

US010977380B2

(12) **United States Patent**  
**Ferrans et al.**

(10) **Patent No.:** **US 10,977,380 B2**  
(45) **Date of Patent:** **Apr. 13, 2021**

(54) <b>HYBRID ROLE AND ATTRIBUTE BASED ACCESS CONTROL SYSTEM</b>	9,774,586 B1	9/2017	Roche et al.
	9,853,993 B1 *	12/2017	Zhou ..... H04L 63/1425
	10,831,904 B2 *	11/2020	Braksator ..... G06F 16/903
(71) Applicant: <b>Uptake Technologies, Inc.</b> , Chicago, IL (US)	2003/0084069 A1	5/2003	Boreham et al.
	2008/0016580 A1	1/2008	Dixit et al.
	2008/0034438 A1	2/2008	Mireku et al.
	2009/0106207 A1	4/2009	Solheim
(72) Inventors: <b>James Ferrans</b> , Wheaton, IL (US); <b>John Berg</b> , Westmont, IL (US)	2009/0150981 A1	7/2009	Amies et al.
	2009/0288136 A1	11/2009	Chang et al.
	2011/0137946 A1 *	6/2011	Siress ..... G06F 21/6227 707/784
(73) Assignee: <b>Uptake Technologies, Inc.</b> , Chicago, IL (US)	2012/0060207 A1 *	3/2012	Mardikar ..... G06F 21/31 726/4

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 106 days.

(Continued)

**OTHER PUBLICATIONS**

(21) Appl. No.: **15/990,276**

International Search Report and Written Opinion for International Application No. PCT/US2019/033562, dated Aug. 13, 2019.

(22) Filed: **May 25, 2018**

(Continued)

(65) **Prior Publication Data**

US 2019/0362087 A1 Nov. 28, 2019

*Primary Examiner* — Sarah Su

(74) *Attorney, Agent, or Firm* — Lee Sullivan Shea & Smith LLP

(51) **Int. Cl.**

**G06F 21/00** (2013.01)  
**G06F 21/62** (2013.01)  
**G06F 21/31** (2013.01)

(57) **ABSTRACT**

A method may include receiving, from a client device, a request for a resource of a computer system, determining one or more roles of a user associated with the client device, and determining one or more attributes of the user. The method may include determining one or more attributes of the resource and determining an access permission based on the one or more roles of the user and the resource. The method may include generating, by a processing device, a modified access permission by modifying the access permission based on at least one of: the one or more attributes of the user or the one or more attributes of the resource and providing or denying access to the resource of the computer system based on the modified access permission.

(52) **U.S. Cl.**

CPC ..... **G06F 21/6218** (2013.01); **G06F 21/31** (2013.01)

(58) **Field of Classification Search**

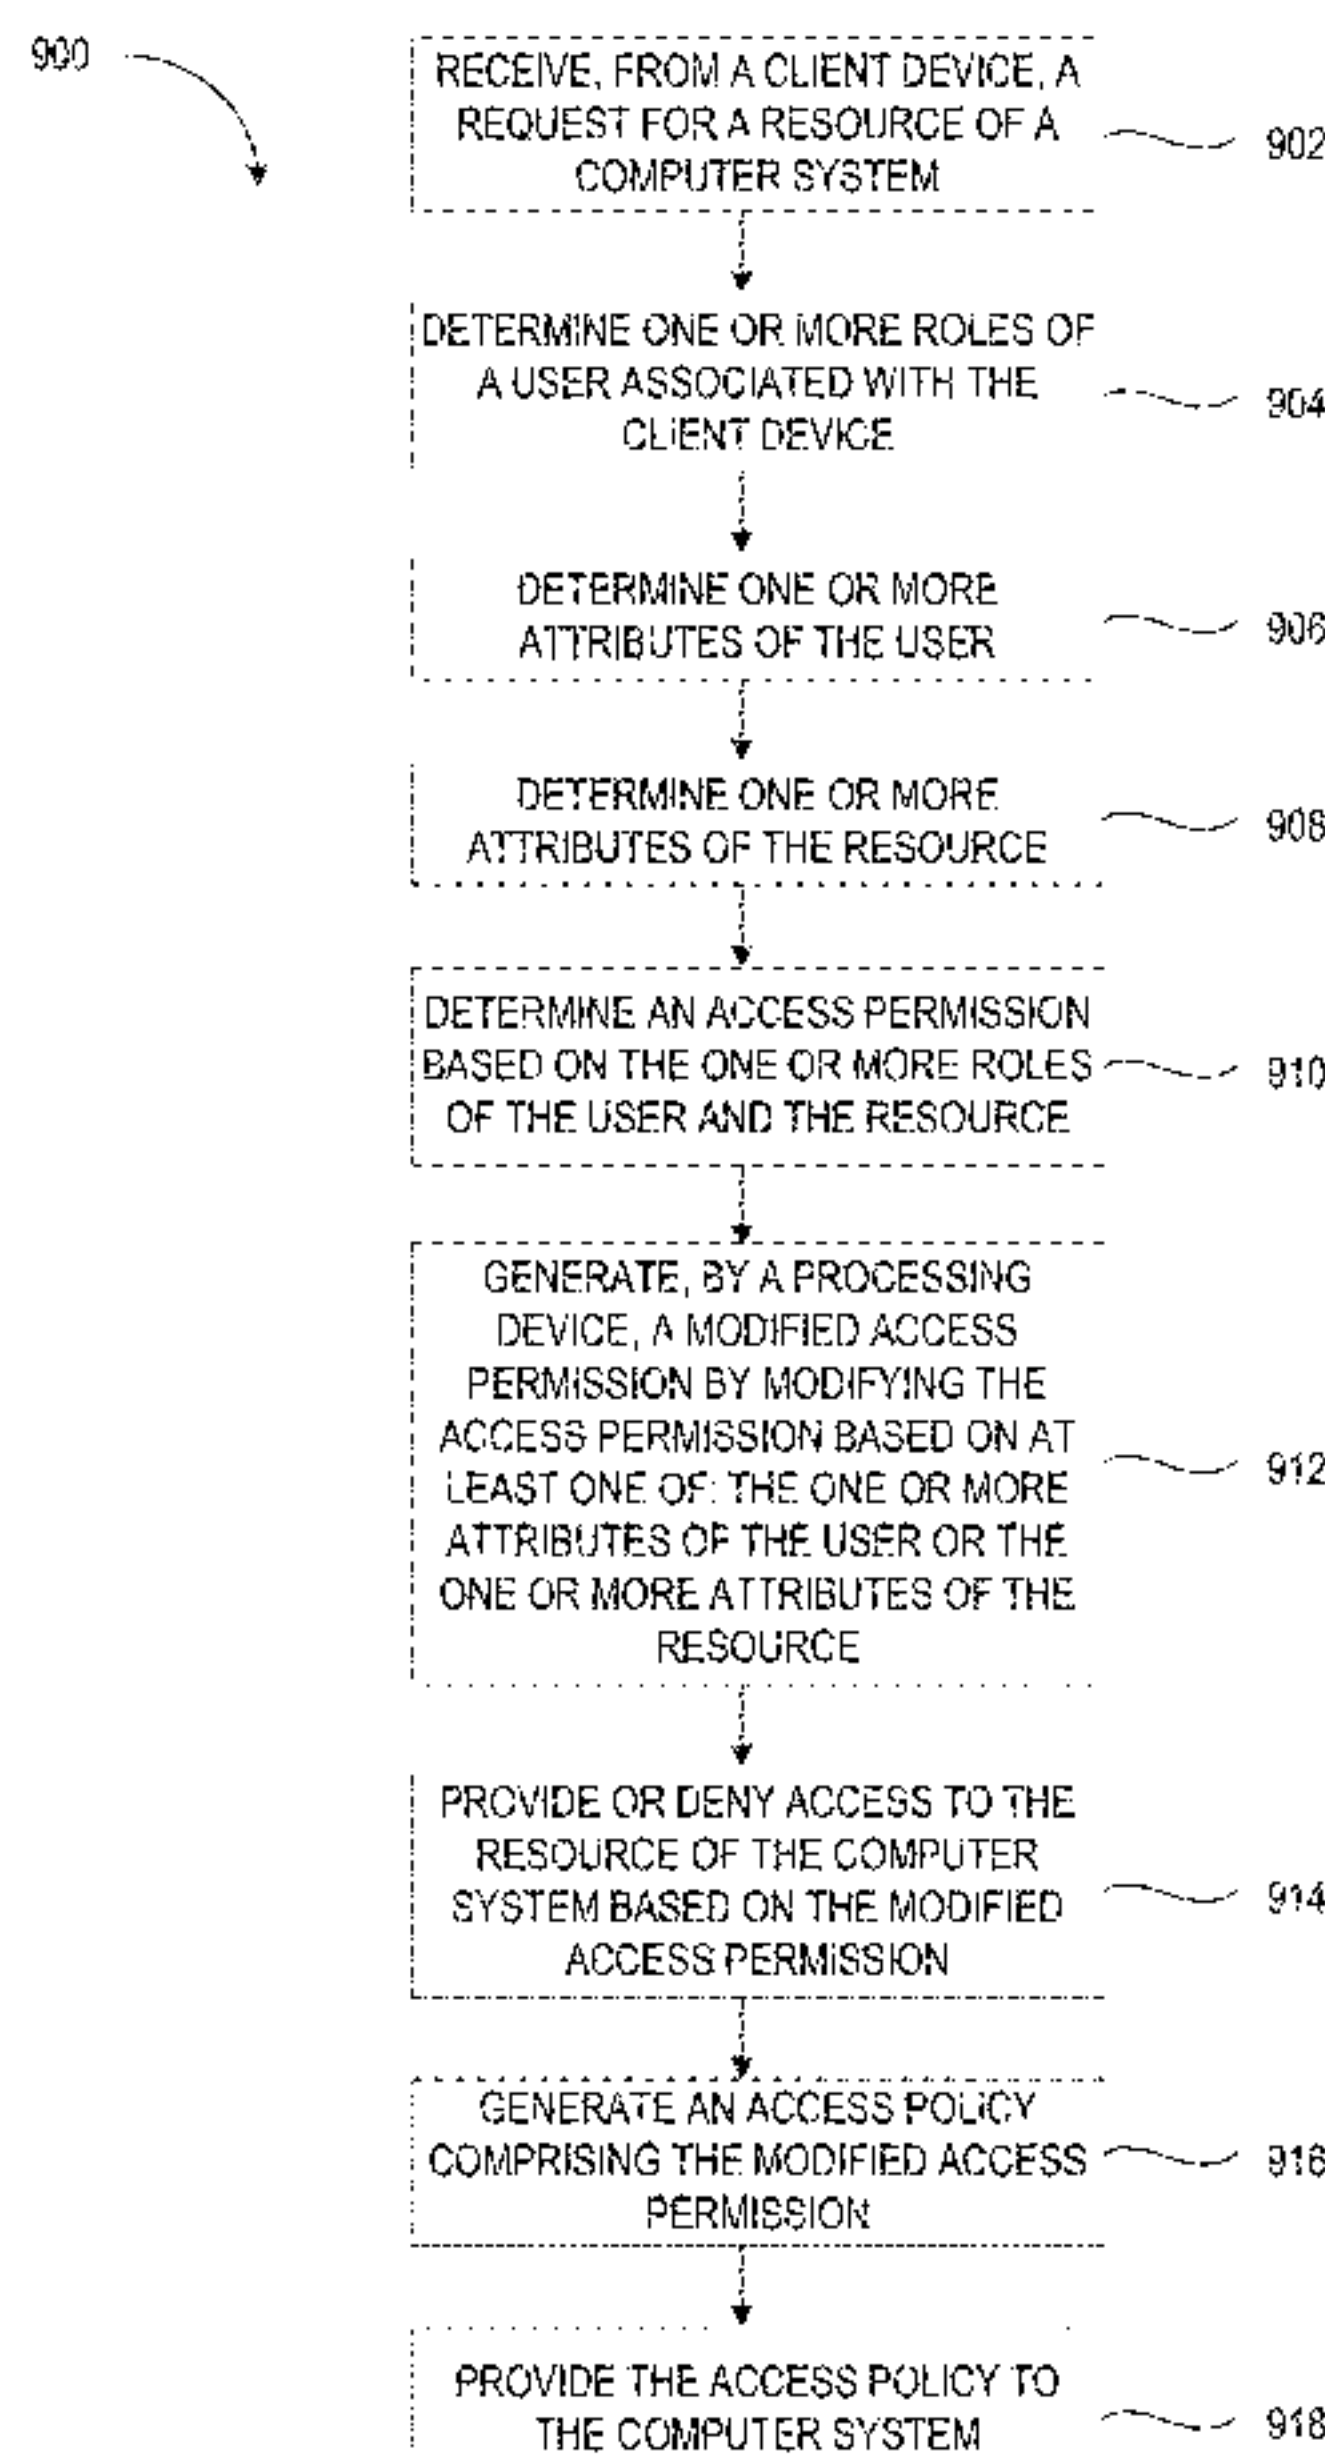
CPC ..... G06F 21/6218; G06F 21/31  
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

7,107,610 B2 \* 9/2006 Lortz ..... H04L 63/08  
709/223  
8,978,122 B1 3/2015 Zolfonoon et al.

**23 Claims, 11 Drawing Sheets**



(56)

References Cited

U.S. PATENT DOCUMENTS

2013/0007891 A1 1/2013 Mogaki  
 2013/0283340 A1 10/2013 Biswas et al.  
 2013/0283350 A1\* 10/2013 Afek ..... G06F 21/6218  
 726/4  
 2014/0115005 A1\* 4/2014 Week ..... G06F 21/6218  
 707/785  
 2014/0282889 A1\* 9/2014 Ishaya ..... H04L 63/0876  
 726/4  
 2015/0101014 A1 4/2015 Giambiagi et al.  
 2016/0072814 A1\* 3/2016 Martinelli ..... H04L 63/101  
 726/1  
 2016/0246983 A1 8/2016 Rissanen et al.  
 2016/0267413 A1\* 9/2016 Liang ..... G06F 21/552  
 2016/0357869 A1 12/2016 Hang  
 2017/0134242 A1\* 5/2017 Ridl ..... G06N 5/025  
 2017/0223138 A1\* 8/2017 Amiri ..... H04L 67/02  
 2017/0339148 A1 11/2017 Syomichev et al.  
 2018/0124031 A1 5/2018 Krishnaprasad  
 2018/0268159 A1\* 9/2018 Yu ..... H04L 9/0819  
 2019/0050476 A1\* 2/2019 Florentino ..... G06F 16/313  
 2019/0306171 A1\* 10/2019 Sisley ..... G06F 21/10  
 2019/0364051 A1\* 11/2019 Ferrans ..... H04L 63/104  
 2020/0120098 A1\* 4/2020 Berg ..... G06F 21/604  
 2020/0186531 A1\* 6/2020 Bugenhagen ..... H04L 67/12

OTHER PUBLICATIONS

Chen Zhen et al. Query rewriting in spatial database for access control, Computing for Geospatial Research & Applications ACM

2 Penn Plaze Suite 701 New York 10121-0701 USA May 23, 2011 (May 23, 2011) pp. 1-4.

Calero et al. Toward a Multi-Tenancy Authorization System for Cloud Services. Security & Privacy, IEEE, Nov. 1, 2010, pp. 1-7. [online]. [retrieved on Mar. 20, 2020]. Retrieved from the Internet <URL:https://www.researchgate.net/publication/224202017\_Tow.

Tang et al. A multi-tenant RBAC model for collaborative cloud services. 2013 Eleventh Annual Conference on Privacy, Security and Trust, IEEE, Jul. 10, 2013, pp. 229-238. [online], [retrieved on Mar. 20, 2020]. Retrieved from the Internet <URL:https://www.profsandhu.com/cs6393\_s16/Tang-et-al-2013b.pdf>.

Ranjbar, Alireza. Domain Isolation in a Multi-Tenant Software-Defined Network. Aalto University, Finland, Dec. 2015, 10 pages. [online], [retrieved on Mar. 20, 2020]. Retrieved from the Internet <URL: http://www.cs.hut.fi/~antikaml/domainIsolation.pdf>.

Ranjbar, Alireza. Domain Isolation in a Multi-Tenant Software-Defined Network. AALTO University: School of Electrical Engineering. Thesis submitted for degree of Master of Science in Technology. Apr. 15, 2015 107 pages.[online], [retrieved on Mar. 20, 2020]. Retrieved from the Internet <URL:https://core.ac.uk/download/pdf/80715216.pdf>.

Tang et al. Multi-Tenancy Authorization Models for Collaborative Cloud Services. Collaboration Technologies and Systems, 2013 International Conference On, IEEE, May 20, 2013, pp. 132-138. [online], [retrieved on Mar. 20, 2020]. Retrieved from the Internet <URL:https://ieeexplore.ieee.org/document/6567218>.

Calero et al. Toward a Multi-Tenancy Authorization System for Cloud Services. Security & Privacy, IEEE, Nov. 1, 2010, pp. 48-55. [online]. [retrieved on Mar. 20, 2020]. Retrieved from the Internet <URL:https://s3.amazonaws.com/academia.edu.documents/45820841/msp.2010.19420160520-7567-107jwuv.pdf>.

