



US 20190319967A1

(19) **United States**

(12) **Patent Application Publication**
Holt et al.

(10) **Pub. No.: US 2019/0319967 A1**

(43) **Pub. Date: Oct. 17, 2019**

(54) **SYSTEM FOR EFFICIENT MANAGEMENT AND STORAGE OF ACCESS TOKENS**

(71) Applicant: **Barclays Services Limited**, London (GB)

(72) Inventors: **Dickon Holt**, Knutsford (GB); **Michael Forrest**, Knutsford (GB)

(73) Assignee: **Barclays Services Limited**, London (GB)

(21) Appl. No.: **16/382,148**

(22) Filed: **Apr. 11, 2019**

(30) **Foreign Application Priority Data**
Apr. 11, 2018 (EP) 18166865.8

Publication Classification

(51) **Int. Cl.**
H04L 29/06 (2006.01)

(52) **U.S. Cl.**
CPC **H04L 63/108** (2013.01); **H04L 63/0853** (2013.01)

(57) **ABSTRACT**
A client system comprises processing circuitry configured to receive, from a user device, a first access request comprising a first instruction to access a protected resource; transmit a token request for an access token to be used for accessing the protected resource; receive an access token in response to the token request, the access token having a corresponding time to expire indicative of a time at which the access token will not be valid for obtaining the protected resource; and transmit the access token to a resource system and, in response, receive the protected resource. The client system further comprises a token storage unit configured to store the access token.

