyword indicating whether the document includes an image, a keyword associated with a geographic location, e.g., global satellite coordinates, associated with an image in the document, a keyword indicating a date and/or time when the image was captured, and/or a keyword indicating the quantity of images in document.

Process 300 may further include determining, based on the metadata, whether a portion of the document includes information only to be accessible by a particular user(s) of the set of users (block 330). For example, server device 220, e.g., the indexer component of search system 225, may analyze the metadata of the document to determine whether any portion of the document includes information that is only to be accessible by one or more particular users of the set of users. In one example, if the document is an image, the indexer component may analyze the metadata of the image to determine whether there are private tags (e.g., private face tags) associated with the image. In some implementations, if the document includes private location information, associated with a user of the set of users, the indexer component may analyze the document metadata to identify the private location information. In some implementations, if the document is a presentation, the indexer component may analyze the metadata of the presentation to determine whether there are private notes associated with the presentation. The private tags, the private location information, and the private notes may be considered portions of a document that include information that is only to be accessible by the particular user(s).

In some implementations, two or more different portions of a document may be marked as private (i.e., restricted) to the same user, to different users, etc. For example, a first portion of

an image may be marked as private to a first user, and a second, different portion of the image may be marked as private to a second user. In some implementations, an entire document may be restricted to a subset of users of the set of users, and one or more portions of the document may be restricted to a subset of the subset of users. For example, if the set of users includes one-hundred (100) people, the entire document may be restricted to ten (10) people from the set of users, and portion(s) of the document may be restricted to three (3) of the ten people.

As further shown in Fig. 3, if the portion of the document includes information only to be accessible by the particular user(s) of the set of users (block 330 - YES), process 300 may include appending a restrict indicator to the metadata (block 340). For example, server device 220, e.g., the indexer component of search system 225, may determine that a portion of the document includes information that is only to be accessible by a particular user of the set of users. Based on this determination, the indexer component may append a restrict indicator to the metadata associated with the document. The restrict indicator may identify the particular user, the portion of the document, and/or the information provided in the document portion. For example, if the document is an image that includes a private tag, the restrict indicator may include the following information:

private_tag=NameOfPerson_only_visible_to=NameOfImageOwner.

If the document is a text document that includes pages that are restricted to particular user(s), the restrict indicator may identify the particular user(s) and the pages of the text document that are restricted.

In some implementations, if two or more different portions of a document are restricted to particular user(s), the indexer component may append a restrict indicator for each of the two or more different portions. In some implementations, the indexer component may append a single restrict indicator for the two or more different portions of the document.

Process 300 may also include storing the metadata and the restrict indicator in a personal search index (block 350). For example, server device 220, e.g., the indexer component of search system 225, may store the metadata and the restrict indicator, associated with the document, in a personal search index. The document metadata and the restrict indicator may be searched by the set of users in the personal search index. The restrict indicator may be used to instruct the personal search index to provide search results, which include the document with the private portion, to only the particular user(s) with access to the private portion of the document. In some

implementations, the search results may provide the document to all users of the set of users but may block out the portion for the set of users, except for the particular user(s). For example, if the document is an image that includes a restricted portion, the search results may provide the image to the set of users but may blur (e.g., via image processing) the restricted portion for the set of users, except for the particular user(s).

As further shown in Fig. 3, if the portion of the document does not include information only to be accessible by the particular user(s) of the set of users (block 330 - NO), process 300 may include storing the metadata in the personal search index (block 350). For example, server device 220, e.g., the indexer component of search system 225, may determine that a portion of the document does not include information that is only to be accessible by a particular user of the set of users. Based on this determination, the indexer component may store the document metadata in the personal search index. In this situation, the indexer component would not associate a restrict indicator with the metadata.

Although Fig. 3 shows example blocks of process 300, in some implementations, process 300 may include additional blocks, different blocks, fewer blocks, and/or differently arranged blocks than those depicted in Fig. 3. Additionally, or alternatively, one or more of the blocks of process 300 may be performed in parallel.

Figs. 4A-4C are diagrams of an example 400 of the process described above with respect to Fig. 3. In example 400, assume that a first user (e.g., USER 1), of a set of users, is associated with user device 210 (not shown in Fig. 4A). Further assume that user device 210 includes an image of the first user, identified as document 410 in Fig. 4A. The first user may provide a private tag 420 in a portion of document 410, and may want document 410 and/or private tag 420 to be visible and/or searchable (i.e., accessible) by the first user and a second user (not shown in Fig. 4A). For example, the first user may provide, in private tag 420, information that indicates that private tag 420 is only to be accessible by the first user and the second user (e.g., USER 2).

The first user may instruct user device 210 to provide document 410 and private tag 420 to a repository of private content. In some implementations, the repository of the private content may be provided in server device 230 (e.g., as shown in Fig. 4A), server device 220, and/or server device 240. The private content may include one or more documents that may be visible and/or searchable (i.e., accessible) by the set of users. Server device 230, e.g., the repository, may receive document 410 and private tag 420, and may store document 410 and private tag 420,

as further shown in Fig. 4A.

In example 400 and as shown in Fig. 4B, further assume that search system 225 includes an indexer component 430 that receives documents from the private content, and extracts metadata from the documents for a personal search index (e.g., not shown in Fig. 4B). As further shown in Fig. 4B, indexer component 430 may receive document 410 and private tag 420 from the private content and may extract metadata 440 from document 410 (e.g., the image of the first user). Metadata 440 may include, for example, information associated with a size of the image, a color of the image, when the image was created, etc.; location information associated with the image; information associated with private tag 420; etc.

Indexer component 430 may analyze metadata 440 to determine whether a portion of document 410 includes information only to be accessible by particular user(s) of the set of users. In example 400 and as shown in Fig. 4C, indexer component 430 may identify the private tag information as a portion of document 410 that includes information 450 only to be accessible by the particular user(s). For example, the private tag information may indicate that private tag 420 is only to be accessible by the first user and the second user.

Based on the identification of information 450, indexer component 430 may provide metadata 440 of document 410 in a personal search index 460, as further shown in Fig. 4C. Personal search index 460 may be provided in server device 230 (e.g., as shown in Fig. 4C), server device 220, and/or server device 240. Personal search index 460 may include an index that is searchable by the set of users associated with the private content. Personal search index 460 may store information (e.g., metadata) associated with other documents, and may associate the metadata from the other documents with information identifying the other documents. Personal search index 460 may store metadata 440 associated with document 410, and may associate metadata 440 with information identifying document 410. Based on the identification of information 450, indexer component 430 may append a restrict indicator 470 to metadata 440 in personal search index 460. Restrict indicator 470 may identify the first user and the second user, private tag 420, and/or the information provided in private tag 420. For example, restrict indicator 470 may include the following information:

private_tag=USER 1 only_visible_to=USER 1 and USER 2.

If document 410 did not include private tag 420, indexer component 430 may determine that document 410 does not include information only to be accessible by particular user(s) of the

set of users. In such a scenario, indexer component 430 may provide metadata 440 of document 410 to personal search index 460, and personal search index 460 may store metadata 440 without a restrict indicator.

As indicated above, Figs. 4A-4C are provided merely as an example. Other examples are possible and may differ from what was described with regard to Figs. 4A-4C.

Fig. 5 is a flowchart of an example process 500 for searching documents with restricted portions. In some implementations, process 500 may be performed by server device 220, such as by search system 225 of server device 220. In some implementations, process 500 may be performed by one or more other devices instead of, or possibly in conjunction with, server device 220.

Process 500 may include receiving a search query from a user (block 510). For example, a user may provide a search query to user device 210, and user device 210 may provide the search query to server device 220, e.g., search system 225 of server device 220. In some implementations, search system 225 may include a search engine component that receives the search query from user device 210. In some implementations, the search query may be directed to a search for one or more documents provided in the private content and indexed in personal search index 460 (Fig. 4C). For example, the search query may be directed to a search for images associated with a particular user of the set of users.

Process 500 may also include adding one or more restrictions to the search query based on information associated with the user (block 520). For example, the search engine component of search system 225 may add one or more restrictions to the search query based on information associated with the user of user device 210. In some implementations, the search engine component may include a data structure that associates the user information (e.g., user identification) with restrictions that are associated with the user. The information associated with the user may include information identifying the user, a location of user device 210, etc. The restriction(s) may include information identifying document portion(s) that are accessible by the user. In some implementations, a single restriction may be provided for the user and may include a list of documents with portions marked as private to the user. For example, the restriction(s) may specify that a portion of a particular document is accessible by the user. In some implementations, if the particular document is an image and the portion of the particular document is a private tag associated with the user, the restriction(s) may include the following

information:

private_tag=User visible to=User.

In some implementations, the user may provide one or more restrictions in the search query, with or without the aid of the search engine component. For example, the user may provide private tags, such as *private_tag=User visible to=User*, in the search query. Such a search query may return search results associated with documents that include private tags that are restricted to the user.

Process 500 may further include identifying search results, based on the search query and the restriction(s), from a personal search index associated with the user (block 530). For example, the user may be associated with the set of users and personal search index 460 (Fig. 4C). The search engine component of search system 225 may compare the search query and the restriction(s) to information provided in personal search index 460, and may identify search results based on the comparison. In some implementations, the search engine component may identify documents by determining which terms of the search query match the document metadata provided in personal search index 460. The search engine component may determine whether any of the identified documents are associated with restrict indicators. If particular documents of the identified documents are associated with restrict indicators, the search engine component may determine whether the restrict indicators associated with the particular documents match the restriction(s). The search engine component may include the particular documents with matching restrict indicators in the search results, and may not include the particular documents without matching restrict indicators in the search results.

For example, if the search query is directed to a search for images of the user, the search engine component may search personal search index 460 for metadata associated with images of the user. The search engine component may also search personal search index 460 for restrict indicators associated with images of the user that are accessible by the user. The search engine component may identify the search results based on the searches of personal search index 460. In some implementations, the search results may include one or more images, e.g., of the user, that do not include private information and/or one or more images, e.g., of the user, that include private information.

Process 500 may additionally include scoring the search results (block 540). For example, the search engine component of search system 225 may provide the search results to a

scoring component of search system 225. The scoring component may score the search results. In some implementations, the scoring component may score the search results based on the search query and/or the restriction(s).

Process 500 may further include providing one or more scored search results (block 550). For example, the scoring component of search system 225 may provide one or more scored search results to user device 210, and user device 210 may display the one or more scored search results to the user. In some implementations, the scored search result(s) may include a ranked list of scored search results that is sorted according to the scores allotted to the search results. In some implementations, at least one scored search result may be associated with a document that includes a portion only to be accessible by particular user(s) of the set of users. In some implementations, the scored search results may include documents with private portions that are not accessible to the particular user. Such documents may be visible to the particular user, but the private portions may not be visible to the particular user. For example, if Bob is the particular user and wants to search for "photos of Matt," the scored search results may include image(s) of Matt that do not include private portions, image(s) of Matt with the private portions that are accessible by Bob, and/or image(s) of Matt with the private portions blocked since that are not accessible by Bob. Alternatively, the scored search results may not include image(s) of Matt that have private portions not accessible to Bob.

In some implementations, search system 225 may pre-filter the search results that are returned to the user based on the restriction(s) added to the search query and/or the restrict indicators provided in personal search index 460. For example, search system 225 may not return, to the user, search results associated with documents that include portions not accessible to the user. In some implementations, search system 225 may return search results associated with documents that include portions not accessible to the user, and may provide such documents to the user. However, search system 225 may block (i.e., remove) the portions not accessible to the user from such documents.

In some implementations, search system 225 may post-filter the search results that are returned to the user based on the restriction(s) added to the search query and/or the restrict indicators provided in personal search index 460. For example, the search engine of search system 225 may return all search results associated with the search query, and may remove search results associated with documents that include portions not accessible to the user (e.g.,

based on the restriction(s)).

Process 500 may also include receiving a selection of a particular scored search result associated with a document with a restricted portion (block 560). For example, the user may utilize user device 210 to select a particular scored search result from the displayed one or more scored search results. In some implementations, the particular scored search result may be associated with a document that includes a portion only to be accessible by particular user(s) of the set of users. For example, the particular scored search result may include information associated with an image with a private tag that is only to be accessible to the user of user device 210. User device 210 may provide the selection to server device 220, e.g., search system 225 of server device 220, and search system 225 may receive the selection.

Process 500 may include providing the document with the restricted portion when the user is permitted access to the restricted portion (block 570). For example, based on the selection of the particular scored search result, search system 225 of server device 220 may retrieve, from the private content, the document that includes the portion only to be accessible by the user of user device 210. Search system 225 may provide the document to user device 210, and user device 210 may display the document to the user. Since the user has access to the restricted portion of the document, the user may view the restricted portion of the document.

Although Fig. 5 shows example blocks of process 500, in some implementations, process 500 may include additional blocks, different blocks, fewer blocks, and/or differently arranged blocks than those depicted in Fig. 5. Additionally, or alternatively, one or more of the blocks of process 500 may be performed in parallel.

Figs. 6A-6D are diagrams of an example 600 of the process described above with respect to Fig. 5. In example 600, assume that a second user (e.g., USER 2) is associated with user device 210, and that user device 210 displays a user interface 610 to the second user, as shown in Fig. 6A. User interface 610 may be provided by the search engine component of search system 225, and may include information that enables the second user to perform a search for documents provided in the private content. As further shown in Fig. 6A, the second user may enter a search query 620 via user interface 610. In example 600, search query 620 may be directed to photos or images of a first user (e.g., USER 1). User device 210 may receive search query 620, and may provide search query 620 to a search engine component 630 of search system 225.