\* cited by examiner

101




FIG. 1

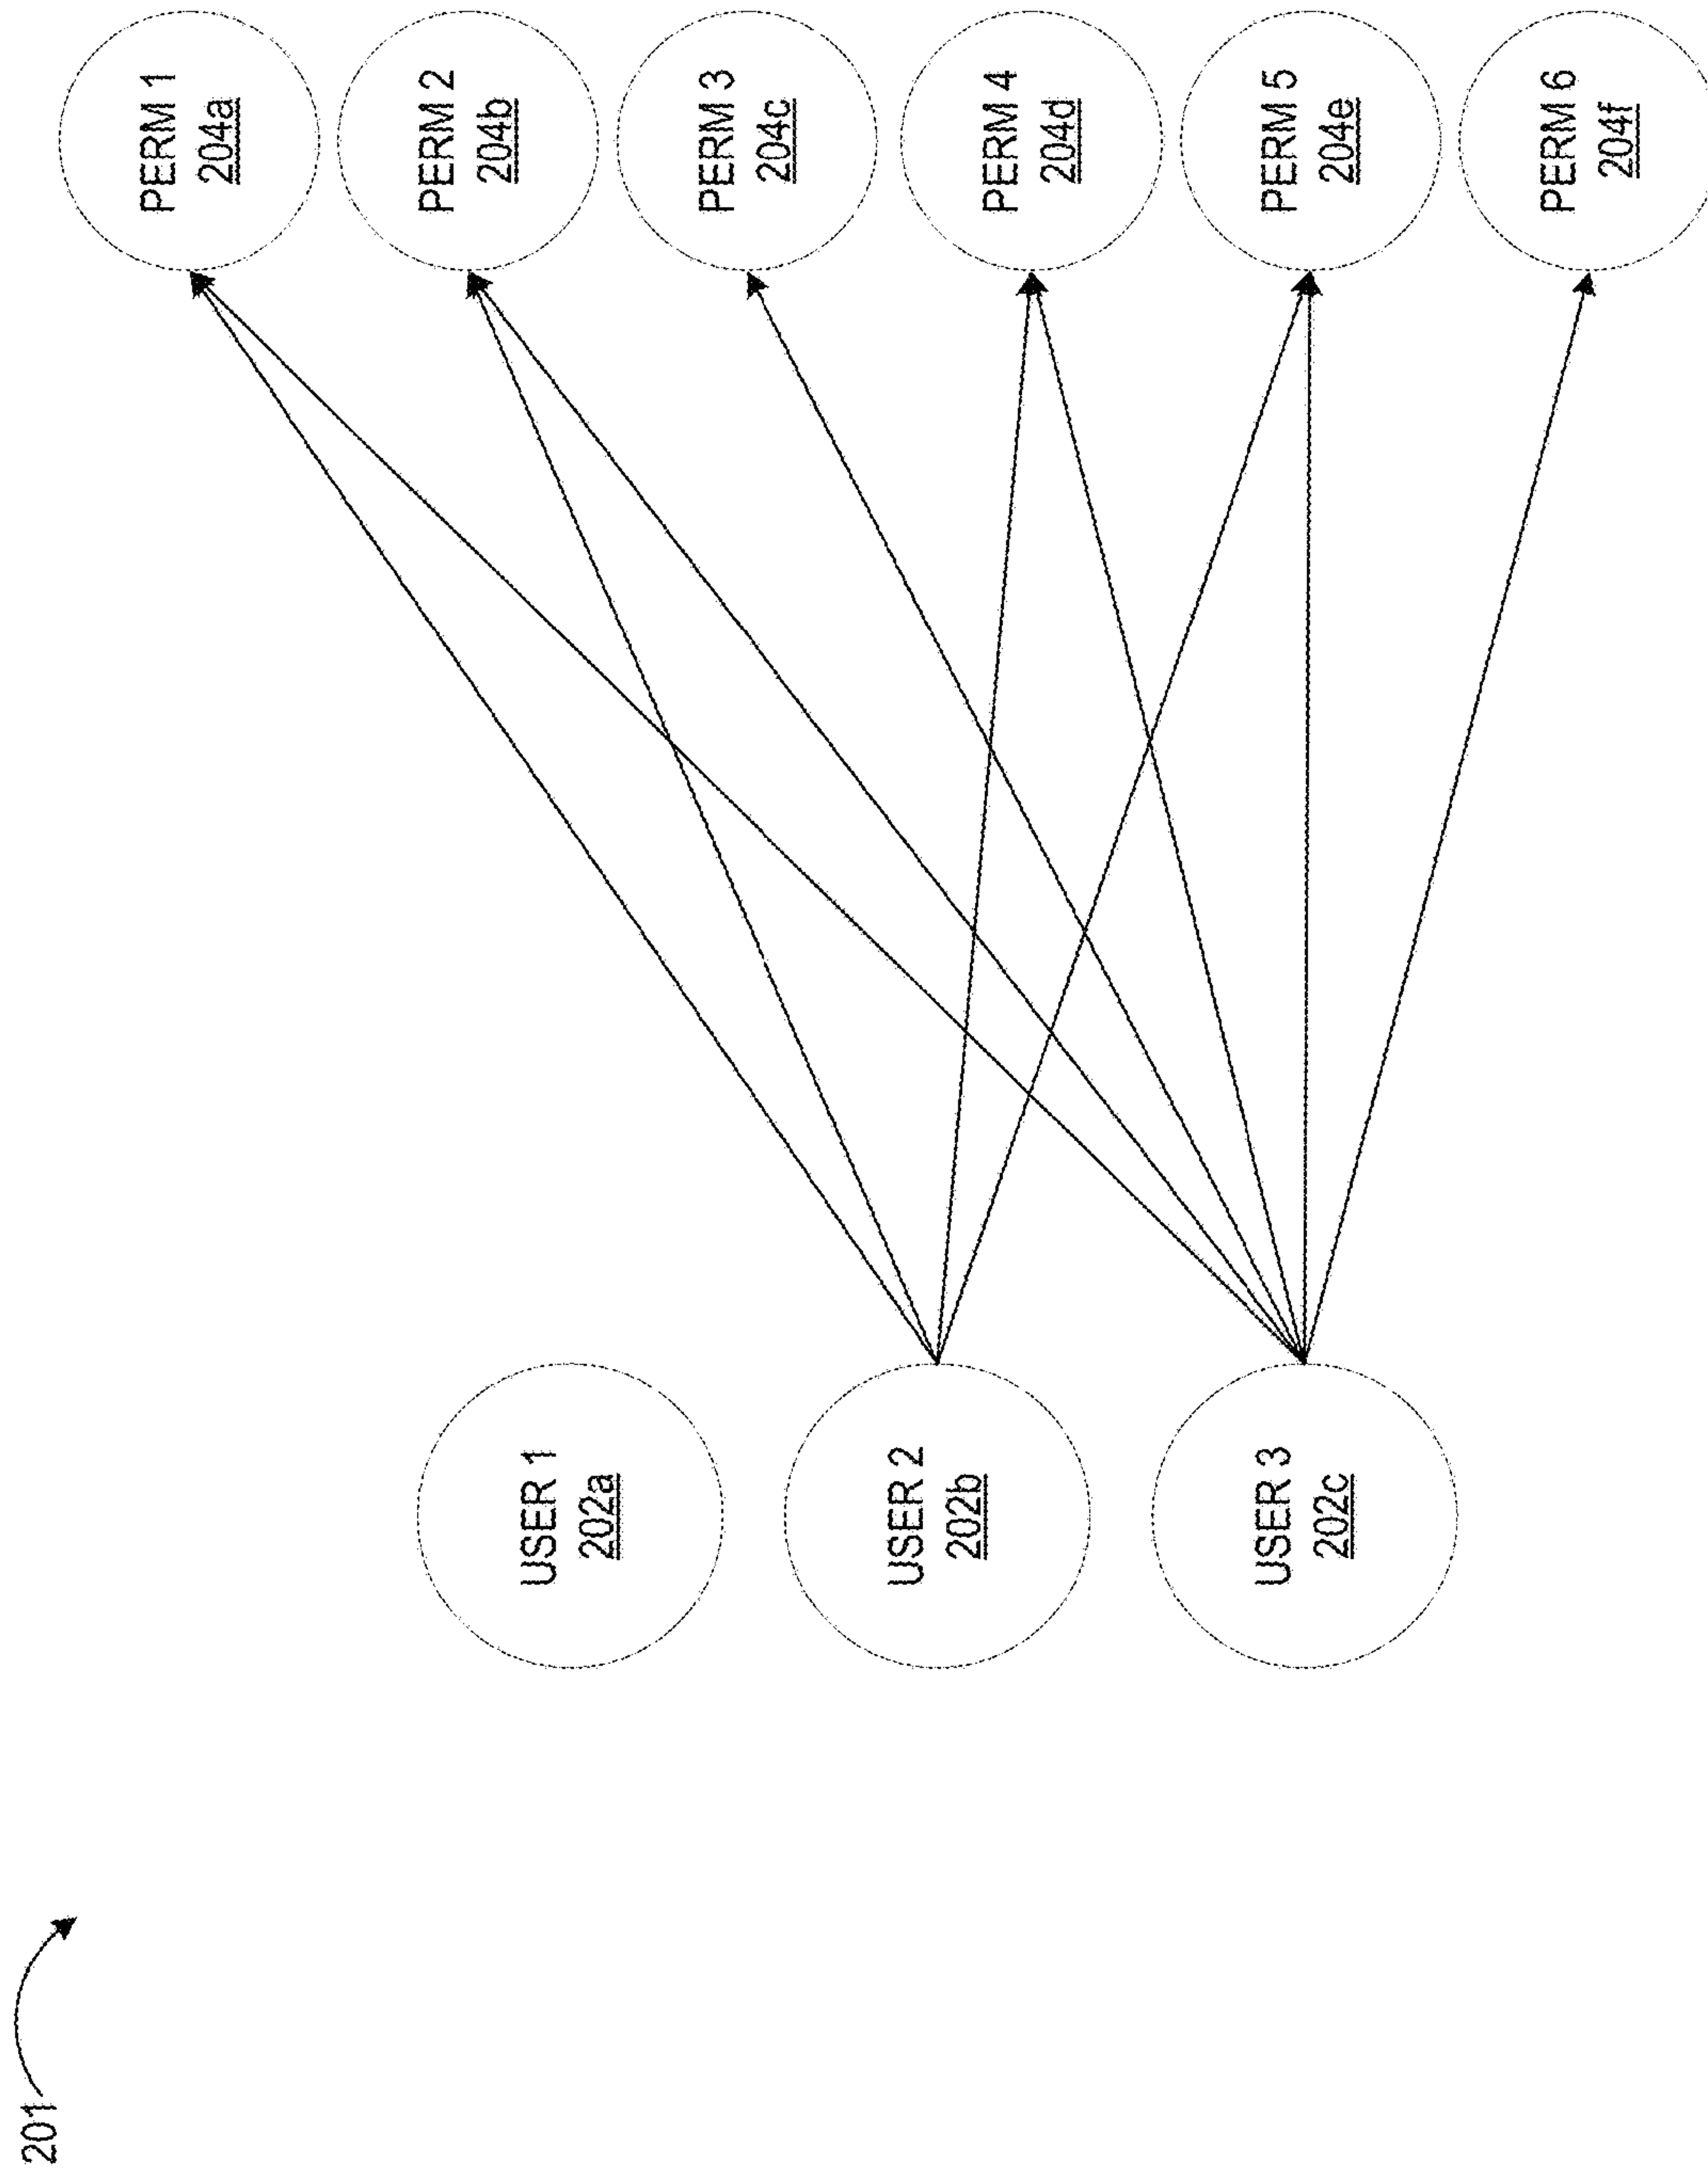


FIG. 2



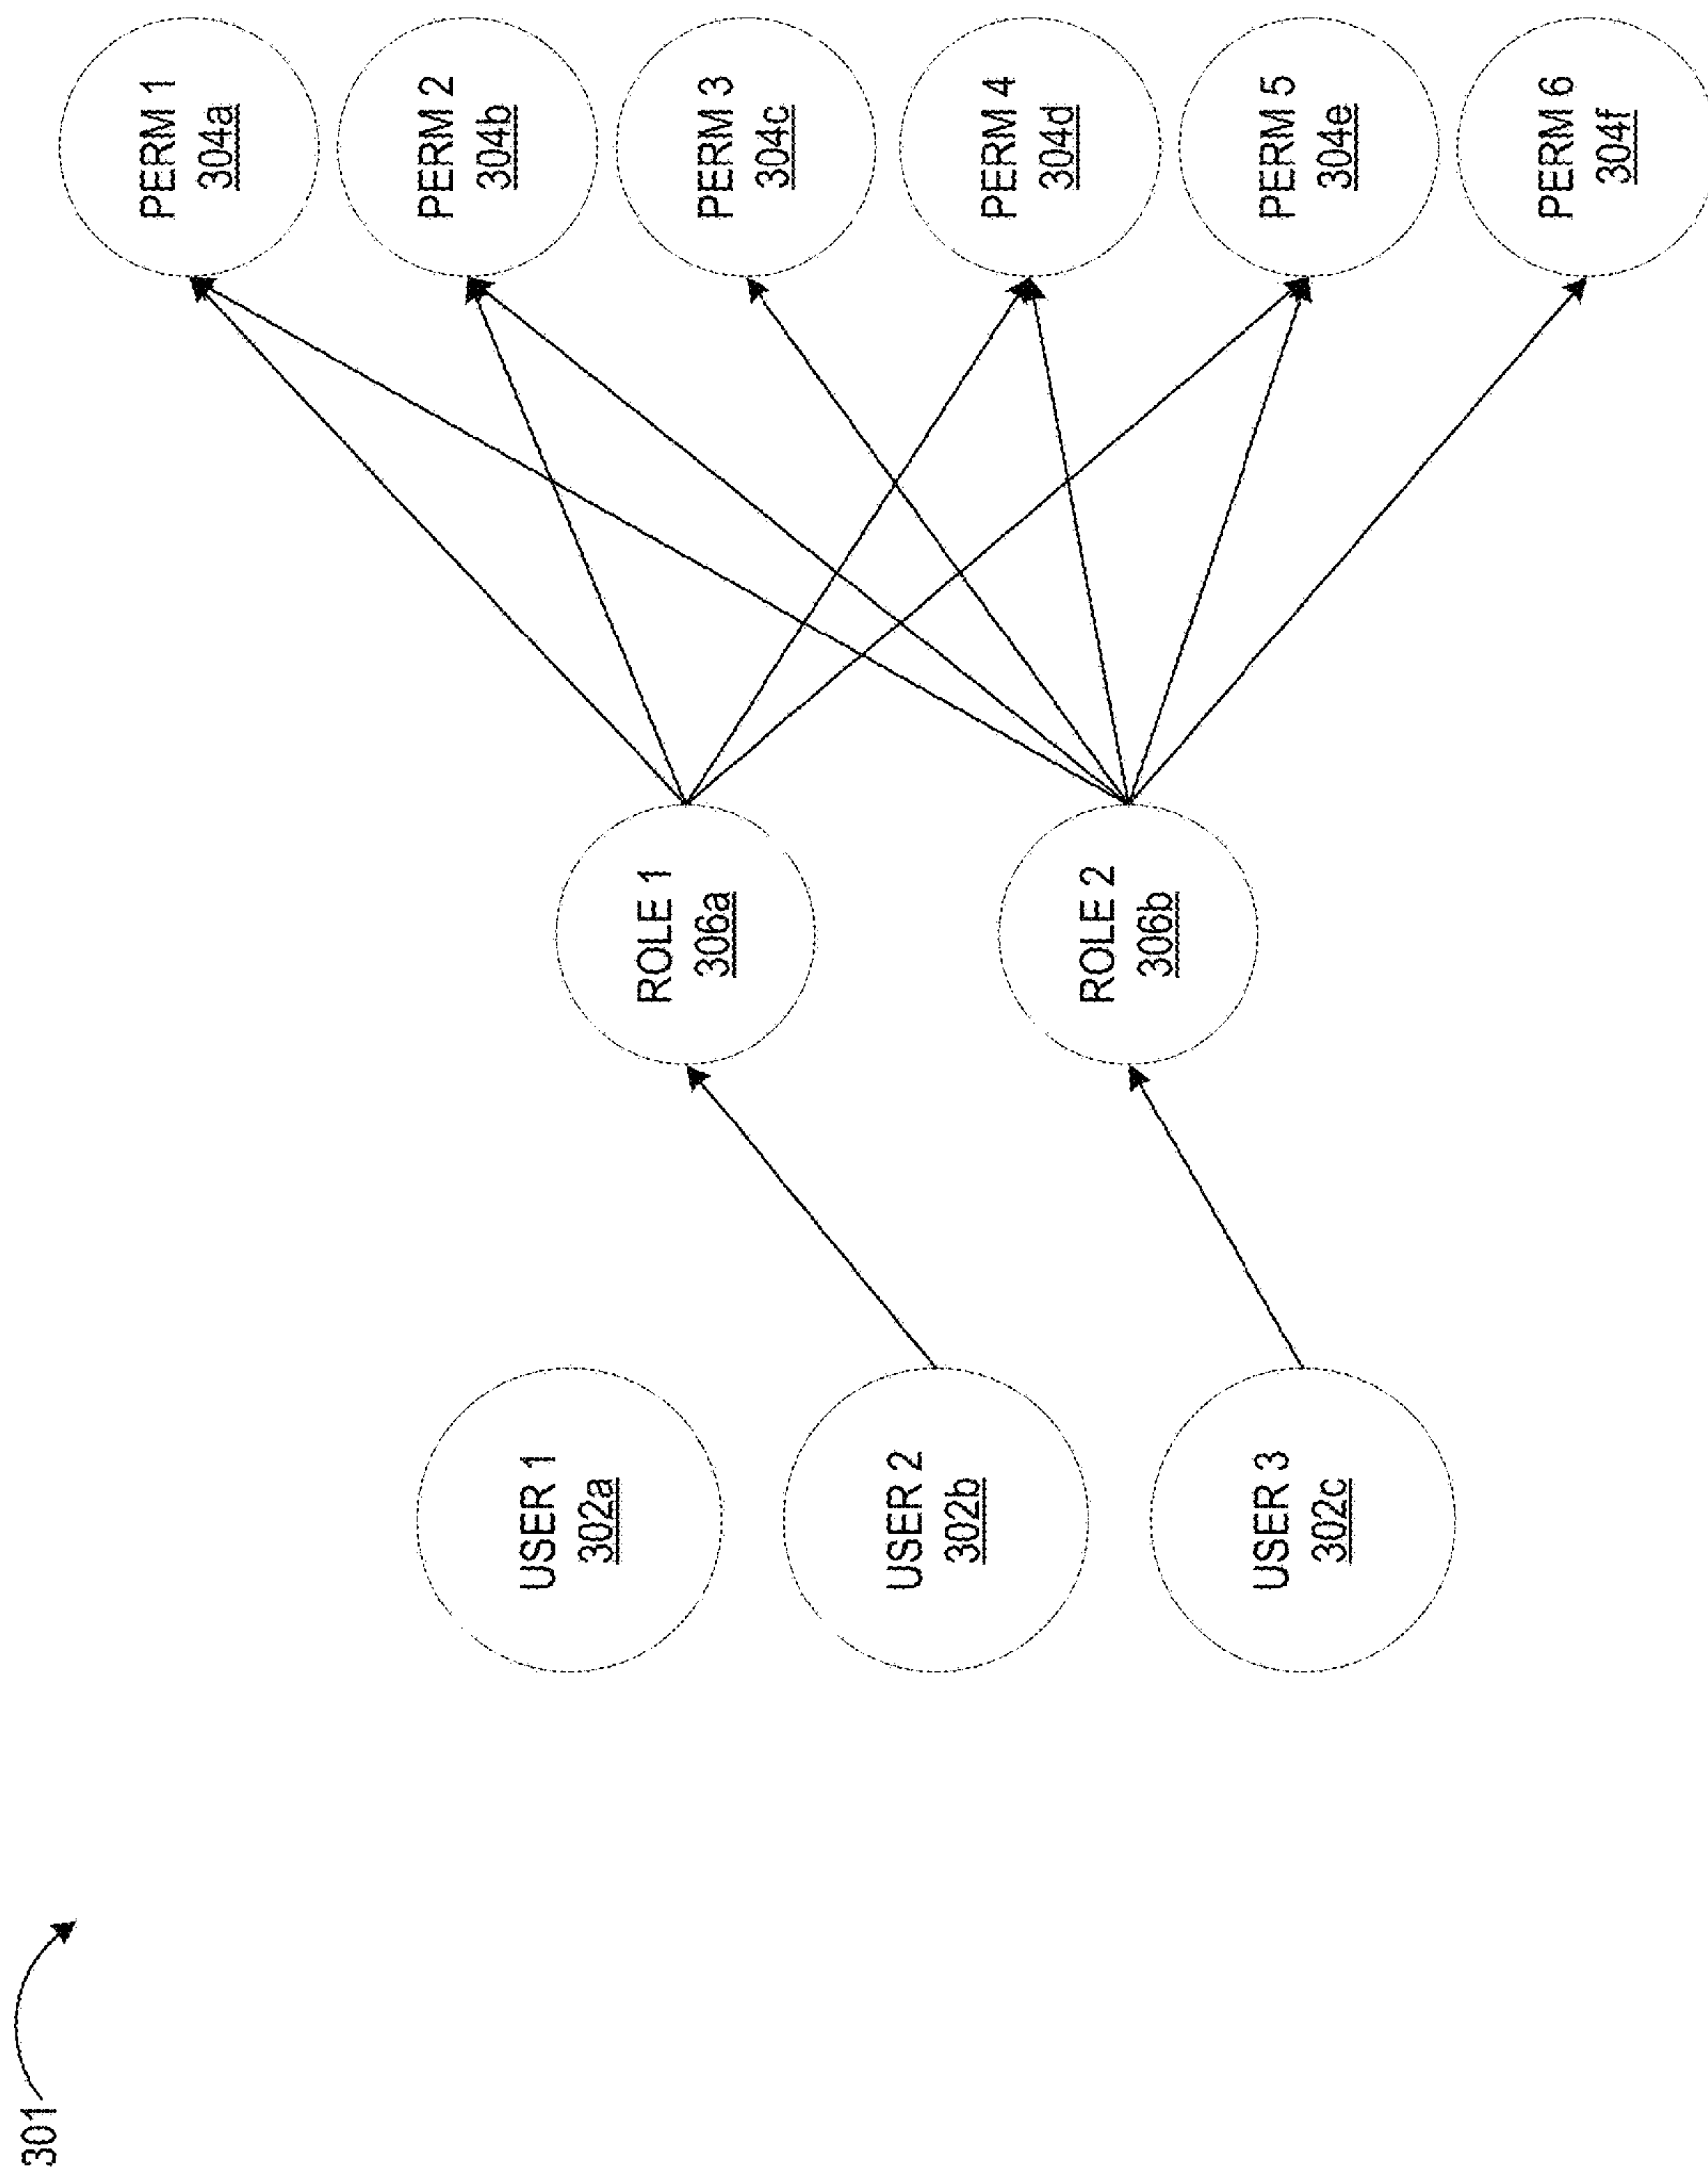


FIG. 3

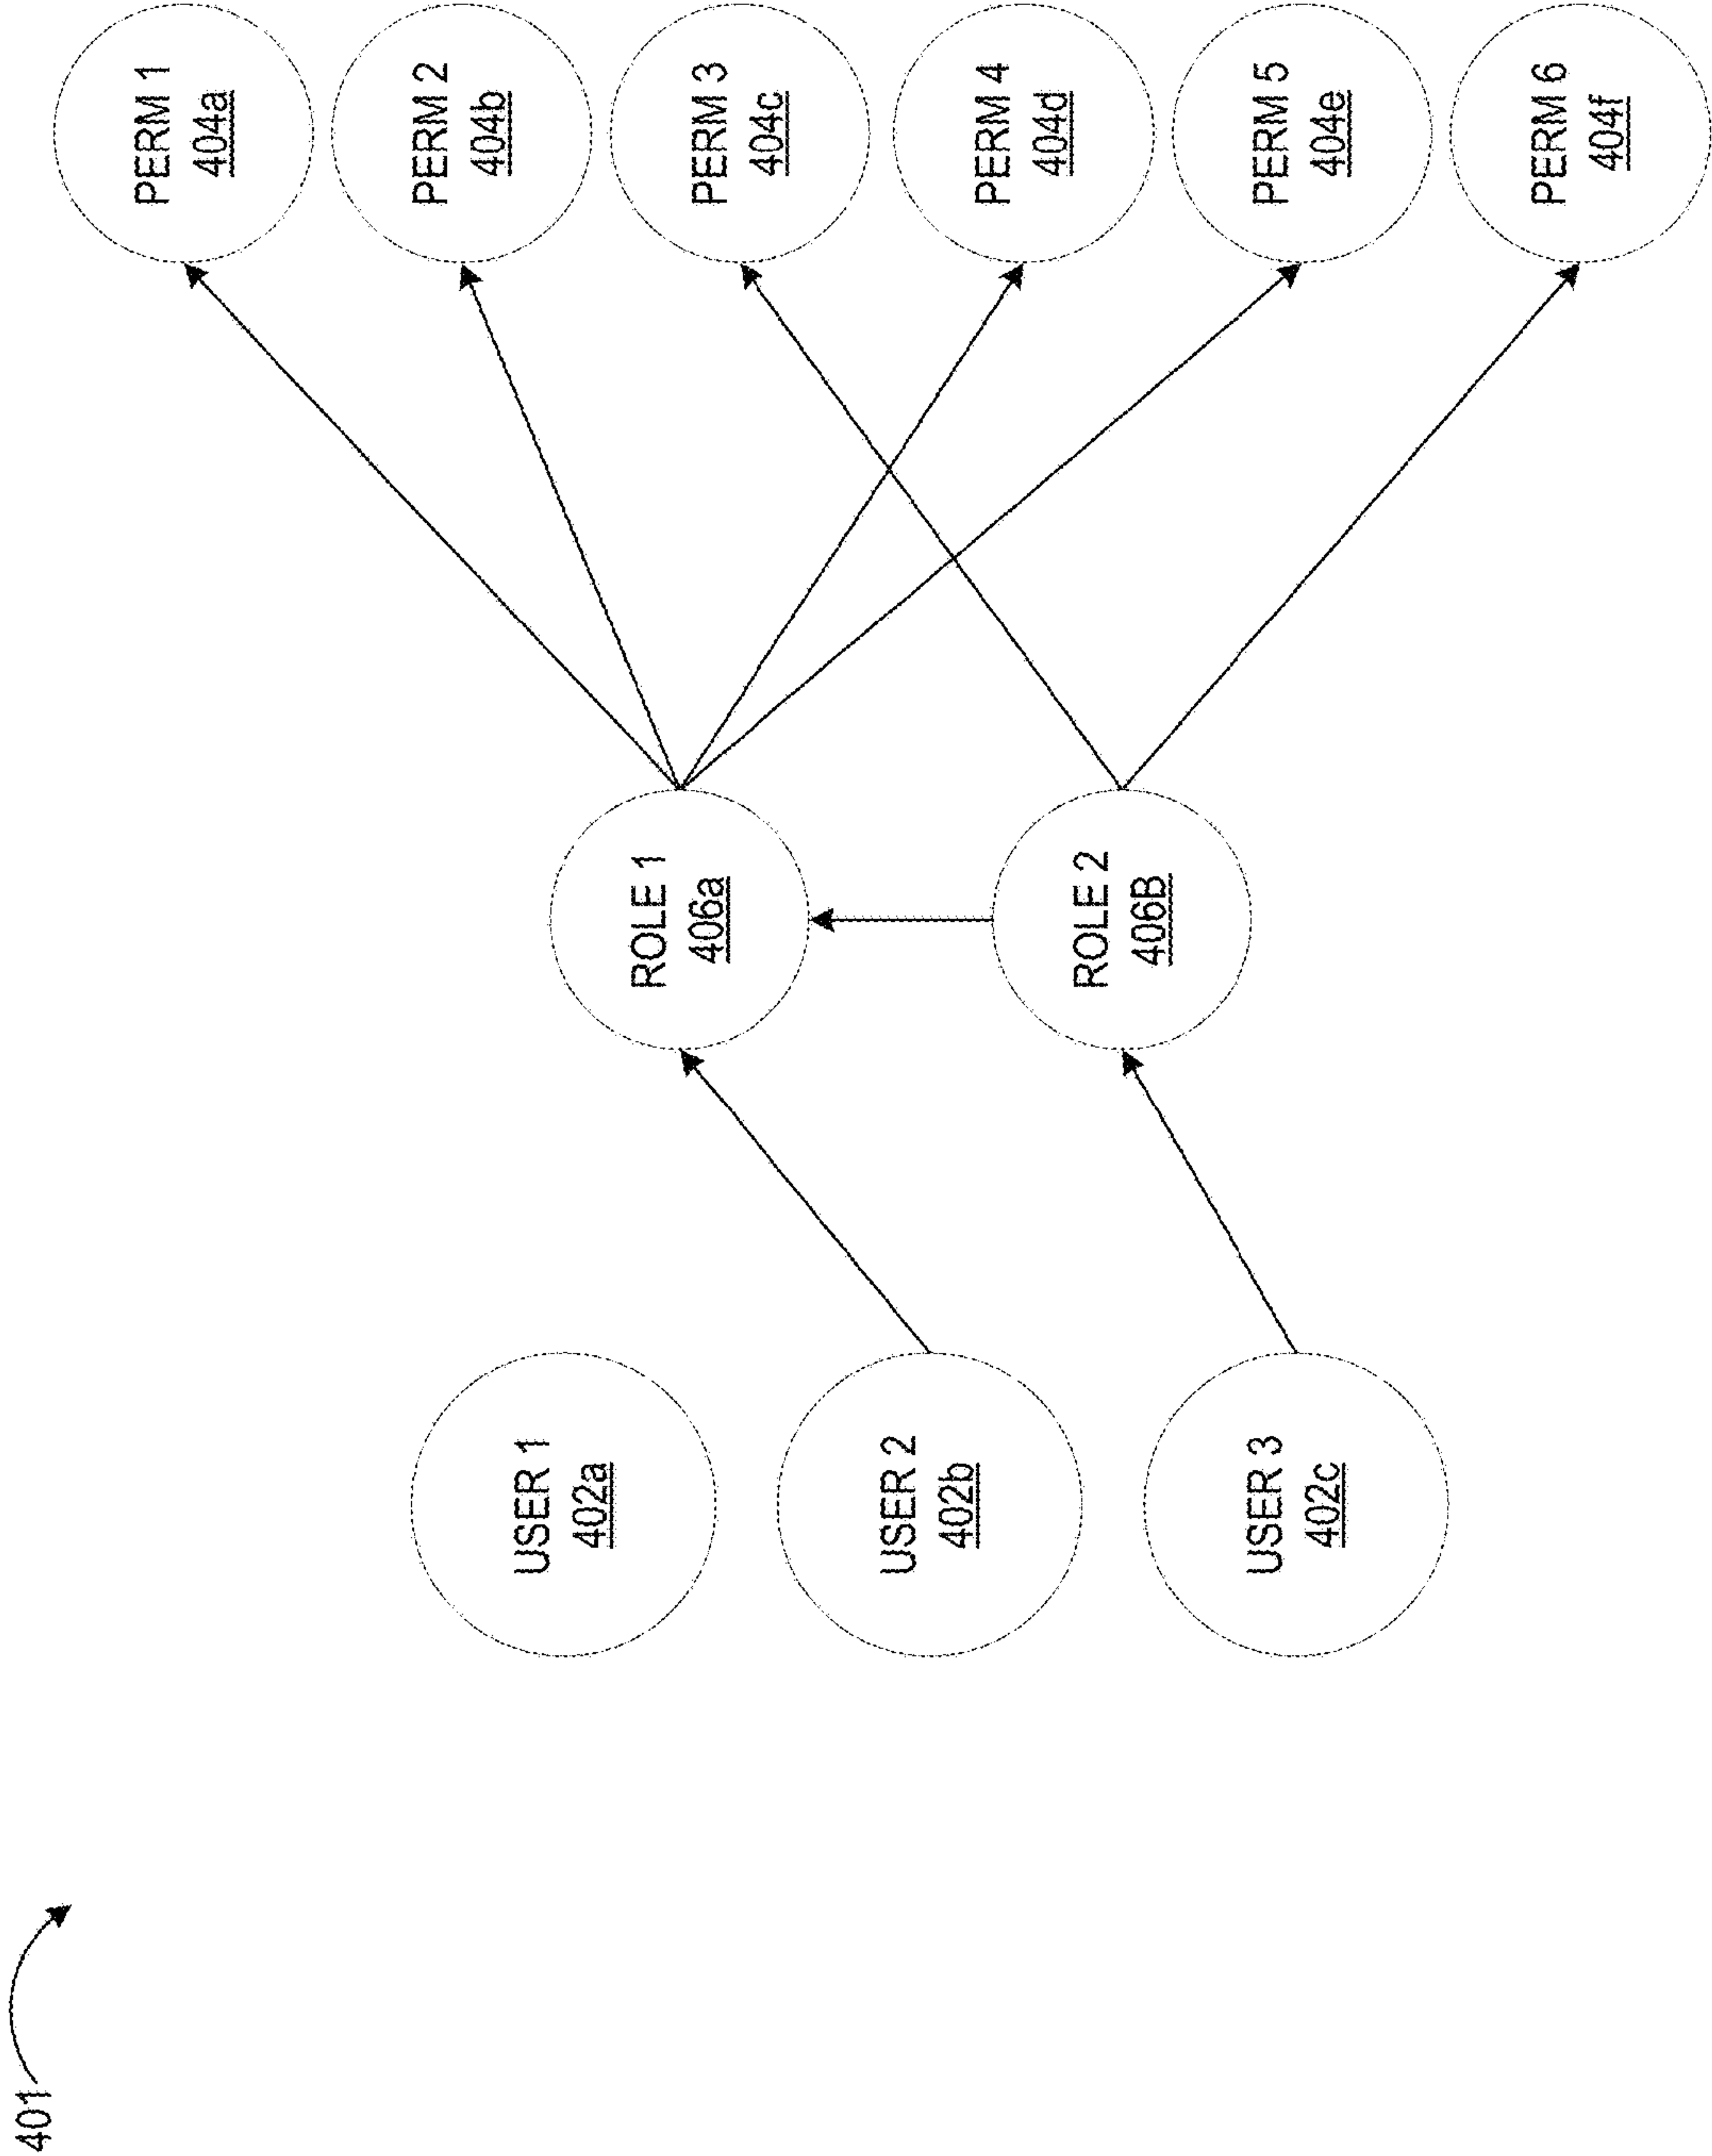


FIG. 4

500




FIG. 5

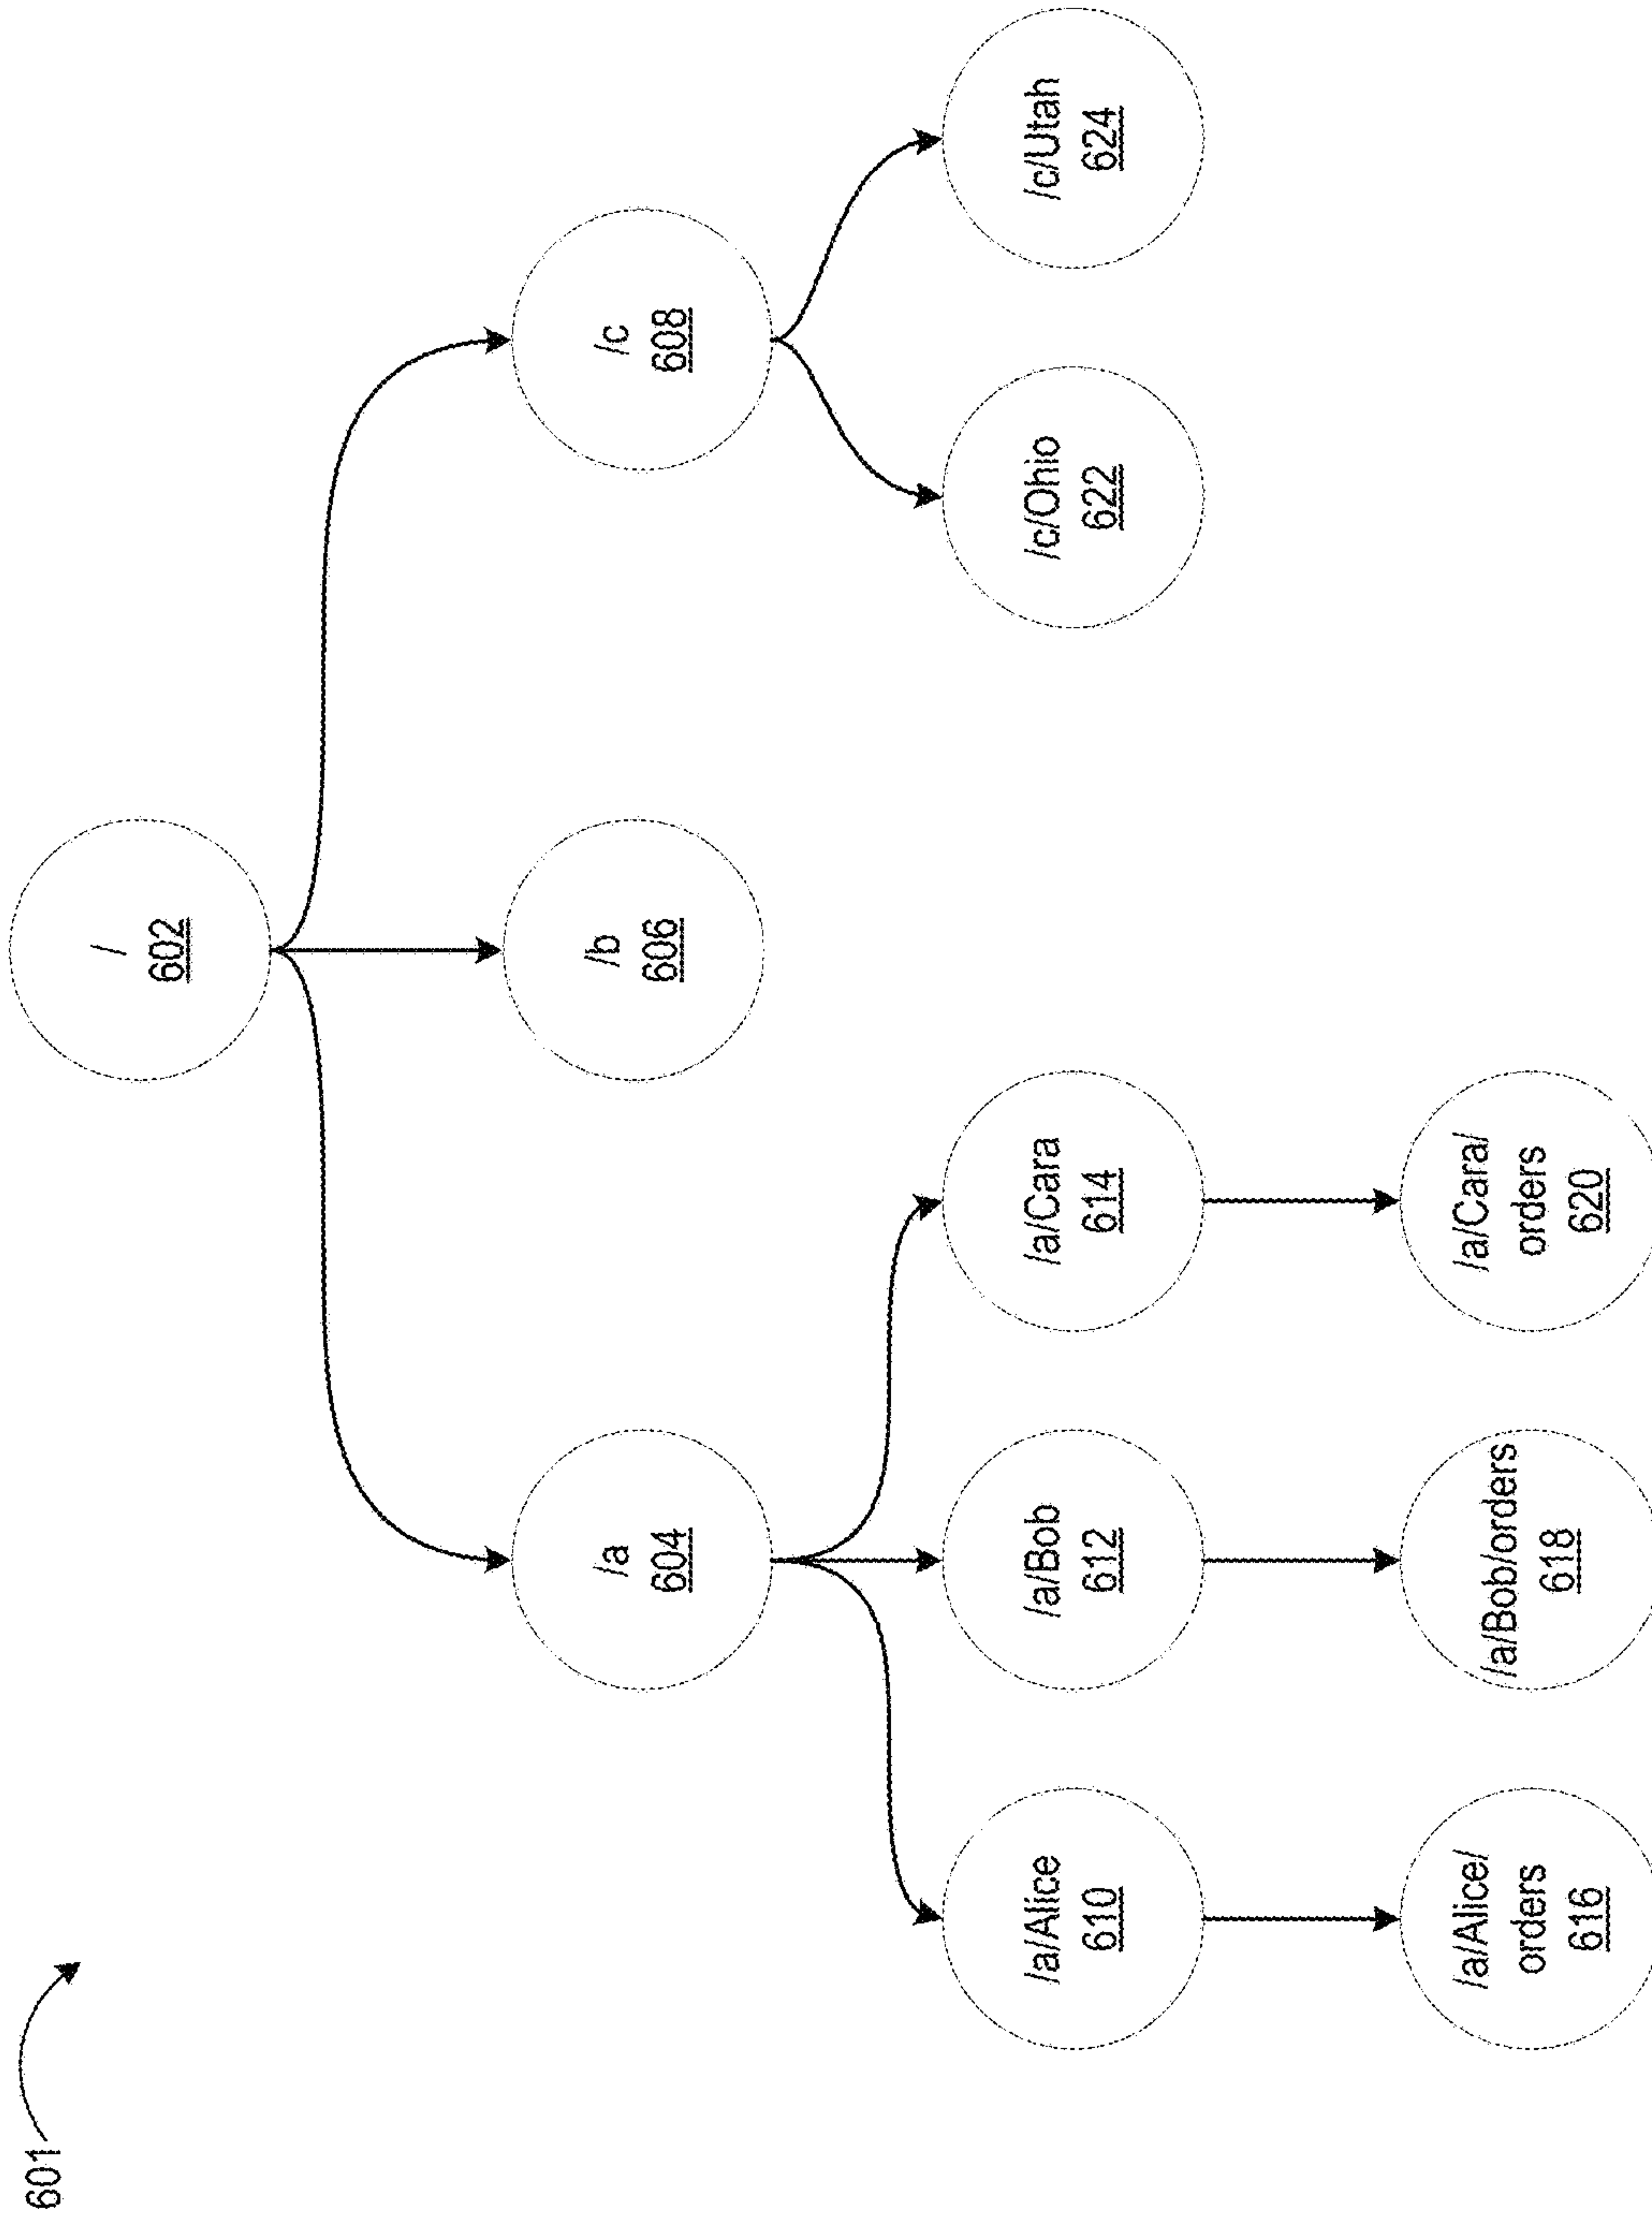


FIG. 6






FIG. 7A



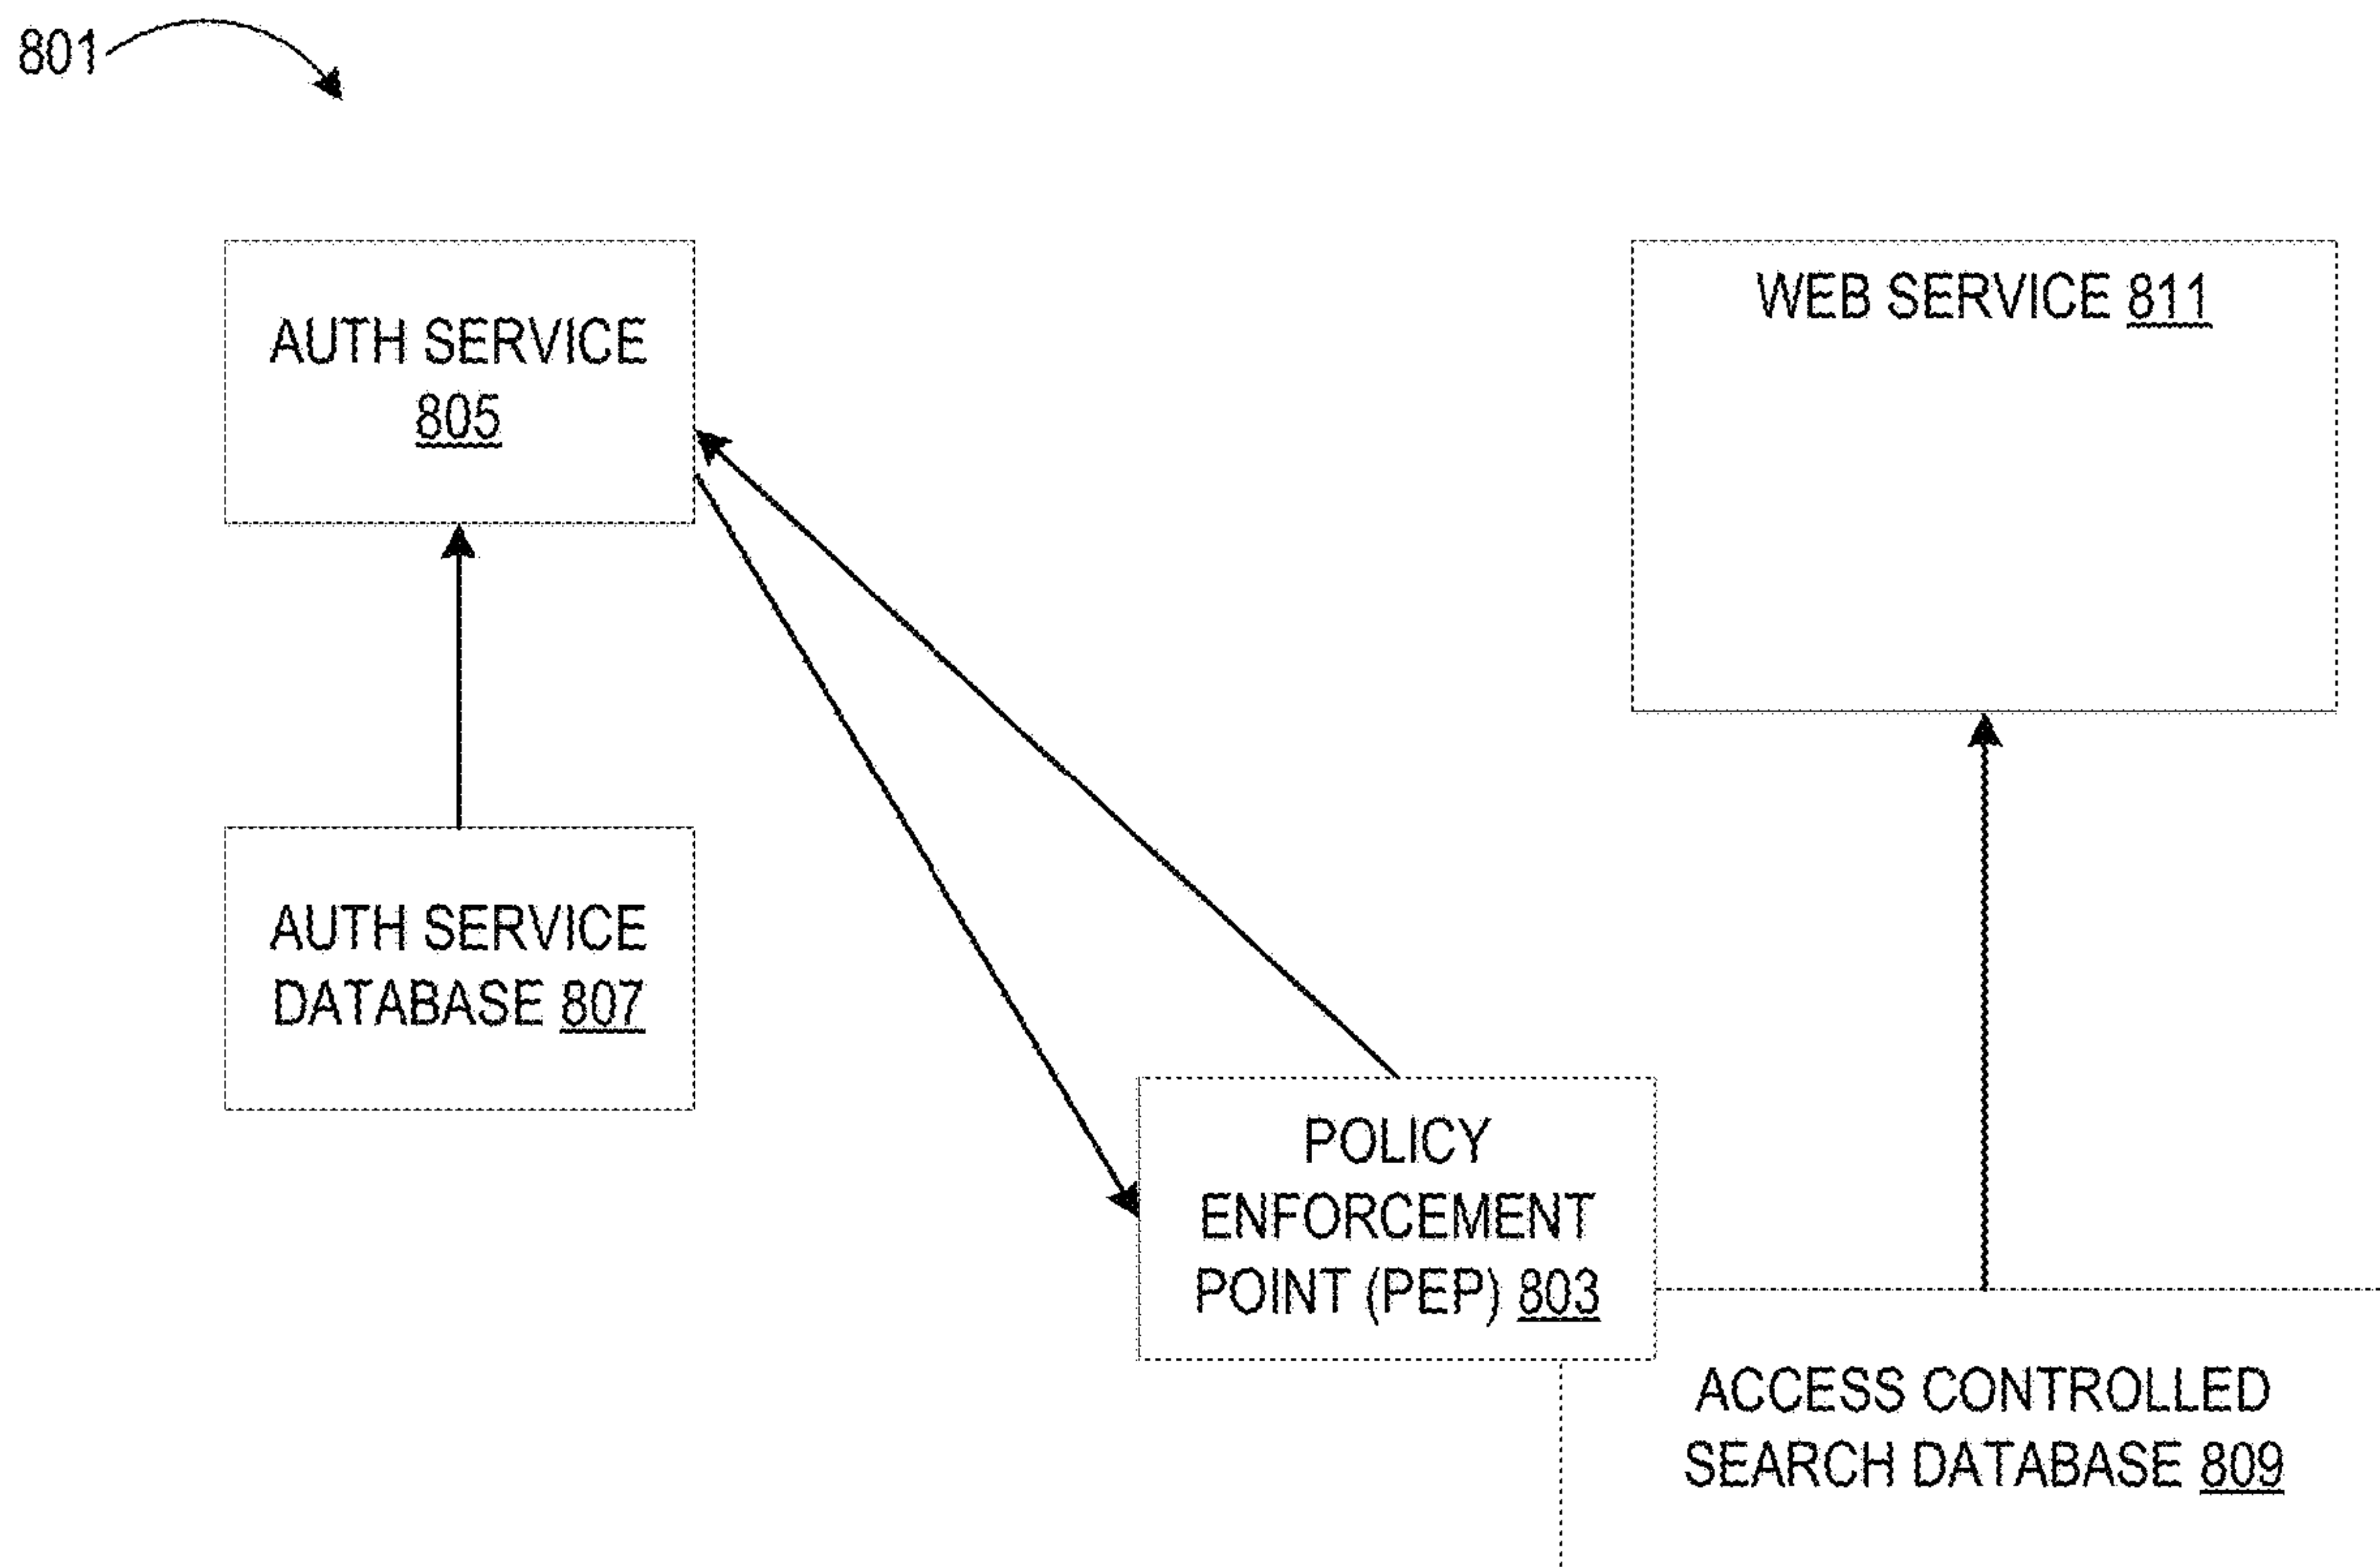


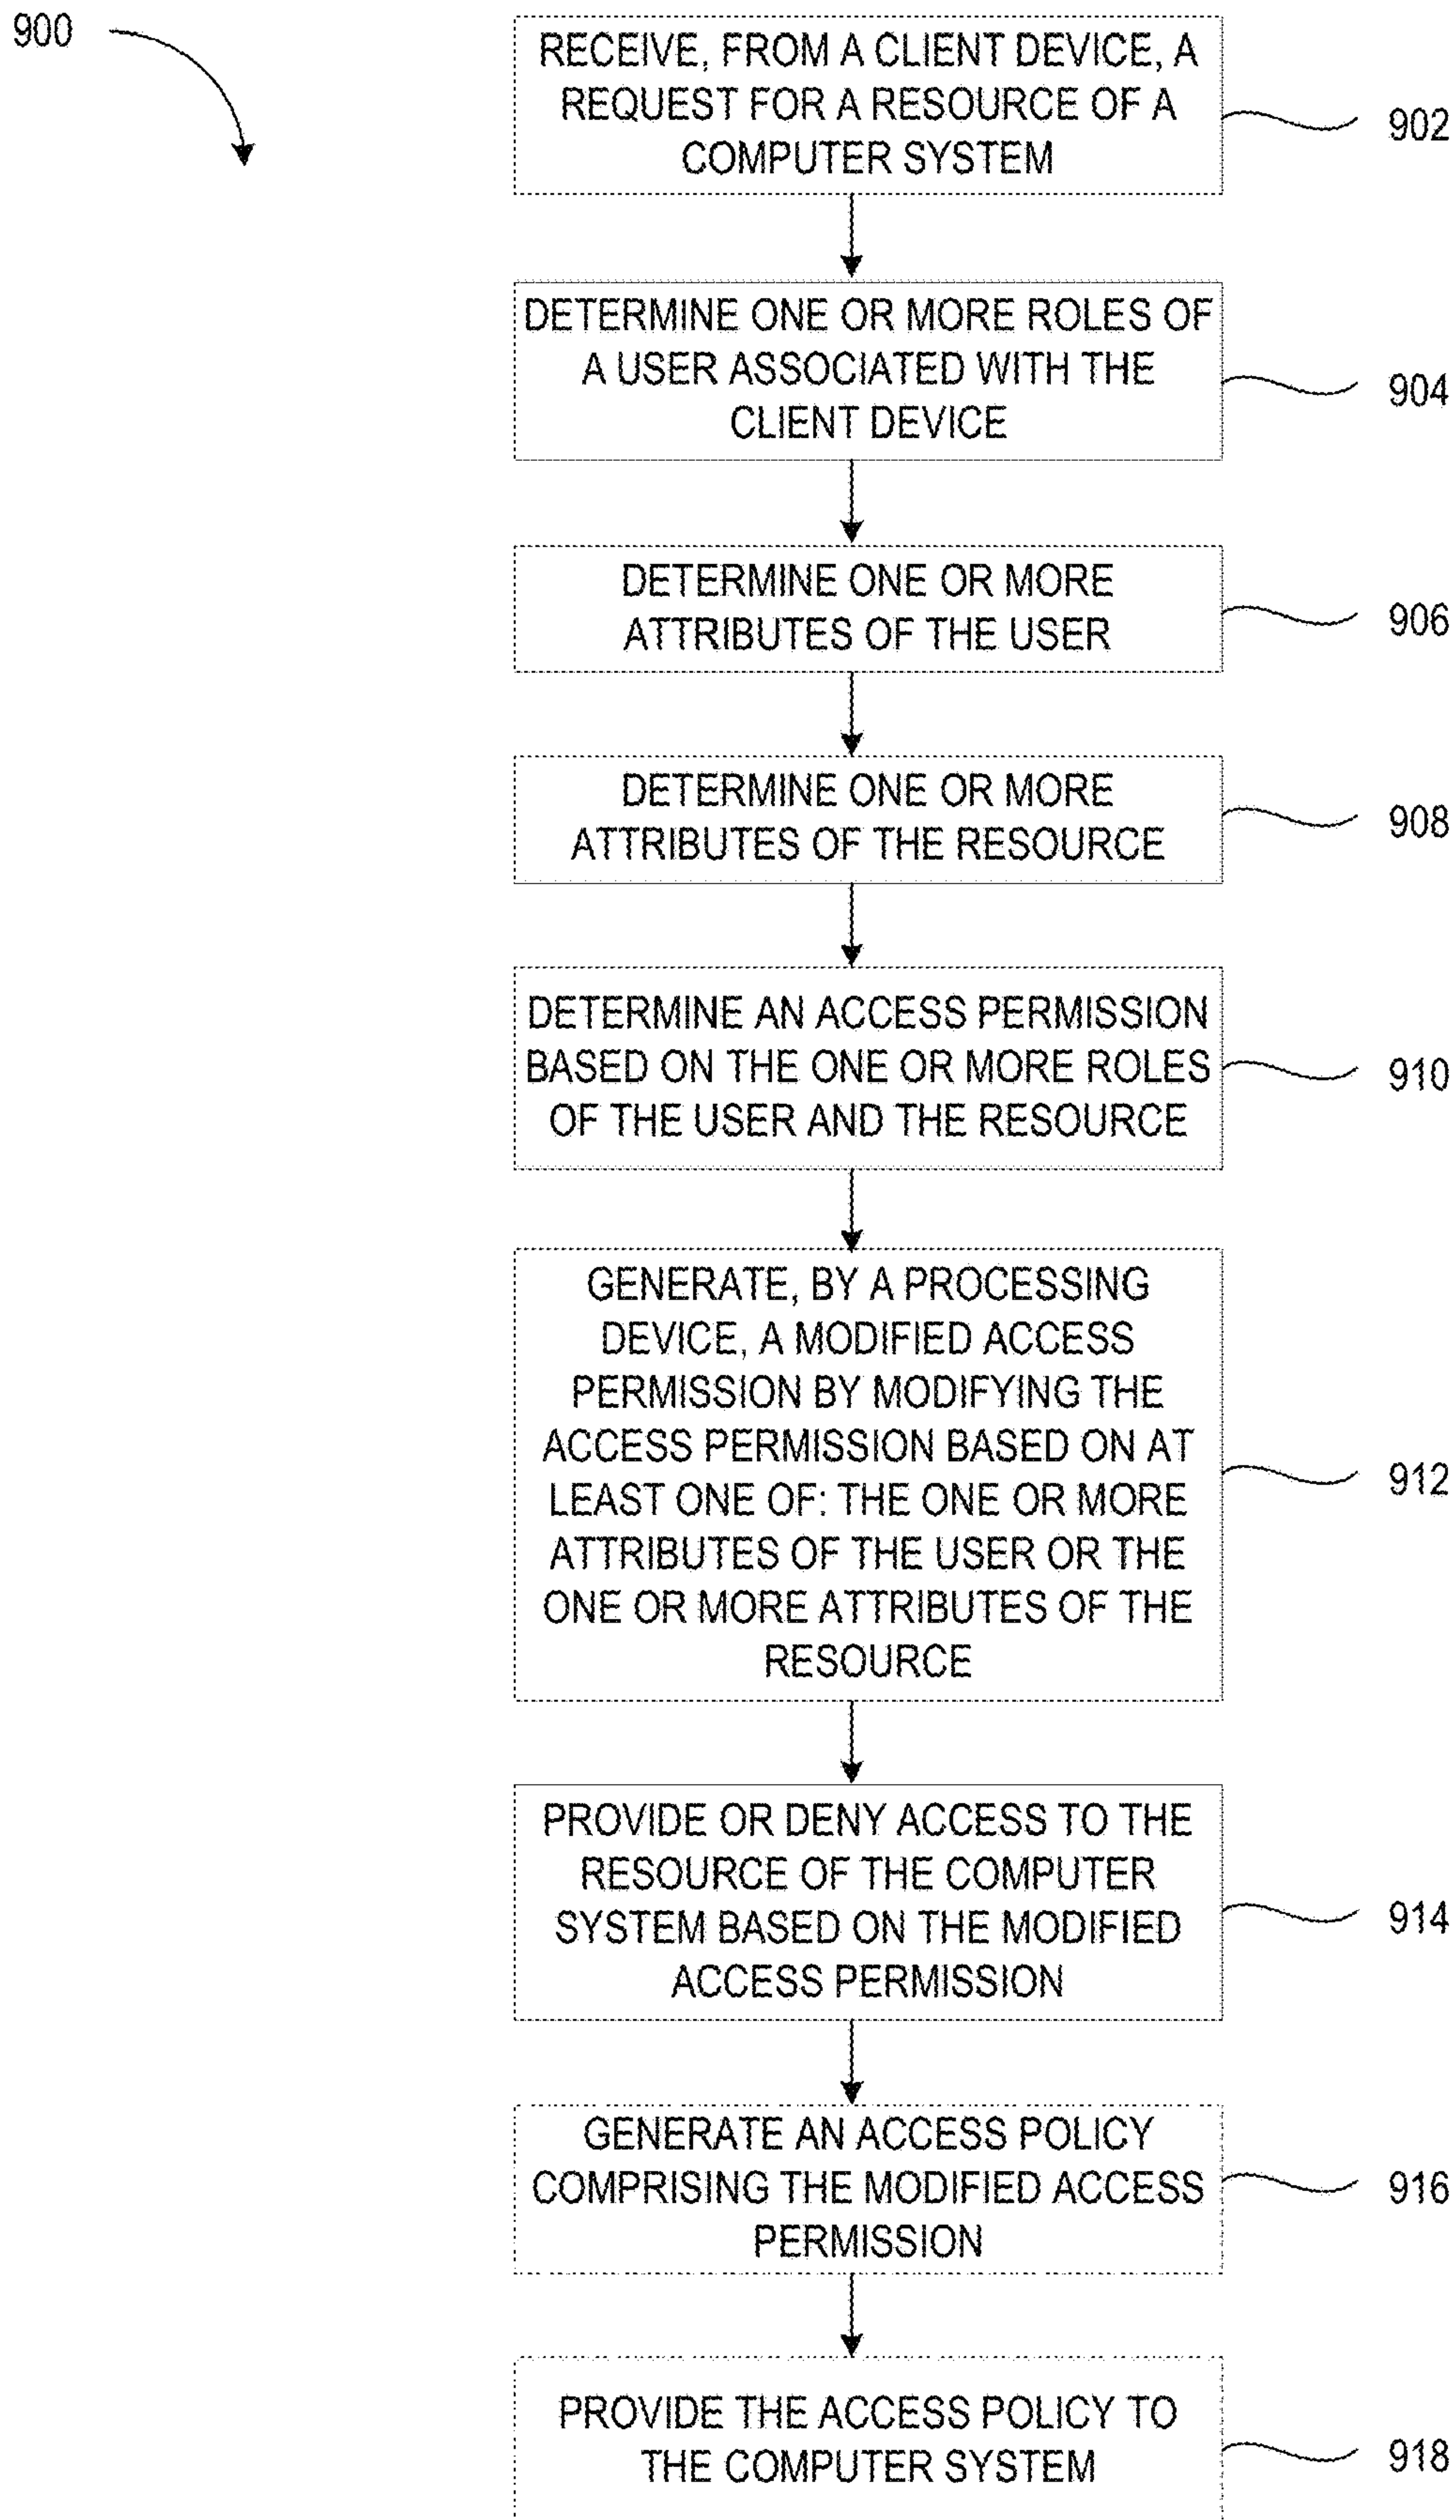
FIG. 7B



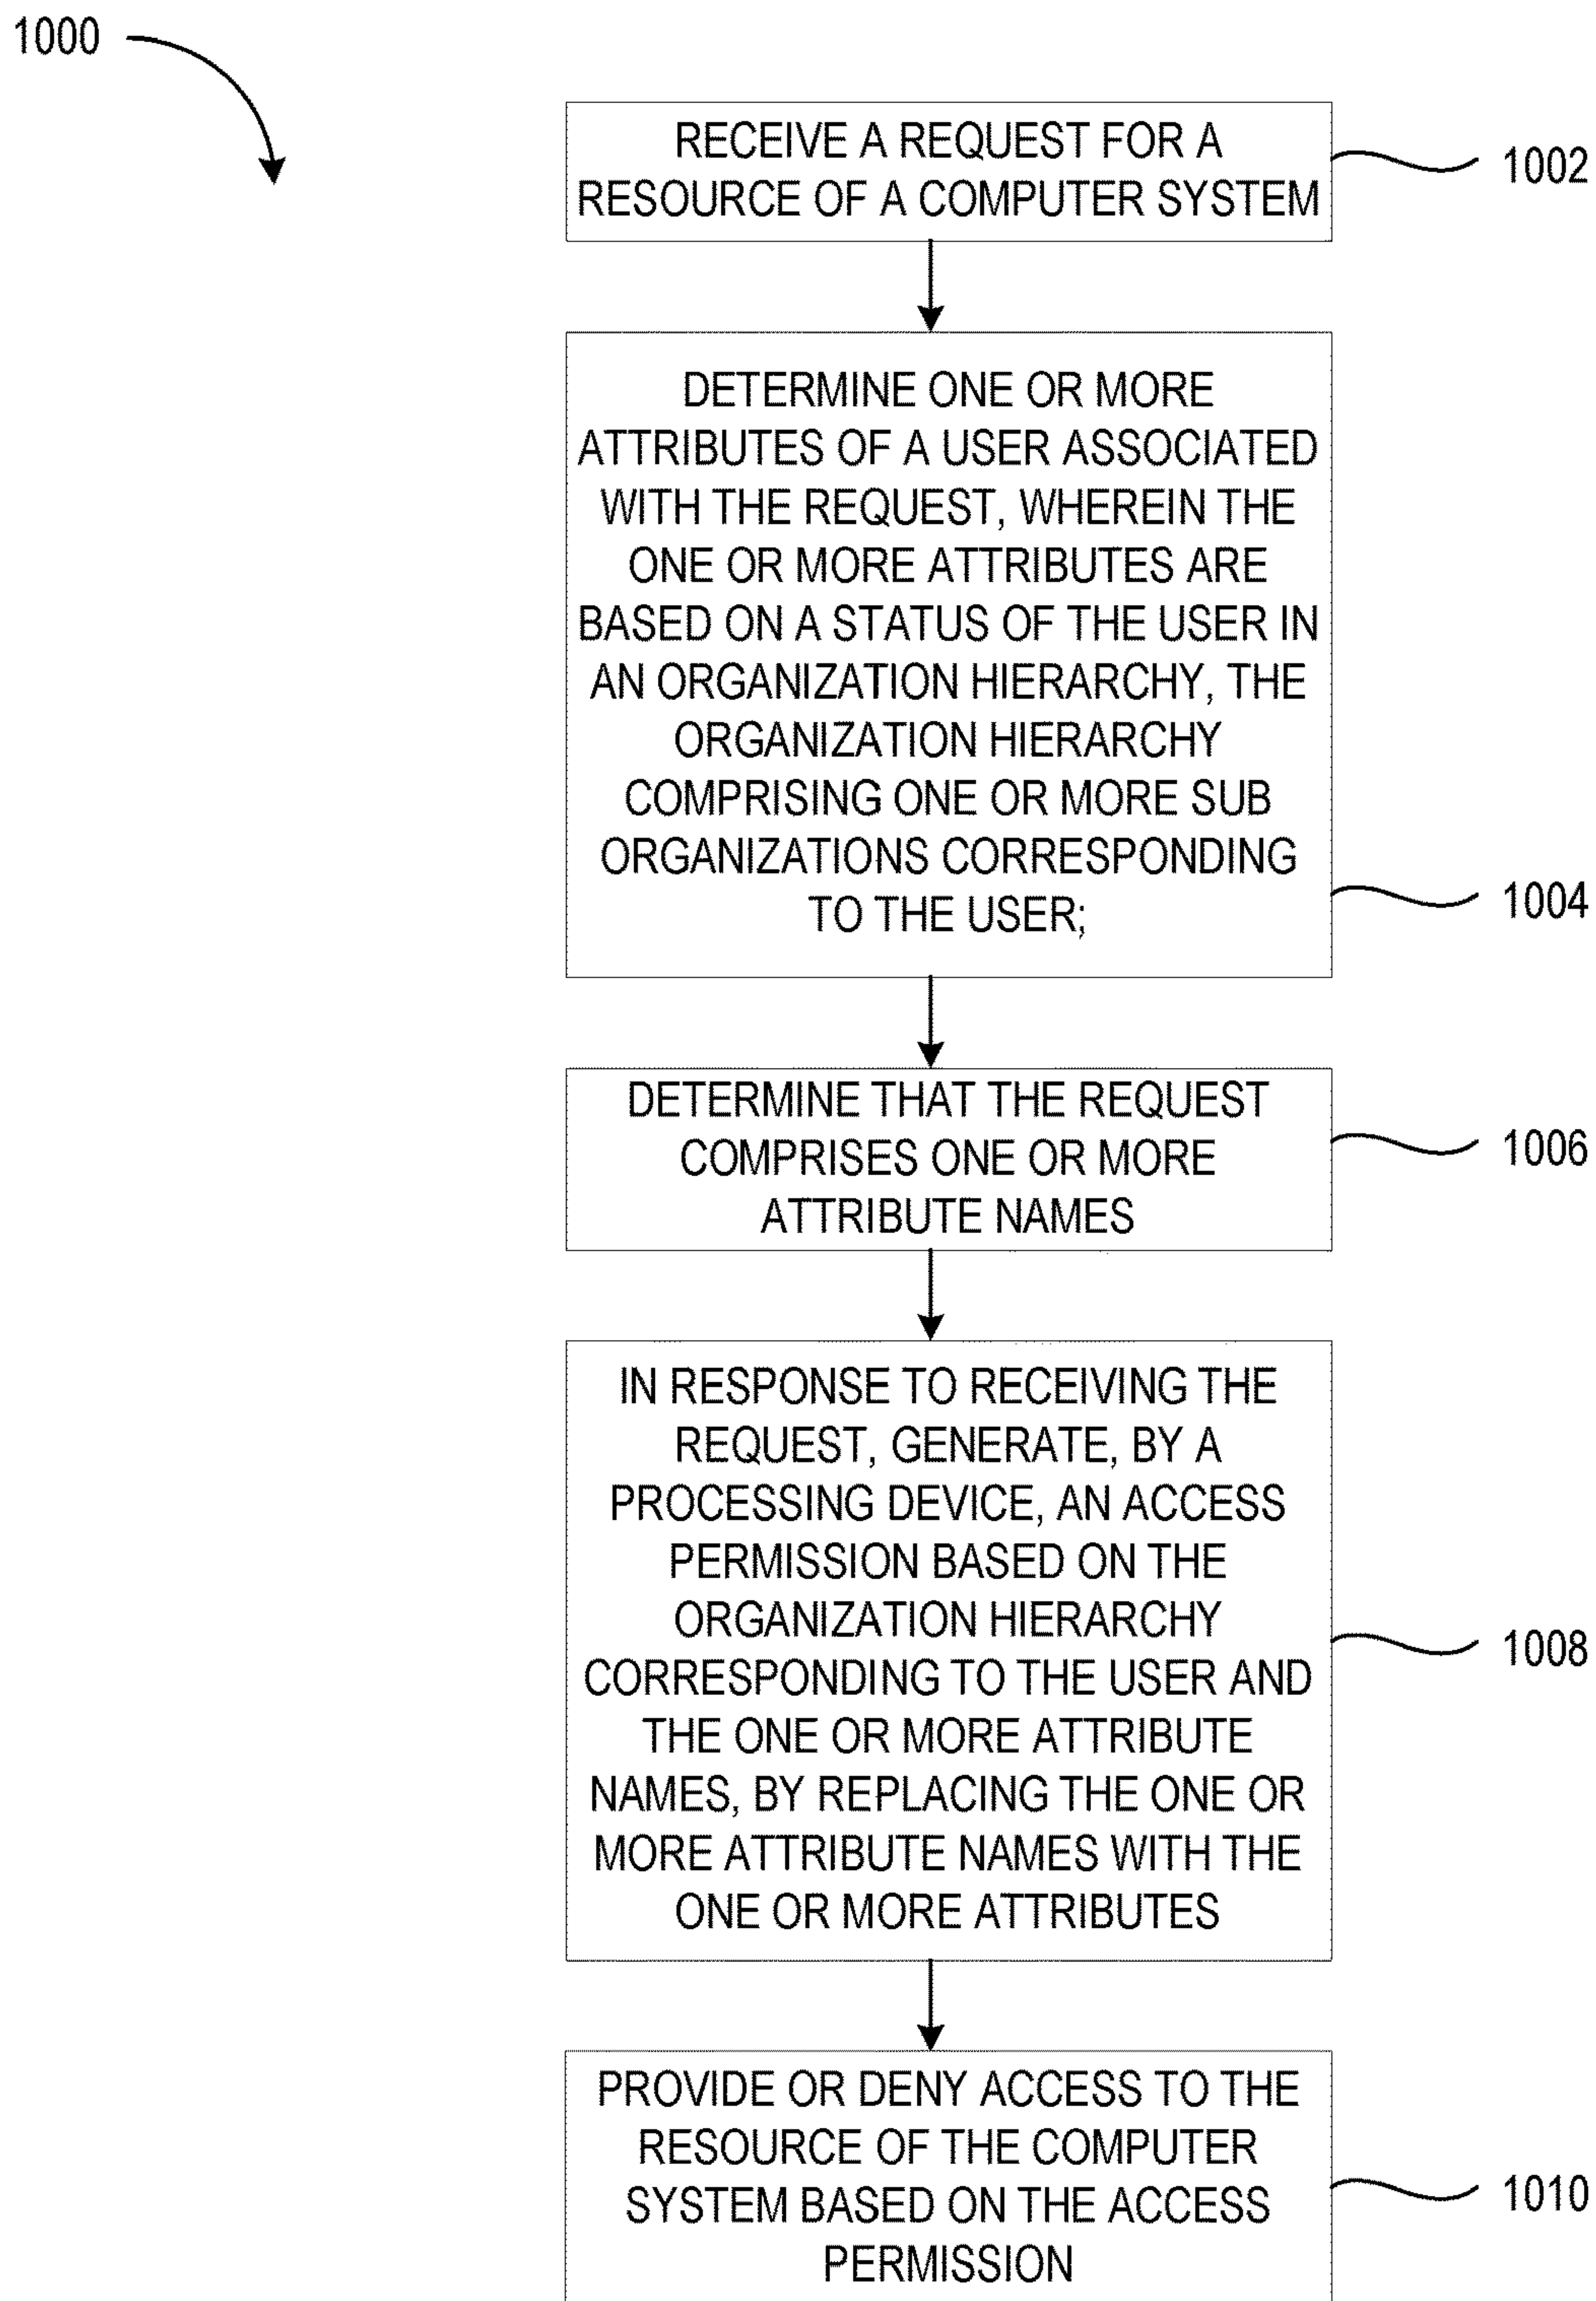
**FIG. 8A**



**FIG. 8B**



**FIG. 9**



**FIG. 10**



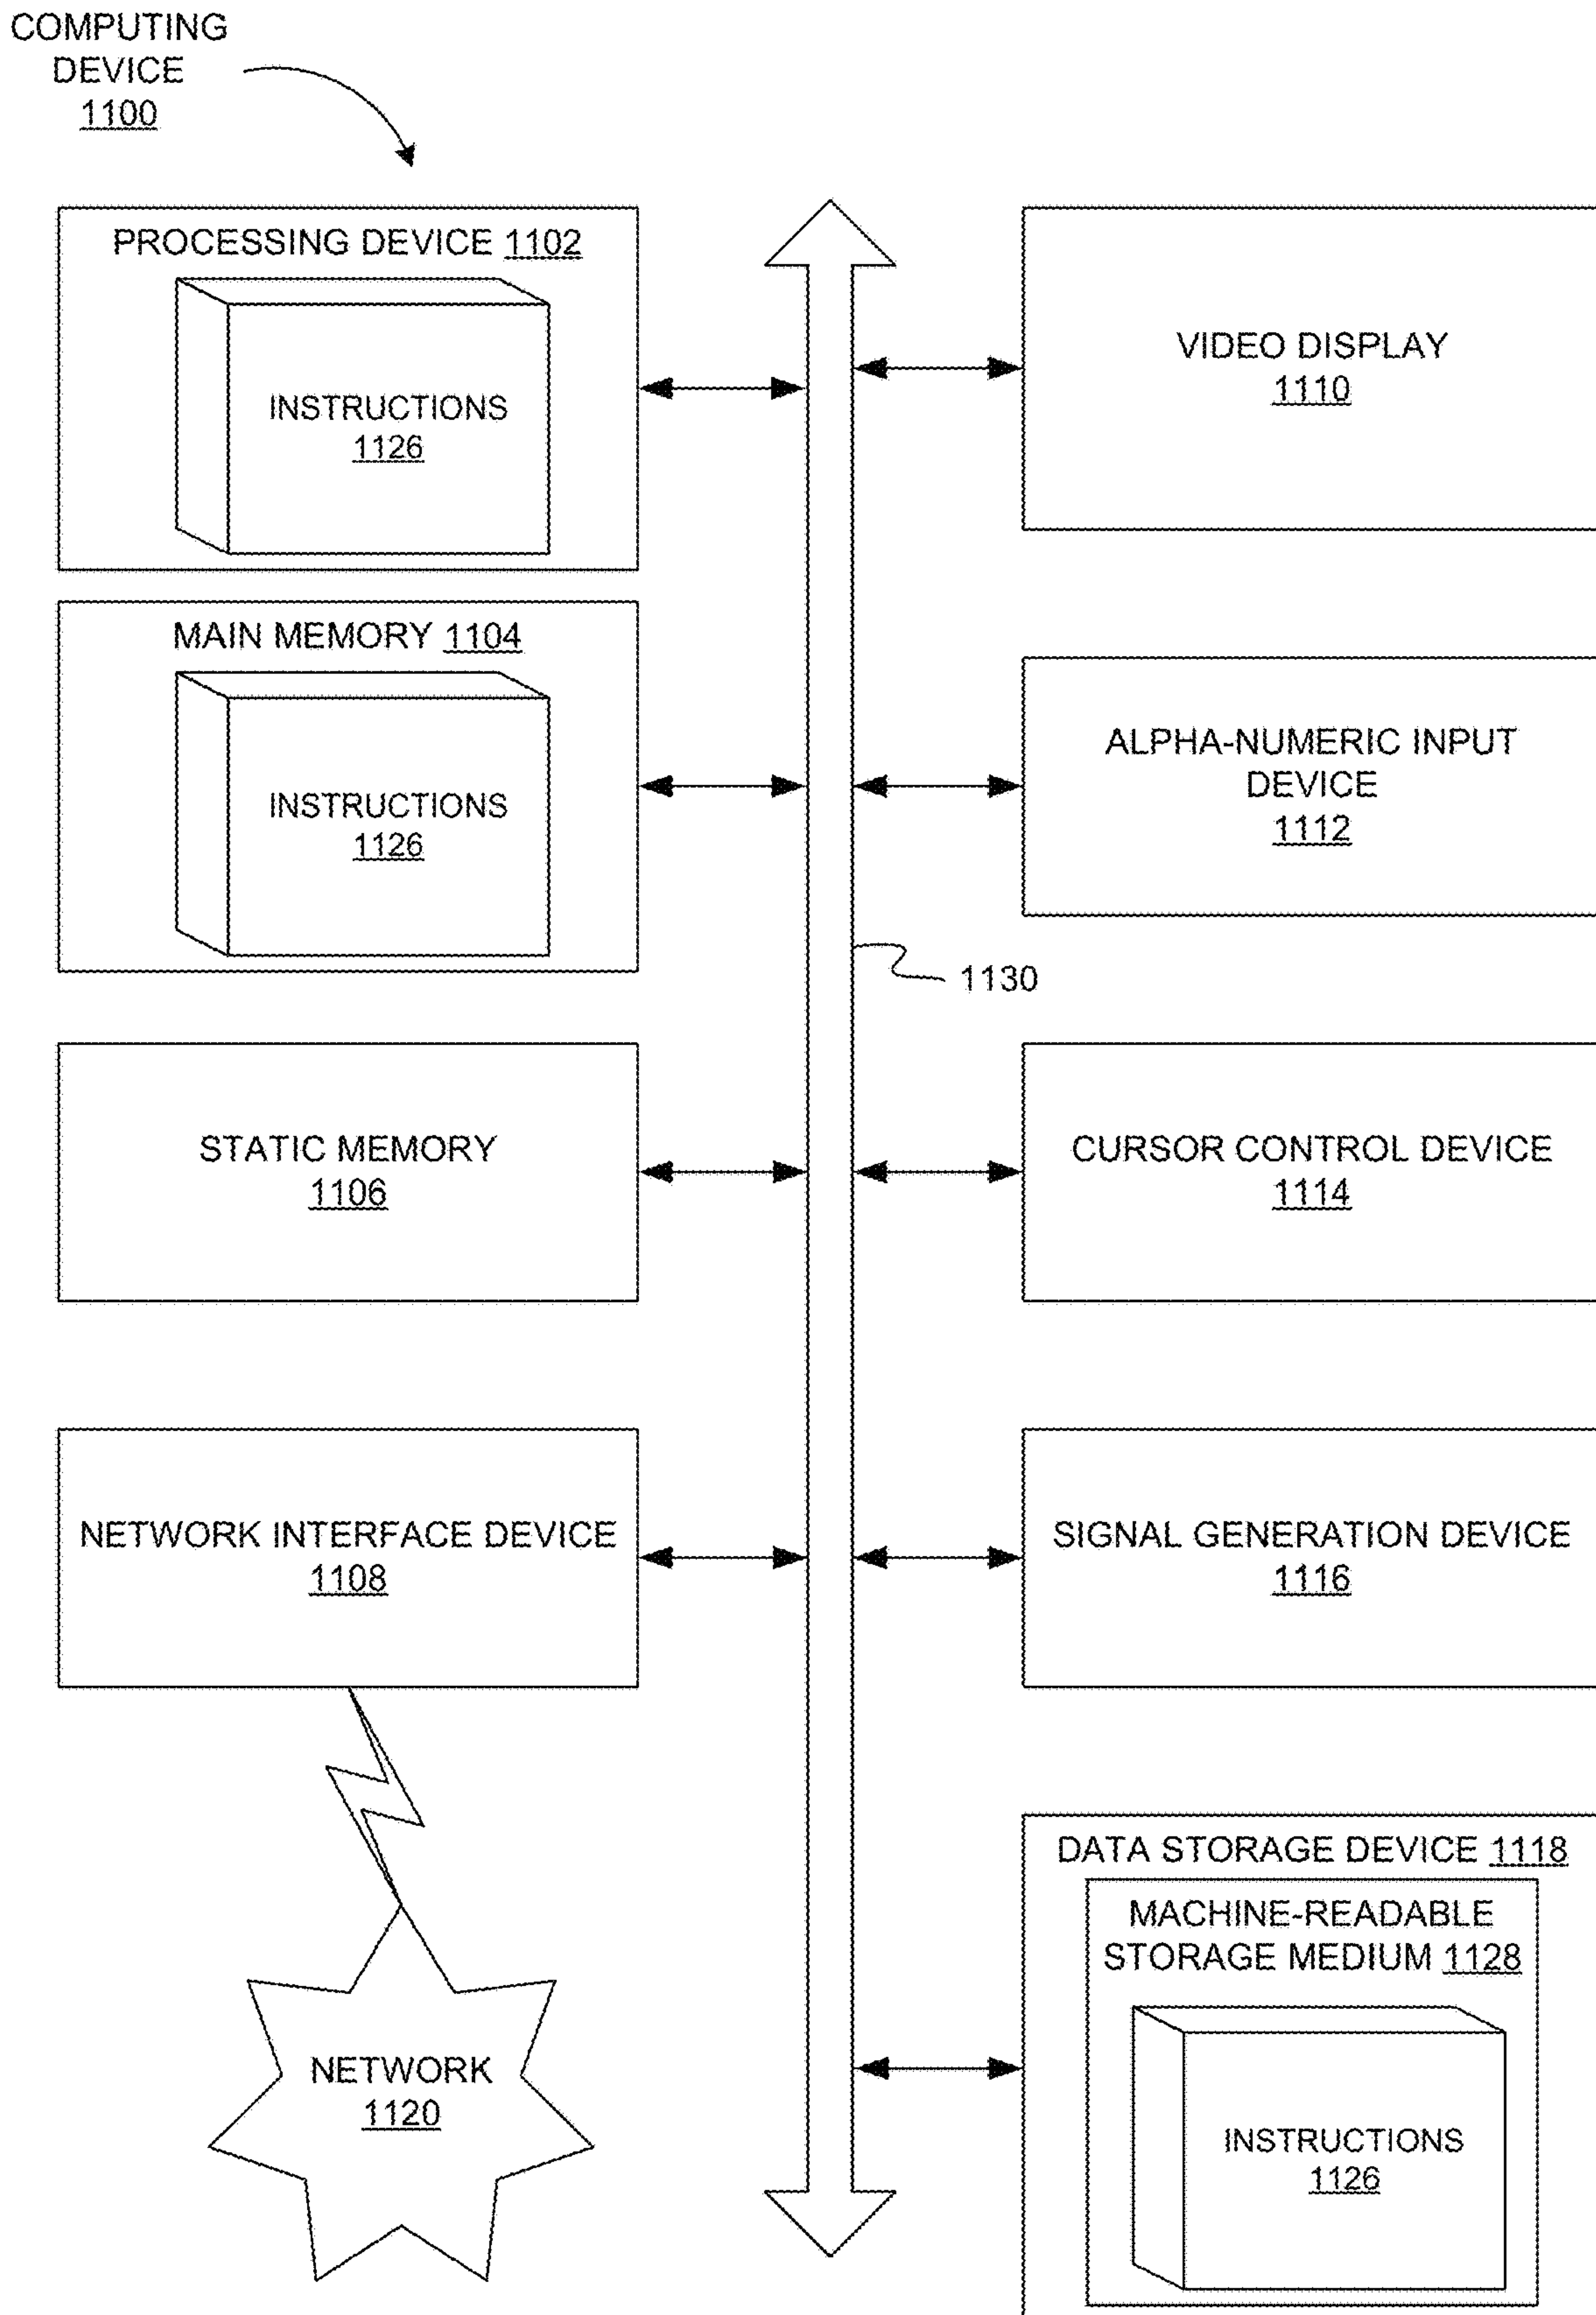


FIG. 11



## HYBRID ROLE AND ATTRIBUTE BASED ACCESS CONTROL SYSTEM

### TECHNICAL FIELD

Aspects of the present disclosure relate to access control systems, and more particularly, to a hybrid role and attribute based access control system.

### BACKGROUND

In general, access control systems allow for the authentication, authorization, and/or audit of client devices and user credentials. In an access control system, the entities that can perform actions on the system are called subjects, and the entities representing resources to which access may need to be controlled are called objects. Subjects and objects may both be considered as software entities. Authorization specifies what a subject can do and identification and authentication ensure that only legitimate subjects can log on to a system. Access approval via an access control system grants access during operations, by association of users with the resources that they are allowed to access, based on the authorization policy. Authentication methods and tokens include passwords, biometric scans, physical keys, electronic keys and devices, hidden paths, social barriers, and monitoring by humans and automated systems.

### BRIEF DESCRIPTION OF THE DRAWINGS

The described embodiments and the advantages thereof may best be understood by reference to the following description taken in conjunction with the accompanying drawings. These drawings in no way limit any changes in form and detail that may be made to the described embodiments by one skilled in the art without departing from the spirit and scope of the described embodiments.

FIG. 1 is a block diagram that illustrates an example access control system, in accordance with some embodiments of the present disclosure.

FIG. 2 is a block diagram that illustrates an example policy graph without roles, in accordance with some embodiments of the present disclosure.

FIG. 3 is a block diagram that illustrates an example policy graph with flat roles, in accordance with some embodiments of the present disclosure.

FIG. 4 is a block diagram that illustrates an example policy graph with a role hierarchy, in accordance with some embodiments of the present disclosure.

FIG. 5 is a block diagram that illustrates an example authorization service data model, in accordance with some embodiments of the present disclosure.

FIG. 6 is a block diagram that illustrates an example resource hierarchy, in accordance with some embodiments of the present disclosure.

FIG. 7A is a block diagram that illustrates an example authentication service user and organization attributes data model, in accordance with some embodiments of the present disclosure.

FIG. 7B is a block diagram that illustrates example user attributes, in accordance with some embodiments of the present disclosure.

FIG. 8A is a block diagram that illustrates an example external policy enforcement point, in accordance with some embodiments of the present disclosure.

FIG. 8B is a block diagram that illustrates an example internal policy enforcement point, in accordance with some embodiments of the present disclosure.

FIG. 9 is a flow diagram of a method of hybrid access control, in accordance with some embodiments.

FIG. 10 is a flow diagram of a method of organization based access control, in accordance with some embodiments.

FIG. 11 is a block diagram of an example computing device that may perform one or more of the operations described herein, in accordance with some embodiments of the present disclosure.