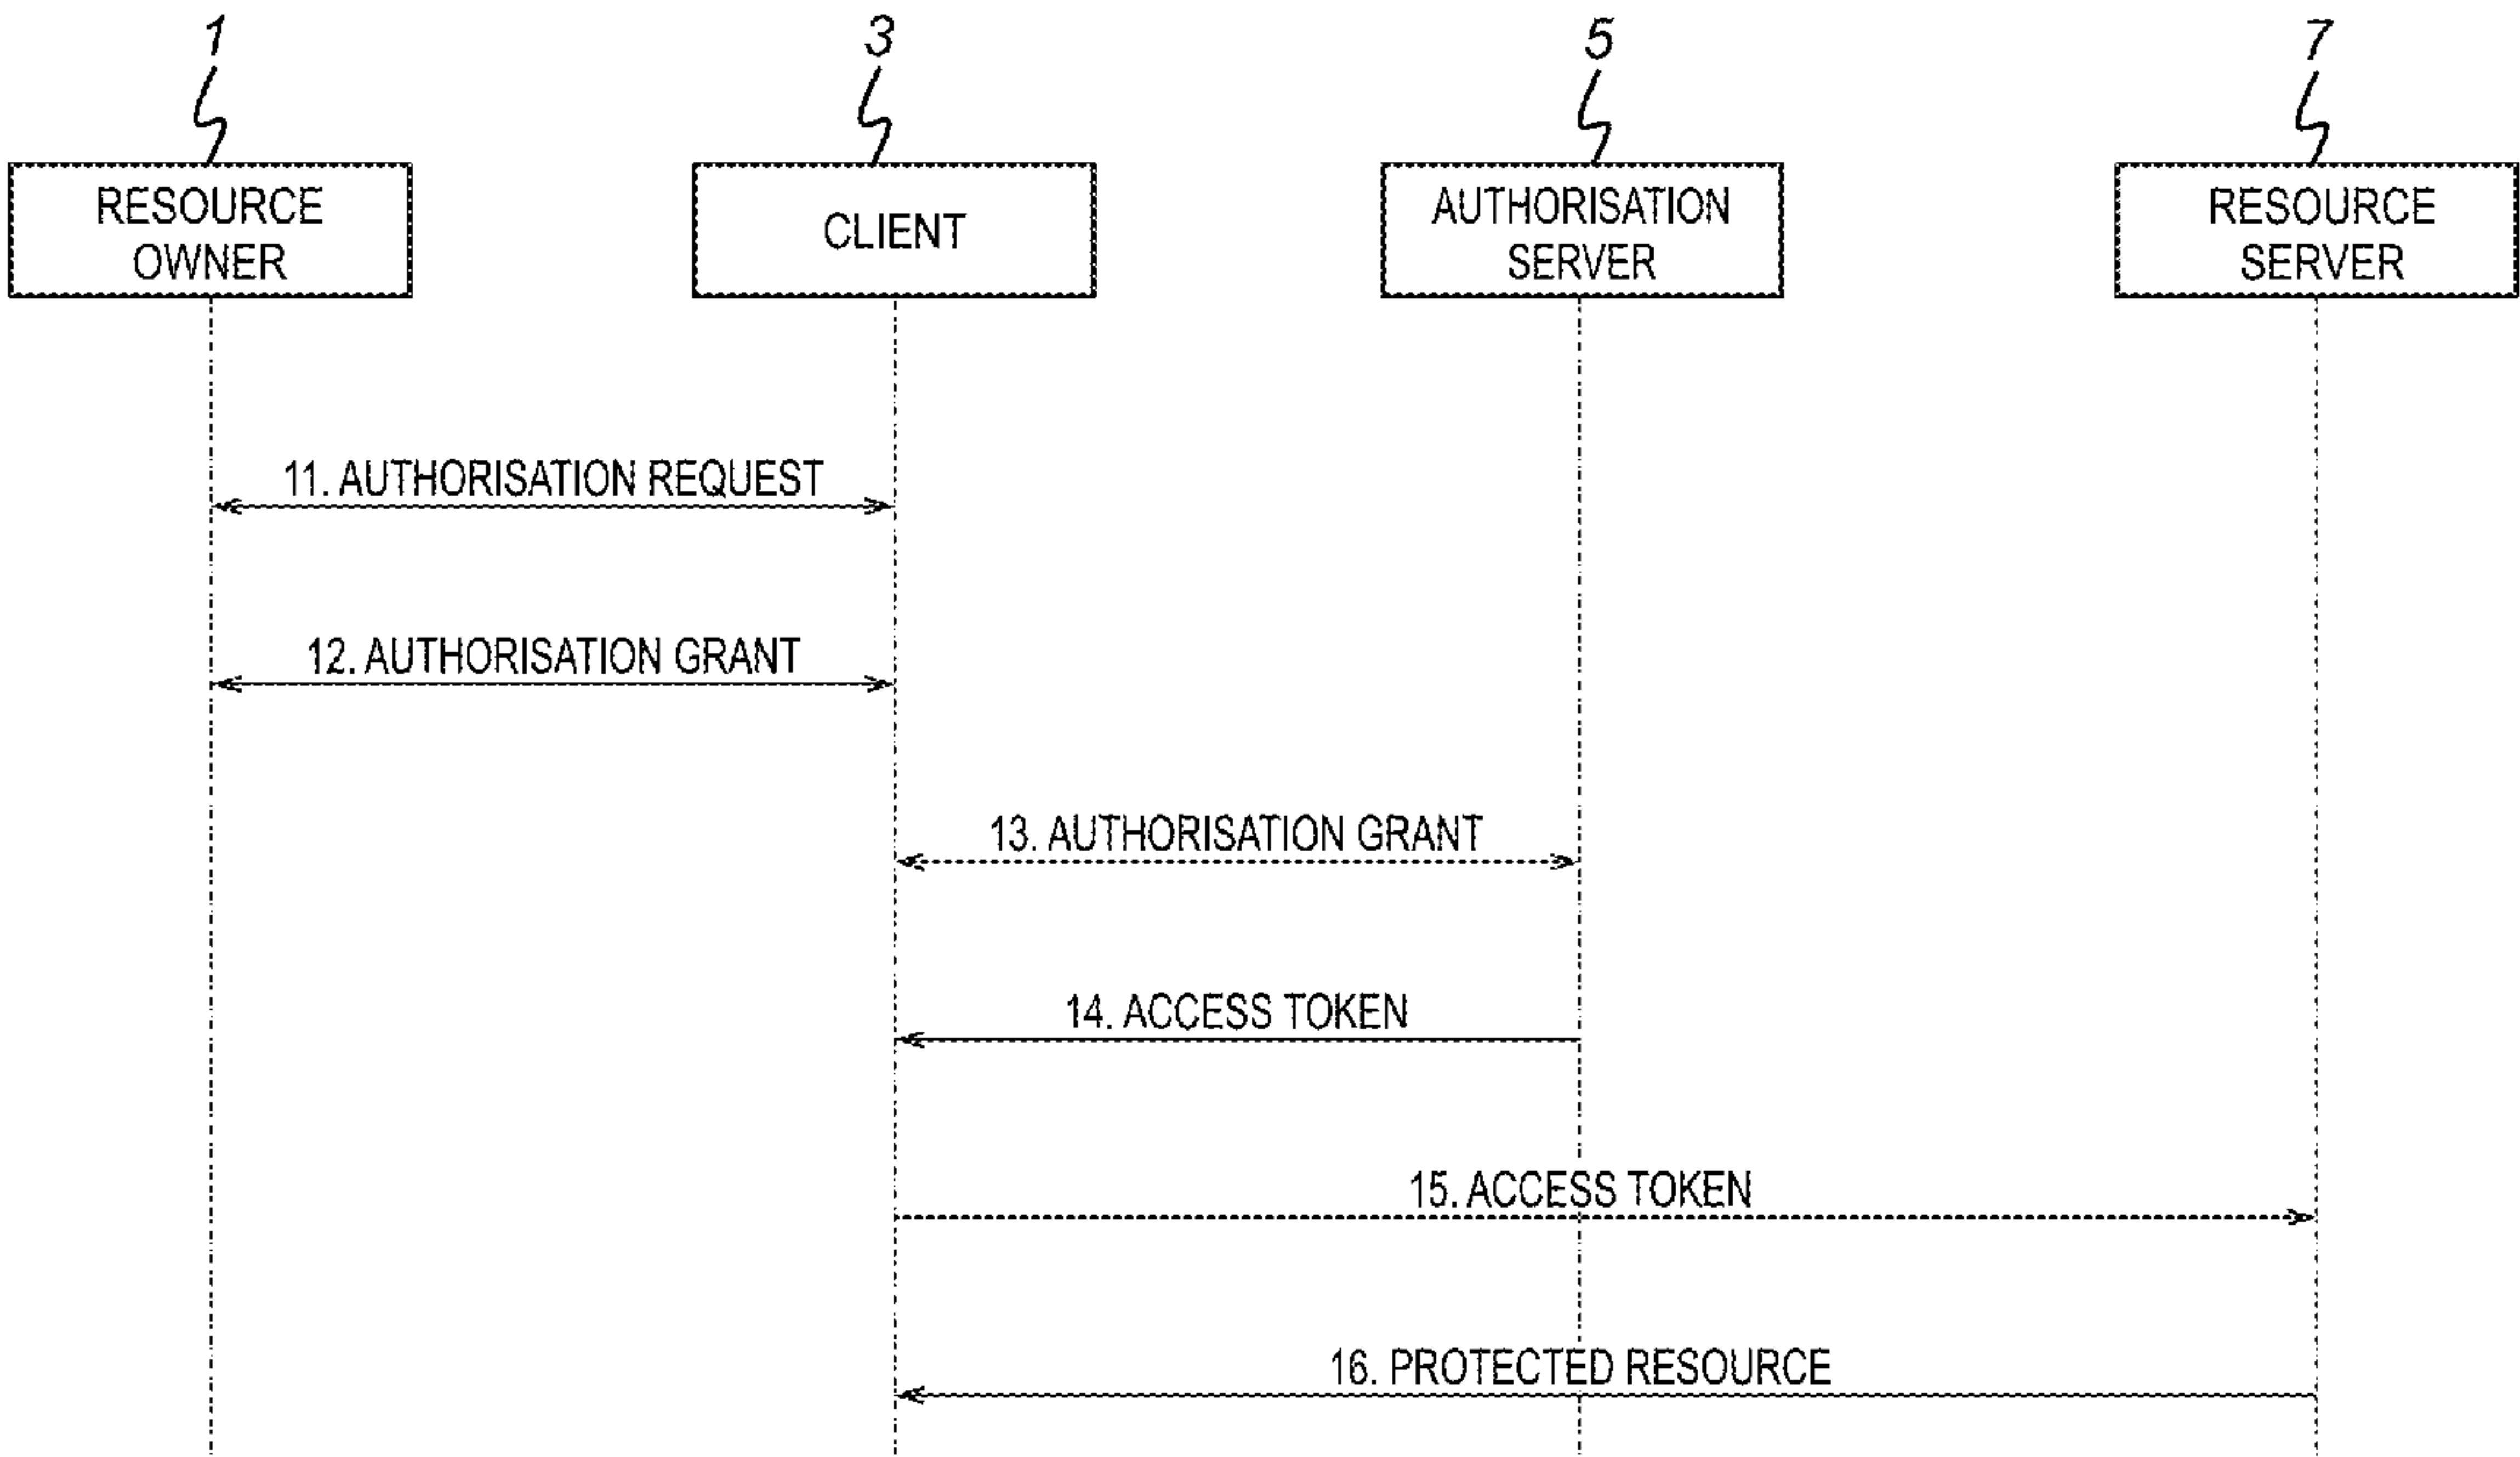


FIG. 1

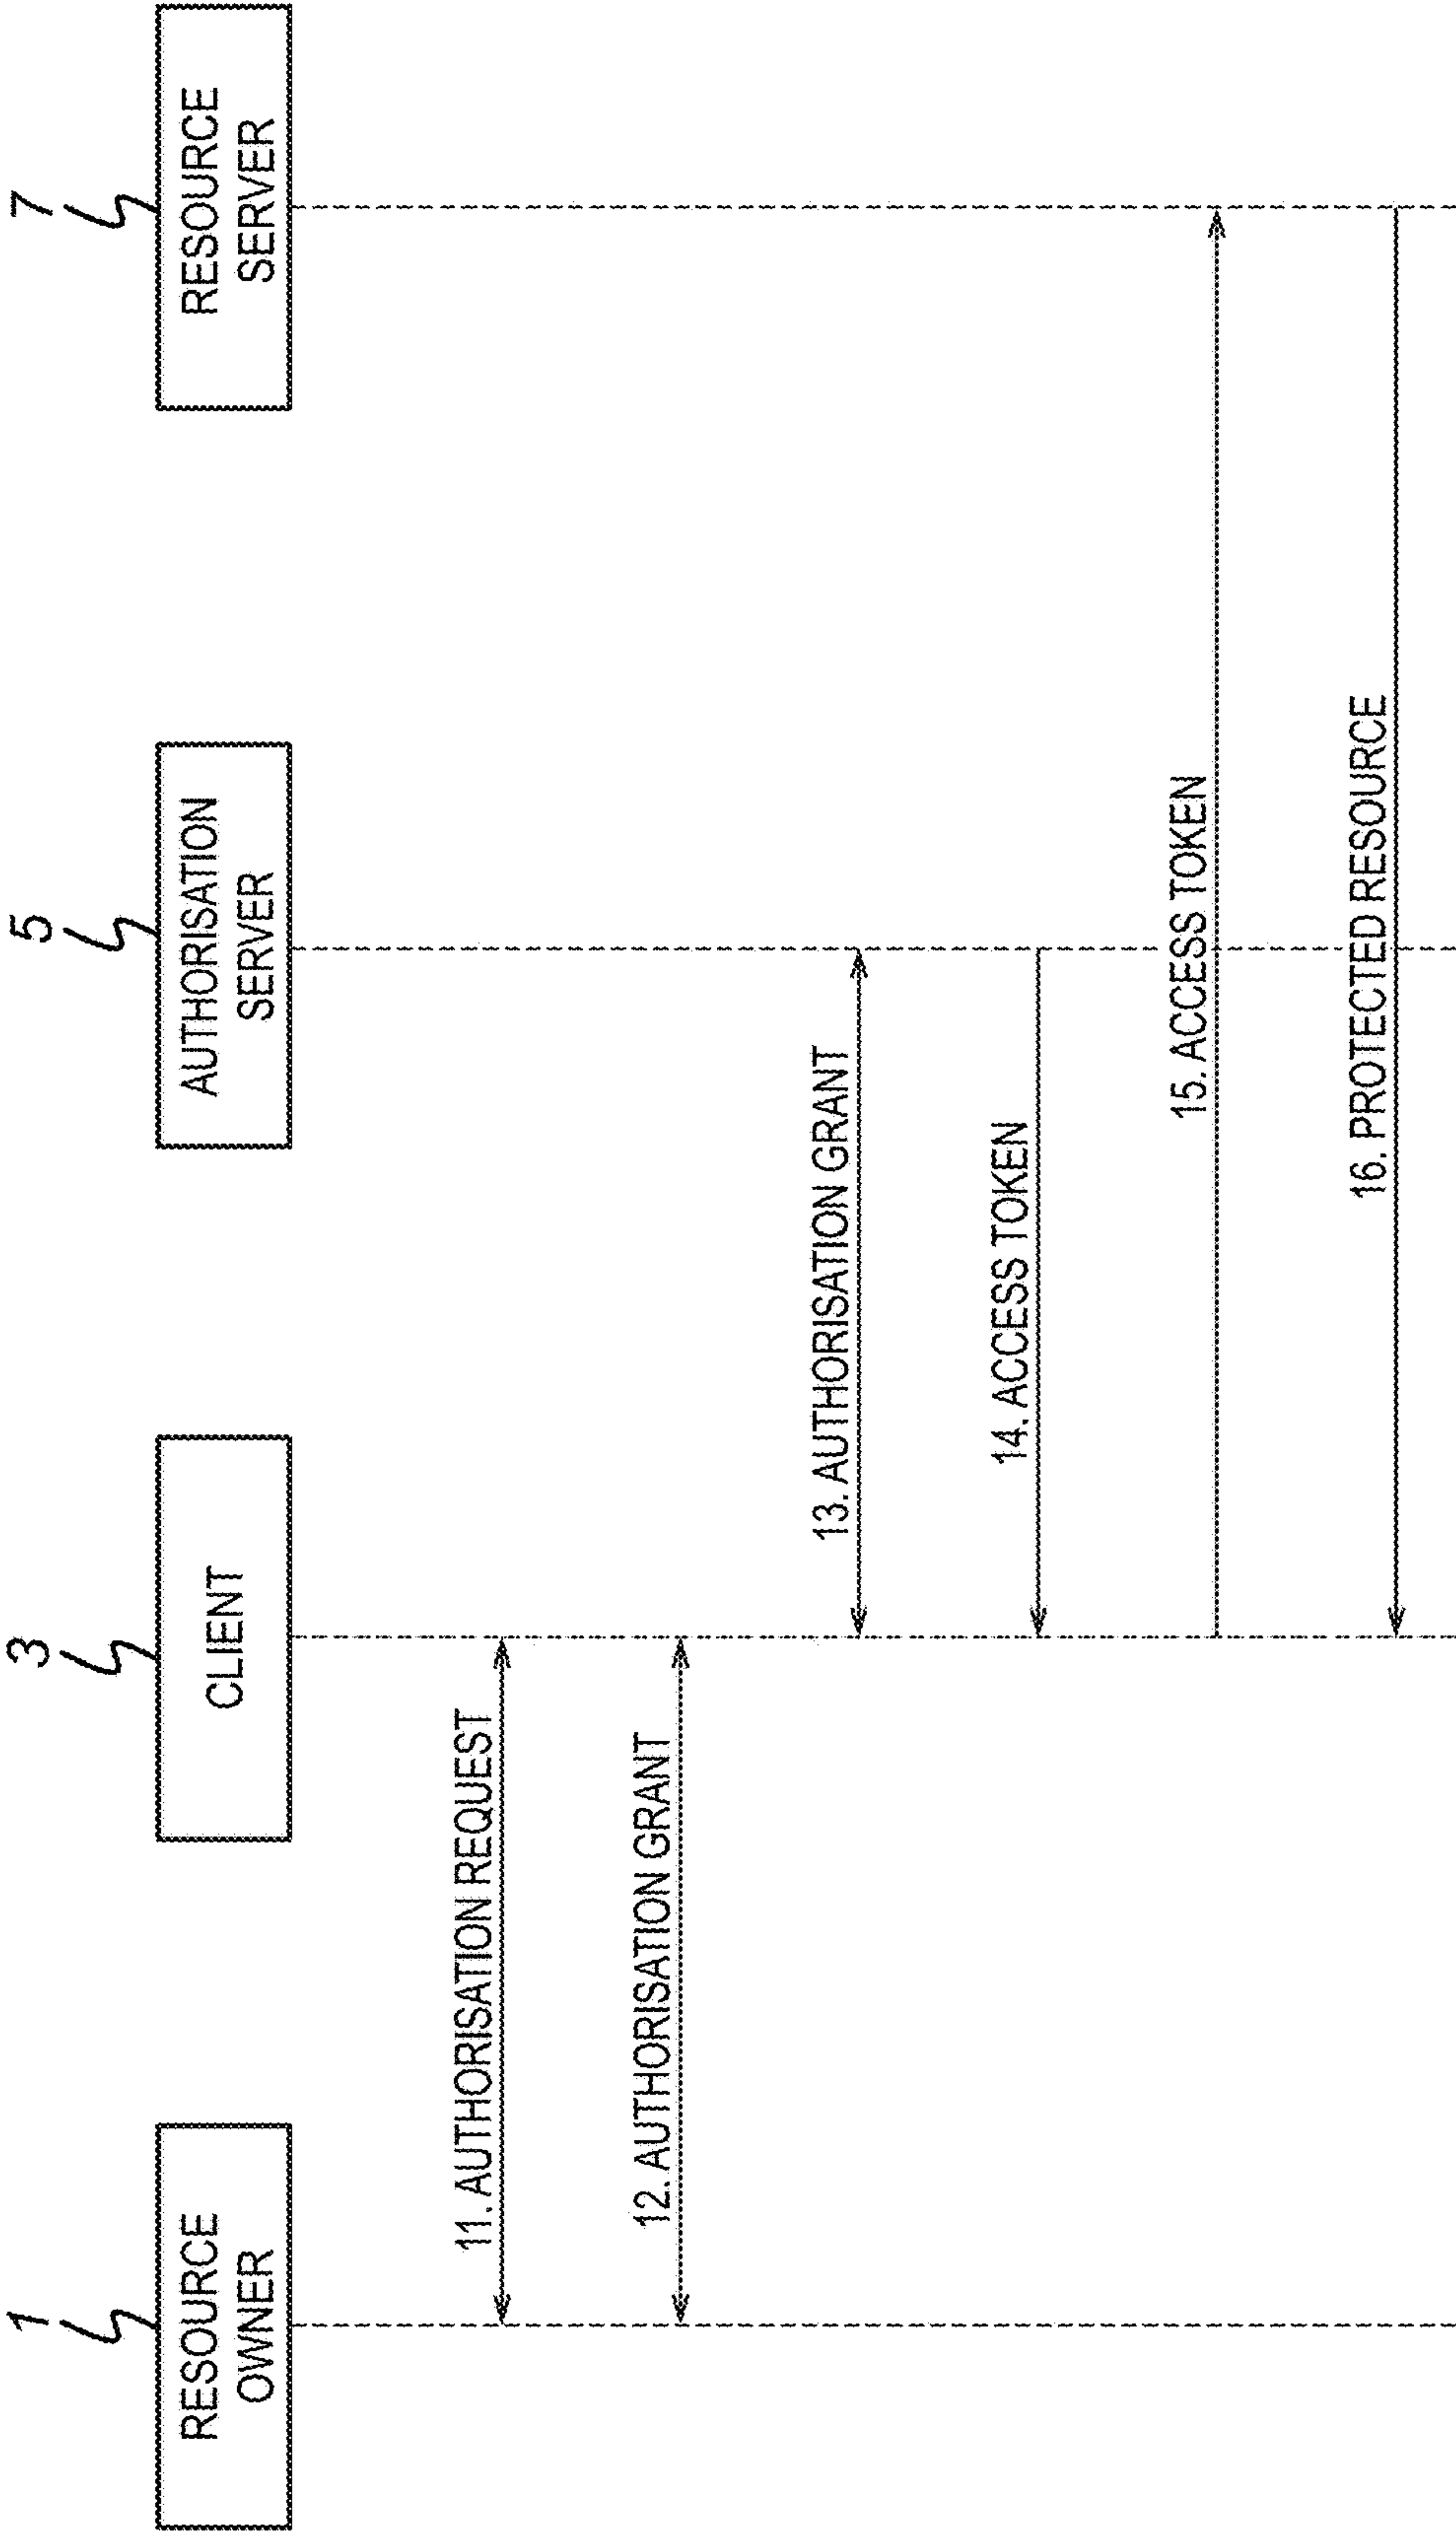
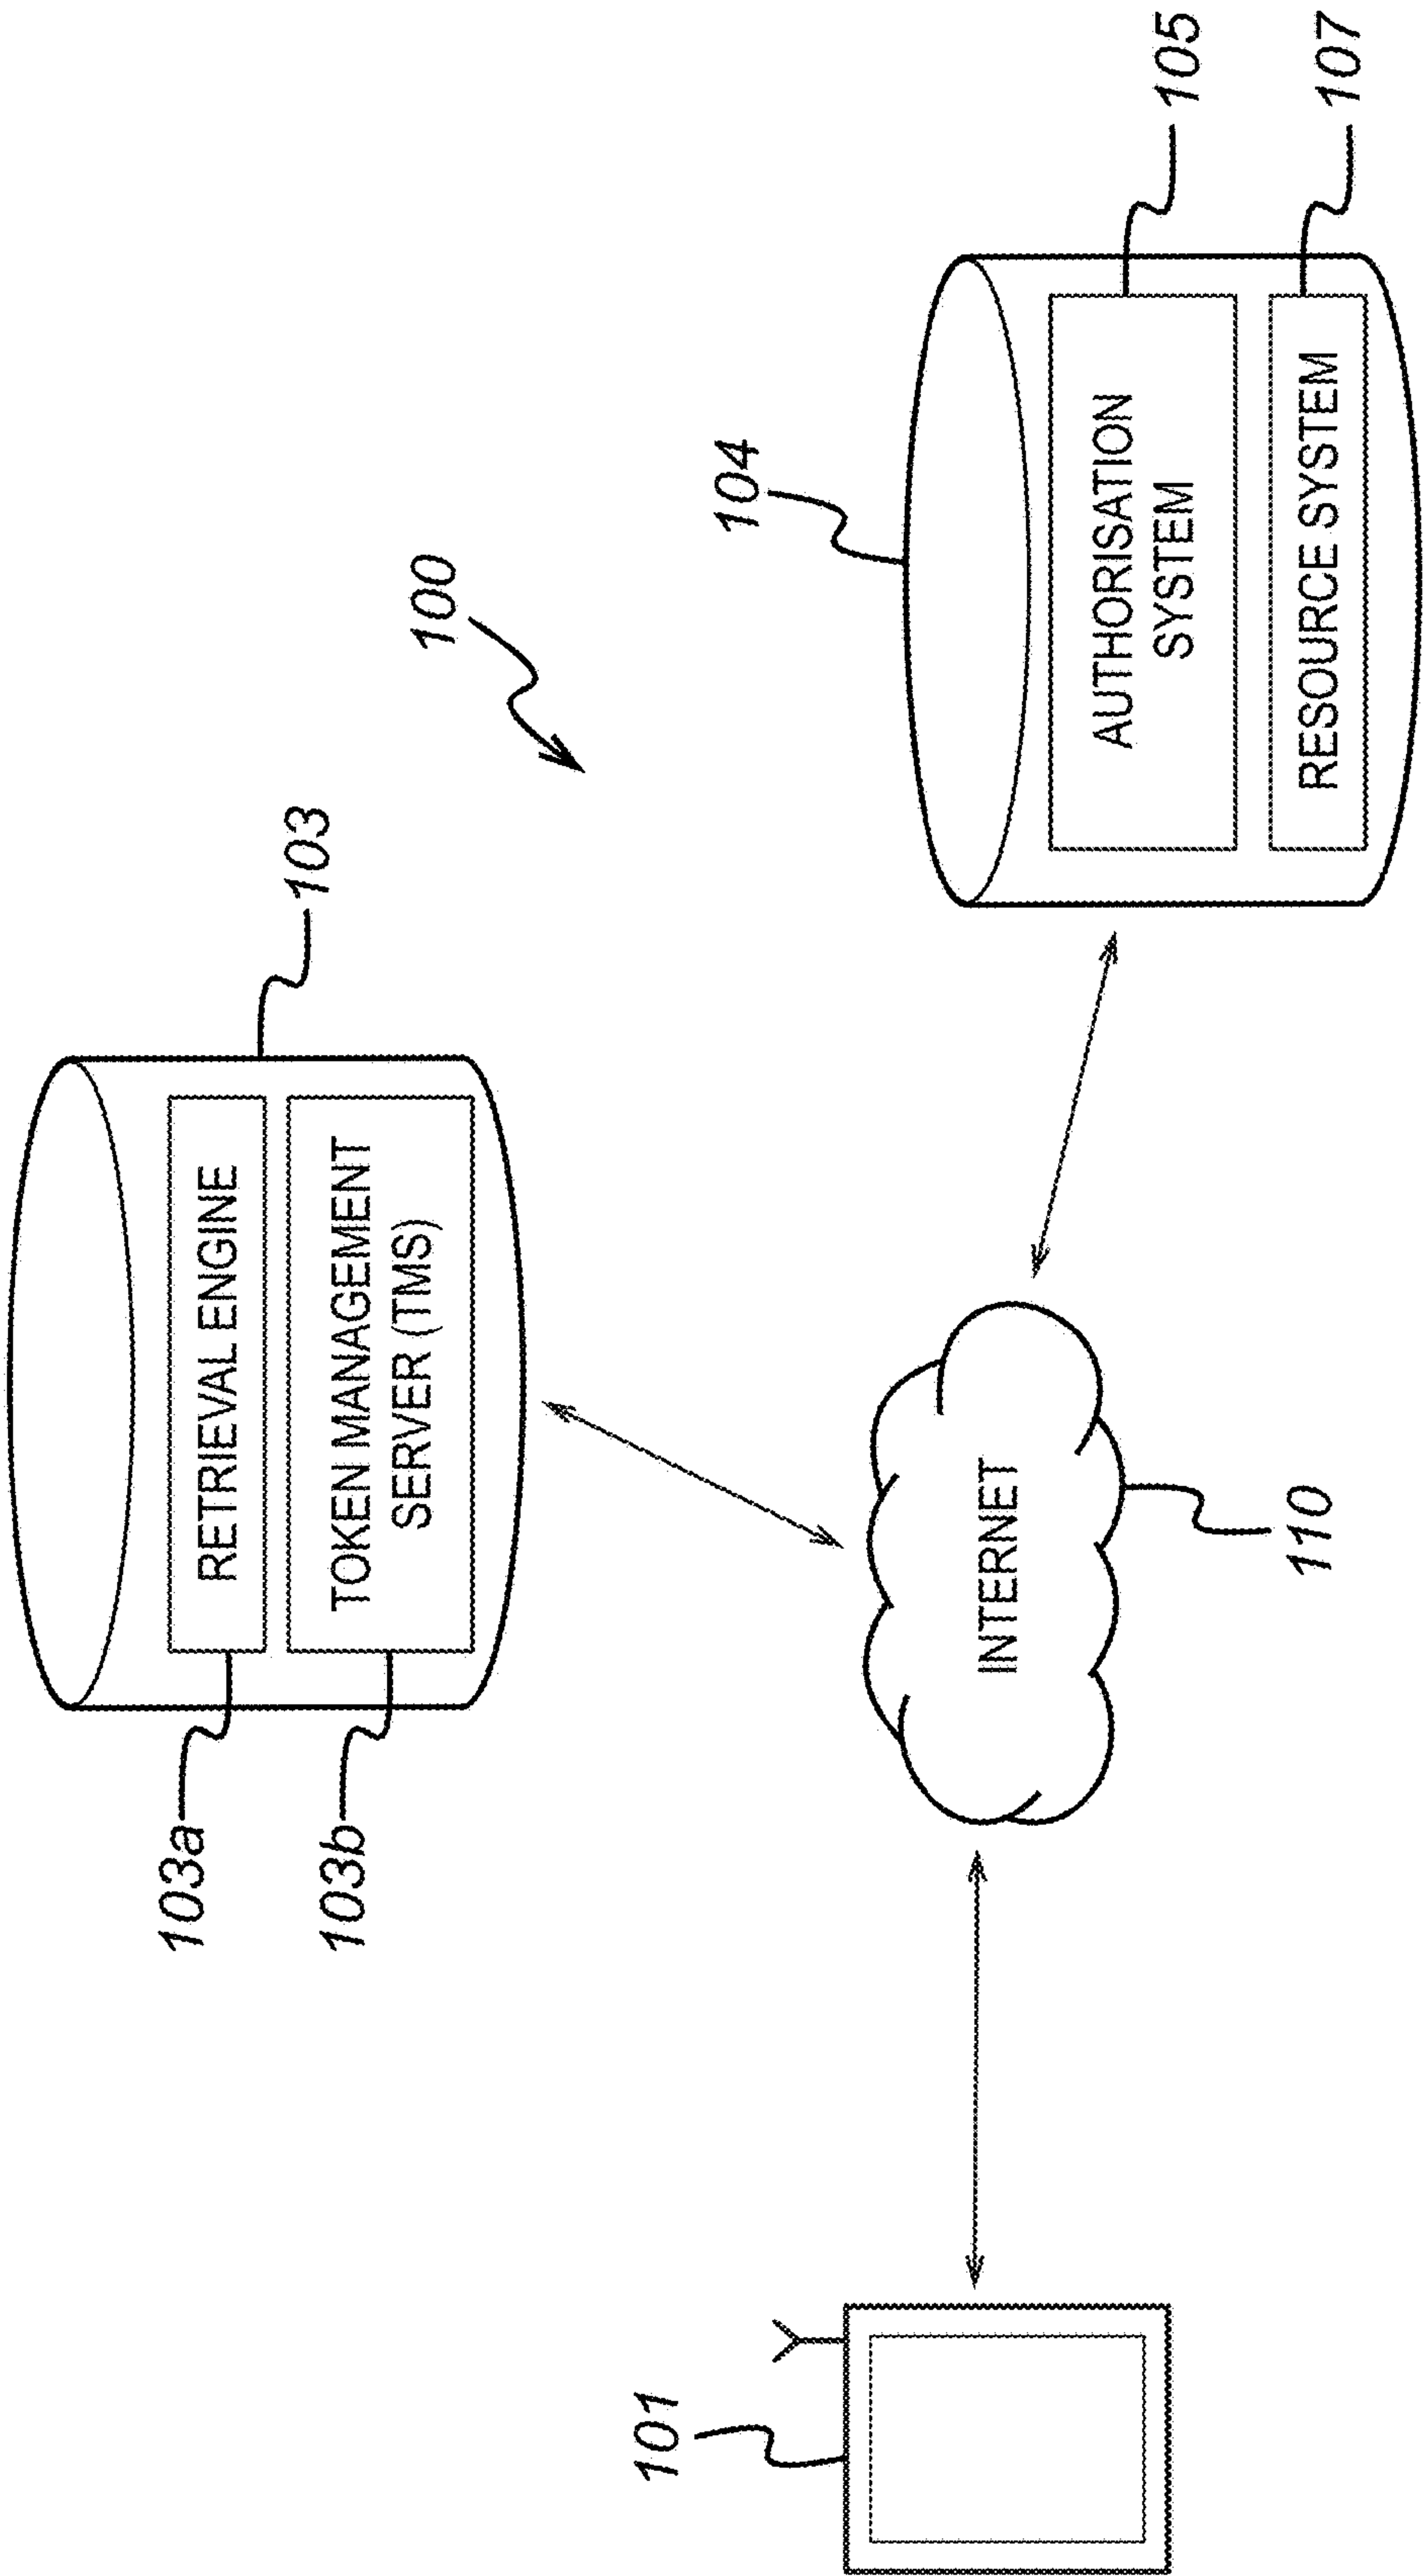