Search engine component 630 may receive search query 620, and may add restrictions

640 to search query 620 based on information associated with the second user, as shown in Fig. 6B. For example, search engine component 630 may add restrictions 640 to search query 620 based on the identification of the second user. In example 600, assume that restrictions 640 specify that portions of documents, related to search query 620, are accessible by the second user, based on the following:

find all photos that have private_tag=USER 1 visible to=USER 2.

Search engine component 630 may compare search query 620 to information contained in personal search index 460, as further shown in Fig. 6B, to identify documents that match search query 620. Search engine component 630 may determine whether any of the identified documents are associated with the restrict indicators. If particular documents of the identified documents are associated with the restrict indicators, search engine component 630 may determine whether the restrict indicators associated with the particular documents match restrictions 640. Search engine component 630 may include the particular documents with matching restrict indicators in search results 650, and may not include the particular documents without matching restrict indicators in search results 650.

In example 600 and as shown in Fig. 6B, assume that search results 650 include metadata of a particular image of the first user (e.g., USER 1), such as document 410 (Fig. 4A), restrict indicator 470 (Fig. 4C) associated with document 410, and metadata for other images of the first user. Further assume that the particular image includes a portion that is only to be accessible to the first user and the second user, and that the particular image may or may not be accessible to other users. For example, the particular image may include private tag 420 provided by the first user in document 410, as described above in connection with Fig. 4A.

Search engine component 630 may provide search results 650 to a scoring component 660 of search system 225, as shown in Fig. 6C. Scoring component 660 may score search results 650 to generate scored search results 670. In example 600, assume that scoring component 660 scores the metadata of document 410 higher than the metadata for other images of the first user. Scoring component 660 may provide scored search results 670 to user device 210, and user device 210 may display scored search results 670 in a user interface 680, as further shown in Fig. 6C. In example 600, user interface 680 may display the metadata of document 410 as a "Restricted photo of USER 1." User interface 680 may also display the metadata for other images of the first user as "Unrestricted photos of USER 1" (i.e., as not including restricted

portions).

The second user may utilize user device 210 to select one of scored search results 670 from user interface 680. For example, the second user may select the "Restricted photo of USER 1" from user interface 680, and user device 210 may receive the selection. User device 210 may provide the selection to search system 225, and search system 225 may retrieve document 410 from the private content based on the selection. Search system 225 may provide document 410 to user device 210, and user device 210 may display document 410 in a user interface 690, as shown in Fig. 6D. The second user may view document 410 (e.g., the image of the first user) and information provided in private tag 420.

In some implementations, if a third user submits search query 620 and does not have access to private tag 420 in document 410, search system 225 may not return the metadata of document 410 in search results 650 to the third user. Alternatively, search system 225 may return the metadata of document 410 in search results 650 to the third user. However, if the third user selects the metadata of document 410, search system 225 may provide document 410 to the third user but may not provide private tag 420 of document 410 to the third user.

As indicated above, Figs. 6A-6D are provided merely as an example. Other examples are possible and may differ from what was described with regard to Figs. 6A-6D.

Fig. 7 is a diagram of an example of a generic computing device 700 and a generic mobile computing device 750, which may be used with the techniques described here. Generic computing device 700 or generic mobile computing device 750 may correspond to, for example, user device 210 and/or a server device 220, 230, or 240. Computing device 700 is intended to represent various forms of digital computers, such as laptops, desktops, workstations, personal digital assistants, servers, blade servers, mainframes, and other appropriate computers. Mobile computing device 750 is intended to represent various forms of mobile devices, such as personal digital assistants, cellular telephones, smart phones, tablet computers, and other similar computing devices. The components shown in Fig. 7, their connections and relationships, and their functions, are meant to be examples only, and are not meant to limit implementations described herein.

Computing device 700 may include a processor 702, a memory 704, a storage device 706, a high-speed interface 708 connecting to memory 704 and high-speed expansion ports 710, and a low-speed interface 712 connecting to a low-speed expansion port 714 and a storage

device 706. Each of components 702, 704, 706, 708, 710, 712, and 714, are interconnected using various buses, and may be mounted on a common motherboard or in other manners as appropriate. Processor 702 can process instructions for execution within computing device 700, including instructions stored in memory 704 or on storage device 706 to display graphical information for a graphical user interface, also referred to as a "GUI," on an external input/output device, such as display 716 coupled to high-speed interface 708. In some implementations, multiple processors and/or multiple buses may be used, as appropriate, along with multiple memories and types of memory. Also, multiple computing devices 700 may be connected, with each device providing portions of the necessary operations, as a server bank, a group of blade servers, or a multi-processor system, etc.

Memory 704 stores information within computing device 700. In some implementations, memory 704 includes a volatile memory unit or units. In some implementations, memory 704 may include a non-volatile memory unit or units. Memory 704 may also be another form of computer-readable medium, such as a magnetic or optical disk. A computer-readable medium may refer to a non-transitory memory device. A memory device may refer to storage space within a single storage device or spread across multiple storage devices.

Storage device 706 is capable of providing mass storage for computing device 700. In some implementations, storage device 706 may be or contain a computer-readable medium, such as a floppy disk device, a hard disk device, an optical disk device, or a tape device, a flash memory or other similar solid state memory device, or an array of devices, including devices in a storage area network or other configurations. A computer program product can be tangibly embodied in an information carrier. The computer program product may also contain instructions that, when executed, perform one or more methods, such as those described herein. The information carrier is a computer or machine-readable medium, such as memory 704, storage device 706, or a memory on processor 702.

High-speed interface 708 manages bandwidth-intensive operations for computing device 700, while low-speed interface 712 manages lower bandwidth-intensive operations. Such allocation of functions is provided for explanatory purposes only. In some implementations, high-speed interface 708 may be coupled to memory 704, display 716, e.g., through a graphics processor or accelerator, and to high-speed expansion ports 710, which may accept various expansion cards. In some implementations, low-speed interface 712 may be coupled to storage

device 706 and low-speed expansion port 714. Low-speed expansion port 714, which may include various communication ports, e.g., USB, Bluetooth, Ethernet, wireless Ethernet, etc., may be coupled to one or more input/output devices, such as a keyboard, a pointing device, a scanner, or a networking device such as a switch or router, e.g., through a network adapter.

Computing device 700 may be implemented in a number of different forms, as shown in the figure. For example, computing device 700 may be implemented as a standard server 720, or multiple times in a group of such servers. Computing device 700 may also be implemented as part of a rack server system 724. In addition, computing device 700 may be implemented in a personal computer, such as a laptop computer 722. Alternatively, components from computing device 700 may be combined with other components in a mobile device, such as mobile computing device 750. Each of such devices may contain one or more of computing devices 700, 750, and an entire system may be made up of multiple computing devices 700, 750 communicating with each other.

Mobile computing device 750 may include a processor 752, a memory 764, an input/output device, such as a display 754, a communication interface 766, and a transceiver 768, among other components. Mobile computing device 750 may also be provided with a storage device, such as a micro-drive or other device, to provide additional storage. Each of the components 752, 764, 754, 766, and 768 are interconnected using various buses, and several of the components may be mounted on a common motherboard or in other manners as appropriate.

Processor 752 can execute instructions within mobile computing device 750, including instructions stored in memory 764. Processor 752 may be implemented as a chipset of chips that include separate and multiple analog and digital processors. Processor 752 may provide, for example, for coordination of the other components of mobile computing device 750, such as control of user interfaces, applications run by mobile computing device 750, and wireless communication by mobile computing device 750.

Processor 752 may communicate with a user through control interface 758 and display interface 756 coupled to a display 754. Display 754 may be, for example, a Thin-Film-Transistor Liquid Crystal Display or an Organic Light Emitting Diode display, or other appropriate display technology. Display interface 756 may include appropriate circuitry for driving display 754 to present graphical and other information to a user. Control interface 758 may receive commands from a user and convert the commands for submission to processor 752.

In addition, an external interface 762 may be provided in communication with processor 752, so as to enable near area communication of mobile computing device 750 with other devices. External interface 362 may provide, for example, for wired communication in some implementations, or for wireless communication in some implementations, and multiple interfaces may also be used.

Memory 764 stores information within mobile computing device 750. Memory 764 can be implemented as one or more of a computer-readable medium or media, a volatile memory unit or units, or a non-volatile memory unit or units. Expansion memory 774 may also be provided and connected to mobile computing device 750 through expansion interface 772, which may include, for example, a Single In Line Memory Module, also referred to as "SIMM," card interface. Such expansion memory 774 may provide extra storage space for device 750, or may also store applications or other information for mobile computing device 750. Specifically, expansion memory 774 may include instructions to carry out or supplement the processes described herein, and may include secure information also. Thus, for example, expansion memory 774 may be provided as a security module for mobile computing device 750, and may be programmed with instructions that permit secure use of mobile computing device 750. In addition, secure applications may be provided via the SIMM cards, along with additional information, such as placing identifying information on the SIMM card in a non-hackable manner.

Expansion memory 774 may include, for example, flash memory and/or NVRAM memory. In some implementations, a computer program product is tangibly embodied in an information carrier. The computer program product contains instructions that, when executed, perform one or more methods, such as those described above. The information carrier is a computer-or machine-readable medium, such as memory 764, expansion memory 774, or a memory on processor 752, that may be received, for example, over transceiver 768 or external interface 762.

Mobile computing device 750 may communicate wirelessly through communication interface 766, which may include digital signal processing circuitry where necessary. Communication interface 766 may provide for communications under various modes or protocols, such as GSM voice calls, SMS, EMS, or MMS messaging, CDMA, TDMA, PDC, WCDMA, CDMA2000, or GPRS, among others. Such communication may occur, for example,

through transceiver 768. In addition, short-range communication may occur, such as using a Bluetooth, WiFi, or other such transceiver. In addition, a Global Positioning System, also referred to as "GPS," receiver module 770 may provide additional navigation- and location-related wireless data to mobile computing device 750, which may be used as appropriate by applications running on mobile computing device 750.

Mobile computing device 750 may also communicate audibly using audio codec 760, which may receive spoken information from a user and convert the received spoken information to digital information. Audio codec 760 may likewise generate audible sound for a user, such as through a speaker, e.g., in a handset of mobile computing device 750. Such sound may include sound from voice telephone calls, may include recorded sound, e.g., voice messages, music files, etc., and may also include sound generated by applications operating on mobile computing device 750.

Mobile computing device 750 may be implemented in a number of different forms, as shown in the figure. For example, mobile computing device 750 may be implemented as a cellular telephone 780. Mobile computing device 750 may also be implemented as part of a smart phone 782, personal digital assistant, or other similar mobile device.

Various implementations of the systems and techniques described herein can be realized in digital electronic circuitry, integrated circuitry, specially designed application specific integrated circuits, computer hardware, firmware, software, and/or combinations thereof. These various implementations can include implementations in one or more computer programs that are executable and/or interpretable on a programmable system including at least one programmable processor, which may be special or general purpose, coupled to receive data and instructions from, and to transmit data and instructions to, a storage system, at least one input device, and at least one output device.

These computer programs, also known as programs, software, software applications, or code, include machine instructions for a programmable processor, and can be implemented in a high-level procedural and/or object-oriented programming language, and/or in assembly/machine language. As used herein, the terms "machine-readable medium" and "computer-readable medium" refer to any apparatus and/or device, e.g., magnetic discs, optical disks, memory, Programmable Logic Devices, used to provide machine instructions and/or data to a programmable processor, including a machine-readable medium that receives machine

instructions as a machine-readable signal. The term "machine-readable signal" refers to any signal used to provide machine instructions and/or data to a programmable processor.

To provide for interaction with a user, the systems and techniques described herein can be implemented on a computer having a display device, e.g., a cathode ray tube or liquid crystal display monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well. For example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback. Also, input from the user can be received in any form, including acoustic, speech, or tactile input.

The systems and techniques described herein can be implemented in a computing system that includes a back end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with some implementations of the systems and techniques described here, or any combination of such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a LAN, a WAN, and the Internet.

The foregoing description provides illustration and description, but is not intended to be exhaustive or to limit the implementations to the precise form disclosed. Modifications and variations are possible in light of the above teachings or may be acquired from practice of the implementations.

As used herein, the term component is intended to be broadly interpreted to refer to hardware or a combination of hardware and software, such as software executed by a processor.

It will be apparent that systems and methods, as described above, may be implemented in many different forms of software, firmware, and hardware in the implementations illustrated in the figures. The actual software code or specialized control hardware used to implement these systems and methods is not limiting of the implementations. Thus, the operation and behavior of the systems and methods were described without reference to the specific software code – it being understood that software and control hardware can be designed to implement the systems and methods based on the description herein.

Even though particular combinations of features are recited in the claims and/or disclosed in the specification, these combinations are not intended to limit the disclosure of the possible implementations. In fact, many of these features may be combined in ways not specifically recited in the claims and/or disclosed in the specification. Although each dependent claim listed below may directly depend on only one other claim, the disclosure of the possible implementations includes each dependent claim in combination with every other claim in the claim set.

No element, act, or instruction used in the present application should be construed as critical or essential unless explicitly described as such. Also, as used herein, the articles "a" and "an" are intended to include one or more items. Where only one item is intended, the term "one" or similar language is used. Further, the phrase "based on" is intended to mean "based, at least in part, on" unless explicitly stated otherwise.