### DETAILED DESCRIPTION

In general, each access control system supports a policy governing what permissions each user of the system has. Permissions may involve an action on a resource (possibly composed of many smaller resources). For example, one user could be granted a permission to “read” all records in a table of a relational database, while another could be granted a permission to “read”, “create”, “update,” and “delete” those records. The range of actions and resources is unbounded and very general: they can be virtual or physical, coarse- or fine-grained, etc.

A variety of access control methodologies exists. Role-based access control (RBAC) introduces a third entity between users and permissions: roles. A role may be an aggregation of permissions (each an action on a resource), and it is given a descriptive name like “Database Reader” or “Loan Approver.” Permissions may be granted to roles and revoked from them. Users may be placed in roles and removed from them. RBAC can grow increasingly complicated when there are many users and many resources, and access to those resources does not follow a regular pattern. Another access control methodology is attribute-based access control (ABAC).

In an ABAC system, when a user requests access to a resource, the ABAC system makes an access control decision based on a policy defined in terms of attributes (name/value pairs). Attributes describe the user, the resource, and the environment (or context) of the access control decision (things like the time of day, day of week, physical location of the user as determined by IP address, etc.). For example, a user might be granted access to a patient’s medical record if (a) the user has the job title “MD” or “Registered Nurse”, and (b) the medical record includes a signed privacy agreement. Like the RBAC system, the ABAC can grow increasingly complicated when there are many users and many resources, and access to those resources does not follow a regular pattern.

The present disclosure addresses the above deficiencies by introducing a new form of hybrid RBAC/ABAC system, called permission-centric. Because the access control community considers permissions “as un-interpreted symbols” they seem to have missed what’s possible if you define a deeper structure for permissions (e.g., one leveraging attributes). In a permission-centric approach, the user’s permissions are gathered just as in a RBAC system—the user’s roles are determined, and all the permissions granted by those roles are gathered together. While other role-centric hybrid approaches may then go on to eliminate some of those permissions based on the set of attributes, the current approach retains all the permissions but modifies them based on the set of attributes.

FIG. 1 is a block diagram that illustrates an example access control system 101, in accordance with some



embodiments of the present disclosure. In one embodiment, access control system functionality may be provided by a server **100**. Server **100** may include various components, which may allow an access control system application to run on a server device. Each component may perform different functions, operations, actions, processes, methods, etc., for an access control system and/or may provide different services, functionalities, and/or resources for the access control system.

Different components of system **101** may be located on and/or may execute on different processing devices (e.g., processing device **120**) of the server **100**, as discussed in more detail below. In one embodiment, authorization service **127** may be implemented on processing device **120** of server **100**. In one embodiment, authorization service **127** performs hybrid role and attribute based access control system operations. In another embodiment, authorization service **127** performs role, attribute, and/or organization based access control system operations.