FIG. 2



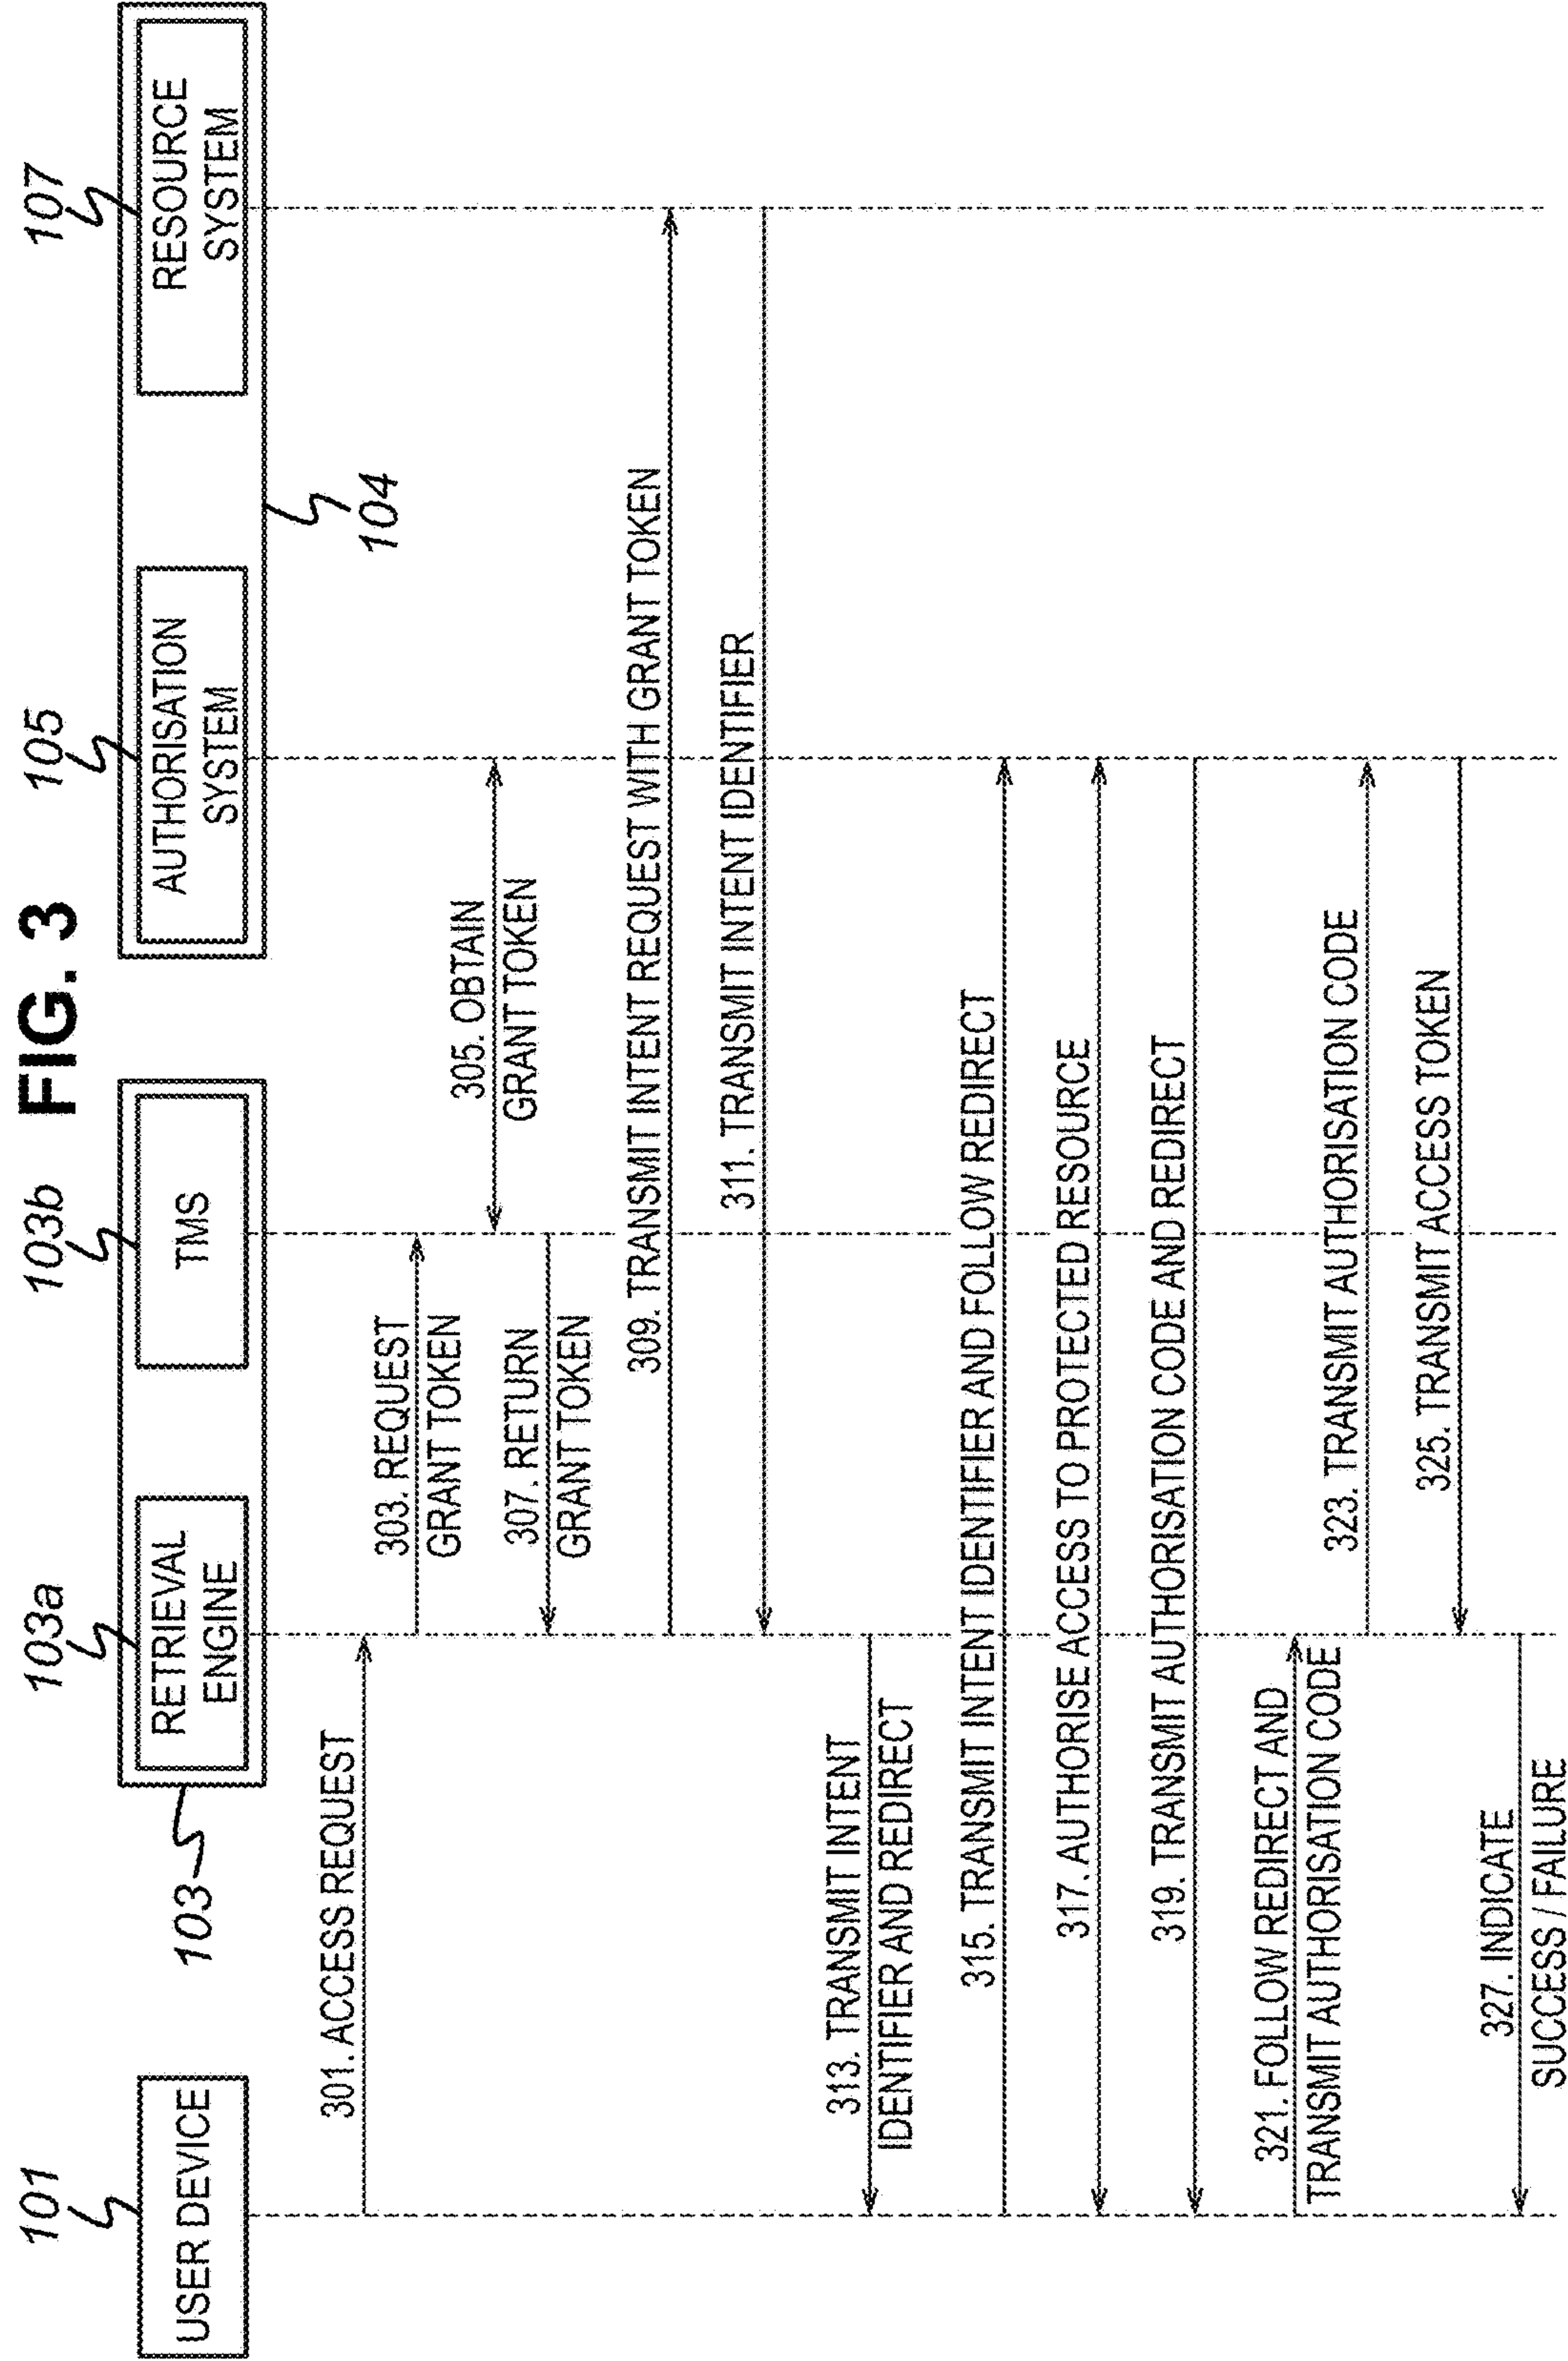
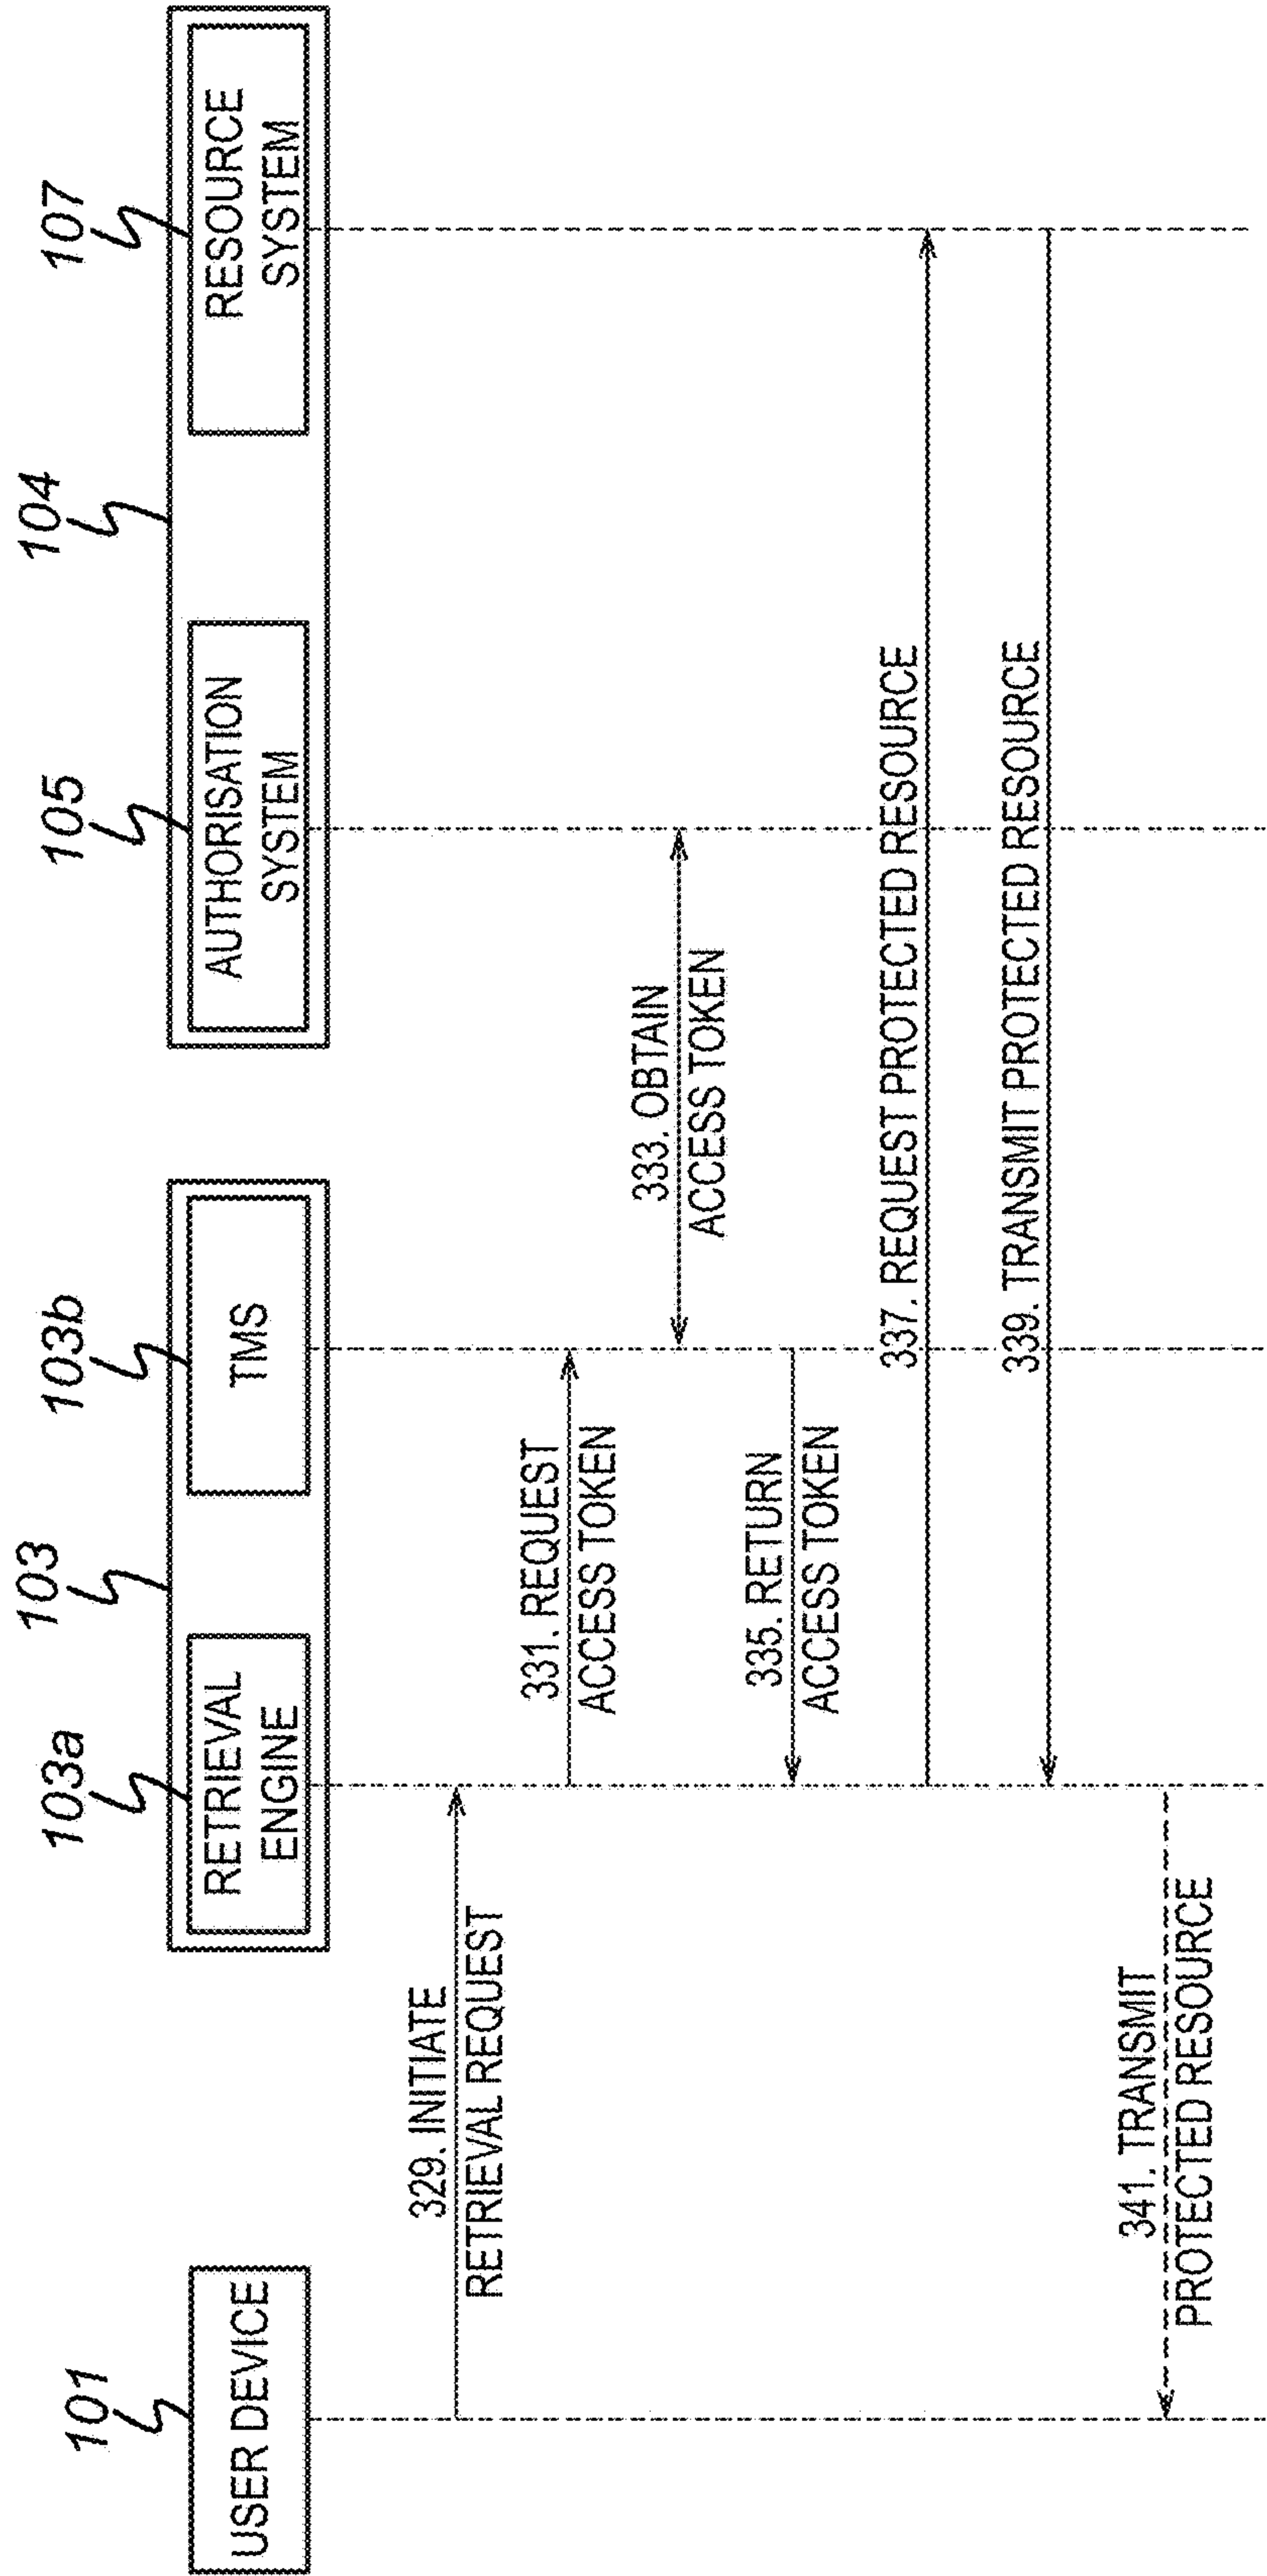


