



- (51) International Patent Classification:
H04L 9/32 (2006.01) *H04L 29/06* (2006.01)
- (21) International Application Number:
PCT/US2015/024760
- (22) International Filing Date:
7 April 2015 (07.04.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant: **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 11445 Compaq Center Drive W., Houston, Texas 77070 (US).
- (72) Inventors: **BEITER, Michael B.**; 3404 E. Harmony Road, Fort Collins, Colorado 80528-9544 (US). **SAUNDERS, Shane I.**; 3404 E. Harmony Road, Fort Collins, Colorado 80528-9544 (US). **KORDIS, Susan J.**; 3404 E. Harmony Road, Fort Collins, Colorado 80528-9544 (US).

- (74) Agents: **ALLEN, Matthew R.** et al.; Hewlett-Packard Company, Intellectual Property Administration, 3404 E. Harmony Road, Mail Stop 35, Fort Collins, Colorado 80528 (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, 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, ST, 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,

[Continued on next page]

(54) Title: PROVIDING SELECTIVE ACCESS TO RESOURCES

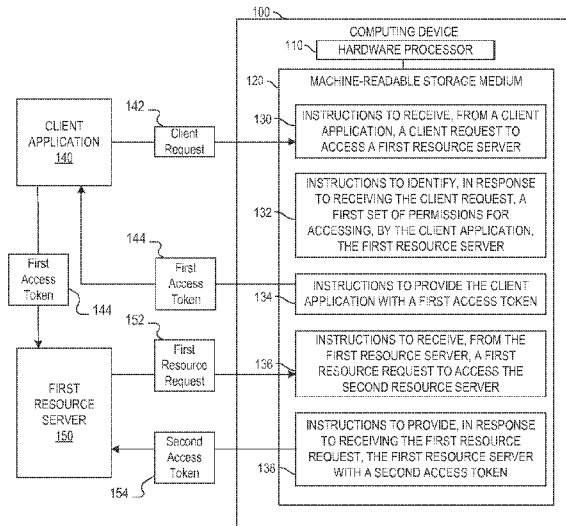



FIG. 1

(57) Abstract: Examples relate to providing selective access to resources. In one example, a computing device may: receive, from a client application, a request to access a first resource server, the request including a client access token; identify a first set of permissions specified by a client topology, the client topology specifying: the first resource server; the first set of permissions for accessing, by the client application, the first resource server; a second resource server; and a second set of permissions for accessing, by the first resource server, the second resource server; provide the client application with a first access token specifying the first set of permissions and the first resource server; receive, from the first resource server, a request to access the second resource server, the request including the first access token; and provide the first resource server with a second access token specifying the second set of permissions.



SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, — *as to applicant's entitlement to apply for and be granted*
GW, KM, ML, MR, NE, SN, TD, TG). *a patent (Rule 4.17(ii))*

Declarations under Rule 4.17:

— *as to the identity of the inventor (Rule 4.17(i))*

Published:

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

PROVIDING SELECTIVE ACCESS TO RESOURCES

BACKGROUND

[0001] Computing services and applications increasingly seek access to resources that are owned and/or managed by other services and applications. For example, a web application running on a website may seek a user's permission to post content from the website onto the user's blog, which may be a resource managed by a separate entity. After obtaining permission to post on the user's behalf, the entity that manages the user's blog may allow the web application to post content on the user's blog.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] The following detailed description references the drawings, wherein:

[0003] FIG. 1 is a block diagram of an example computing device for providing selective access to resources.

[0004] FIG. 2A is an example data flow for providing selective access to resources.

[0005] FIG. 2B is an example topology used for providing selective access to resources.

[0006] FIG. 3 is a flowchart of an example method for providing selective access to resources.

DETAILED DESCRIPTION

[0007] Resource and service providers often delegate various tasks to other resource and service providers. For example, a social media website may make use of a data storage service to provide dynamic storage based on the social media website's needs, and may use an advertisement provider to provide advertisements to social media website users. These services or services, or cascaded services, may often be managed by a third party that specializes in providing a particular service. In situations where a user allows a client service or application to access the user's resources managed by a particular resource

provider, e.g., a situation where the user permits a website to post to the user's blog, the permissions granted to the cascaded services of the particular resource provider may be controlled in a manner designed to ensure that cascaded services and resource providers are only provided with the permissions needed to perform their particular task.

[0008] In general, for a client service or application to access a user's resources managed by a separate resource server, the user grants the client access by communicating with an authorization server associated with the resource server. The authorization server, upon receiving authentication from the user and, in some implementations, also from the client – such as a user name and password, a client ID, and/or a secret – issues the client a token that allows the client to request access to the user's resources on the resource server, as authorized by the user. The token provided to the client application identifies the client application and generally grants no permissions. To interact with the resource server, the client application may exchange the token for a second token that specifies the resource server and the scope of permissions granted to the client application with respect to the resource server. The scope of the permissions granted to the client application may vary, e.g., depending upon what access the user determined to provide the client application. For example, a user may grant a web application permission to post content on the user's blog, but not permission to modify the user's profile information.

[0009] For the resource server to perform various actions, e.g., to post an article on the user's blog, the resource server may need to communicate with one or more cascaded services. For example, to ensure that content posted on a user's blog doesn't violate the resource provider's terms of service or other content restrictions, the resource provider may use, as a cascaded service, a content verification service that verifies the content provided by the client application, e.g., the content verification service may attempt to ensure the content doesn't violate any copyrights or restrictions on content deemed inappropriate by the resource server. Another example cascaded service may be used to render the content in a particular format, for instance, as specified in the user's profile.