WHAT IS CLAIMED IS:

1. A method comprising:

receiving, by one or more computer devices, a document from private content,

the private content being accessible by a set of users;

extracting, by the one or more computer devices, metadata from the document,

the metadata providing information relating to one or more aspects of the document;

determining, by one or more computer devices and based on the metadata, whether a portion of the document includes information that is only to be accessible by a particular user from the set of users;

appending, by the one or more computer devices, a restrict indicator to the metadata when the portion of the document includes information that is only to be accessible by the particular user,

the restrict indicator identifying the particular user and the portion of the document; and

storing, by the one or more computer devices, the metadata and the restrict indicator in a personal search index.

2. The method of claim 1, further comprising:

storing only metadata in the personal search index for a second document that is accessible to the set of users.

3. The method of claim 1, further comprising:

receiving a search query from the particular user;

adding a restriction to the search query based on information associated with the particular user; and

identifying search results, based on the search query and the restriction, from the personal search index,

the search results including information identifying the document.

4. The method of claim 3, further comprising:

comparing the restriction and the restrict indicator; and

identifying, from the personal search index, the document in the search results based on the comparing.

5. The method of claim 3, further comprising:

scoring the search results; and

providing the scored search results to the particular user.

6. The method of claim 5, further comprising:

receiving a selection of a particular scored search result of the scored search results,

the particular scored search result being associated with the document; and

providing the document and the portion of the document to the particular user.

7. The method of claim 5, further comprising:

receiving a selection of a particular scored search result of the scored search results,

the particular scored search result being associated with a second document that includes a restricted portion that is not accessible by the particular user; and

providing the second document, without the restricted portion, to the particular user.

8. A device comprising:

one or more processors to:

receive a document from private content,

the private content being accessible by a set of users,

extract metadata from the document,

the metadata providing information relating to one or more aspects of the document,

determine, based on the metadata, whether a portion of the document includes information that is only to be accessible by a particular user from the set of users,

append a restrict indicator to the metadata when the portion of the document includes information that is only to be accessible by the particular user,

the restrict indicator identifying the particular user and the portion of the document, and

store the metadata and the restrict indicator in a personal search index.

9. The device of claim 8, where the one or more processors are further to:

store only metadata in the personal search index for a second document that is accessible to the set of users.

10. The device of claim 8, where the one or more processors are further to:

receive a search query from the particular user,

add a restriction to the search query based on information associated with the particular user, and

identify search results, based on the search query and the restriction, from the personal search index,

the search results including information identifying the document.

11. The device of claim 10, where the one or more processors are further to:

compare the restriction and the restrict indicator, and

identify, from the personal search index, the document in the search results based on the comparing.

12. The device of claim 10, where the one or more processors are further to:

score the search results, and

provide the scored search results to the particular user.

13. The device of claim 12, where the one or more processors are further to:

receive a selection of a particular scored search result of the scored search results,

the particular scored search result being associated with the document, and

provide the document and the portion of the document to the particular user.

14. The device of claim 12, where the one or more processors are further to:

receive a selection of a particular scored search result of the scored search results,

the particular scored search result being associated with a second document that includes a restricted portion that is not accessible by the particular user, and

provide the second document, without the restricted portion, to the particular user.

15. A non-transitory computer-readable medium for storing instructions, the instructions comprising:

one or more instructions that, when executed by a processor of a device, cause the processor to:

receive a document from private content,

the private content being accessible by a set of users,

extract metadata from the document,

the metadata providing information relating to one or more aspects of the document,

determine, based on the metadata, whether a portion of the document includes information that is only to be accessible by a particular user from the set of users,

append a restrict indicator to the metadata when the portion of the document includes information that is only to be accessible by the particular user,

the restrict indicator identifying the particular user and the portion of the document, and

store the metadata and the restrict indicator in a personal search index.

16. The computer-readable medium of claim 15, where the instructions further comprise:

one or more instructions that, when executed by the processor, cause the processor to:

store only metadata in the personal search index for a second document that is accessible to the set of users.

17. The computer-readable medium of claim 15, where the instructions further comprise:

one or more instructions that, when executed by the processor, cause the processor to:

receive a search query from the particular user,

add a restriction to the search query based on information associated with the particular user, and

identify search results, based on the search query and the restriction, from the personal search index,

the search results including information identifying the document.

18. The computer-readable medium of claim 17, where the instructions further comprise:

one or more instructions that, when executed by the processor, cause the processor to:

compare the restriction and the restrict indicator, and

identify, from the personal search index, the document in the search results based on the comparing.

19. The computer-readable medium of claim 17, where the instructions further comprise:

one or more instructions that, when executed by the processor, cause the processor to:

score the search results, and

provide the scored search results to the particular user.

20. The computer-readable medium of claim 19, where the instructions further comprise:

one or more instructions that, when executed by the processor, cause the processor to:

receive a selection of a particular scored search result of the scored search results,

the particular scored search result being associated with the document, and

provide the document and the portion of the document to the particular user.

FIG. 1A

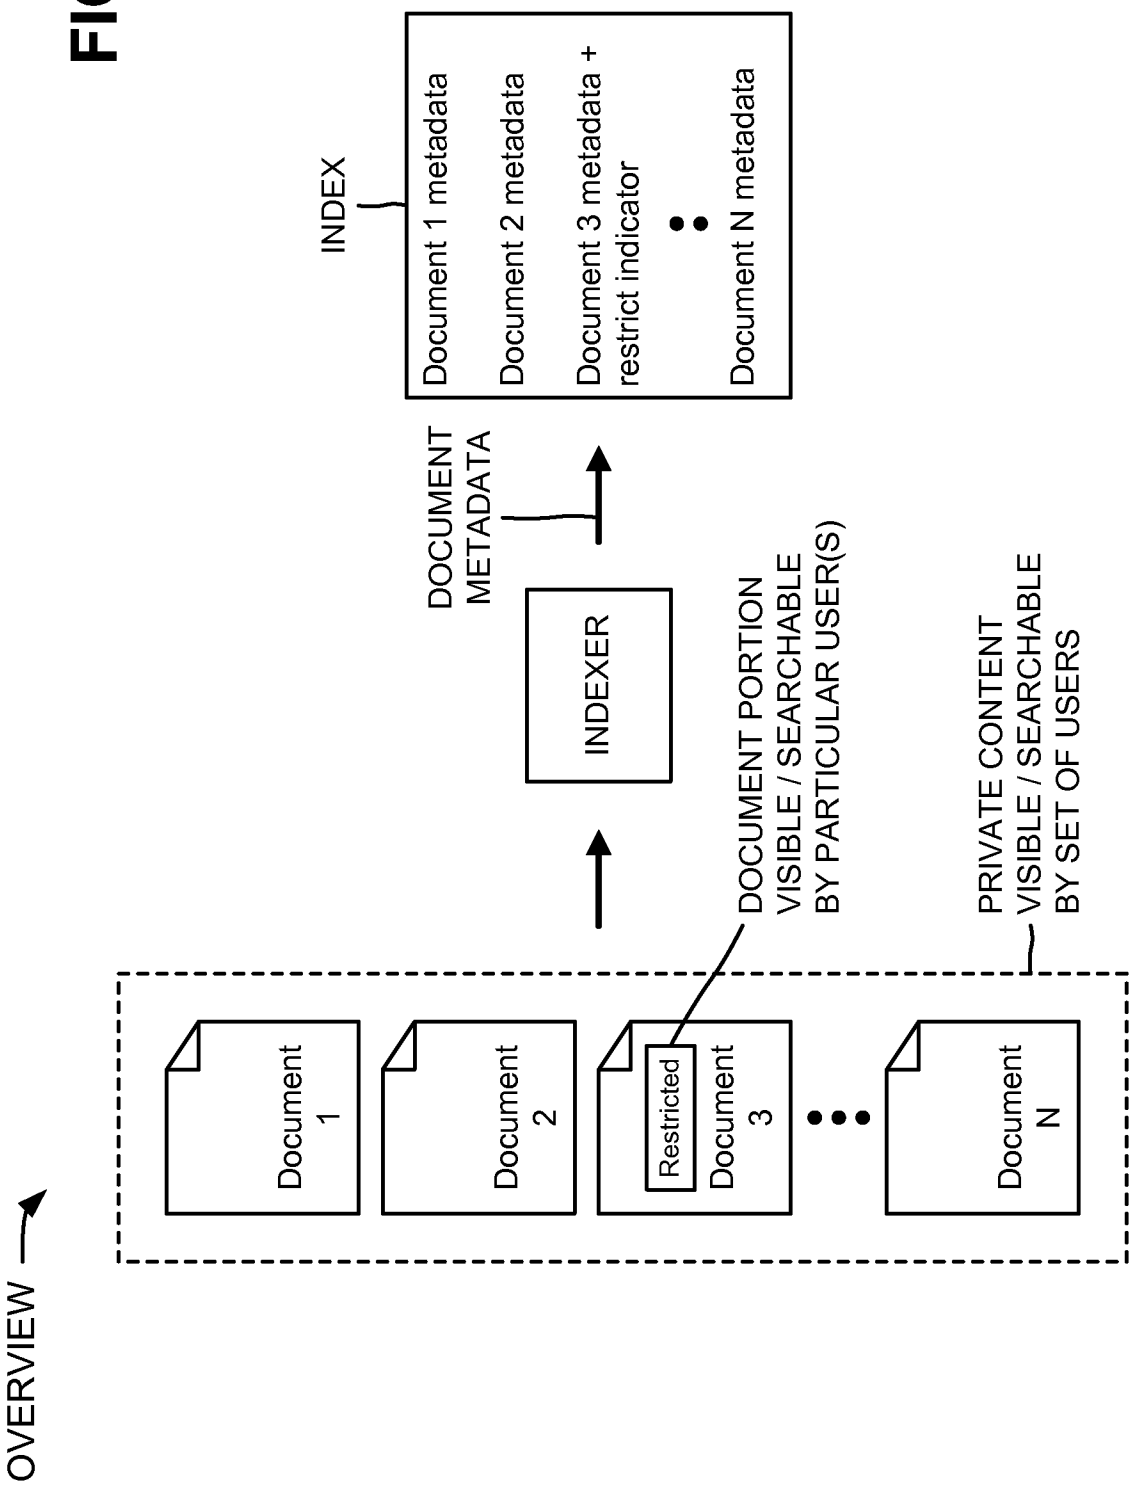

OVERVIEW
Document 1
Document 2
Restricted
Document 3
Document N
DOCUMENT PORTION VISIBLE / SEARCHABLE BY PARTICULAR USER(S)
PRIVATE CONTENT VISIBLE / SEARCHABLE BY SET OF USERS
INDEXER
DOCUMENT METADATA
INDEX
Document 1 metadata
Document 2 metadata
Document 3 metadata + restrict indicator
Document N metadata