FIG. 4



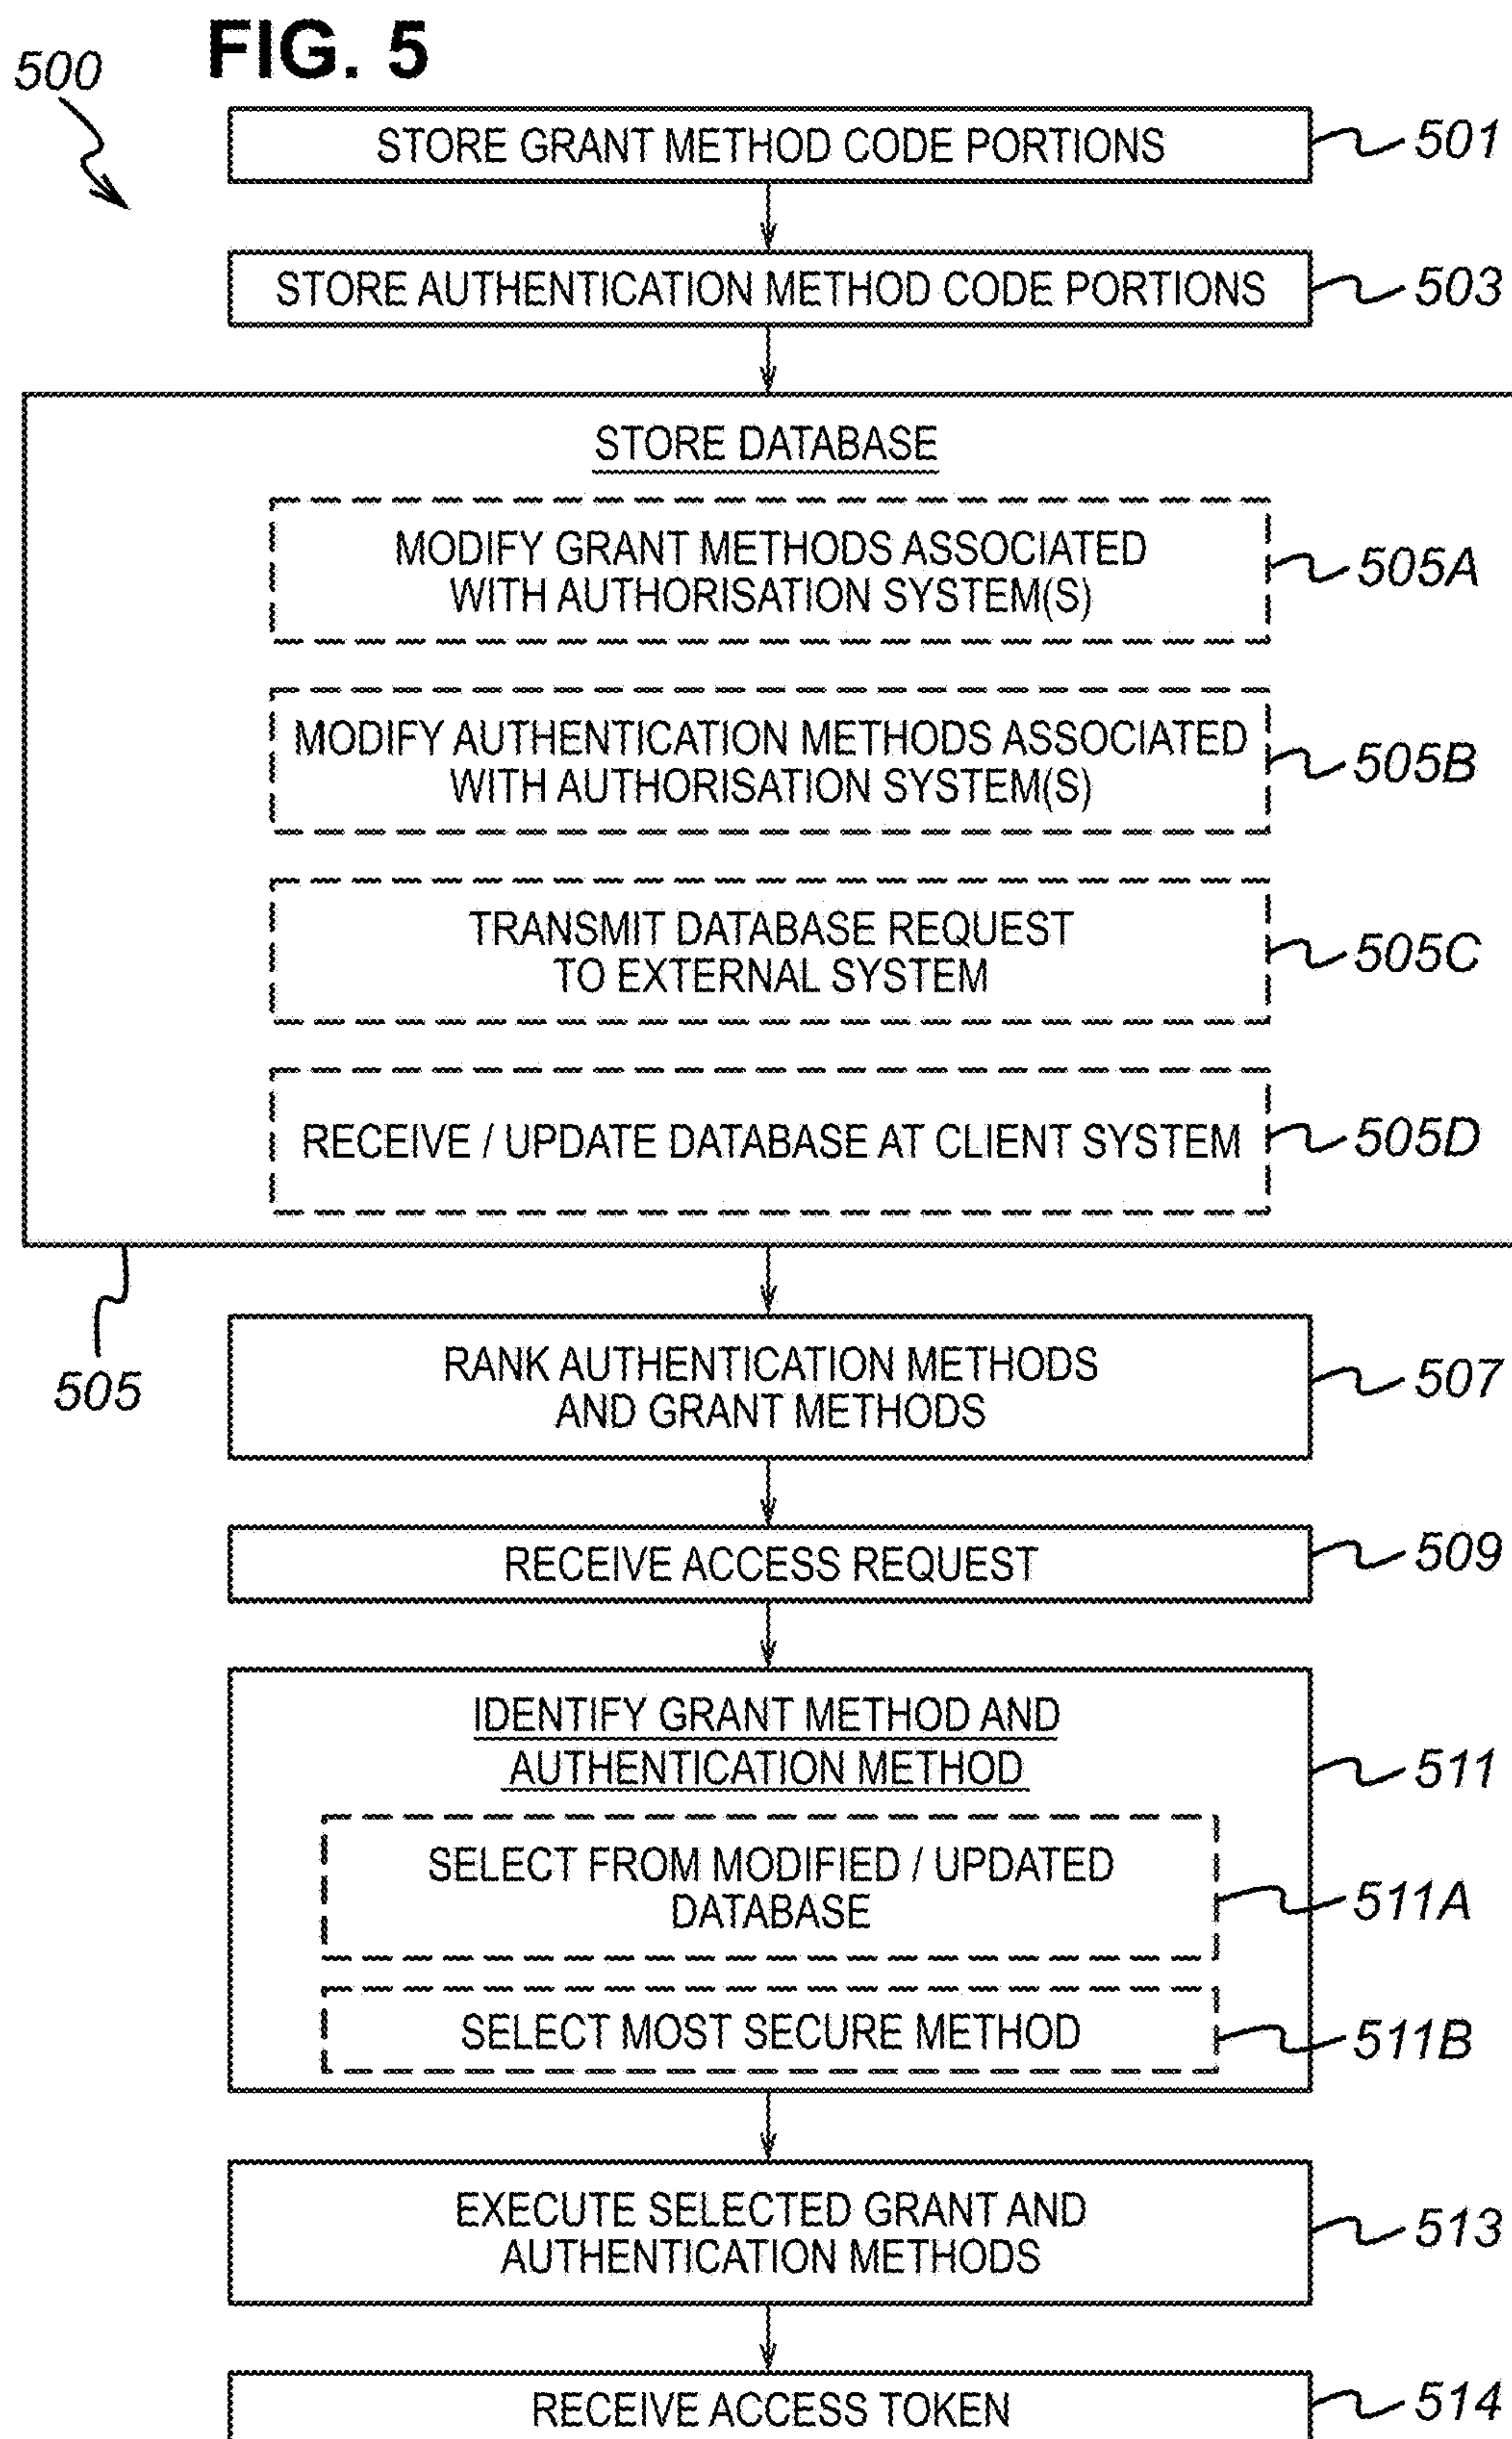


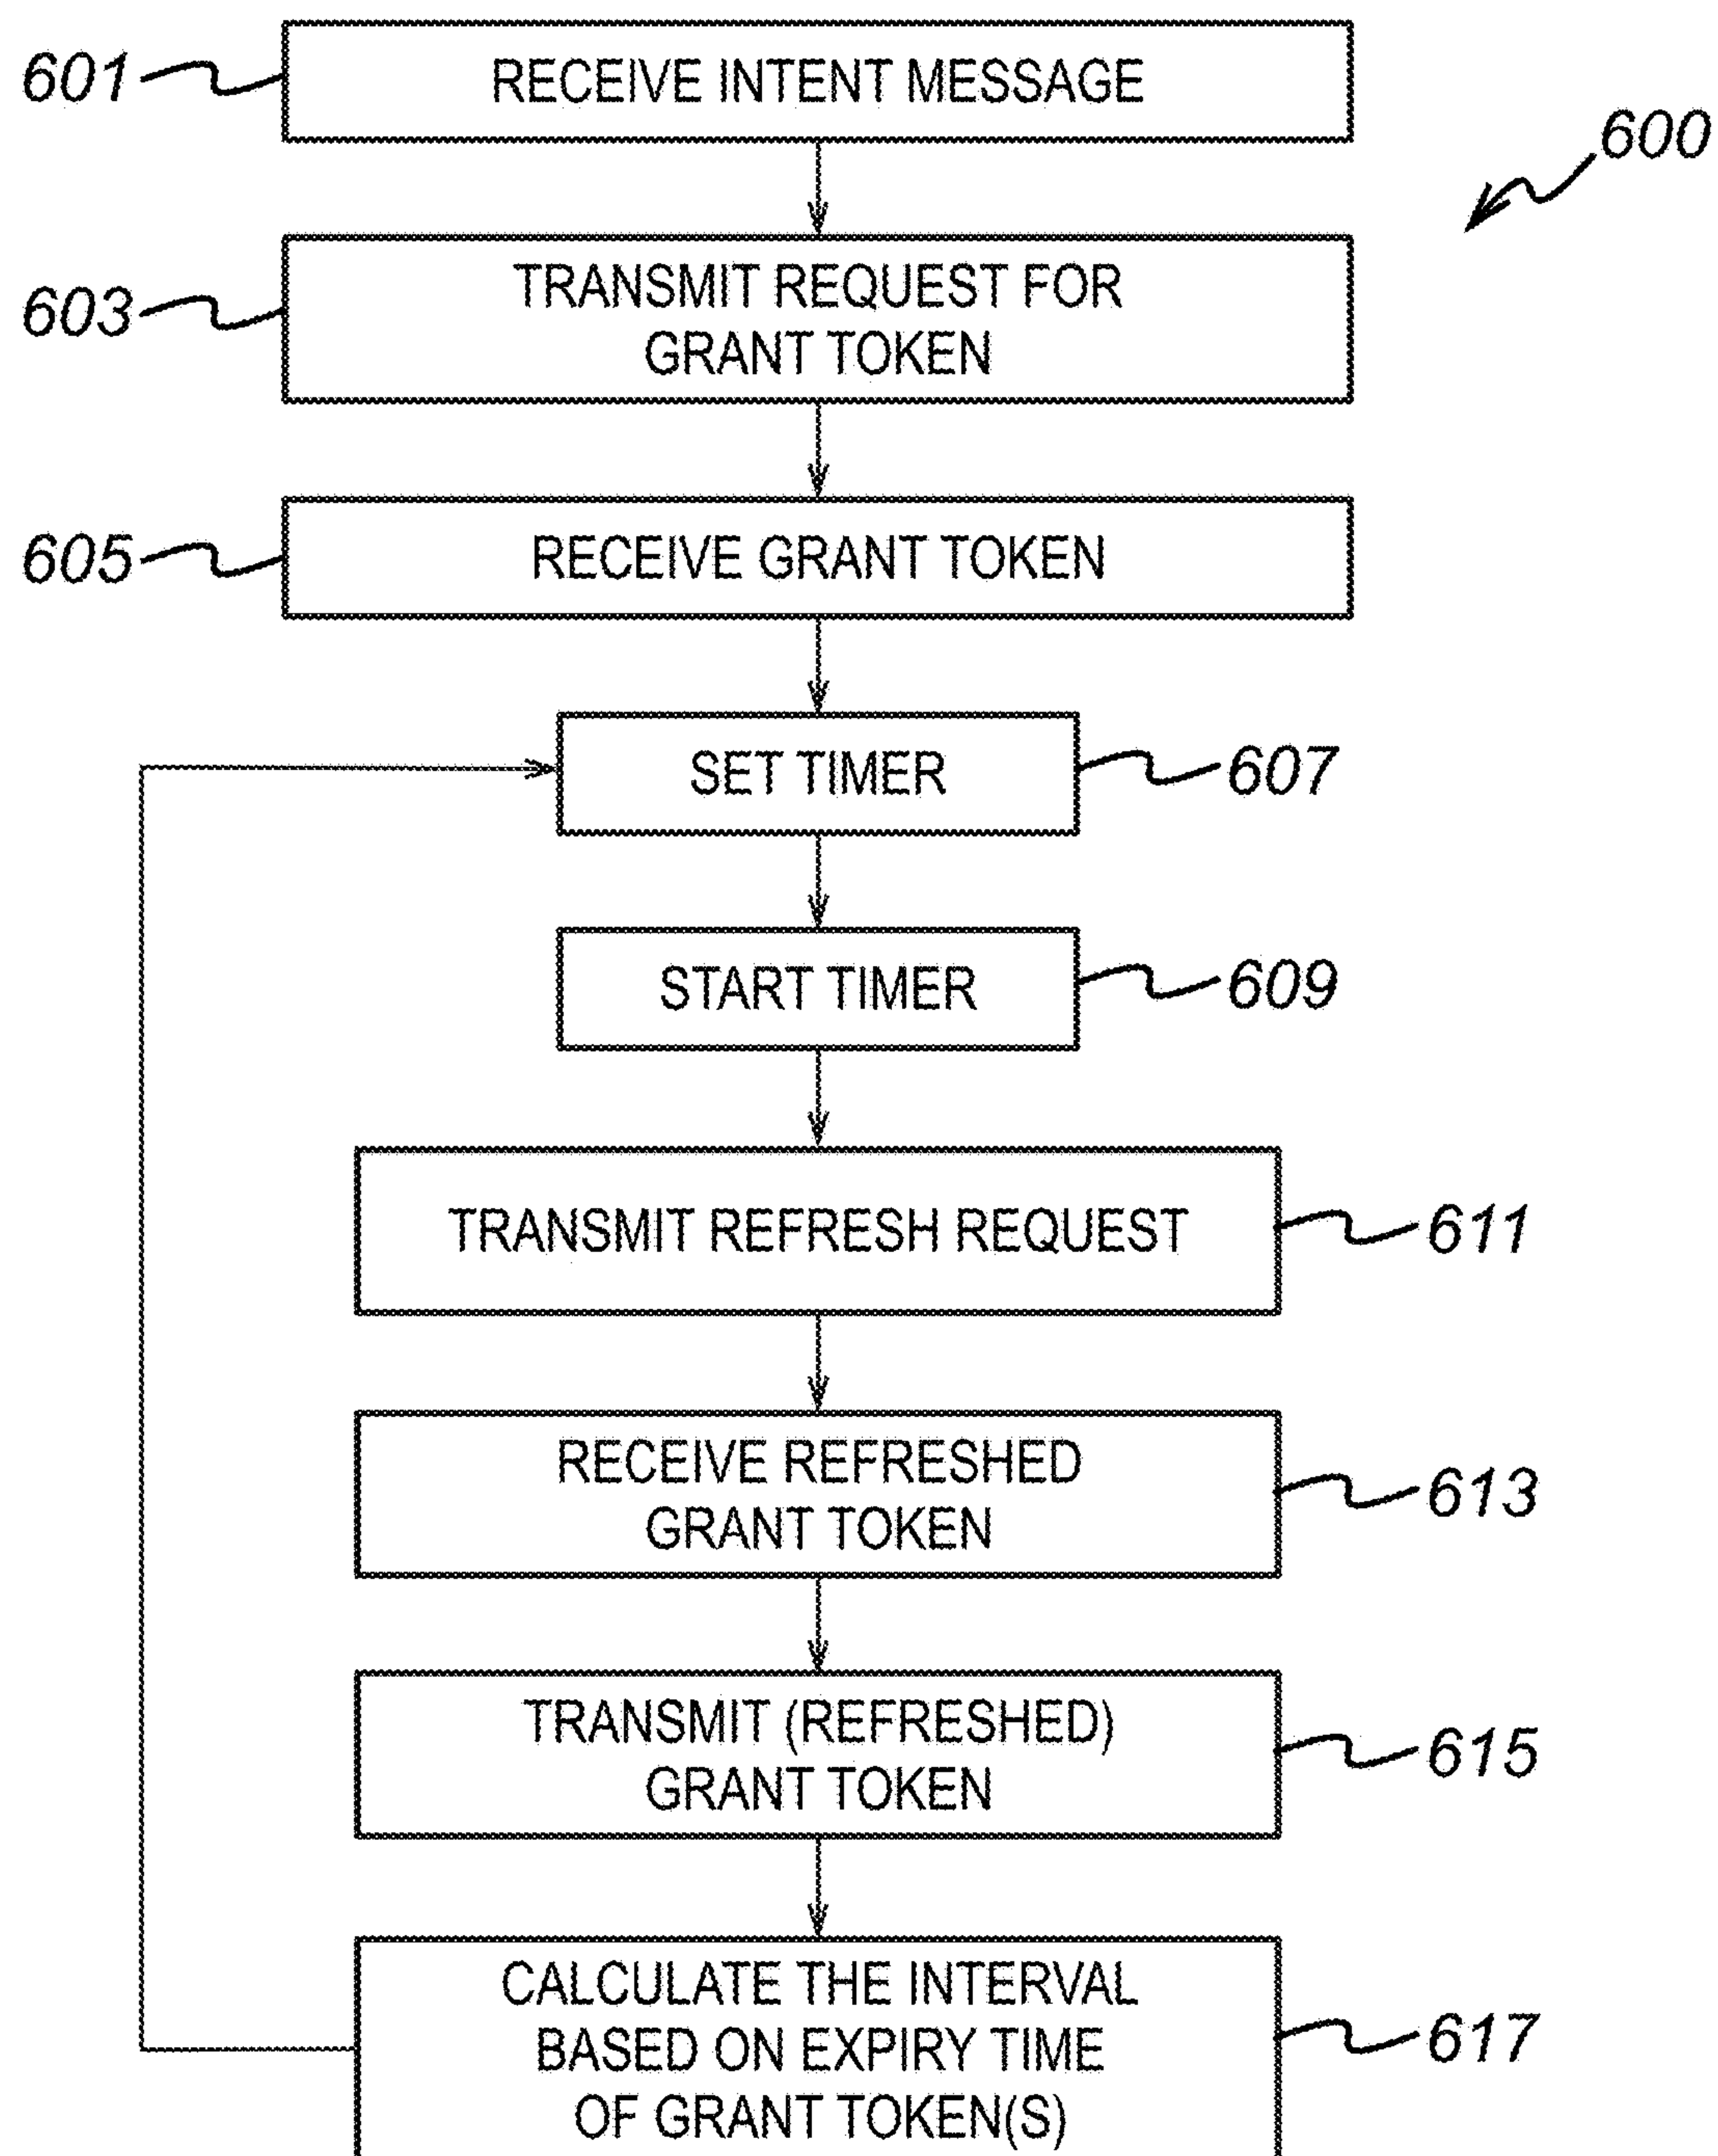
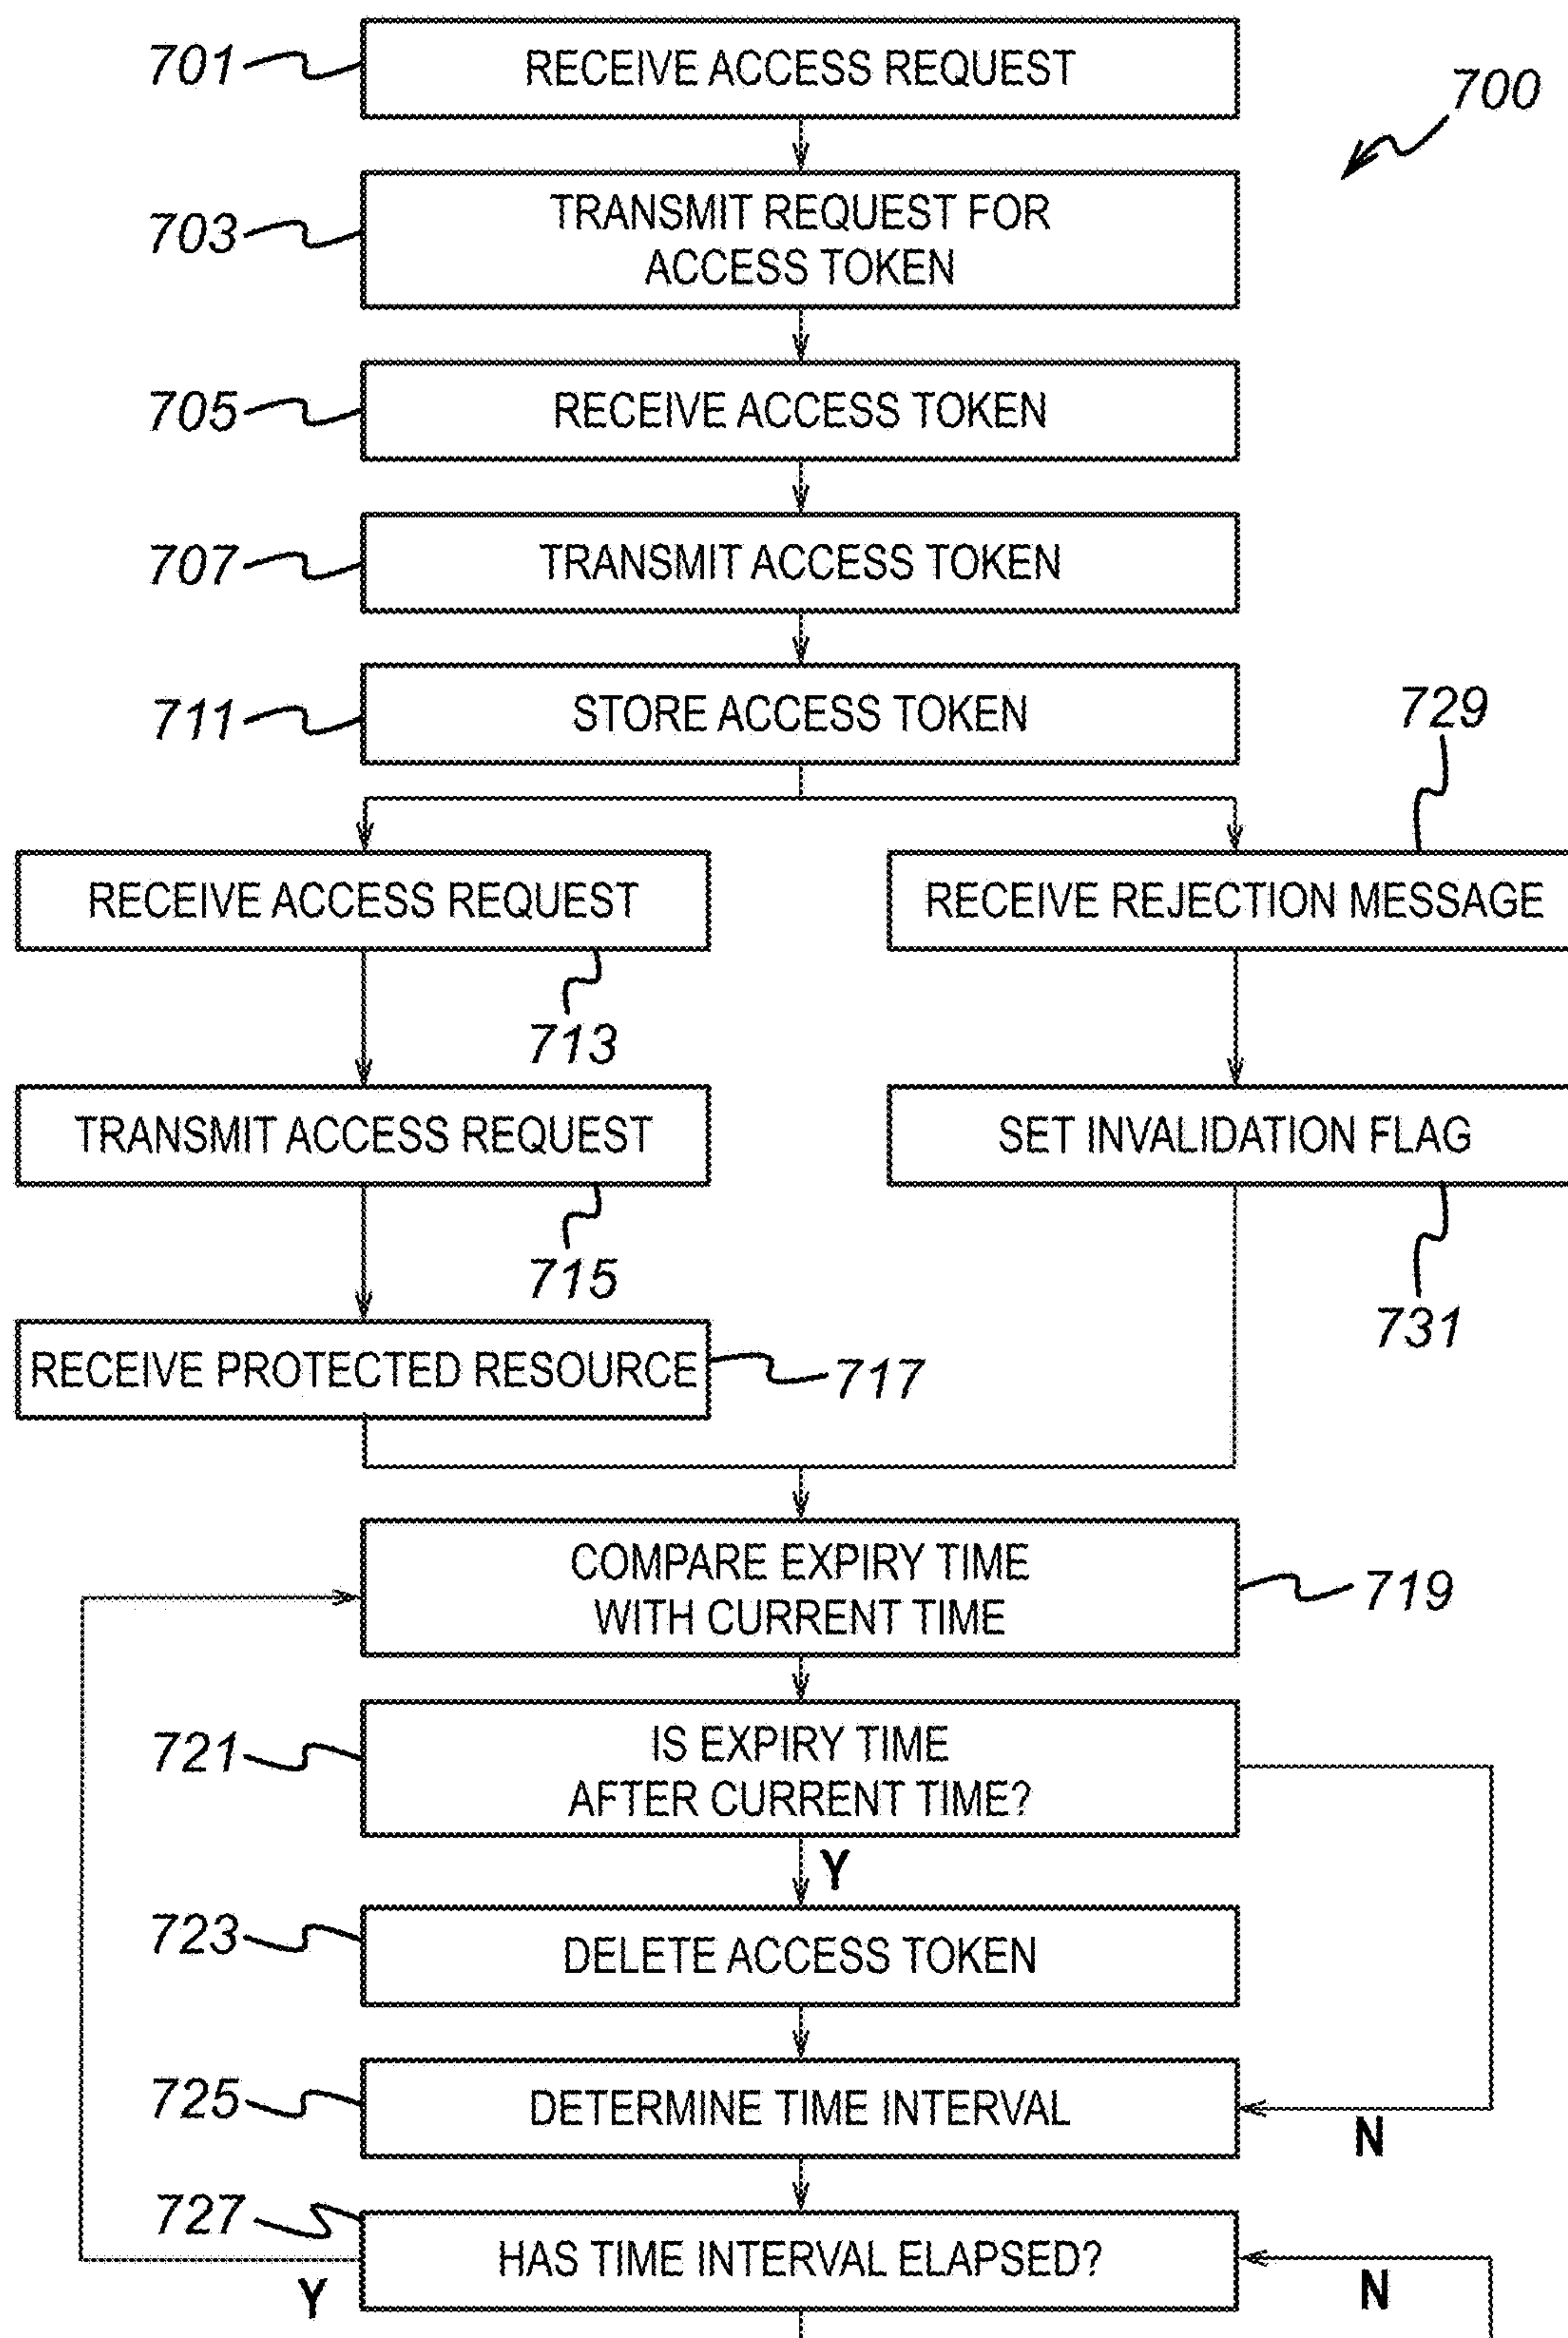
FIG. 6

