



- (51) International Patent Classification:
G06F 21/00 (2013.01)
- (21) International Application Number:
PCT/US2014/015414
- (22) International Filing Date:
7 February 2014 (07.02.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/765,239 12 February 2013 (12.02.2013) US
- (71) Applicant: AMAZON TECHNOLOGIES, INC.
[US/US]; P.O. Box 8102, Reno, Nevada 89507 (US).
- (72) Inventors: ROTH, Gregory Branche; 410 Terry Avenue North, Seattle, Washington 98109-5210 (US). WREN, Matthew James; 410 Terry Avenue North, Seattle, Washington 98109-5210 (US). BRANDWINE, Eric Jason; 410 Terry Avenue North, Seattle, Washington 98109-5210 (US). PRATT, Brian Irl; 410 Terry Avenue North, Seattle, Washington 98109-5210 (US).
- (74) Agents: ADAMS, Scott S. et al.; Davis Wright Tremaine LLP, c/o IP Docketing Dept., 1201 Third Avenue, Suite 2200, Seattle, Washington 98101-3045 (US).

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

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

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

[Continued on next page]

(54) Title: DELAYED DATA ACCESS

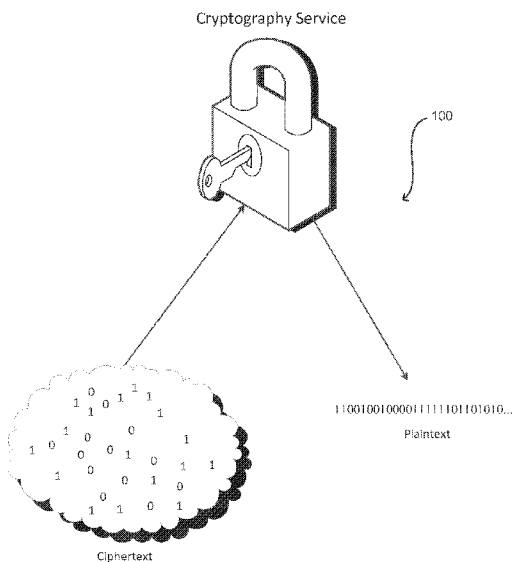


FIG. 1

(57) Abstract: A system uses information submitted in connection with a request to determine if and how to process the request. The information may be electronically signed by a requestor using a key such that the system processing the request can verify that the requestor has the key and that the information is authentic. The information may include information that identifies a holder of a key needed for processing the request, where the holder of the key can be the system or another, possibly third party, system. Requests to decrypt data may be processed to ensure that a certain amount of time passes before access to the decrypted data is provided, thereby providing an opportunity to cancel such requests and/or otherwise mitigate potential security breaches.

WO 2014/126816 A1

Published:

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

DELAYED DATA ACCESS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to U.S. Patent Application No. 13/765,239, filed February 12, 2013, the contents of which are incorporated by reference herein in its entirety.

5 This application incorporates by reference for all purposes the full disclosure of co-pending U.S. Patent Application No. 13/764,944, filed concurrently herewith, entitled "AUTOMATIC KEY ROTATION, co-pending U.S. Patent Application No. 13/764,963, filed concurrently herewith, entitled "DATA SECURITY SERVICE, co-pending U.S. Patent Application No. 13/765,020, filed concurrently herewith, entitled "DATA SECURITY WITH A SECURITY
10 MODULE, co-pending U.S. Patent Application No. 13/764,995, filed concurrently herewith, entitled "POLICY ENFORCEMENT WITH ASSOCIATED DATA, co-pending U.S. Patent Application No. 13/765,209, filed concurrently herewith, entitled "FEDERATED KEY MANAGEMENT, co-pending U.S. Patent Application No. 13/764,963, filed concurrently herewith, entitled "DATA SECURITY SERVICE, and co-pending U.S. Patent Application
15 No. 13/765,283, filed concurrently herewith, entitled "SECURE MANAGEMENT OF INFORMATION USING A SECURITY MODULE".

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Various embodiments in accordance with the present disclosure will be described with reference to the drawings, in which:

20 [0003] FIG. 1 shows an illustrative diagram representing various aspects of the present disclosure in accordance with various embodiments;

[0004] FIG. 2 shows an illustrative example of an environment in which various aspects of the present disclosure may be implemented;

[0005] FIG. 3 shows an illustrative example of an environment in which various aspects of
25 the present disclosure may be implemented and an example flow of information among the various components of the environment in accordance with at least one embodiment;

[0006] FIG. 4 shows example steps of an illustrative process for storing ciphertext, in accordance with at least one embodiment;

[0007] FIG. 5 shows an illustrative example of an environment in which various aspects of the present disclosure may be implemented and an example flow of information among the various components of the environment in accordance with at least one embodiment;

5 [0008] FIG. 6 shows example steps of an illustrative process for responding to a request to retrieve data, in accordance with at least one embodiment;

[0009] FIG. 7 shows an illustrative example of an environment in which various aspects of the present disclosure may be implemented and an example flow of information among the various components of the environment in accordance with at least one embodiment;

10 [0010] FIG. 8 shows example steps of an illustrative process for responding to a request to store data, in accordance with at least one embodiment;

[0011] FIG. 9 shows an illustrative example of an environment in which various aspects of the present disclosure may be implemented and an example flow of information among the various components of the environment in accordance with at least one embodiment;

15 [0012] FIG. 10 shows example steps of an illustrative process for responding to a request to retrieve data, in accordance with at least one embodiment;

[0013] FIG. 11 shows an illustrative example of an environment in which various aspects of the present disclosure may be implemented;

20 [0014] FIG. 12 shows an illustrative example of an environment in which various aspects of the present disclosure may be implemented and an example flow of information among the various components of the environment in accordance with at least one embodiment;

[0015] FIG. 13 shows example steps of an illustrative process for responding to a request to retrieve data, in accordance with at least one embodiment;

[0016] FIG. 14 shows example steps of an illustrative process for responding to a request to decrypt data, in accordance with at least one embodiment;

25 [0017] FIG. 15 shows example steps of an illustrative process for obtaining decrypted data, in accordance with at least one embodiment;

[0018] FIG. 16 shows a diagrammatic representation of an example cryptography service, in accordance with at least one embodiment;

30 [0019] FIG. 17 shows example steps of an illustrative process for configuring policy, in accordance with at least one embodiment;

[0020] FIG. 18 shows example steps of an illustrative process for performing cryptographic operations while enforcing policy, in accordance with at least one embodiment;

[0021] FIG. 19 shows an illustrative example of a process for encrypting data in accordance with at least one embodiment;

5 [0022] FIG. 20 shows an illustrative example of using a security module to encrypt data in accordance with at least one embodiment;

[0023] FIG. 21 shows an illustrative example of using a security module to encrypt a key used to encrypt data in accordance with at least one embodiment;

10 [0024] FIG. 22 shows an illustrative example of a process for enforcing policy using associated data in accordance with at least one embodiment;

[0025] FIG. 23 shows an illustrative example of a process for enforcing policy using associated data and a security module in accordance with at least one embodiment;

[0026] FIG. 24 shows an illustrative example of a state diagram for a policy, in accordance with at least one embodiment;

15 [0027] FIG. 25 shows an illustrative example of another state diagram for a policy, in accordance with at least one embodiment;

[0028] FIG. 26 shows an illustrative example of a process for automatically rotating keys in accordance with at least one embodiment;

20 [0029] FIG. 27 shows an illustrative example of a process for automatically rotating keys in accordance with at least one embodiment;

[0030] FIG. 28 shows an illustrative example of a representation of a database that may be used to track key use in accordance with at least one embodiment;

[0031] FIG. 29 shows an illustrative example of an environment in which various embodiments can be implemented;

25 [0032] FIG. 30 shows a diagrammatic representation of information that may be associated with a key in accordance with at least one embodiment;

[0033] FIG. 31 shows a diagrammatic representation of a key access annotation that may be included in connection with a request, in accordance with at least one embodiment;

30 [0034] FIG. 32 shows an illustrative example of a process for processing requests, in accordance with at least one embodiment;

[0035] FIG. 33 shows an illustrative example of a process for processing requests, in accordance with at least one embodiment; and

[0036] FIG. 34 illustrates an environment in which various embodiments can be implemented.

5

DETAILED DESCRIPTION

[0037] In the following description, various embodiments will be described. For purposes of explanation, specific configurations and details are set forth in order to provide a thorough understanding of the embodiments. However, it will also be apparent to one skilled in the art that the embodiments may be practiced without the specific details. Furthermore, well-known features may be omitted or simplified in order not to obscure the embodiment being described.

10

[0038] Techniques described and suggested herein allow for enhanced data security in environments involving distributed computing resources. In one example, a distributed computing environment includes one or more data services which may be implemented by appropriate computing resources. The data services may allow various operations to be performed in connection with data. As one illustrative example, the distributed computing environment includes one or more data storage services. Electronic requests may be transmitted to the data storage service to perform data storage operations. Example operations are operations to store data using the data storage service and using the data storage service to retrieve data stored by the data storage service. Data services, including data storage services, may also perform operations that manipulate data. For example, in some embodiments, a data storage service is able to encrypt data.

15

20

[0039] Various embodiments of the present disclosure include distributed computing environments that include cryptography services that are implemented using appropriate computing resources. A cryptography service may be implemented by a distributed system that receives and responds to electronic requests to perform cryptographic operations, such as encryption of plaintext and decryption of ciphertext. In some embodiments, a cryptography service manages keys. In response to a request to perform a cryptographic operation, the cryptography service may execute cryptographic operations that use the managed keys. For example, the cryptography service can select an appropriate key to perform the cryptographic operation, perform the cryptographic operation, and provide one or more results of the cryptographic operation in response to the received request. In an alternative configuration, the cryptography service can produce an envelope key (e.g., a session key that is used to

25

30

encrypt specific data items) and return the envelope key to the system invoking the cryptographic operations of the service. The system can then use the envelope key to perform the cryptographic operation.

[0040] In some embodiments, the cryptography service manages keys for multiple tenants of a computing resource service provider. A tenant of the computing resource may be an entity (e.g. organization or individual) operating as a customer of the computing resource provider. The customer may remotely and programmatically configure and operate resources that are physically hosted by the computing resource provider. When a customer submits a request to the cryptography service to perform a cryptographic operation (or when an entity submits a request to the cryptography service), the cryptography service may select a key managed by the cryptography service for the customer to perform the cryptographic operation. The keys managed by the cryptography service may be securely managed so that other users and/or data services do not have access to the keys of others. A lack of access by an entity (e.g., user, customer, service) to the key of another entity may mean that the entity does not have an authorized way of obtaining the key of the other and/or that the entity does not have an authorized way of causing a system that manages the key of the other of using the key at the entity's direction. For example, the cryptography service may manage keys so that, for a customer, other customers both do not have access to the customer's key(s) and are unable to cause the cryptography service to use the customer's key(s) to perform cryptographic operations. As another example, the cryptography service may manage keys so that other services, such as a data storage service, are unable to cause the cryptography service to use some or all keys to perform cryptographic operations. Unauthorized access to a key may be prevented by appropriate security measures so that, for example, unauthorized access is difficult or impossible. Difficulty may be due to computational impracticality and/or due to a need for an unauthorized (e.g., illegal, tortious and/or otherwise disallowed such as a. compromise of authorization credentials) to occur for access to be gained. Systems in accordance with the various embodiments may be configured to ensure an objective measure of computational impracticality to gain access to a key. Such measures may be, for example, measured in terms of an amount of time, on average, it would take a computer having a defined unit of computational ability (e.g. certain operations per unit of time) to crack encrypted information needed for authorized access to the key.

[0041] As noted, a cryptography service may receive requests from various entities, such as customers of a computing resource provider. A cryptography service may also receive requests from entities internal to the computing resource provider. For example, in some

embodiments, data services implemented by the computing resource provider may transmit requests to a cryptography service to cause the cryptography service to perform cryptographic operations. As one example, a customer may transmit a request to a data storage service to store a data object. The request may indicate that the data object should be encrypted when stored. The data storage service may communicate a request to a cryptography service to perform a cryptographic operation. The cryptographic operation may be, for example, encryption of a key used by the data storage service to encrypt the data object. The cryptographic operation may be encryption of the data object itself. The cryptographic operation may be to generate an envelope key that the data storage service can use to encrypt the data object.

[0042] Systems in accordance with the various embodiments implement various security measures to provide enhanced data security. For example, in various embodiments, the manner in which a cryptography service can utilize keys it manages is limited. For example, in some embodiments, a cryptography service is configured to only use a key corresponding to a customer upon appropriate authorization. If a request to use a customer's key purportedly originates from the customer (i.e., from a computing device operating on behalf of the customer), the cryptography service may be configured to require that the request be electronically (digitally) signed using appropriate credentials owned by the customer. If the request to use the customer's key originated from another data service, the cryptography service may be configured to require that the data service provide proof that a signed request to the data service has been made by the customer. In some embodiments, for example, the data service is configured to obtain and provide a token that serves as proof of an authenticated customer request. Other security measures may also be built into a configuration of an electronic environment that includes a cryptography service. For example, in some embodiments, a cryptography service is configured to limit key use according to context. As one illustrative example, the cryptography service may be configured to use a key for encryption for requests from a customer or from a data service acting on behalf of the customer. However, the cryptography service may be configured to only use a key for decryption for requests from the customer (but not from another data service). In this manner, if the data service is compromised, the data service would not be able to cause the cryptography service to decrypt data.

[0043] Various security measures may be built into a cryptography service and/or its electronic environment. Some security measures may be managed according to policies which, in some embodiments, are configurable. As one example, a cryptography service may

utilize an application programming interface (API) that enables users to configure policies on keys. A policy on a key may be information that, when processed by a cryptography service, is determinative of whether the key may be used in certain circumstances. A policy may, for instance, limit identities of users and/or systems able to direct use of the key, limit times
5 when the key can be used, limit data on which the key can be used to perform cryptographic operations on, and provide other limitations. The policies may provide explicit limitations (e.g., who cannot use the keys) and/or may provide explicit authorizations (e.g., who can use the keys). Further, policies may be complexly structured to generally provide conditions for when keys can and cannot be used. When a request to perform a cryptographic operation
10 using a key is received, any policies on the key may be accessed and processed to determine if the request can, according to policy, be fulfilled.

[0044] Various embodiments of the present disclosure are related to the enforcement of policy associated with keys, where the keys may be managed by a cryptography service. Users of the cryptography service, such as customers of a computing resource provider
15 hosting the cryptography service, may specify policies on keys to be enforced by the cryptography service. The policies may encode restrictions and/or privileges in connection with who can direct the cryptography service to use the keys, what operations the keys can be used to perform, in which circumstances the key can be used and/or other key uses.

[0045] In an embodiment, data associated with ciphertext is used in the enforcement of
20 policy. The data associated with the ciphertext may be data that is obtained through use of a cipher, such as modes of the advanced encryption standard (AES). For example, input to a cryptographic algorithm may include plaintext to be encrypted and associated data. The cryptographic algorithm may use a key to encrypt the plaintext and provide authentication output, such as a message authentication code (MAC) that enables a determination whether
25 the associated data has been altered. The authentication output may be determined based at least in part on the associated data and the plaintext.

[0046] Policy enforcement may be based at least in part on associated data. For example, some policies may require associated data to have a particular value before decrypted
30 ciphertext (i.e., plaintext) is provided. The authentication output (e.g., MAC) may be used to ensure that the associated data has not been altered and, therefore, that enforcement of the policy is performed correctly. The associated data may be any suitable data, and the data itself may be specified explicitly or implicitly by policy. For example, a policy may specify that decrypted ciphertext (plaintext) can be provided only if a request to decrypt the ciphertext is submitted by a user having a user identifier encoded in associated data used to

encrypt the ciphertext. In this manner, if another user requests decryption of the ciphertext (without impersonating the user having the user identifier), the request will not be fulfilled due to a conflict with policy. As another example, a policy may state that decrypted ciphertext can be provided only if the ciphertext is tagged with specified information. As yet
5 another example, a policy may state that decrypted ciphertext may be provided if tagged with associated data equal to a hash of the plaintext, hash of the ciphertext, or other specified values. Generally, embodiments of the present disclosure allow for rich policy enforcement around the input or output of a cryptographic algorithm before output of the cryptographic algorithm is revealed. In some embodiments, associated data can itself represent policy.

10 [0047] Various embodiments of the present disclosure also allow for policy around key use. For instance, in some embodiments, keys are automatically rotated to prevent the keys from being used enough time to enable successful cryptographic attacks that can reveal the keys. To prevent a key from being used enough times to result in a potential security breach, a cryptography service or other system utilizing keys may track operations performed with
15 keys. When a key identified by a key identifier (KeyID) is used in a threshold number of operations, the key may be retired (e.g., unusable for future encryption operations, but usable for future decryption operations) and replaced with a new key to be identified by the KeyID. In this manner, new keys are produced in a timely basis. Further, various embodiments of the present disclosure perform such key rotation in a manner that is transparent to certain entities.
20 As one example, a customer of a computing resource provider or other entity may submit requests to a cryptography service to perform operations using a key identified by a KeyID. The cryptography service may perform key rotation independent of any request from the entity to perform the key rotation. From the point of view of the customer or other entity, requests may still be submitted using a specified KeyID without any reprogramming or other
25 reconfiguration necessary due to a key having been retired and replaced with a new key.

[0048] In some embodiments, multiple systems supporting a cryptography or other service simultaneously have access to keys and are used to fulfill requests to perform cryptographic operations. For example, a cryptography service may utilize a cluster of security modules, at least some of which redundantly store one or more keys. The service may allocate operations
30 to the security modules and maintain its own counter. When a security module uses its allocation (e.g., performs an allocated number of operations using a key), the service may check whether the key is still usable or whether the key should be retired. It should be noted that security modules (or other computer systems) may be configured to perform multiple types of operations using a key, such as encryption, decryption, electronic signature

generation, and the like. In some embodiments, not all types of operations cause a security module to use a portion of an allocation of operations. For example, decryption operations may not result in allocated operations being used whereas encryption operations may result in allocated operations used. Generally, in various embodiments, cryptographic operations that result in generation of new information (e.g., ciphertexts and/or electronic signatures) may result in allocated operations being used whereas cryptographic operations that do not result in the generation of new information may not result in allocated operations being used. Further, different types of operations may result in different numbers of cryptographic operations being performed. As one example, encryption of plaintexts may vary in the amount of cryptographic operations required based at least in part on the size of the plaintexts. Use of block ciphers, for instance, may cause an allocated cryptographic operation to be used for each block of ciphertext generated.

[0049] If a total number of operations available for a key is still usable, the service may allocate additional operations to the security module. If the key should be retired (e.g., because a counter so indicates), the service may cause security modules that redundantly store the key to retire the key and replace the key with a new key, where the new key may be generated or otherwise obtained by one security module and securely passed to the remaining security modules. In some embodiments, other security modules instead use up their allocated operations under the older key. If a security module malfunctions, becomes inoperable, is intentionally taken offline (e.g., for maintenance) and/or otherwise becomes unavailable for performing cryptographic operations without providing information regarding how many operations it has performed using one or more keys, the service may treat the unavailability as use of its allocation. For example, if a security module is allocated one million operations for each key in a set of keys, and the security module becomes inoperable, the service may operate as if the security module performed one million operations for each of the set of keys. For example, the service may allocate additional operations to the security module or another security module, adjusting a counter accordingly, and/or may cause one or more of the keys to be retired and replaced, should corresponding counters indicate that replacement is necessary.

[0050] Embodiments of the present disclosure also allow for enhanced data security through annotations and/or federated key management techniques. In some embodiments, requests submitted to a service (such as a cryptography service or other data service) may include or otherwise be associated with an annotation (referred to also as a key access annotation) containing information that enables enforcement of policy. In an embodiment,

the key access annotation must satisfy one or more conditions before a corresponding request can be fulfilled. In some embodiments, the one or more conditions include a condition that the annotation be electronically signed using a key that is associated with a KeyID, where the KeyID identifies a different key usable for fulfilling the request. The presence of a valid
5 electronic signature may both indicate that the information in the annotation has not been modified and also prove possession of the key used to generate the electronic signature.

[0051] In some embodiments, a key access annotation may include an identifier of a holder of a key usable to fulfill the request. The holder of the key may be an entity hosting a system that receives the request or may be another system, such as a system of a third party. A
10 system receiving the request may detect the presence of the identifier and, as appropriate according to the identifier, either process the request itself or transmit the request to the identified key holder for processing. The entity receiving the request and/or the key holder may verify the electronic signature for the enforcement of policy, such as described above and elsewhere herein. For example, if the electronic signature is not valid, the receiver of the
15 request may not pass the request on to a key holder identified in the annotation. Similarly, the key holder may deny the request should it determine that the electronic signature be invalid. The entity receiving the request and the key holder may verify the same electronic signature or, in some embodiments, the request includes at least two signatures, one for the receiver of the request and one for the key holder. Each signature may be generated using a
20 key corresponding to the entity intended to verify the signature.

[0052] Embodiments of the present disclosure also allow for enhanced data security through forced delays before certain types of requests are fulfilled. For instance, in some embodiments, decryption of certain data requires a delay before plaintext is provided in response to a corresponding request. During the delay, various actions may be taken to notify
25 interested parties of the pending request to decrypt the information. In this manner, interested parties (e.g., a compliance officer of an organization or other person authorized to allow the plaintext to be provided) are provided an opportunity to cancel the request before plaintext is provided. In various embodiments a request is easily cancelled. For example, the requirements for canceling the request may be less stringent than the requirements for the
30 request to be fulfilled. In this manner, unauthorized data leaks are easily detectable and/or preventable.

[0053] FIG. 1 is an illustrative diagram 100 demonstrating various embodiments of the present disclosure. In an embodiment, a cryptography service performs cryptographic operations which may include application of one or more calculations in accordance with one

or more cryptographic algorithms. As illustrated in FIG. 1, the cryptography service enables a user or a service to generate plaintext from ciphertext. In an example configuration the cryptographic service can be used to encrypt/decrypt keys and these keys can be used to encrypt/decrypt data, such as data stored in a data storage service. For example, the

5 cryptography service receives a request to generate plaintext from ciphertext encrypted under a key. The cryptography service determines that the requestor is an authorized entity; decrypts the key using a master key and returns the now decrypted key to the service, which can generate plaintext from the ciphertext using the decrypted key. In another configuration, the cryptography service receives ciphertext and processes the received ciphertext into

10 plaintext which is provided as a service by the cryptography service. In this example, the ciphertext may be provided to the cryptography service as part of an electronic request to the cryptography service from an authorized entity, which may be a customer of a computing resource provider that operates the cryptography service and/or which may be another service of the computing resource provider. The cryptography service illustrated in FIG. 1 may

15 utilize one or more cryptographically strong algorithms to encrypt data. Such cryptographically strong algorithms may include, for example, Advanced Encryption Standard (AES), Blowfish, Data Encryption Standard (DES), Triple DES, Serpent or Twofish, and depending on the specific implementation selected, may be either asymmetric or symmetric key systems. Generally, the cryptography service may utilize any encryption

20 and/or decryption algorithm (cipher) or combination of algorithms that utilizes data managed by the cryptography service.

[0054] As will be discussed below in more detail, the cryptography service can be implemented in a variety of ways. In an embodiment, the cryptography service is implemented by a computer system configured in accordance with the description below.

25 The computer system may itself comprise one or more computer systems. For example, a cryptography service may be implemented as a network of computer systems collectively configured to perform cryptographic operations in accordance with the various embodiments. Or put another way, the computer system could be a distributed system. In an embodiment, ciphertext is information that has been encrypted using a cryptographic algorithm. In the

30 example of FIG. 1, the ciphertext is the plaintext in an encrypted form. The plaintext may be any information and while the name includes no word text, plaintext and ciphertext may be information encoded in any suitable form and does not necessarily include textual information but it may include textual information. For example, as illustrated in FIG. 1, plan text and ciphertext comprise sequences of bits. Plaintext and ciphertext may also be

represented in other ways and generally in any manner by which encryption and decryption can be performed by a computer system.

[0055] FIG. 2 shows an illustrative example of an environment 200 in which a cryptography service, such as illustrated in FIG. 1, may be implemented. In the environment of 200, various components operate together in order to provide secure data related services. In this particular example, the environment 200 includes a cryptography service, an authentication service, a data service frontend and a data service backend storage system. In an embodiment, the cryptography service is configured in the environment 200 to perform cryptographic operations, such as by receiving plaintext from a data service frontend and providing ciphertext in return or providing envelope keys to services, so that the services can use the envelope keys to perform encryption operations. The cryptography service may perform additional functions, such as described below, such as secure storage of keys for performance of cryptographic operations, such as converting plaintext into ciphertext and decrypting ciphertext into plaintext. The cryptography service may also perform operations involved in policy enforcement, such as by enforcing policy associated with keys stored therein. Example policies that may be enforced by a cryptography service are provided below. The data service frontend in an embodiment is a system configured to receive and respond to requests transmitted over a network from various users. The requests may be requests to perform operations in connection with data stored or to be stored in the data service backend storage system. In the environment 200, the authentication service, cryptography service, data service frontend and data service backend storage system may be systems of a computing resource provider that utilizes the systems to provide services to customers represented by the users illustrated in FIG. 2. The network illustrated in FIG. 2 may be any suitable network or combination of networks, including those discussed below.

[0056] The authentication service in an embodiment is a computer system configured to perform operations involved in authentication of the users. For instance, the data service frontend may provide information from the users to the authentication service to receive information in return that indicates whether or not the user requests are authentic. Determining whether user requests are authentic may be performed in any suitable manner and the manner in which authentication is performed may vary among the various embodiments. For example, in some embodiments, users electronically sign messages transmitted to the data service frontend. Electronic signatures may be generated using secret information (e.g., a private key of a key pair associated with a user) that is available to both an authenticating entity (e.g., user) and the authentication service. The request and signatures

for the request may be provided to the authentication service which may, using the secret information, compute a reference signature for comparison with the received signature to determine whether the request is authentic. If the request is authentic, the authentication service may provide information that the data service frontend can use to prove to other services, such as the cryptography service, that the request is authentic, thereby enabling the other services to operate accordingly. For example, the authentication service may provide a token that another service can analyze to verify the authenticity of the request. Electronic signatures and/or tokens may have validity that is limited in various ways. For example, electronic signatures and/or tokens may be valid for certain amounts of time. In one example, electronic signatures and/or tokens are generated based at least in part on a function (e.g., a Hash-based Message Authentication Code) that takes as input a timestamp, which is included with the electronic signatures and/or tokens for verification. An entity verifying a submitted electronic signature and/or token may check that a received timestamp is sufficiently current (e.g. within a predetermined amount of time from a current time) and generate a reference signature/token using for the received timestamp. If the timestamp used to generate the submitted electronic signature/token is not sufficiently current and/or the submitted signature/token and reference signature/token do not match, authentication may fail. In this manner, if an electronic signature is compromised, it would only be valid for a short amount of time, thereby limiting potential harm caused by the compromise. It should be noted that other ways of verifying authenticity are also considered as being within the scope of the present disclosure.

[0057] The data service backend storage system in an embodiment is a computer system that stores data in accordance with requests received through the data service frontend. As discussed in more detail below, the data service backend storage system may store data in encrypted form. Data in the data service backend storage system may also be stored in unencrypted form. In some embodiments, an API implemented by the data service frontend allows requests to specify whether data to be stored in the data service backend storage system should be encrypted. Data that is encrypted and stored in the data service backend storage system may be encrypted in various ways in accordance with the various embodiments. For example, in various embodiments, data is encrypted using keys accessible to the cryptography service, but inaccessible to some or all other systems of the environment 200. Data may be encoded by the cryptography service for storage in the data service backend storage system and/or, in some embodiments, data may be encrypted by another system, such as a user system or a system of the data service frontend, using a key that was

decrypted by the cryptography service. Examples of various ways by which the environment 200 may operate to encrypt data are provided below.

[0058] Numerous variations of the environment 200 (and other environments described herein) are considered as being within the scope of the present disclosure. For example, the environment 200 may include additional services that may communicate with the cryptography service and/or authentication service. For example, the environment 200 may include additional data storage services (which may each comprise a frontend system and a backend system) which may store data in different ways. For instance, one data storage service may provide active access to data where the data storage service performs data storage services in a synchronous manner (e.g., a request to retrieve data may receive a synchronous response with the retrieved data). Another data storage service may provide archival data storage services. Such an archival data storage service may utilize asynchronous request processing. For example, a request to retrieve data may not receive a synchronous response that includes the retrieved data. Rather, the archival data storage service may require a second request to be submitted to obtain the retrieved data once the archival data storage service is ready to provide the retrieved data. As another example, the environment 200 may include a metering service which receives information from the cryptography service (and/or other services) and uses that information to produce accounting records. The accounting records may be used to bill customers for usage of the cryptography service (and/or other services). Further, information from the cryptography service may provide an indication of how charges should be incurred. For instance, in some instances customers may be provided bills for use of the cryptography service. In other instances, charges for use of the cryptography service may be rolled into usage charges of other services, such as a data service that utilizes the cryptography service as part of its operations. Usage may be metered and billed in various ways, such as per operation, per period of time, and/or in other ways. Other data services may also be included in the environment 200 (or other environments described herein).