FIG. 1B

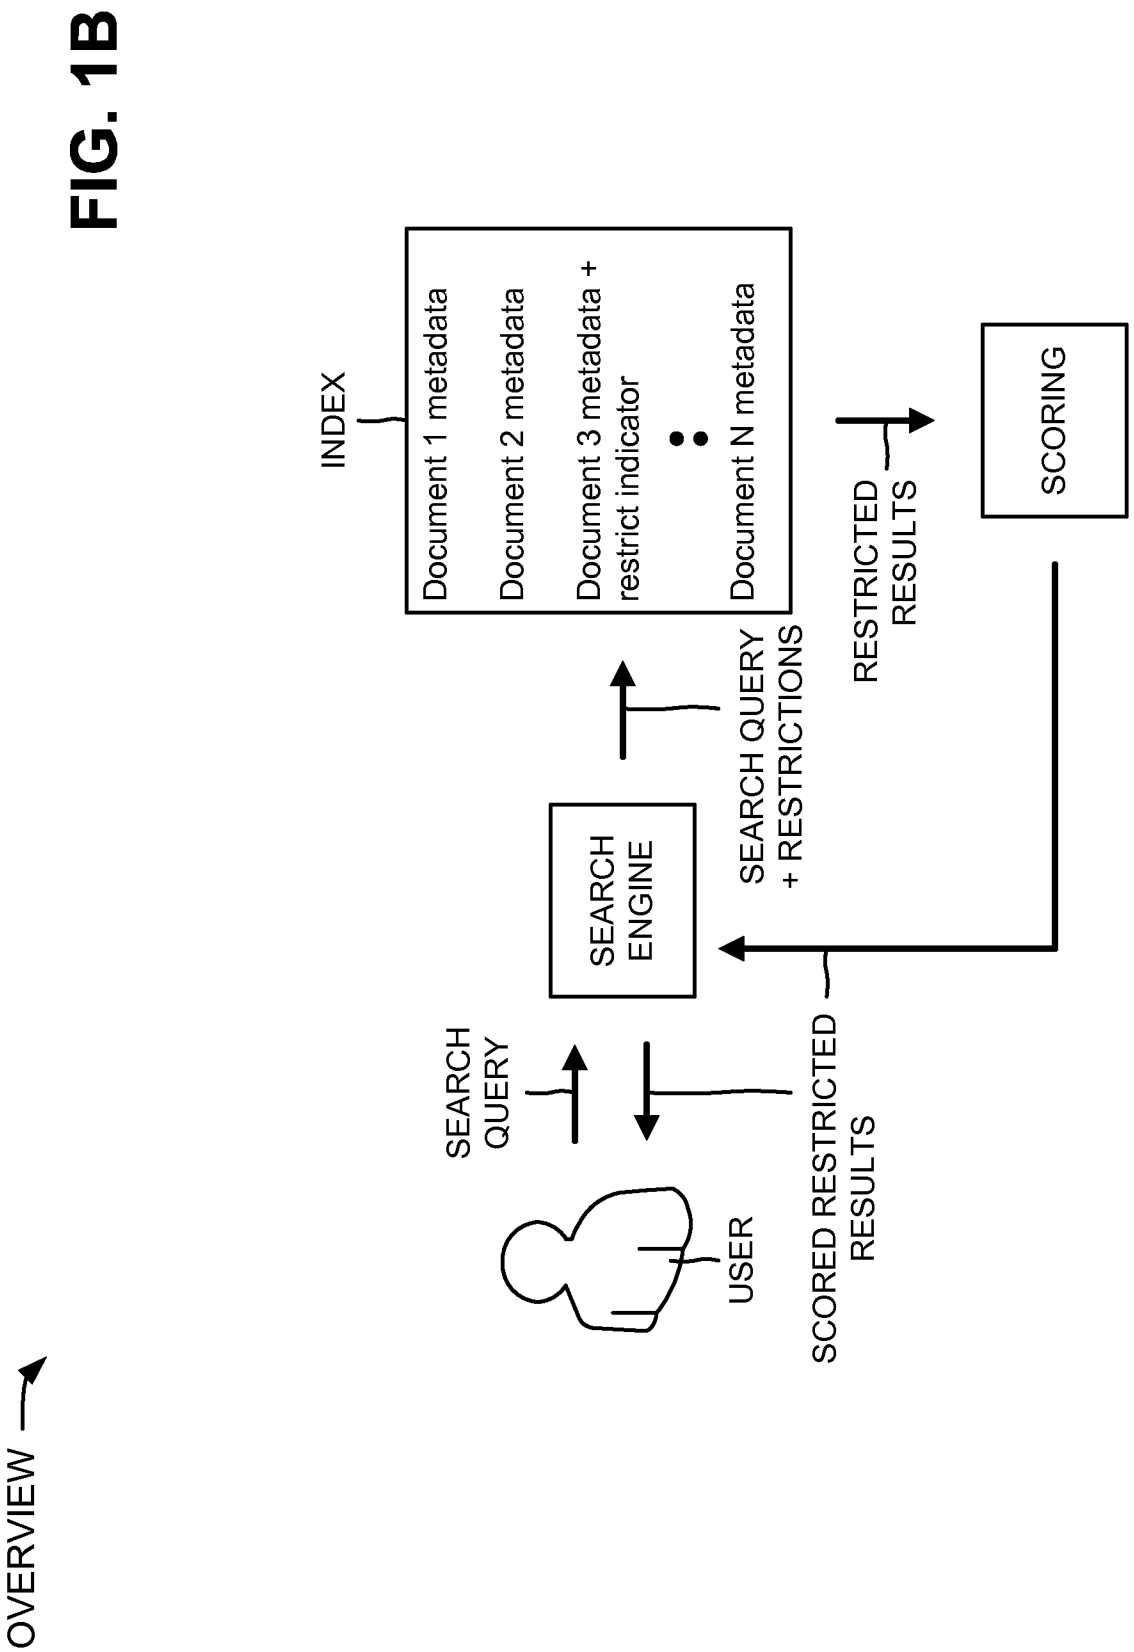

OVERVIEW
USER
SEARCH QUERY
SEARCH ENGINE
INDEX
Document 1 metadata
Document 2 metadata
Document 3 metadata + restrict indicator
Document N metadata
SEARCH QUERY + RESTRICTIONS
RESTRICTED RESULTS
SCORING
SCORED RESTRICTED RESULTS