FIG. 7

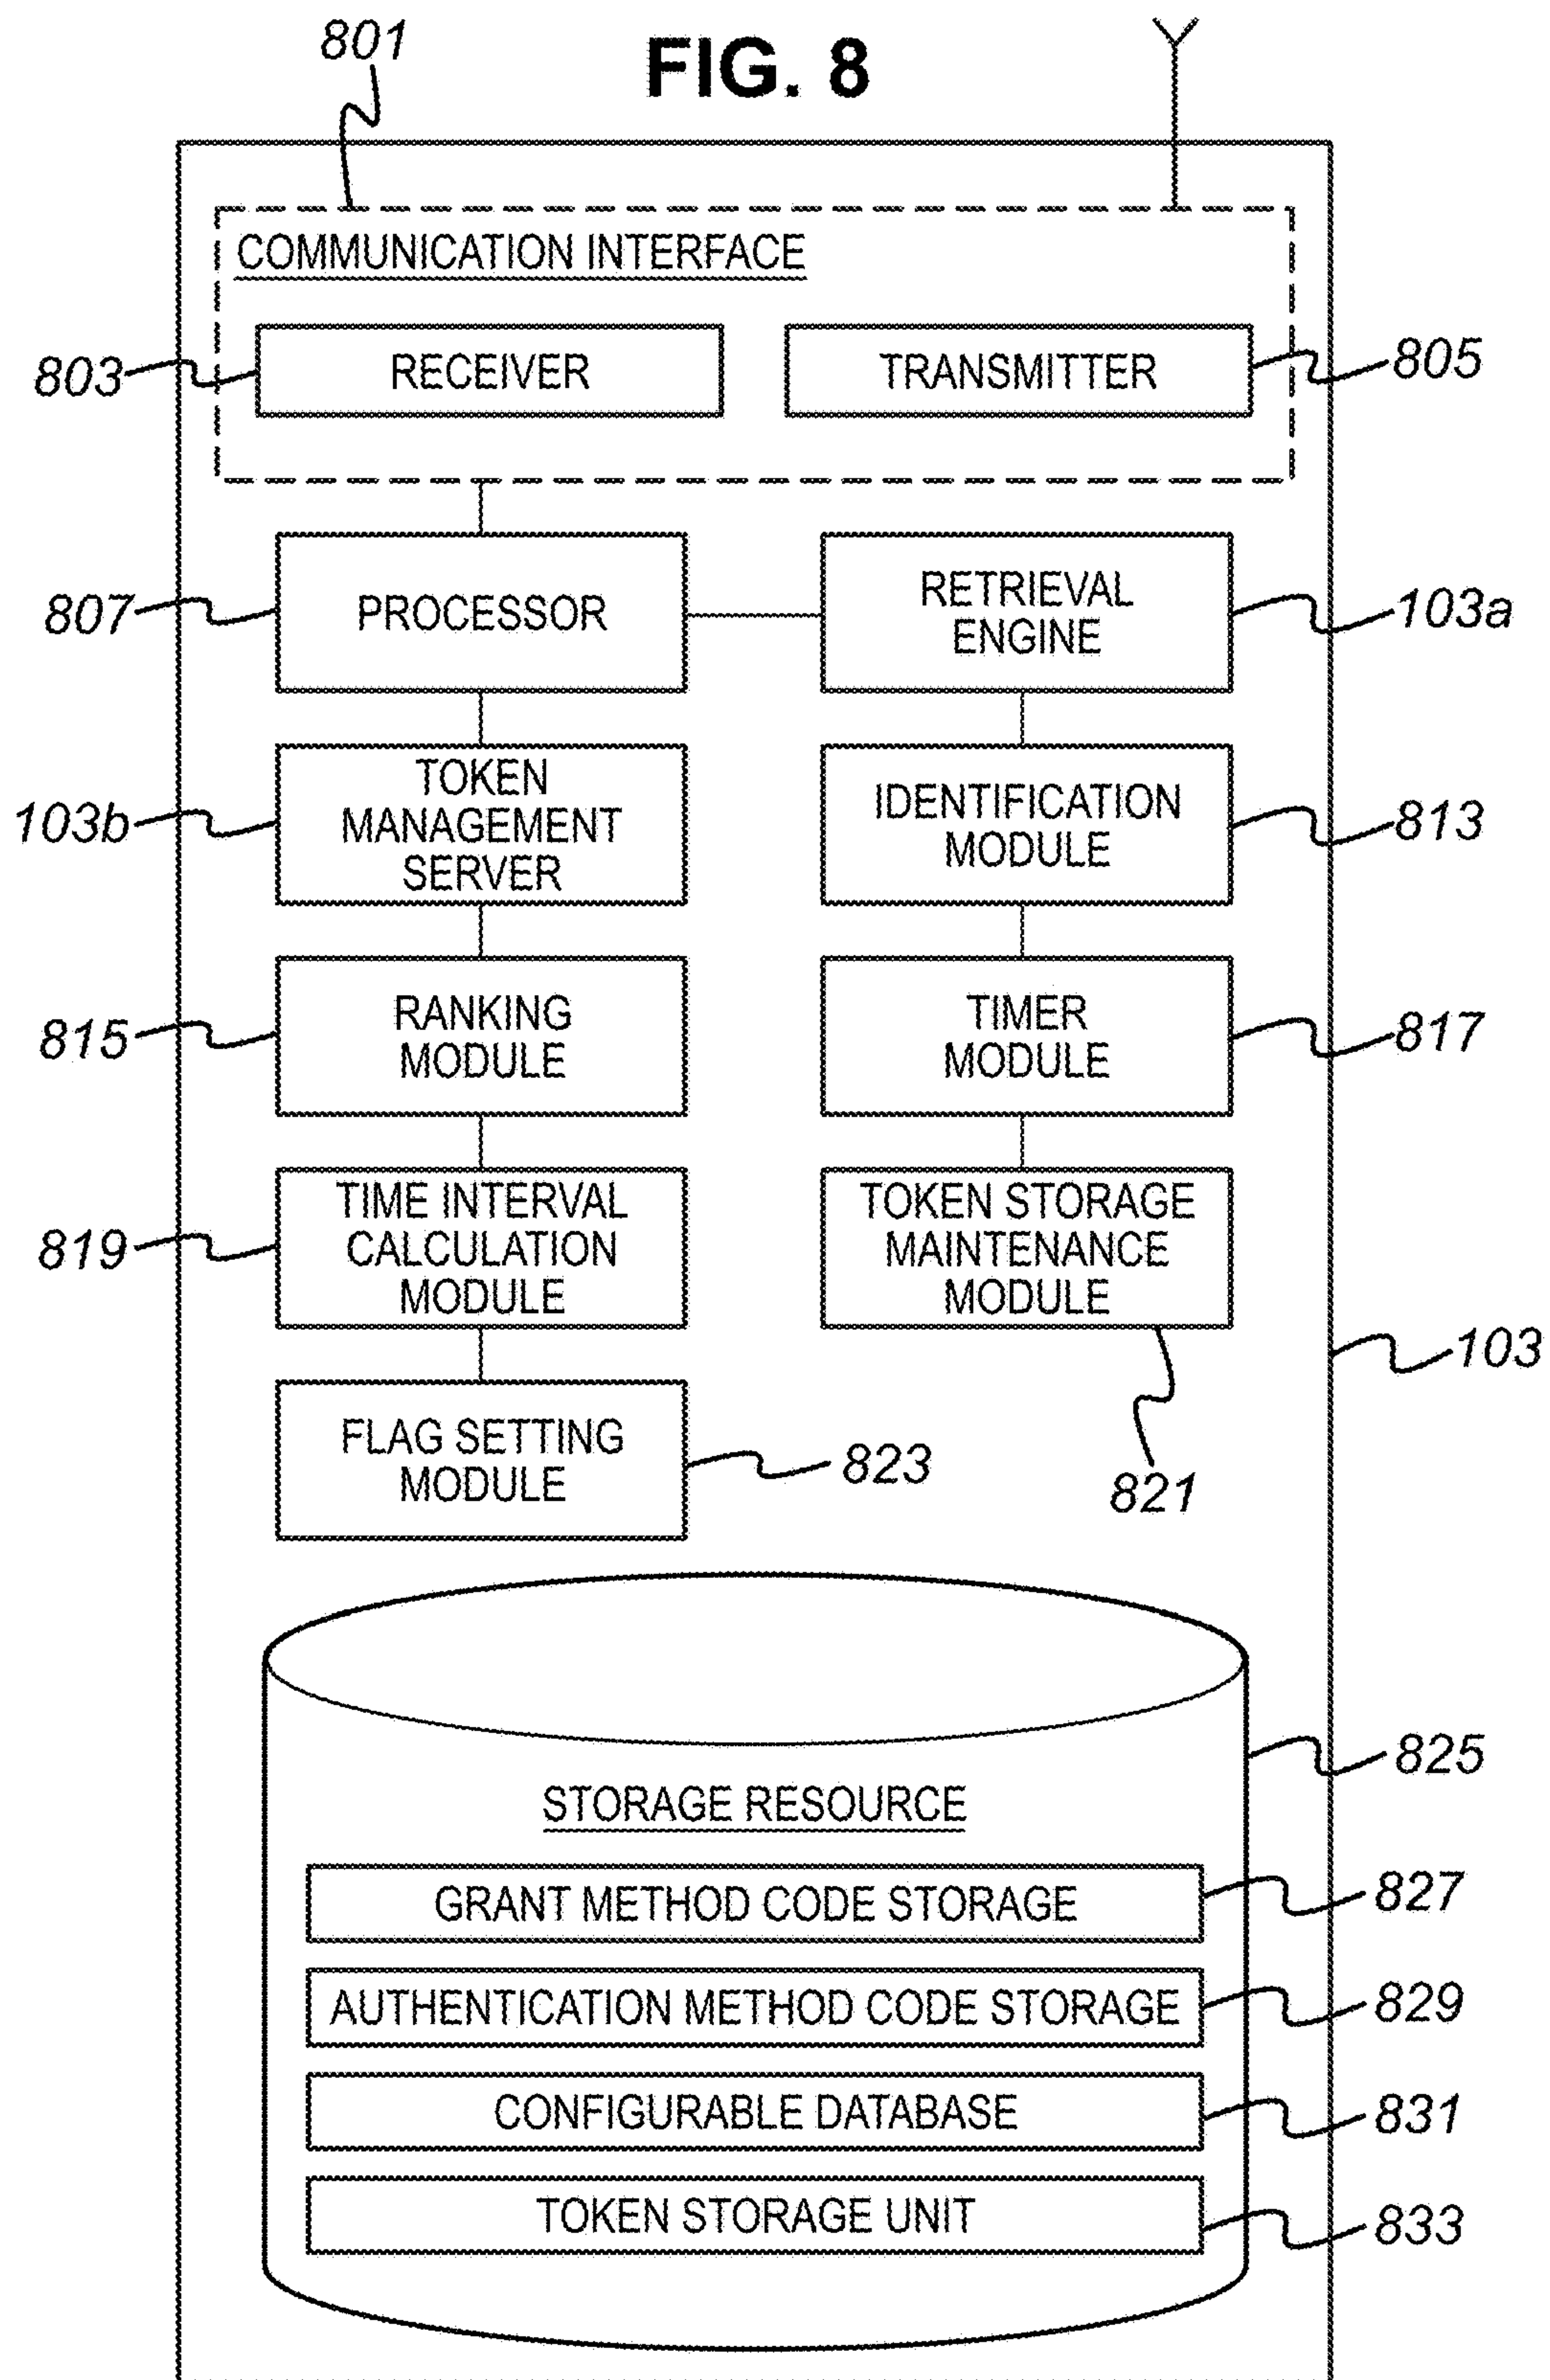
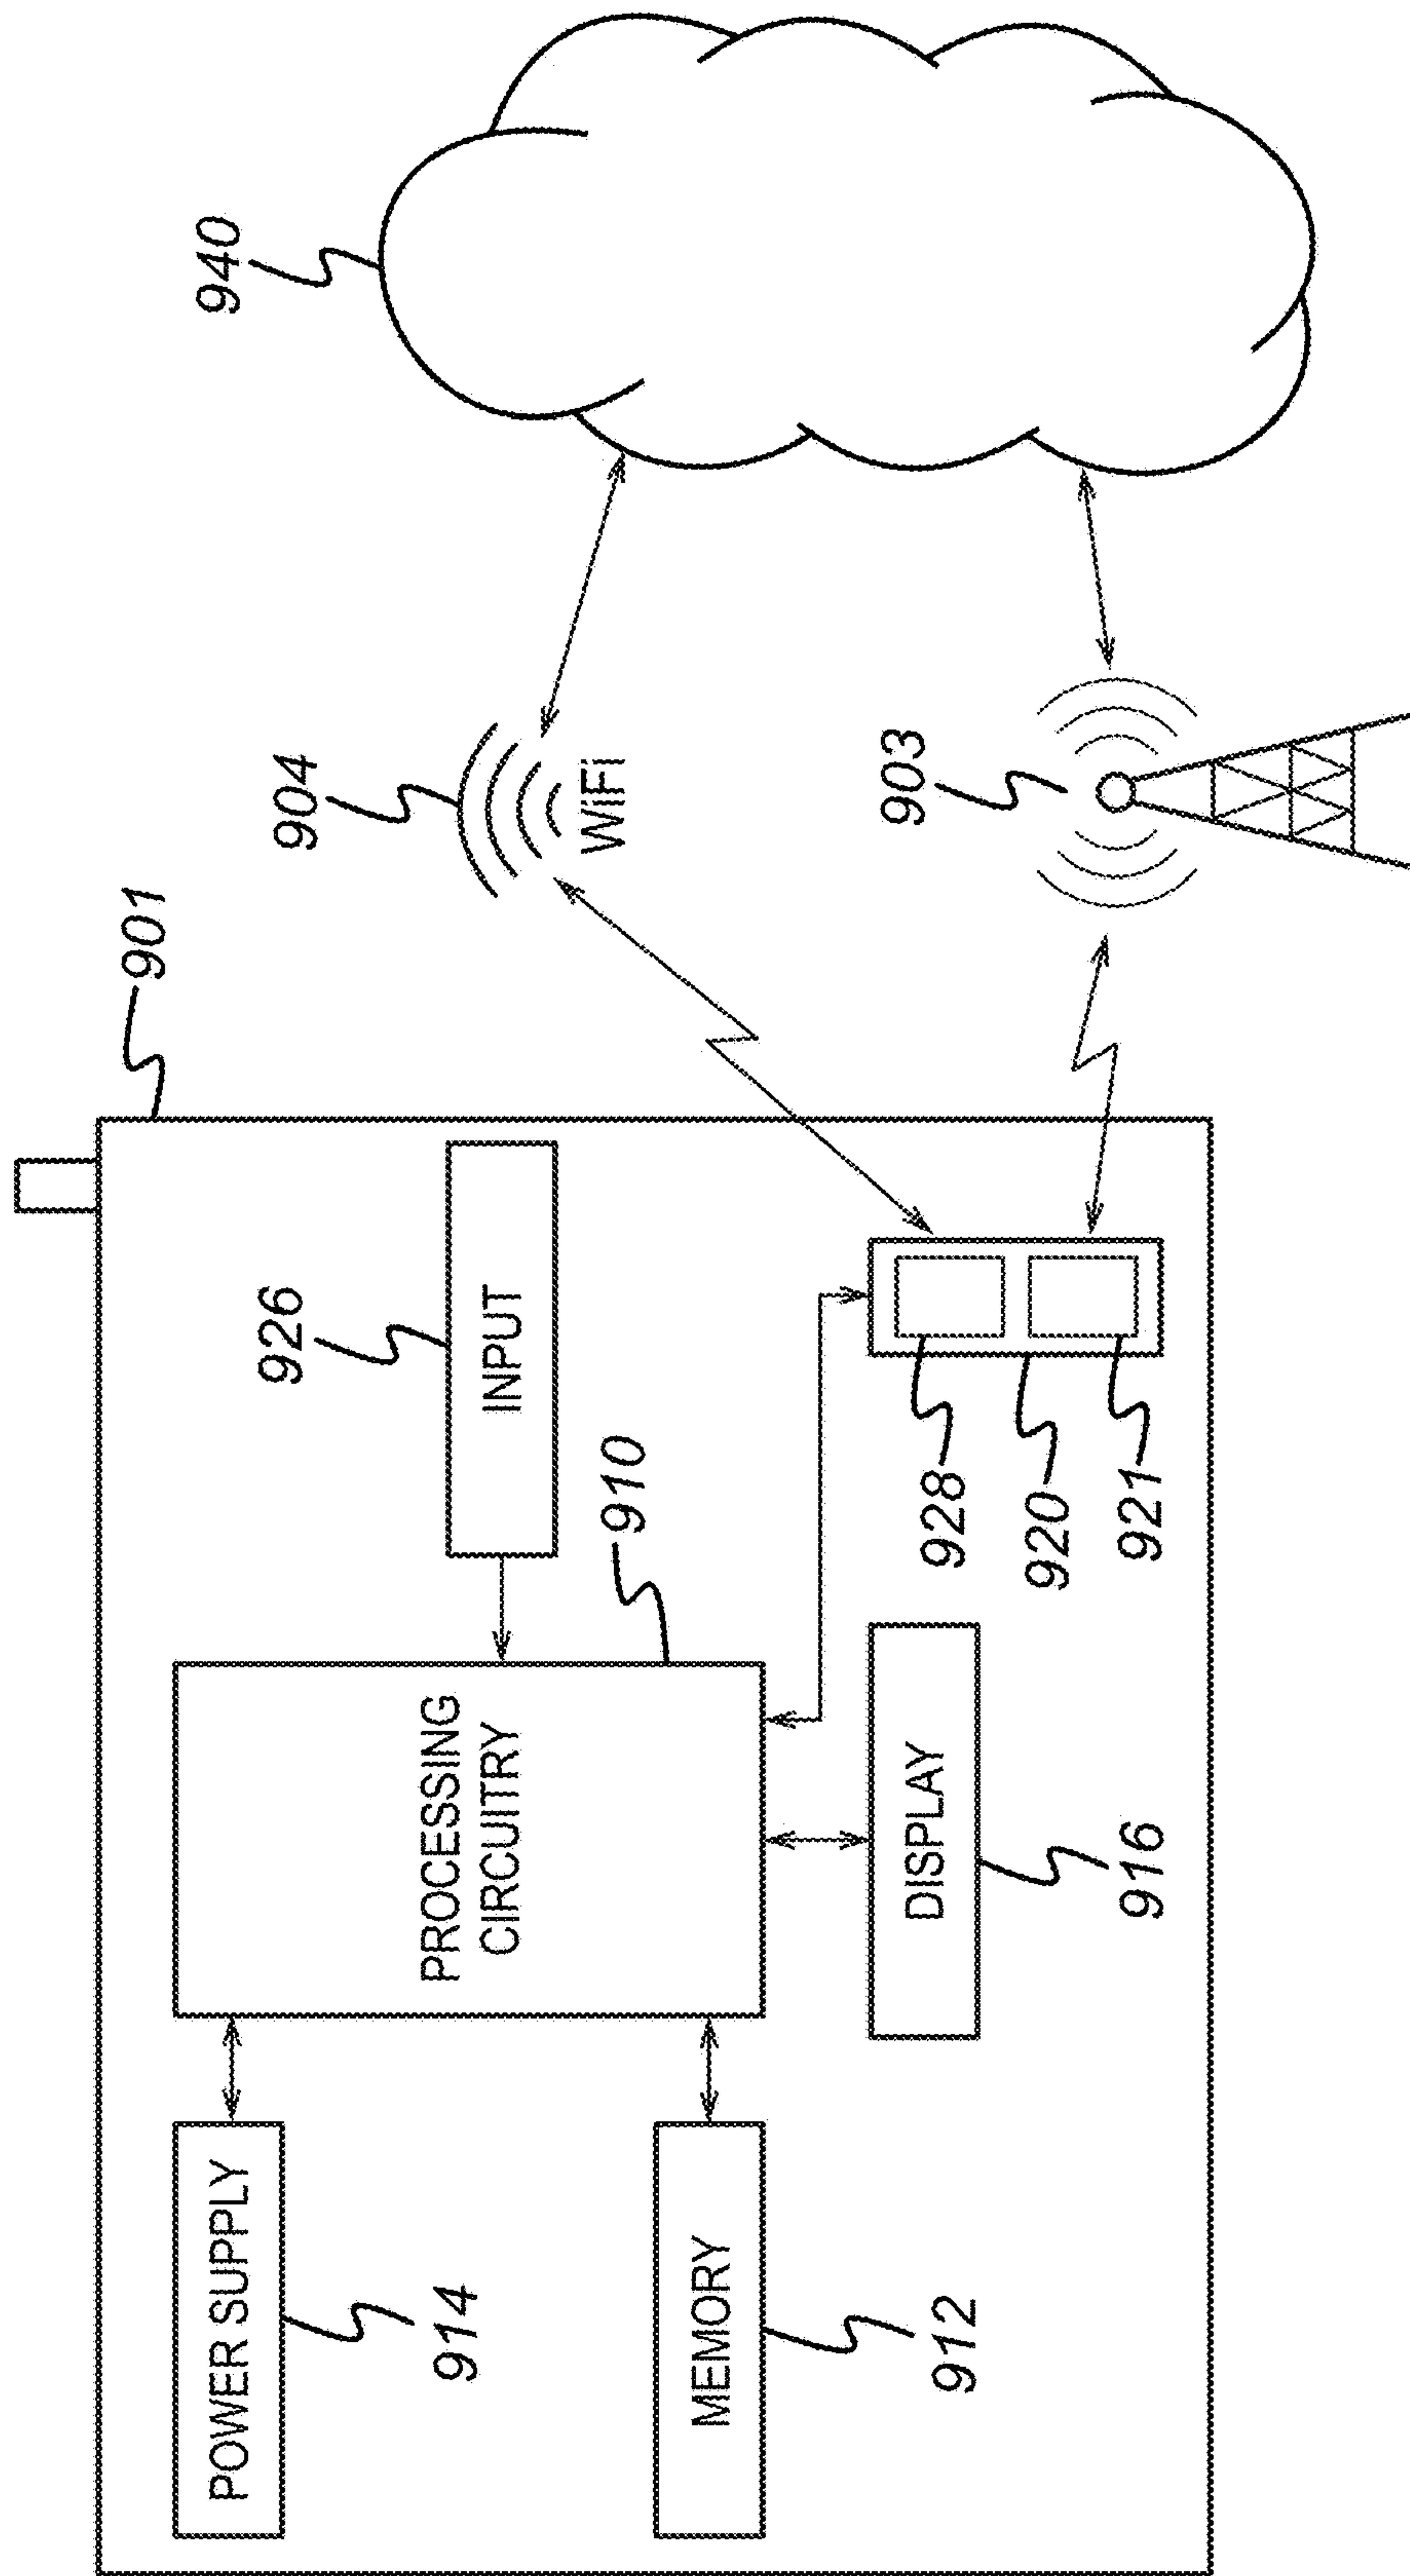