As illustrated in FIG. 1, server **100** includes a computing processing device **120**, a data store **130**, and a network **105**. The processing device **120** and the data store **130** are coupled to each other (e.g., may be operatively coupled, communicatively coupled, may communicate data/messages with each other) via network **105**. Network **105** may be a public network (e.g., the internet), a private network (e.g., a local area network (LAN) or wide area network (WAN)), or a combination thereof. In one embodiment, network **105** may include a wired or a wireless infrastructure, which may be provided by one or more wireless communications systems, such as a wireless local area network (WLAN) (e.g., WiFi®) hotspot connected with the network **105** and/or a wireless carrier system that can be implemented using various data processing equipment, communication towers (e.g. cell towers), etc. The network **105** may carry communications (e.g., data, message, packets, frames, etc.) between the various components of server **100**. The data store **130** may be a persistent storage that is capable of storing data. A persistent storage may be a local storage unit or a remote storage unit. Persistent storage may be a magnetic storage unit, optical storage unit, solid state storage unit, electronic storage units (main memory), or similar storage unit. Persistent storage may also be a monolithic/single device or a distributed set of devices.

Each component may include hardware such as processing devices (e.g., processors, central processing units (CPUs), memory (e.g., random access memory (RAM), storage devices (e.g., hard-disk drive (HDD), solid-state drive (SSD), etc.), and other hardware devices (e.g., sound card, video card, etc.). The servers **100**, **103** may comprise any suitable type of computing device or machine that has a programmable processor including, for example, server computers, desktop computers, laptop computers, tablet computers, smartphones, set-top boxes, etc. In some examples, the server **100** may comprise a single machine or may include multiple interconnected machines (e.g., multiple servers configured in a cluster). The server **100** may be implemented by a common entity/organization or may be implemented by different entities/organizations. For example, a server **100** may be operated by a first company/corporation and a second server (not pictured) may be operated by a second company/corporation. Each server may execute or include an operating system (OS), as discussed in more detail below. The OS of a server may manage the execution of other components (e.g., software, applications, etc.) and/or may manage access to the hardware (e.g., processors, memory, storage devices etc.) of the computing

device. Further implementation details of the operations performed by server **100** are described with respect to FIGS. 2-11.

FIG. 2 is a block diagram that illustrates an example policy graph **201** without roles, in accordance with some embodiments of the present disclosure. In one embodiment, a role-based access control system provides authorization to perform particular actions by particular people or client devices. In other words, authorization (e.g., access control) is a process that determines whether an authenticated user has the right to perform some requested action (e.g., “Get”) on a given resource (e.g., a particular web service endpoint). For example, an access control system may provide responses to questions of the form: “Can subject S perform action A on resource R?” Some examples, of these types of questions include: Can Bob read document /db/es/orders/123?; Can Alice do a GET on web service endpoint /service/auth/users?; Can Alice see user interface control /ui/wiki/page\_list/add\_page?; and Can service account NBSSVC01 update the database table /db/orders\_db/customer\_schema/address? To provide responses to such questions, the access control system may be configured with a set of rules. This set of rules may be known as an access control policy. The policy may be sliced into smaller chunks, such as the policy for a particular user or the policy for a given web service API.

In one embodiment, a permission may be a grouping of an action and a resource. For example, “read document /db/es/orders/123” may be considered a permission with “read” being the action and “document /db/es/orders/123” being the resource. Accounting for permissions, an access control system may provide responses to queries such as, Does Bob have the permission “read document /db/es/orders/123”?; and Does Alice have the permission “get endpoint /service/auth/users”? Thus, in one embodiment, a policy may be a group of permissions.