[0010] To ensure that cascaded services are only provided with the permissions needed to perform their task or tasks, the resource server does not provide the cascaded service with the token provided by the client application. The resource server may, instead, provide its token to the authorization server with a request to access the cascaded service. The authorization server may provide the resource server with a token specific to the cascaded service, and this token may have permissions that are of a different scope than the permissions specified by the token given to the resource server. For example, a content rendering service may not need a token that includes permissions to post on a user's blog and may only need permissions to read the format of the blog from the user's profile to appropriately render the content to be posted by the resource server.

[0011] The permissions to be granted by a token between any services associated with the resource server are specified in a client application topology accessible to the authorization server. The authorization server may use the topology to determine which permissions are to be granted by tokens issued to requesting services. Further details regarding the manner in which an authorization server, client, and resource providers interact to selectively delegate access to resources are provided in the paragraphs that follow.

[0012] Referring now to the drawings, FIG. 1 is a block diagram of an example computing device 100 for selectively delegating access to resources. Computing device 100 may be, for example, a server computer, a personal computer, a mobile computing device, or any other electronic device suitable for processing data. In the embodiment of FIG. 1, computing device 100 includes hardware processor 110 and machine-readable storage medium 120.

[0013] Hardware processor 110 may be one or more central processing units (CPUs), semiconductor-based microprocessors, and/or other hardware devices suitable for retrieval and execution of instructions stored in machine-readable storage medium 120. Hardware processor 110 may fetch, decode, and execute instructions, such as 130-138, to control the process for selectively delegating

access to resources. As an alternative or in addition to retrieving and executing instructions, hardware processor 110 may include one or more electronic circuits that include electronic components for performing the functionality of one or more of instructions.

[0014] A machine-readable storage medium, such as 120, may be any electronic, magnetic, optical, or other physical storage device that contains or stores executable instructions. Thus, machine-readable storage medium 120 may be, for example, Random Access Memory (RAM), an Electrically Erasable Programmable Read-Only Memory (EEPROM), a storage device, an optical disc, and the like. In some implementations, storage medium 120 may be a non-transitory storage medium, where the term "non-transitory" does not encompass transitory propagating signals. As described in detail below, machine-readable storage medium 120 may be encoded with a series of executable instructions: 130-138, for selectively delegating access to resources.

[0015] As shown in FIG. 1, the computing device 100 executes instructions to receive, from a client application 140, a client request 142 to access a first resource server 150 (130). The client request 142 includes an access token specifying the client application 140 as the audience, e.g., the intended recipient. In some implementations, the access token may have been previously provided to the client application 140 by the computing device 100. For example, the access token may have been provided to the client application 140 as the result of an authorization process where a user authorized the client application 140 to access a protected resource on the user's behalf. The access token alone, in some implementations, grants no permissions to access resources, e.g., it must be exchanged for the client application 140 to access a particular protected resource.

[0016] The computing device 100 executes instructions to identify, in response to receiving the client request 142, a first set of permissions for accessing, by the client application 140, the first resource server 150(132). The first set of permissions may be specified by a topology for the client application 140. The topology may specify, for example: the first resource server 150 and the first set of permissions for accessing the first resource server 150 by the client

application 140; and a second resource server with a second set of permissions for accessing the second resource server by the first resource server 150.

[0017] By way of example, the client application 140 may be a news feed application running on a user's mobile device, and the first resource server 150 may be the host of a blog managed by the user. In this example, the topology for the news feed application specifies the blog host and a first set of permissions granted to the news feed application on the user's behalf. For example, the user may have previously authorized the news feed application to post news stories to the user's blog, e.g., upon the user's interaction with a news article presented within the news feed application. In this example, the topology for the news feed application also specifies a cascaded service, such as a rendering service for formatting a news article being posted to the user's blog, and a set of permissions for the blog host to interact with the cascaded service. A rendering service may have permission to access the user's profile information, such as the user's blog layout, to perform the rendering service, but may not have permission to post anything to the user's blog. As the rendering service may be managed by a third party – outside of the control of the blog host – limiting the permissions granted to cascaded services may prevent unauthorized and/or un-intended access to the user's protected resources.

[0018] In some implementations, the computing device 100 may receive the topology from the client application 140. For example, a developer of the client application 140 may generate a topology based on data obtained from the first resource server 150 and its associated resource servers and services, cascaded or otherwise.

[0019] In some implementations, the computing device 100 may generate the topology for the client application 140 using sub-topologies received from the client application 140 and/or the first resource server 150. For example, rather than relying on the developer of the client application 140 to know the topology of the first resource server, including all of the associated resource servers and permissions, the developer of the client application 140 may provide a sub-topology specifying the permissions associated with the direct

interactions of the client application 140 with the resources associated with the first resource server 150 which, in this example, is only the first resource server 150. The first resource server 150 may provide a sub-topology specifying the permissions associated with the direct interactions of the first resource server 150 and, in some implementations, any cascaded resource servers and services associated with the first resource server 150. Using the sub-topologies provided by the client application 140 and the first resource server 150, the computing device 100 may generate the entire topology specifying the permissions of interactions between the client application 140, the first resource server 150, and resource servers and services associated with the first resource server 150.