FIG. 2

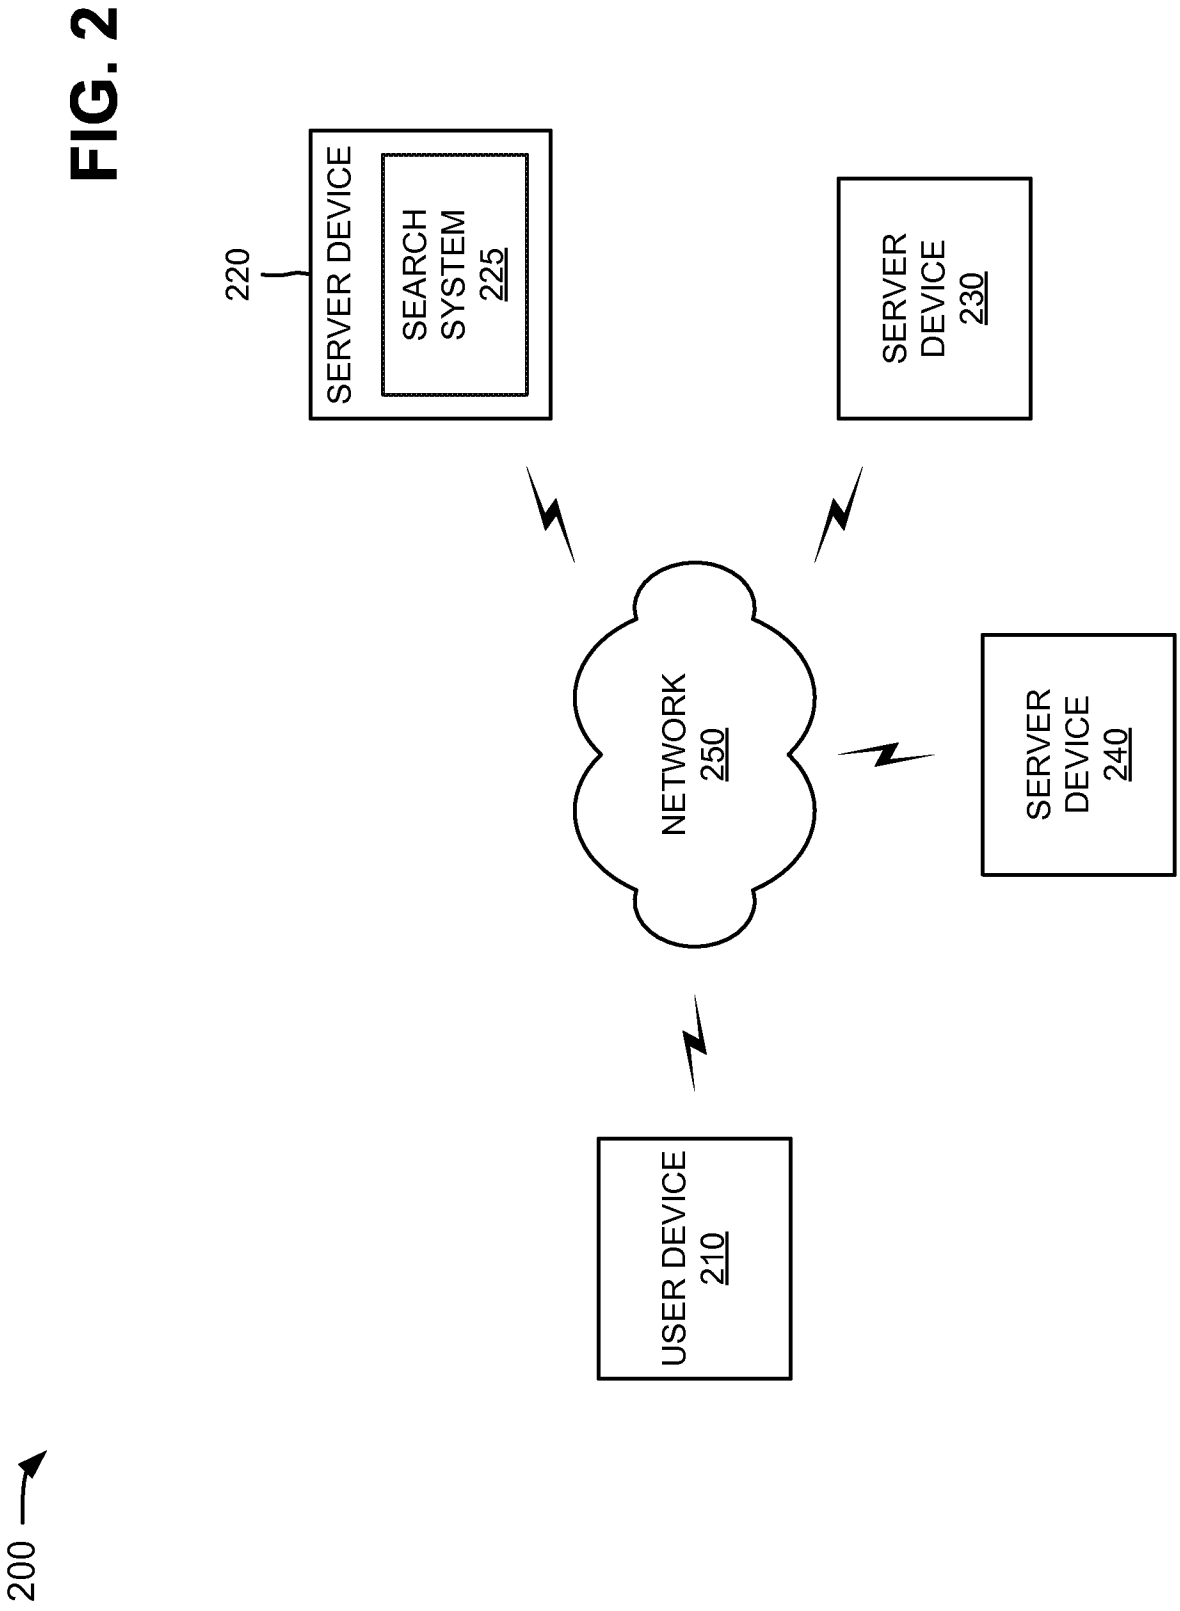

200
220
SERVER DEVICE
SEARCH SYSTEM 225
SERVER DEVICE 230
NETWORK 250
SERVER DEVICE 240
USER DEVICE 210

FIG. 3

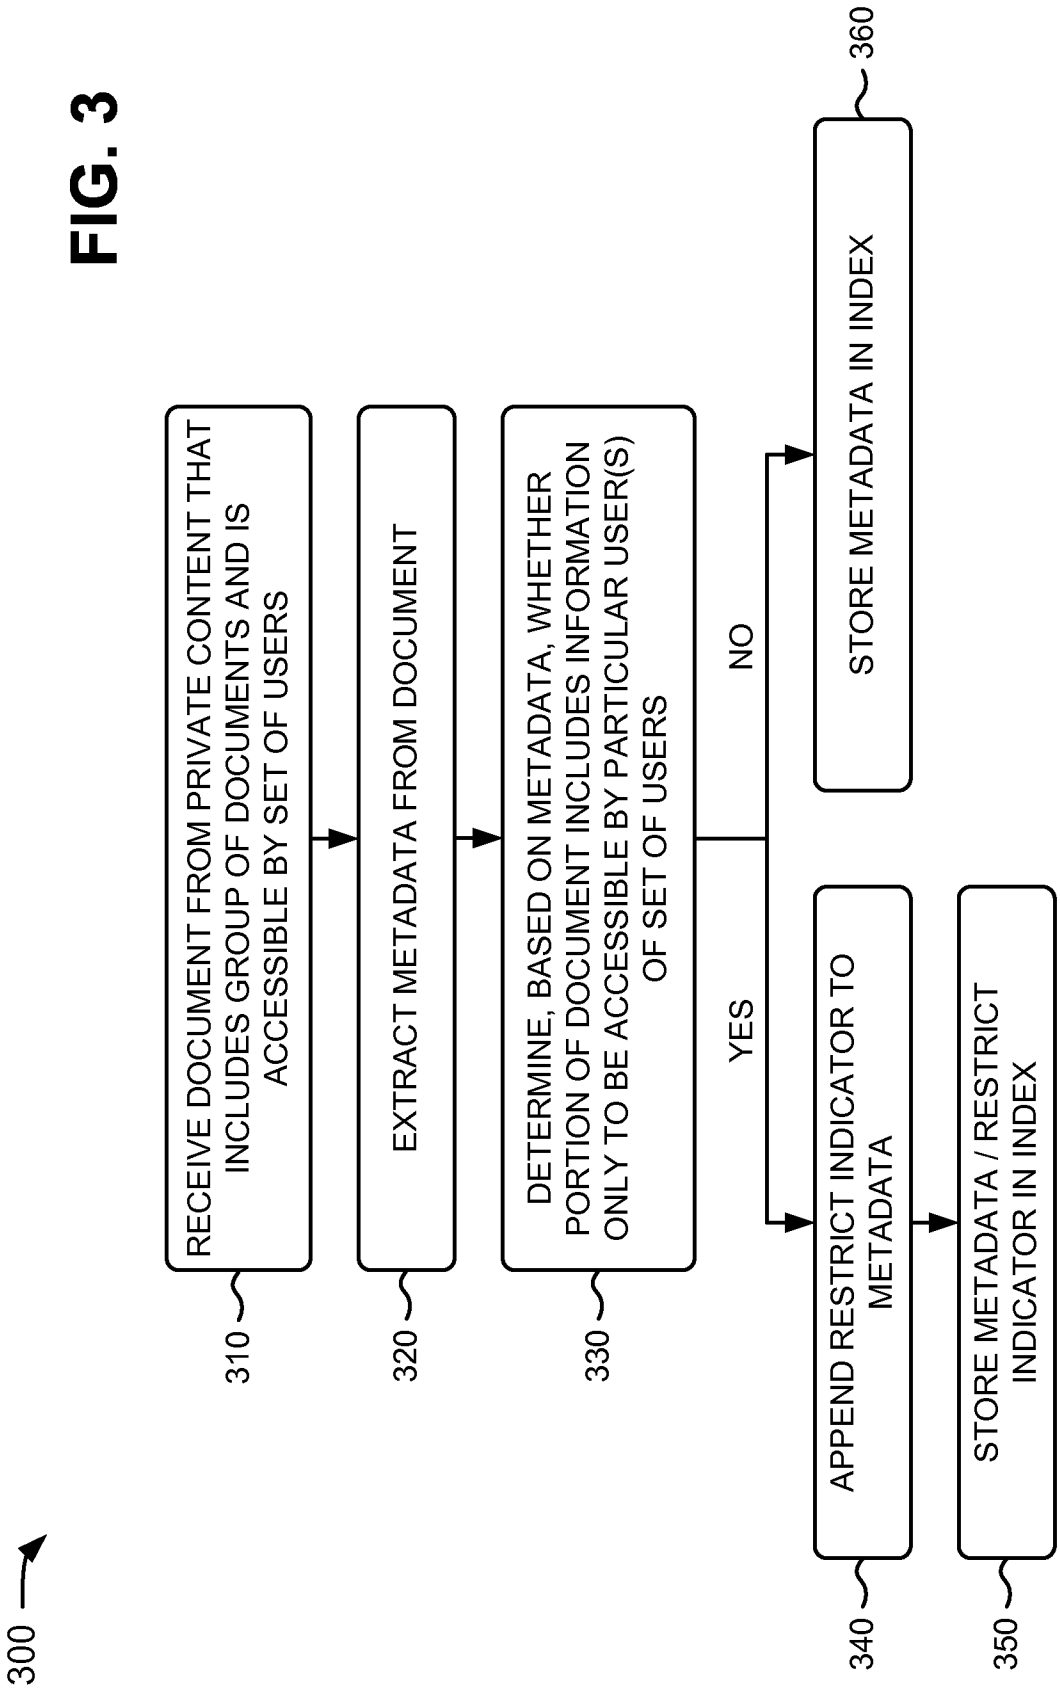

300
RECEIVE DOCUMENT FROM PRIVATE CONTENT THAT INCLUDES GROUP OF DOCUMENTS AND IS ACCESSIBLE BY SET OF USERS
310
EXTRACT METADATA FROM DOCUMENT
320
DETERMINE, BASED ON METADATA, WHETHER PORTION OF DOCUMENT INCLUDES INFORMATION ONLY TO BE ACCESSIBLE BY PARTICULAR USER(S) OF SET OF USERS
330
YES
APPEND RESTRICT INDICATOR TO METADATA
340
STORE METADATA / RESTRICT INDICATOR IN INDEX
350
NO
STORE METADATA IN INDEX
360

FIG. 4A

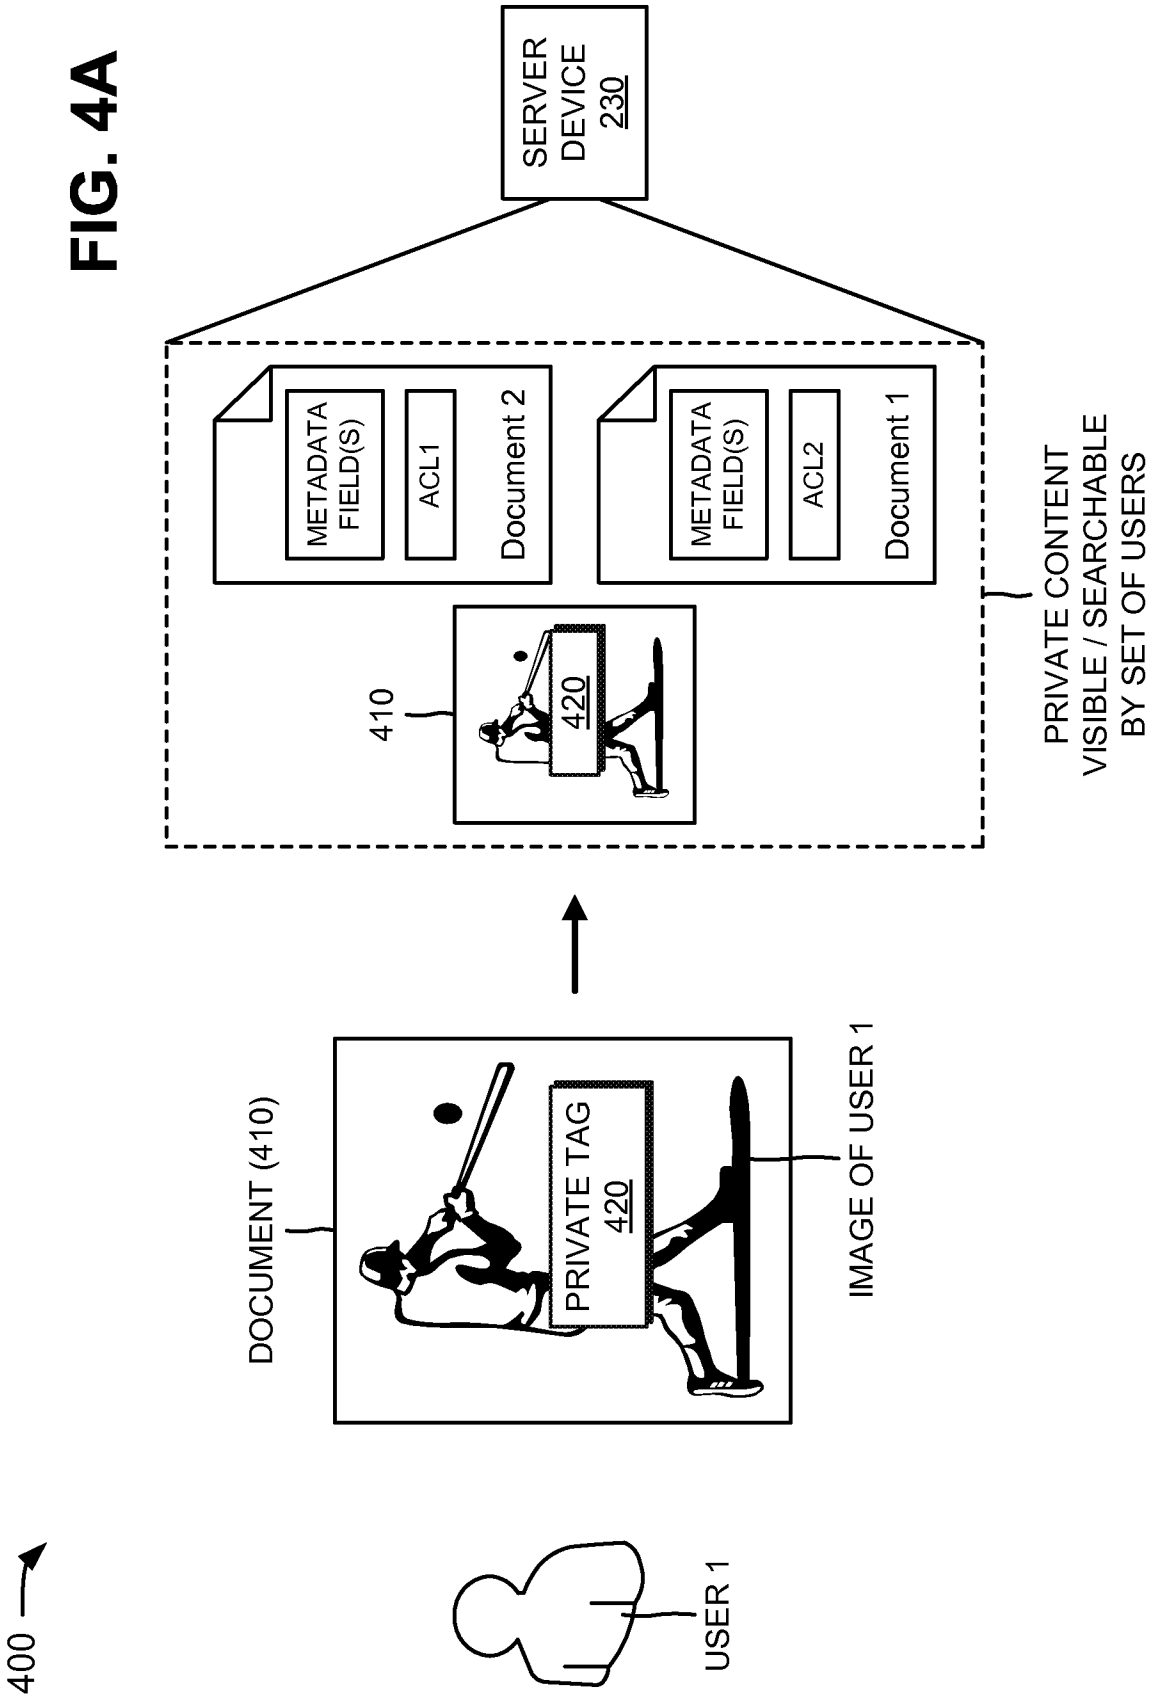

400
USER 1
DOCUMENT (410)
PRIVATE TAG
420
IMAGE OF USER 1
SERVER DEVICE 230
METADATA FIELD(S)
ACL1
Document 2
METADATA FIELD(S)
ACL2
Document 1
410
420
PRIVATE CONTENT VISIBLE / SEARCHABLE BY SET OF USERS

FIG. 4B

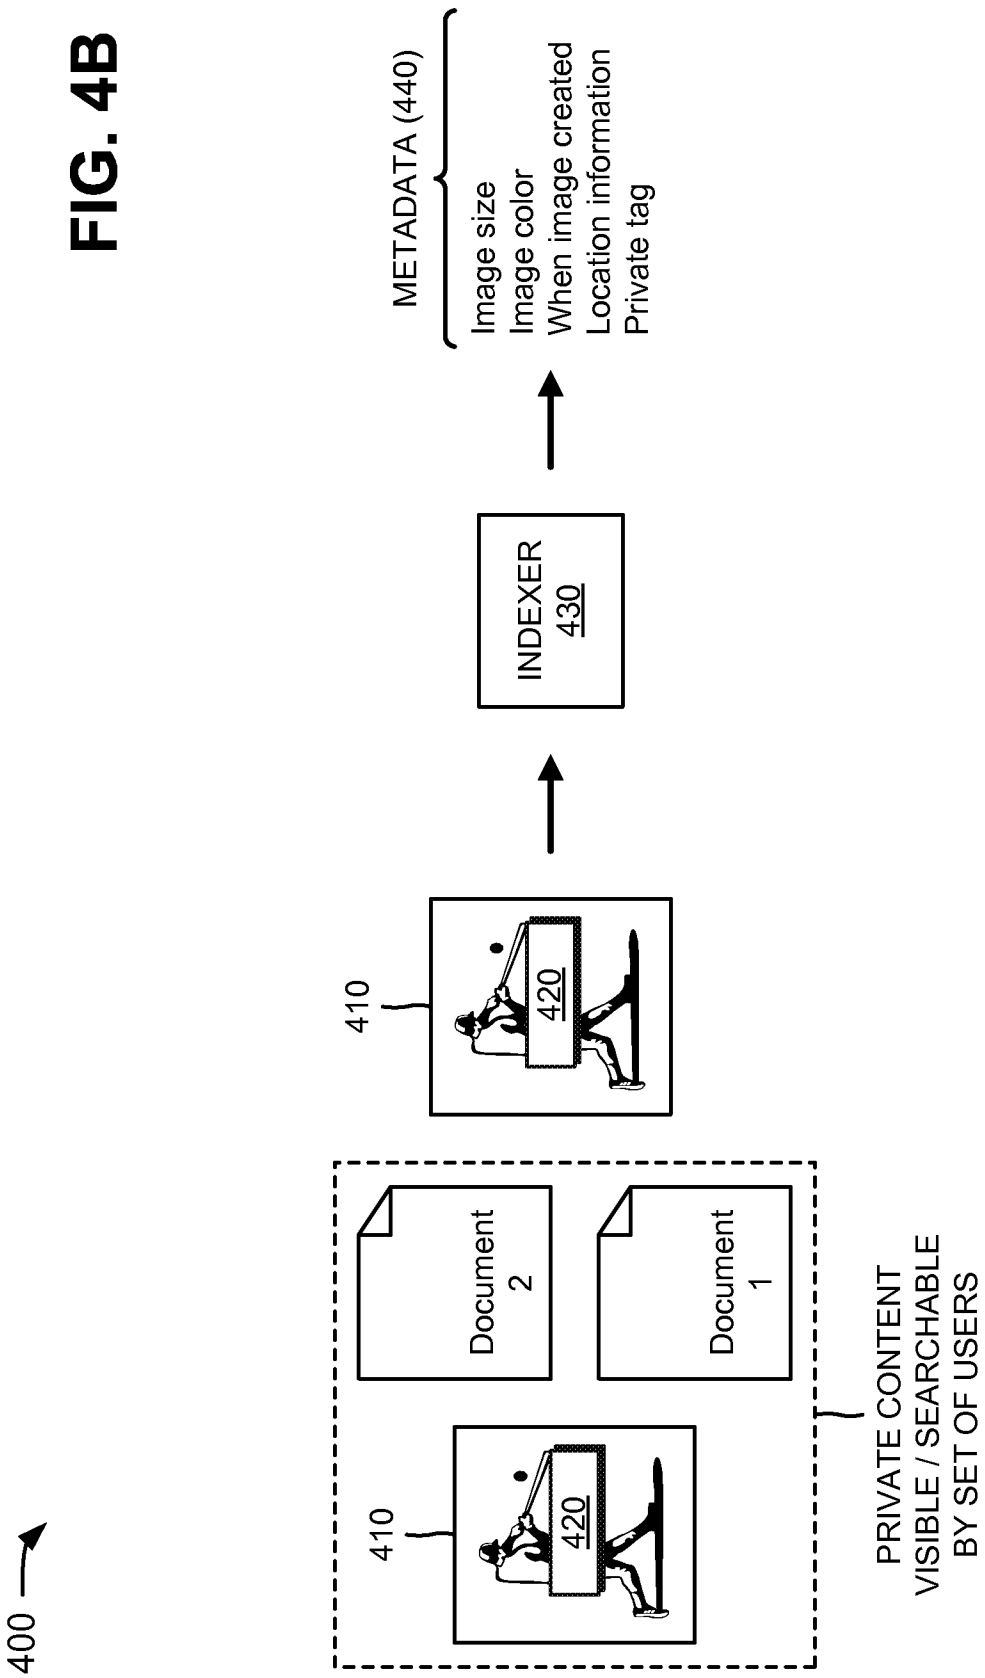

400
410
420
Document 2
Document 1
PRIVATE CONTENT
VISIBLE / SEARCHABLE
BY SET OF USERS
INDEXER
430
METADATA (440)
Image size
Image color
When image created
Location information
Private tag