In one embodiment, if the rules forming an access control policy are too specific then a large number of rules may be required to be maintained. Disadvantageously, the policy may become harder to maintain as rules continue to be added. Auditing such a policy may also prove difficult and time consuming. In one embodiment, such a policy may include many very specific rules of the form: Subject S has permission to perform action A on resource R. Whenever a new user is added to the policy, then the policy administrator has to create a rule for every permission with which the user should be entitled. Likewise, whenever a user is dropped from the policy, the administrator has to delete the permissions of that user. In one embodiment, a policy graph representing such a policy may appear similar to policy graph **201** in FIG. 2. In policy graph **201**, Permissions **204a-f** are assigned to Users **202a-c**, according to individual authorization allowances. For example, as shown, User 1 **202a** is not assigned any Permissions **204a-f** and therefore is not authorized to perform any actions in the system. User 2 **202b** is assigned Permissions 1, 2, 4, and 5 (**204a**, **204b**, **204d**, and **204e**, respectively), while User 3 **202c** is assigned each of the Permissions **204a-f**, allowing each user to perform actions corresponding to their respective permissions.

In one embodiment, real-world policies may be orders of magnitude more complex than that shown in policy **201**. Even if there were just a hundred users and a thousand resources, and each user had access to 200 resources on average, the policy graph would have 20,000 edges. The policy administrator would have a difficult time managing the connections and adding and deleting users and resources. The administrator may make mistakes and be swamped with



## 5

access control issues. The security auditor may have a difficult time determining who has which permissions. Therefore, an access control system may be more efficient using fewer abstract, universal rules instead of many concrete, specific ones.

FIG. 3 is a block diagram that illustrates an example policy graph 301 with flat roles, in accordance with some embodiments of the present disclosure. Advantageously, a Role-Based Access Control (RBAC) system may increase the efficiency of a system without roles, by introducing intermediate role nodes (e.g., Role 1 306a and Role 2 306b) to the policy graph 301. In one embodiment, a role is a related group of permissions. Users may be assigned to roles instead of being directly granted permissions. Adding roles the example policy 201 of FIG. 2 results in User 2 302b being assigned to Role 1 306a, and User 3 302c assigned to Role 2 306b. Each being assigned to a specific role, User 1 302a and User 2 302b may be assigned the respective Permissions 304a-304f assigned to each Role 306a-b. User 1 302a is not assigned to a role, and therefore is not assigned any of Permissions 304a-f.

In one embodiment, if the administrator needs to add a new user with the same permissions as User 2, the access control system may add one new edge (connecting the user to Role 1 306a) to the graph. When User 2 leaves the company, only one edge needs deletion. An access control system supporting only one level of roles, such as that of FIG. 3, is a Flat RBAC system. Access control policies may be simplified further by introducing role hierarchies, where roles can be assigned to other roles. FIG. 4 is one example of an access control system having role hierarchies.

FIG. 4 is a block diagram that illustrates an example policy graph 401 with a role hierarchy, in accordance with some embodiments of the present disclosure. RBAC systems may include role hierarchies in a variety of ways. For example, a Restricted Hierarchical RBAC system may support a tree of roles. In another example, a General Hierarchical RBAC system supports a more general directed acyclic graph (DAG) of roles. One example of a General Hierarchical RBAC system is demonstrated by policy graph 401.