[0020] The computing device 100 executes instructions to provide the client application 140 with a first access token 144 (134). The first access token 144 specifies the first set of permissions and, as an audience, the first resource server 150. For example, as specified by the example topology, the news feed application may be provided with an access token that specifies the blog host as the audience and the permission to post to the user's blog in the first set of permissions.

[0021] In the example environment depicted in FIG. 1, the client application 140 is depicted as communicating with the first resource server 150 using the first access token 144. The communications may include, for example, a communication that includes the news feed application's news article or a link to the article, a request to post the news article on the user's blog, and the token indicating that the news feed application has been authorized to post to the user's blog.

[0022] The computing device 100 executes instructions to receive, from the first resource server 150, a first resource request 152 to access the second resource server (136). For example, upon receipt of the first access token 144 and data indicating that the news feed application has a news article and authorization to post the news article on the user's blog, the blog host may need to communicate with a rendering service to format the news article in a manner that

meets the design standards of the blog host and/or the user's blog layout. The first access token 144 does not specify the second resource server, e.g., the rendering service, as an audience, and may not include the necessary permissions, e.g., access to user's blog layout; accordingly, in this example, the first resource server 150 sends the first resource request 152 for the proper token to communicate with and make use of the second resource server.

[0023] The computing device 100 executes instructions to provide, in response to receiving the first resource request 152, the first resource server 150 with a second access token 154 (138). The second access token 154 specifies the second set of permissions and, as a second audience, the second resource server. As with the first access token 144, the second access token 154 may be provided to the first resource server 150 based on the client application topology. Using the example provided above, the second access token 154 may specify the rendering service as the audience and the permissions may include access to the user's blog layout.

[0024] As discussed above, the permissions provided by any given token may vary. Resource servers, cascaded or otherwise, may have any combination of permissions as specified in the topology for a given client application. While the example described with reference to FIG. 1 includes a single resource server and a single cascaded service, many other resource servers – which may include cascaded services – may be included in a topology. In addition, a variety of configurations are possible for communications between resource servers, cascaded or otherwise, e.g., cascaded servers may also use cascaded services, and permissions between services may be uni-directional – where the permissions granted by a first service to a second service are different from the permissions granted by the second service to the first service.

[0025] Each client application may be associated with multiple topologies, e.g., one topology for each different protected resource to which the client application may access on behalf of a user. For example, the news feed application may have authorization to access the user's second blog on a second blog host, authorization to the user's email contacts to send a news article to one of the

user's contacts, and/or authorization to the user's advertising content profile to add topics of interest to the user's profile based on the articles in which the user is interested. In addition, each resource server, cascaded or otherwise, may belong to multiple topologies. For example, a third party rendering service may be used by multiple resource servers, and belong to multiple topologies associated with the same client application. As another example, each client application that is authorized to access a resource server may have a topology for the interactions with the resource server, and the topology may be separate from the topologies of other client applications for the same resource server.

[0026] FIG. 2A is an example data flow 200 for providing selecting access to resources, while FIG. 2B is an example topology used for providing selective access to resources. The data flow 200 depicts an authorization server 230, which may be implemented by a computing device, such as the computing device 100 described above with respect to FIG. 1. The user device 210, client application 220, first resource server 240, and second resource server 250 may be implemented by any computing device suitable for processing data, such as a personal computer, mobile computer, or server computer. For example, the user device 210 may be a user's personal computer, the client application 220 may be a web application implemented on a web server, the authorization server 230, first resource server 240, and second resource server 250 may be implemented on one or more server devices, virtual machines, and the like.

[0027] The example data flow 200 depicts an example exchange of information that may occur between the depicted devices, e.g., using example topology 280. In general, and as described above, the topology 280 may be generated by the developer of the client application 220, alone or in combination with the entity that manages the first resource server 240. For example, in some implementations, the client topology may be generated based on permissions data received from an entity that manages the first resource server 240, where the permissions data specifies, for each of the resource servers associated with the first resource server, e.g., second 250, third 260, and fourth 270 resource servers, at least one of the other resource servers and, for each of the other resource servers, permissions granted to the other resource server. The topology is

generally accessible by the authorization server 230, which may use the topology to distribute tokens for accessing protected resources.

[0028] In the example data flow 200, a user of the user device 210 may wish to grant access to the user's protected resources, e.g., associated with the first resource server 240, to a client application 220. By way of example, the first resource server 240 may provide a digital photography management service to which the user of user device 210 is subscribed. The user's photographs, e.g., protected resources, may be managed by the first resource server 240 in conjunction with other resource servers, e.g., providing storage, user profile management, and printing services. Access to the photographs may be protected, e.g., by a user name and password. The client application 220, which by way of example may be an image tagging web application running on a web server, may give the user device 210 redirect data 212 that redirects a web browser of the user device 210 to the authorization server 230.

[0029] The user device 210 sends a user request 214 to the authorization server 230 requesting that the client application 220 be given access to the user's images managed by the first resource server 240. The authorization server 230 may, in the example given above, require that the user authenticate by providing a user name and password and, once authenticated, provide the user device 210 with an access code 216. The user device 210 provides the access code 216 to the client application 220, which provides the access code 216 and, in some implementations, client authentication credentials, to the authorization server in exchange for the access token 218. The access token 218 is provided by the authorization server 230 to allow the client application to access resources associated with the first resource server 240. In some implementations, the access token 218 does not, by itself, grant access to any individual resources or resource servers.