[0059] In addition, FIG. 2 depicts users interacting with the data service frontend. It should be understood that users may interact with the data service frontend through user devices (e.g. computers) which are not illustrated in the figure. Further, the users depicted in FIG. 2 (and elsewhere in the figures) may also represent non-human entities. For example, automated processes executing on computer systems may interact with the data service frontend as described herein. As one illustrative example, an entity represented by a user in FIG. 2 may be a server that, as part of its operations, uses the data service frontend to store and/or retrieve

data to/from the data service backend storage system. As yet another example, an entity represented by a user in FIG. 2 may be an entity provided as a service of a computing resource provider that operates one or more of the services in FIG. 2. For instance, a user in FIG. 2 may represent a virtual or other computer system of a program execution service offered by the computing resource provider. Other variations, including variations of other environments described below, are also considered as being within the scope of the present disclosure.

[0060] For example, FIG. 3 shows an illustrative example of an environment 300 in which various embodiments of the present disclosure may be implemented. As with FIG. 2, the environment in FIG. 3 includes an authentication service, a data service frontend system (data service front end), a cryptography service and a data service backend storage system. The authentication service, data service frontend, cryptography service and data service backend storage system may be configured such as described above in connection with FIG. 2. For example, users may access the data service frontend through a suitable communications network, although such network is not illustrated in the figure. In the example environment 300 illustrated in FIG. 3, arrows representing a flow of information are provided. In this example, a user transmits a PUT request to the data service frontend. The PUT request may be a request to store specified data in the data service backend storage system. In response to the PUT request, the data service frontend may determine whether the PUT request is authentic, that is if the user has submitted the request in the manner the requested operation can be performed in accordance with authentication policy implemented by the system.

[0061] In FIG. 3, an illustrative example of how such authentication decisions may be made is illustrated. In this particular example, the data service frontend submits an authentication request to the authentication service. The authentication service may use the authentication request to determine if the PUT request from the user is authentic. If the request is authentic, the authentication service may provide authentication proof to the data service frontend. The authentication proof may be an electronic token or other information that is usable by another service, such as the cryptography service, to independently determine that an authentic request was received. In one illustrative example, the PUT request is transmitted with a signature for the PUT request. The PUT request and its signature are provided through the authentication service, which independently computes what a signature should be if authentic. If the signature generated by the authentication service matches that signature provided by the user, the authentication service may determine that the PUT request was authentic and may provide authentication proof in response.

Determining whether the PUT request is authentic may also include one or more operations in connection with the enforcement of policy. For example, if the signature is valid but policy otherwise indicates that the PUT request should not be completed (e.g. the request was submitted during a time disallowed by policy), the authentication service may provide

5 information indicating that the request is not authentic. (It should be noted, however, that such policy enforcement may be performed by other components of the environment 300.) The authentication service may generate the signature, such as by using a key shared by the authentication service and the user. The authentication proof, as noted, may be information from which another service, such as the cryptography service, can independently verify that a
10 request is authentic. For example, using the example of the cryptography service illustrated in FIG. 3, the authentication proof may be generated based at least in part on a key shared by both the authentication service and the cryptography service, such as a key that is inaccessible to other services.

[0062] As illustrated in FIG. 3, the data service frontend, upon receipt of authentication
15 proof from the authentication service, provides plaintext and authentication proof to the cryptography service. The plaintext and authentication proof may be provided according to an API call or other electronic request to the cryptography service (e.g., an Encrypt API call). The cryptography service may analyze the authentication proof to determine whether to encrypt the plaintext.

20 [0063] It should be noted that additional information may be provided to the cryptography service. For example, an identifier of a key to be used to encrypt the plaintext may be provided as an input parameter to the API call from the data service frontend (which, in turn, may have received the identifier from the user). However, it should be noted that an identifier may not be transmitted to the cryptography service. For instance, in various
25 embodiments, it may be otherwise determinable which key to use to encrypt the plaintext. For example, information transmitted from the data service frontend to the cryptography service may include information associated with the user, such as an identifier of the user and/or an organization associated with the user, such as an identifier of a customer on behalf of which the user has submitted the PUT request. Such information may be used by the
30 cryptography service to determine a default key to be used. In other words, the key may be implicitly specified by information that is usable to determine the key. Generally, determination of the key to be used may be performed in any suitable manner. Further, in some embodiments, the cryptography service may generate or select a key and provide an identifier of the generated or selected key to be used later. Another example API parameter

can be an identifier for a master key for the customer account the encryption operation is being performed for.

[0064] As illustrated in FIG. 3, if the authentication proof is sufficient to the cryptography service for the plaintext to be encrypted, the cryptography service can perform one or more cryptographic operations. In an embodiment, the one or more cryptographic operations can include an operation to generate an envelope key to be used to encrypt the plaintext. The envelope key can be a randomly generated symmetric key or a private key of a key pair. After the envelope key is generated, the cryptographic service can encrypt the envelope key with the master key specified in the API call and cause the encrypted envelope key to be persistently stored (e.g., by storing the encrypted key in a storage service or some other durable storage) or discarded. In addition, the cryptographic service can send a plaintext version of the envelope key as well as and the encrypted envelope key to the data service frontend. The data service can then use the plaintext version of the envelope key to encrypt the plaintext (i.e., the data associated with the encryption request) and cause the envelope key to be stored in persistent storage in association with an identifier for the master key used to encrypt the envelope key. In addition, the data service can discard the plaintext version of the envelope key. As such, in an embodiment after the data service discards the plaintext version of the envelope key it will no longer be able to decrypt the ciphertext.

[0065] In an alternative embodiment, the cryptographic operation can involve encrypting the plaintext. For example, the cryptographic service encrypts the plaintext and provides ciphertext to the data service frontend storage system. The data service frontend may then provide the ciphertext to the data service backend storage system for persistent storage in accordance with its operation. Other information may also be transmitted from the data service frontend to the data service backend storage system. For example, an identifier of the key used to encrypt the plaintext to generate ciphertext may be provided with the ciphertext for storage by the data service backend storage system. Other information, such as metadata identifying the user and/or the user's organization, may also be provided.

[0066] As with all environments described herein, numerous variations are considered as being within the scope of the present disclosure. For example, the flow of information among the various components of the environment 300 may vary from that which is shown. For example, information flowing from one component to another component through an intermediate component (e.g. data from the authentication service to the cryptography service and/or data from the cryptography service to the data service backend storage system) may be provided directly to its destination and/or through other intermediate components of the

environment 300 (which are not necessarily included in the figure). As another example, PUT requests (and, below, GET requests) are provided for the purpose of illustration. However, any suitable request for performing the operations described may be used.

[0067] FIG. 4 shows an illustrative example of a process 400 which may be used to store data in a data storage service in accordance with an embodiment. The process 400 may be performed, for example, by the data service frontend illustrated in FIG. 3. Some or all of the process 400 (or any other processes described herein, or variations and/or combinations thereof) may be performed under the control of one or more computer systems configured with executable instructions and may be implemented as code (e.g., executable instructions, one or more computer programs or one or more applications) executing collectively on one or more processors, by hardware or combinations thereof. The code may be stored on a computer-readable storage medium, for example, in the form of a computer program comprising a plurality of instructions executable by one or more processors. The computer-readable storage medium may be non-transitory.

[0068] As illustrated in FIG. 4, the process 400 includes receiving 402 a PUT request. The PUT request may be electronically received over a network and may include information associated with the request, such as information required for authentication, such as an electronic signature of the PUT request. In response to having received the PUT request, the process 400 may include submitting 404 an authentication request. For example, the system performed in the process 400 may submit (e.g., via an appropriately configured API call) an authentication request to a separate authentication service, such as described above in connection with FIG. 3. Similarly, a data service frontend that performs its own authentication may submit the authentication request to an authentication module implemented by the data service frontend. Generally, the authentication request may be submitted in any suitable manner in accordance with the various embodiments.

[0069] Upon submission of the authentication request, an authentication response is received 406 by the entity to which the authentication request was submitted 404. For example, referring to FIG. 3, the authentication service may provide a response to the data service frontend that includes proof of the authentication for use by other services. Other information, such as an indication of whether or not authentication was successful, may also be transmitted. A determination may be made 408 whether the request is authentic. Authenticity of the request may depend from one or more factors checked by an entity, such as by an authentication service, or a combination of entities that collectively perform such checks. Authenticity may, for example, require that the request provide necessary valid

credentials (e.g., an electronic signature generated by a secret key shared by the checking entity) and/or that policy allows the request to be fulfilled. From the perspective of a system that submits 404 an authentication request and receives an authentication response, authenticity may depend from the received authentication response. Accordingly, in an embodiment, the determination 408 whether the request is authentic may be performed based at least in part of the received authentication response. For example, if authentication was not authentic, the authentication response so indicates and the determination 408 may be made accordingly. Similarly, the response may implicitly indicate that the authentication request is authentic, such as by not including information that would be included if the request was authentic. If it is determined 408 that the PUT request is not authentic, then the PUT request may be denied 410. Denying the PUT request may be performed in any suitable manner and may depend upon the various embodiments in which the process 400 is being performed. For example, denying 410, the PUT request may include transmitting a message to a user that submitted the PUT request. The message may indicate that the request was denied. Denying the request may also include providing information about why the request was denied, such as an electronic signature not being correct or other reasons that may be used for determining how to resolve any issues that resulted in the PUT request not being authentic or authorized.

[0070] If it is determined 408 that the PUT request is authentic and authorized, then, in an embodiment, the process 400 includes performing 412 one or more cryptographic operations that result in the plaintext being encrypted. For example, a request (e.g., an appropriately configured API call) may be submitted to a cryptography service to provide a key to be used for performing the one or more cryptographic operations. The request provided to the cryptography service may be provided with proof of the PUT request being authentic so that the cryptography service can independently determine whether to perform the cryptographic operation (e.g., to encrypt the plaintext and provide ciphertext or generate an envelope key that can be used to encrypt the plaintext). However, in various embodiments, authentication proof may not be provided to the cryptography service and, for example, the cryptography service may operate in accordance with the request that it receives. For example, if the cryptography service receives a request from the data service frontend, the cryptography service may rely on the fact that the data service frontend has already independently verified authentication of the request. In such an embodiment and other embodiments, the data service frontend may authenticate itself with the cryptography service to provide an additional layer of security. The cryptography service may generate or otherwise obtain a key, encrypt the obtained key or otherwise obtain the encrypted key (e.g., from memory), and

provide the obtained key and the encrypted obtained key in response to the request. The obtained key may be encrypted using a key identified in the request to the cryptography service. The obtained key may be used to encrypt the plaintext and, after encrypting the plaintext, the obtained key may be discarded (e.g., irrevocably removed from memory). In

5 alternate embodiment, a system performing the process 400 may generate or otherwise obtain the key used to perform the one or more cryptographic operations, provide the obtained key to the cryptography service to be encrypted.

[0071] In some embodiments, performing the one or more cryptographic operations may result in ciphertext being generated. Ciphertext generated as a result of one or more

10 cryptographic operations may be stored 414 for possible retrieval at a later time. As noted above, storage of the ciphertext may include storage of additional information that would enable the decryption of the ciphertext at later time. For instance, the ciphertext may be stored with an identifier of a key used to encrypt the plaintext into the ciphertext so that the key having that identifier may later be used to decrypt the ciphertext to obtain the plaintext.

15 Storage of the ciphertext may also be performed in any suitable manner. For example, storage of the ciphertext may be performed by a data service backend storage system, such as described above.

[0072] FIG. 5, accordingly, shows an illustrative example of an environment 500 and the flow of information illustrating how plaintext may be obtained. The environment 500 in this

20 example includes an authentication service, a cryptography service, a data service frontend and a data service backend storage system. The authentication service, cryptography service, data service frontend and a data service backend storage system may be systems such as described above. As illustrated in FIG. 5, the data service frontend is configured to receive a GET request from a user and provide plaintext in response. In order to do this, the data

25 service frontend may also be configured to submit authentication requests to an authentication service which itself may be configured to, if appropriate, provide to the data service frontend authentication proof. The data service frontend may also be configured to send a request to the cryptographic service to cause it to execute one or more cryptographic operations related to decrypting the data. In an embodiment where envelope keys are used,

30 the data service can submit a request (e.g., an API call) to the cryptographic service that includes or specifies the encrypted envelope key (or an identifier for the encrypted envelope key) authentication proof, and an identifier of the master key used to encrypt the envelope key to the cryptographic service. The cryptographic service can determine whether the authentication proof is sufficient to allow the operation and, if the authentication proof is

sufficient, decrypt the envelope key. The decrypted envelope key can be sent back to the data service, which can use the key to decrypt the encrypted plaintext. The data service can then discard the decrypted plaintext key.

[0073] In an alternative embodiment, the data service frontend may be configured to provide the received authentication proof to the cryptography service with ciphertext for the cryptography service to decrypt. The cryptography service may, accordingly, be configured to determine whether the authentication proof is sufficient to allow decryption of the ciphertext and, if the authentication proof is sufficient, decrypt the ciphertext using an appropriate key (which may be identified to the cryptography service by the data service frontend), and provide the decrypted ciphertext (plaintext) to the data service frontend. To provide ciphertext to the cryptography service, the data service frontend may be configured to obtain (e.g., via an appropriately configured API call) the ciphertext from the data service backend storage system.

[0074] FIG. 6 shows an illustrative example of a process 600 which may be used to obtain plaintext in accordance with various embodiments. The process 600 may be performed, for example, by the data service frontend system (data service frontend) illustrated above in connection FIG. 5, although the process 600 and variations thereof may be performed by any suitable system. In an embodiment, the process 600 includes receiving 602 a GET request (or other appropriate request) from a user. Receiving the GET request may be performed such as described above in connection with other types of requests. Upon receipt 602 of the GET request, an authentication request may be submitted 604 to an authentication service or in any manner such as described above. An authentication response may, accordingly, be received. Based at least in part on the receive authentication response, a determination may be made 608 whether the GET request is authentic. If it is determined 608 that the GET request is not authentic, the process 600 may include denying 610 the request which, as described above, may be performed in various manners in accordance with the various embodiments.

[0075] If it is determined 608 that the GET request is authentic, the process 600 may include retrieving ciphertext from storage. Retrieving 612 ciphertext from storage may be performed in any suitable manner. For example, referring to the environment 500 discussed above in connection with FIG. 5, the data service frontend may submit a request for the ciphertext to the data service backend storage system and may receive the ciphertext in response. Generally, ciphertext may be obtained from storage in any suitable manner. Upon receipt of the ciphertext, the process 600 may include performing 614 one or more operations

relating to decrypting the ciphertext. For example, in an embodiment, the data storage service may send a request to the cryptographic service to perform one or more cryptographic operations relating to decrypting the ciphertext 614. In one example configuration, the data service can send the cryptographic service an API call that includes the encrypted envelope key (or an identifier for the encrypted envelope key) authentication proof, and an identifier of the master key used to encrypt the envelope key to the cryptographic service. The cryptographic service can determine whether the authentication proof is sufficient to allow the operation and, if the authentication proof is sufficient, decrypt the envelope key. The decrypted envelope key can be sent back to the data service, which can use the key to decrypt the encrypted plaintext.

[0076] In another configuration, the ciphertext may be provided to a cryptography service such as the cryptography service described above in connection with FIG. 5. Other information may also be provided to the cryptography service such as proof of authentication that can be used by the cryptography service to determine whether or not to decrypt the ciphertext. In addition, in some embodiments, an identifier of a key to be used by the cryptography service to decrypt the ciphertext may be provided to the cryptography service. However, in other embodiments, the key may be implicitly indicated to the cryptography service. For example, the cryptography service may use a default key associated with a customer that is indicated to the cryptography service. Generally, any manner by which the cryptography service can determine which key to use to decrypt the ciphertext may be used.

[0077] As illustrated in FIG. 6, after the ciphertext is decrypted, the process 600 may include providing 616 a response to the GET request. Providing a response to the GET request may be performed in various ways in accordance with the various embodiments. For example, providing the response to the GET request may include providing the plaintext. In other embodiments, the plaintext may be a key that is used to decrypt other encrypted information which is then provided in response to the GET request. Generally, depending on the role of the plaintext in a particular embodiment of the disclosure, providing a response to the GET request may be performed in various ways.

[0078] As noted, various embodiments of the present disclosure allow for data to be stored by a data storage service in various ways. FIG. 7 shows an illustrative example of an environment 700 with arrows indicating information flow in accordance with such embodiment. As illustrated in FIG. 7, environment 700 includes an authentication service, a cryptography service, a data service frontend and a data service backend storage system, such as described above. In this particular example, the data service frontend is a computer system

configured to receive PUT requests from various users. PUT requests may include or specify data objects to be stored by a data service backend storage system. PUT requests may also specify a key identifier for a key to be used to encrypt the data object. The data service frontend may also be configured to interact with an authentication service, such as described above, in order to provide authentication proof to a cryptography service that is operable to receive keys and key identifiers and provide in response keys encrypted by the keys identified by the key identifiers. The data service frontend may then cause storage in a data service backend storage system. The data that may be stored may include a data object encrypted by the key. The data that may be stored may also include the key encrypted by the key identified by the key identifier. As discussed elsewhere, herein, the encrypted data object and encrypted key may be stored in different services.

[0079] As illustrated in FIG. 7, the data service frontend is configured to provide encrypted information to the data service backend storage system for storage. In this example, the data service frontend is configured to provide a data object encrypted under a key and the key encrypted under another key having a KeyID. It should be noted that, for the purpose of illustration, curly bracket notation is used to denote encryption. In particular, information inside of curly brackets is the information that is encrypted under a key specified in subscript. For example, {Data Object}_{Key} denotes that the data “Data Object” is encrypted under the key “Key.” It should be noted that a key identifier may also appear in subscript using this curly bracket notation. When a key identifier appears in subscript, the information inside the curly brackets is encrypted under a key identified by the key identifier. For example, {Data Object}_{KeyID} denotes that the data object “Data Object” is encrypted under a key identified by the key identifier “KeyID.” Similarly, {Key}_{KeyID} denotes that the key “Key” is encrypted under the key identified by the key identifier “KeyID.” In other words, this disclosure makes use of both keys and key identifiers in subscript and the meaning of the subscript should be clear from context. The ciphertext may include additional metadata usable to determine the identity of the associated decryption key.

[0080] FIG. 8 shows an illustrative example of a process 800 which may be performed to store a data object in a data storage system, such as the data service backend storage system described above in connection with FIG. 7. The process 800 may be performed by any suitable system, such as by the data service frontend system described above in connection with FIG. 7. In an embodiment, the process 800 includes receiving 802 a PUT request for a data object. Receiving the PUT request for the data object may be performed in any suitable manner, such as described above. It should be noted that the data object can be received in

connection with the request or may be received from another service. For example, the request may include an identifier for a data object that may be obtained from another service using the identifier. As with other processes described above, the process 800 in an embodiment includes submitting 804 an authentication request and receiving 806 an authentication response. The authentication response that was received 806 may be used to determine 808 whether the PUT request is an authentic request. If it is determined 808 that the PUT request is not authentic, the process 800 may include denying 810 the request such as described above. If it is determined 808 that the PUT request is authentic, the process 800 may include obtaining 812 a key identifier (KeyID), such as a keyID for a master key used to encrypt envelope keys. Obtaining 812 the KeyID may be performed in any suitable manner and the manner in which the KeyID is obtained may vary in accordance with the various embodiments. For example, as illustrated in FIG. 7, the PUT request may specify the KeyID. As another example, the identity of the user, or otherwise associated with the user, may be used to obtain an identifier or a default key. As another example, the ciphertext may provide an indication of the associated key ID. As yet another example, one or more policy determinations may be used to determine which key identifier to obtain.

[0081] In an embodiment, the process 800 also includes generating 814 a key, such as an envelope key. Generating the key may be performed in any suitable manner by, for example, the cryptographic service or the service requesting encryption operations from the cryptographic service (e.g., a data storage service). For example, the key may be generated using a key derivation function using appropriate input to the key derivation function. Example key derivation functions include KDF1, defined in IEEE Std 1363 2000, key derivation functions defined in ANSI X9.42, and HMAC-based key derivation functions, such as the HMAC-Based Extract-and-Expand Key Derivation Function (HKDF) specified in RFC 5869. As another example, the key may be generated by a random or pseudo random number generator, a hardware entropy source or a deterministic random bit generator such as is specified by National Institute of Standards and Technology Special Publication (NIST SP) 800-90A. It should be noted that while FIG. 8 shows the process 800 including generating 814 a key, the key may be obtained in other ways such as by retrieval from storage. In other words, the key may have been pre-generated.

[0082] Continuing with the process 800 illustrated in FIG. 8, in an embodiment, the process 800 includes using 816 the generated key to encrypt a data object. For example, in an embodiment where the cryptographic service generates the key, the cryptographic service can provide the key, the KeyID, and an encrypted copy of the key to the data service. For

example, referring to FIG. 7, the data service frontend may receive the envelope key and the KeyID for the master key used to encrypt the envelope key from the cryptography service with any other relevant information, such as authentication proof. The plaintext copy of the encryption key may then be used to encrypt the data object. The plaintext copy of the encryption key can be discarded and the encrypted data object as well as the encrypted key may then be stored 818. For example, referring to FIG. 7, the data service frontend may transmit to the encrypted data object and the encrypted key to the data service backend storage system for storage. In a configuration where the service generates the key, the service can provide the key and the KeyID to the cryptography service. For example, the data service frontend may send the envelope key and the KeyID for the master key used to encrypt the envelope key to the cryptography service with any other relevant information, such as authentication proof. The plaintext copy of the encryption key may then be used to encrypt the data object. The service can discard the plaintext copy of the encryption key and the encrypted data object as well as the encrypted key may then be stored. For example, referring to FIG. 7, the data service frontend may transmit to the encrypted data object and the encrypted key to the data service backend storage system for storage.

[0083] The encrypted data object and the encrypted envelope key may be stored without the plaintext version of key, that is, the plaintext key may be inaccessible to the data service backend storage system and one or more other systems. The key under which the data object is encrypted (e.g., the master key) may be made inaccessible in any suitable manner. In some embodiments this is achieved by storing it in memory accessible only to the cryptographic service. In some other embodiments this can be achieved by storing the master key in a hardware or other security module or otherwise under the protection of a hardware or other security module. In some embodiments, the memory location storing the plaintext envelope key (e.g., memory of the data service) may be allowed to be overwritten or memory location storing the key may be intentionally overwritten to render inaccessible the key to the data service frontend. As another example, the plaintext envelope key may be maintained in volatile memory that eventually ceases to store the key. In this manner, the envelope key is only accessible if it is decrypted using a key identified by the KeyID or otherwise obtained in an unauthorized manner, such as by cracking the key without the key identified by the KeyID, which may be computationally impractical. In other words, the key identified by the KeyID is required for authorized access to the key under which the data object is encrypted. Thus, if the data service backend storage system of FIG. 7 is compromised, such compromise would not provide access to the unencrypted data object because decrypting the data object

would require access to the key, which is only obtainable through decryption using the key identified by the KeyID or through other ways which are not computationally feasible.

[0084] As noted, various embodiments of the present disclosure allow users to both store data objects and retrieve them in a secure manner. FIG. 9, accordingly, shows an illustrative example of an environment 900 which may be used to obtain data objects from storage. As illustrated in FIG. 9, the environment 900 includes an authentication service, a cryptography service, a data service frontend system and a data service backend storage system. The authentication service, cryptography service, data service frontend and data service backend storage system may be computer systems such as described above. As illustrated in FIG. 9, the data service frontend system is configured to receive data object requests and provide data objects in response. In order to provide the data objects in response, the data storage frontend system in this embodiment is configured to interact with the authentication service, the cryptography service, and the data service backend storage system as illustrated in FIG. 9. For example, in various embodiments, the data service frontend system is configured to submit authentication requests to the authentication service and receives authentication proof in response to the requests. As another example, the data service frontend is configured to provide keys encrypted by a key identified by a KeyID and authentication proof to a cryptography service which is operable to determine whether to provide the key based, at least in part, on the authentication proof and, if determined to provide the key, then provide the key to the data service frontend. The data service frontend may also be configured to provide other information, such as the KeyID, to the cryptography service. Although, in some embodiments, the KeyID may be implicitly indicated to the cryptography service, such as through association with other information provided to the cryptography service. It should also be noted that, in some embodiments, the user provides the KeyID to the data service frontend in connection with submitting requests to the data service frontend. Also, as illustrated in FIG. 9, the data service frontend in an embodiment is configured to request the data object from the data service backend storage system and receive in response the data object encrypted by the key and the key encrypted by a key identified by the KeyID. In some embodiments the cryptographic service may be operable to refuse to perform decryption of a ciphertext not generated using a key associated with a specified KeyID.

[0085] The data service frontend, in an embodiment, is configured to use the key received from the cryptography service to decrypt the data object and provide the decrypted data object to the user. FIG. 10, accordingly, shows an illustrative example of a process 1000 which may be used to provide the decrypted object in accordance with various embodiments.

The process 1000 may be performed by any suitable system such as the data service frontend system described in connection with FIG. 9. In an embodiment, the process 1000 includes receiving 1002 a GET request for a data object. Receiving the GET request for the data object may be performed in any suitable manner such as described above in connection with
5 other types of requests. For example, the GET request for the data object may include information used to authenticate the request and/or other information. The process 1000, accordingly, in an embodiment, as with other processes described herein, includes submitting 1004 an authentication request to an authentication system and receiving 1006 an authentication response. Submitting the authentication request and receiving the
10 authentication response may be performed in any suitable manner, such as described above. The authentication response may be used to determine 1008 whether or not the GET request is authentic. If it is determined 1008 that the GET request is not authentic, the process 1000 in an embodiment includes denying 1010 the request. If, however, it is determined 1008 that the GET request is authentic, the process 1000 in an embodiment includes retrieving 1012 the
15 encrypted data object and an encrypted key from storage. For example, the data service frontend system may obtain the encrypted data object and encrypted key from the data service backend storage system illustrated above in connection with FIG. 9.

[0086] In an embodiment, the process 1000 includes providing 1014 the encrypted envelope key to a cryptography service. Providing 1014 the encrypted envelope key to the
20 cryptography service may be performed in any suitable manner and may be provided with other information, such as authentication proof that enables the cryptography service to determine whether to decrypt the encrypted key. In addition, providing 1014 the encrypted envelope key to the cryptography service may include providing an identifier of a key required for authorized decryption of the encrypted envelope key to enable the cryptography
25 service to select a key identified by the identifier from among multiple keys managed by the cryptography service. As noted above, however, keys may be implicitly identified. The cryptography service may, accordingly, select an appropriate key and decrypt the encrypted key. Accordingly, in an embodiment, the process 1000 includes receiving 1016 the decrypted envelope key from the cryptography service. For example, if the cryptography service
30 determines that the authentication proof is valid and/or that decryption of the encrypted is allowable according to any applicable policies, the cryptography service may provide the decrypted key to a system attempting to decrypt the data object. The data object may then be decrypted 1018 using the decrypted envelope key. The decrypted data object may then be provided 1020 to the requestor, such as the user or other system that submitted the GET
35 request.

[0087] In many instances, it is desirable for users (i.e., generally devices utilizing a cryptography service) to interact directly with the cryptography service. FIG. 11 accordingly shows an illustrative example of an environment 1100 which allows for direct user access to a cryptography service. In environment 1100, included is an authentication service, a data service frontend and a data service backend storage system. The authentication service, data service frontend and data service backend storage system may be as described above. For example, the data service frontend may be configured to receive and respond to requests from users as illustrated in FIG. 11 over a suitable network. As part of responding to requests from the users over the network, the data service frontend may also be configured to interact with the authentication service in order to determine whether user requests are authentic and/or enforce policy on the requests. The data service frontend may also be configured to interact with the data service backend storage system as part of fulfilling user requests. User requests may include, for example, PUT requests to store data in the backend storage system and GET requests to retrieve data from the data service backend storage system. As above, other requests may also be used in accordance with the various embodiments, such as requests to delete data stored in the data service backend storage system, requests to update the data stored in the data service backend storage system and the like.

[0088] In the particular example of FIG. 11, in the environment 1100, the cryptography service includes a cryptography service frontend and a data service backend. As with the data service frontend, the cryptography service frontend is configured to receive and respond to requests from users over the network. The cryptography service frontend is also configured to interact with the authentication service to determine whether user requests are authentic. Determining whether user requests are authentic may be performed in a simple manner such as described above. It should be noted that, while the cryptography service frontend and the data service frontend interact with the same authentication service, the cryptography service frontend and the data service frontend may interact with different authentication services. Further, the cryptography service frontend may be configured to enforce policy when responding to user requests.