FIG. 9



SYSTEM FOR EFFICIENT MANAGEMENT AND STORAGE OF ACCESS TOKENS

RELATED APPLICATIONS

[0001] Foreign priority benefits are claimed under 35 U.S.C. § 119 to European application number 18166865.8, filed Apr. 11, 2018, the entire contents of which is hereby incorporated herein by reference in its entirety.

TECHNICAL FIELD

[0002] This disclosure relates to a system, a method and a computer program for efficiently managing and storing access tokens for accessing a protected resource at a client system on behalf of a user.

BACKGROUND

[0003] The OAuth (“Open Authorization”) 2.0 authorisation framework enables a third-party application to obtain limited access to an HTTP service, such as access to a protected resource. The third-party application may obtain access to the HTTP service on behalf of a resource owner by orchestrating an approval interaction between the resource owner and the HTTP service. Alternatively, the third-party application can obtain access to the HTTP service on its own behalf.

[0004] FIG. 1 illustrates an overview of the OAuth 2.0 protocol flow, involving the following four entities: a resource owner 1, a client 3, an authorisation server 5 and a resource server 7. The resource owner 1 is an entity that is capable of granting access to a protected resource. The resource owner 1 may be a person, referred to as a user operating a user device. The client 3 is an entity, or an application, that can make a request for access to the protected resource on behalf of the resource owner 1, when authorised by the resource owner 1. The authorisation server 5 is an entity that grants and issues access tokens to the client after successfully authenticating the resource owner and obtaining authorisation. The resource server 7 is an entity that hosts the protected resource and is capable of accepting and responding to requests for the protected resource using access tokens.

[0005] Access tokens, such as those issued by the authorisation server 5, are credentials used to access protected resources. An access token is a string representing an authorisation issued to the client. The string is usually opaque to the client. Tokens represent specific scopes and durations of access, granted by the resource owner, and enforced by the resource server and authorisation server.

[0006] An access token may denote an identifier used to retrieve the authorisation information or may self-contain the authorisation information in a verifiable manner (i.e., a token string consisting of some data and a signature). Access tokens can have different formats, structures, and methods of utilization (e.g., cryptographic properties) based on the resource server security requirements.

[0007] An access token can be used to identify a client. When an access token is used in this context it may be referred to herein as a grant token, rather than an access token. However, a grant token could equally be described as an access token.

[0008] The authorisation server 5 can issue refresh tokens, which are credentials used to obtain access tokens. Refresh tokens are issued to the client by the authorisation server and

are used to obtain a new access token when the current access token becomes invalid or expires, or to obtain additional access tokens with identical or narrower lifetime and fewer permissions than authorised by the resource owner). Issuing a refresh token is optional at the discretion of the authorisation server. If the authorisation server issues a refresh token, it is included when issuing an access token.

[0009] A refresh token is a string representing the authorisation granted to the client by the resource owner. The string is usually opaque to the client. The token denotes an identifier used to retrieve the authorisation information.

[0010] In step 11, client 3 requests authorisation from the resource owner 1 for access to the protected resource. The authorisation request can be made directly to the resource owner 1, or indirectly via the authorisation server 5 as an intermediary.

[0011] In step 12, the client 3 receives an authorisation grant, which is a credential representing the authorisation provided by the resource owner 1. This authorisation may be expressed using one of a plurality of “grant types” described in the Authorisation Framework. The authorisation grant type depends on the method(s) used by the client 3 to request authorisation and the type(s) supported by the authorisation server 5.

[0012] In step 13, the client 3 requests an access token by authenticating with the authorisation server 5 and presenting the authorisation grant. The authentication may be expressed using one of a plurality of “authentication types” for authenticating identifying client 3 at the authorisation server 5. The authentication type depends on the method(s) used by the client 3 to authenticate itself and the type(s) supported by the authorisation server 5.

[0013] In step 14, the authorisation server 5 authenticates the client 3 and validates the authorisation grant, and if valid, issues an access token.

[0014] In step 15, the client 3 requests the protected resource from the resource server 7 and authenticates itself at the resource server 7 by presenting the access token.

[0015] In step 16, the resource server 7 validates the access token. If the access token is valid, the resource server 7 serves the request by transmitting the protected resource to the client 3.

[0016] There may be an expiry time associated with each access token, where an access token will be deemed to be invalid by the resource server 7 after the expiry time. If the access token is valid, it can be used to access the protected resource. However, if the access token is invalid, it cannot be used to access the protected resource successfully.

[0017] In conventional systems, since access tokens expire, the client 3 may transmit a request for an access token in response to each request from the resource owner 1 for the client 3 to access a protected resource. However, this requires the use of additional bandwidth and processing resources. Therefore, there is a need for a system that makes more efficient use of bandwidth and processing resources.

SUMMARY

[0018] In one aspect of the invention there is provided a computer-implemented method comprising: receiving, at a client system from a user device, a first access request comprising a first instruction to access a protected resource stored at a resource system; transmitting, from the client system to an authorisation system, a token request for an access token to be used for accessing the protected resource,

in response to the first access request; receiving, at the client system, an access token in response to the token request, the access token having a corresponding time to expire indicative of a time at which the access token will not be valid for obtaining the protected resource from the resource system; transmitting the access token from the client system to the resource system and, in response, receiving the protected resource; storing the access token at a token storage unit of the client system; receiving, at the client system from the first user device, a second access request comprising a second instruction to access a protected resource stored at the resource system; and in response to the second access request, transmitting the stored access token to the resource system for receiving the protected resource.

[0019] In this way, since the access token is stored and used at least one more time in response to an access request from a user, the client system is able to reduce the amount of requests for access tokens made. This makes more efficient use of bandwidth and processing resources.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] Embodiments of the invention will be described, by way of example, with reference to the following drawings, in which:

[0021] FIG. 1 illustrates a protocol sequence diagram of the OAuth 2.0 framework;

[0022] FIG. 2 illustrates the general architecture of a system for accessing a protected resource;

[0023] FIGS. 3 and 4 illustrate a protocol sequence diagram of a computer-implemented invention for accessing a protected resource;

[0024] FIG. 5 illustrates a flow chart of a computer-implemented method in which a client system is configured to execute a selected authentication method and grant method for obtaining an access token used for accessing a protected resource;

[0025] FIG. 6 illustrates a flow chart of a computer-implemented method in which a client system refreshes a grant token based on its expiry time;

[0026] FIG. 7 illustrates a flow chart of computer-implemented method for managing and storing valid and invalid access tokens;

[0027] FIG. 8 illustrates a schematic diagram of a client system; and

[0028] FIG. 9 illustrates a schematic diagram of an example device in the system.

DETAILED DESCRIPTION

[0029] Referring to FIG. 2, there is a system 100 for managing access to protected resources. The system 100 comprises one or more user devices 101, a client system 103 and an external system 104, which comprises an authorisation system 105 and a resource system 107.

[0030] The user device 101 may be considered as the resource owner 1, the client system 103 may be considered as the client 3, the authorisation system 105 may be considered as the authorisation server 5, and the resource system 107 may be considered as the resource server 7 when considering the methods and systems described herein in terms of the OAuth 2.0 protocol.

[0031] The external system 104 is configured to store data associated with users of devices, such as the user device 101.

The data stored at the external system 104 may comprise protected resources, such as one or more secure data items.

[0032] Each one of the protected resources may be indicative of private information relating to a user of the external system 104. In one specific example, each data item stored at the external system 104 comprises financial data relating to the user, such as details that enable the user to make payments or the details of previous financial transactions made by the user.

[0033] The following systems and methods are described in the context of managing access to financial data in terms of the OAuth 2.0 protocol. However, these systems and methods could be used to manage access to any type of protected resource for which access by unauthorised third parties is to be restricted using another suitable protocol, such as SAML, OpenID and the like.

[0034] The client system 103 is configured to access protected resources stored at the external system 104, upon request from a user.

[0035] The external system 104 may comprise a plurality of sub-systems, for instance, a plurality of servers. In the example described herein, the external system 104 comprises the authorisation system 105 and the resource system 107. The authorisation server 105 is configured to authorise requests for access to protected resources that are stored at the resource system 107. Alternatively, the external system 104 may comprise a single server for performing these functions.

[0036] The client system 103 may comprise a plurality of sub-systems, for instance, a plurality of servers. In the example described herein, the client system 103 comprises a retrieval engine 103a and a token management server (TMS) 103b. The retrieval engine 103a is configured to interface with the user device 101, and the TMS 103b is configured to interface with the external system 104. Alternatively, the client system 103 comprises a single server for performing these functions.

[0037] In the following examples, secure data and protected resources, are referred to as being accessible by a user, and that the data is associated with the user. For example, the user may have access to an online user account, such as an online banking account, via an account interface provided by the external system 104. In this scenario, the user may be assigned a unique username and a shared secret (e.g. login information), such as a password, that can be used to access the user account via the account interface. Once the user has accessed the user account, that user is able to access the data via the user account. Therefore, the data is accessible by the user via login information that is unique to the user.

[0038] The secure data and protected resources that are accessible by the first user may be accessible by the external system 104 itself. The secure data and protected resources may be accessible by the user only, unless otherwise authorised by the user. In other words, the data/protected resource are prevented from being sent to a device or a system that is remote and distinct from the external system 104, such as the client system 103, without the corresponding user providing authorisation to the external system 104 for the data to be sent to a remote device or system.

[0039] Each one of the external system 104, the client system 103 and the user device 101 are arranged to communicate with one another via a communications network 110. The communications network 110, in this example, is

the Internet 110. However, it will be appreciated that any suitable form of communications network 110 could be used.

[0040] Each one of the external system 104, the client system 103, and the user device 101 are web-enabled and may comprise a display, a user interface, a processor and memory. The devices and systems 101, 103, 104 can be arranged to communicate data between one another via any suitable communications protocol or connection. For instance, the devices and systems 101, 103, 104 may communicate with one another via a wired and/or a wireless connection.

[0041] The user device 101 may be any suitable type of personal computing device, such as a laptop computer, a desktop computer, a web-enabled telephone, such as a smartphone, or a tablet device. The client system 103 and the external system 104 may be any suitable type of computing system or collection of computing systems, such as a server or a collection of servers.

[0042] Referring to FIGS. 3 and 4, there is a computer-implemented method in which the client system 103 obtains access to a protected resource associated with the user of the user device 101 from the external system 104.

[0043] In step 301, the user device 101 transmits an access request to the retrieval engine 103a at the client system 103. The access request comprises an instruction for the client system 104 to initiate accessing the protected resource stored at the external system 104, in other words the access request indicates the user's intent for the client system 103 to access the protected resource.

[0044] In response to the access request received from the user device 101, the client system 103 performs steps 303 to 307 for obtaining a grant token from the external system 104. The grant token can be used by the client system 103 to identify itself to the external system 104 in order to initiate requests for the protected resource.

[0045] In step 303, the retrieval engine 103a transmits a request for a grant token to the TMS 103b.

[0046] In step 305, the TMS 103b forwards the request for the grant token to the authorisation system 105 at the external system 104. The authorisation sever 105 validates the request for the grant token. If the request for the grant token is valid, the authorisation system 105 responds by transmitting the grant token to the TMS 103b.

[0047] In step 307, if the grant token has been received, the TMS 103b forwards the grant token to the retrieval engine 103a.

[0048] Once the client system 103 has received the grant token, steps 309 and 311 can be performed in order to receive an intent identifier from the external system 104. The intent identifier is an identification label that can be used to identify a specific request from the user for the client system 103 to access a protected resource associated with the user from the external system 104 (i.e. the user's intent for the client system 103 to access the protected resource). The client system 103 is able to receive the intent identifier, if it has been issued with a valid grant token previously.

[0049] In step 309, the retrieval engine 103a transmits an intent request to the resource system 107 at the external system 104. The intent request is sent with the grant token received previously, which identifies the client system 103 to the external system 104.

[0050] In step 311, in response to the intent request, the external system 104 validates the grant token. If the intent

request is valid, the resource system 107 responds by transmitting an intent identifier to the TMS 103a.

[0051] In step 313, the retrieval engine 103a forwards the intent identifier to the user device 101. In addition, an instruction to redirect to the authorisation system 105 is sent to the user device 101 along with the intent identifier. The redirect instruction may comprise a URL corresponding with the authorisation system 105.

[0052] In step 315, the user device 101 is redirected to the authorisation system 105 using the URL, and the user device 101 transmits the intent identifier to the authorisation system 105. The intent identifier is used by the authorisation system 105 to identify the user's intention for the client system 103 to access the protected resource.

[0053] In step 317, the user device 101 and the authorisation system 105 communicate with one another in order for the user to provide authorisation for the client system 103 to access the protected resource. This step may involve the user selecting the protected resource, or a portion of the protected resource such as a specific sub-set of data relating a specific user account (or group of accounts) stored at the resource system 107 that are accessible by the user.

[0054] In step 319, once the user has provided authorisation for the client system 104 to access the protected resource, the authorisation system 105 transmits an authorisation code to the user device 101. The authorisation code comprises an indicator that the user has provided authorisation. In this step, the authorisation system 105 transmits an instruction to the user device 101 to redirect to the client system 103. This redirect instruction may comprise a URL corresponding with the retrieval engine at the client system 103.

[0055] In step 321, the user device 101 is redirected to the retrieval engine 103a, and the user device 101 transmits the authorisation code to the retrieval engine 103a.

[0056] In step 323, the retrieval engine 103a transmits a request for an access token by transmitting the authorisation code to the authorisation system 105. The authorisation code may be transmitted directly from the retrieval engine 103a to the authorisation system 105, or indirectly via the TMS 103b.

[0057] In step 325, the authorisation system 105 validates the authorisation code. If the authorisation code is valid, the authorisation system 105 responds by transmitting an access token to the retrieval engine 103a. The access token can be used by the retrieval engine 103a to access the protected resource stored at the resource system.

[0058] In step 327, if the access token has been received successfully, the retrieval engine 103a transmits a success message indicating that the access token has been received. Alternatively, if the access token has not been received, the retrieval engine 103a transmits a failure message indicating that the access token has not been received.

[0059] Referring to FIG. 4, in step 329 the user device 101 transmits a request to the retrieval engine 103a for the client system 103 to retrieve the protected resource. If the access token has been received previously by the client system 103, the method proceeds to step 337. Alternatively, if the access token has not been received previously, steps 331 to 335 are performed in order for the client system 103 to obtain the access token in a similar manner to that described above.

[0060] In step 331, the retrieval engine 103a forwards a request for an access token to the TMS 103b. Then, in step 333, the TMS 103b forwards the access token request to the

authorisation system **105** and, in response, receives the access token. In step **335**, the access token is transmitted from the TMS **103b** to the retrieval engine **103a**.

[0061] In step **337**, the retrieval engine **103a** transmits a request for the protected resource by transmitting the access token to the resource system **107**. In step **339**, the resource system **104** the access token. If the access token is valid, the resource system responds by transmitting the protected resource to the retrieval engine **103a**.

[0062] In step **341**, once the protected resource has been received by the client system **103**, the retrieval engine transmits the protected resource to the user device **101**. Once received, the protected resource can be displayed at the user device **101** via an app or a browser at the user device **101**.

[0063] Referring to FIG. **5**, there is a computer-implemented method **500** performed by the client system **103** that enables the client system **103** to flexibly adapt to the grant and authentication methods supported by the external system **104**. This method can be used in conjunction with the method described above with reference to FIGS. **3** and **4**.

[0064] In step **501**, the client system **103** stores a plurality of grant method code portions. Each of the grant method code portions are stored in memory at the client system **103** and are executable by a processor of the client system **103**. Execution of any one of the grant method code portions causes the client system **103** to obtain an access token from the authorisation system **105** of the external system **104** using a specific type of grant method. The grant method corresponding to a particular grant method code portion is different to the grant methods corresponding to other the grant method code portions.

[0065] As discussed previously, the authorisation system **105** may support one or more of a plurality of different types of grant method, each of which enable the client system **103** to access a protected resource. In the OAuth 2.0 framework there are a number of available types of grant method, (or “grant types” as referred to in the OAuth 2.0 framework). The grant methods may comprise types such as the “authorisation code”, “implicit”, “resource owner password credentials” and “client credentials” grant types.

[0066] An overview of the “authorisation code” grant type is described above with reference to steps **313** to **327** of FIG. **3**. In these steps the authorisation system **105** provides an authorisation code in response to the user authorising access the protected resource. The client system **103** is then able to exchange the authorisation token for the access token at the authorisation system **105**.

[0067] The “client credentials” grant type involves fewer steps than the “authorisation code” grant type. Instead of redirecting the user device **101** to the authorisation system **105** in order to obtain the authorisation code, the client system **101** transmits a request for the access token comprises a client identifier and a client secret, which is some form of shared secret such as a password. The request for the access token comprises an indication of the grant type being used which, in this case is the “client credentials” grant type. The authorisation system **105** validates the client identifier and the client secret. If the client identifier and the client secret are valid, the authorisation system **105** responds with the access token.

[0068] Although only certain grant types have been described herein, any other suitable grant type(s) could be used in addition to or instead of the grant type(s) described.