In policy graph 401, to add new permissions (e.g., Permission 1 404a, Permission 2 404b, Permission 4 404d, Permission 5 404e) for a User 3 402c (who already has Permission 3 404c and Permission 6 404f), an access control system may only have to grant that permission to Role 1 406a, and the User 3 assigned to Role 2 406B will get it automatically. User 2's 402b permissions (Permission 1 404a, Permission 2 404b, Permission 4 404d, Permission 5 404e) are maintained by the single assignment to Role 1 406a while User 1 402 still has zero permissions. Roles and hierarchies of roles can shrink a policy graph by orders of magnitude.

Role-based hierarchical access control systems may also include Constrained RBACs and Symmetric RBACs. In Constrained RBACs, users may be constrained to not have certain combinations of roles. For example in a loan application, users might be allowed to have either the "Loan Officer" or "Loan Approver" roles, but not both at the same time, to ensure that each loan has at least two people making the decision—achieving separation of duty (SOD). These SOD constraints can be static or dynamic. In static separation of duty, a user can never have certain combinations of roles. In dynamic separation of duty, if a user has a conflicting combination, they must choose which of the conflicting roles must be inactivated to resolve the conflict.

## 6

Symmetric RBACs have backwards and forwards symmetry. In a normal RBAC, one can find out which roles and permissions a user has, but not easily discover things in the reverse direction (e.g., given a particular permission, which roles grant that permission and who has been assigned those roles). A Symmetric RBAC, as described herein, is one that makes it easy for security auditors to answer such questions.

In one embodiment, in Attribute-Based Access Control (ABAC) systems, access control policies are composed of rules that grant or revoke access to resources based on various attributes. These attributes may pertain to the user, the resource, and/or the overall context of the request. For example, if the resource is a brokerage account, an ABAC policy may grant access only to employees that: (1) have a certain job category, (2) belong to the wealth management group assigned to the account, (3) are logged in from an IP address indicating that the request is coming from within the company VPN, and (4) are making the request during ordinary business hours. In some ways, ABAC can be more flexible than RBAC.

In one embodiment, hybrid approaches melding both RBAC and ABAC can result in smaller policies. RBAC can be used to model the more static attributes, while ABAC can be used for the more dynamic ones. A variety of hybrid access control systems may be described. In one approach, dynamic roles may be implemented. For example, an activation expression may be supplied, referencing one or more attributes to each role. If this expression evaluates to true, then the role is active and the user gains all its permissions. If the expression evaluates to false, then the user does not gain those permissions.

A second approach may be defined as attribute-centric. In this approach, a user's role names are additional boolean-valued attributes that can be referenced like any others in an ABAC system. The result of this second approach is essentially a pure ABAC system, since access is not being controlled by roles formed by collections of permissions.

A third approach may be role-centric. In this approach, permissions are granted by roles, as in RBAC, but those permissions are further constrained by rules referencing attributes. For example roles assigned to a user may grant the user ten permissions, but the user's attributes make four of the ten inactive.

A fourth approach, as described herein, may be permission-centric. In this approach, roles grant users permissions whose effects are modified by the values of attributes of the users. A data model describing this fourth approach is described in FIG. 5.