FIG. 4C

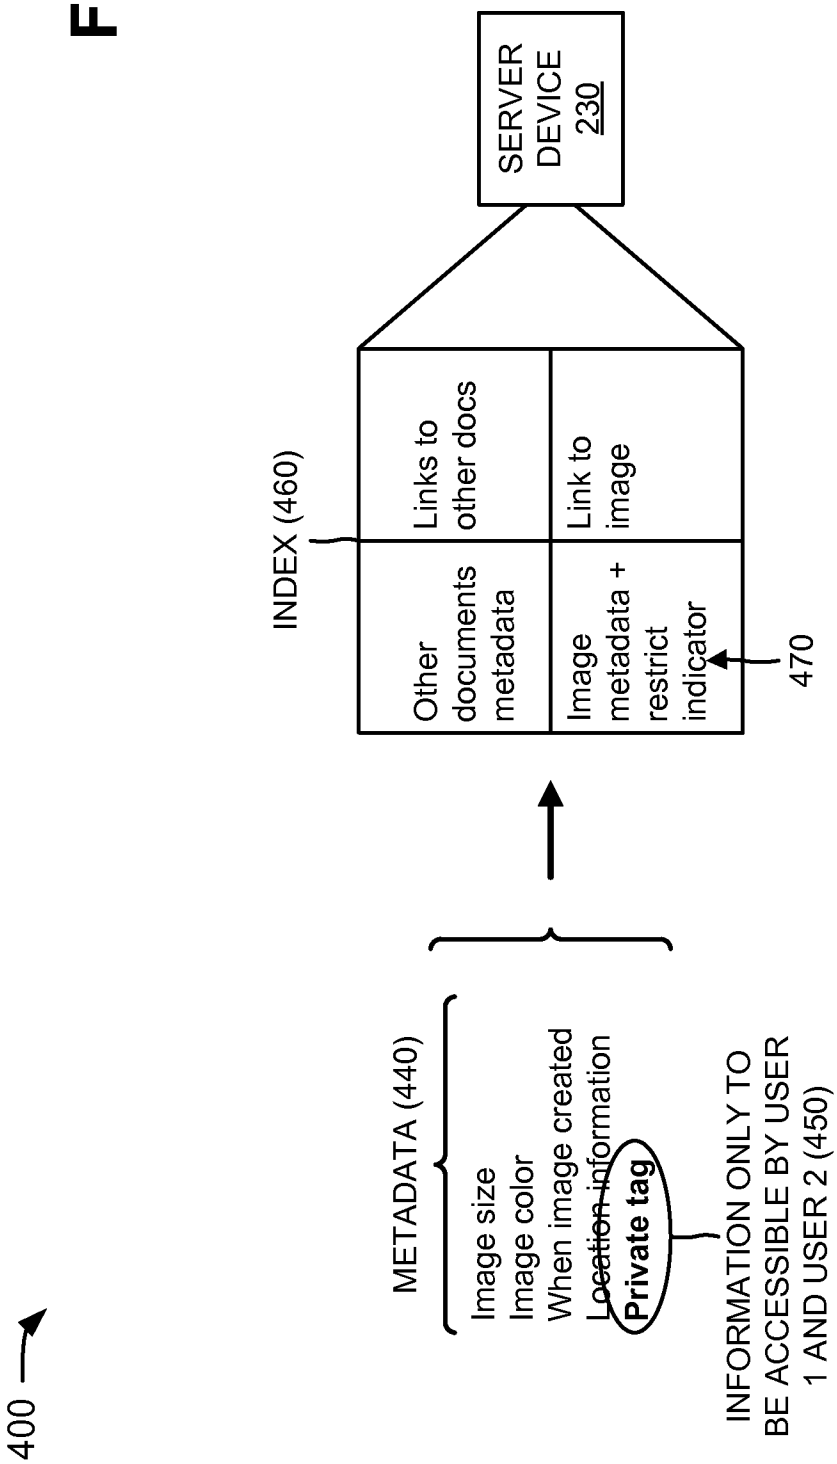

400
METADATA (440)
Image size
Image color
When image created
Location information
Private tag
INFORMATION ONLY TO BE ACCESSIBLE BY USER 1 AND USER 2 (450)
INDEX (460)
Other documents metadata
Links to other docs
Image metadata + restrict indicator
Link to image
470
SERVER DEVICE 230

FIG. 5

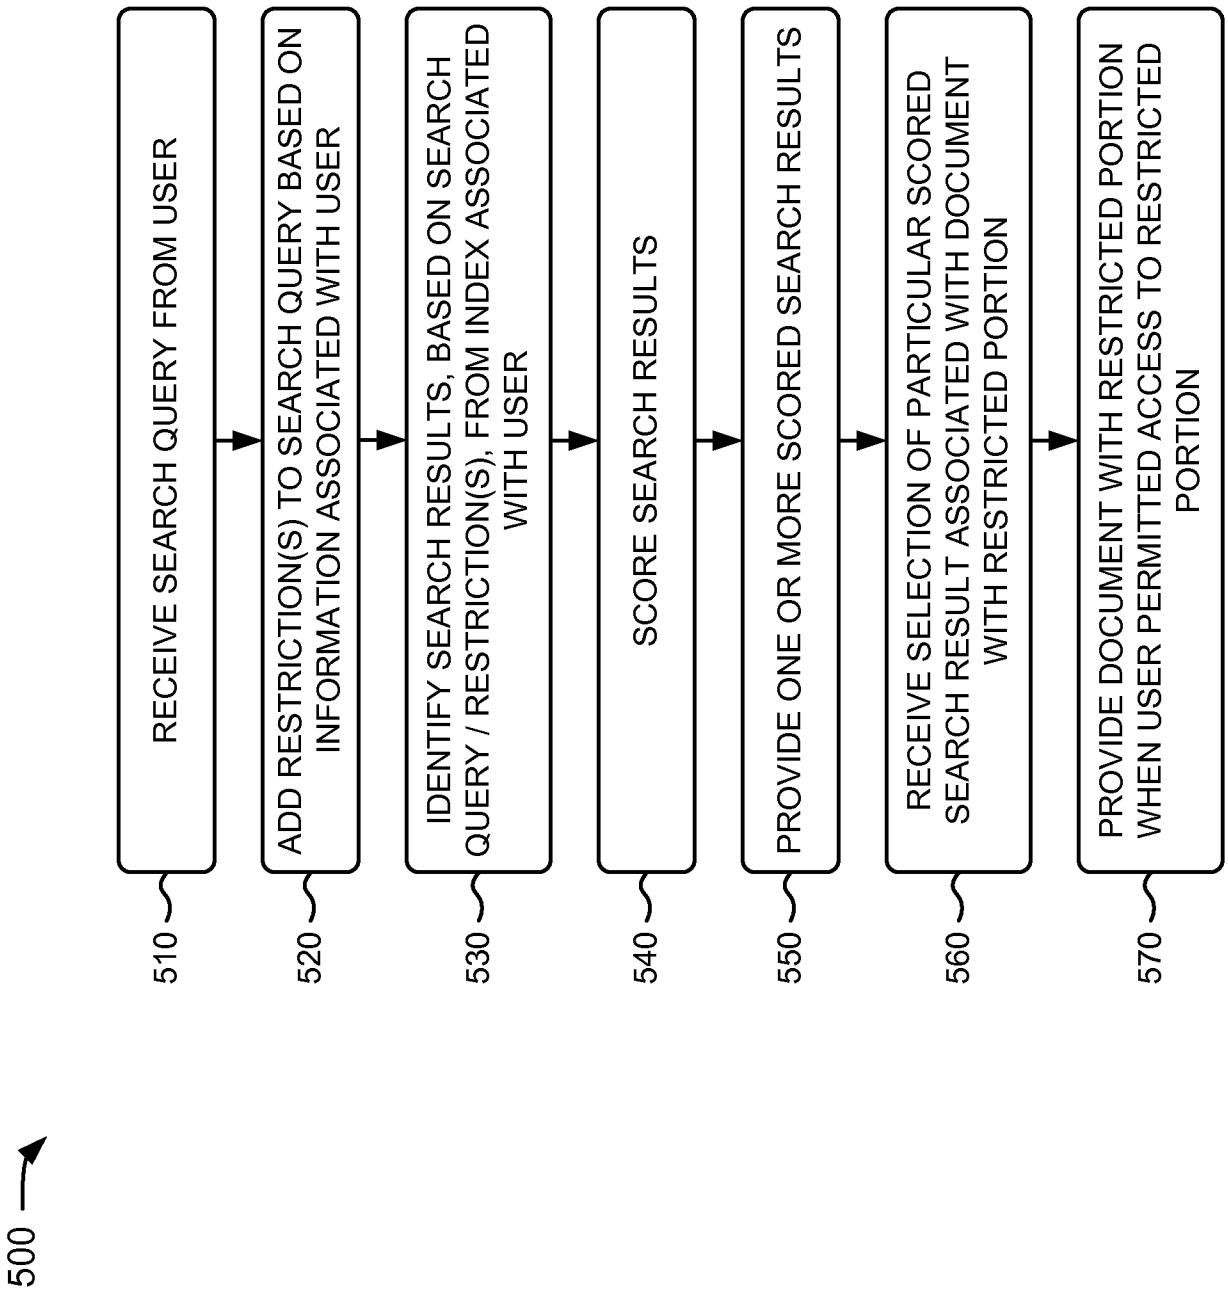

500
510
RECEIVE SEARCH QUERY FROM USER
520
ADD RESTRICTION(S) TO SEARCH QUERY BASED ON INFORMATION ASSOCIATED WITH USER
530
IDENTIFY SEARCH RESULTS, BASED ON SEARCH QUERY / RESTRICTION(S), FROM INDEX ASSOCIATED WITH USER
540
SCORE SEARCH RESULTS
550
PROVIDE ONE OR MORE SCORED SEARCH RESULTS
560
RECEIVE SELECTION OF PARTICULAR SCORED SEARCH RESULT ASSOCIATED WITH DOCUMENT WITH RESTRICTED PORTION
570
PROVIDE DOCUMENT WITH RESTRICTED PORTION WHEN USER PERMITTED ACCESS TO RESTRICTED PORTION