[0030] The client application 220 provides the authorization server 230 with the access token 218 along with a request to access the first resource server 240. The authorization server 230, upon receipt of the request and the access token 218, may identify, in the client application topology 280, a first token 222 that is

defined for accessing the first resource server 240. In the example topology 280, the first token 222 is represented by the string, "TID_01: {P3, P4}." This example string indicates the topology, "TID," the audience, "01," which indicates the first resource server 240, and a set two permissions, P3 and P4. The manner in which a token may be presented may vary, and tokens generally specify the audience of the token, permissions indicating what can be accessed at the audience, and – in some implementations – the topology to which the token belongs. Using the image tagging application example, the first token 222 provided to the image tagging application may specify the audience, which may be the first resource server 240, and the set of permissions may specify that the image tagging application may access general images and image attributes.

[0031] The client application 220 provides the first token to the first resource server 240 with a request for a particular resource, e.g., an image to which the image tagging application is seeking to tag. In the example data flow 200, the first resource server 240 relies on the second resource server 250 to help fulfill the resource request. For example, to tag an image, an image tagging application may need access to an actual image file, and/or may need permission to store an image file, on a storage device that is managed by the second resource server 250.

[0032] To communicate with the second resource server 250 with the proper permissions to access the requested resource, the first resource server provides the first token 222 to the authorization server 230 in exchange for the second token 224. As shown in the example topology 280, the second token 224 is represented by the string, "TID_02: {p5}," indicating that the first resource server 240 is authorized to access "02" of "TID" with permission, p5. In the image tagging example, the permission may be, for example, file permissions to retrieve and store images.

[0033] In the example data flow 200, upon receiving an authorized request, e.g., authorized due to the second token 224, the second resource server 250 provides data 226 representing the requested resource to the first resource server 240, which may then provide the data 226 to the client application 220. In the

image tagging application example, the data 226 may be an image file to be analyzed and tagged by the client application, or may be confirmation that an image file stored by the second resource server 250 has changed.

[0034] The example topology 280 depicted in FIG. 2B also indicates a third token 232 for granting the client application 220 access to a third resource server 260, a fourth token 234 for granting the first resource server 240 access to the third resource server 260, and a fifth token 236 for granting the first resource 240 access to the fourth resource server 270. The topology 280 also indicates, using arrows, the direction(s) in which tokens provide access. The access granted from one resource server to another may not be the same as the access granted in the reverse direction – if any access in the reverse direction is granted at all. For example, the first resource server 240 may access the third resource server with permission, P2, but no token exists in the topology for providing the third resource server 260 access to the first resource server. Many other topology variations may be used to selectively provide access to various related and cascaded resource servers.

[0035] FIG. 3 is a flowchart of an example method 300 for providing selective access to resources. The method 300 may be performed by a computing device, such as a computing device described in FIG. 1. Other computing devices may also be used to execute method 300, such as a computing device used to implement a resource server. Method 300 may be implemented in the form of executable instructions stored on a machine-readable storage medium, such as the storage medium 120, and/or in the form of electronic circuitry.

[0036] A client request for user data and a first token specifying a first audience and a first set of permissions is received from a client device (302). The client device may be, for example, a device implementing a client application to which a user has delegated access, e.g., access to the user's protected resources managed by a first server. The first token and client request may be received by a first server, which manages the user's protected resources.

[0037] An authorization server is provided with a token request that includes i) the first token, and ii) a request to access a resource server (304). For example,

the first server that manages the user's protected resources may need access to a cascaded resource server, e.g., a resource server to which the first server has delegated some functionality. To access the cascaded resource server, the first server exchanges the first token for a token providing access to the cascaded resource server.

[0038] A resource server token is received from the authorization server, the resource server token specifying a second audience and a second set of permissions, the second audience being the resource server (306). For example, the resource server token may specify the cascaded resource server as the audience, and a set of permissions to access the cascaded resource and fulfill the client's request for user data. In some implementations, the set of permissions specified by the resource server token are different from the first set of permissions specified by the first token. For example, the cascaded resource server may require different permissions to perform its service than the permissions required by the first server.

[0039] The resource server is provided with i) a resource request for the user data, and ii) the resource server token (308). For example, the resource request for the user data may be similar to the client's request for user data, the resource server token that was received from the authorization server specifies that the requesting server, e.g., the first server, has the permissions to request the user data and that the intended audience for the resource server token is the resource server.

[0040] The user data is received from the resource server (310). For example, using the resource server token and the resource request, the resource server is able to provide the first server with the requested user data.

[0041] The user data is provided to the client device (312). For example, the first server, upon receiving the requested data, may provide the data to the client device from which the client request originated.

[0042] The foregoing disclosure describes a number of example implementations for providing selective access to resources. As detailed above, examples provide a mechanism for using client application and resource server topology to selectively provide access to various resource servers with various permissions and potential applications of a system that is capable of providing selective access to resources.

CLAIMS

We claim:

1. A non-transitory machine-readable storage medium encoded with instructions executable by a hardware processor of a computing device for providing selective access to resources, the machine-readable storage medium comprising instructions to cause the hardware processor to:

receive, from an authorization server, an access token for accessing resources associated with a first resource server;

provide the authorization server with the access token and a request to access the first resource server;

receive, from the authorization server, a first token, the first token specifying a first set of permissions for accessing the first resource server and, as a first audience, the first resource server, the first set of permissions being specified by a client topology, the client topology specifying:

the first resource server;

the first set of permissions for accessing the first resource server;

a second resource server; and

a second set of permissions for accessing, by the first resource server, the second resource server; and

provide the first resource server with a resource request for a resource, the resource request including the first token.