FIG. 5 is a block diagram that illustrates an example authorization service data model 500, in accordance with some embodiments of the present disclosure. As shown in model 500, Users 502 are granted Roles 504, and the Roles 504 are granted Permissions 506. The relationship 508 from Role 504 back to itself is the role hierarchy relationship (e.g., roles can be granted other roles), and the relationship 512 from Organization 510 back to itself is the organization hierarchy relationship.

In one embodiment, User 502 is the subject of access control requests. In various embodiments, the user may be a human, a service account, or some other entity. The access control system may store basic information on each user, such as, for example:

id	A unique identifier key (UUID)
organizationId	The user's organization's key (UUID).



-continued

firstName	E.g., Abraham.
lastName	E.g., Lincoln
email	E.g., AbrahamLincoln@whitehouse.gov
code	A human readable identifier, as defined by the tenant.
locale	"en_US" (any valid Java locale)
timezone	EST (any valid Java timezone name)
isLocked	True if the User has been locked out for some reason (e.g., due to inactivity).
lastAuthenticatedAt	The time the User last successfully authenticated, or null if the User has never logged in.

In other embodiments, the access control system may store less or more information than that described above. In one embodiment, textual fields in the access control system (e.g., displayName, firstName, lastName, email, etc.) are Unicode.

In one embodiment, each user (e.g., **502**) belongs to one organization (e.g., **510**). Organizations may be arranged in hierarchies, where each root organization is called a tenant. Root organizations (tenants) may have one or more sub-organizations beneath them. The access control system may store basic information on each organization, such as, for example:

id	A unique identifier key (UUID).
code	A short organization code, e.g., "Acme."
name	The name of the organization, e.g., "Acme, Inc."
description	A description of the organization (optional).

In one embodiment, parent/child relationships between organizations may be managed in a separate SQL closure table. Organization hierarchies of arbitrary depth are supported. The access control system described herein may be multi-tenant: it can support multiple tenant organizations at the same time, while giving each tenant the perception that it is the only tenant. Each tenant has its own users, roles, and permissions, and is rigorously isolated from every other tenant and has no way to tell if there are other tenants at all, except indirectly as a function of overall system performance. In one embodiment, there may be an exception to this isolation.

In one embodiment, an organization may be treated as a special kind of role. For example, if you are a member of the HR department, then you are a member of the HR role and are permitted to read all candidate interview results and comment on them. This technique can be supported in any RBAC system as is, or by adding a graphical user interface that makes some roles appear to look like organizations. In another embodiment, a group may parallel a role, and can be populated automatically from a directory of employee data. In this case, when an access control policy is compiled, the group/role distinction may be erased (e.g., both behave identically).

In one embodiment, each user has a list of the organizations to which they belong. One set might refer to the single formal organization the user belongs to, while another enumerates all of the organizations of which the user is a member, both formal and ad hoc. For example, the user's formal organization's attributes could be:

userOrgId: 1140  
 userOrgCode: 'OPS'  
 userOrgName: 'Operations'

In one embodiment, the current user's attributes may be referred to in a resource path. For example, if there is a web service that gives access to documents managed by an organization, there could be a permission that allows all members of an organization to read that organization's documents by default.

In another embodiment, the access control system may use list-valued attributes in resource paths. Embodiments can also allow expressions of attributes to appear in resource paths. In another embodiment, the access control system may support built in functions in resource paths. Any use of programming language expressions attributes may be used with respect to the embodiments disclosed herein.

In one embodiment, permission constraint bodies can reference the current user's attributes. For example, these two permissions grant access to all organization subtree documents no matter what the user's organization is:

```
May(Read, '/db/es', Row, '{organization:
suborgNames}')
```

```
May(Read, '/db/es', Column, '*')
```

In a similar manner to the resource path above, embodiments can use any programming language expression involving attributes. In one example, suppose that some documents have a "location" field whose value is a location code, and others where the "location" field is an location name. A combined permission may look like:

```
May(Read, '/db/es', Row, {location: locationCodes.concat(locationNames)}')
```

```
May(Read, '/db/es', Column, '*')
```

The expression locationCodes.concat(locationNames) may produce a list including all the user's location names and their codes.

In one embodiment, a permission (e.g., **506**) consists of: an action the user may take, a resource (or set of resources) the user may take that action on, and/or an optional constraint on the action. Permissions in the access control system described herein may be purely positive: a user with no permissions at all can do nothing, while each new permission only grants the user new authorization while not removing any. Some examples of permission fields are:

id	A unique identifier key (UUID).
roleID	The UUID of the Role the Permission belongs to. A permission may only belong to one Role.
domain	The Permission's "domain" type, either "Tabular" if it is a permission on some table-like resource, or "HTTP" if it is a permission on some resource that looks like a URI.
action	The action that is being permitted on the resource.
resource	The hierarchical path denoting a particular resource or subtree of resources.
constraintType	An optional constraint type.
constraintBody	An optional constraint that qualifies the action permitted on the resource.

In various embodiments, some examples of permissions are:



Example	Domain	Action	Resource	Constraint Type	Constraint Body
GET from endpoints starting with /service/auth/users?	HTTP	GET	/service/auth/users	n/a	n/a
GET a UI page?	HTTP	GET	/ui/data-viewer/home-page	n/a	n/a
Read all documents in index db/es/orders	Tabular	READ	/db/es/orders	ROW	{ }
Read all fields in documents from index /db/es/orders	Tabular	READ	/db/es/orders	COLUMN	*

15

In one embodiment, a user having the first permission may do GETs on all endpoints whose URIs (after a bit of abstraction) start with “/service/auth/users.” The second permissions form a pair, and together they allow users to read all rows and all columns from a particular tabular structure.

In one embodiment, each permission has a domain (e.g., the broad family to which it belongs). In some embodiments, there are two domains. In other embodiments, any number of domains may exist. In the present example, a “Tabular” domain may be used when the resource has a tabular structure of some sort and it makes sense to control access at the row and column level. A relational database table is one example of a tabular of course, and so is a document store index. When a permission is tabular, its action may be one of the “CRUD” verbs, either “Create”, “Read”, “Update”, or “Delete,” and it may have a constraint type of “ROW” or “COLUMN”. The “HTTP” domain may be used when the resource can be viewed as a unitary whole and the user is being granted permission to operate on it in its entirety. Its resource can be thought of as a URI. The permitted actions may be the HTTP verbs—“GET”, “PUT”, “POST”, “PATCH”, “DELETE”, and “HEAD,” and there is no constraint. In one embodiment it may seem that limiting the verbs to HTTP is too restrictive, but advantageously, because HTTP supports the REST architectural style it has a lot of expressive power. For instance, a set of UI widget permissions maybe modeled by coming up with a RESTful resource design that organizes them in a hierarchy, say by application, then page, then component hierarchy on the page. In one embodiment, granting a user the “GET” permission on a component may mean that component is enabled, otherwise it may be disabled.

In one embodiment, an Action is a verb. As discussed herein, Tabular permissions may support CRUD verbs and HTTP permissions support HTTP methods. The access control system herein does may not impose semantics on the actions (e.g., other than to limit them based on the domain type). To the access control system, actions may simply be strings.

In one embodiment, a resource of a permission may be a path string addressing some entity or aggregate entity of an application (e.g., /ui/checkout/order\_review\_page/buttons/one\_click). Resources are described in greater detail below. In one embodiment, tabular domain permissions may be constrained as to which rows and columns may be read. The row and column constraint languages are also described in further detail below.

In one embodiment, a role includes the following fields:

id	A unique identifier key (UUID)
name	The name of the role, e.g., “Loan Officer.”
description	A description of the role (optional)

20

In one embodiment, a role may have many permissions granted to it, and many roles granted to it. A permission may be shared by many roles, and a role may be granted to many roles.

FIG. 6 is a block diagram that illustrates an example resource hierarchy 601, in accordance with some embodiments of the present disclosure. In FIG. 6, example resource hierarchy 601 is shown to include nodes/602, /a 604, /b 606, /c 608, /a/Alice 610, /a/Bob 612, /a/Cara 614, /a/Alice/orders 616, /a/Bob/orders 618, /a/Cara/orders 620, /c/Ohio 622, and/c/Utah 624. In one embodiment, a resource is a slash-separated path, e.g., a Unix file name or a key in a key-value store. To the access control system described herein, a resource is simply a path. In one embodiment, resources do not have to correspond to anything concrete. Some possible resources may include: /db/es/my\_index/my\_document\_type; /service/family/users/Bob/children/Rachel/pets/Clover; and /ui/checkout/order\_review\_page/buttons/one\_click.

Permissions can be used to reference resources at any level of granularity. It is possible to grant access to every row and column of an entire database all at once, and it is also possible to grant access to individual table columns or particular table rows based on their unique key values. To keep policies small, expressive and efficient, the access control system described herein introduces various mechanisms for addressing groups of resources.

In one embodiment, a permission may be created that grants access to a subtree of resources instead of one permission per resource in that tree. For example, to allow a user to wear any clown nose, the access control policy may set the permission resource to /db/costume/clown/noses. Or, to enable the wearing of, any clown costume components, the access control system may set the permission resource to /db/costume/clown. Turning to the example resource hierarchy 601, to grant a permission on the subtree containing /c 608, /Ohio 622, and /Utah 624, the access control system may set the permission’s resource to /c 608.

In another embodiment, the access control system described herein may utilize wildcard segments. The access control system may use wildcards in resource names, which may be especially useful for granting permissions to web service endpoints, for example. To match URI path segments containing arbitrary characters, the access control system may use the wildcard { }. For example, the access control system itself may have one endpoint for getting a user’s permissions and another for getting the user’s roles: (1)



## 11

/users/{userId}/permissions and (2) /users/{userId}/roles. The access control system may also have endpoints for returning information on a particular user (e.g., email): /users/{userId}, another for getting pages of user information: /users, and another for getting a user's compiled policy: /users/{userId}/policy.

In one embodiment, to grant access to just the user roles endpoint without the others the access control system may use the resource: /users/{ }/roles. To make wildcard resources more descriptive the access control system can understand any text between the braces. For example, these resources are all the same: /users/{ }/roles; /users/{userId}/roles; and /users/{user UUID}/roles. Turning to example resource hierarchy 601, to select the nodes /Alice 610, /Bob 612, /Cara 614, /Alice/orders 616, /Bob/orders 618, and /Cara/orders 620, the access control system may use /a/{userId} (or just /a/{ }).

In another embodiment, the access control system described herein may utilize attributes. The access control system's resources and constraints can make reference to attributes of a user and the user's organization. In one example in which the access control system is to grant each user the right to see their own roles but not the roles of other users, the access control system may insert a reference to a user-specific attribute in the resource: /users/\${userId}/roles.

When it comes time for this permission to be checked in an access control request, the string \${userId} may be replaced by the current user's actual id value. One user would be granted access to only: /users/5e670257-bd15-4f75-91c3-5fedd02f49d1/roles, while another would be granted access to only: /users/c828d5e1-86e3-4a36-ae02-bfed02aad11b/roles. Subtree references, wildcards, and attributes allow the access control system described herein to carefully control the size of an access control policy.

In one embodiment, the total namespace containing all resources in a policy should be properly segregated to leverage subtree and wildcard references. In the previous example, access was granted to every user's subtree. To grant only the current user's subtree, the access control policy may use /a/\${userId}: (e.g., resulting in /Bob 612 and /Bob/orders 618 being selected).

In one embodiment, the access control system may allow for the specification of a list of possible values for a resource segment. For example, comma-separated values inside a pair of square brackets may be used. For example, the resource /a/[Alice,Cara,Xavier,Yolanda,Zelda] may match two subtrees of example hierarchy 601 (e.g., /Alice 610, Alice/orders 616, /Cara 614, and Cara/orders 620).

In one embodiment, a resource may include regular expression segments. These may be delimited by a leading r{ and a trailing }, for example. In this embodiment, the resource /a/r{Alice|Cara|Xavier|Yolanda|Zelda} may match the same two subtrees of example hierarchy 601 (e.g., /Alice 610, Alice/orders 616, /Cara 614, and Cara/orders 620).

In one embodiment, a constraint identifies which part of a resource a user may act on. In one embodiment, permissions in the HTTP domain do not have constraints, meaning the user can act on the whole resource (e.g., endpoint, user interface component, etc.). In another embodiment, permissions in the tabular domain do have constraints—each tabular permission has either a column constraint or a row constraint.

In one embodiment, by default, a user is not permitted to read or update any column in a row (or document) in a tabular resource. Column access may be granted through a permission with a column constraints. In a column con-

## 12

straint, the action must be READ or UPDATE (CREATE and DELETE don't apply). The constraint itself may be a string that takes one of the following forms:

Column Constraint	Notes
A, b.b1, c	A list of the columns the user is permitted to read (or update). Here the user is granted access to a, b.b1, and c. The dot notation supports nested aggregates in document stores.
*	A shorthand name for "all columns."
-d, -e, -f.fl	A shorthand for "all columns except."

To determine which columns a user had access to in a tabular resource (e.g., 'R'): the access control system may gather all permissions the user holds through any role where: (1) the permission's constraint is a column constraint; (2) the permission's action is the desired operation, e.g., "READ"; and (3) the permission's resource name "matches" R. For example, if R is "/sales/orders/order\_item," then each matching permission will have a resource name of "/sales/orders/order\_item," "/sales/orders," or "/sales." The access control system may then merge these matching permissions into one, whose column constraint is the combination of the individual column constraints. For example, if one permission grants READ access to fields a, b, and c, while a second grants READ access to a, d, and e, then the combined permission grants READ access to a, b, c, d, and e. If there are no matching permissions, the user is not allowed to read any columns.

The following on-limiting examples demonstrate how two column constraints may be added together:

Column Constraint 1	Column Constraint 2	Merged Constraint	Notes
*	Anything	*	If one column constraint grants all access, it does not matter what the other constraint grants.
a, b	b, c	a, b, c	Two grants of explicit columns are unioned.
-a, -b	-b, -c	-b	Two "all-but" grants are intersectioned. Here 'b' is the only column not granted in the combined constraint.
a	-a	*	One constraint gives access to a single column, and another grants access to all other columns, so all columns are granted.
a, b, c	-a, -b, -c	-d, -e	Another example of an explicit list of columns being added to an "all-but" grant.

In one embodiment, by default a user is not permitted to read or modify any rows (documents) in a tabular resource. Row access may be granted through row constraints. A row constraint may be a serialized Javascript object (aka, map or dictionary). Each key/value pair in the constraint maps a column name to a condition on that column. For instance {foo: 1} grants access to every row (document) whose foo column has the value 1. A row constraint may be analogous to a SQL WHERE clause. For example, to grant the user the



right to read rows where the organizationCode column has a particular value, and the customerCode column has another value: {organizationCode: 'ACME', customerCode: '12341234'}. This corresponds to the SQL condition: (organizationCode='ACME' AND customerCode='12341234'). Any client of the access control system making a SQL query on this table must therefore ensure this SQL condition is AND-ed to the WHERE clause.

When determining which rows a user has access to in a particular tabular resource (e.g., 'R'): the access control system gathers all permissions that this user holds through any role where: (1) the permission's resource name "matches" R (for example, if R is "/sales/orders/order\_item", then each matching permission will have a resource names of "/sales/orders/order\_item," "/sales/orders," or "/sales"); (2) the permission's constraint is a row constraint; and (3) the permission's action is the desired operation, eg, "READ."

The access control system then tries to merge the matching permissions by coalescing their constraints. For example, if one of the permissions is unconstrained, then all

In one embodiment, row constraints may be free to leverage more of Javascript's power. For instance, if organization codes are sometimes stored with their proper case intact, and sometimes stored in lower case, this may be captured with the following row constraint: {locationCode: user.location.codes.join('').toLowerCase().split('')} In this embodiment, the user may read documents whose location code is any of the user's configured location codes converted to lower case.

In one embodiment, attributes in the access control system may have a variety of characteristics. For example, attributes may grouped into three types: (1) standard attributes for the current user and their organization, managed by the access control system; (2) optional named bundles of user attributes, managed by software outside of the access control system (or some component thereof); and (3) optional named bundles of organization attributes, also managed by software external to the access control system (or some component thereof). The following are non-limiting examples of standard attributes:

Attribute Name	Sample Value	Meaning
userId	"278254a0-92e3-4e9b-b7bf-5c04dae52eaa"	An Auth Service UUID uniquely identifying the current user.
userCode	"bob@email.example.com"	A unique code identifying the current user. The code may be an email address or some other value assigned by the user's identity provider (IDP).
organizationID	"f9d22251-7c9e-490c-aa17-26004e21466d"	The Auth Service UUID for the user organization.
organizationCode	"Example"	A code identifying the user's organization.
organizationName	"Example Technologies, Inc."	The human oriented name of the user's organization

the other permissions can be ignored. If two permissions have identical constraints, one can be discarded. If one constraint is more restrictive than another, it can be discarded. The result is a set of optimized row permissions. If there are no matching permissions, the user may not be allowed to read any rows.

In one example, a tenant has a tabular resource containing rows from many organizations, and it wants to restrict access to those rows based on the user's organization code. A user from organization "a" would only be able to read rows marked with organization code "a", while a user from organization "x" would only be able to read the "x" rows. In one embodiment, the access control system may allow for the definition of a role for each organization. Users in the "OrganizationA" role would have row permission: {organizationCode: 'a'} and users in the "OrganizationX" role would have: {organizationCode: 'x'}. Disadvantageously, this may results in the creation of a large number roles and permissions. It also may require a lot of administration to ensure that organizational changes are reflected in the access control policy's roles and user assignments. Advantageously, in another embodiment, the policy may be simplified and rendered more efficient by allowing it to reference user-specific attributes such as the organization code. In this embodiment, there may be only one role needed: "OrganizationReader", with the row permission: {organizationCode: organizationCode}. Advantageously, no matter what organization a user moves to, the policy adapts automatically.

Optional attributes for the current user and their organization are described with respect to FIGS. 7A and 7B.

FIG. 7A is a block diagram that illustrates an example authentication service user and organization attributes data model 700, in accordance with some embodiments of the present disclosure. In one embodiment, the access control service does not need to be the only source of user and organization attributes. Different applications can create their own bundles of attributes. For example, an equipment rental application may need to create user attributes specifying which clients each user is set up to rent from, and which clients each user is allowed to administer. To specify these attributes, the rental application may manage a group of "rental" attributes associated with each user. Likewise, applications can manage groups of attributes at the organization level. These groups may be called namespaces. FIG. 7A shows part of the access control data model extended with a User 702, User Attributes 704, Organization 706, and Organization Attributes 708.

FIG. 7B is a block diagram 701 that illustrates example user attributes, in accordance with some embodiments of the present disclosure. FIG. 7B shows two users (e.g., Alice 703 and Bob 705) with attribute namespaces (e.g., 707 and 709, respectively). The user Bob 705 has a group of attributes in the "rental" namespace 709. Bob is only allowed to rent from the client whose code is "c001," and he is not a rental administrator at any client. Alice 703 also has a group of attributes in the "rental" namespace 707. She is an admin-



istrator at two clients, and is also a customer of three other clients. In addition, Alice 703 has a collection of attributes in the “app\_store” namespace 711. In one embodiment, each namespace is represented with a JSON object consisting of zero or more attribute name/value pairs. In one embodiment, attribute names are strings and their values can be numbers, strings, lists, and sub-objects.

In one embodiment, standard attributes have simple names like `userId` or `tenantCode` and optional attributes have compound names. The first segment may be the user or organization, the second may be the namespace, and the third may be the attribute within the namespace. So, in FIG. 7B, the list of clients the current user can rent from is indicated by `user.rental.client_codes`.

In one embodiment, a policy is a collection of related permissions. It has been organized, compressed, and optimized so that a Policy Enforcement Point (PEP) can get it from the access control system more quickly, and so that the PEP’s questions to it can be answered in a time-effective manner. The access control system may store information about organizations, users, roles, and permissions and can form access control policies out of this data. Policy enforcement may be performed by Policy Enforcement Points (PEPs) (e.g., a service running inside or outside of the access control system). A distributed system may have many PEPs.

In one embodiment, there are two kinds of PEPs for databases. The first type is an Internal PEP: it is part of the database system itself. The second is an External PEP: it is part of the process accessing the database. An internal PEP may require less work to use and may be more difficult to bypass. An external PEP may be harder to write and easier to bypass. External PEPs are described with respect to FIG. 8A and internal PEPs are described with respect to FIG. 8B.