[0069] In step **503**, the client system **103** stores a plurality of authentication method code portions. Each of the authentication method code portions are stored in memory at the client system **103** and are executable by a processor of the client system **103**. Execution of any one of the authentication method code portions causes the client system **103** to authenticate itself to the external system **104** during the process of obtaining the access token. The authentication method corresponding to a particular authentication method code portion is different to the authentication methods corresponding to the other authentication method code portions.

[0070] As discussed previously, an authorisation system may support one or more of a plurality of different types of authentication method, each of which enabling the client system **103** to authenticate itself to the external system **104**. In the OAuth 2.0 framework there are a number of available types of authentication methods, (or “authentication types”). The authentication methods may comprise types such as the “client secret” authentication method and the “client assertion” authentication method.

[0071] The “client secret” authentication method defines the way in which the client system **103** authenticates itself at the authorisation system **105** when obtaining the access token. If the “client secret” authentication method is used, the client system **103** will transmit a client secret to the authorisation system **103** when requesting the access token. The client secret is a shared secret, such as a password, assigned to the client system **103**. The authorisation system **105** uses the client secret to authenticate the client system **103** in validating requests for the access token.

[0072] The “client assertion” authentication method is similar to the “client secret” method. However, in the “client assertion” authentication method the client system **103** transmits an integrity protected version of the client secret. For instance, the client secret may be integrity protected using a digital signature or Message Authentication Code (MAC). In this way, the client secret can be protected from eavesdropping and tampering.

[0073] Although only certain authentication types have been described herein, any other suitable authentication type(s) could be used in addition to or instead of the authentication type(s) described.

[0074] In step **505**, the client system **103** stores a configurable database that identifies which authorisation systems support which grant and authentication method types. The configurable database comprises a plurality of authorisation system identifiers that are each indicative of a particular authorisation system. Thus, each authorisation system identifier enables a specific authorisation system to be identified, such as the authorisation system **105** described with reference to FIG. **2**.

[0075] Each one of the authorisation system identifiers is stored in association with one or more of the plurality of types of grant method. In other words, each authorisation system identifier is logically linked with one or more of the types of grant method in the database. This enables the database to indicate which grant method type(s) are supported by a specific authorisation system.

[0076] In addition, each one of the authorisation system identifiers is stored in association with one or more of the plurality of types of authentication method. In other words, each authorisation system identifier is logically linked with one or more of the types of authentication method in the

database. This enables the database to indicate which authentication method type(s) are supported by a specific authorisation system.

[0077] The database is configurable such that it can be updated in order to modify the types of grant method and the authentication method that are associated with each authorisation system identifier. Since the client system 103 is capable of executing a plurality of different types of the authentication and grant methods by virtue of the grant and authentication method code portions, the client system 103 can adapt to changes in the types of method supported by an external system in a quick and simple fashion.

[0078] In addition, the configurable database can be updated in order to store additional authorisation system identifiers and associated grant and authentication method types. This allows the client system 103 to be configured for communicating with an external system with which the client system 103 has not previously communicated.

[0079] In steps 505A-B, the grant and authentication methods associated with one or more of the authorisation systems can be modified. This may occur in response the client system 103 receiving a message that the grant and/or authentication methods supported by an authorisation system have changed. For instance, the client system 103 may receive a message from the external system 104, or any other system, that a particular grant or authentication method is no longer supported by its authorisation system 105. The client system 103 may disassociate the unsupported grant or authentication method with the authorisation system identifier in response to this message.

[0080] In another example, the client system may receive a message from the external system 104, or any other system, that a particular grant or authentication method that was not previously supported by the authorisation system 105 is now supported. The client system 103 may associate the newly supported grant or authentication method with the authorisation system identifier in response to this message.

[0081] Instead of receiving a message from another system regarding the supported grant/authentication types and responding accordingly, the database may be configured by an administrator at the client system 103 to reflect which methods are supported by a selected authorisation system.

[0082] The client system 103 may perform steps 505C-D in order to maintain an accurate copy of the configurable database. This allows the client system 103 to ensure that the correct grant/authentication methods are used when communicating with the external system 104.

[0083] In step 505C, the client system 103 transmits a database request to a system that hosts the database which indicates the grant/authentication systems that are supported by a selection of authorisation systems. The database hosting system responds to the database request by transmitting at least a portion of the database to the client system 103. Then, in step 505D, the client system 103 updates the configurable database using the database received from the database hosting system. In this step, the configurable database is configured to store the one or more of the grant/authentication methods in the received database in association with the corresponding authorisation system identifier.

[0084] The client system 103 may execute step 505C in an intermittent fashion. This allows the client system 103 to conserve bandwidth usage by reducing amount of communication between the client system 103 and the system that stores the database. The client system 103 may transmit the

database request in accordance with a predetermined schedule. For instance, the predetermined schedule may define a time interval between consecutive database requests. In one example, the time interval between each data request is defined as 24 hours in the predetermined schedule, such that a single database request is sent once a day. This achieves a balance between the aim of maintaining an accurate version of the database at the client system 103 and bandwidth usage. The predetermined schedule may be configurable. For instance, the time interval between adjacent database requests may be configurable by an administrator of the client system 103. This allows the client system 103 to be tuned in order for the optimum number of database requests to be sent within a given time period.

[0085] As discussed above, the authorisation system 105 may support different types of grant method and authentication method. However, one of these different types of method may be more or less secure than the other types. For example, the “authorisation code” grant method may be more secure than the “client credentials” grant method. In another example, the “client secret” authentication method may be less secure than the “client assertion” authentication method.

[0086] In step 507, the client system 103 ranks the types of grant method and authentication method in the configurable database based on the security strength of each grant type. This step may involve associating a score with each one of the grant method types and each one of the authentication method types, which indicates the security strength of each method. For example, the “authorisation code” grant method may be associated with a score of “10” and the “client credentials” grant method may be associated with a score of “5”. In this way, the respective scores indicate that the “authorisation code” grant method is more secure than the “client credentials” grant method.

[0087] In another example, the “client secret” authentication method may be associated with a score of “4” and the “client assertion” authentication method may be associated with a score of “9”. In this way, the respective scores indicate that the “client assertion” authentication method is more secure than the “client secret” authentication method.

[0088] Ranking the grant and authentication methods based on their security strength allows the client system 103 to choose the more secure method, where a choice is available. This enhances the security of the system as a whole, since using less secure methods may expose the system to issues such as eavesdropping.

[0089] In step 509, the client system 103 receives an access request from the user device 101 in a similar manner to that described in step 301 of FIG. 3. The access request comprises an instruction for the client system 103 to access a protected resource. The access request also comprises a request identifier which indicates the authorisation system with which the client system 103 must communicate with before accessing the protected resource at the corresponding resource system. In this example, the authorisation system is the authorisation system 105 of the external system 104.

[0090] In step 511, the client system 103 compares the authorisation system 105 indicated by the access request against the authorisation systems indicated by the authorisation system identifiers in the configurable database. If a match is found, the grant and authentication method types associated with the matching authorisation system are identified.

[0091] The identified authorisation and grant method types are the methods that are supported by the authorisation system 105 with which the client system 103 must communicate in order to service the access request. There may be a plurality of grant method types supported by the authorisation system 105, or there may be only a single grant method type supported by the authorisation system 105. There may be a plurality of authentication method types supported by the authorisation system 105, or there may be only a single authentication method type supported by the authorisation system 105.

[0092] Steps 511A-B refer to the scenario in which a plurality of grant and/or authentication types are supported by the authorisation system 104. In step 511A, the client system 103 selects one of the grant method types and one of the authentication method types identified in step 511. The method type selected may be based on a variety of different criteria, for instance by choosing the fastest or the most efficient method of each type available.

[0093] Step 511B refers to a specific example in which the most secure grant and/authentication method is chosen. In this step, the client system 103 selects the most secure method based on the ranking process performed in step 507. For instance, the grant method or the authentication method with the highest score is selected for execution.

[0094] In step 513, the grant method code portion corresponding with the selected grant method is executed, and the authentication method code portion corresponding with the selected authentication method is executed.

[0095] In step 514, if the client system 103 is successfully validated at the authorisation system 104 using the selected grant and authentication methods, the client system 103 receives the access token. Subsequently, the client system 103 can transmit the access token to the resource system 107 in order to obtain the protected resource corresponding with the access request in step 501.

[0096] Referring to FIG. 6, there is a computer-implemented method 600 that enables the client system 103 to ensure that it has access to a valid grant token, which is required in order to successfully respond to requests from the user to access a protected resource. This method can be used in conjunction with the method described above with reference to FIGS. 3 and 4.

[0097] Each of the grant tokens provided to the client system 103 by the authorisation system 105 may have a corresponding time to expire which is the time at which the identification will not be considered valid at the authorisation system. Thus, an expired grant token cannot be used in the process of obtaining a protected resource from the resource system 107. If the grant token stored at the client system 103 is not valid (i.e. the token has expired) at a time when an access request (such as the request described with reference to step 301 in FIG. 3) is received, the client system 103 may not be able to service the access request successfully, or at the least there will be a delay in servicing the request. Thus, the method described with reference to FIG. 6 allows user access requests to be serviced reliably and quickly, by ensuring that a valid grant token is held by the client system 103.

[0098] In step 601, the client system 103 receives an intent message from the user device 101. The intent message may be the access message described previously. Thus, step 601 may be performed in a similar manner to step 301 described with reference to FIG. 3. The access message may be

referred to as an intent message because this message indicates the user's intent for the client system 103 to access the protected resource.

[0099] In step 603, the client system 103 transmits a request for a grant token to the authorisation system 105. As explained above with reference to FIG. 3, the grant token is a token that enables the external system 104 to authenticate the identity of the client system 103.

[0100] In step 605, the authorisation system 105 validates the client system's 103 request for an identification code. If the request is valid, the authorisation system 105 responds by transmitting the grant token to the client system 103. The authorisation system 105 may transmit an expiry time indicator to the client system 103 that indicates the expiry time of the grant token. In addition, the authorisation system 105 may transmit a refresh token to the client system 103 that can be used to obtain a new (i.e. an unexpired) grant token.

[0101] Steps 603-605 may be performed in a similar manner to that described with reference to steps 303-307 in FIGS. 3 and 4. Also, steps 603-605 may be performed before, after or at the same time as receiving the intent message of step 601.

[0102] In step 607, a time interval is determined which defines the time between receiving the grant token and requesting a new one, and a timer is set using the time interval. The time interval may be a predetermined time interval, such as a 1 minute, 5 minutes or 10 minutes etc. This time interval may be based on the expiry time of the received grant token. For instance, the time interval may be equal to the expiry time of the grant token, such that the client system 103 can initiate the process of requesting a new grant token at the moment the grant token expires. In a specific example, the expiry time of the grant token may be 5 minutes and the time interval may, therefore, be set to 5 minutes. Thus, the client system 103 will request a new grant token 5 minutes after the previous grant token has been received, which is the moment at which the previous token expires.

[0103] In another example, the time interval may be set as an amount of time less than the expiry time of the grant token. In this way, the client system 103 can initiate the process of requesting a new grant token before the grant token expires. In a specific example, the expiry time of the grant token may be 5 minutes and the time interval may be set to 4 minutes. Thus, the client system 103 can ensure that there is only a small window of time between requesting a new grant token and expiry of the previous token. The time interval may be set such that this window is not greater than a predefined window length (e.g. 1 minute). This window can be configured by configuring the time interval to optimise the number of refresh requests sent, while ensuring that a valid grant token is held at the client system 103.

[0104] In another example, the time interval may be set as an amount of time greater than the expiry time of the grant token. In this way, the client system 103 can ensure that the process of requesting a new grant token will occur at a precise moment after the previous grant token has expired. In a specific example, the expiry time of the grant token may be 5 minutes and the time interval may, therefore, be set to 6 minutes. Thus, the client system 103 can ensure that there is only a small window of time between expiry of the grant token and the request for a new one. The time interval may be set such that this window is not greater than a predefined window length (e.g. 1 minute). Again, this window can be

configured by configuring the time interval to optimise the number of refresh requests sent, while ensuring that a valid grant token is held at the client system 103.

[0105] In step 609, the client system 103 starts a timer in response to receiving the grant token and using the time interval set in step 607. The timer is used to determine the moment at which the time interval has elapsed.

[0106] In step 611, once the time interval has elapsed the client system 103 transmits a refresh request to the authorisation system 105. The refresh request may comprise an instruction for a new (i.e. unexpired) grant token to be provided to the client system 103 from the authorisation system 105.

[0107] In step 613, the client system 103 receives the new grant token in response to the refresh request.

[0108] In step 615, the client 103 may transmit the most recently received grant token to the resource system 107 in order to initiate a request for the protected resource. Step 615 may be performed in a similar manner that described with reference to step 309 of FIG. 3.

[0109] Steps 601 to 615 may be repeated such that many different grant tokens are received from the authorisation system 107. In addition, steps 601 to 615 may be repeated for other authorisation system such that many different grant tokens are received from a variety of authorisation systems. The grant tokens received from a particular authorisation system may have similar expiry times.

[0110] In step 617, the time interval discussed in connection with steps 607 and 609 is calculated based on the expiry times of the grant tokens received and the specific authorisation systems from which the grant tokens are received.

[0111] In one example, the client system 103 may calculate a predicted expiry time of the grant tokens received from a particular authorisation system. The predicted expiry time may be calculated by calculating an average of the expiry times of the grant token received from a specific authorisation system, or a plurality of different authorisation systems.

[0112] The predicted expiry time calculated in step 617 may be used to determine the time interval in step 607. For instance, the predicted expiry time may equal to, a time less than or a pre-set time greater than the time interval in step 607.

[0113] Referring to FIG. 7, there is a computer-implemented method 700 that enables the client system 103 reduce the bandwidth and processing resources used by minimising the number of requests for access tokens sent, such as the requests for access tokens described with reference to steps 323 and 325 in FIG. 3. Thus, the method described with reference to FIG. 7 can be used in conjunction with the method described above with reference to FIGS. 3 and 4.

[0114] In step 701, the client system 103 receives an access request from the user device 101. This step may be performed in a similar manner to that described with reference to step 301 in FIG. 3.

[0115] In step 703, the client system 103 transmits a request for an access token to the authorisation server 105 in response to the access request. This step may be performed in a similar manner to step 323, for instance after steps 303 to 321 have been performed, as described with reference to FIG. 3.

[0116] In step 705, the client system 103 receives the access token from the authorisation system 105. This step may be performed in a similar manner to step 325, as

described with reference to FIG. 3. In this step, the authorisation system 105 may transmit an expiry time indicator to the client system 103. The expiry time indicator corresponds to the access token and indicates the expiry time of the access token. After the expiry time indicated by the expiry time the access token cannot be used to access the protected resource. In other words, the expiry time indicator is indicative of the time at which the corresponding access token will not be valid for obtaining the protected resource from the resource system 107. In step 505, the authorisation system 105 can transmit a refresh token corresponding with the access token. The refresh token can be used by the client system 3 obtain an unexpired access token from the authorisation system 105.

[0117] The expiry time indicator may be indicative of a length of time during which the access token will be valid. For instance, this length of time may be expressed as a number of second or minutes. Alternatively, the expiry time may indicate a point in time at which the access token will no longer be valid. For instance, the expiry time may indicate a specific time during the day. The client system 103 may determine the expiry time of the access token based on the length of time or the point in time indicators by the expiry time indicator.

[0118] In step 707, the client system 103 transmits the access token to the resource system 107 in a request to receive the protected resource. This step may be performed in a similar manner to step 337 as described with reference to FIG. 4.

[0119] In step 711, the access token received from the authorisation system 107 is stored at a token storage unit at the client system 103. The expiry time indicated by the expiry time indicator may be stored at the token storage unit also. The access token and/or the expiry time indicated by the expiry time indicator may be encrypted and stored at the token storage unit to enhance security. In this step, the access token may be stored independently of the expiry time of the corresponding access token. In other words, the client system 103 stores the access token irrespective of whether the expiry time is short or long. This reduces the processing effort required in analysing the expiry time of each access token. Alternatively, the client system 103 may not store, or may delete, the access token if its expiry time is less than a predetermined threshold. If the access token is stored, the access token may be stored after its expiry time.

[0120] In step 713, the client system 103 receives another access request similar to the access request received in step 701.

[0121] In step 715, the client system 103 transmits the stored access token to the authorisation system 107 in another request to receive the protected resource similar to the request sent in step 707. Thus, the client system 103 uses the stored access token, rather than requesting a new access token to service the user's access request. This reduces the bandwidth and processing resources used by the client system 103.

[0122] In steps 719 to 727, token storage maintenance requests performed at the client system 103. These steps assist in conserving storage resources, while making more efficient use of bandwidth and processing resource by minimising the number of requests for new access tokens.

[0123] In step 719, the client system 103 compares the time indicated by the expiry time indicator corresponding with one or more of the access tokens stored.

[0124] In step 721, if the time indicated by an expiry time indicator is after the current time, the method proceeds to step 723. Alternatively, if the time indicated by the expiry time indicator is not after the current, the method proceed to step 725.

[0125] In step 723, the access token corresponding with the expiry time that is after the current time is deleted.

[0126] Steps 719 to 723 may be performed intermittently. For instance, steps 719 may be executed in accordance with a predetermined schedule. In one example, the predetermined schedule defines a time interval between adjacent executions of steps 719 to 723. This time interval may be configurable at the client system 103 based on a user input received from an administrator. The time interval may be configurable at the client system 103 based on monitored performance of the client system 103. Configuration of the time interval may occur automatically.

[0127] Steps 725 and 727 may be executed In order to perform steps 719 to 723 in accordance with the predetermined schedule. In step 725, the time interval of the predetermined schedule is determined. Then, in step 727 the method proceed to repeat steps 719 to 723 once the time interval has elapsed.

[0128] Steps 729 and 731 may be performed in the method 700 in order to manage a situation in which an access token has become invalid and, therefore, will not be usable for obtaining the protected resource. When an access token has become invalid it should not be used again, as this will involve unnecessary communications being transmitted in between the client system 103 and the external system 104. However, deleting access tokens each time they are deemed to be invalid involves a processing burden that it would preferable to avoid. This is particularly relevant when the client system 103 handles a large number of access token for a large number of users with protected resources stored at a variety of external systems. Steps 729 and 731 allow the transmission of invalid access to be prevented, while avoiding the processing burden managing the storing of access tokens an on individual basis.

[0129] In step 729, the client system 103 receives a rejection message indicating that the access token is not valid. This message may be received from the external system 104, for instance via the authorisation system 105 or the resource system 107. Alternatively, the rejection message may be received from any other system, or may be received via an input at the client system 103.

[0130] The rejection message may be received in response to the external system 104 determining that the access token has expired. For example, the client system 103 may transmit the stored access token to the authorisation system 105 in an attempt to access the protected resource. However, the authorisation system 107 may determine that the access token has expired. In response to determining the access token has expired, the authorisation system 107 transmits a rejection message to the client system 103.

[0131] In another example, the rejection message may be received in response to the user of the user device 101 revoking their authorisation for the client system 103 to access the protected resource. The user may inform the authorisation system 107 that their authorisation has been revoked. This will invalidate the corresponding access token, and its corresponding refresh token. However, in this scenario the client system 103 will be unaware that access token and the refresh token have been invalidated. There-

fore, the client system 103 may continue to transmit requests for access to the protected resource using the invalid access token. The client system 103 will receive a rejection message from the external system 104 in response to each one of these requests because the access token is invalid. However, the client system 103 might still transmit refresh requests using the invalid refresh token, under the assumption that the access token has expired rather than the user's authorisation having been revoked. All of these processes represent an unnecessary load on the processing resources and bandwidth of the client system 103. If the access token has a corresponding refresh token, the rejection message comprises an indication that the refresh token is invalid.

[0132] In step 731, rather than deleting the access token and the refresh token directly in response to the rejection message, the client system 103 sets an invalidation flag in association with the corresponding access token and refresh token. Setting the invalidation flag involves a lower processing overhead than deleting the tokens. Thus, when the client system 103 is managing many tokens, this will equate to a significant enhancement in efficiency.

[0133] The invalidation flag may be an addition bit field stored in association with the access token and the refresh token. For instance, if the bit field is set to "1", this may indicate that the corresponding token is invalid, and, if the bit field is set to "0", this may indicate that the corresponding token is invalid (or vice versa).

[0134] In another example, the invalidation flag may be set by modifying the expiry time indicator associated with the access token and/or the refresh token. In this example, the expiry time indicator is set in the past in order to indicate that the token(s) are no longer valid. Thus, when steps 719 to 727 are performed subsequently, the tokens will be deleted as part of the token storage maintenance steps

[0135] Referring to FIG. 8, the client system 103 comprises a communication interface 801 comprising a receiver 803 and a transmitter 805. The client system 103 comprises a processor 807, an identification module 813, a ranking module 815, a timer module 817, a time interval calculation module 819, a token storage maintenance module and a flag setting module 823. The client system 103 also comprises the retrieval engine 103a and the token management server 103b described above.

[0136] There is a storage resource 825 at the client system 103 which comprises a grant method code portion storage resource 827, an authentication method code storage resource 829, a configurable database storage resource 831 and a token storage unit 833.

[0137] The receiver 803 and the transmitter 805 are configured to receive and transmit message, instructions and tokens to and from the client system 103 as explained above.

[0138] The storage resource 825 is configured to store grant method code portions, as described above with reference to step 501, at the grant method code portions storage resource 827. The storage resource 825 is configured to store authentication method code portions, as described above with reference to step 503, at the authentication method code portion storage resource 829. The storage resource 825 is configured to store the configurable database, as described above with reference to step 505, at the configurable database storage resource 831. The token storage unit 833 is arranged for storing tokens, such as the access tokens, grant tokens and refresh tokens described above.

[0139] The ranking module **817** is configured to perform the ranking processes described above, such as those described above with reference to step **507**. The identification module **813** is configured to identify the authentication method and grant methods supported by an authorisation server, as described above with reference to step **511**. The processor **807** is configured to execute instructions, such as the instruction of the selected grant and authentication method code portions, as described with reference to step **513**. The timer module **817** is configured to activate a timer for a time period, for instance as described above with reference to step **607**, **609**. In addition, the timer module **817** can be used to monitor a current time for comparison with the expiry time of a token, as described above with reference to step **719**. The time interval calculation module **819** is configured to calculate a time interval for setting the timer, as described above with reference to step **725**. The token storage management module **821** is configured to compare the current time indicated by the timer with the expiry time of a token, and to delete a token in response, as described above with reference to steps **721** and **723** above. The flag setting module **823** is configured to set an invalidation flag in association with a token, as described above with reference to step **731**.