2. The storage medium of claim 1, wherein the access token grants no permissions.

3. The storage medium of claim 1, wherein the instructions further cause the hardware processor to:

provide the authorization server with the client topology.

4. The storage medium of claim 1, wherein:

the resource request specifies a request for particular user data stored by the second resource server specified by the client topology; and

the second set of permissions includes at least one permission that is different from the permissions included in the first set of permissions.

5. The storage medium of claim 4, wherein:

the second set of permissions is an only set of permissions, specified by the client topology, for accessing the second resource server; and

the client topology further specifies, as an only source associated with the second set of permissions, the first resource.

6. The storage medium of claim 4, wherein the instructions further cause the hardware processor to:

receive, from the first resource, the particular user data.

7. The storage medium of claim 1, wherein the instructions further cause the hardware processor to:

generate the client topology based on permissions data received from an entity that manages the first resource server, the permissions data specifying, for each of a plurality of resource servers associated with the first resource server, at least one other resource server of the plurality of resource servers and, for each of the at least one other resource servers, permissions granted to the other resource server.

8. A computing device for providing selective access to resources, the computing device comprising:

a hardware processor; and

a data storage device storing instructions that, when executed by the hardware processor, cause the hardware processor to:

receive, from a client application, a client request to access a first resource server, the client request including a client access token specifying, as a client audience, the client application;

in response to receiving the client request, identify a first set of permissions for accessing, by the client application, the first resource server, the first set of permissions being specified by a client topology for the client application, the client topology specifying:

the first resource server;

the first set of permissions for accessing, by the client application, the first resource server;

a second resource server; and

a second set of permissions for accessing, by the first resource server, the second resource server;

provide the client application with a first access token, the first access token specifying the first set of permissions and, as a first audience, the first resource server;

receive, from the first resource server, a first resource request to access the second resource server, the first resource request including the first access token; and

in response to receiving the first resource request, provide the first resource server with a second access token, the second access token specifying the second set of permissions and, as a second audience, the second resource server.

9. The computing device of claim 8, wherein the instructions further cause the processor to:

receive, from the client application, the client topology;

10. The computing device of claim 8, wherein the instructions further cause the processor to:

receive, from the client application, a client sub-topology, the client sub-topology specifying:

the first resource server; and

the first set of permissions;

receive, from the first resource server, a first resource sub-topology, the first resource sub-topology specifying:
the second resource server; and
the second set of permissions; and
generate the client topology using the client sub-topology and the first resource sub-topology.

11. The computing device of claim 8, wherein the second set of permissions includes at least one permission that is different from permissions included in the first set of permissions.

12. A method for providing selective access to resources, implemented by a hardware processor, the method comprising:

receiving, from a client device, i) a client request for user data, and ii) a first token specifying a first audience and a first set of permissions;

providing an authorization server with a token request, the token request including i) the first token, and ii) a request to access a resource server;

receiving, from the authorization server, a resource server token, the resource server token specifying a second audience and a second set of permissions, the second audience being the resource server;

providing the resource server with i) a resource request for the user data, and ii) the resource server token;

receiving, from the resource server, the user data; and

providing the user data to the client device.

13. The method of claim 12, wherein the second set of permissions includes at least one permission that is different from permissions specified by the first set of permissions.

14. The method of claim 12, further comprising:

providing the authorization server with a resource sub-topology that specifies, for each of a plurality of resource servers, at least one set of permissions for accessing the resource server, and

wherein the first set of permissions is specified by the resource sub-topology.

15. The method of claim 12, further comprising:

providing a resource sub-topology to a client associated with the client device, the resource sub-topology specifying, for each of a plurality of resource servers, at least one set of permissions for accessing the resource server, and

wherein the first set of permissions is specified by the resource sub-topology.