FIG. 8A is a block diagram 800 that illustrates an example external policy enforcement point 802, in accordance with some embodiments of the present disclosure. In the example FIG. 8A, the Web Service 804 wants to make a database read (e.g., from Access Controlled Database 806) request on behalf of the effective user (Bob). In this example, the Web Service 804 requests Bob’s access control policy for that database 806 from the Auth Service 808 client library. The client library translates this call into the corresponding Auth Service HTTP request. The Auth Service 808 queries its database 810 for all the relevant permissions, e.g., the ones granted to user Bob, pertaining to the tabular resource “/db/ordersm,” and granting the READ action. The Auth Service 808 compiles the selected permissions into Bob’s policy for that resource and action. This compilation may include: collecting the permissions by action and type of constraint; stripping off unnecessary parts of the permissions (e.g., timestamps and UUID); subdividing each set of permissions by the resource they apply to; propagating permissions down the resource tree (e.g., if Bob has a permission for resource /db/orders/a, it also applies to /deb/orders/a/b); and optimizing each set of permissions by striking out duplicates and merging less powerful permissions into more powerful ones.

The Auth Service 808 then serializes the optimized policy into JSON and returns it as the response to the Web Service’s 804 request. The Web Service 804 (with the help of a future Auth Service client library) gets Bob’s access control grants from the Policy, and uses them to restrict queries it makes on the access-controlled database on Bob’s behalf. This may involve: modifying the SELECT clause (or No-SQL equivalent) to only read the columns Bob has access to and AND-ing the WHERE clause (or equivalent) with Bob’s row constraints.

FIG. 8B is a block diagram 801 that illustrates an example internal policy enforcement point 803, in accordance with some embodiments of the present disclosure. In one embodiment the Auth Service 805 has facilities for mapping user policies into internal PEP 803 policies and transmitting them to the internal PEP 803. In this scenario, the Auth Service 805 runs a batch job periodically to push each user’s policy to PEP 803. This batch job iterates over each user and: asks the Auth Service 805 to query its database 807 for all the relevant permissions: the ones granted to that user, pertaining to the tabular resource “/db/es” (809), and granting the READ action (PEP 803 may be read-only); asks the Auth Service 805 to compile the selected permissions the user’s policy (as described above); maps the Auth Service 805 policy into an PEP 803 policy and sends that policy to PEP 803. When the Web Service 811 makes a query on PEP 803 it needs to identify the effective user via a header on the query request. PEP 803 enforces the policy.

In one embodiment, an Auth Service 805 policy is an outer map (e.g., object) whose attributes have objects (e.g., maps) as their values. The outer map (e.g., object) may have keys that indicate the action and the constraint type. The following example shows two keys, both pertaining to READ actions:

```

{ 'READ_ROW':
  { '/db/orders/a': [ {orgCode:'IL99'} ],
    '/db/orders/a/b': [ {orgCode:['IL01','IL02','IL03','IL99']} ] }
  'READ_COLUMN':
  { '/db/orders/a': [ '-commission,-price' ],
    '/db/orders/a/b': [ '**' ] }
}

```

The first key gathers all the row constraints for READs while the second gathers all the column constraints. Each group of constraints is organized with an inner Map (Object). The inner map’s keys are resource names, and the values are the lists of constraints to apply on each resource name. In the above example, there are two READ\_ROW constraints. The first says that throughout the resource subtree rooted at /db/orders/a the user may read any row (e.g., document) whose orgCode field is ‘IL99’. The second says that in the smaller subtree rooted at /db/orders/a/b the user may read any row whose orgCode is either ‘IL01’ or ‘IL99’. Likewise there are two READ\_COLUMN constraints: at /db/orders/a the user can read any field except commission and price, while at /db/orders/b the user can read any field period.

When looking up permissions in a policy, the search may start at the deepest resource level (e.g., /db/orders/a/b/c) and then see if there are any constraints at that level. If not, the search continues looking for the next higher level (/db/orders/a/b), and so on until constraints are found or the topmost level of the resource is reached without finding constraints for it. In one embodiment, the search for row constraints and the search for column constraints may be conducted independently.

FIG. 9 is a flow diagram of a method 900 of hybrid access control, in accordance with some embodiments. The method 900 may be performed by processing logic that comprises hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions run on a processing device to perform hardware simulation), or a combination thereof.

Referring to FIG. 9, at block 902 processing logic receives, from a client device, a request for a resource of a computer system. In one embodiment, the request for the



resource includes a resource path. In another embodiment, the resource path includes at least one of: a variable or a textual pattern. At block **904**, processing logic determines one or more roles of a user associated with the client device, at block **906**, processing logic determines one or more attributes of the user, and at block **908**, processing logic determines one or more attributes of the resource. In one embodiment, the one or more attributes of the user and the one or more attributes of the resource correspond to one or more key value pairs.

In one embodiment, at block **910**, processing logic determines an access permission based on the one or more roles of the user and the resource. At block **912**, processing logic generates, by a processing device, a modified access permission by modifying the access permission based on at least one of: the one or more attributes of the user or the one or more attributes of the resource. In one embodiment, the modified access permission identifies a resource expression and an action. In another embodiment, the modified access permission identifies a constraint, which limits access to a subpart of the resource. In one embodiment, the constraint includes at least one of: a variable or a textual pattern.

At block **914**, processing logic provides or denies access to the resource of the computer system based on the modified access permission. For example, if it is determined based on the modified access permission that the user is authorized to access the resource, processing logic will provide the access. If it is determined based on the modified access permission that the user is not authorized to access the resource, processing logic will deny the access.

Optionally, at block **916**, processing logic generates an access policy comprising the modified access permission, and, at block **918**, processing logic provides the access policy to the computer system.

FIG. **10** is a flow diagram of a method **1000** of organization based access control, in accordance with some embodiments. The method **1000** may be performed by processing logic that comprises hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (e.g., instructions run on a processing device to perform hardware simulation), or a combination thereof.

Referring to FIG. **10**, at block **1002** processing logic receives a request for a resource of a computer system. In one embodiment, the request is received from a client device of the computer system. In another embodiment, the request is received from a client device external to the computer system. The request may include a resource path, and wherein the resource path comprises at least one of: a variable or a textual pattern, as described herein.

At block **1004**, processing logic determines one or more attributes of a user associated with the request, wherein the one or more attributes are based on a status of the user in an organization hierarchy, the organization hierarchy comprising one or more sub organizations corresponding to the user. In one embodiment, the one or more attributes are not roles in the access control system. The one or more attributes based on the status of the user in an organization hierarchy may be a separate and distinct classification of organization-based attributes, compared to other attributes of the access control system. In one embodiment, the organization hierarchy is represented by a tree-type data structure. In another embodiment, the organization hierarchy is represented by a directed acyclic graph data structure. In a variety of other embodiments, the organization hierarchy may be any other data structure capable of describing the organization hierarchy.

At block **1006**, processing logic determines that the request comprises one or more attribute names. In one embodiment, attribute names may be a string of text (e.g., a variable, expression, pattern, etc.) that represents a classification of attribute to be inserted into a resource path of the request, as described herein. In one embodiment, the request for the resource includes a resource path. In one embodiment, the resource path includes at least one of: a variable or a textual pattern, as described herein.

At block **1008**, in response to receiving the request, processing logic generates, by a processing device, an access permission based on the organization hierarchy corresponding to the user and the one or more attribute names. In one embodiment, processing logic generates the access permission by replacing the one or more attribute names with the one or more attributes. The access permission may identify a resource expression, an action, and a constraint, wherein the constraint limits access to a subpart of the resource.

At block **1010**, processing logic provides or denies access to the resource of the computer system based on the access permission. Optionally, processing logic may generate an access policy comprising the access permission and provide the access policy to the computer system. It should be noted that the above operations may be performed in conjunction with a hybrid role and attribute based access control system, a strictly access based control system, or any other access control system that allows for the distinct definition of organizational attributes.

FIG. **11** is a block diagram of an example computing device **1100** that may perform one or more of the operations described herein, in accordance with some embodiments. Computing device **1100** may be connected to other computing devices in a LAN, an intranet, an extranet, and/or the Internet. The computing device may operate in the capacity of a server machine in client-server network environment or in the capacity of a client in a peer-to-peer network environment. The computing device may be provided by a personal computer (PC), a set-top box (STB), a server, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while only a single computing device is illustrated, the term "computing device" shall also be taken to include any collection of computing devices that individually or jointly execute a set (or multiple sets) of instructions to perform the methods discussed herein.

The example computing device **1100** may include a processing device (e.g., a general purpose processor, a PLD, etc.) **1102**, a main memory **1104** (e.g., synchronous dynamic random access memory (DRAM), read-only memory (ROM)), a static memory **1106** (e.g., flash memory and a data storage device **1118**), which may communicate with each other via a bus **1130**.

Processing device **1102** may be provided by one or more general-purpose processing devices such as a microprocessor, central processing unit, or the like. In an illustrative example, processing device **1102** may comprise a complex instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or a processor implementing other instruction sets or processors implementing a combination of instruction sets. Processing device **1102** may also comprise one or more special-purpose processing devices such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor (DSP), network processor, or the like. The processing device **1102** may be configured to



execute the operations described herein, in accordance with one or more aspects of the present disclosure, for performing the operations and steps discussed herein.

Computing device **1100** may further include a network interface device **1108** which may communicate with a network **1120**. The computing device **1100** also may include a video display unit **1110** (e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT)), an alphanumeric input device **1112** (e.g., a keyboard), a cursor control device **1114** (e.g., a mouse) and an acoustic signal generation device **1116** (e.g., a speaker). In one embodiment, video display unit **1110**, alphanumeric input device **1112**, and cursor control device **1114** may be combined into a single component or device (e.g., an LCD touch screen).

Data storage device **1118** may include a computer-readable storage medium **1128** on which may be stored one or more sets of instructions, e.g., instructions for carrying out the operations described herein, in accordance with one or more aspects of the present disclosure. Instructions **1126** implementing access control components, etc., may also reside, completely or at least partially, within main memory **1104** and/or within processing device **1102** during execution thereof by computing device **1100**, main memory **1104** and processing device **1102** also constituting computer-readable media. The instructions may further be transmitted or received over a network **1120** via network interface device **1108**.

While computer-readable storage medium **1128** is shown in an illustrative example to be a single medium, the term “computer-readable storage medium” should be taken to include a single medium or multiple media (e.g., a centralized or distributed database and/or associated caches and servers) that store the one or more sets of instructions. The term “computer-readable storage medium” shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by the machine and that cause the machine to perform the methods described herein. The term “computer-readable storage medium” shall accordingly be taken to include, but not be limited to, solid-state memories, optical media and magnetic media.

Unless specifically stated otherwise, terms such as “determining,” “generating,” “storing,” “calculating,” “requesting,” “allowing,” “preventing,” or the like, refer to actions and processes performed or implemented by computing devices that manipulates and transforms data represented as physical (electronic) quantities within the computing device’s registers and memories into other data similarly represented as physical quantities within the computing device memories or registers or other such information storage, transmission or display devices. Also, the terms “first,” “second,” “third,” “fourth,” etc., as used herein are meant as labels to distinguish among different elements and may not necessarily have an ordinal meaning according to their numerical designation.

Examples described herein also relate to an apparatus for performing the operations described herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computing device selectively programmed by a computer program stored in the computing device. Such a computer program may be stored in a computer-readable non-transitory storage medium.

The methods and illustrative examples described herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used in accordance with the teachings described herein, or it may prove convenient to construct more specialized apparatus to

perform the required method steps. The required structure for a variety of these systems will appear as set forth in the description above.

The above description is intended to be illustrative, and not restrictive. Although the present disclosure has been described with references to specific illustrative examples, it will be recognized that the present disclosure is not limited to the examples described. The scope of the disclosure should be determined with reference to the following claims, along with the full scope of equivalents to which the claims are entitled.

As used herein, the singular forms “a,” “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms “comprises,” “comprising,” “includes,” and/or “including,” when used herein, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof. Therefore, the terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting.

It should also be noted that in some alternative implementations, the functions/acts noted may occur out of the order noted in the figures. For example, two figures shown in succession may in fact be executed substantially concurrently or may sometimes be executed in the reverse order, depending upon the functionality/acts involved.

Although the method operations were described in a specific order, it should be understood that other operations may be performed in between described operations, described operations may be adjusted so that they occur at slightly different times or the described operations may be distributed in a system which allows the occurrence of the processing operations at various intervals associated with the processing.

Various units, circuits, or other components may be described or claimed as “configured to” or “configurable to” perform a task or tasks. In such contexts, the phrase “configured to” or “configurable to” is used to connote structure by indicating that the units/circuits/components include structure (e.g., circuitry) that performs the task or tasks during operation. As such, the unit/circuit/component can be said to be configured to perform the task, or configurable to perform the task, even when the specified unit/circuit/component is not currently operational (e.g., is not on). The units/circuits/components used with the “configured to” or “configurable to” language include hardware—for example, circuits, memory storing program instructions executable to implement the operation, etc. Reciting that a unit/circuit/component is “configured to” perform one or more tasks, or is “configurable to” perform one or more tasks, is expressly intended not to invoke 35 U.S.C. 112, sixth paragraph, for that unit/circuit/component. Additionally, “configured to” or “configurable to” can include generic structure (e.g., generic circuitry) that is manipulated by software and/or firmware (e.g., an FPGA or a general-purpose processor executing software) to operate in manner that is capable of performing the task(s) at issue. “Configured to” may also include adapting a manufacturing process (e.g., a semiconductor fabrication facility) to fabricate devices (e.g., integrated circuits) that are adapted to implement or perform one or more tasks. “Configurable to” is expressly intended not to apply to blank media, an unprogrammed processor or unprogrammed generic computer, or an unprogrammed programmable logic device, programmable gate array, or other



unprogrammed device, unless accompanied by programmed media that confers the ability to the unprogrammed device to be configured to perform the disclosed function(s).

The foregoing description, for the purpose of explanation, has been described with reference to specific embodiments. However, the illustrative discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The embodiments were chosen and described in order to best explain the principles of the embodiments and its practical applications, to thereby enable others skilled in the art to best utilize the embodiments and various modifications as may be suited to the particular use contemplated. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

What is claimed is:

1. A computer-implemented method comprising:
  - receiving, from a client device, a request to access a resource of a computer system;
  - determining one or more roles of a user associated with the client device, wherein each of the one or more roles has a corresponding set of one or more permissions;
  - based on the one or more roles of the user, identifying, as being applicable to the received request, at least one predefined access permission for performing a given action on a given resource, wherein the predefined access permission comprises (i) a definition of the given action and (ii) a definition of the given resource, and wherein at least one of the definition of the given action or the definition of the given resource includes a given attribute name that serves as a placeholder to be replaced by a corresponding attribute value prior to using the predefined access permission to make an access control decision;
  - dynamically determining an attribute value that corresponds to the given attribute name included in the predefined access permission, wherein the dynamically-determined attribute value comprises at least one of (i) information about the user associated with the client device from which the request was received, (ii) information about the resource for which access is being requested, or (iii) contextual information about the receipt of request;
  - dynamically generating a modified access permission that is specific to the received request by replacing the given attribute name included in the predefined access permission with the dynamically-determined attribute value; and
  - providing or denying access to the resource of the computer system based on the modified access permission.
2. The computer-implemented method of claim 1, further comprising:
  - generating an access policy comprising the modified access permission; and
  - providing the access policy to the computer system.
3. The computer-implemented method of claim 1, wherein the request for the resource comprises a resource path.
4. The computer-implemented method of claim 1, wherein the definition of the given action comprises a constraint on the given action that is permitted to be performed on the given resource.

5. The computer-implemented method of claim 1, wherein the predefined permission is retained for use in handling future access requests.

6. The computer-implemented method of claim 1, wherein the information about the user associated with the client device from which the request was received comprises an identifier of the user.

7. The computer-implemented method of claim 1, wherein the information about the user associated with the client device from which the request was received comprises information about a status of the user in an organizational hierarchy, wherein the status of the user in the organizational hierarchy is distinct from the one or more roles of the user.

8. The computer-implemented method of claim 7, wherein the information about the status of the user in the organizational hierarchy comprises an indication of one or both of an organization within the organizational hierarchy to which the user belongs or a sub organization within the organizational hierarchy to which the user belongs.

9. The computer-implemented method of claim 7, wherein the organizational hierarchy is represented by one of (i) a tree-type data structure or (ii) a directed acyclic graph (DAC)-type data structure.

10. The computer-implemented method of claim 1, wherein the given attribute name comprises a first attribute name in an attribute-based expression that also includes at least a second attribute name, wherein the method further comprises dynamically determining an attribute value that corresponds to the second attribute name, and wherein dynamically generating the modified access permission further comprises replacing the second attribute name in the attribute-based expression with the dynamically-determined attribute value that corresponds to the second attribute name.

11. The computer-implemented method of claim 1, wherein the contextual information about the receipt of request comprises information about when the request was received.

12. A computing platform comprising:

a network interface;

at least one processor; and

at least one computer-readable storage medium comprising program instructions that are executable by the at least one processor such that the computing platform is configured to:

receive, from a client device via the network interface,

a request to access a resource of a computer system; determine one or more roles of a user associated with the client device, wherein each of the one or more roles has a corresponding set of one or more permissions;

based on the one or more roles of the user, identify, as being applicable to the received request, at least one predefined access permission for performing a given action on a given resource, wherein the predefined access permission comprises (i) a definition of the given action and (ii) a definition of the given resource, and wherein at least one of the definition of the given action or the definition of the given resource includes a given attribute name that serves as a placeholder to be replaced by a corresponding attribute value prior to using the predefined access permission to make an access control decision;

dynamically determine an attribute value that corresponds to the given attribute name included in the predefined access permission, wherein the dynamically-determined attribute value comprises at least one of (i) information about the user associated with



23

the client device from which the request was received, (ii) information about the resource for which access is being requested, or (iii) contextual information about the receipt of request;

dynamically generate a modified access permission that is specific to the received request by replacing the given attribute name included in the predefined access permission with the dynamically-determined attribute value; and

provide or deny access to the resource of the computer system based on the modified access permission.

13. The computing platform of claim 12, wherein the at least one computer-readable storage medium further comprises program instructions that are executable by the at least one processor such that the computing platform is configured to:

generate an access policy comprising the modified access permission; and

provide the access policy to the computer system.

14. The computing platform of claim 12, wherein the request for the resource comprises a resource path.

15. The computing platform of claim 12, wherein the definition of the given action comprises a constraint on the given action that is permitted to be performed on the given resource.

16. The computing platform of claim 12, the predefined permission is retained for use in handling future access requests.

17. The computing platform of claim 12, wherein the information about the user associated with the client device from which the request was received comprises an identifier of the user.

18. The computing platform of claim 12, wherein the information about the user associated with the client device from which the request was received comprises information about a status of the user in an organizational hierarchy, wherein the status of the user in the organizational hierarchy is distinct from the one or more roles of the user.

19. The computing platform of claim 18, wherein the information about the status of the user in the organizational hierarchy comprises an indication of one or both of an organization within the organizational hierarchy to which the user belongs or a sub organization within the organizational hierarchy to which the user belongs.

20. The computing platform of claim 18, wherein the organizational hierarchy is represented by one of (i) a tree-type data structure or (ii) a directed acyclic graph (DAC)-type data structure.

21. The computing platform of claim 12, wherein the given attribute name comprises a first attribute name in an attribute-based expression that also includes at least a second attribute name, wherein the at least one computer-readable storage medium further comprises program instruc-

24

tions that are executable by the at least one processor such that the computing platform is configured to dynamically determine an attribute value that corresponds to the second attribute name, and wherein the program instructions that are executable by the at least one processor such that the computing platform is configured to dynamically generate the modified access permission further comprise program instructions that are executable by the at least one processor such that the computing platform is configured to replace the second attribute name in the attribute-based expression with the dynamically-determined attribute value that corresponds to the second attribute name.

22. The computing platform of claim 12, wherein the contextual information about the receipt of request comprises information about when the request was received.

23. A non-transitory computer readable medium comprising instructions that, when executed by at least one processor of a computing platform, cause the computing platform to:

receive, from a client device, a request to access a resource of a computer system;

determine one or more roles of a user associated with the client device, wherein each of the one or more roles has a corresponding set of one or more permissions;

based on the one or more roles of the user, identify, as being applicable to the received request, at least one predefined access permission for performing a given action on a given resource, wherein the predefined access permission comprises (i) a definition of the given action and (ii) a definition of the given resource, and wherein at least one of the definition of the given action or the definition of the given resource includes a given attribute name that serves as a placeholder to be replaced by a corresponding attribute value prior to using the predefined access permission to make an access control decision;

dynamically determine an attribute value that corresponds to the given attribute name included in the predefined access permission, wherein the dynamically-determined attribute value comprises at least one of (i) information about the user associated with the client device from which the request was received, (ii) information about the resource for which access is being requested, or (iii) contextual information about the receipt of request;

dynamically generate a modified access permission that is specific to the received request by replacing the given attribute name included in the predefined access permission with the dynamically-determined attribute value; and

provide or deny access to the resource of the computer system based on the modified access permission.

\* \* \* \* \*