FIG. 6A

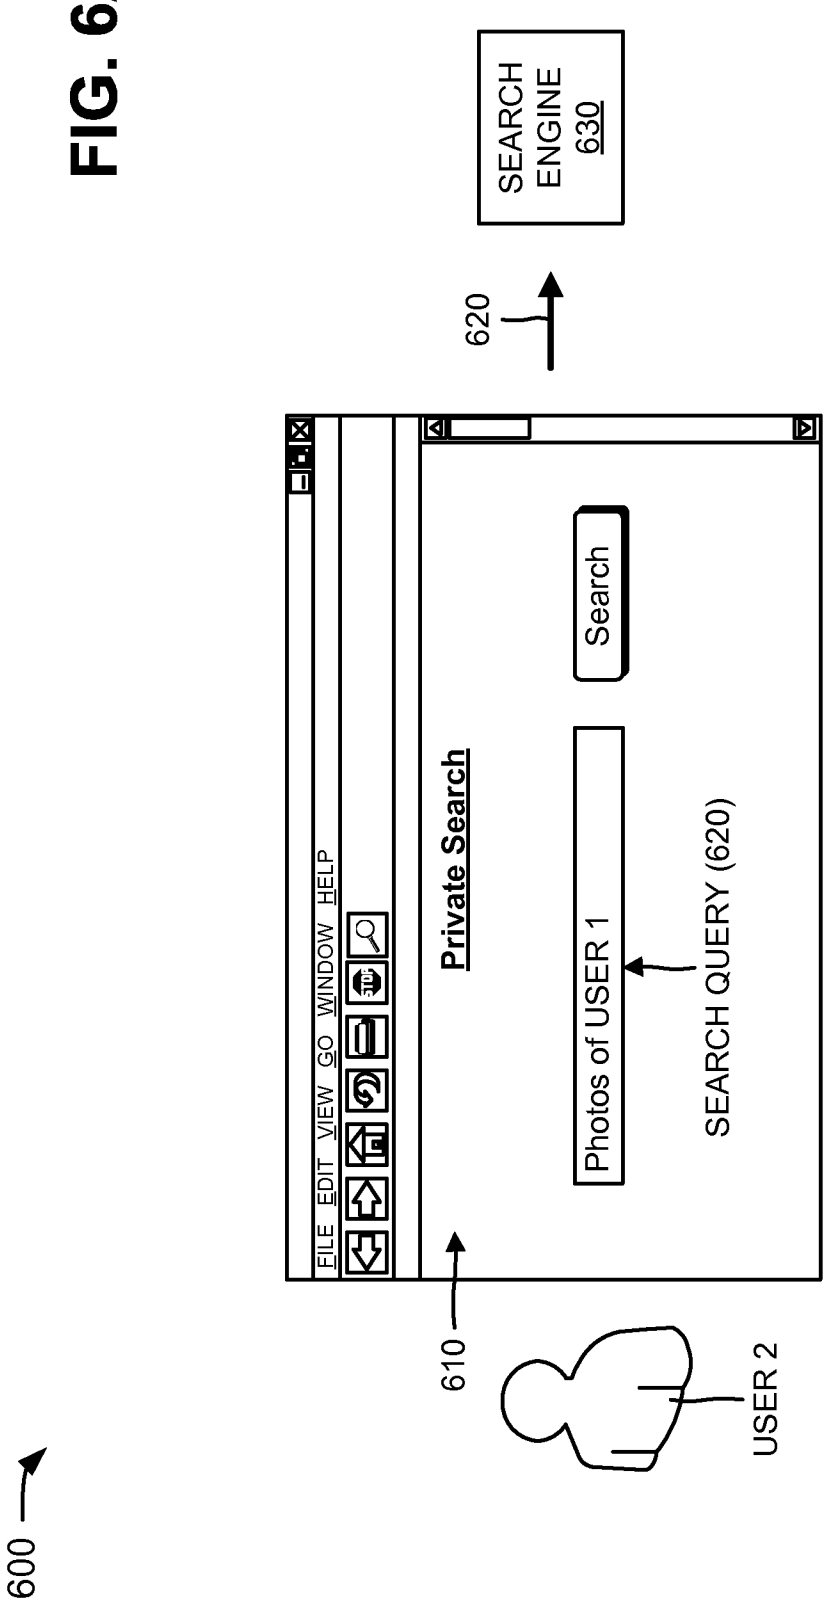

600
610
620
SEARCH ENGINE 630
USER 2
FILE EDIT VIEW GO WINDOW HELP
Private Search
Photos of USER 1
Search
SEARCH QUERY (620)

FIG. 6B

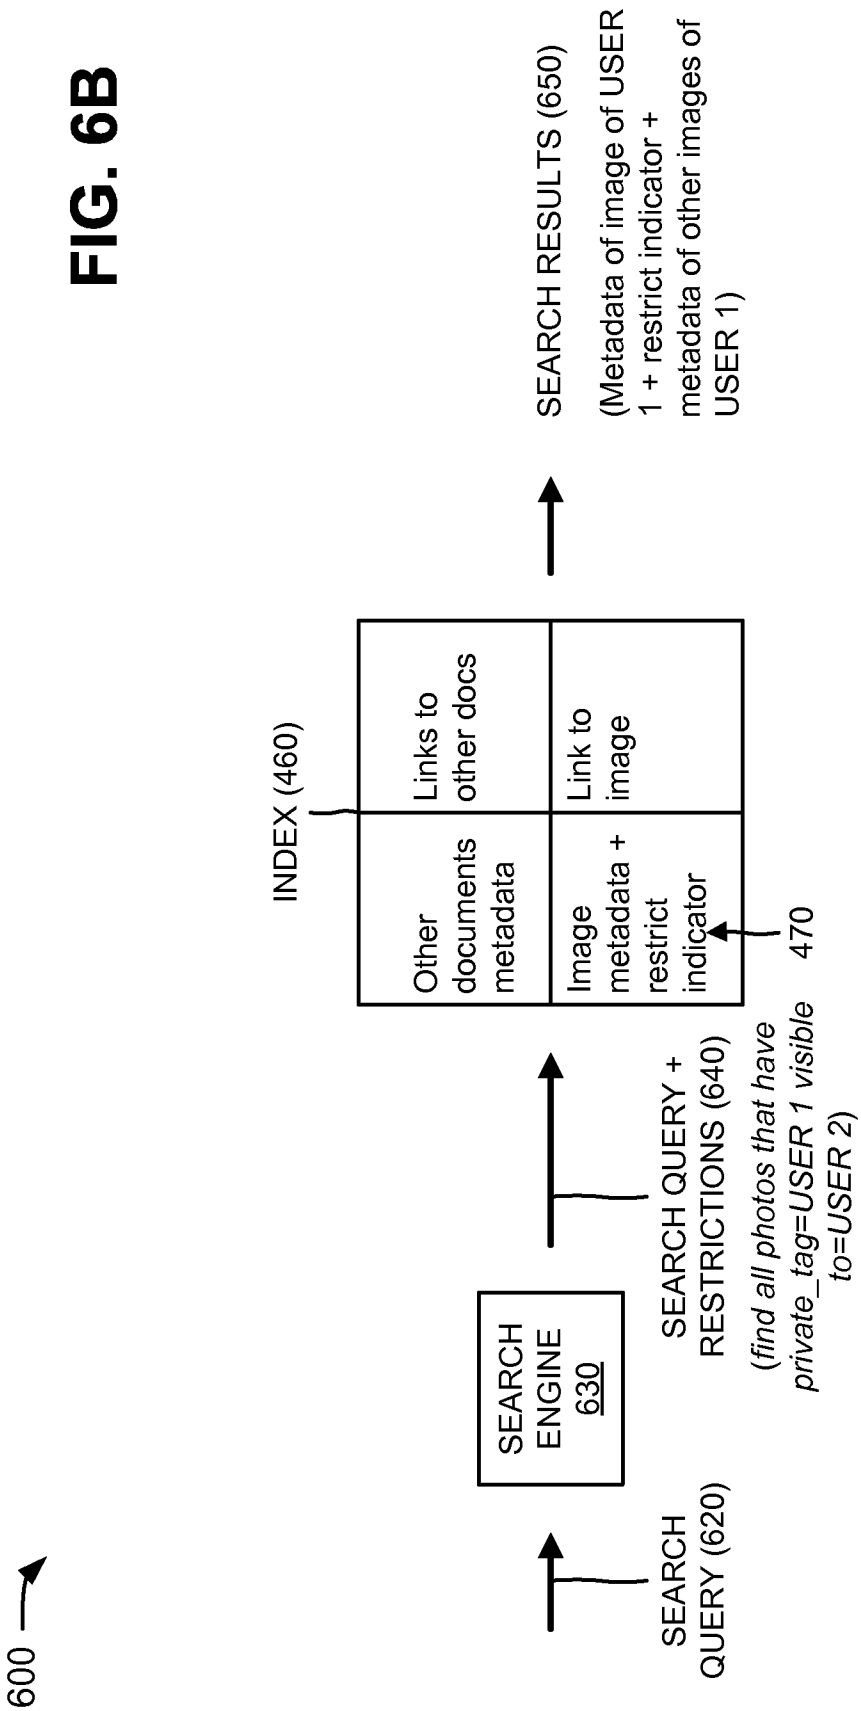

600
SEARCH QUERY (620)
SEARCH ENGINE 630
SEARCH QUERY + RESTRICTIONS (640)
(find all photos that have private_tag=USER 1 visible to=USER 2)
INDEX (460)
Other documents metadata
Links to other docs
Image metadata + restrict indicator
Link to image
470
SEARCH RESULTS (650)
(Metadata of image of USER 1 + restrict indicator + metadata of other images of USER 1)

FIG. 6C

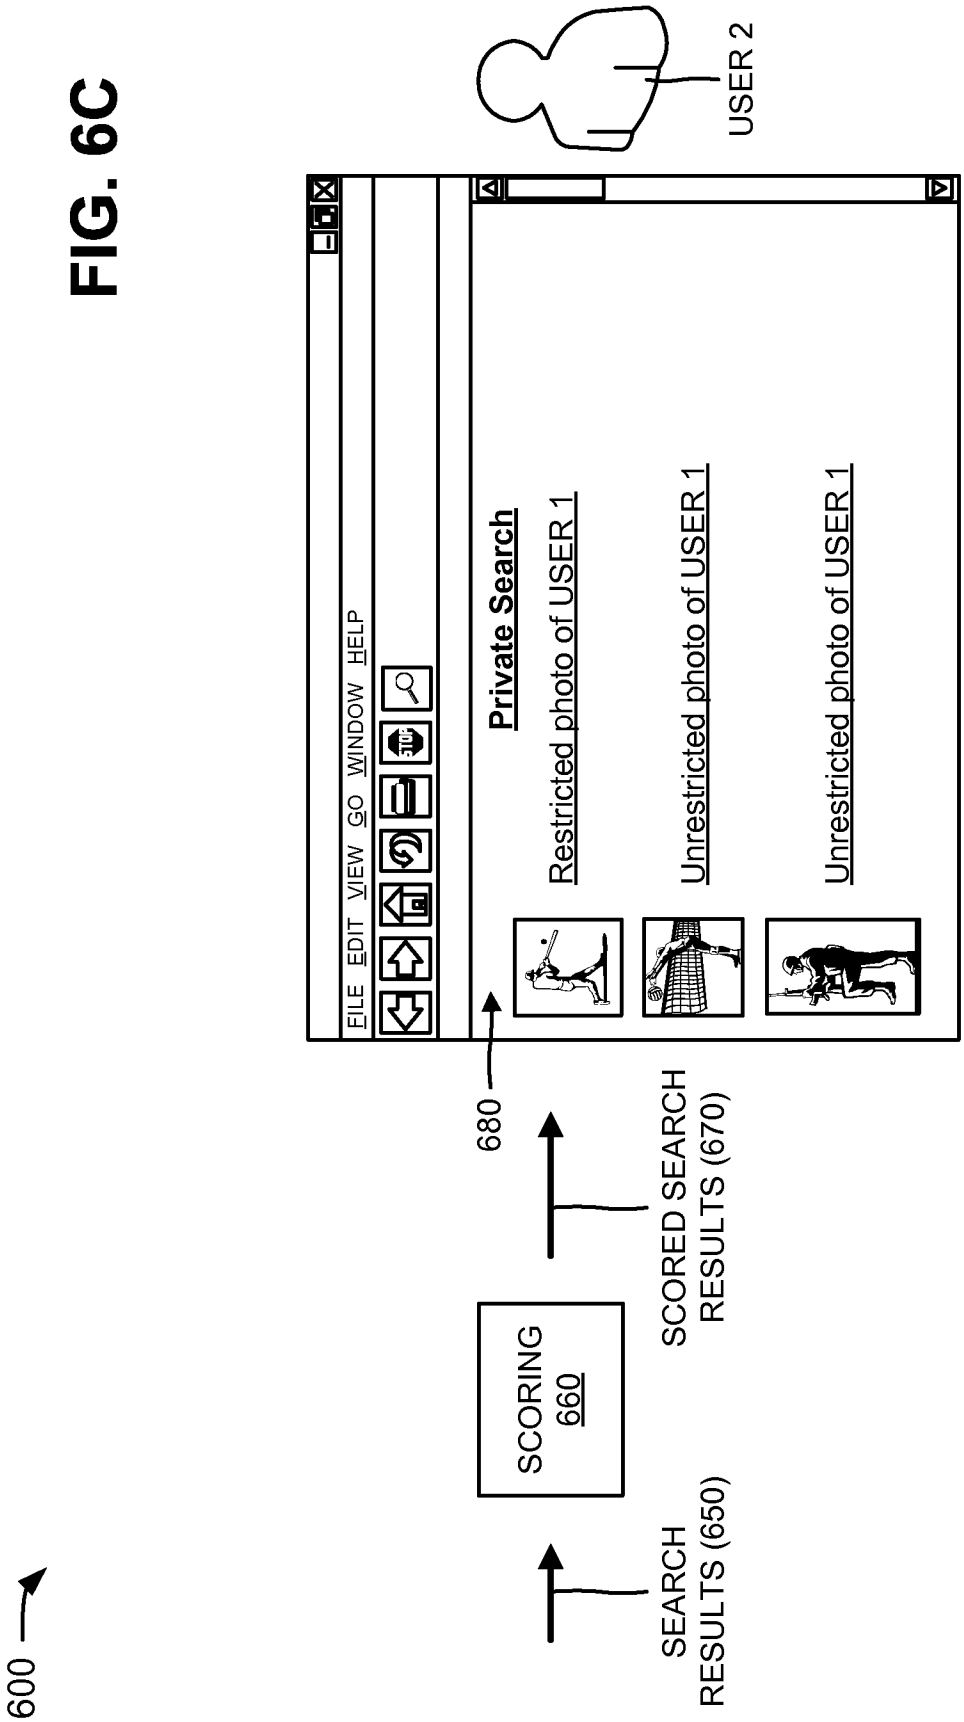

600
SEARCH RESULTS (650)
SCORING 660
SCORED SEARCH RESULTS (670)
680
FILE EDIT VIEW GO WINDOW HELP
Private Search
Restricted photo of USER 1
Unrestricted photo of USER 1
Unrestricted photo of USER 1
USER 2