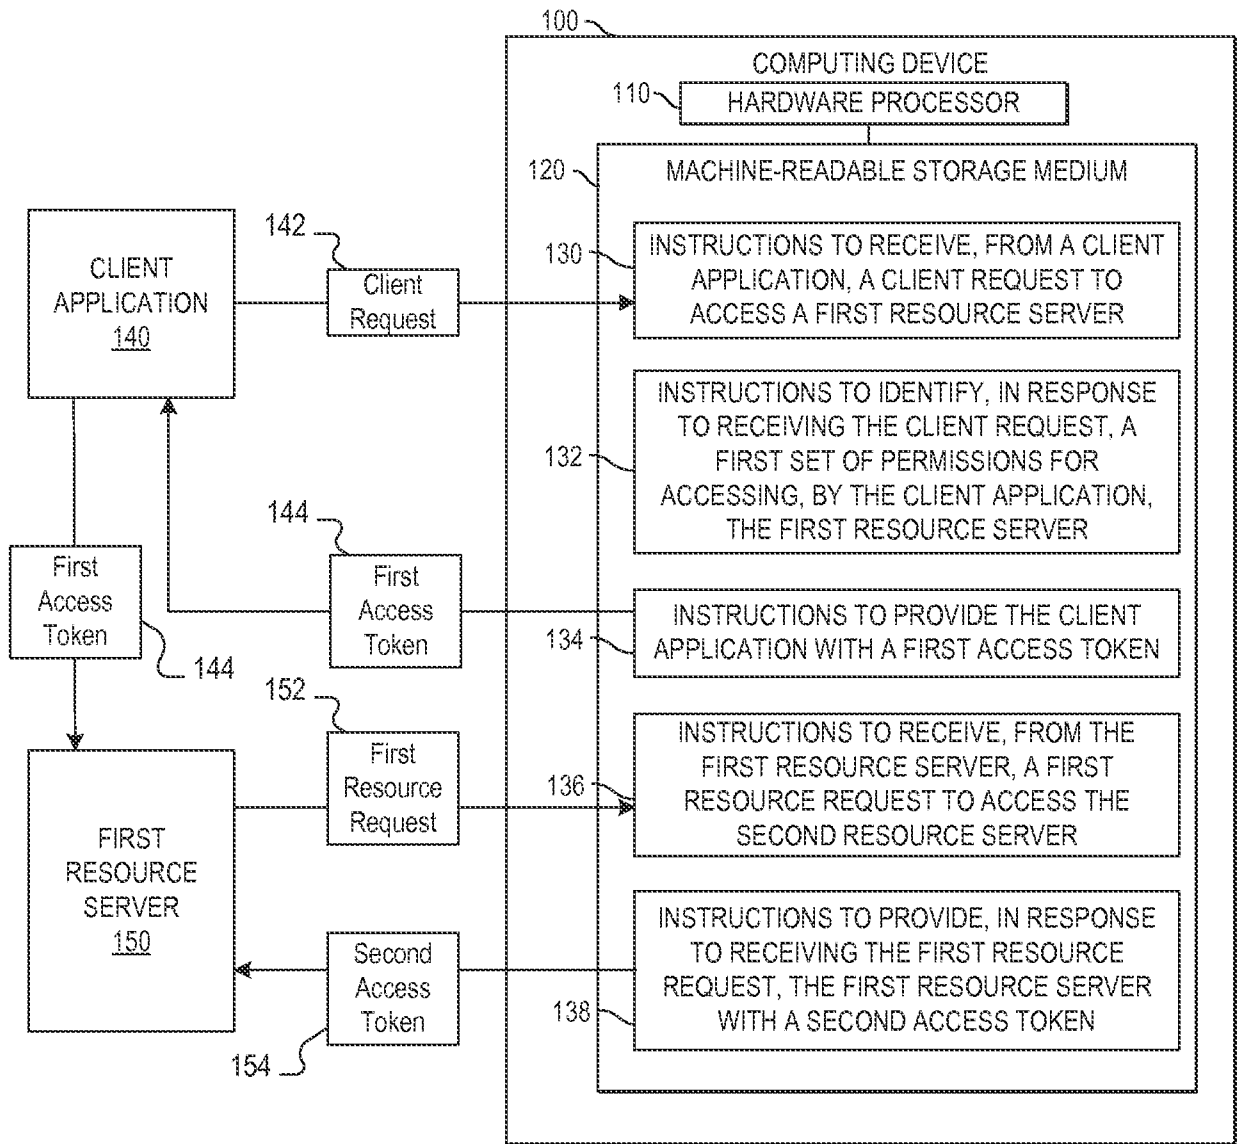


FIG. 1

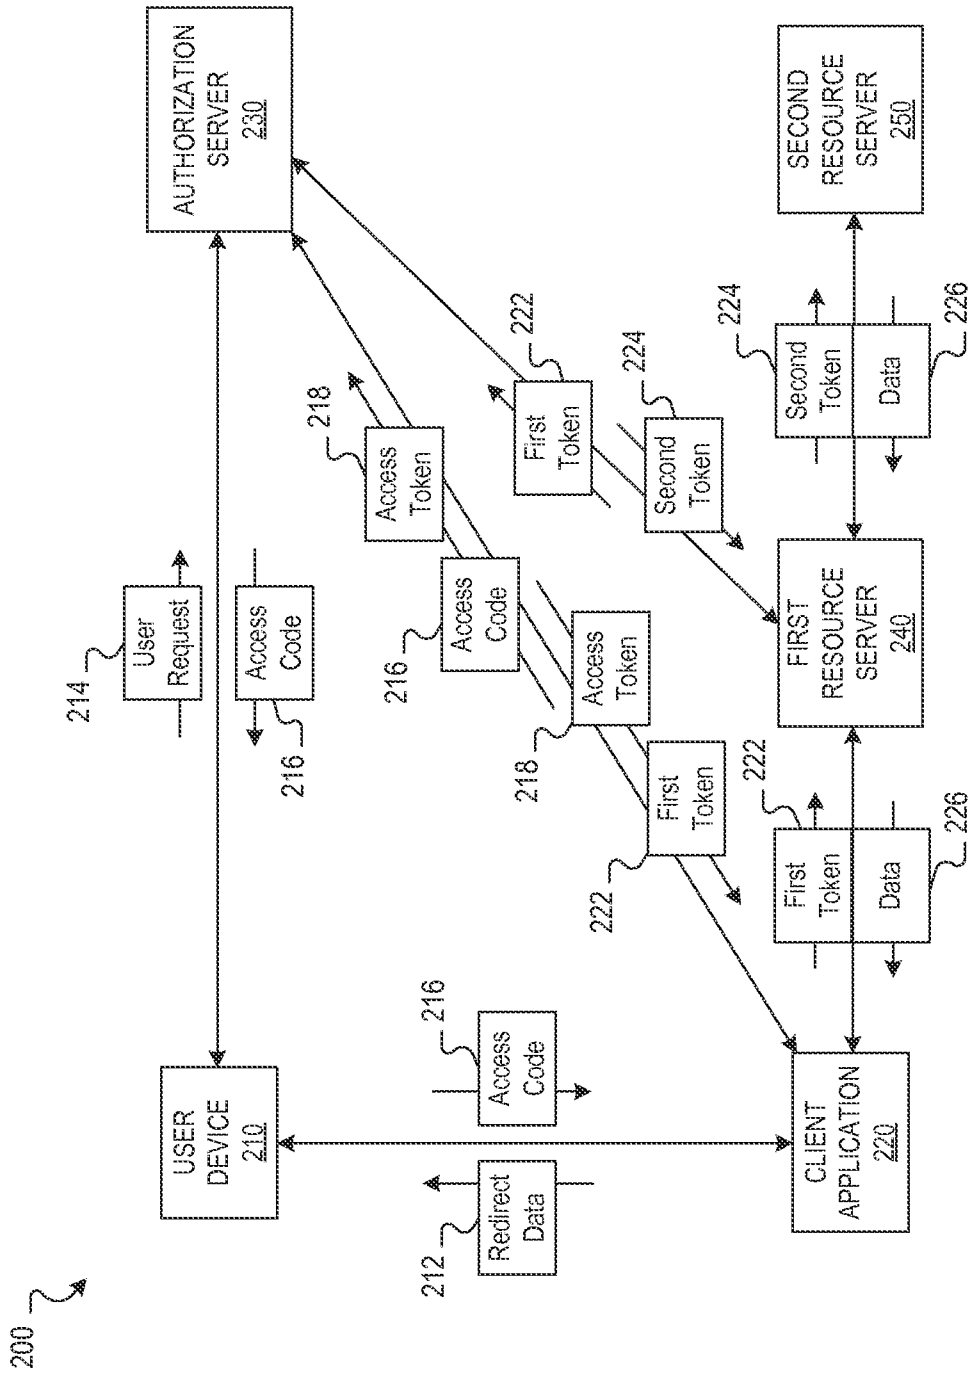


FIG. 2A

280 ↗

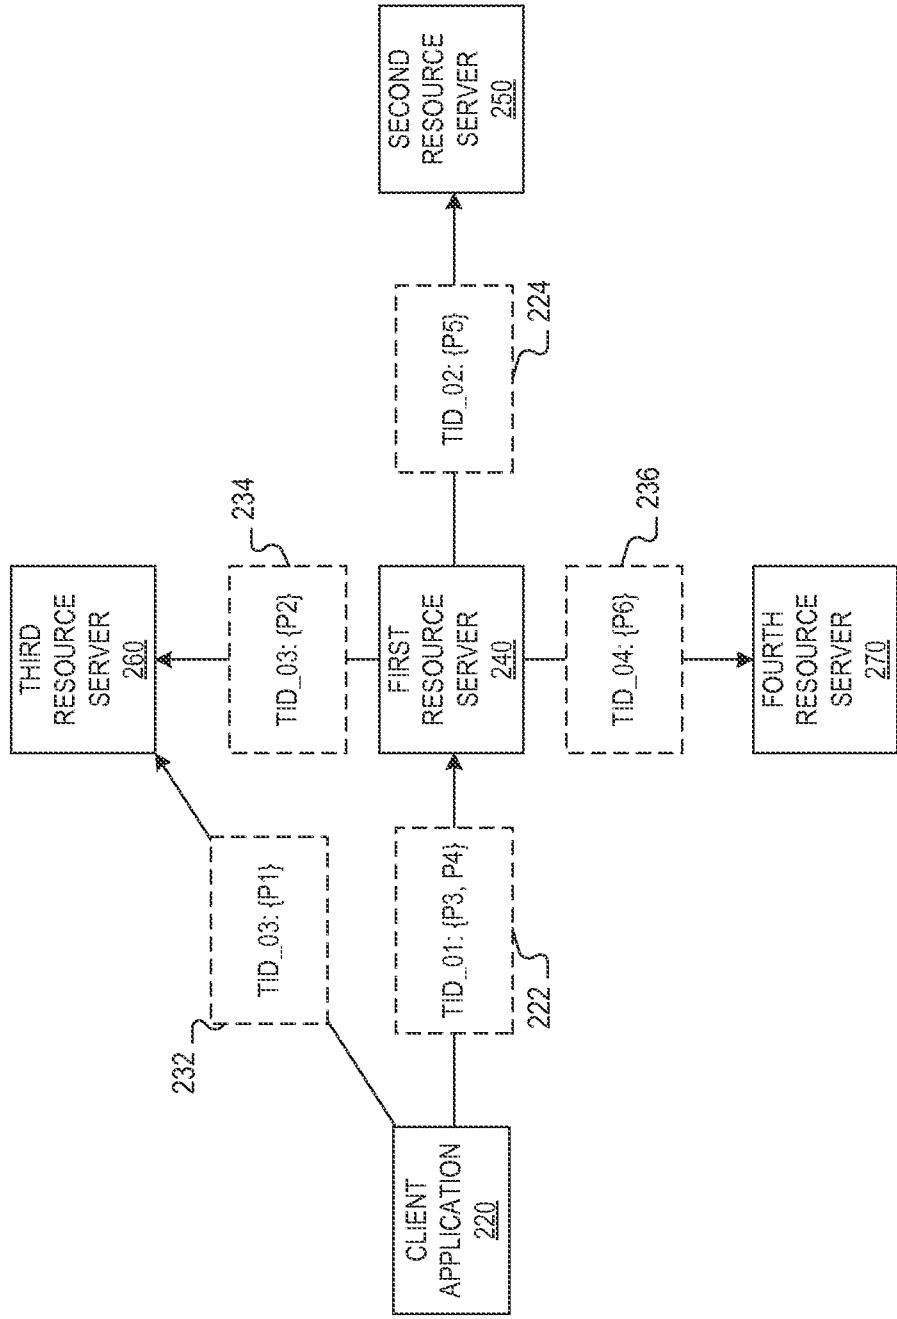
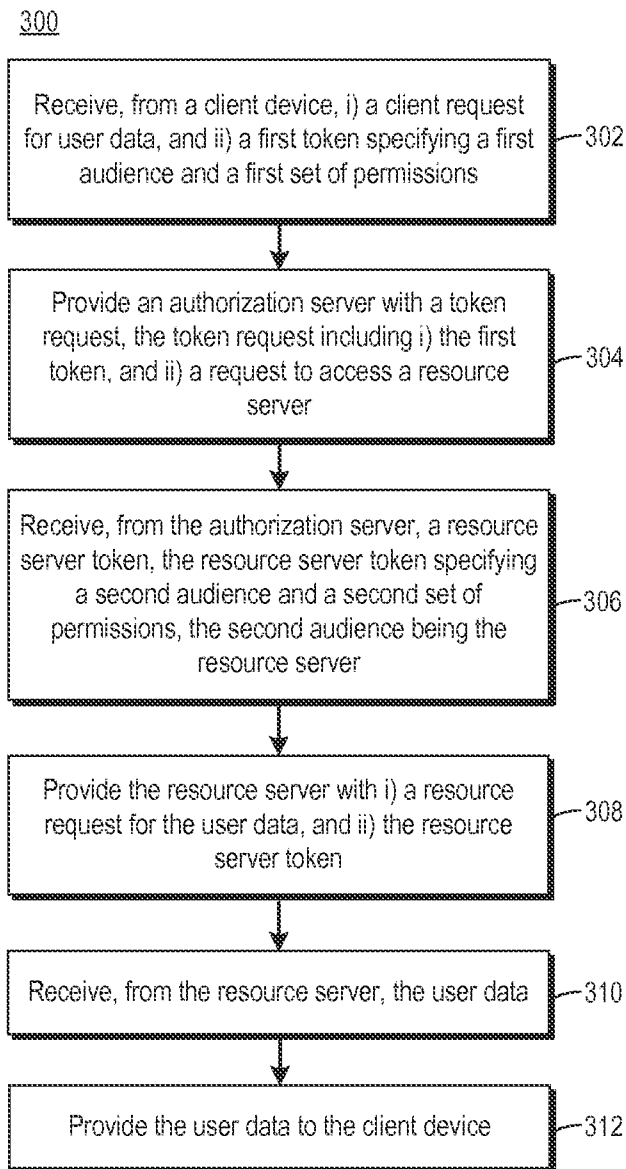


FIG. 2B

**FIG. 3**

A. CLASSIFICATION OF SUBJECT MATTER**H04L 9/32(2006.01)i, H04L 29/06(2006.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
H04L 9/32; G06F 15/16; H04L 9/00; G06F 21/00; H04L 29/06Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: computing device, authorization server, resource server, access token, resource request, permission**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2004-0054791 A1 (KRISHNENDU CHAKRABORTY et al.) 18 March 2004 See paragraph [0048]; claim 1; and figure 4B.	1-15