[0089] The cryptography service frontend, in an embodiment, is configured to interact with the cryptography service backend. The cryptography service backend is configured, in accordance with instructions received from the cryptography service frontend, to perform cryptographic operations. Cryptographic operations include encryption, decryption and hash calculations and others. The environment 1100 may be used, for example, by users to have plaintext encrypted by the cryptography service so that the encrypted data can be stored in the

data service backend storage system. Examples of such use of the environment 1100 are provided below. In addition, example details of an example cryptography service are also provided below.

[0090] Data may be stored in the data service backend storage system in any suitable manner such as described above. For example, techniques for storing encrypted data in the backend storage system described above may be used in the environment 1100. For example, while not illustrated, the data service frontend may communicate with the cryptography service frontend to cause the cryptography service backend to encrypt data which may then be stored in the data service backend storage system. The encrypted data may be a data object and/or an encrypted key that was used to encrypt a data object. In the environment 1100, data may be placed into the data service backend storage system in other ways as well. For example, users may provide plaintext to be encrypted by the cryptography service and may receive ciphertext in response. The users may then interact or may submit a request to the data service frontend to request that the ciphertext be stored in the data service backend storage system. The data service frontend, in this example, may store the ciphertext in any manner. For instance, the data service frontend and backend storage systems may be configured to be indifferent to whether data is encrypted.

[0091] In addition, as with all environments illustrated herein, an additional frontend system may be logically located between the users and the data service frontend and the cryptography service frontend and possibly other frontend systems in order to coordinate actions between the systems. For example, in some embodiments, users may interact with a frontend system that itself interacts with the cryptography service frontend and data service frontend so that operations from the perspective of the user are simpler. For example, a user may request that a data object be encrypted and stored and a frontend system responds to the request by appropriate interactions with the cryptography service frontend and data service frontend. From the perspective of the user, however, such may be performed by a single request. Other variations are also within the scope of the present disclosure.

[0092] FIG. 12 shows an illustrative example of an environment 1200 which may be used to implement various embodiments of the present disclosure. In FIG. 12, the environment 1200 is configured to enable users to store ciphertext in a data service backend storage system. As illustrated in FIG. 12, accordingly, the environment 1200 includes a data service frontend, a data service backend storage system, an authentication service, a cryptography service frontend and a cryptography service backend. The data service backend storage system, the data service frontend, the authentication service, the cryptography service

frontend and the cryptography service backend may be systems such as described above in connection with FIG. 11. For example, as illustrated in FIG. 12, the data service frontend is configured to receive and respond to user requests and may also be configured to enforce policy on user requests. The data service frontend, as part of responding to requests, may be configured to submit authentication requests to the authentication service and receive authentication proof in response in response. Upon successful authentication, the data service frontend may be further configured to interact with the data service backend storage system to obtain encrypted data objects and possibly unencrypted data objects from the data service backend storage system, which may be then provided to a user.

5 [0093] As illustrated in FIG. 12, the cryptography service frontend is also configured to submit authentication requests to the authentication service and receive, in response, authentication proof. Authentication proof may be used to obtain services from the cryptography service backend. For example, the cryptography service frontend may be configured to provide ciphertext to the cryptography service backend with authentication proof and the cryptography service backend may be configured to decrypt the ciphertext and provide the ciphertext in return. As illustrated in FIG. 12, the ciphertext may be an encrypted key and the cryptography service backend may decrypt the encrypted key and provide the decrypted key, that is a plaintext key, to the cryptography service frontend, which is further configured to provide the plaintext key to the user. The user may then use the key to decrypt the encrypted data object received from the data service frontend or to decrypt encrypted data objects stored within the user's domain (e.g., within a user operated or controlled data center or computer system). In this example, the user may have obtained the encrypted key from the data service frontend. For example, the user may have submitted a request to the data service frontend for the data object and/or the key used to encrypt the data object. While illustrated in FIG. 11 as a single request, the separate requests may be made for both the data object and the key. As illustrated in FIG. 11, the data service frontend may obtain the encrypted data object and encrypted key from the data service backend storage system and provide the encrypted data object and encrypted key to the user.

[0094] It should be noted that, as with all environments illustrated herein, variations are considered as being within the scope of the present disclosure. For example, FIG. 12 shows a data object encrypted under a key and the key encrypted by another key identified by a key identifier being provided to the user. Further levels of encryption may also be used. For example, the data object may be encrypted under a key that is only accessible to the user (and/or that is not accessible by the other components of the environment 1200). A key used

to encrypt the data object may also be encrypted under the key that is only accessible to the user. In this example, unauthorized access to the components of the environment 1200 (absent the user) still does not provide access to the unencrypted contents of the data object since access to the user's key is still required for authorized decryption.

5 [0095] As another example, in the environment 1200 illustrated in FIG. 12, the data service frontend and the data service backend storage system do not have access to the plaintext data stored by the data service backend storage system because the data service frontend and the data service backend storage system do not have access to a key needed to decrypt the encrypted data. In some embodiments, however, access may be granted to the data service
10 frontend and/or the data service backend storage system. For instance, in an embodiment, temporary access to the key may be provided to the data service frontend to enable the data service frontend to obtain encrypted data, decrypt the encrypted data, use the decrypted data for a particular purpose (e.g., indexing), and then delete or otherwise lose access to the decrypted data. Such actions may be governed by policy enforced by the data service
15 frontend and/or cryptography service and may require authorization from the user.

[0096] FIG. 13 shows an illustrative example of a process 1300 which may be used to obtain an encrypted data object and an encrypted key such as from a data service backend storage system such as described above. The process 1300, for example, may be performed by the data service frontend system described above in connection with FIG. 12. In an
20 embodiment, the process 1300 includes receiving 1302 a GET request for an encrypted data object. Receiving the GET request may be performed in any suitable manner such as by receiving the request via an API call to the data service frontend system. As a result of having received the GET request, process 1300 may include submitting 1304 an authentication request and receiving 1306 an authentication response. Submitting 1304 the
25 authentication request and receiving 1306 the authentication response may be performed in any suitable manner such as described above. The authentication response may be used to determine 1308 whether the GET request is authentic. If it is determined 1308 that the GET request is not authentic, the process 1300 may include denying 1310 the GET request. Denying 1310 the GET request may be performed in any suitable manner such as described
30 above. If, however, it is determined 1308 that that GET request is authentic, the process 1300 may include providing 1312 the encrypted data object with an encrypted key that, when decrypted, is usable to decrypt the encrypted data object. It should be noted that, as with all processes described herein, numerous variations are considered as being within the scope of the present disclosure. For example, the process 1300 may be configured to respond to the

GET request, when authentic, by providing the encrypted data object but without providing an encrypted key. A requester, that is a user or system that submitted the GET request, may obtain the encrypted key in other ways. For example, in some embodiments, users may store encrypted keys themselves in a data storage system under the user's control. As another
5 example, one storage service may store the encrypted data object and another service may store the encrypted key and the user may obtain the encrypted data object and encrypted key from the respective services. As another example, another service or a third party to the user may be used to store encrypted keys and users may obtain encrypted keys upon request. Generally, any way by which encrypted keys may be provided may be used.

10 [0097] As illustrated in FIG. 13, the process 1300 may result in an entity having been provided a data object and an encrypted key usable to decrypt the data object. In various embodiments, the encrypted key must be decrypted in order to decrypt the data object. FIG. 14, accordingly, shows an illustrative example of a process 1400 which may be used to
15 provide a decrypted key to an entity in need of such a decrypted key in order to use the decrypted key for decryption of an encrypted data object. The process 1400 may be performed by any suitable system such as by the cryptography service frontend system described above in connection with FIG. 12. In an embodiment, the process 1400 includes receiving 1402 a decryption to decrypt a key using another key having a specified KeyID. While the process 1400 is described in connection with decryption of a key, it should be
20 noted that the process 1400 may be adapted for decryption of data in general. The decryption request may be received 1402 in any suitable manner such as described above (e.g., via an appropriately configured API call). Further, the decryption request may be received by any entity appropriate to the context in which the process 1400 is being performed. For instance, the decryption request may originate from the user or from another system, such as the data
25 service frontend discussed above. The decryption request may also include data (e.g. a key) to be decrypted or a reference thereto. The KeyID may be specified also in any suitable manner. For example, in some embodiments, the decryption request includes the KeyID or a reference to the KeyID, that is, information that can be used to determine the KeyID. As discussed above, the KeyID may also be implicitly specified. For example, the KeyID may
30 be obtained through association with available data such as an identity of a requester that submitted the decryption request. For instance, the key corresponding to the KeyID may be a default key for the requestor or for the entity on behalf of which the request was submitted.

[0098] The process 1400, in an embodiment, includes submitting 1404 an authentication request and receiving 1406 an authentication response. Submitting 1404 the authentication

request and receiving 1406 the authentication response may be performed in any suitable manner such as described above. Further, as described above, the received authentication response may be used to determine 1408 whether the GET request is authentic. If it is determined 1408 that the GET request is not authentic, the process 1400 may include denying 5 1410 the GET request. Denying 1410 the GET request may be performed in any suitable manner such as is described above. If, however, it is determined 1408 that the GET request is authentic, the process 1400 may include accessing policy information for the specified KeyID and/or for the requester. Policy information may include information including one or more policies on the KeyID and/or the requester.

10 [0099] In an embodiment, the accessed policy information is used to determine 1414 whether any applicable policy allows decryption of the key having the specified KeyID. If it is determined 1414 that the policy does not allow decryption of the key specified by the KeyID, the process 1400 may include denying 1410 the GET request such as described above. If, however, it is determined 1414 that policy allows decryption of the key having the 15 specified KeyID, the process 1400 may include decrypting 1416 the key using the key identified by the KeyID. Once the key has been decrypted, using a key having the KeyID, the decrypted key may then be provided 1418 such as by transmission over a network to the requester that submitted the decryption request (or, in some embodiments, another authorized destination).

20 [0100] As illustrated in the environment 1200 discussed above, a user may obtain encrypted data objects and keys for decrypting the data objects in various ways. FIG. 15 shows an illustrative example of a process 1500 which may be used to obtain plaintext in accordance with various embodiments. The process 1500 may be performed by any suitable system such as by a system being operated and/or hosted by a user such as described in 25 connection with FIG. 12. Other suitable systems include systems operating on behalf of users and not necessarily according to real time user input provided but perhaps according to preprogrammed processes.

[0101] In an embodiment, the process 1500 includes receiving 1502 ciphertext from a data storage service. Requesting 1502 ciphertext from a data storage service may be performed in 30 any suitable manner such as described above. For example, a system performing the process 1500 may request 1502 ciphertext, using an appropriately configured API call in the environment 1200 illustrated above in connection with FIG. 12 and/or by the process 1300 described above in connection with FIG. 13.

[0102] The process 1500 may also include receiving ciphertext and an encrypted key. Receiving ciphertext and encrypted key may be performed in any suitable manner. For example, the ciphertext and encrypted key may be received in a response to the request for the ciphertext from a data storage service. Generally, however, the ciphertext and encrypted key may be received 1504 in other suitable ways. For example, the request to receive ciphertext from the data storage service may be an asynchronous request and ciphertext may be received 1504 pursuant to another request that is subsequently submitted. Further, the ciphertext and encrypted key may be provided in a single response or may be obtained separately such as by different responses (which may be from the same or from different systems). As another example, a system performing a process 1500 may locally or otherwise store encrypted keys and the encrypted key may be received from local memory.

[0103] In an embodiment, the process 1500 includes requesting decryption of the encrypted key, using a key having a specified KeyID. The KeyID may be specified in any suitable manner such as described above. Further, it should be noted that the system performing the process 1500 may be able to specify the KeyID in any suitable manner. For example, the encrypted key and/or information provided therewith may specify the KeyID. As another example, the system performing the process 1500 may have local or remote access to information that enables determining the KeyID. A local or remote database, for instance, may associate data object identifiers with key identifiers for keys used to encrypt the data objects. Generally, any manner in which the system can be enabled to specify the KeyID may be used. Further, in some embodiments, the KeyID need not be specified, such as when information provided to a cryptography service is sufficient to determine the KeyID. The request 1506 for decryption of the encrypted key may be performed in any suitable manner such as in connection with an environment discussed above in connection with FIG. 12 and/or by performance of the process 1400 described above in connection with FIG. 14.

[0104] The process 1500, in an embodiment, includes receiving 1508 the decrypted key. Receiving 1508 the decrypted key may be performed in any suitable manner. For example, the decrypted key may be received in response to a request for decryption of the encrypted key. As another example, the request for decryption of the encrypted key may be an asynchronous request and another request may have been submitted for receiving the decrypted key. Generally, the decrypted key may be received in any suitable manner. Further, as with all information flowing from one device to another, the passage of information may be performed using secure channels. For instance, the decrypted key may be once again encrypted for decryption by an entity receiving the decrypted key. Generally,

any manner of secure communication may be used to pass information from one entity to another.

[0105] Once the decrypted key has been received 1508, the process 1500 may include using 1510 the decrypted key to decrypt 1510 ciphertext and therefore obtain plaintext. It should be noted that, as with all processes described herein, variations are considered as being within the scope of the present disclosure. For example, the process 1500 shows a request for ciphertext and a request for decryption of an encrypted key being performed sequentially. However, as with many operations described herein in connection with the various processes, operations need not be performed sequentially in various embodiments.

For example, if a system performing the process 1500 has access to the encrypted key prior to requesting the ciphertext, or is otherwise able to do so, the system may request ciphertext and request decryption of the encrypted key in parallel or in an order different from that which is illustrated. Other variations are also considered as being within the scope of the present disclosure.

[0106] As discussed above, various embodiments of the present disclosure are directed to providing cryptography services. Cryptography services may be provided by a cryptography service system such as described above. FIG. 16 accordingly shows an illustrative example of a cryptography service 1600 in accordance with various embodiments. As illustrated in FIG. 16 and as discussed above, the cryptography service 1600 is logically comprised of a frontend system and a backend system. Both the frontend system and the backend system may be implemented by one or more computer systems configured to perform operations described herein. For example, as illustrated in FIG. 16, the frontend system of the cryptography service 1600 implements a request API and a policy configuration API. The request API, in an embodiment, is an API configured for requesting cryptographic and other operations to be performed by the cryptography service. Thus, requests may be made to the frontend system via the request API in order for such cryptographic operations to be performed by the cryptography service.

[0107] The request API may be configured with the following example, high-level, requests available:

CreateKey(KeyID)

Encrypt(KeyID, Data, [AAD])

Decrypt(KeyID, Ciphertext, [AAD])

Shred(KeyID)

ReKey(Ciphertext, OldKeyID, NewKeyID).

[0108] A CreateKey(KeyID) request, in an embodiment, causes the cryptography service to create a key identified by the KeyID identified in the request. Upon receipt of a request, the
5 cryptography service may generate a key and associate the key with the KeyID. It should be known that KeyID's may be, but are not necessarily unique identifiers. For instance, a KeyID may identify a family of keys. For example, in some embodiments, key rotation is performed. Key rotation may involve replacing keys with other keys to prevent collection of enough decrypted data to allow practical cracking of a cipher used. If performed at the
10 direction of an entity different from the cryptography service, use of the CreateKey(KeyID) request may cause the cryptography service to create a new key to replace an old key identified by the KeyID. The old key may remain identified by the KeyID, but may, for instance, be only used for decryption (of data that has already been encrypted using the old key) and not for future encryption. As another example, in some embodiments, users of the
15 cryptography service provide their own key identifiers and there is a possibility that two different customers may provide the same identifier. In such instances, the identifier may not uniquely identify a key or even uniquely identify a family of keys. Various measures may be in place to address this. For example, an identity or other information associated with a user of the cryptography service may be used to identify the proper key or family of keys. In still
20 other embodiments the cryptographic service may assign a KeyID randomly, sequentially, or using any other method.

[0109] It should be noted that, when a KeyID does not uniquely identify a key, various systems may be in place to enable proper functionality. For example, in various
25 embodiments, a family of keys identified by a KeyID is finite. If a decryption operation using a key identified by a KeyID is requested, additional data (e.g., a time stamp of when the encryption was performed) may enable determining the proper key to use. In some embodiments, ciphertexts may include information indicating a key version. In some embodiments, all possible keys are used to provide different decryptions of the data. Since there are a finite number of keys, the proper decryption may be selected from those provided.
30 In some embodiments, decryption with a key is performed in a manner that enables the cryptographic service to detect that the ciphertext was not generated based at least in part on the key, such as by using authenticated encryption. Other variations are also considered as being within the scope of the present disclosure.

[0110] An Encrypt(KeyID, Data, [AAD]) request may be used to cause the cryptography service to encrypt the specified data using a key identified by the KeyID. Additional Authenticated Data (AAD) may be used for various purposes and may be data that is not necessarily encrypted, but that is authenticated, e.g. by an electronic signature, a message authentication code or, generally, a keyed hash value included with the AAD. In some 5 embodiments, the ciphertext is generated including at least a portion of the AAD. In some other embodiments the AAD is provided separately during decryption. In some other embodiments, the AAD is generated at decryption time based at least in part on the request and or other metadata such that decryption will only succeed when the metadata passes. In 10 some embodiments, policy may constrain whether a cryptographic operation can be performed with respect to particular AAD. Processing of Encrypt(KeyID, Data, [AAD]) requests may require, by programming logic and/or policy enforced by the cryptography service, both that the AAD contain particular values and that the AAD be authentic (e.g., not modified since original transmission). Similarly, a Decrypt(KeyID, Ciphertext, [AAD]) 15 request may be used to cause the cryptography service to decrypt the specified ciphertext using a key identified by the KeyID. The AAD in the Decrypt(KeyID, Ciphertext, [AAD]) request may be used such as described above. For instance, processing of the Decrypt(KeyID, Ciphertext, [AAD]) may require, by programming logic and/or policy enforced by the cryptography service, both that the AAD contain particular values and that 20 the AAD be authentic (e.g., not modified since original transmission).

[0111] The Shred(KeyID), in an embodiment, may be used to cause the cryptography service to electronically shred a key or family of keys identified by the specified KeyID. Electronic shredding may include making the key no longer accessible. For example, use of the Shred(KeyID) request may cause the cryptography system to command one or more 25 hardware devices to perform a SecureErase operation on one or more keys identified by the specified KeyID. Generally, the key(s) identified by the KeyID may be electronically shredded in any suitable manner, such as by overwriting data encoding the key with other data (e.g., a series of zeroes or ones or a random string). If the key(s) are stored encrypted under a key, the key used to encrypt the keys may be electronically shredded, thereby causing 30 a loss of access to the key(s). In some embodiments, the shred operation may cause decrypt operations indicating the shredded KeyID to fail at some determined point in the future. Other manners of securely and permanently destroying any possible access to the key(s) may be used.

[0112] The ReKey(Ciphertext, OldKeyID, NewKeyID) request, in an embodiment, may be used to cause the cryptography service to encrypt ciphertext under a different key. When the cryptography service receives a ReKey(Ciphertext, OldKeyID, NewKeyID) request, it may use a key identified by the OldKeyID to decrypt the specified ciphertext and then use a key identified by the NewKeyID to encrypt the decrypted ciphertext. If a key identified by the NewKeyID does not yet exist, the cryptography service may generate a key to use and associate the generated key with the specified NewKeyID, such as described in connection the Create(KeyID) request described above. In some embodiments, the ReKey operation may be operable to cause data to be transferrable between isolated instances of a cryptography service. In some embodiments, a policy might permit a rekey operation to be performed on a ciphertext but might not permit the same requestor to directly decrypt the ciphertext. In some embodiments, ReKey might support rekeying a ciphertext from a key identified by a first KeyID within a first account to a key identified by a KeyID within a second account.

[0113] Similarly, the frontend system may implement a policy configuration API which, in an embodiment, enables users to submit requests for configuring policies for the performance of cryptographic operations and for other policy-related operations. Policies may be associated with keys, groups of keys, accounts, users and other logical entities in various embodiments. Example policies, which may be configured via the policy configuration API, are provided below. In an embodiment, the cryptography service policy configuration API includes the following requests:

SetKeyPolicy(KeyID, Policy)

Suspend(KeyID, Public Key)

Reinstate(KeyID, Private Key)

[0114] In an embodiment, the SetKeyPolicy(KeyID, Policy) request may be used to cause the cryptography service to store a policy on the key (or family of keys) identified by the KeyID. A policy may be information that is determinative of whether a requested cryptographic operation can be performed in a particular context. The policy may be encoded in a declarative access control policy language, such as eXtensible Access Control Markup Language (XACML), Enterprise Privacy Authorization Language (EPAL), Amazon Web Services Access Policy Language, Microsoft SecPol or any suitable way of encoding one or more conditions that must be satisfied for a cryptographic operation to be performed. Policies may define what operations can be performed, when the operations can be

performed, which entities can make authorized requests for operations to be performed, which information is required for a particular request to be authorized, and the like. In addition, policies may be defined and/or enforced using access control lists, privileges associated with users, and/or operation bitmasks in addition to or instead of the examples
5 given above. Example policies appear below.

[0115] In some embodiments the cryptographic service may support a suspend operation, e.g., using a Suspend(KeyID, Public Key) API call. A suspend operation enables the customer of the cryptographic service to deny the operator of the cryptographic service use of
10 or access to a key. This can be useful to customers concerned about covert lawful orders or other circumstances in which the operator of the cryptographic service might be compelled to perform some operation using a key. It may also be useful to customers that wish to lock particular data and render it inaccessible online. In some embodiments, a suspend operation might include receiving a public key from a customer and encrypting the key specified by a
15 given KeyID with the received public key and shredding the key specified by the KeyID, such that the provider is not able to access the suspended key unless the private key associated with the public key is provided, e.g., using a Reinstate(KeyID, Private Key) API Call that both specifies the KeyID and includes the private key. In some other embodiments, a suspend operation might involve encrypting a key associated with a specified KeyID using
20 another key managed by the cryptographic service, including without limitation one created for the purpose of the instant suspend operation. The ciphertext produced by this operation can be provided to the customer and not retained within the cryptographic service. The original key identified by the KeyID can then be shredded. The cryptographic service may be operable to receive the provided ciphertext and re-import the suspended key. In some
25 embodiments the ciphertext may be generated in a manner that will prevent the cryptographic service from returning a decrypted version to the customer.

[0116] As illustrated in FIG. 16, the cryptography service 1600 includes a backend system that itself comprises various components in some embodiments. For example, the backend system in this example includes a request processing system which may be a subsystem of
30 the cryptography service 1600 that is configured to perform operations in accordance with requests received through either the request API or the policy configuration API. For example, the request processing component may receive requests received via the request API and the policy configuration API determines whether such requests are authentic and are therefore fulfillable and may fulfill the requests. Fulfilling the request may include, for example, performing and/or having performed cryptographic operations. The request

processing unit may be configured to interact with an authentication interface which enables the request processing unit to determine whether requests are authentic. The authentication interface may be configured to interact with an authentication system such as described above. For example, when a request is received by the request processing unit, the request processing unit may utilize the authentication interface to interact with an authentication service which may, if applicable, provide authentication proof that may be used in order to cause a performance of cryptographic operations.

[0117] The backend system of the cryptography service 1600 also, in this illustrative example, includes a plurality of a security modules (cryptography modules) and a policy enforcement module. One or more of the security modules may be hardware security modules although, in various embodiments, a security module may be any suitable computer device configured according to have capabilities described herein. Each security module in an embodiment stores a plurality of keys associated with KeyIDs. Each security module may be configured to securely store the keys so as to not be accessible by other components of the cryptography service 1600 and/or other components of other systems. In an embodiment, some or all of the security modules are compliant with at least one security standard. For example, in some embodiments, the security modules are each validated as compliant with a Federal Information Processing Standard (FIPS) outlined in FIPS Publication 140-1 and/or 140-2, such as one or more security levels outlined in FIPS Publication 140-2. In addition, in some embodiments, each security module is certified under the Cryptographic Module Validation Program (CMVP). A security module may be implemented as a hardware security module (HSM) or another security module having some or all capabilities of an HSM. In some embodiments, a validated module is used to bootstrap operations. In some embodiments, customers can configure some keys that are stored in and operated on only by validated modules and other keys that are operated on by software. In some embodiments, the performance or cost associated with these various options may differ.

[0118] The security modules may be configured to perform cryptographic operations in accordance with instructions provided by the request processing unit. For example, the request processing unit may provide ciphertext and a KeyID to an appropriate security module with instructions to the security module to use a key associated with the KeyID to decrypt the ciphertext and provide in response the plaintext. In an embodiment, the backend system of the cryptography service 1600 securely stores a plurality of keys forming a key space. Each of the security modules may store all keys in the key space; however, variations are considered as being within the scope of the present disclosure. For example, each of the

security modules may store a subspace of the key space. Subspaces of the key space stored by security modules may overlap so that the keys are redundantly stored throughout the security modules. In some embodiments, certain keys may be stored only in specified geographic regions. In some embodiments, certain keys may be accessible only to operators
5 having a particular certification or clearance level. In some embodiments certain keys may be stored in and used only with a module operated by a particular third party provider under contract with the provider of data storage services. In some embodiments, constructive control of security modules may require that lawful orders seeking to compel use of keys other than as authorized by the customer to involve either additional entities being compelled
10 or additional jurisdictions compelling action. In some embodiments, customers may be offered independent options for the jurisdiction in which their ciphertexts are stored and their keys are stored. In some embodiments, security modules storing keys may be configured to provide audit information to the owner of the keys, and the security modules may be configured such that the generation and providing of audit information not suppressible by
15 the customer. In some embodiments, the security modules may be configured to independently validate a signature generated by the customer such that the provider (e.g., hosting the security modules) is not able to perform operations under keys stored by the security modules. In addition, some security models may store all of the key space and some security modules may store subspaces of the key space. Other variations are also considered
20 as being the scope of the present disclosure. In instances where different security modules store different subspaces of the key space, the request processing unit may be configured such as with a relational table or other mechanism to determine which security module to instruct to perform cryptographic operations in accordance with various requests.

[0119] In an embodiment, the policy enforcement module is configured to obtain
25 information from a request processing unit and determine, based at least in part on that information, whether the request received through the API may be performed. For example, when a request to perform cryptographic operation is received through the request API, the request processing unit may interact with the policy enforcement module to determine whether fulfillment of the request is authorized according to any applicable policy such as
30 policy applicable to a specified KeyID in the request and/or other policies such as policy associated with the requestor. If the policy enforcement module allows fulfillment of the request, the request processing unit may, accordingly, instruct an appropriate security module to perform cryptographic operations in accordance with fulfilling the request.

[0120] As with all figures described herein, numerous variations are considered as being within the scope of the present disclosure. For example, FIG. 16 shows the policy enforcement module separate from security modules. However, each security module may include a policy enforcement module in addition to or instead of the policy enforcement module illustrated as separate. Thus, each security module may be independently configured to enforce policy. In addition, as another example, each security module may include a policy enforcement module which enforces policies different from policies enforced by a separate policy enforcement module. Numerous other variations are considered as being within the scope of the present disclosure.

[0121] As discussed above, various policies may be configured by users in or association with KeyID such that when requests specify cryptographic operations being performed in connection with keys corresponding to KeyIDs, policy may be enforced. FIG. 17 provides an illustrative example of the process 1700 for updating policy in accordance with various embodiments. Process 1700 may be performed by any suitable system, such as by a cryptography service system, such as described above in connection with FIG. 16. In an embodiment, the process 1300 includes receiving 1302 a request to update policy for a KeyID. The request may be received 1302 in any suitable manner. For example, referring to FIG. 16 as an example, a request may be received through a policy configuration API of the frontend system of the cryptography service 1600 described above. The request may be received in any suitable manner.

[0122] The process 1700 in an embodiment includes submitting 1704 an authentication request and receiving 1706 an authentication response. Submitting 1704 the authentication request and receiving 1706 the authentication response may be performed in any suitable manner such as described above. Also as described above, the received authentication response may be used to determine 1708 whether the request to update policy for the KeyID is authentic. If it is determined 1708 that the received request to update policy for the KeyID is not authentic, the request may be denied 1710. Denying 1710 the request may be performed in any suitable manner as described above. If, however, it is determined 1708 that the received request to update policy for the KeyID is authentic, the process 1700 may include accessing 1712 policy information applicable to the requestor. The policy information may be information from which any policy applicable to the requestor may be enforced. For example, within an organization that utilizes a cryptography service performed by process 1700, only certain users of the organization may be allowed to update policy for KeyIDs. Policy information may indicate which users are able to cause the cryptography

service to update policy for the KeyID and/or even whether the policy is updatable according to an existing policy. For example, a cryptography service may, in some embodiments, receive request to enforce a new policy. The cryptography service may check whether any existing policy allows the new policy to be put into place. If the cryptography service
5 determines that the existing policy does not allow enforcement of the new policy, the request may be denied. Generally, the policy information may be any information usable for the enforcement of policy applicable to the requestor.