FIG. 6D

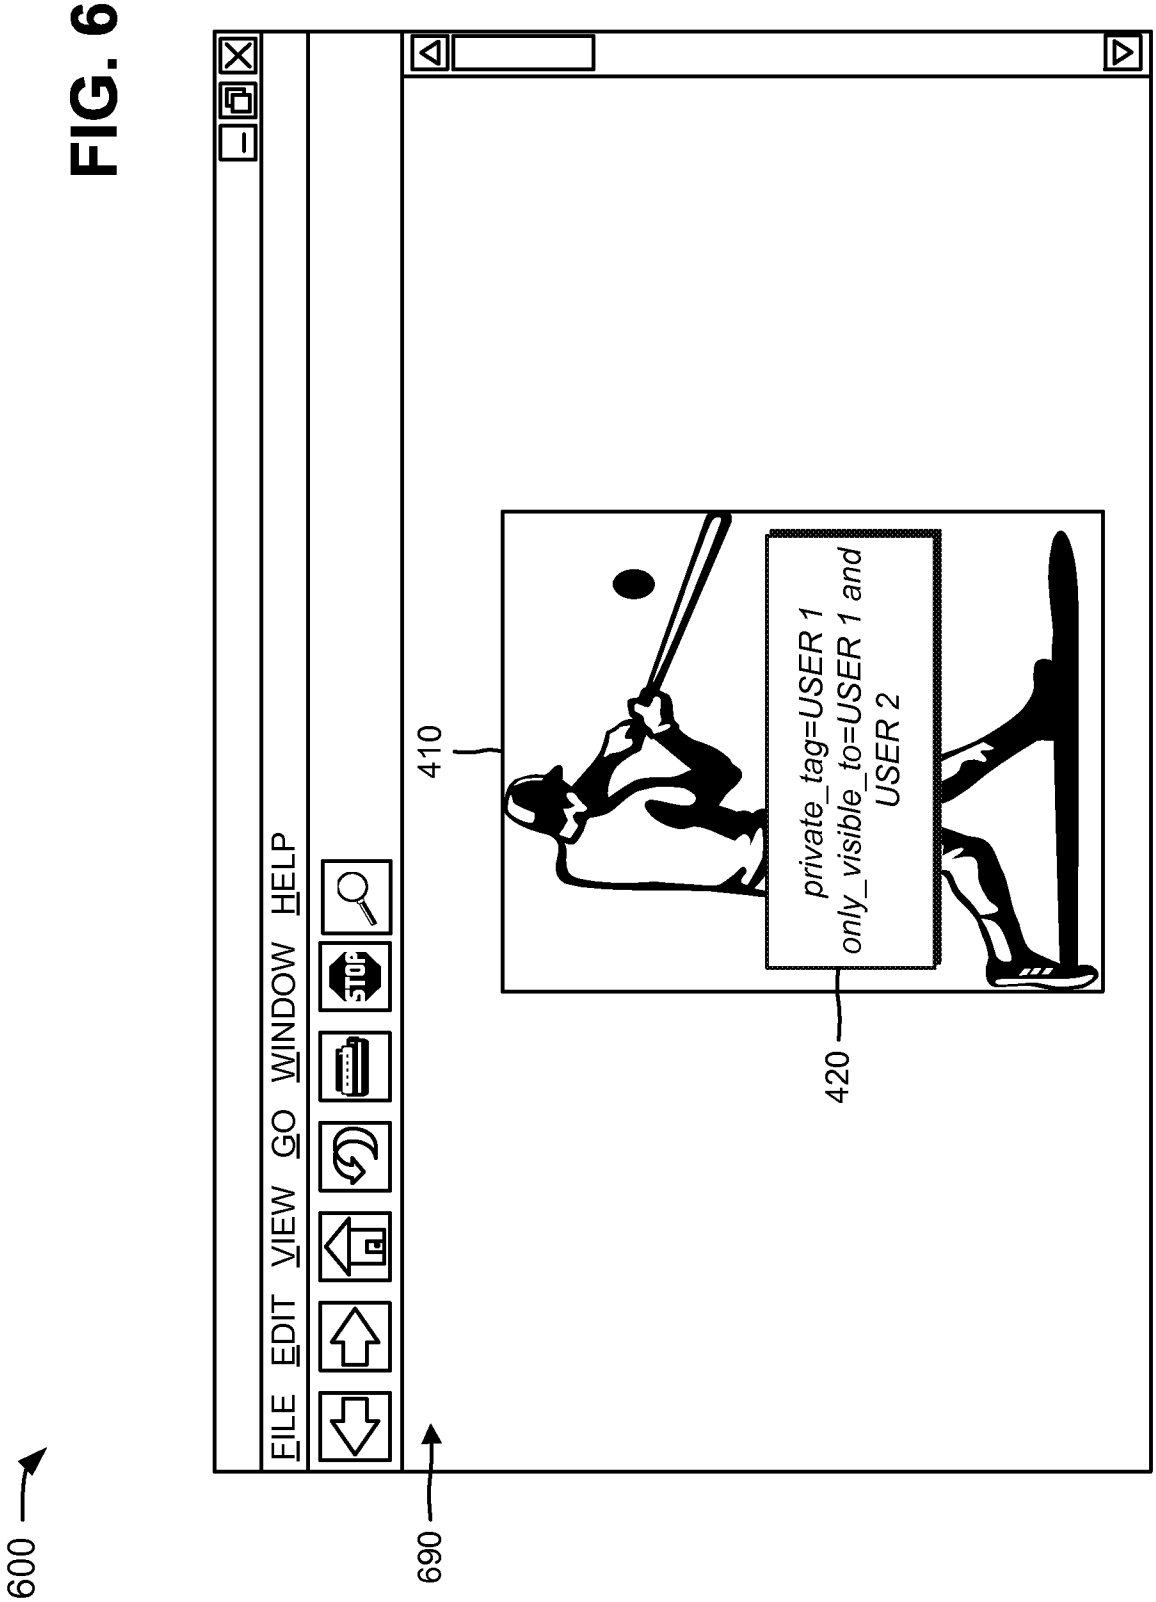
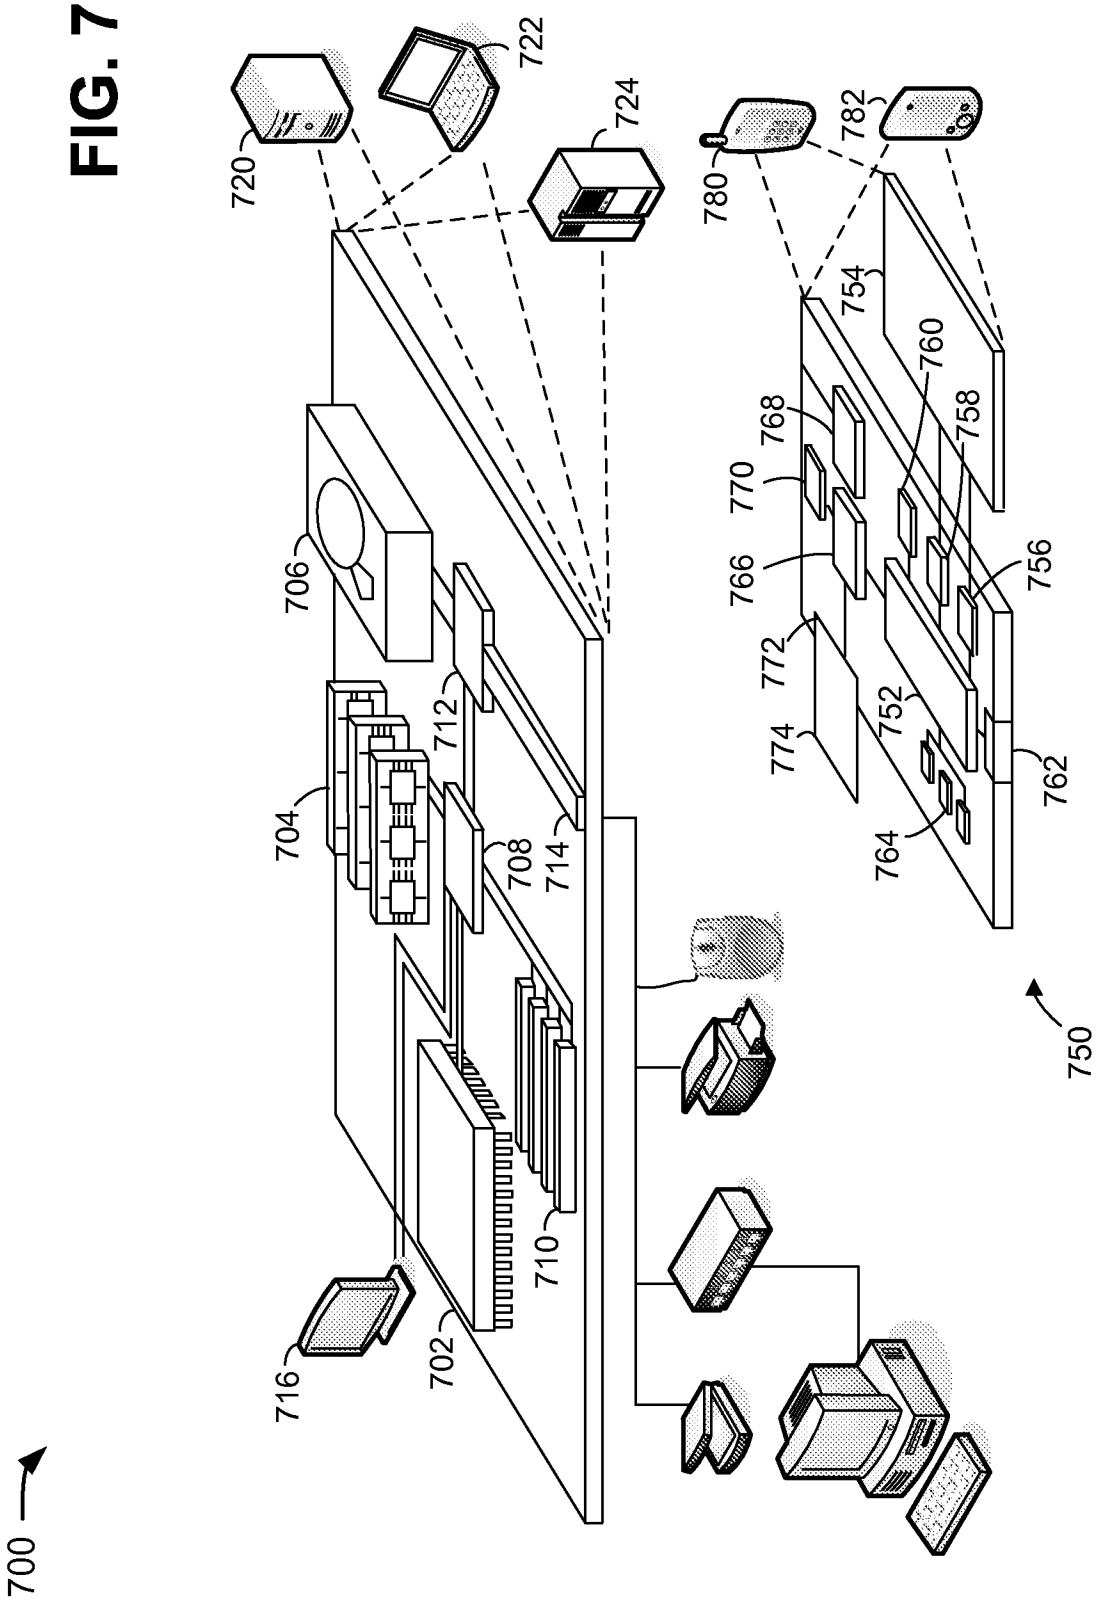

600
690
FILE EDIT VIEW GO WINDOW HELP
410
private_tag=USER 1
only_visible_to=USER 1 and
USER 2
420

FIG. 7


700
702
704
706
708
710
712
714
716
720
722
724
750
752
754
756
758
760
762
764
766
768
770
772
774
780
782

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2014/037973

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

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

B. FIELDS SEARCHED

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

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

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012/117626 A1 (YATES ROBERT LESLIE [US] ET AL) 10 May 2012 (2012-05-10) paragraphs [0023] - [0039]; figures 2-3	1-20
A	US 2013/054976 A1 (BROWN DOUGLAS S [US] ET AL) 28 February 2013 (2013-02-28) paragraphs [0011] - [0024]	1-20
A	WO 2012/058075 A1 (APPLE INC [US]; COLCLASURE KAELIN L [US]; HOERNKVIST JOHN M [US]; THOM) 3 May 2012 (2012-05-03) abstract	1-20

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
29 August 2014	04/09/2014

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Moon, Timothy

5

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2014/037973

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012117626 A1	10-05-2012	NONE	
US 2013054976 A1	28-02-2013	CA 2776182 A1	23-02-2013
		CN 103051600 A	17-04-2013
		DE 102012213807 A1	28-02-2013
		GB 2494022 A	27-02-2013
		US 2013054976 A1	28-02-2013
WO 2012058075 A1	03-05-2012	AU 2011320742 A1	09-05-2013
		CN 103250161 A	14-08-2013
		EP 2633469 A1	04-09-2013
		JP 2014503865 A	13-02-2014
		KR 20130067304 A	21-06-2013
		US 2012110323 A1	03-05-2012
		WO 2012058075 A1	03-05-2012