A	WO 99-60458 A2 (DESKGATE TECHNOLOGIES, INC.) 25 November 1999 See page 13, lines 25-30; claim 1; and figure 5.	1-15
A	D. MITTON et al., `Network Access Server Requirements Next Generation (NASREQNG) NAS Model`, RFC-2881, July 2000 See section 9.	1-15
A	US 2012-0198512 A1 (DENNIS ZHOU et al.) 02 August 2012 See paragraphs [0056]-[0058]; and figure 7.	1-15
A	US 7346930 B1 (KEN BOYDSTUN et al.) 18 March 2008 See column 3, line 43 - column 6, line 19; and figure 1.	1-15

 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

19 February 2016 (19.02.2016)

Date of mailing of the international search report

19 February 2016 (19.02.2016)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

COMMISSIONER

Telephone No. +82-42-481-5916



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/024760

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2004-0054791 A1	18/03/2004	None	
WO 99-60458 A2	25/11/1999	AU 4187999 A CA 2332447 A1 CA 2332447 C CA 2634662 A1 EP 1064592 A2 EP 1064592 A4 US 6389541 B1 WO 99-60458 A3 WO 99-60458 A9	06/12/1999 25/11/1999 07/10/2008 25/11/1999 03/01/2001 29/03/2006 14/05/2002 17/02/2000 14/12/2000
US 2012-0198512 A1	02/08/2012	CN 103404103 A IL 227671 D0 TW 201242312 A TW I467982 B WO 2012-103495 A1	20/11/2013 30/09/2013 16/10/2012 01/01/2015 02/08/2012
US 7346930 B1	18/03/2008	None	