[0123] As illustrated in FIG. 17, the process 1700 includes using the access policy information to determine 1704 whether policy allows the requested update to be performed.

10 If it is determined 1714 that the policy does not allow the requested update to be performed, the process 1700 may include denying 1710 the request such as described above. If, however, it is determined 1714 the policy allows the requested update to be performed, the process 1700 may include updating 1716 policy for the KeyID. Updating policy for the KeyID may include updating policy information and having the updated policy being stored
15 in accordance with or in association with the KeyID. The updated policy information may be stored, for example, by a policy enforcement module of the cryptography service such as described above, in connection with FIG. 16.

[0124] Policy may also be enforced by other components of an electronic environment that operate in connection with a cryptography service. For instance, referring FIG. 2, discussed
20 above, a cryptography service may provide an electronic representation of a policy to the data service frontend for the data service frontend to enforce. Such may be useful in circumstances where the data service is better suited to enforce the policy. For example, whether an action is allowed by policy may be based at least in part on information accessible to the data service frontend and not to the cryptography service. As one example, a policy
25 may depend on data stored by the data service backend storage system on behalf of a customer associated with the policy.

[0125] As discussed above, a cryptography service may include various systems that allow for the enforcement of policy in accordance with policy on keys having KeyIDs. FIG. 18, accordingly, shows an illustrated example of a process 1800 which may be used to enforce
30 policy. The process 1800 may be performed by any suitable system such as by a cryptography service system such as described above in connection with FIG. 16. In an embodiment, the process 1800 includes receiving 1802 a request to perform one or more cryptographic operations using a key having a KeyID. While FIG. 18 illustrates the process 1800 as being performed in connection with a request to perform one or more cryptographic

operations, it should be noted that the process 1800 may be adapted for use with any request to perform an operation which is not necessarily cryptographic. Example operations are described above.

[0126] A determination may be made 1804 whether the received request is authentic.

5 Determining whether the received request is authentic may be performed in any suitable manner such as described above. For instance, determining 1804 whether the request is authentic may include submitting an authentication request and receiving an authentication response such as described above. If it is determined 1804 that the request is not authentic, the process 1800 may include denying 1806 the request. Denying 1806 the request may be
10 performed in any suitable manner such as described above. If, however, it is determined 1804 that the request is authentic, the process 1800 may include accessing 1808 policy information for the KeyID and/or the requestor. Accessing policy information for the KeyID and/or request may be performed in any suitable manner. For example, accessing policy information for the KeyID and/or requestor may be performed by accessing the storage
15 policy information from one or more storage systems that store such policy information. The access policy information may be used to determine 1810 whether policy allows the one or more operations to be performed.

[0127] If it is determined 1810 that policy does not allow the one or more operations to be performed, the process 1800 may include denying 1806 the request. If, however, it is
20 determined that policy does allow the one or more operations to be performed, the process 1800 may include performing 1812 the requested one or more cryptographic operations. One or more results of performance of the one or more cryptographic operations may be provided 1814 such as provided to the requestor that submitted the received 1802 request to perform the one or more cryptographic operations. In some embodiments information derived at least
25 in part from allowed requests and or denied requests may be provided through an audit subsystem.

[0128] As discussed, embodiments of the present disclosure allow for flexible policy configuration and enforcement. In some embodiments, policies may state which services can perform which operations in which contexts. For example, a policy on a key may allow a
30 data storage service to cause the cryptography service to perform encryption operations but not decryption operations. A policy on a key may also include one or more conditions on the ciphertext and/or decrypted plaintext. For example, a policy may require that the ciphertext and/or plaintext produce a certain hash value (which may be a keyed hash value) before results of an operation be provided in response to a request. Policies may specify one or

more restrictions and/or permissions that are based at least in part on time, Internet Protocol (IP) from which requests originate, types of content to be encrypted/decrypted, AAD, and/or other information.

[0129] Numerous variations are considered as being within the scope of the present disclosure. For example, various embodiments discussed above discuss interaction with a separate authentication service. However, components of the environments discussed above may have their own authorization components and determining whether requests are authentic may or may not involve communication with another entity. Further, each of the environments discussed above are illustrated in connection with particular operations and capabilities enabled by the environments. The techniques discussed above in connection with different environments may be combined and, generally, environments in accordance with the present disclosure may allow for flexible use of the various techniques. As just one example, a cryptography service may be used to, upon request, encrypt both keys and other content, such as non-key data objects. As another example, a cryptography service may be configured to receive and respond to requests from both users (e.g., customers of a computing resource provider) and other services (e.g., data storage services). In some embodiments, a cryptography service and/or associated authentication service may be configured for use with a mobile device to perform encryption of stored data. In some embodiments at least one unlock pin may be validated by a cryptography service. In still other embodiments a cryptographic service, as part of an operation, may receive information generated by a hardware attestation. In some embodiments a cryptography service may be operable to provide digital rights management services with respect to content.

[0130] As discussed above, various embodiments of the present disclosure allow for rich policy enforcement and configurability. Many cryptosystems provide for authenticated encryption modes of operation where cryptographic operations can be performed to simultaneously provide confidentiality, integrity, and authenticity assurances on data. Confidentiality may be provided by encryption of plaintext data. Authenticity may be provided for both the plaintext and for associated data which may remain unencrypted. With such systems, changes to either the ciphertext or associated data may cause decryption of the ciphertext to fail.

[0131] In an embodiment, data associated with plaintext is used in the enforcement of policy. FIG. 19, accordingly, shows an illustrative example of a process 1900 for encrypting data in a manner that allows for policy enforcement using associated data in accordance with various embodiments. The process 1900 may be performed by any suitable system, such as a

cryptography service and/or a security module. As illustrated, the process 1900 includes obtaining 1902 plaintext. The plaintext may be obtained in any suitable manner. For example, in a service provider environment such as described above, a user (e.g., customer) may provide data to be encrypted. As another example, obtaining 1902 may include
5 generating a key (to be encrypted) and/or obtaining a key to be encrypted. The key may be used, such as described above.

[0132] As shown, the process 1900 includes obtaining associated data. The associated data may be any data that is associated with or is to be associated with the plaintext. The associated data may be any data on which one or more policies are, at least in part, based.

10 Examples appear below. Further, the associated data may be encoded in any suitable manner, such as in eXtensible Markup Language (XML), JavaScript Object Notation (JSON), Abstract Syntax Notation One (ASN1), YAML Ain't Markup Language (also referred to as Yet Another Markup Language) (YAML) or another structured extensible data format. In an embodiment, the process 1900 includes generating 1906, based at least in part on the
15 plaintext and the associated data, a message authentication code (MAC) and ciphertext. The combination of a MAC and ciphertext, such as the output of the AES-GCM cipher, may be referred to as authenticated ciphertext. Generating the MAC and ciphertext may be performed in any suitable manner and generation of the MAC and ciphertext may depend on which cryptosystem(s) is/are used. For example, in an embodiment, the advanced encryption
20 standard (AES) supports associated authenticated data (AAD) when operated in either CCM mode or GCM mode, where CCM represents Counter with CBC-MAC, GCM represents Galois/Counter Mode, and CBC represents a cipher block chaining. Using AES in either CCM or GCM mode, the plaintext and associated data may be provided as inputs to obtain
25 output of the concatenated pair of the ciphertext and a MAC of both the plaintext and associated data. It should be noted that while AES-CCM and AES-GCM are provided for the purpose of illustration, other authenticated encryption schemes may be used and the techniques described explicitly herein can be modified accordingly. For example, techniques of the present disclosure are applicable, in general, to symmetric block ciphers supporting
30 authenticated encryption modes. In addition, other encryption schemes are combinable with MAC functions in accordance with various embodiments of the present disclosure. Suitable encryption scheme and MAC function combinations include, but are not limited to, those where the encryption scheme is semantically secure under a chosen plaintext attack and the MAC function is not forgeable under chosen message attack. Further, while various embodiments of the present disclosure utilize ciphers that result in a single output that
35 encodes both the ciphertext and the MAC, the MAC and ciphertext may be generated using

different ciphers. Further, while MACs are used as illustrative examples, other values not generally referred to as MACs may also be used, such as general hashes, checksums, signatures, and/or other values that can be used in place of MACs. Accordingly, ciphers with automated encryption modes that support associated data include ciphers that use other cryptographic primitives in addition to or as an alternative to MACs.

[0133] Further, generating the MAC and ciphertext may be performed in various ways in accordance with various embodiments. For example, in an embodiment, the plaintext is provided to a security module, such as described above. The security module may be configured to generate the MAC. In other embodiments, a component of an electronic environment other than a security module generates the MAC and ciphertext. In such embodiments, a security module may be used to decrypt a key that, when in plaintext form, is used to generate the MAC and ciphertext. Once generated, the MAC and ciphertext (i.e., authenticated ciphertext) may be provided 1908. In some embodiments, the associated data is also provided. The MAC and ciphertext may be provided in various ways in various implementations making use of the process 1900 and variations thereof. For example, in some embodiments, the MAC and ciphertext are provided to a user, such as described above, or provided to a data service, such as described above, for processing by the data service. Further, as noted, while the associated data may be provided, in various embodiments, the associated data is not provided and/or it is generally persisted in plaintext form. As an example, the associated data may not be provided if it is independently obtainable. As an illustrative example, if the associated data is a persistent identifier of a device (e.g., an identifier of a storage device), the associated data may be obtained at a later time when needed for policy enforcement and/or for other purposes.

[0134] As discussed above, various embodiments of the present disclosure utilize security modules to provide enhanced data security. FIG. 20 provides an illustrative example of a process 2000 that may be used to encrypt data in a manner allowing for novel and rich policy enforcement, in accordance with various embodiments. The process 2000 may be performed by any suitable system, such as a cryptography service and/or a security module. As illustrated in FIG. 20, the process 2000 includes obtaining plaintext and associated data. As above, the plaintext and associated data may be received in a single communication, in separate communications and/or from separate entities. Once obtained, the plaintext, associated data, and a KeyID are provided 2004 to a security module. The security module may be such as described above. Further, the security module may be selected from multiple security modules participating in an electronic environment, such as an environment

supporting a cryptography service, such as described above. The KeyID may be as described above and may be specified in a request to encrypt the plaintext submitted to a cryptography service or may otherwise be specified. Further, in alternate embodiments of the process 2000, a KeyID may not be specified. For instance, in some embodiments, a security module
5 may select a KeyID and/or may generate a key that is later assigned a KeyID. In such embodiments, the process 2000 may be modified to provide the KeyID from the security module.

[0135] Returning to the illustrated embodiment, the process 2000 may include receiving
2006 from the security module ciphertext and a MAC. The ciphertext may be encrypted
10 under the key identified by the KeyID. The MAC may be a MAC over a combination of both the plaintext and associated data such that a change to the ciphertext or associated data will cause a check on the MAC to fail. As above, it should be noted that variations include those where the MAC is generated based at least in part on the associated data but independent of the plaintext. Further, as discussed above, the ciphertext and MAC may be provided together
15 (such as from output of use of an AES-CCM or AES-GCM cipher) or may be provided separately. Once received from the security module, the MAC and ciphertext are provided 2008 to an appropriate entity, such as a user of a cryptography service or a data service that operates in connection with a cryptography service, such as described above.

[0136] As discussed above, a security module may be used in various ways to enhance the
20 protection of data. As discussed above, in some embodiments, security modules are used to encrypt keys that are used (in their plaintext form) to encrypt other data. FIG. 21 accordingly shows an illustrative example of a process 2100 that can be used in such circumstances. The process 2100 may be performed by any suitable system, such as a cryptography service and/or a security module. The process 2100, in an embodiment, includes obtaining 2102
25 plaintext and associated data, such as described above. As illustrated, the process 2100 includes providing 2104, to a security module, an encrypted key, associated data, and KeyID identifying a key usable by the security module to decrypt the encrypted key. Accordingly, the process 2100 includes obtaining a decrypted key from the security module that used the
30 key identified by the KeyID to decrypt the encrypted key. Once obtained, the key can be used to encrypt the plaintext, thereby calculating 2108 ciphertext and a MAC. The ciphertext may be an encryption of the plaintext and the MAC may be over (i.e., based at least in part on) the associated data or both the associated data and plaintext, such as described above. Once encrypted, the process 2100 may include providing 2110 the MAC and ciphertext such as described above. Further, the process may also include losing 2112 access to the

decrypted key, which may be performed in any suitable manner, such as by a SecureErase operation, overwriting memory storing the decrypted key, removing power to volatile memory storing the key and/or in any other way in which a system performs the process 2100 (e.g., a cryptography system absent its security modules). While illustrated in parallel, providing the associated data, MAC, and/or ciphertext and losing access to the key may be performed in sequence, the order of which may vary among the various embodiments.

[0137] FIG. 22 shows an illustrative example of a process 2200 that may be used to enforce policy using associated data, in accordance with various embodiments. The process 2200 may be performed by any suitable system, such as a cryptography service and/or a security module. In an embodiment, the process 2200 includes receiving 2202 a request to perform an operation. The request may be any request submitted to a service that processes requests. In an embodiment, the request is a request to perform a cryptographic operation submitted to a cryptography service. In response to receiving 2202 the request, the process 2200 may include obtaining 2204 ciphertext, a MAC, and expected associated data. Obtaining 2204 the ciphertext, MAC, and expected associated data may be performed in any suitable manner. For example, in some embodiments one or more of the ciphertext, MAC, and expected associated data are received in the request. Two or more of the ciphertext, MAC and expected associated data may be received in separate requests or other communications and/or accessed from a data store, such as a local data store. For example, in an embodiment, the ciphertext and MAC are received as a concatenated pair (possibly generated from output of an AES-GCM or AES-CCM cipher) as part of the request. The expected associated data may also be part of the request or may be identified in other ways. For example, an identity of the requestor may be used, directly or indirectly, to determine the associated data. As a specific example, if the request is to perform an operation in connection with data stored in a storage device, obtaining 2204 the associated data may include obtaining an identifier of the data storage device. The identifier may be identified explicitly (e.g., as part of the request) or implicitly (e.g., because other information is available to determine that the data is stored in the data storage device). The associated data may be or may otherwise be based at least in part on the identifier of the data storage device. As discussed above, the associated data may vary greatly among the various embodiments.

[0138] In an embodiment, the process 2200 includes generating 2206 a reference MAC usable to determine the authenticity of the expected associated data. For example, the ciphertext, associated data, and an appropriate key (which may be identified in the request or may be otherwise determined) are used to generate 2206 a reference MAC. Generating the

MAC may be performed in any suitable manner, such as by using the same cipher that was used to obtain the ciphertext. A determination may be made 2208 whether the reference MAC and the obtained MAC match. For example, in many cryptosystems, MACs match when they are equal, although it is contemplated that other types of matching may be used in various embodiments. If it is determined 2208 that the reference MAC and obtained MAC match, in an embodiment, the process 2200 includes accessing 2210 policy information based at least in part on the associated data. Accessing 2210 the policy information may include accessing one or more policies (i.e., electronic representations of the one or more policies) from a remote or local data store based at least in part on the one or more policies associated with a KeyID used to generate the reference MAC and/or perform another cryptographic operation.

[0139] A determination may then be made 2212, based at least in part on the accessed policy information, whether policy allows the requested operation to be performed (e.g., whether policy allows the request to be fulfilled). Determining whether the policy allows the requested operation to be performed may include determining whether the ciphertext is tagged with associated data specified by the accessed policy information. Further, while not illustrated, policy information that is not based at least in part on associated data (e.g., policies based on information other than associated data) may also be used to determine whether policy allows the operation to be performed. If determined 2212 that policy allows the operation, the process 2200 may include performing 2214 the operation. If, however, it is determined 2212 that policy does not allow the operation, and/or if determined 2208 that the reference MAC and obtained MAC do not match, the process 2200 may include denying 2216 the request, such as described above.

[0140] Various policies are enforceable using the techniques described above. For example, as noted, policies may be associated with keys such that, when enforced, the policies determine what can and/or cannot be done with the keys. As one example, a policy may state that a data service can use a key only for certain types of operations specified by the policy (or, alternatively, that certain operations are prohibited to the data service). A policy may also specify conditions on use, time of use, IP address, what may be encrypted, what may be decrypted, and the like. As one illustrative example, one policy may specify that providing a decrypted result is allowed only if a hash of the decryption matches a specified value. Thus, a cryptography or other service enforcing the policy would not provide plaintext if a hash of the plaintext was not in accordance with the policy. As another example, a policy may specify that decryption of ciphertexts is allowed only if the ciphertexts

are tagged with associated data equal to or beginning with a specified value. As yet another example, a policy may specify that decryption of ciphertexts is allowed only if the ciphertexts are tagged with an identifier of a storage device encoded in associated data.

[0141] Generally, policies may specify restrictions and/privileges that are based at least in part on the value of data associated with a ciphertext (i.e., authenticated associated data). Some additional policies include policies that specify decryption is only allowed on ciphertext tagged with an identifier of a computer requesting decryption, ciphertext tagged with an identifier of a storage volume mounted (operationally connected) to a computer requesting decryption, and/or ciphertext tagged with identifiers of other computing resources. The computing resources may also be computing resources hosted by a computing resource provider that enforces the policies. Other policies are considered as being within the scope of the present disclosure, such as policies that are based at least in part on the input and/or output of a cryptographic algorithm before the output of the cryptographic algorithm is revealed to an entity outside of the entity performing the cryptographic algorithm (e.g., revealed to a user and/or other data service outside of a cryptography service enforcing the policies). As noted above, policies may also specify conditions, which may be based at least in part on associated data, for when policies may be modified.

[0142] FIG. 23 shows an illustrative example of a process 2300 which is a variation of the process 2200 discussed above in connection with FIG. 22, where the variation illustrates the use of a security module in enforcing policy, in accordance with various embodiments. In an embodiment, the process 2300 includes receiving 2302 a request to decrypt ciphertext, which may be an encrypted key or other encrypted data. The process 2300 also includes obtaining 2304 the ciphertext, MAC, and expected associated data, such as described above in connection with FIG. 22. As illustrated in FIG. 23, in an embodiment, the process 2300 includes using 2306 a security module to decrypt the ciphertext. Using 2306 a security module may also include selecting a security module from a plurality of security modules operable to decrypt the ciphertext, thereby producing plaintext. The security module may also be used 2308 to generate a reference MAC based at least in part on the plaintext and the expected associated data. It should be noted that, while shown as two separate steps in FIG. 23, using the security module to decrypt the ciphertext and generate the reference MAC may be performed in a single operation (e.g., a single request to the security module). Once obtained from the security module, the process 2300 includes determining 2310 whether the reference MAC and the obtained MAC match, such as described above in connection with FIG. 22. It should be noted, however, that in some embodiments, the process 2300 may be

modified so that the security module is provided the reference MAC and determines whether the reference MAC and the obtained MAC match. In this variation, the security module may provide a response indicating whether there is a match.

[0143] Returning to the embodiment illustrated in FIG. 23, if it is determined 2310 that the reference MAC and obtained MAC match, the process 2300 includes accessing 2312 policy information based at least in part on the associated data, such as described above in connection with FIG. 22. Also, as above, while not illustrated as such, additional policy information regarding policies not based at least in part on associated data may also be accessed. A determination may be made 2314 whether policy allows the operation. If determined 2314 that policy does allow the operation, the plaintext may be provided 2316. As above in connection with FIG. 22, if determined 2314 that policy does not allow the operation and/or if determined that the reference MAC does not match the obtained MAC, the process may include denying 2318 the request, such as described above.

[0144] While various embodiments of the present disclosure are illustrated using associated data of an authentication mode of a cipher, other embodiments are also considered as being within the scope of the present disclosure. For example, embodiments of the present disclosure generally apply to the use of data that is verifiable with ciphertext to enforce policy. As an illustrative example, a representation of a policy can be combined with a first plaintext to generate a new plaintext (e.g., a new plaintext comprising the plaintext and the policy). The new plaintext can be encrypted using a suitable cipher, such as AES, to produce a ciphertext. When a request to decrypt the ciphertext is received, a system receiving the request may decrypt the ciphertext, extract the policy from the new plaintext, and check whether the policy allows the first plaintext to be provided. If policy does not allow for the first plaintext to be provided, the request may be denied. Such embodiments may be used instead of or in addition to the embodiments discussed above in connection with associated data of an authentication mode of a cipher.

[0145] Various embodiments of the present disclosure also allow for policies on keys that specify conditions for how auditing takes place. For example, a policy on a key may specify an audit level for a key, where the audit level is a parameter for a cryptography service that is determinative of how the cryptography service audits key use. Auditing may be performed by any suitable system. For example, referring to FIG. 16, the request processing unit may communicate with an auditing system (not shown) that may be part of or separate from the cryptography service. When events occur in connection with the performance of cryptographic operations, relevant information may be provided to the auditing system that

logs the information. Events may be requests to perform cryptographic operations and/or information indicating whether the requested operations were performed. For example, if a user successfully requests the cryptography service to perform a decryption operation, the cryptography service may provide the auditing system information to enable the request and that the operation was performed. Administrative access events and, generally, any interaction or operation of the cryptography service may be logged with relevant information which may identify entities involved in the events, information describing the events, a timestamp of the events, and/or other information.

[0146] In an embodiment, audit levels include a high-durability level and a low-durability level. For a low-durability level, audit operations for a key may be performed by a cryptography service on a best-effort basis. With auditing according to a low-durability level, during normal operation, all operations are audited but, in the event of a failure of a component of the cryptography service, some audit data may be lost. With auditing according to a high-durability level, before revealing the result of a cryptographic operation, an assurance, that an audit record that the operation has occurred has been durably committed to memory, is obtained. Because of the acknowledgment needed, operations in high durability audit mode are slower than operations in low durability audit mode. The assurance that the audit record has been durably committed to memory may include an acknowledgment from one or more other systems used to store audit records. Thus, referring to the previous paragraph, a cryptography service may delay providing plaintext to a user until an acknowledgement from the auditing system that a record of the decryption resulting in the plaintext has been durably committed to memory. Durably committed to memory may mean that the data has been stored in accordance with one or more conditions for durability. For example, the data may be durably committed to memory when the data is written to non-volatile memory and/or the data has been redundantly stored among a plurality of data storage devices (e.g., using an erasure coding or other redundancy encoding scheme).

[0147] In an embodiment, a cryptography service uses sub-levels for the low-durability and high-durability audit levels. For example, in an embodiment, each level corresponds to two separate states, an immutable state and a mutable state. Whether the state is immutable or mutable may determine how and whether transitions between states are to occur. For example, using the illustrative example of audit durability, policy on a key may be able to change between low-durability mutable and high-durability mutable, from low-durability mutable to low-durability immutable, and from high-durability mutable to high-durability immutable. However, a cryptography service may be configured such that, once the policy

for the key is in either low-durability immutable or high-durability immutable, transition is prohibited. Thus, once the policy for a key is in an immutable state, the policy cannot be changed.

[0148] FIG. 24 shows an illustrative example of a state diagram for such a system, generalized to a policy that can be on (enforced) and off (unenforced). As illustrated in FIG. 24, the policy for a key can be on or off. When on and immutable, the policy can be changed to on and immutable (unchangeable) or to off and mutable (changeable). Similarly, when the policy is off but mutable, the policy can be changed to be on but mutable or to be off and immutable. It should be noted that other transitions may also be available, such as a direct transition from the policy being off but mutable to on and immutable. Further, not all transitions shown may be available. For example, the key may not have an off and immutable state in some instances.

[0149] FIG. 25 shows a generalized state diagram showing how a system may allow transitions among various policies applicable to a key. In this example, three policies, Policy A, Policy B and Policy C are shown. Each of these policies has a mutable and immutable state and transitions allowable between the states and the policies are shown. For example, transitions out of an immutable state are not allowed. However, a policy in a mutable state can be changed to another policy in a mutable state. For instance, the policy on the key may be changed from Policy A (mutable) to Policy B (mutable). As illustrated with Policy B, there may be transitions available to multiple policies. For example, from Policy B, the policy can be changed to either Policy C or Policy A. As with FIG. 24, other transitions and policies may be included and not all policies may have all states. Further, while various examples show a policy with an immutable and mutable state, policies may have more than two states, where each state corresponds to a set of actions that can or cannot be performed. For example, a semi-mutable state may allow some, but not all transitions that would be available under a mutable state.

[0150] As noted, policies can be used for various operations in addition to auditing. For example, the above restrictions on policy transitions may be applied to key shreddability. For instance, a policy may indicate whether a key may be shredded (irrevocably lost). The policy may have four states: shreddable-mutable; shreddable-immutable; unshreddable-mutable; and unshreddable-immutable. As above, when in an immutable state, the policy may not be changed. As another example, a policy on whether a key can be exported from a security module may also have such a quad-state policy.

[0151] Policies may also be associated with keys to prevent key use from enabling vulnerabilities to security attacks. For example, in an embodiment, one or more keys are associated with an automatic rotation policy that causes keys to be retired (e.g., marked so as to be no longer usable for encryption) after a certain amount of use. Such a policy may be a user-configurable (e.g., customer-configurable) policy that users (e.g., customers) activate and/or provide parameters for. The policy may also be a global policy applicable to a larger set of keys (such as a set comprising at least all keys managed by a cryptography service on behalf of its customers). In this manner, keys can be retired before they are used enough times to enable cryptographic attacks where knowledge of enough plaintexts and corresponding ciphertexts provide an ability to determine the key.

[0152] FIG. 26 shows an illustrative example of a process 2600 which may be used to rotate keys at appropriate intervals, in accordance with various embodiments. The process 2600 may be performed by any suitable device, such as a security module discussed above. In an embodiment, the process 2600 includes receiving 2602 a request to perform a cryptographic operation using a key identified by a KeyID. The request may be a request received from a cryptography service request processor, such as described above. The request may be a request to encrypt or decrypt data or, generally, to perform any cryptographic operation using the key identified by the KeyID, such as generation of an electronic signature, another key, or other information based at least in part on the key. Upon receipt 2602 of the request, the process 2600 includes performing 2604 the requested operation. Performing the requested operation may include additional operations, such as selecting an appropriate version of the key to perform the operation. For example, if the operation is encryption, a key marked as active may be used to encrypt. If the operation is decrypted, performing the operation may include selecting an appropriate version of the key identified by the KeyID to decrypt, which, in various embodiments, is the key originally used to encrypt the data. The key may be selected through various ways. For example, in some embodiments, the ciphertext may include metadata that identifies a version, serial number, date, or other information that enables selection of the key. In some embodiments, each possible key may be tried until data is properly decrypted, where proper decryption may be determined by a hash of plaintext output associated with the ciphertext or by the correctness of associated data.

[0153] Once the cryptographic operation has been performed 2604, the process 2600 includes updating 2606 a key use counter for an active key identified by the KeyID. For example, if the cryptographic operation results in a single use of the key, the counter may be

incremented by one. Similarly, if the cryptographic operation results in N (N a positive integer) uses of the key, the counter may be incremented by N. A determination may be made 2608 whether the counter exceeds a threshold. The threshold may be a number of uses allocated to a version of the key identified by the KeyID. The threshold may be provided by
5 a component of a cryptography service that manages allocation of operations for keys. The threshold may also be a default number of operations. If determined 2608 that the counter exceeds the threshold, in an embodiment, the process 2600 includes obtaining 2610 a new key. Obtaining the new key may be performed in any suitable manner. For example, if the process 2600 is performed by a security module, obtaining the new key may include
10 generating the new key or obtaining the new key from another security module, which may be orchestrated by an operator of the cryptography service. Passing keys from one security module to the other may be performed by encrypting the key with a key to which the providing and receiving security modules have access. The security module performing the process 2600 may receive and decrypt the encrypted key. Public-key key exchange
15 techniques may also be used.

[0154] Once a new key has been obtained, in an embodiment, the process 2600 may include marking 2612 the current active key as retired. Marking the current active key as retired may be performed in any suitable manner, such as by changing an appropriate value in a database maintained by the security module. Further, the process 2600 may include
20 associating 2614 the new key with the KeyID and marking the new key as active, such as by updating a database maintained by the security module. While not illustrated, the process 2600 may also include providing the new key for use by another security module. As indicated by the dashed line, at some point after the new key is ready for use by the security module (or other system) performing the process 2600, the process may include receiving
25 2602 a request to perform another cryptographic operation and the process 2600 may proceed as discussed above. Further, if it is determined 2608 that the counter does not exceed the threshold, the process 2600 may end and/or repeat once another request is received 2602.

[0155] FIG. 27 shows an illustrative example of a process 2700 that may be used to perform automatic rotation of keys in a cryptography service or other environment, in
30 accordance with various embodiments. The process 2700 may be performed by any suitable system, such as a component of a cryptography service that tracks key usage and orchestrates key rotation in accordance with the various embodiments. As illustrated in FIG. 27, the process 2700 includes allocating 2702 a number of key operations for keys (e.g., a number of operations for each of a plurality of keys) to one or more security modules. As a specific

example, in an environment utilizing five security modules to redundantly store/use a set of keys, each security module may be allocated one million operations for each of the keys it manages. Security modules (or other computer systems) to which operations are allocated may be hosted in the same data center or among multiple data centers. For example, in some
5 embodiments, a computing resource provider utilizes security modules in multiple data centers in multiple geographic areas to implement a geographically distributed cryptography or other service.