[0140] FIG. **9** shows an exemplary electronic device **901** according to any of the electronic devices or systems of this disclosure (such as the user device **101**, client system **103**, the external system **104**, the authorisation system **105**, the resource system **107**, the retrieval engine **103a** or the TMS **103b**). The electronic device **901** comprises processing circuitry **910** (such as a microprocessor) and a memory **912**. Electronic device **901** may also comprise one or more of the following subsystems: a power supply **914**, a display **916**, a transceiver **920**, and an input **926**.

[0141] Processing circuitry **910** may control the operation of the electronic device **901** and the connected subsystems to which the processing circuitry is communicatively coupled. Memory **912** may comprise one or more of random access memory (RAM), read only memory (ROM), non-volatile random access memory (NVRAM), flash memory, other volatile memory, and other non-volatile memory.

[0142] Display **916** may be communicatively coupled with the processing circuitry **910**, which may be configured to cause the display **916** to output images representative of the secure data, or protected resources, shared between the entities in the system **100**.

[0143] The display **916** may comprise a touch sensitive interface, such as a touch screen display. The display **916** may be used to interact with software that runs on the processor **910** of the electronic device **901**. The touch sensitive interface permits a user to provide input to the processing circuitry **910** via a discreet touch, touches, or one or more gestures for controlling the operation of the processing circuitry and the functions described herein. It will be appreciated that other forms of input interface may additionally or alternatively be employed for the same purpose, such as the input **926** which may comprise a keyboard or a mouse at the input device.

[0144] The transceiver **920** may be one or more long-range RF transceivers that are configured to operate according to communication standard such as LTE, UMTS, 3G, EDGE, GPRS, GSM, and Wi-Fi. For example, electronic device **901** may comprise a first wireless transceiver **921**, such as a cellular transceiver, that is configured to commu-

nicate with a cell tower **903** via to a cellular data protocol such as LTE, UMTS, 3G, EDGE, GPRS, or GSM, and a second transceiver **928**, such as a Wi-Fi transceiver, that is configured to communicate with a wireless access point **904** via to a Wi-Fi standard such as 802.11 ac/n/g/b/a. In this regard and for the purposes of all embodiments herein concerning a long-range wireless protocol, a long-range wireless protocol may be a protocol which is capable and designed for communication over 5, 10, 20, 30, 40, 50, or 100 m. This is in contrast to short-range wireless protocol mentioned above. The long-range wireless protocol may communicate utilizing higher power than the short-range wireless protocol. The range (e.g. line of sight distance) between the long-range end nodes (electronic device and router or base station) for the long-range wireless protocol may be greater than the range (e.g. line of sight distance) between the short-range end nodes (e.g. electronic device and wireless beacon).

[0145] Electronic device **901** may be configured to communicate via the transceiver **920** with a network **940**. Network **940** may be a wide area network, such as the Internet, or a local area network. Electronic device **901** may be further configured to communicate via the transceiver **920** and network **940** with one or more systems or user devices. These servers or user devices may be any one of those described herein.

[0146] The term “comprising” encompasses “including” as well as “consisting” e.g. a composition “comprising” X may consist exclusively of X or may include something additional e.g. X+Y.

[0147] The word “substantially” does not exclude “completely” e.g. a composition which is “substantially free” from Y may be completely free from Y. Where necessary, the word “substantially” may be omitted from the definition of the invention.

[0148] Unless otherwise indicated each embodiment as described herein may be combined with another embodiment as described herein.

[0149] The methods described herein may be performed by software in machine readable form on a tangible storage medium e.g. in the form of a computer program comprising computer program code means adapted to perform all the steps of any of the methods described herein when the program is run on a computer and where the computer program may be embodied on a computer readable medium. Examples of tangible (or non-transitory) storage media include disks, thumb drives, memory cards etc. and do not include propagated signals. The software can be suitable for execution on a parallel processor or a serial processor such that the method steps may be carried out in any suitable order, or simultaneously. This acknowledges that firmware and software can be valuable, separately tradable commodities. It is intended to encompass software, which runs on or controls “dumb” or standard hardware, to carry out the desired functions. It is also intended to encompass software which “describes” or defines the configuration of hardware, such as HDL (hardware description language) software, as is used for designing silicon chips, or for configuring universal programmable chips, to carry out desired functions.

[0150] It will be appreciated that the modules described herein may be implemented in hardware or in software. Furthermore, the modules may be implemented at various locations throughout the system.

[0151] Those skilled in the art will realise that storage devices utilised to store program instructions can be distributed across a network. For example, a remote computer may store an example of the process described as software. A local or terminal computer may access the remote computer and download a part or all of the software to run the program. Alternatively, the local computer may download pieces of the software as needed, or execute some software instructions at the local terminal and some at the remote computer (or computer network). Those skilled in the art will also realise that by utilizing conventional techniques known to those skilled in the art that all, or a portion of the software instructions may be carried out by a dedicated circuit, such as a DSP, programmable logic array, or the like.

[0152] Any range or device value given herein may be extended or altered without losing the effect sought, as will be apparent to the skilled person.

[0153] It will be understood that the benefits and advantages described above may relate to one embodiment or may relate to several embodiments. The embodiments are not limited to those that solve any or all of the stated problems or those that have any or all of the stated benefits and advantages.

[0154] Any reference to ‘an’ item refers to one or more of those items. The term ‘comprising’ is used herein to mean including the method blocks or elements identified, but that such blocks or elements do not comprise an exclusive list and a method or apparatus may contain additional blocks or elements.

[0155] The steps of the methods described herein may be carried out in any suitable order, or simultaneously where appropriate. Additionally, individual blocks may be deleted from any of the methods without departing from the spirit and scope of the subject matter described herein. Aspects of any of the examples described above may be combined with aspects of any of the other examples described to form further examples without losing the effect sought. Any of the module described above may be implemented in hardware or software

[0156] It will be understood that the above description of a preferred embodiment is given by way of example only and that various modifications may be made by those skilled in the art. Although various embodiments have been described above with a certain degree of particularity, or with reference to one or more individual embodiments, those skilled in the art could make numerous alterations to the disclosed embodiments without departing from the scope of this invention.

List of Numbered Embodiments

[0157] 1. A computer-implemented method comprising:

[0158] receiving, at a client system from a user device, a first access request comprising a first instruction to access a protected resource stored at a resource system;

[0159] transmitting, from the client system to an authorisation system, a token request for an access token to be used for accessing the protected resource, in response to the first access request;

[0160] receiving, at the client system, an access token in response to the token request, the access token having a corresponding time to expire indicative of a time at which the access token will not be valid for obtaining the protected resource from the resource system;

[0161] transmitting the access token from the client system to the resource system and, in response, receiving the protected resource;

[0162] storing the access token at a token storage unit of the client system;

[0163] receiving, at the client system from the first user device, a second access request comprising a second instruction to access a protected resource stored at the resource system; and

[0164] in response to the second access request, transmitting the stored access token to the resource system for receiving the protected resource.

2. The computer-implemented method of embodiment 1 further comprising:

[0165] in response to the token request, receiving and storing at the token storage unit an expiry time indicator corresponding to the received access token’s time to expire;

[0166] wherein the expiry time indicator is indicative of the time at which the corresponding access token will not be valid for obtaining the protected resource from the resource system.

3. The computer-implemented method of embodiment 2 wherein the expiry time indicator is indicative of a length of time that the access token will be valid for; and the method further comprises:

[0167] determining the received access token’s time to expire based on the length of time that the access token will be valid for and a time at which the access token is received at the client system.

4. The computer-implemented method of embodiment 2 or embodiment 3 wherein the access token is stored at the token storage unit independently of the time indicated by the expiry time indicator.

5. The computer-implemented method of any one of embodiments 2 to 4 further comprising wherein storage of the access token at the token storage unit is maintained until a time after the time indicated by the expiry time indicator.

6. The computer-implemented method according to any one of embodiments 2 to 5 further comprising token storage maintenance steps performed by the client system comprising:

[0168] comparing the time indicated by the stored expiry time indicator with a current time;

[0169] deleting the access token corresponding with the stored expiry time if the time indicated by the stored expiry time indicator is after the current time.

7. The computer-implemented method of embodiment 6 wherein the token storage maintenance steps are executed intermittently.

8. The computer-implemented method of embodiment 6 or 7 wherein the token storage maintenance steps are executed according to a predetermined schedule.

9. The computer-implemented method of embodiment 8 wherein the predetermined scheduled defines a time interval between adjacent executions of the token storage maintenance steps.

10. The computer-implemented method of embodiment 9 wherein the time interval is configurable at the client system based on a user input received at the client system.

11. The computer-implemented method of embodiment 9 or embodiment 10 wherein the time interval is configurable at the client system based on monitored performance of the client system.

12. The computer-implemented method of any one of embodiments 2 to 11 further comprising:

[0170] the client system comparing the time indicated by the stored expiry time indicator with a current time, and refreshing the access token if the time indicated by the stored expiry time indicator is after the current time.

13. The computer-implemented method of any one of embodiments 2 to 12 further comprising:

[0171] transmitting the access token from the client system to the resource system and, in response, receiving a message indicating that the access token is not valid; and

[0172] refreshing the access token, at the token storage unit, in response to receiving the message indicating that the access token is not valid.

14. The computer-implemented method according to any one of the preceding embodiments comprising:

[0173] receiving, at the client system, a plurality of access tokens from the authorisation system, wherein each of the plurality of access tokens is received in response to a token request transmitted to the authorisation system;

[0174] receiving and storing an expiry time indicator in association with each access token received at the client system, wherein each of the stored expiry time indicators is indicative of a time at which the corresponding access token will not be valid.

15. The computer-implemented method of embodiment 14 further comprising:

[0175] the client system determining a predicted expiry time for the access tokens received from the authorisation system based on the stored expiry time indicators.

16. The computer-implemented method of embodiment 15 further comprising:

[0176] the client system determining that the predicted expiry time does not meet a predetermined threshold and, in response, not storing further access tokens received from the authorisation system.

17. The computer-implemented method of embodiment 15 or embodiment 16 wherein the predicted expiry time is determined by calculating an average of the times indicated by the stored expiry time indicators.

18. The computer-implemented method according to any one of the preceding embodiments comprising:

[0177] receiving, at the client system from the first user device, a plurality of access requests subsequent to the first access request, each of the plurality of subsequent access requests comprising a subsequent instruction to access a protected resource stored at the resource system; and

[0178] in response to the subsequent access request, transmitting the stored access token to the resource system from the client system for receiving the protected resource.

19. The computer-implemented method according to any one of the preceding embodiments comprising:

[0179] receiving, at the client system from each of a plurality of user devices, an access request comprising an instruction to access a protected resource stored at a resource system;

[0180] transmitting, from the client system to an authorisation system, a token request for an access token to be used for accessing the protected resource, in response to each of the access requests received from the plurality of user devices;

[0181] receiving, at the client system, an access token in response to each of the token requests, each access token having a corresponding time to expire indicative of a time at

which the access token will not be valid for obtaining the protected resource from the resource system;

[0182] transmitting each of the access tokens from the client system to the resource system and, in response, receiving the protected resource;

[0183] storing each of access tokens at the token storage unit.

20. The computer-implemented method according to embodiment 19 further comprising:

[0184] in response to each token request, receiving and storing an expiry time indicator in association with each access token received at the token storage unit;

[0185] wherein each expiry time indicator is indicative of a time at which the corresponding access token will not be valid for obtaining the protected resource from the resource system.

21. The computer-implemented method according to embodiment 20 further comprising token storage maintenance steps performed by the client system comprising:

[0186] comparing the time indicated by each stored expiry time indicator with a current time;

[0187] deleting the access token corresponding with the stored expiry time if the time indicated by the stored expiry time indicator is after the current time.

22. The computer-implemented method of embodiment 21 wherein the token storage maintenance steps are executed intermittently or according to a predetermined schedule.

23. The computer-implemented method of embodiment 22 wherein the predetermined schedule defines a time interval between adjacent executions of the token storage maintenance steps.

24. The computer-implemented method of embodiment 23 wherein the time interval is configurable at the client system based on the times indicated by the expiry time indicators.

21. The computer-implemented method according to any one of the preceding embodiments wherein the client system comprises one or more client servers in communication with one another.

22. The computer-implemented method according to any one of the preceding embodiments wherein the authorisation system comprises one or more authorisation servers in communication with one another.

23. The computer-implemented method according to any one of the preceding embodiments wherein the resource system comprises one or more resources servers in communication with one another.

24. A computer program comprising instructions which, when the program is executed by a computer, cause the computer to carry out the method of any one of the preceding embodiments.

25. A data carrier signal carrying the computer program of embodiment 24.

26. A computer readable medium comprising instructions which, when executed by a computer, cause the computer to carry out the method of any one of embodiments 1 to 23.

27. A client system comprising processing circuitry configured to:

[0188] receive, from a user device, a first access request comprising a first instruction to access a protected resource stored at a resource system;

[0189] transmit to an authorisation system, a token request for an access token to be used for accessing the protected resource, in response to the first access request;

[0190] receive an access token in response to the token request, the access token having a corresponding time to expire indicative of a time at which the access token will not be valid for obtaining the protected resource from the resource system; and

[0191] transmit the access token from the client system to the resource system and, in response, receiving the protected resource;

[0192] wherein the client system further comprises a token storage unit configured to store the access token; and

[0193] the processing further configured to:

[0194] receive, from the first user device, a second access request comprising a second instruction to access a protected resource stored at the resource system; and

[0195] transmit the stored access token to the resource system for receiving the protected resource, in response to the second access request.

28. A client system comprising:

[0196] a receiver configured to receive, from a user device, a first access request comprising a first instruction to access a protected resource stored at a resource system;

[0197] a transmitter configured to transmit to an authorisation system, a token request for an access token to be used for accessing the protected resource, in response to the first access request;

[0198] wherein the receiver is configured to receive an access token in response to the token request, the access token having a corresponding time to expire indicative of a time at which the access token will not be valid for obtaining the protected resource from the resource system; and

[0199] wherein the transmitter is configured to transmit the access token from the client system to the resource system and, in response, receiving the protected resource;

[0200] wherein the client system further comprises a token storage unit configured to store the access token;

[0201] the receiver configured to receive, from the first user device, a second access request comprising a second instruction to access a protected resource stored at the resource system; and

[0202] the transmitter configured to transmit the stored access token to the resource system for receiving the protected resource, in response to the second access request.

What is claimed:

1. A computer-implemented method comprising:

receiving, at a client system from a user device, a first access request comprising a first instruction to access a protected resource stored at a resource system;

transmitting, from the client system to an authorisation system, a token request for an access token to be used for accessing the protected resource, in response to the first access request;

receiving, at the client system, an access token in response to the token request, the access token having a corresponding time to expire indicative of a time at which the access token will not be valid for obtaining the protected resource from the resource system;

transmitting the access token from the client system to the resource system and, in response, receiving the protected resource;

storing the access token at a token storage unit of the client system;

receiving, at the client system from the first user device, a second access request comprising a second instruction to access a protected resource stored at the resource system; and

in response to the second access request, transmitting the stored access token to the resource system for receiving the protected resource.

2. The computer-implemented method of claim 1 further comprising:

in response to the token request, receiving and storing at the token storage unit an expiry time indicator corresponding to the received access token's time to expire;

wherein the expiry time indicator is indicative of the time at which the corresponding access token will not be valid for obtaining the protected resource from the resource system.

3. The computer-implemented method of claim 2 wherein the expiry time indicator is indicative of a length of time that the access token will be valid for; and the method further comprises:

determining the received access token's time to expire based on the length of time that the access token will be valid for and a time at which the access token is received at the client system.

4. The computer-implemented method of claim 2 wherein the access token is stored at the token storage unit independently of the time indicated by the expiry time indicator; and/or wherein storage of the access token at the token storage unit is maintained until a time after the time indicated by the expiry time indicator.

5. The computer-implemented method of claim 2 further comprising token storage maintenance steps performed by the client system comprising:

comparing the time indicated by the stored expiry time indicator with a current time;

deleting the access token corresponding with the stored expiry time if the time indicated by the stored expiry time indicator is after the current time and, optionally, wherein the token storage maintenance steps are executed intermittently.

6. The computer-implemented method of claim 5 wherein the token storage maintenance steps are executed according to a predetermined schedule.

7. The computer-implemented method of claim 8 wherein the predetermined schedule defines a time interval between adjacent executions of the token storage maintenance steps wherein, optionally, the time interval is configurable at the client system based on a user input received at the client system and/or wherein the time interval is configurable at the client system based on monitored performance of the client system.

8. The computer-implemented method of claim 2 further comprising:

the client system comparing the time indicated by the stored expiry time indicator with a current time, and refreshing the access token if the time indicated by the stored expiry time indicator is after the current time.

9. The computer-implemented method of claim 2 further comprising:

transmitting the access token from the client system to the resource system and, in response, receiving a message indicating that the access token is not valid; and

refreshing the access token, at the token storage unit, in response to receiving the message indicating that the access token is not valid.

10. The computer-implemented method of claim **1** comprising:

receiving, at the client system, a plurality of access tokens from the authorisation system, wherein each of the plurality of access tokens is received in response to a token request transmitted to the authorisation system;
receiving and storing an expiry time indicator in association with each access token received at the client system, wherein each of the stored expiry time indicators is indicative of a time at which the corresponding access token will not be valid.

11. The computer-implemented method of claim **10** further comprising:

the client system determining a predicted expiry time for the access tokens received from the authorisation system based on the stored expiry time indicators.

12. The computer-implemented method of claim **11** further comprising:

the client system determining that the predicted expiry time does not meet a predetermined threshold and, in response, not storing further access tokens received from the authorisation system and/or wherein the predicted expiry time is determined by calculating an average of the times indicated by the stored expiry time indicators.

13. The computer-implemented method of claim **1** comprising:

receiving, at the client system from the first user device, a plurality of access requests subsequent to the first access request, each of the plurality of subsequent access requests comprising a subsequent instruction to access a protected resource stored at the resource system; and

in response to the subsequent access request, transmitting the stored access token to the resource system from the client system for receiving the protected resource.

14. The computer-implemented method of claim **1** comprising:

receiving, at the client system from each of a plurality of user devices, an access request comprising an instruction to access a protected resource stored at a resource system;

transmitting, from the client system to an authorisation system, a token request for an access token to be used for accessing the protected resource, in response to each of the access requests received from the plurality of user devices;

receiving, at the client system, an access token in response to each of the token requests, each access token having a corresponding time to expire indicative of a time at which the access token will not be valid for obtaining the protected resource from the resource system;

transmitting each of the access tokens from the client system to the resource system and, in response, receiving the protected resource;

storing each of access tokens at the token storage unit.

15. A computer program comprising instructions which, when the program is executed by a computer, cause the computer to carry out the method of:

receiving, at a client system from a user device, a first access request comprising a first instruction to access a protected resource stored at a resource system;

transmitting, from the client system to an authorisation system, a token request for an access token to be used for accessing the protected resource, in response to the first access request;

receiving, at the client system, an access token in response to the token request, the access token having a corresponding time to expire indicative of a time at which the access token will not be valid for obtaining the protected resource from the resource system;

transmitting the access token from the client system to the resource system and, in response, receiving the protected resource;

storing the access token at a token storage unit of the client system;

receiving, at the client system from the first user device, a second access request comprising a second instruction to access a protected resource stored at the resource system; and

in response to the second access request, transmitting the stored access token to the resource system for receiving the protected resource.

16. A data carrier signal carrying the computer program of claim **15**.

17. A computer readable medium comprising instructions which, when executed by a computer, cause the computer to carry out the method of:

receiving, at a client system from a user device, a first access request comprising a first instruction to access a protected resource stored at a resource system;

transmitting, from the client system to an authorisation system, a token request for an access token to be used for accessing the protected resource, in response to the first access request;

receiving, at the client system, an access token in response to the token request, the access token having a corresponding time to expire indicative of a time at which the access token will not be valid for obtaining the protected resource from the resource system;

transmitting the access token from the client system to the resource system and, in response, receiving the protected resource;

storing the access token at a token storage unit of the client system;

receiving, at the client system from the first user device, a second access request comprising a second instruction to access a protected resource stored at the resource system; and

in response to the second access request, transmitting the stored access token to the resource system for receiving the protected resource.

18. A client system comprising processing circuitry configured to:

receive, from a user device, a first access request comprising a first instruction to access a protected resource stored at a resource system;

transmit to an authorisation system, a token request for an access token to be used for accessing the protected resource, in response to the first access request;

receive an access token in response to the token request, the access token having a corresponding time to expire

indicative of a time at which the access token will not be valid for obtaining the protected resource from the resource system; and
 transmit the access token from the client system to the resource system and, in response, receiving the protected resource;
 wherein the client system further comprises a token storage unit configured to store the access token; and
 the processing further configured to:
 receive, from the first user device, a second access request comprising a second instruction to access a protected resource stored at the resource system; and
 transmit the stored access token to the resource system for receiving the protected resource, in response to the second access request.

19. A client system comprising:

a receiver configured to receive, from a user device, a first access request comprising a first instruction to access a protected resource stored at a resource system;
 a transmitter configured to transmit to an authorisation system, a token request for an access token to be used for accessing the protected resource, in response to the first access request;

wherein the receiver is configured to receive an access token in response to the token request, the access token having a corresponding time to expire indicative of a time at which the access token will not be valid for obtaining the protected resource from the resource system; and

wherein the transmitter is configured to transmit the access token from the client system to the resource system and, in response, receiving the protected resource;

wherein the client system further comprises a token storage unit configured to store the access token;

the receiver configured to receive, from the first user device, a second access request comprising a second instruction to access a protected resource stored at the resource system; and

the transmitter configured to transmit the stored access token to the resource system for receiving the protected resource, in response to the second access request.

* * * * *