[0156] It should be noted, however, that allocations may be made for some, but not all keys, and the allocation for each key may not be equal. Allocating the key operations may
10 include providing notice of the allocation to each security module to which key operations have been allocated. The notice may specify the number that has been allocated for each key or, in some embodiments, the notice to a security module may indicate to the security module to reinitialize a counter that is pre-programmed into the security module or to add a
preprogrammed number to a counter. Upon allocating 2702 the key operations to the security
15 modules, a key use counter for each of the keys for which operations have been allocated may be updated. Continuing the above concrete example, if each of the five security modules has been allocated one million operations for a key identified by a particular KeyID, a counter for the KeyID may be incremented (upward or downward, depending on whether counting is performed upward or downward) by five million.

[0157] Upon having been allocated key operations, security modules may perform
20 cryptographic operations such as described above. Security modules may maintain their own counters that are based, at least in part, on the allocations that have been made. With the above example, if a security module has been allocated one million operations for a key identified by a particular KeyID, the security module may set the counter to one million (or
25 more than a million if an existing counter had remaining operations). Upon performing cryptographic operations using the key, the security module may increment its own counter accordingly.

[0158] At some point, an allocation exhaustion event for one or more KeyIDs may be detected 2706. An exhaustion event may be any event for which one or more security
30 modules loses or exhausts its allocation. As one example, a security module may use its allocation of operations for a key identified by a particular KeyID and detection of an exhaustion event may include receiving, from the security module, a notification that the security module has exhausted its allocation for the KeyID by performing a corresponding number of operations (or is within some predetermined threshold of exhausting its allocation

or is otherwise predicted to soon exhaust its allocation). As another example, in some embodiments, security modules are configured to lose access to keys stored therein (and data associated with the keys, such as counters) upon certain events, such as a malfunction, a detection of intrusion or other tampering, or when an operator needs access to the security module for maintenance. Accordingly, an exhaustion event may comprise a loss (possibly temporary) of a security module due to malfunction or detection of tampering/intrusion an intentional action (such as taking the security module out of commission temporarily for maintenance). In this example, a security module may be treated as if it used its allocation even though it did not necessarily perform the allocated number of operations. It should be noted, however, that all such events may not comprise exhaustion events in certain embodiments, such as when counters are persistently stored and, therefore, recoverable even upon loss to access to corresponding keys. It should also be noted that an exhaustion event may affect more than one security module. For example, a power outage affecting multiple security modules in a data center may result in an exhaustion event affecting multiple security modules.

[0159] Upon detection of the exhaustion event, a determination may be made 2710 whether a counter exceeds a threshold for any of the keys associated with the exhaustion event. The threshold may be a predetermined number of operations based at least in part on mathematical properties of a cipher used to perform cryptographic operations. For example, for ciphers with CBC modes, the number of operations may be or may otherwise be based at least in part on the size of the key space divided by the product of the squares of: (1) the length of the ciphertexts, expressed in blocks; and (2) the number of ciphertexts. For ciphers with CTR modes (such as AES-GCM), the number of operations may be or may otherwise be based at least in part on the size of the key space divided by the product of: (1) the square of the length of the ciphertext in blocks; and (2) the number of ciphertexts. If determined 2710 that the counter exceeds the threshold for any of the keys affected by the exhaustion event, the process 2700 may include instructing 2712 one or more security modules (i.e., the one or more security modules associated with the exhaustion event) to obtain one or more new keys for which the allocation of operations has been exhausted and replace the affected key(s) with the new key(s). For example, if a security module went offline temporarily (thereby causing the exhaustion event) and, as a result, caused a counter for a KeyID (but not necessarily all KeyIDs) to exceed the threshold, the security module may be instructed to obtain a new key, such as described above (e.g., by generating a new key, accessing a pre-generated key from data storage or obtaining a new key from another security module). It should be noted, however, that a different security module from the security module affected by the exhaustion

event may be instructed to obtain a new key, such as when one security module is brought offline and a new security module is brought online. Replacing an affected key (identified by a KeyID) with a new key may include marking the affected key as retired, associating the new key with the KeyID (e.g., in a database) and marking the new key as active. Replacing the affected key with the new key may also include initializing a counter (maintained by the security module replacing the affected key with the new key) for the new key, which may be a preprogrammed value or may be a value obtained from a system performing the process 2700. The process 2700 may also include marking 2714 the affected key(s) as retired and the new key(s) as active, such as by updating a database accordingly. It should be noted that a system performing the process 2700 may not have access to the affected and/or new key(s) and, as a result, marking the affected key(s) as retired and the new key(s) as active may include associating identifiers of the keys with values indicative of being retired or new, as appropriate.

[0160] In an embodiment, upon marking 2714 if it is determined 2710 that none of the affected keys have a counter exceeding the threshold, the process 2700 may include 2716 allocating additional key operations for the still active affected key(s) and/or any new key(s) that were obtained as a result of performing operations of the process 2700. Upon allocation of additional key operations, appropriate key use counter(s) may be updated 2704, such as described above.

[0161] As with all process described herein, variations are considered as being within the scope of the present disclosure. For example, in some embodiments, security modules do not track their use of a key, but another component of a cryptography or other service updates the counter for a key for each request submitted to any security module to perform one or more operations using the key. In such embodiments, for each key of a plurality of keys, the component of the security module may track requests to perform operations using the key (or track the operations performed, e.g., through acknowledgements or other indications that the requests were successfully fulfilled), updating a counter for the key accordingly. When the counter reaches the threshold, the system maintaining the counter may cause all appropriate security modules to retire the key and replace the key with a new key (or perform some other operation that causes the key to be retired, such as causing the security modules having the key to retire the key and causing one or more other security modules to start using a new key). As another example of a variation being within the scope of the present disclosure, in some embodiments, when a security module uses its allocation of key operations and causes the counter to exceed the threshold, the security module may be instructed to obtain a new

key, such as described above. Other security modules may continue using the key until they use their allocations. When a security module uses its allocation, it can obtain the new key that has replaced the key in the security module that caused the counter to exceed the threshold. In other words, security modules may be allowed to use up their allocations of key operations before having to obtain a new key.

[0162] FIG. 28 shows an illustrative example representation of a database used to maintain track of key usage. The database may be maintained by an appropriate system, such as a system performing the process 2700. In the database illustrated, columns correspond to, respectively, KeyIDs, key versions, usability, and a counter. The KeyIDs and Key Versions may be as described above. A value in the usability column may indicate whether the key is retired or active (or whether the key has another state, should such other states be supported by various implementations of the present disclosure). As illustrated in FIG. 28, the database has a row for each version of a key identified by a KeyID, including all retired versions and the active version. It should be noted, however, that a database may lack all versions of a key. For instance, keys may be permanently removed from storage for various security reasons. Removal may be, for instance, pursuant to a customer request or enforcement of a policy.

[0163] As illustrated in FIG. 28, the database illustrated also includes a counter for each active key. Also, in this particular example, the database includes a counter for inactive keys (e.g., showing a value for each key that exceeds a threshold, thereby causing a new key to be obtained). Counter values for inactive keys, however, may not be retained in some embodiments. A counter may be updated by a system that maintains the database as key operations are allocated to security modules. Once a value in the counter row exceeds a threshold value, a new row may be added to the database to accommodate a new key to replace the key whose counter value exceeded the threshold.

[0164] Security modules may maintain similar databases for their own purposes. For example, a security module may track its own use of keys and, when use of a key by the security module exhausts the number allocated to the security module, the security module may notify a key rotation management component of a cryptography service (or of another service using a security module) which may, for instance, perform a process such as the process 2700 described above to either reallocate additional operations to the security module or to cause the security module to obtain a new key if a number of key operations available for a key have been exhausted.

[0165] Databases used to track key use may vary from that which is illustrated in FIG. 28 and described above. For example, additional information may be included in a database, such as metadata associated with the key, such as a time of generation, a time of retirement, information about a customer for whom the key is used, and/or other information which may be useful in the various embodiments. Further, while a relational table is provided for the purpose of illustration, other ways of storing data in support of the various embodiments may be used.

[0166] Information used to enforce policy may be obtained by a cryptography service in various ways. FIG. 29 shows an illustrative example of an environment 2900 in which various embodiments of the present disclosure may be practiced. Components of the environment 2900 may include components such as those discussed above. For example, as illustrated, the environment 2900 includes a user that, through a user device, interacts with a data service frontend system (data service frontend), although in place of a user may be any computer system configured to interact with the data service frontend. Further, as illustrated, the environment includes a data service backend storage system and an authentication service, which may be as described above. In addition, the environment includes a cryptography service which, in FIG. 29, is labeled as an internal cryptography service to distinguish from another cryptography service, labeled in the figure as an external cryptography service. Either or both of the internal and external cryptography services may operate as a cryptography service, such as described above. As with all environments described herein, variations are considered within the scope of the present disclosure. For example, while FIG. 29 shows a user communicating with the data service frontend, the user may communicate with the cryptography service without the data service frontend, such as described above.

[0167] Further included in the environment 2900 is a message service which, in an embodiment, is a computer system (e.g., computer or network of computers) configured to at least, in response to a request from the internal cryptography service, publish messages to one or more subscribers, which may be represented by accounts. As discussed in more detail below, some requests may trigger components in the environment 2900 to operate so as to use the message service to transmit messages about the requests. For example, some data may be stored encrypted under a key that has an associated policy that responses to decryption requests must be delayed by a specified amount of time. If a request to decrypt the data is submitted, enforcement of policy may result in a delayed response and notifications to be sent by the message service. In this manner, a user or computer system

monitoring the messages may have, during the delay, the opportunity to submit a request to cancel the decryption request or otherwise cause the decryption request to be cancelled. The message service may transmit messages to one or more subscribers using one or more technologies including, but not limited to, electronic mail, short message service (SMS),
5 telephone, paging, facsimile, and/or any suitable technology usable for the transmission of information from the message service to another entity.

[0168] As illustrated in FIG. 29, requests submitted to the data service frontend may include a key access annotation. For example, as shown, a get request submitted by the user includes a key access annotation. A key access annotation may include information
10 necessary for the enforcement of policy by one or more components of the environment 2900. The key access annotation may include or otherwise be provided with an electronic signature generated using a key. The key used to generate the signature may be a key shared between the user and a component of the environment 2900 enforcing policy or, in some
15 embodiments that utilize a public key digital signature scheme, the key may be a private key of the user, thereby rendering the signature verifiable using a corresponding public key. A component of the environment 2900 enforcing policy may use the electronic signature to verify the authenticity of the information in the key access annotation and may use the information in the key access annotation to determine whether policy allows the request to be filled.

[0169] As illustrated in FIG. 29, any of several components of the environment 2900 may receive the key access annotation and use the key access annotation for enforcement in policy. For example, the data service frontend may pass the key access annotation to the internal cryptography service along with ciphertext and authorization proof, which may be used as described above. The internal cryptography service may use the key access
25 annotation to verify policy on a key managed by the internal cryptography service. In some embodiments, the key access annotation includes information identifying another service, such as the external cryptography service. The internal cryptography service may, upon detecting the information identifying the external cryptography service, provide the key access annotation to the external cryptography service, which may be one of multiple
30 cryptography or other services that uses the key access annotation, although such other services are not illustrated in FIG. 29.

[0170] As discussed above, keys managed in accordance with the various embodiments may be associated with policies that encode restrictions and/or privileges of the keys' use. Other information relevant to policy enforcement may also be associated with such keys, in

accordance with the various embodiments. FIG. 30 shows an illustrative example of a key 3000 (such as a key managed by a cryptography service, as described above) and example information that may be maintained (e.g., by a cryptography service) in association with the key 3000. As illustrated, policy information encoding one or more policies for the key 3000 is maintained in association with the key 3000, such as described above. In an embodiment, sets of keys are also maintained in association (e.g., associated with the key 3000 by a database or other data storage system) with the key 3000. The association may be maintained in any suitable manner. For example, a relational table or other mechanism may directly relate the key 3000 with keys in the sets of keys. As another example, a relational table or other mechanism may relate the key 3000 with identifiers of and/or references to the keys in the sets associated with the key 3000.

[0171] For instance, in some embodiments, one or more keys from a use key set and one or more keys from an administrative key set are maintained in association with the key 3000. In an embodiment, a cryptography service is configured such that keys in the use key set, in some embodiments, are necessary for at least some requests to be fulfilled. In particular, in some embodiments, for a request to perform an operation using a key identified by a KeyID, fulfillment of the request may require, possibly among other things, a key access annotation to be electronically (i.e., digitally) signed by a key corresponding to a key in the use key set (e.g., any key in the use key set or a key in the use key set specified by or otherwise in connection with the key access annotation). If such a signature is not present, a request may be denied. The key corresponding to a key in the use key set may be a key in the use key set or a reference thereto, or may be a private key corresponding to a public key in the use key set (where the public key and private key are used in a public key digital signature scheme).

[0172] In an embodiment, keys in the administrative key set are used for the purpose of modifying the sets of keys associated with the key 3000. For example, in some embodiments, for a request to modify either the use key set or the administrative key set (or both) to be fulfilled, an electronic signature using a key corresponding to a key in the administrative key set may be required. Modification may include adding a key to a key set, removing a key from a key set, replacing a key in a key set with another key, reassociating a key in the key set with another entity, and/or other modifications. In other words, proof of access to an administrative key is required to modify the use key set and/or the administrative key set. Administrative keys may also be used in a similar way for other purposes, such as for modifications of policy applicable to the key 3000 and/or other administrative actions. As illustrated in FIG. 30, the key 3000 may also be associated with an operation bitmask. The

information bitmask may be information that encodes a set of operations for which the key 3000 may be used, although such information may be encoded by policy.

[0173] Use of various key sets in association with a key managed by a security module provides numerous technical advantages. For example, referring to FIG. 30, if a key in a key set associated with the key 3000 is compromised, the key set can be modified to protect against data insecurity without having to change the key 3000. As another example, third-party keys can be used so that no single entity has access to both ciphertext and keys needed to decrypt ciphertext so that collusion between entities is required to access plaintext.

[0174] FIG. 31 shows an illustrative representation of an example key access annotation 3100, in accordance with an embodiment. In FIG. 30, the key access annotation 3100 includes information relevant to the enforcement of policy in connection with a request for which the key access annotation is included. In this example, the key access annotation includes a time stamp (TimeStamp), and operation descriptor (OperationDescriptor), key holder identifier (KeyHolderID), a KeyID, and an identifier of a key from a use key set or an administrative key set, although less or more information may be included.

[0175] In an embodiment, information in the key access annotation 3100 is used to determine whether a request can be fulfilled. In particular, in some embodiments, information in the key access annotation 3100 must satisfy one or more conditions (in addition to the requisite electronic signature being present) for the request to be fulfilled. For example, in some embodiments, the time stamp must indicate a time that is within some threshold amount of time of when the request is received. As another example, the key access annotation must be within an operation count window, must authorize a particular operation in connection with a particular piece of data, or must otherwise satisfy one or more conditions enforced by an entity that has received the key access annotation.

[0176] The KeyHolderID may enable a federated key management in accordance with various embodiments of the present disclosure. Referring to FIG. 29, for example, the KeyHolderID may be information that identifies or otherwise enables identification of a key specified by the KeyID. The KeyHolderID may be, for instance, a uniform resource locator (URL), Internet Protocol (IP) address, or other information that enables use of the KeyHolderID to provide the key access annotation (e.g., by forwarding a request including the annotation) to a corresponding entity. In some embodiments, the KeyHolderID is encoded in the KeyID so that a receiver of a request can extract the KeyHolderID and, if appropriate, forward the request to an appropriate third party. The internal cryptography

service of the environment 2900 may, for instance, be configured to detect whether the KeyHolderID of a request specifies a different entity and, upon detection of the KeyHolderID specifying a different entity, forward the request (or other information based at least in part on the request) to the different entity, which may be an external cryptography service, such as

5 described above (e.g., a cryptography service of a different computing resource provider). The different entity may utilize various embodiments of the present disclosure to determine whether to respond to the request and, if applicable, provide a response to the request, which may be provided through the cryptography service. If the key access annotation lacks a KeyHolderID (or the KeyHolderID specifies the internal cryptography service), the internal

10 cryptography service may fulfill the request itself, assuming compliance with policy and other prerequisites for request fulfillment.

[0177] FIG. 32 shows an illustrative example of a process 3200 which may be used to enforce policy in connection with requests submitted in connection with key access annotations, such as described above. The process 3200 may be performed by any suitable

15 device, such as a device of a cryptography service involved in request processing (e.g., a device implementing a frontend API for the cryptography service). The cryptography service may be an internal or external cryptography service, such as described above.

[0178] In an embodiment, the process 3200 includes receiving 3202 a key access annotation in connection with a request. The key access annotation may be an annotation

20 included with the request, or may be received in other ways. For example, a request with the key access annotation may be processed and, as part of processing, the key access annotation may be provided to a device performing the process 3200 (or a variation thereof). Once the key access annotation is received, in an embodiment, the process 3200 includes using 3204 the annotation to generate a reference signature, such as described above. A determination

25 may be made 3206 whether the reference signature matches the signature of the annotation. It should be noted that, in alternate embodiments, the process 3200 includes providing the key access annotation to a security module with information that enables the security module to generate the reference signature and check if the reference signature matches the signature of the key access annotation. In other words, determining whether the signature of the key

30 access annotation is valid may be performed by obtaining such a determination from a security module.

[0179] If determined 3206 that the reference signature matches the signature of the annotation, the process 3200 may include accessing 3208 any policy associated with a KeyID specified in the request. Assuming that policy exists for the KeyID, the process 3200

includes determining 3210 whether the accessed policy allows fulfillment of the request. If determined 3210 that policy allows fulfillment of the request, the process 3200 may include determining 3212 whether a KeyHolderID (such as described above) identifies an external entity. Determining 3212 whether a KeyHolderID identifies an external entity may be
5 performed in various ways in accordance with the various embodiments. For example, the absence of a KeyHolderID in the key access annotation may indicate a negative determination whereas the presence of a KeyHolderID in the key access annotation may indicate a positive determination. As another example, the KeyHolderID may be a required value that either indicates an internal entity (e.g., internal cryptography service) or external
10 entity (e.g., external cryptography service).

[0180] If it is determined 3212 that a KeyHolderID does not identify an external entity, the process 3200 may include fulfilling 3214 the request, such as described above. If, however, it is determined 3212 that a KeyHolderID does identify an external entity, the process 3200 may include forwarding 3216 the request to the identified external entity and receiving 3216,
15 in response, information enabling request fulfillment. The information may be, for example, plaintext obtained by decrypting ciphertext by the external entity and/or information based at least in part on a cryptographic operation performed by the external entity. While not illustrated, other information may also be received in addition or as an alternative. For example, if the external entity independently determines not to allow fulfillment of the
20 request (e.g., because policy enforced by the external entity does not allow the request to be fulfilled), information indicating a denial may be provided. Assuming that the information enabling request fulfillment is received 3216, the process may include fulfilling 3214 the request, such as described above. For instance, the received information enabling request fulfillment may be provided to the requestor.

[0181] Variations in addition to those explicitly described above are also considered as being within the scope of the present disclosure. For example, FIG. 32 illustrates, for a request, verifying an electronic signature and policy and subsequent forwarding of the request (or information based at least in part thereon) to an external entity. As with all processes illustrated herein, the order of operations may vary. For instance, the process 3200 may be
30 modified such that the request is forwarded before any signature and/or policy verification takes place. The external entity may perform signature and/or policy verification in addition to or instead of a system performing the process 3200. For example, in some embodiments, a system performing operations of the process 3200 does not perform signature and/or policy verification, but leaves such verification to the external entity. In other embodiments, a

system performing operations of the process 3200 performs signature and/or policy verification and the external entity performs additional signature and/or policy verification. The external entity may, for instance, have a different key and/or different policies in place that are checked before a request is fulfilled.

5 [0182] Embodiments of the present disclosure also utilize techniques that provide enhanced data security. For example, despite best efforts, security breaches often occur in connection with various systems. To mitigate the effects of such breaches, various embodiments of the present disclosure allow for request processing in a manner that provides time to detect breaches and respond accordingly. For example, in some embodiments, decryption of
10 encrypted data and/or access to data that has been decrypted requires a delay. The delay may be required, for instance, by a system's configuration or by a policy enforced by the system where the policy is configured in accordance with the various embodiments described herein). Such delays may be useful, for example, in circumstances where quick access to encrypted data is unnecessary, such as where regulatory or other requirements require that
15 certain data be archived, but not necessarily that the data be immediately available.

[0183] FIG. 33 shows an illustrative example of a process 3300 for providing such a delay in accordance with various embodiments. The process 3300 may be performed by a cryptography service, such as an internal or external cryptography service, such as described above. In an embodiment, the process 3300 includes receiving 3302 a request to decrypt
20 ciphertext using a key identified by a KeyID. The request may be received in any suitable manner. For example, the request may be received from a user or from a data service, such as described above. In an embodiment, policy associated with the KeyID is accessed 3304. It should be noted, however, that policy may be enforced automatically in some embodiments and, as a result, policy may not be accessed in some implementations. With respect to the
25 embodiment illustrated in FIG. 33, the process 3300 includes determining 3306 whether policy allows request fulfillment. The determination 3306 whether the policy allows request fulfillment may be performed in any suitable manner, such as described above.

[0184] If determined 3306 that the accessed policy allows fulfillment of the request, the process 3300 may include detection 3308 of a time delay policy among one or more of the
30 accessed policies. As noted above, however, in some embodiments, time delay may be automatic and, as a result, detection of such a policy may not occur. The process 3300 may include initiating 3310 a timer and performing an alarm process. The timer may be a timer for a determined amount of time, which, in some embodiments, is a configurable amount that may be, for instance, configured as part of a time delay policy. The timer may also be

predetermined into system configuration settings for a system performing the process 3300.

The alarm process may include execution of a workflow configured to cause the performance of various alarm actions. The alarm actions may include, for example, causing a notification system (such as described above) to transmit one or more messages to one or more

5 individuals or systems. For example, a compliance officer of an organization associated with the encrypted data may be notified. As another example, an alarm action may comprise causing an auditing system to perform enhanced auditing where, for example, more information is collected in connection with various accesses to the system. In some embodiments, alarm actions may also include alarm signaling which may include audible and/or visual indications of the alarm. Other alarm actions are considered as being within the scope of the present disclosure.

[0185] Upon initiation of the timer, the process 3300 may include performing checks 3312 whether the timer has expired until a determination is made 3312 that the timer has, in fact expired. In some embodiments, the request may be cancelled before its fulfillment.

15 Cancelling the request may require less stringent requirements than the request to decrypt the information. For example, in some embodiments, the request may be cancelled trivially by anybody with access to the system. As another example, an alarm action may include transmission of a message including a hyperlink that, when selected, causes a request to cancel the request to be cancelled, in some embodiments, without authentication by the receiver of the message. Other methods of enabling the request to be cancelled may also be used. Accordingly, in an embodiment, once the timer has expired (or at other times, such as before timer expiration) a determination may be made 3314 whether the request has been cancelled. If determined 3314 that the request has been cancelled, the process 3300 may include fulfilling 3316 the request, e.g., by using the key identified by the KeyID to decrypt the ciphertext and provide the resulting plaintext or by causing a security module to decrypt the ciphertext and provide the corresponding plaintext and then transferring the plaintext to the requestor or otherwise providing access to the plaintext. If, however, determined 3314 that the request has not been cancelled and/or if determined 3306 that the policy does not allow fulfillment of the request, the process 3300 may include denying the request, such as described above.

[0186] Numerous other variations are within the scope of the present disclosure. For example, in some embodiments keys are associated with policies that specify effective dates or otherwise indicate some determined period of validity that is not indefinite. For instance, in some embodiments, a key may be associated with a policy that specifies that the key

should be destroyed after a certain amount of time, on a certain date, or otherwise upon the satisfaction of one or more conditions. As an example, an organizational and/or governmental rule, law, policy or regulation may specify that data should be retained for a certain amount of time. A policy on a key used to encrypt such data (e.g., a key used in a cipher to encrypt the data or a key used in a cipher to encrypt another key under which the data is encrypted) may specify that the key should be destroyed (i.e., caused to be irrevocably inaccessible to anyone) upon passage of a required amount of time. A cryptography service or other system may detect passage of the required amount of time and destroy the key accordingly. The cryptography service may detect the passage of the required amount of time in many ways. For example, in some embodiments, the cryptography service maintains a database having stored therein information indicating times for key destruction. A periodic process executed by the cryptography service may destroy keys indicated in the database as destructible.

[0187] Policy effective dates may also be used in other ways. For example, time delays enforced for decryption of data may change as time passes. For instance, as data encrypted under a key ages, the time delay that is enforced for decrypting the data may decrease. As another example, principals may have privileges to have cryptographic operations performed using a key only for a specified period of time, which may be renewable through an API of the cryptography service. In this manner, the ability for individuals to cause cryptographic operations to be performed in connection with a key may expire automatically.

[0188] In some embodiments, requests may be made for other purposes in addition to or as an alternative to requests to decrypt ciphertext. For example, the various techniques described above may be adapted for various types of policy enforcement. For instance, policies utilized in connection with various embodiments of the present disclosure may utilize effective times for policies. A policy update may have an effective time that indicates a point in the future when the policy update become effective in a system enforcing the policy. An existing policy may require that some or all policy changes are allowable under particular conditions, one of which is that a policy update have an effective time that is a specified amount of time into the future (e.g., 48 hours). A request may be made to change a policy to allow for the retrieval of data, which may involve decryption of the data. The request may, for example, be a request to weaken conditions required to be satisfied before the data may be retrieved. The request may specify an effective time for the policy change. A system enforcing policy in connection with the principals to which the policy applies may approve or deny the request based at least in part on whether the effective time is in compliance with

existing policy. If the system approves the policy change, the system may initiate a timer (which may be based at least in part on the effective time) and perform an alarm process, such as described above. Further, policy change may be cancellable until the request is fulfilled, such as described above. In this manner, interested parties may be provided notice and opportunity to prevent unauthorized access to data.

[0189] Embodiments of the present disclosure can be described in view of the following clauses:

1. A computer-implemented method, comprising:

under the control of one or more computer systems configured with executable instructions, receiving, to a service of a computing resource provider, a request to decrypt ciphertext using a key specified by the request, the request satisfying a set of one or more conditions sufficient for the request to be fulfilled;

at a time during which the request is pending, transmitting one or more notifications of the request to one or more computer systems of a customer corresponding to the key;

enabling the request to be aborted during pendency of the request;

decrypting the ciphertext to obtain plaintext; and

providing the plaintext as a result of at least a predetermined amount of time having passed and the request not having been aborted.

2. The computer-implemented method of clause 1, wherein the predetermined time is configurable by the customer.

3. The computer-implemented method of clause 1 or 2, wherein transmitting the one or more notifications includes sending an electronic message to one or more individuals designated to receive notifications.

4. The computer-implemented method of any one of the preceding clauses, wherein satisfaction of one or more conditions sufficient for aborting the request are insufficient for the request to be fulfilled.

5. The computer-implemented method of any one of the preceding clauses, further comprising, as a result of the request, performing additional audit functions in connection with the ciphertext.

6. The computer-implemented method of any one of the preceding clauses, wherein: providing the plaintext includes enforcing a delay in providing the plaintext; and transmitting the one or more notifications and enforcing the delay is performed as a result of a configuration of a policy corresponding to the key.

5

7. A computer-implemented method, comprising:
under the control of one or more computer systems configured with executable instructions,

10 receiving an authenticated request to access plaintext, the fulfillment of which requiring one or more cryptographic operations;

processing the request such that a preprogrammed delay is required before a response to the request corresponding to fulfillment of the request is provided; and providing a response to the request after the delay.

15 8. The computer-implemented method of clause 7, wherein the one or more cryptographic operations include decryption of ciphertext.

9. The computer-implemented method of clause 7 or 8, wherein the one or more cryptographic operations include generating an electronic signature.

20

10. The computer-implemented method of any one of clauses 7 to 9, wherein the predetermined delay is required as a result of a policy corresponding to a key for decrypting ciphertext.

25 11. The computer-implemented method of any one of clauses 7 to 10, further comprising sending one or more notifications of the request to one or more entities with authority to abort the request.

30 12. The computer-implemented method of clause 11, further comprising enabling an ability to abort the request during the delay.

13. The computer-implemented method of any one of clauses 7 to 12, wherein the request being authenticated is sufficient to cause the response to be provided unless the request is aborted.

35

14. The computer-implemented method of any one of clauses 7 to 13, further comprising causing increased auditing as a result of the request.

15. A computer system, comprising:

5 one or more processors; and

memory including instructions that, when executed by the one or more processors, cause the computer system to:

10 detect a trigger in connection with an authenticated request, from a requestor, to access data, where accessing the data requires the performance of one or more cryptographic operations; and

as a result of detecting the trigger, cause a preprogrammed amount of time to pass before the data is accessible to the requestor, the request being abortable by an entity different from the requestor during passage of the amount of time.

15 16. The computer-system of clause 15, wherein:

the system further comprises a messaging subsystem; and

the instructions further cause the computer system to cause the messaging subsystem to transmit a notification of the request.

20 17. The computer-system of clause 15 or 16, wherein the trigger is a policy on a key required for the one or more cryptographic operations.

18. The computer-system of any one of clauses 15 to 17, wherein:

the computer system is hosted by a computing resource provider; and

25 the request is received by the computing resource provider from a customer of the computing resource provider.

19. The computer-system of clause 18, wherein the predetermined amount of time is configurable by the customer.

30

20. The computer-system of any one of clauses 15 to 19, wherein the instructions further cause the computer system to perform the one or more cryptographic operations.

21. One or more computer-readable media having collectively stored thereon instructions that, when executed by one or more processors of a computer system, cause the computer system to:

detect a trigger in connection with an authenticated request, from a requestor, to access data, where accessing the data requires the performance of one or more cryptographic operations; and

as a result of detecting the trigger, cause the request to be abortable by an entity different from the requestor, for an amount of time.

22. The one or more computer-readable storage media of clause 21, wherein the amount of time is a predetermined amount of time.

23. The one or more computer-readable storage media of clause 21 or 22, wherein the instructions further cause the computer system to cause transmission of one or more notifications of the request to another computer system.

24. The one or more computer-readable storage media of any one of clauses 21 to 23, wherein:

the one or more cryptographic operations use a key; and

the trigger is a policy for the key requiring the request to be cancellable for the amount of time.

25. The one or more computer-readable storage media of any one of clauses 21 to 24, wherein:

validity of the request requires satisfaction of one or more conditions; and

cancellation of the request is possible with at least one of the one or more conditions unsatisfied.

26. The one or more computer-readable storage media of any one of clauses 21 to 25, wherein the instructions further cause the system to notify one or more individuals specified for receiving notifications of requests to access the data.

27. The one or more computer-readable storage media of any one of clauses 21 to 26, wherein the request to access the data is a request to decrypt ciphertext using a key specified by the request.

28. A system, comprising:
one or more processors; and
memory including instructions that, when executed by the one or more processors,
5 cause the system to:

receive a request for a policy to be effective at a time indicated by the request;
determine, based at least in part on the indicated time, whether fulfillment of the request
complies with a currently effective policy;

10 at a time between receipt of the request and the indicated time, enable another
request to be received to abort the request.

29. The system of clause 28, wherein the request is to change one or more conditions
required to be satisfied for access to data.

15 30. The system of clause 28 or 29, wherein the policy corresponds to a key managed by
a cryptography service and indicates one or more conditions required to be satisfied for the
cryptography service to perform one or more cryptographic operations.

20 31. The system of any one of clauses 28 to 30, wherein the instructions further cause the
system to, as a result of receiving the request, transmit one or more notifications to one or
more entities each with authority to abort the request.

[0190] Other variations are also considered as being within the scope of the present
disclosure. For instance, in some embodiments, a data storage system may have an inherent
25 latency for request processing that is sufficient for providing time to cancel requests and/or
otherwise protect against potential security breaches. For instance, an archival data storage
system may utilize asynchronous request processing where data responsive to a request may
be available hours after the request is submitted. Further, data responsive to requests may be
available after different amounts of time, which may depend on factors such as system. With
30 such systems, a timer may be unnecessary due to the inherent latency.

[0191] FIG. 34 illustrates aspects of an example environment 3400 for implementing
aspects in accordance with various embodiments. As will be appreciated, although a Web-
based environment is used for purposes of explanation, different environments may be used,
as appropriate, to implement various embodiments. The environment includes an electronic
35 client device 3402, which can include any appropriate device operable to send and receive

requests, messages or information over an appropriate network 3404 and convey information back to a user of the device. Examples of such client devices include personal computers, cell phones, handheld messaging devices, laptop computers, set-top boxes, personal data assistants, electronic book readers and the like. The network can include any appropriate network, including an intranet, the Internet, a cellular network, a local area network or any other such network or combination thereof. Components used for such a system can depend at least in part upon the type of network and/or environment selected. Protocols and components for communicating via such a network are well known and will not be discussed herein in detail. Communication over the network can be enabled by wired or wireless connections and combinations thereof. In this example, the network includes the Internet, as the environment includes a Web server 3406 for receiving requests and serving content in response thereto, although for other networks an alternative device serving a similar purpose could be used as would be apparent to one of ordinary skill in the art.

[0192] The illustrative environment includes at least one application server 3408 and a data store 3410. It should be understood that there can be several application servers, layers, or other elements, processes or components, which may be chained or otherwise configured, which can interact to perform tasks such as obtaining data from an appropriate data store. As used herein the term "data store" refers to any device or combination of devices capable of storing, accessing and retrieving data, which may include any combination and number of data servers, databases, data storage devices and data storage media, in any standard, distributed or clustered environment. The application server can include any appropriate hardware and software for integrating with the data store as needed to execute aspects of one or more applications for the client device, handling a majority of the data access and business logic for an application. The application server provides access control services in cooperation with the data store, and is able to generate content such as text, graphics, audio and/or video to be transferred to the user, which may be served to the user by the Web server in the form of HyperText Markup Language ("HTML"), Extensible Markup Language ("XML") or another appropriate structured language in this example. The handling of all requests and responses, as well as the delivery of content between the client device 3402 and the application server 3408, can be handled by the Web server. It should be understood that the Web and application servers are not required and are merely example components, as structured code discussed herein can be executed on any appropriate device or host machine as discussed elsewhere herein.

[0193] The data store 3410 can include several separate data tables, databases or other data storage mechanisms and media for storing data relating to a particular aspect. For example, the data store illustrated includes mechanisms for storing production data 3412 and user information 3416, which can be used to serve content for the production side. The data store also is shown to include a mechanism for storing log data 3414, which can be used for reporting, analysis or other such purposes. It should be understood that there can be many other aspects that may need to be stored in the data store, such as for page image information and to access right information, which can be stored in any of the above listed mechanisms as appropriate or in additional mechanisms in the data store 3410. The data store 3410 is operable, through logic associated therewith, to receive instructions from the application server 3408 and obtain, update or otherwise process data in response thereto. In one example, a user might submit a search request for a certain type of item. In this case, the data store might access the user information to verify the identity of the user, and can access the catalog detail information to obtain information about items of that type. The information then can be returned to the user, such as in a results listing on a Web page that the user is able to view via a browser on the user device 3402. Information for a particular item of interest can be viewed in a dedicated page or window of the browser.

[0194] Each server typically will include an operating system that provides executable program instructions for the general administration and operation of that server, and typically will include a computer-readable storage medium (e.g., a hard disk, random access memory, read only memory, etc.) storing instructions that, when executed by a processor of the server, allow the server to perform its intended functions. Suitable implementations for the operating system and general functionality of the servers are known or commercially available, and are readily implemented by persons having ordinary skill in the art, particularly in light of the disclosure herein. In some embodiments, an operating system may be configured in accordance with or validated under one or more validation regimes such as Evaluation Assurance Level (EAL) level 4.

[0195] The environment in one embodiment is a distributed computing environment utilizing several computer systems and components that are interconnected via communication links, using one or more computer networks or direct connections. However, it will be appreciated by those of ordinary skill in the art that such a system could operate equally well in a system having fewer or a greater number of components than are illustrated in FIG. 34. Thus, the depiction of the system 3400 in FIG. 34 should be taken as being illustrative in nature, and not limiting to the scope of the disclosure.

[0196] The various embodiments further can be implemented in a wide variety of operating environments, which in some cases can include one or more user computers, computing devices or processing devices which can be used to operate any of a number of applications. User or client devices can include any of a number of general purpose personal computers, such as desktop or laptop computers running a standard operating system, as well as cellular, wireless and handheld devices running mobile software and capable of supporting a number of networking and messaging protocols. Such a system also can include a number of workstations running any of a variety of commercially-available operating systems and other known applications for purposes such as development and database management. These devices also can include other electronic devices, such as dummy terminals, thin-clients, gaming systems and other devices capable of communicating via a network.

[0197] Most embodiments utilize at least one network that would be familiar to those skilled in the art for supporting communications using any of a variety of commercially-available models and protocols, such as Transmission Control Protocol/Internet Protocol (“TCP/IP”), Open System Interconnection (“OSI”), File Transfer Protocol (“FTP”), Universal Plug and Play (“UpnP”), Network File System (“NFS”), Common Internet File System (“CIFS”) and AppleTalk. The network can be, for example, a local area network, a wide-area network, a virtual private network, the Internet, an intranet, an extranet, a public switched telephone network, an infrared network, a wireless network and any combination thereof.

[0198] In embodiments utilizing a Web server, the Web server can run any of a variety of server or mid-tier applications, including Hypertext Transfer Protocol (“HTTP”) servers, FTP servers, Common Gateway Interface (“CGI”) servers, data servers, Java servers and business application servers. The server(s) also may be capable of executing programs or scripts in response requests from user devices, such as by executing one or more Web applications that may be implemented as one or more scripts or programs written in any programming language, such as Java[®], C, C# or C++, or any scripting language, such as Perl, Python or TCL, as well as combinations thereof. The server(s) may also include database servers, including without limitation those commercially available from Oracle[®], Microsoft[®], Sybase[®] and IBM[®].

[0199] The environment can include a variety of data stores and other memory and storage media as discussed above. These can reside in a variety of locations, such as on a storage medium local to (and/or resident in) one or more of the computers or remote from any or all of the computers across the network. In a particular set of embodiments, the information may reside in a storage-area network (“SAN”) familiar to those skilled in the art. Similarly, any

necessary files for performing the functions attributed to the computers, servers or other network devices may be stored locally and/or remotely, as appropriate. Where a system includes computerized devices, each such device can include hardware elements that may be electrically coupled via a bus, the elements including, for example, at least one central
5 processing unit (“CPU”), at least one input device (*e.g.*, a mouse, keyboard, controller, touch screen or keypad), and at least one output device (*e.g.*, a display device, printer or speaker). Such a system may also include one or more storage devices, such as disk drives, optical storage devices, and solid-state storage devices such as random access memory (“RAM”) or read-only memory (“ROM”), as well as removable media devices, memory cards, flash cards,
10 etc. Various embodiments of the present disclosure may also be implemented using custom hardware including, but not limited, custom cryptographic processors, smart cards and/or hardware security modules.

[0200] Such devices also can include a computer-readable storage media reader, a communications device (*e.g.*, a modem, a network card (wireless or wired), an infrared
15 communication device, etc.) and working memory as described above. The computer-readable storage media reader can be connected with, or configured to receive, a computer-readable storage medium, representing remote, local, fixed and/or removable storage devices as well as storage media for temporarily and/or more permanently containing, storing, transmitting and retrieving computer-readable information. The system and various devices
20 also typically will include a number of software applications, modules, services or other elements located within at least one working memory device, including an operating system and application programs, such as a client application or Web browser. It should be appreciated that alternate embodiments may have numerous variations from that described above. For example, customized hardware might also be used and/or particular elements
25 might be implemented in hardware, software (including portable software, such as applets) or both. Further, connection to other computing devices such as network input/output devices may be employed.

[0201] Storage media and computer readable media for containing code, or portions of code, can include any appropriate media known or used in the art, including storage media
30 and communication media, such as but not limited to volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or transmission of information such as computer readable instructions, data structures, program modules or other data, including RAM, ROM, Electrically Erasable Programmable Read-Only Memory (“EEPROM”), flash memory or other memory technology, Compact Disc

Read-Only Memory (“CD-ROM”), digital versatile disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices or any other medium which can be used to store the desired information and which can be accessed by the a system device. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

[0202] The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims.

[0203] Other variations are within the spirit of the present disclosure. Thus, while the disclosed techniques are susceptible to various modifications and alternative constructions, certain illustrated embodiments thereof are shown in the drawings and have been described above in detail. It should be understood, however, that there is no intention to limit the invention to the specific form or forms disclosed, but on the contrary, the intention is to cover all modifications, alternative constructions and equivalents falling within the spirit and scope of the invention, as defined in the appended claims.

[0204] The use of the terms “a” and “an” and “the” and similar referents in the context of describing the disclosed embodiments (especially in the context of the following claims) are to be construed to cover both the singular and the plural, unless otherwise indicated herein or clearly contradicted by context. The terms “comprising,” “having,” “including,” and “containing” are to be construed as open-ended terms (i.e., meaning “including, but not limited to,”) unless otherwise noted. The term “connected” is to be construed as partly or wholly contained within, attached to, or joined together, even if there is something intervening. Recitation of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within the range, unless otherwise indicated herein, and each separate value is incorporated into the specification as if it were individually recited herein. All methods described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. The use of any and all examples, or exemplary language (e.g., “such as”) provided herein, is intended merely to better illuminate embodiments of the invention and does not pose a limitation on the scope of the invention unless otherwise claimed. No language in the specification should be construed as indicating any non-claimed element as essential to the practice of the invention.

[0205] Preferred embodiments of this disclosure are described herein, including the best mode known to the inventors for carrying out the invention. Variations of those preferred embodiments may become apparent to those of ordinary skill in the art upon reading the foregoing description. The inventors expect skilled artisans to employ such variations as appropriate, and the inventors intend for the invention to be practiced otherwise than as specifically described herein. Accordingly, this invention includes all modifications and equivalents of the subject matter recited in the claims appended hereto as permitted by applicable law. Moreover, any combination of the above-described elements in all possible variations thereof is encompassed by the invention unless otherwise indicated herein or otherwise clearly contradicted by context.

[0206] All references, including publications, patent applications and patents, cited herein are hereby incorporated by reference to the same extent as if each reference were individually and specifically indicated to be incorporated by reference and were set forth in its entirety herein.

CLAIMS

WHAT IS CLAIMED IS:

- 1 1. A computer-implemented method , comprising:
2 under the control of one or more computer systems configured with executable
3 instructions,
4 receiving an authenticated request to access plaintext, the fulfillment of which
5 requiring one or more cryptographic operations;
6 processing the request such that a preprogrammed delay is required before a
7 response to the request corresponding to fulfillment of the request is provided; and
8 providing a response to the request after the delay.
- 1 2. The computer-implemented method of claim 1, wherein the one or
2 more cryptographic operations include decryption of ciphertext.
- 1 3. The computer-implemented method of claim 1 or 2, wherein the one or
2 more cryptographic operations include generating an electronic signature.
- 1 4. The computer-implemented method of any one of the preceding
2 claims, wherein the predetermined delay is required as a result of a policy corresponding to a
3 key for decrypting ciphertext.
- 1 5. The computer-implemented method of any one of the preceding
2 claims, further comprising sending one or more notifications of the request to one or more
3 entities with authority to abort the request.
- 1 6. The computer-implemented method of claim 5, further comprising
2 enabling an ability to abort the request during the delay.
- 1 7. The computer-implemented method of any one of the preceding
2 claims, wherein the request being authenticated is sufficient to cause the response to be
3 provided unless the request is aborted.
- 1 8. The computer-implemented method of any one of the preceding
2 claims, further comprising causing increased auditing as a result of the request.

1 9. A computer system, comprising:
2 one or more processors; and
3 memory including instructions that, when executed by the one or more
4 processors, cause the computer system to:
5 detect a trigger in connection with an authenticated request, from a
6 requestor, to access data, where accessing the data requires the performance of one or more
7 cryptographic operations; and
8 as a result of detecting the trigger, cause a preprogrammed amount of
9 time to pass before the data is accessible to the requestor, the request being abortable by an
10 entity different from the requestor during passage of the amount of time.

1 10. The computer-system of claim 9, wherein:
2 the system further comprises a messaging subsystem; and
3 the instructions further cause the computer system to cause the messaging
4 subsystem to transmit a notification of the request.

1 11. The computer-system of claim 9 or 10, wherein the trigger is a policy
2 on a key required for the one or more cryptographic operations.

1 12. The computer-system of any one of claims 9 to 11, wherein:
2 the computer system is hosted by a computing resource provider; and
3 the request is received by the computing resource provider from a customer of
4 the computing resource provider.

1 13. The computer-system of claim 12, wherein the predetermined amount
2 of time is configurable by the customer.

1 14. The computer-system of any one of claims 9 to 13, wherein the
2 instructions further cause the computer system to perform the one or more cryptographic
3 operations.

1 15. The computer system of any one of claims 9 to 14, wherein:
2 the one or more cryptographic operations use a key; and
3 the trigger is a policy for the key requiring the request to be cancellable for the
4 amount of time.

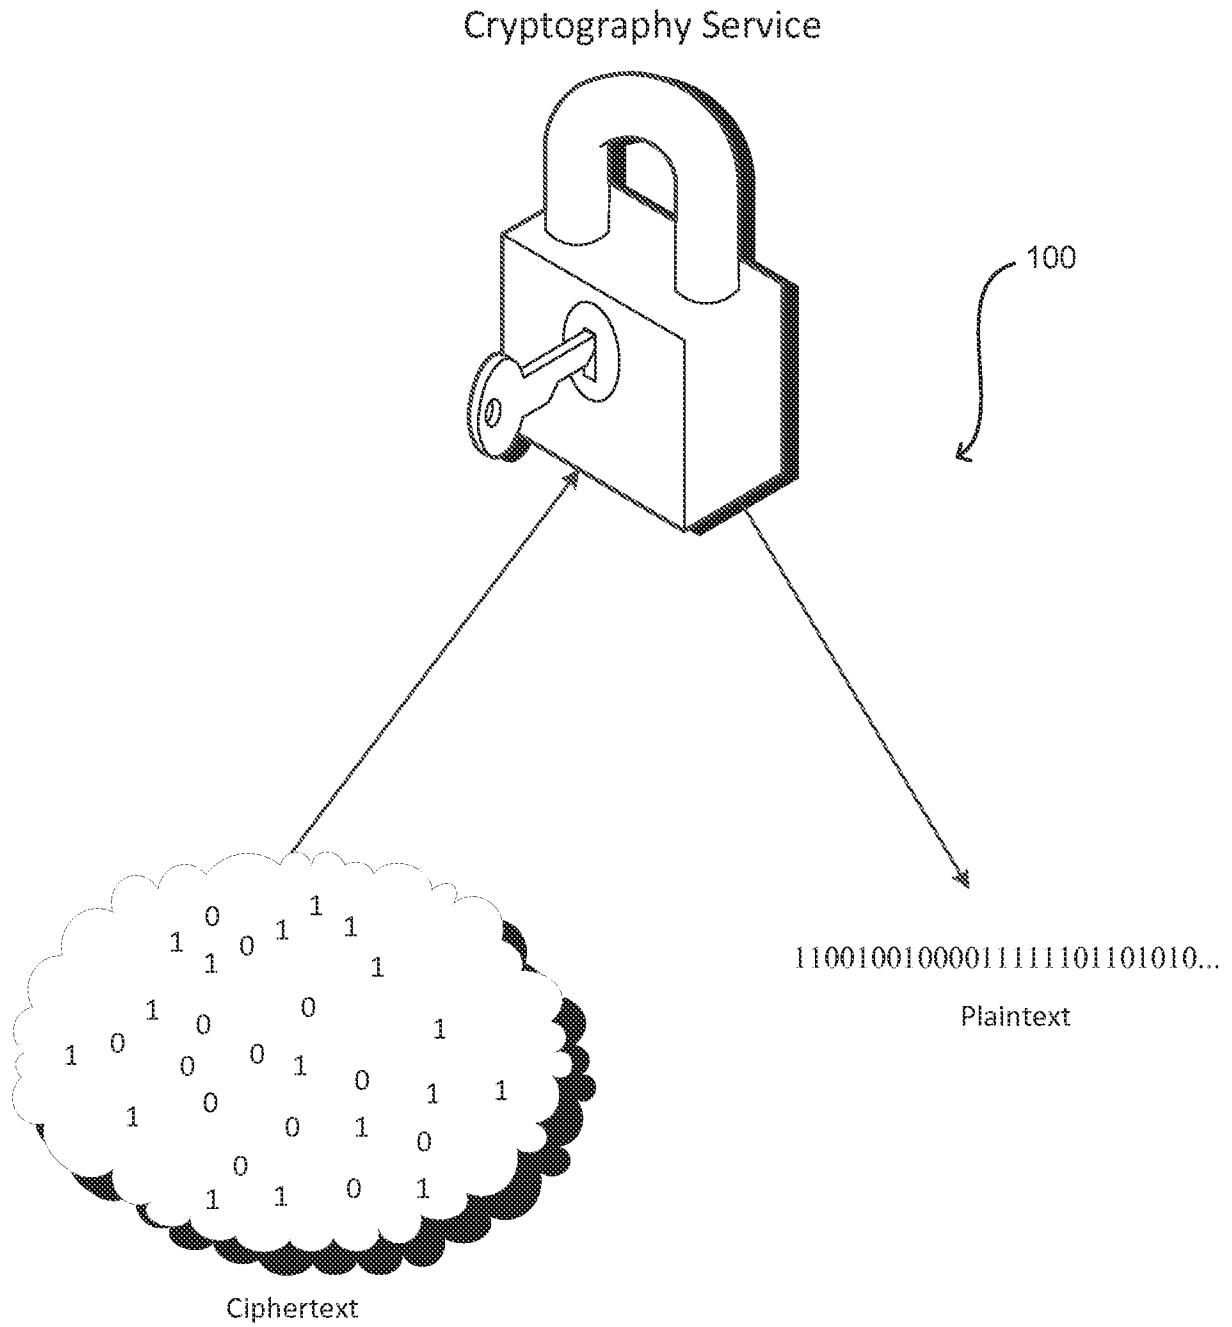


FIG. 1

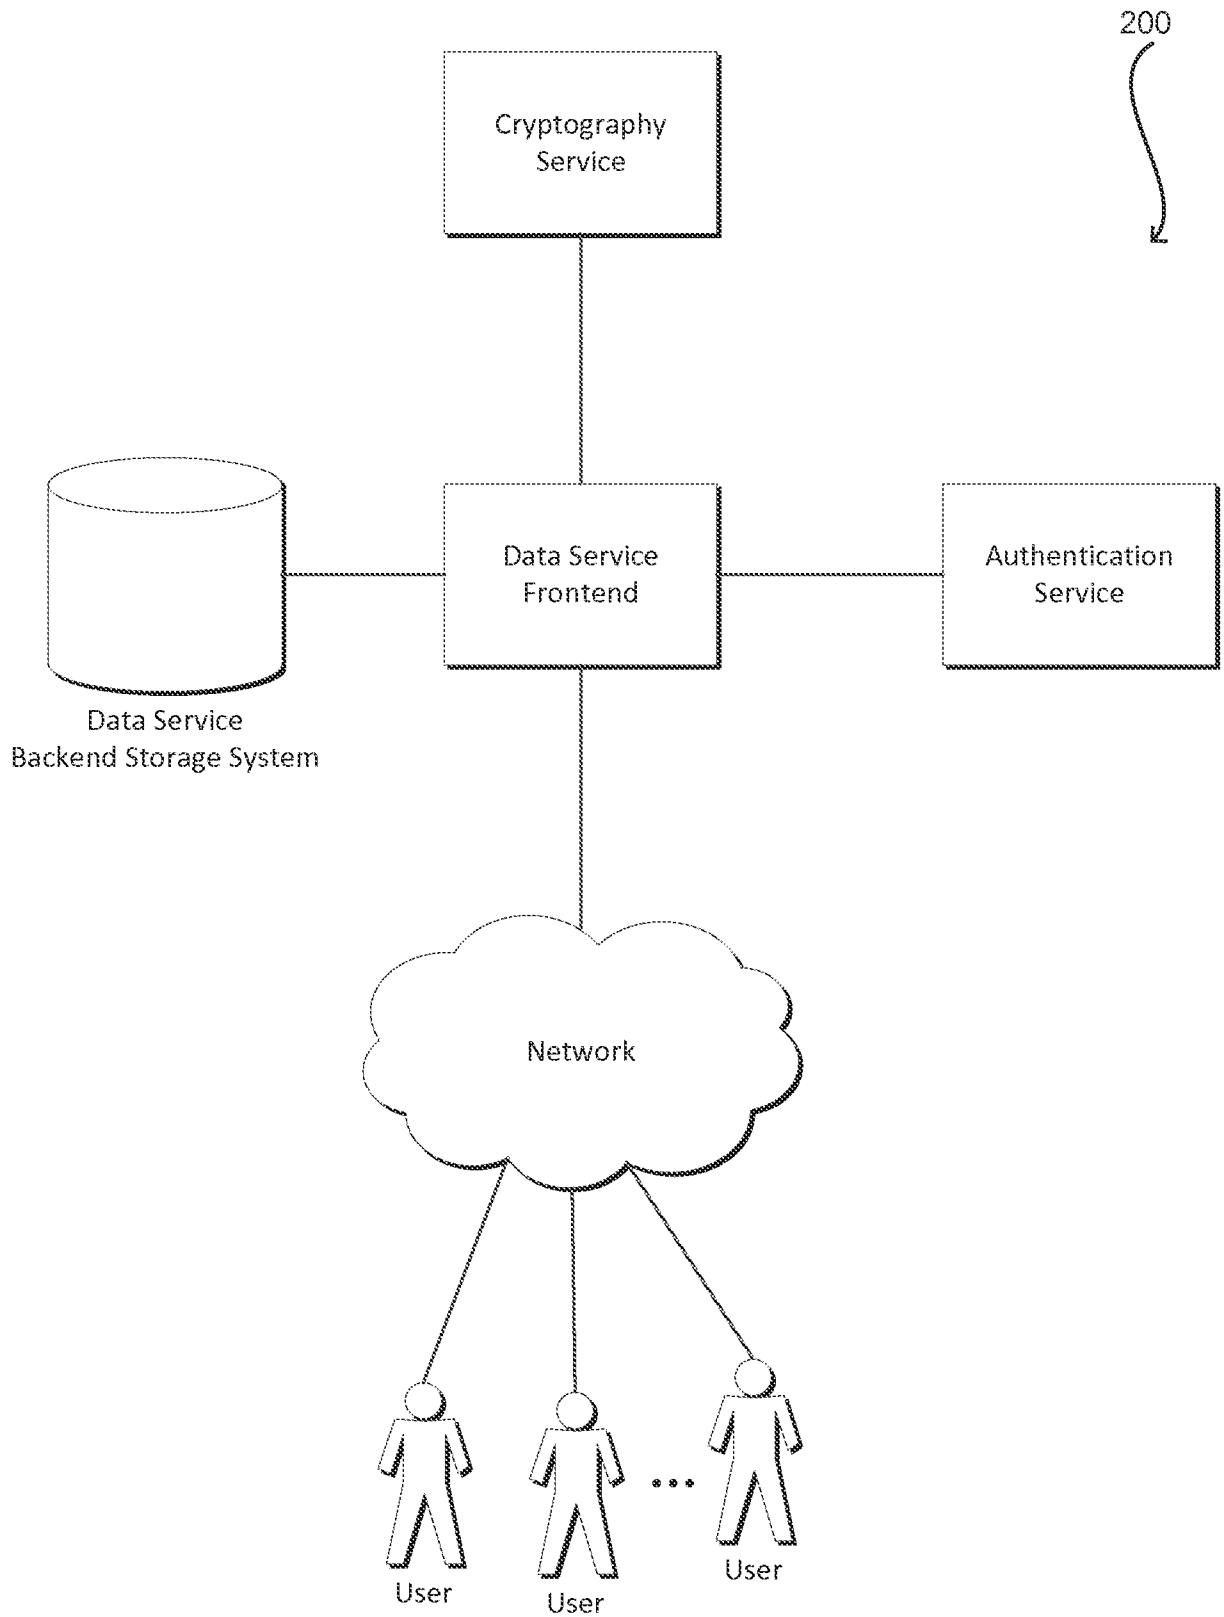


FIG. 2

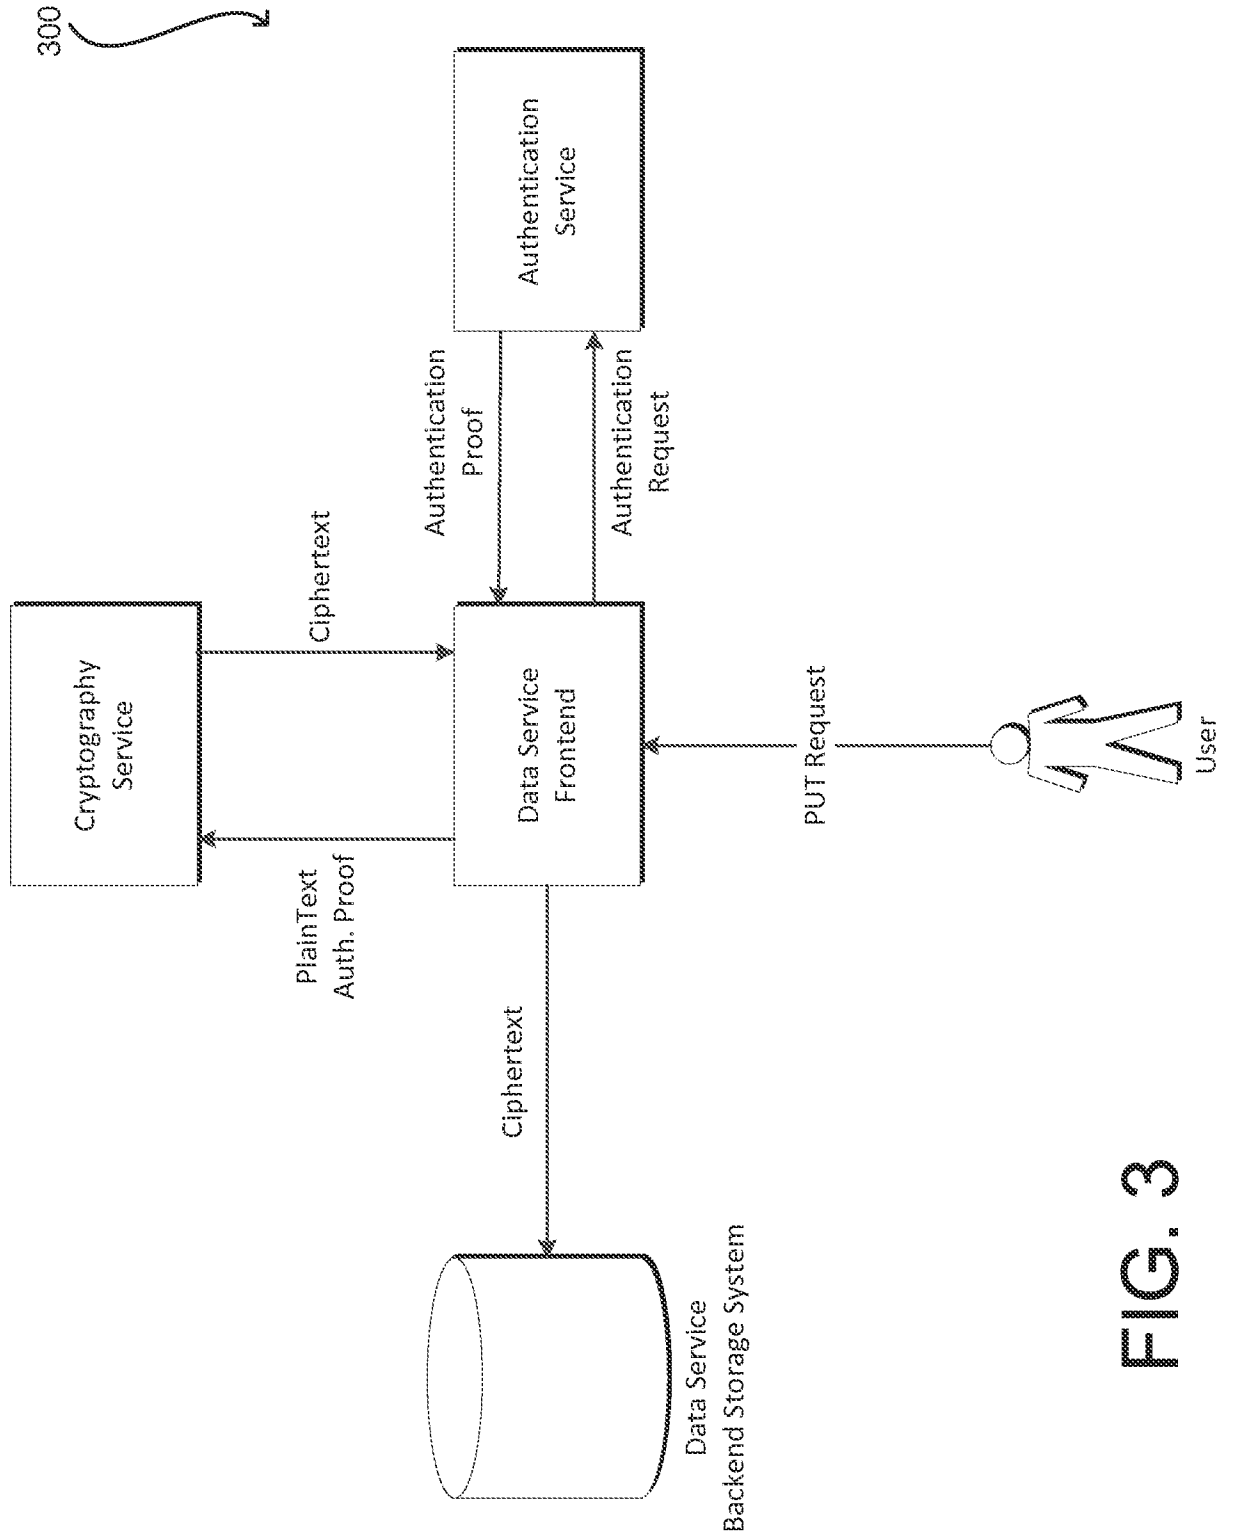


FIG. 3

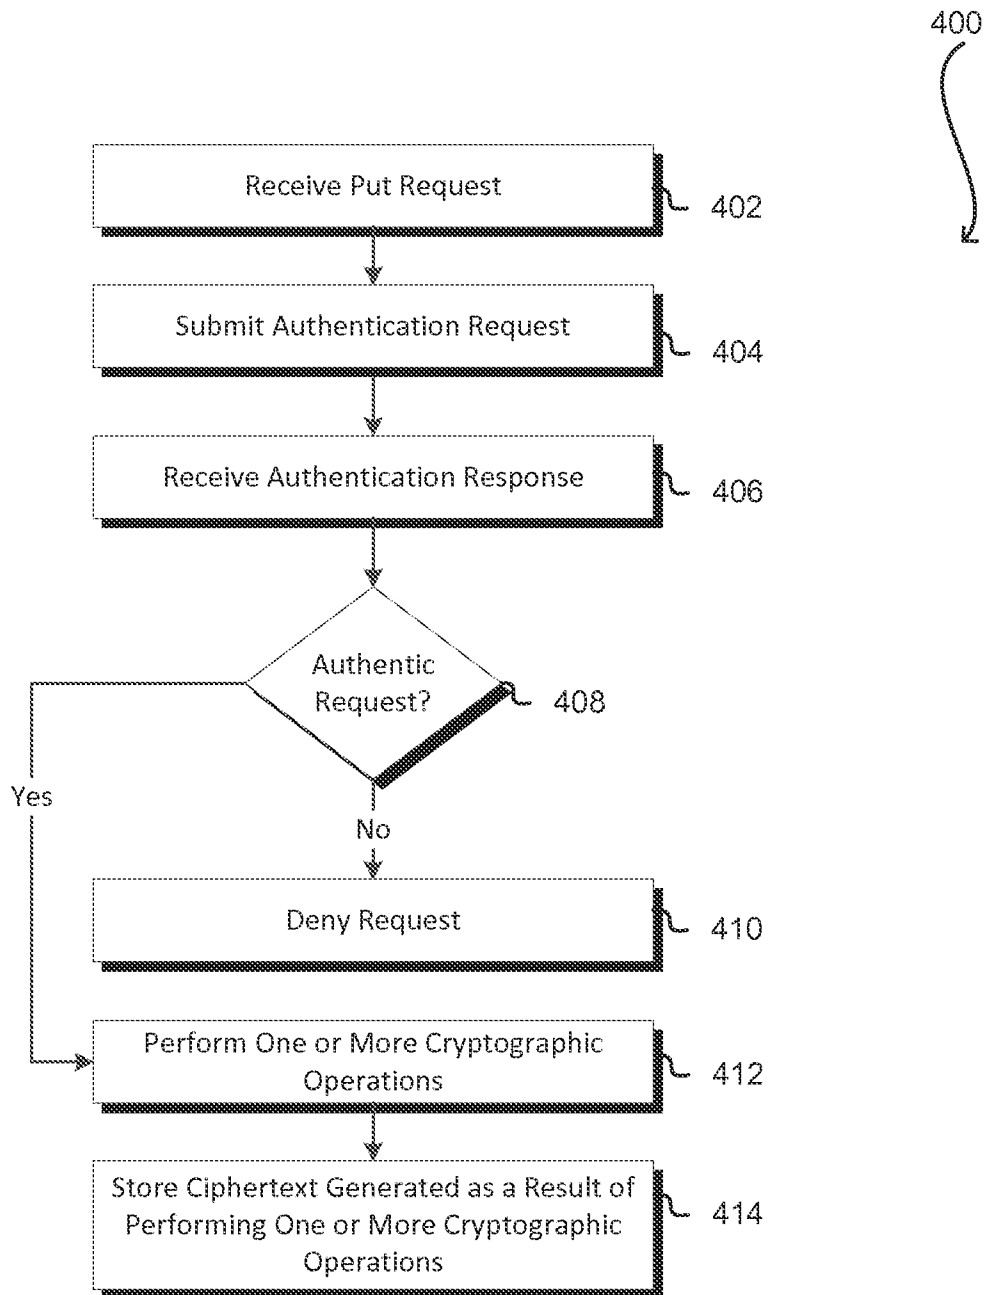


FIG. 4

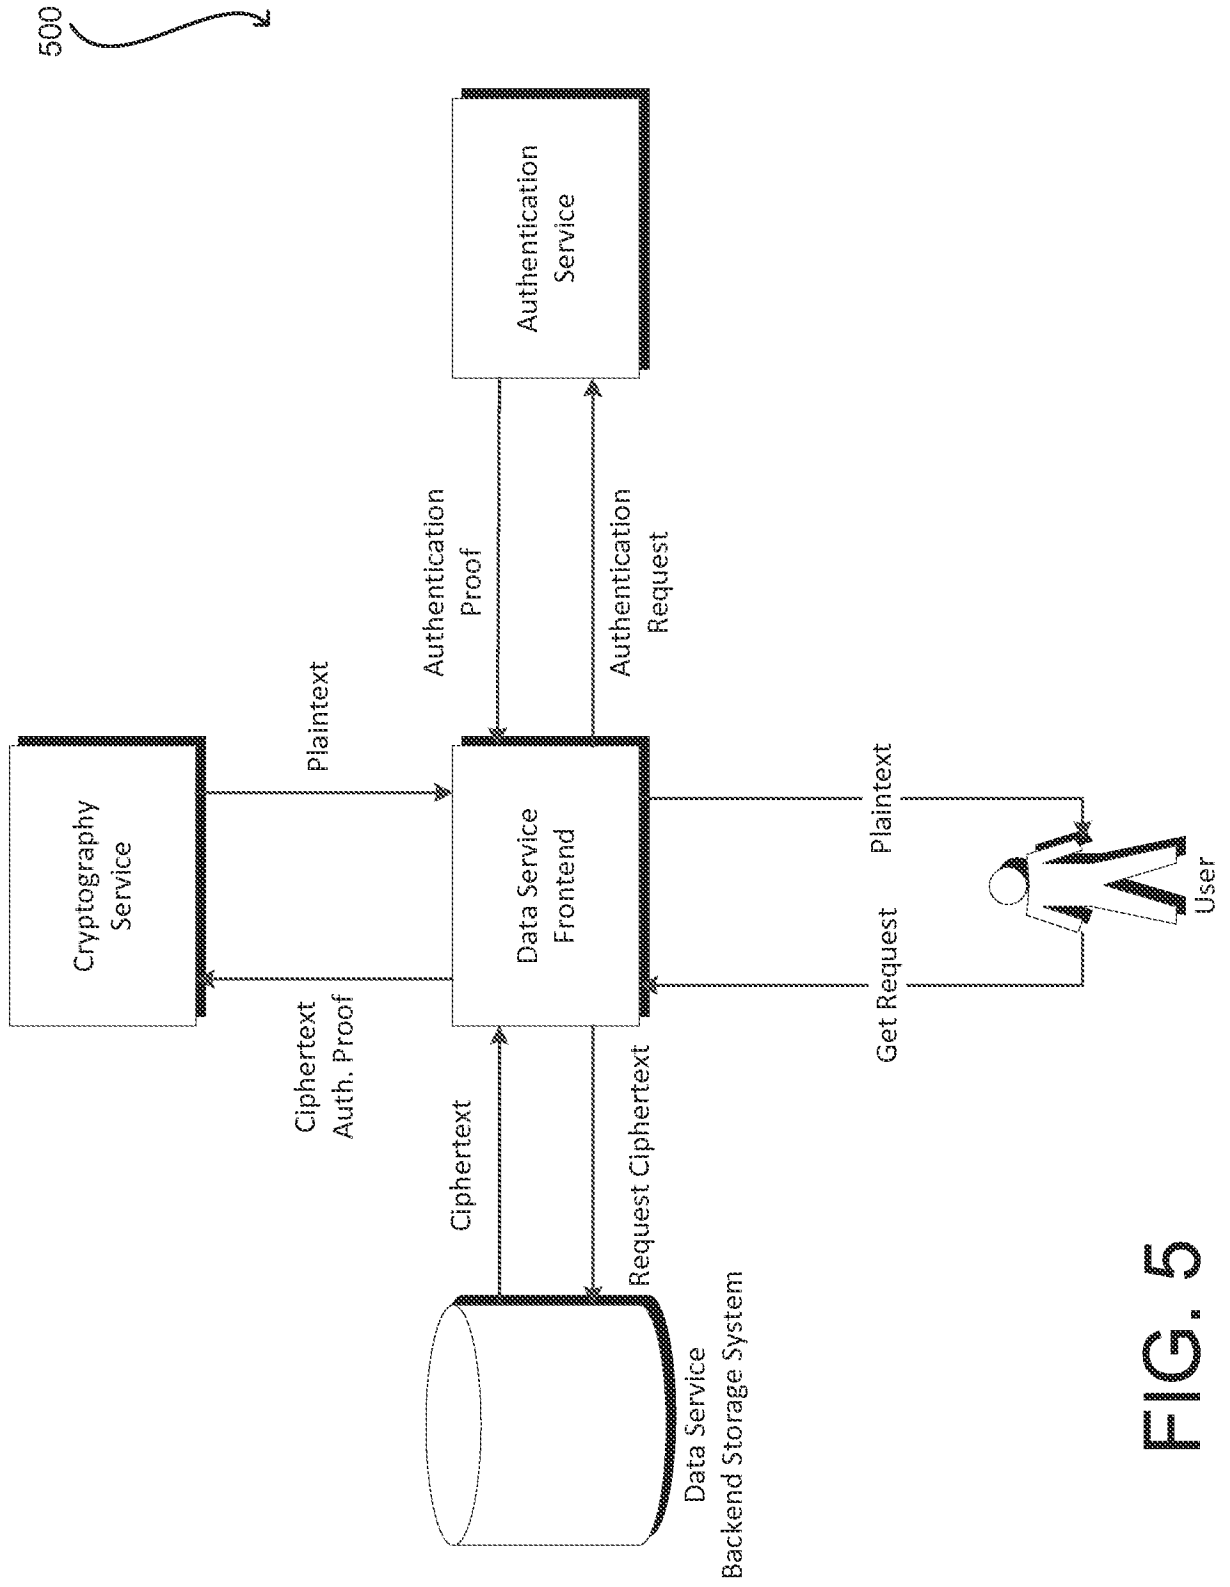


FIG. 5

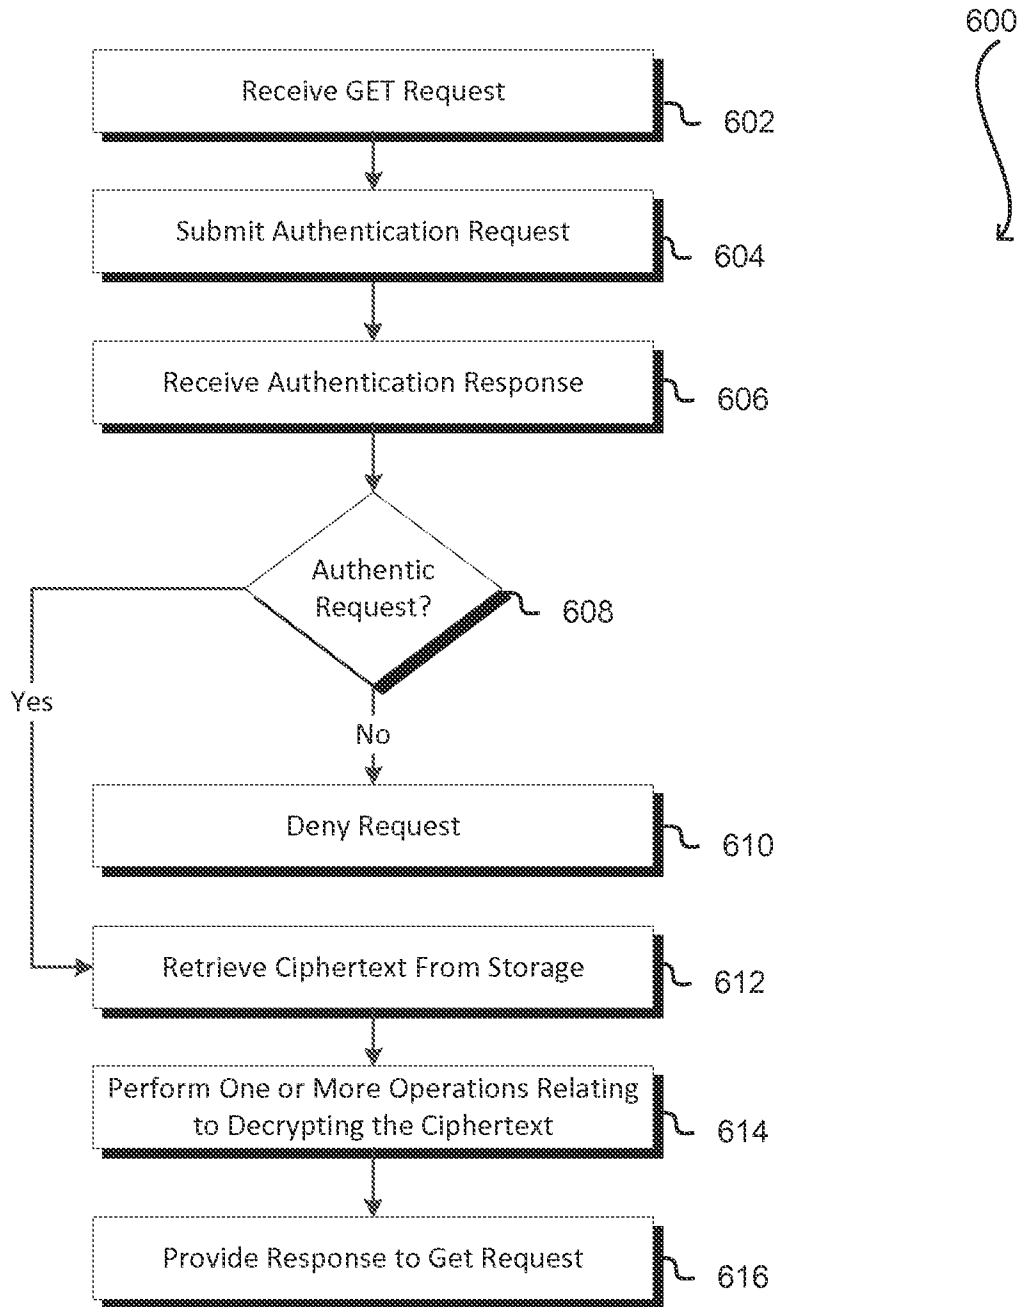


FIG. 6

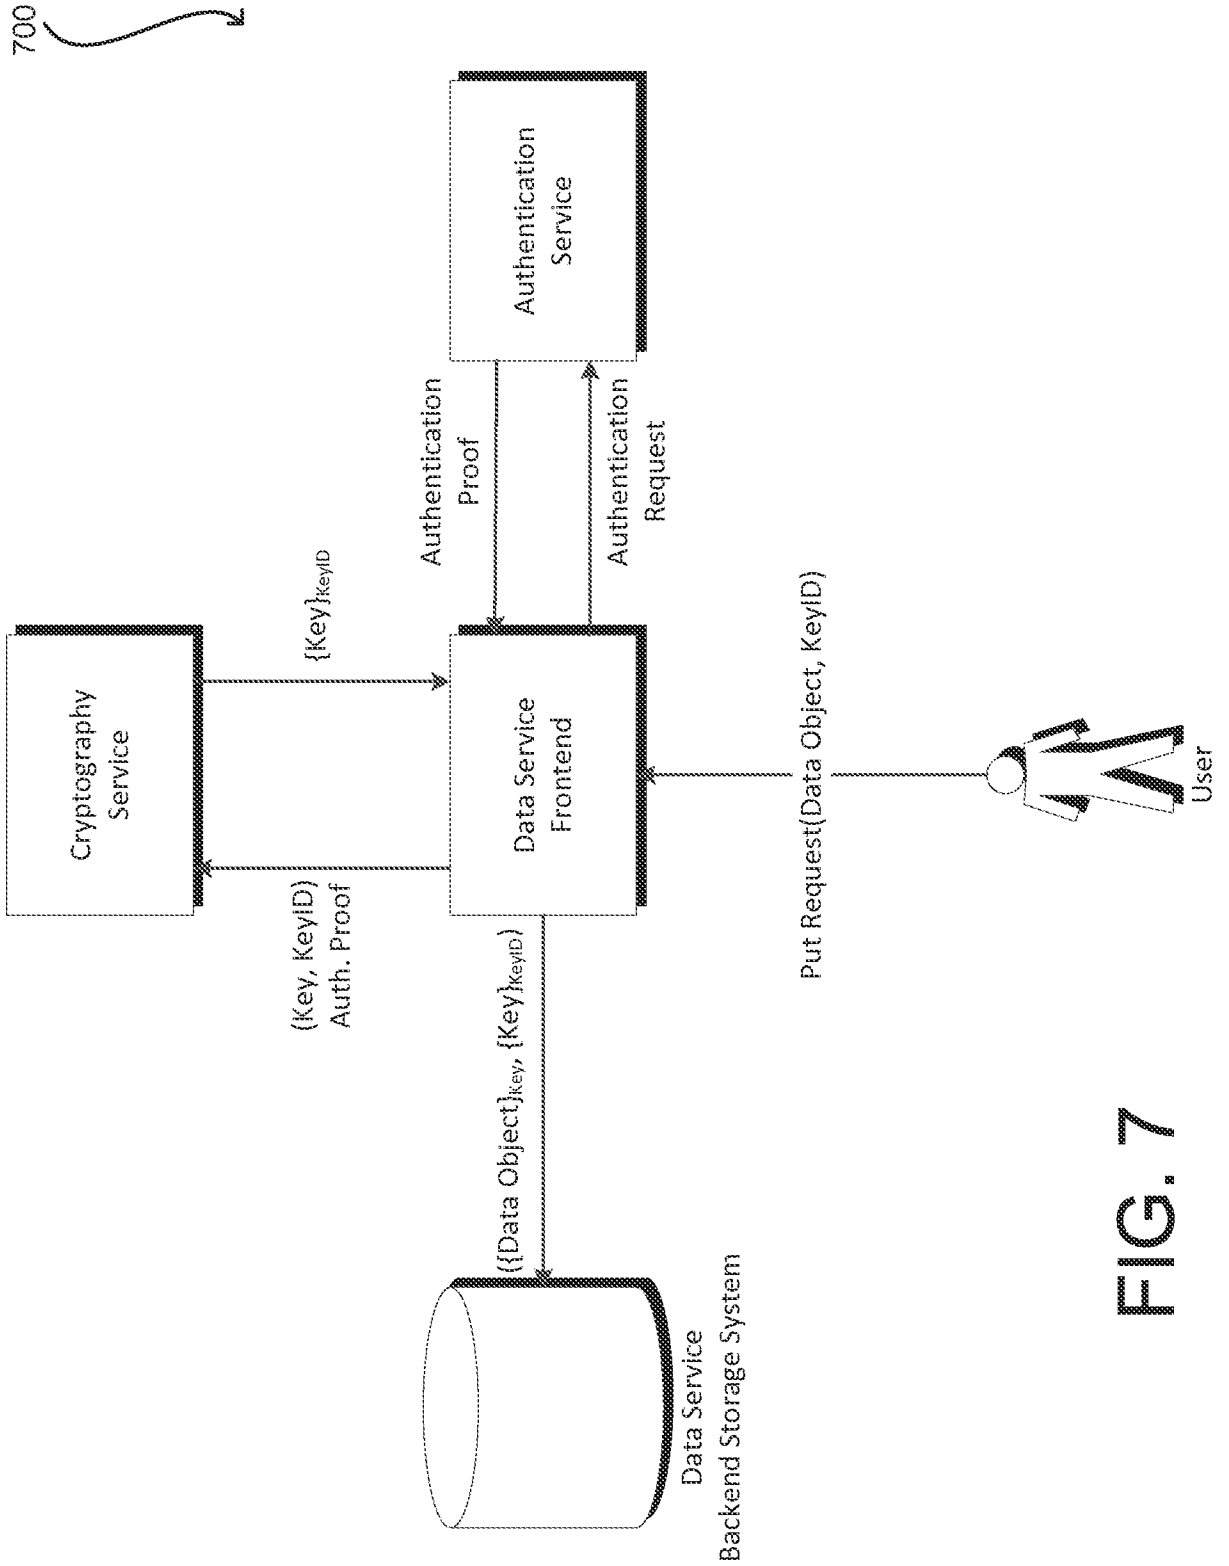


FIG. 7

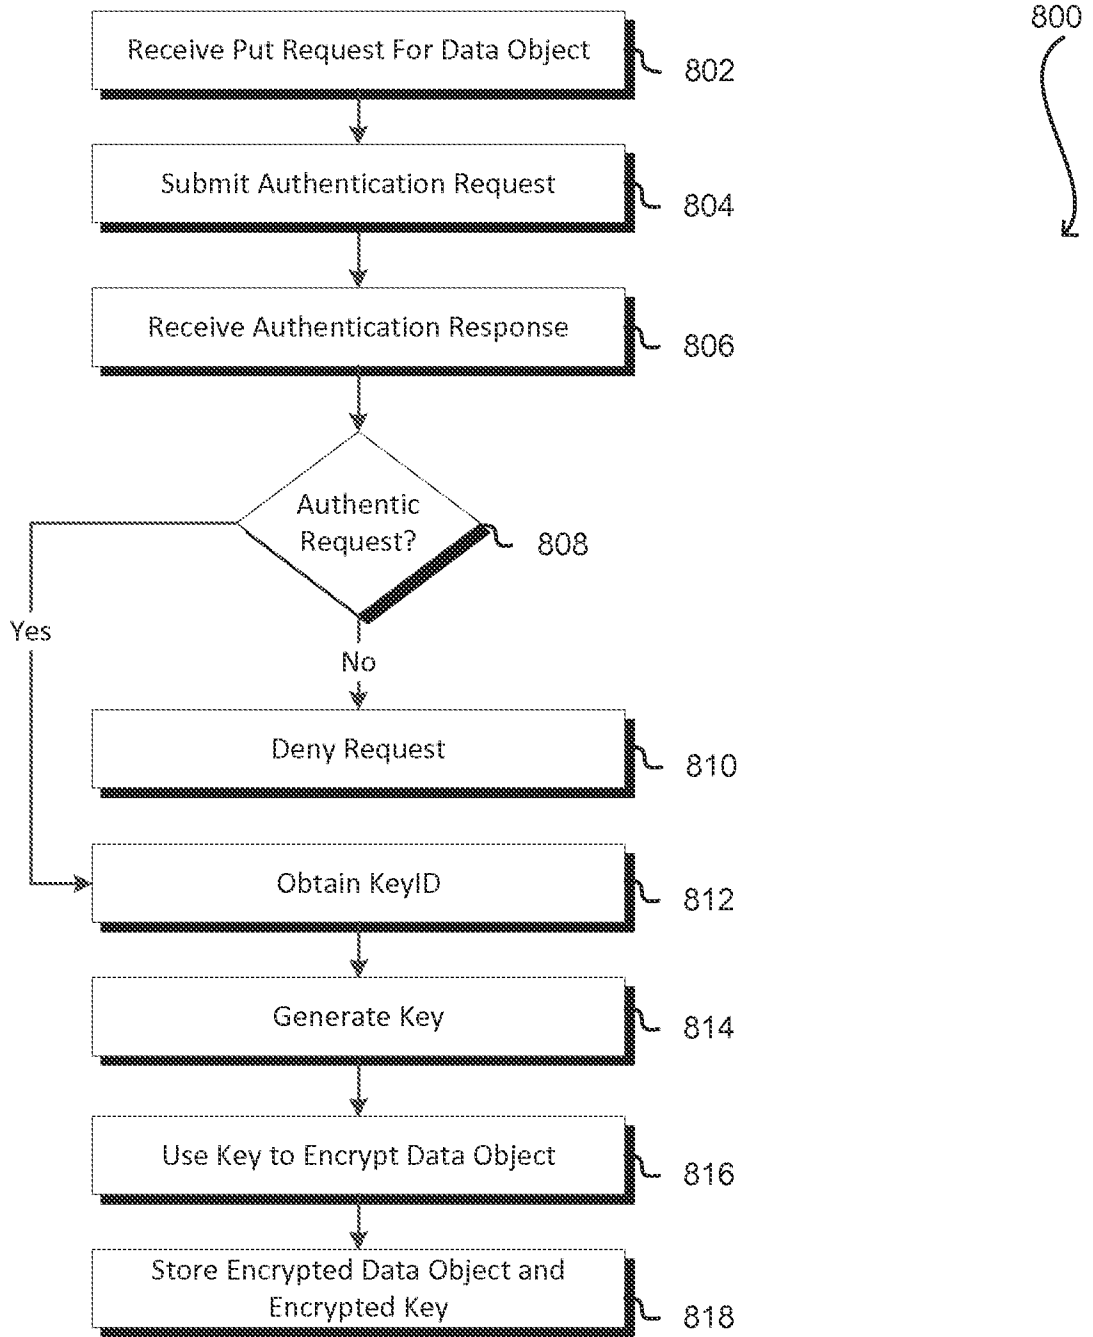


FIG. 8

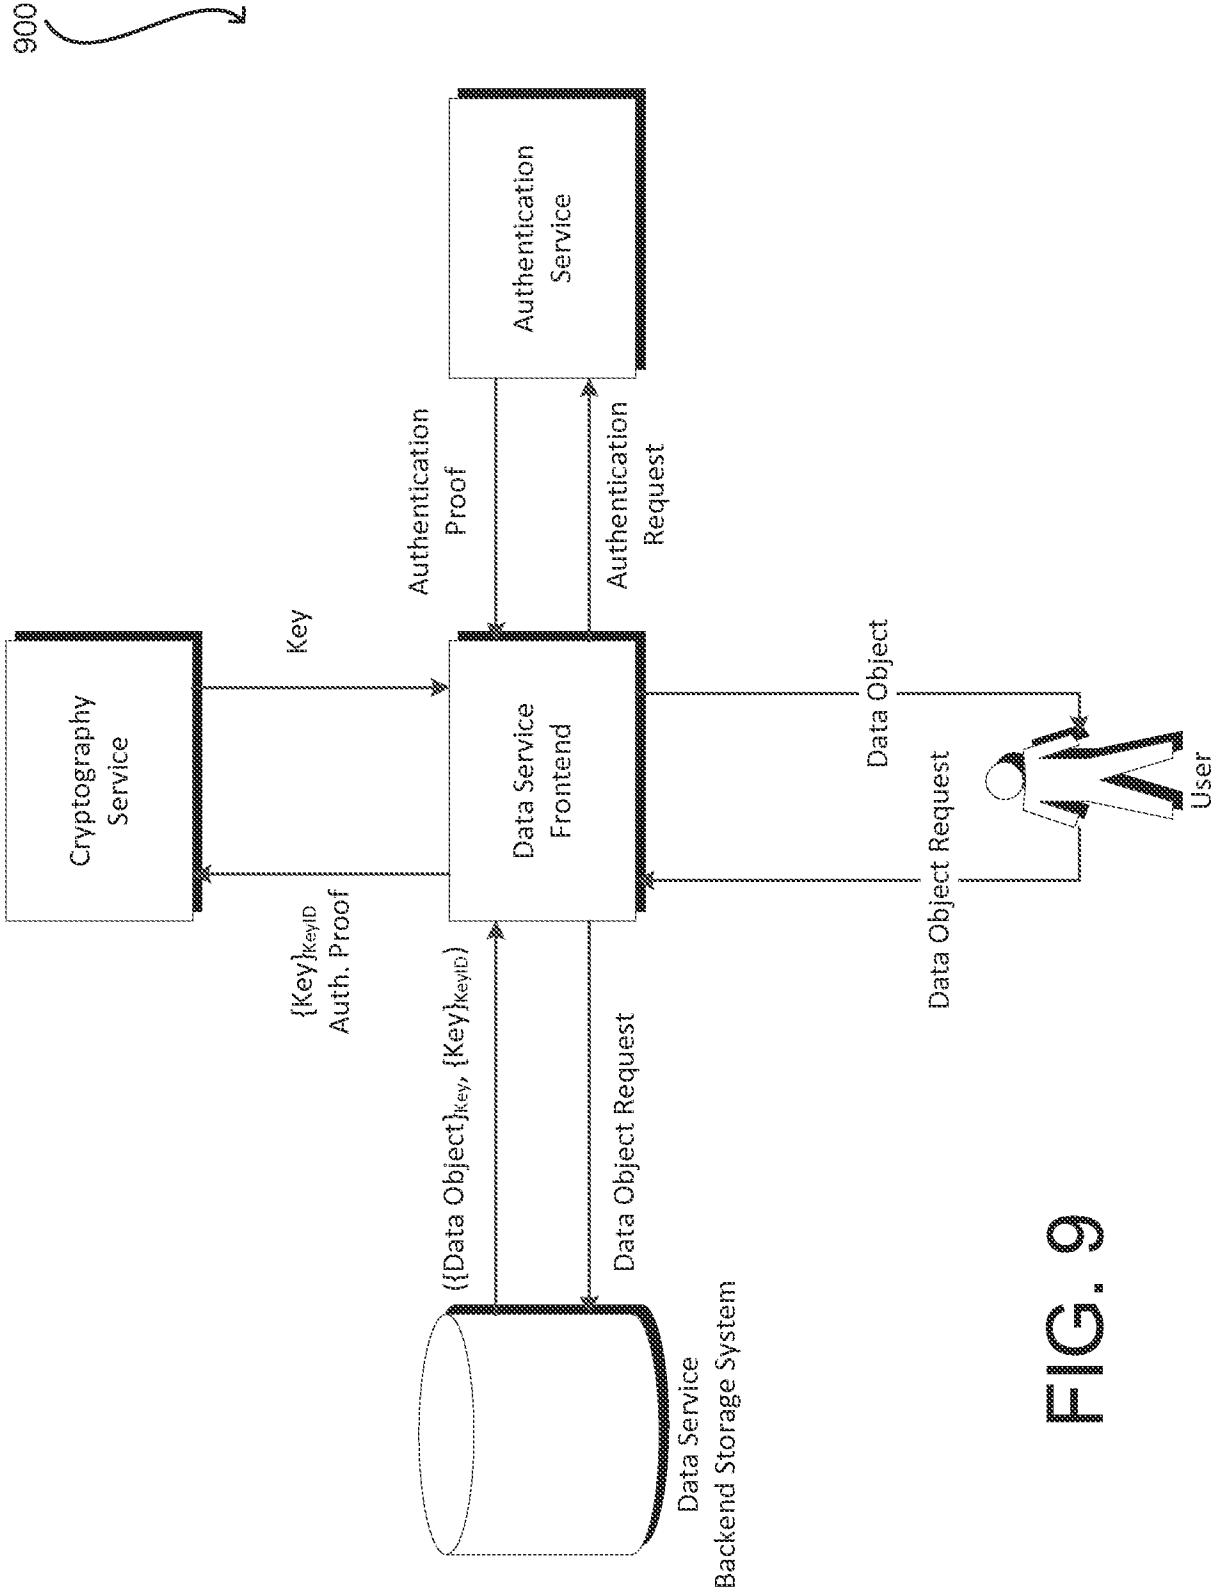


FIG. 9

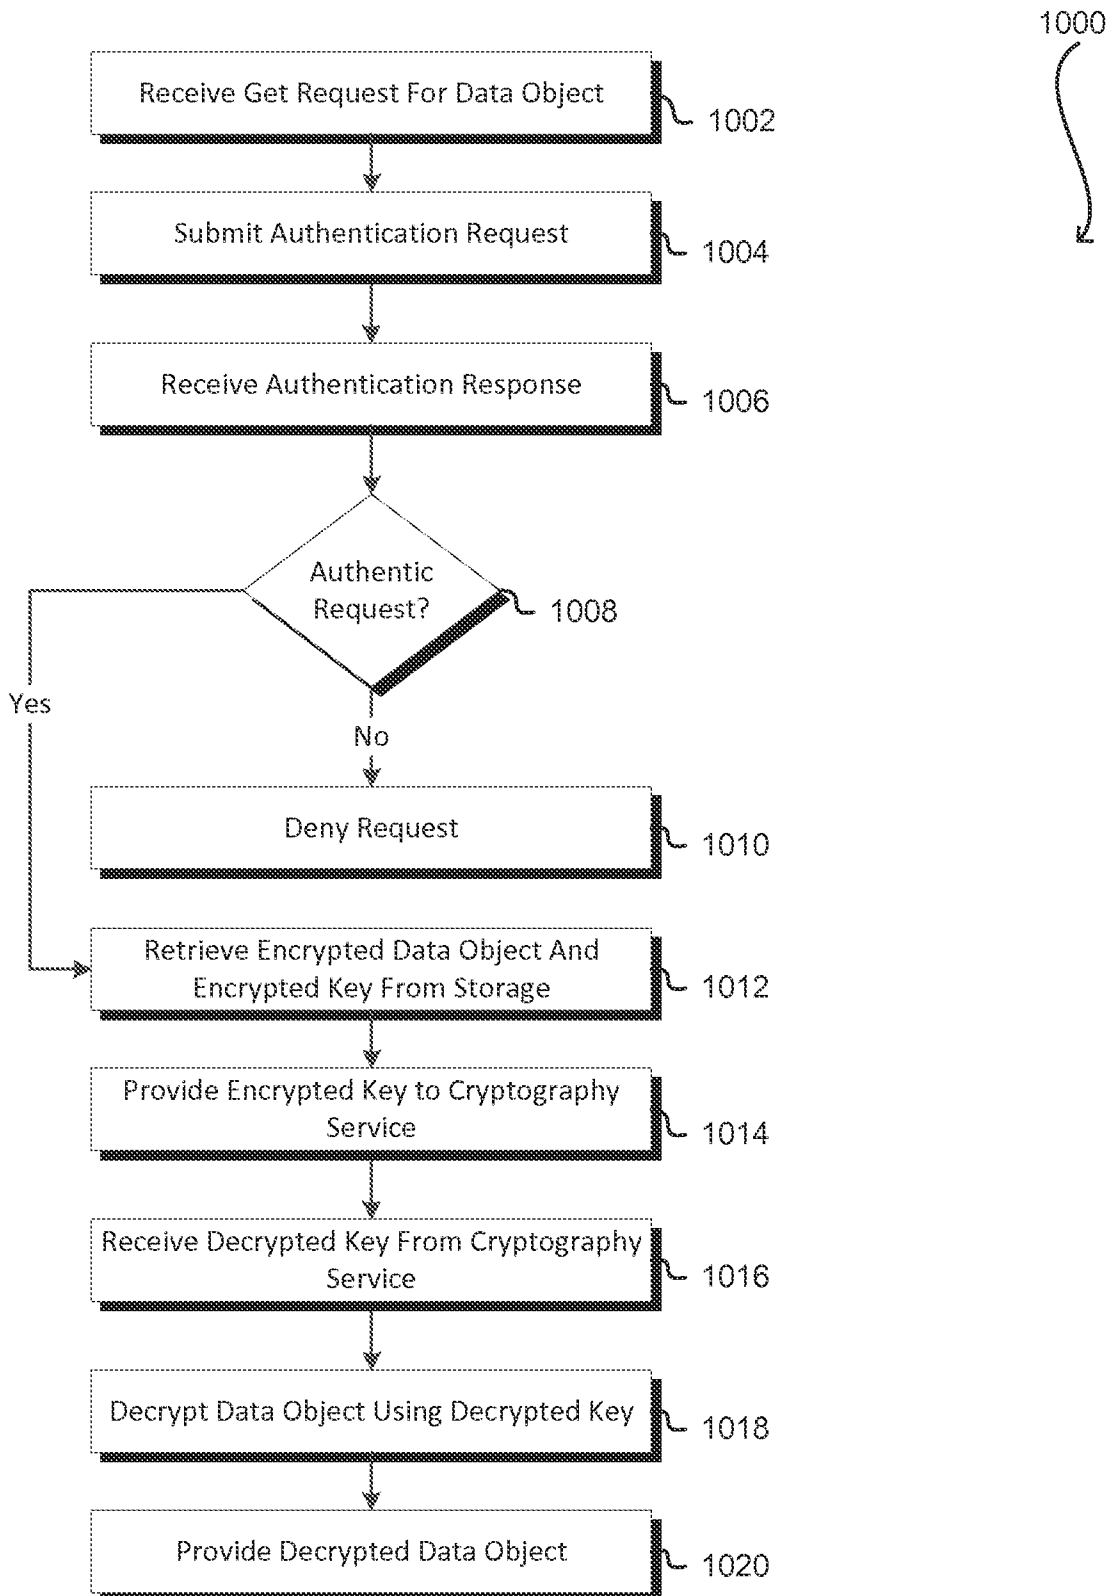


FIG. 10

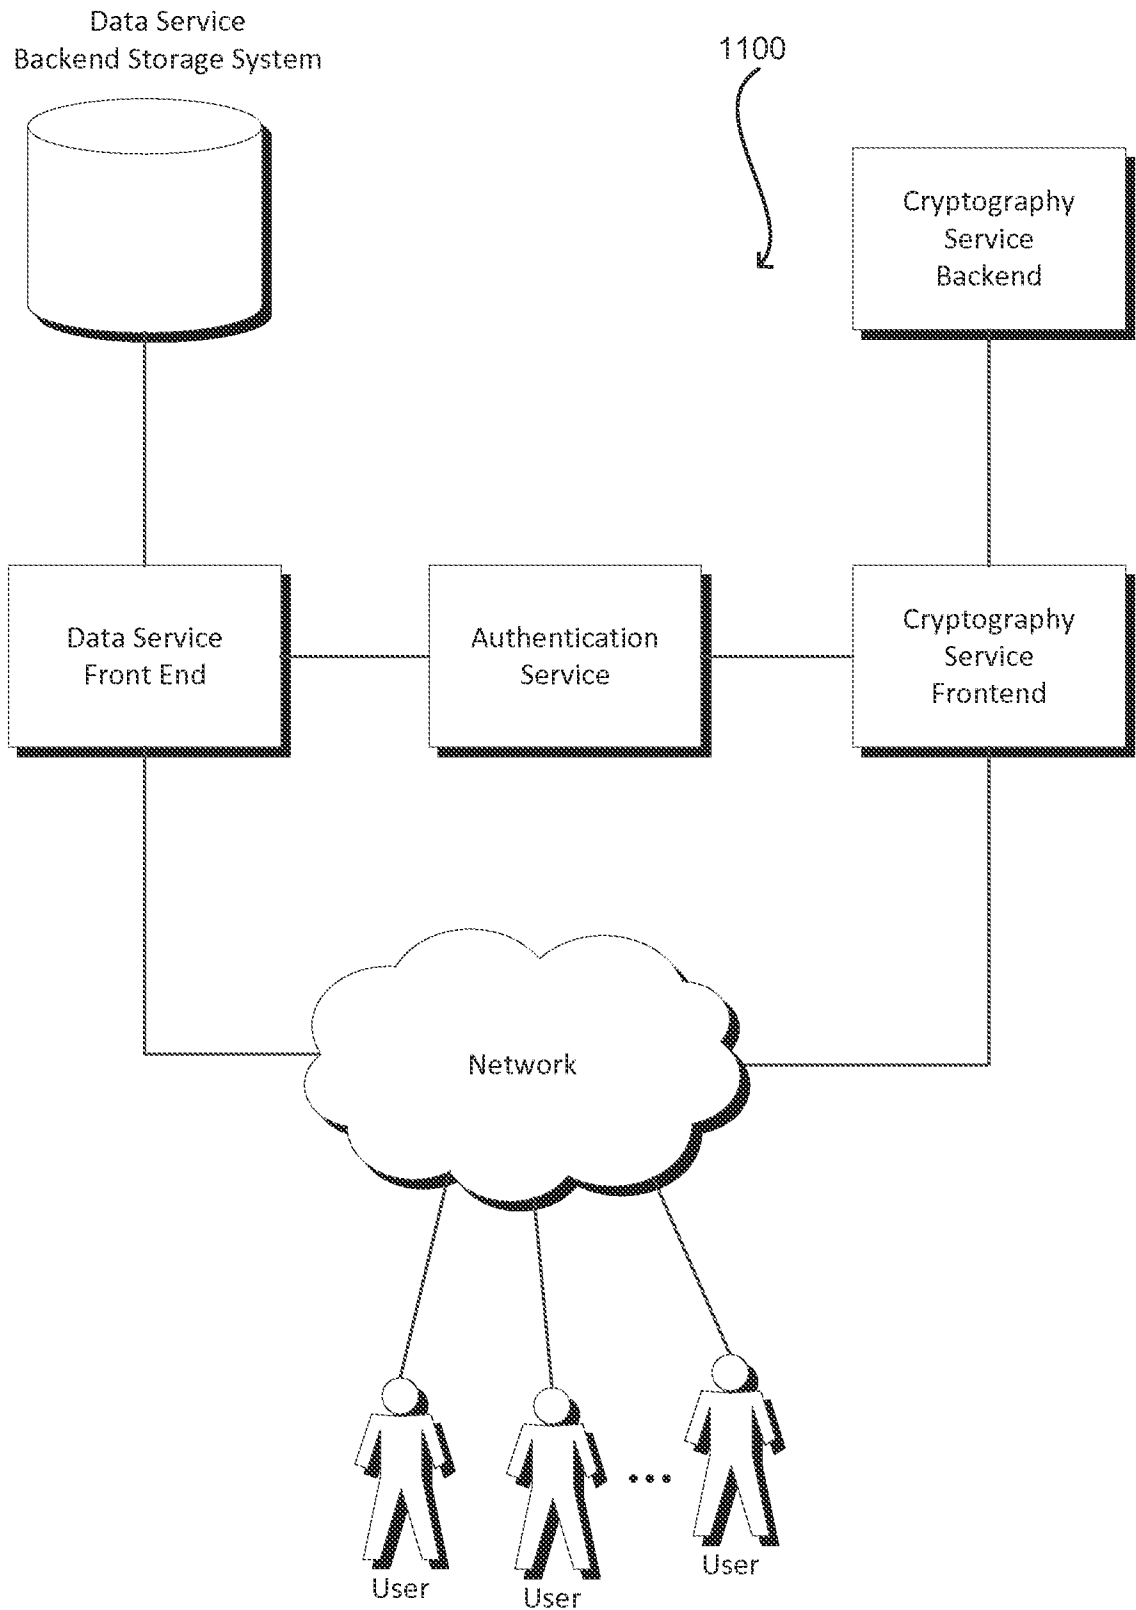


FIG. 11

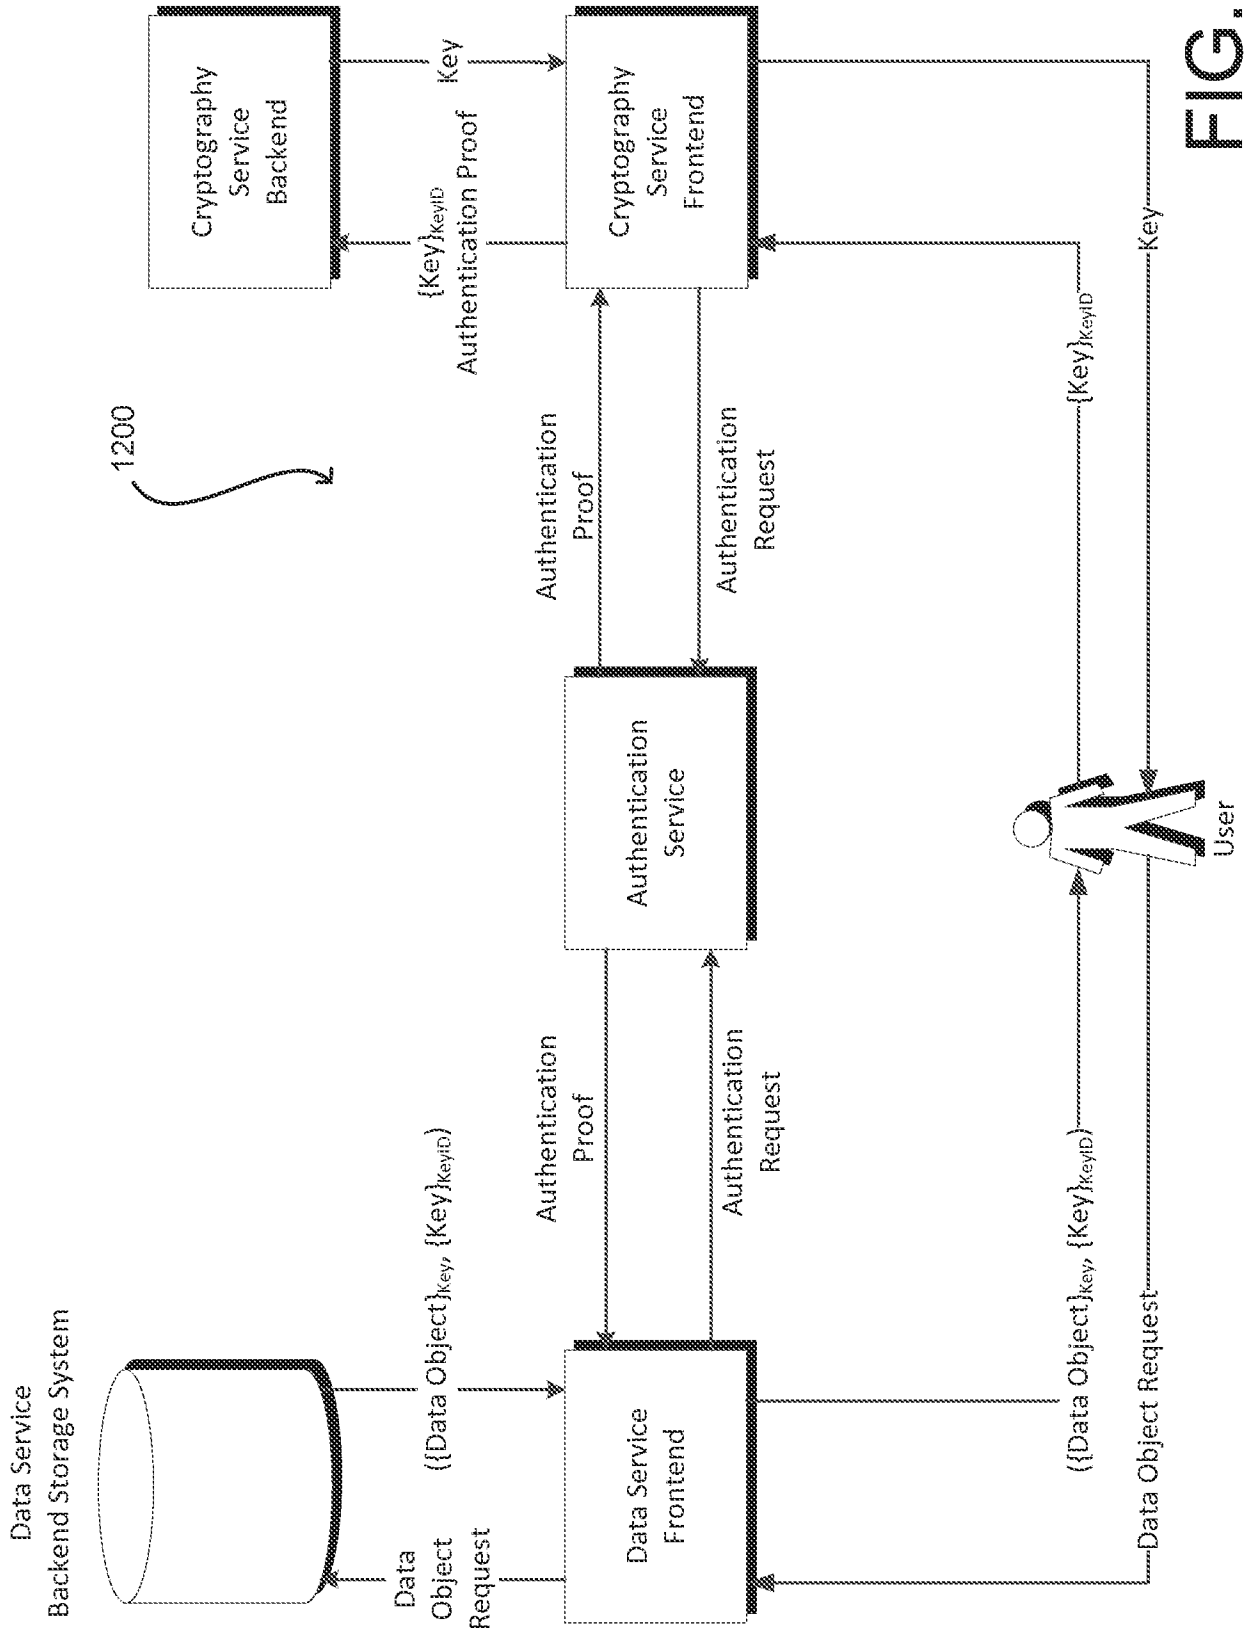


FIG. 12

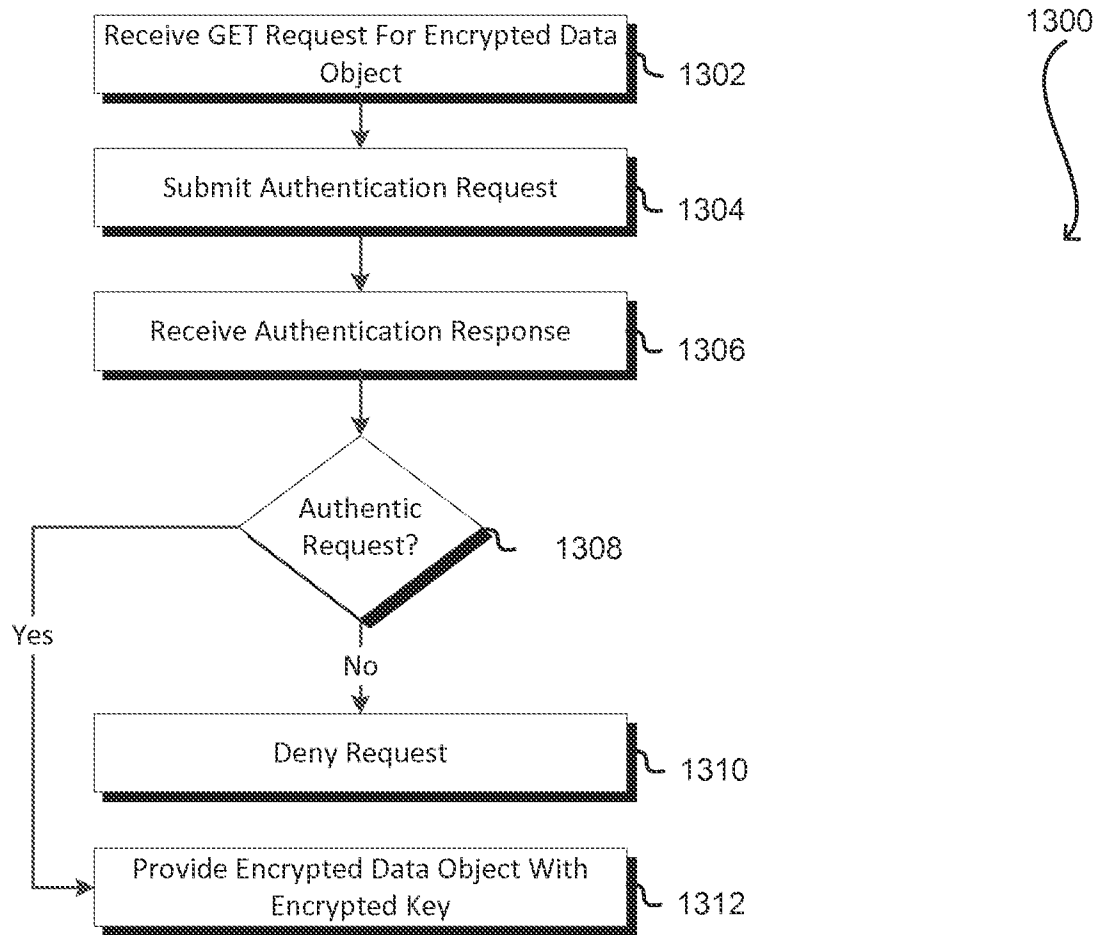
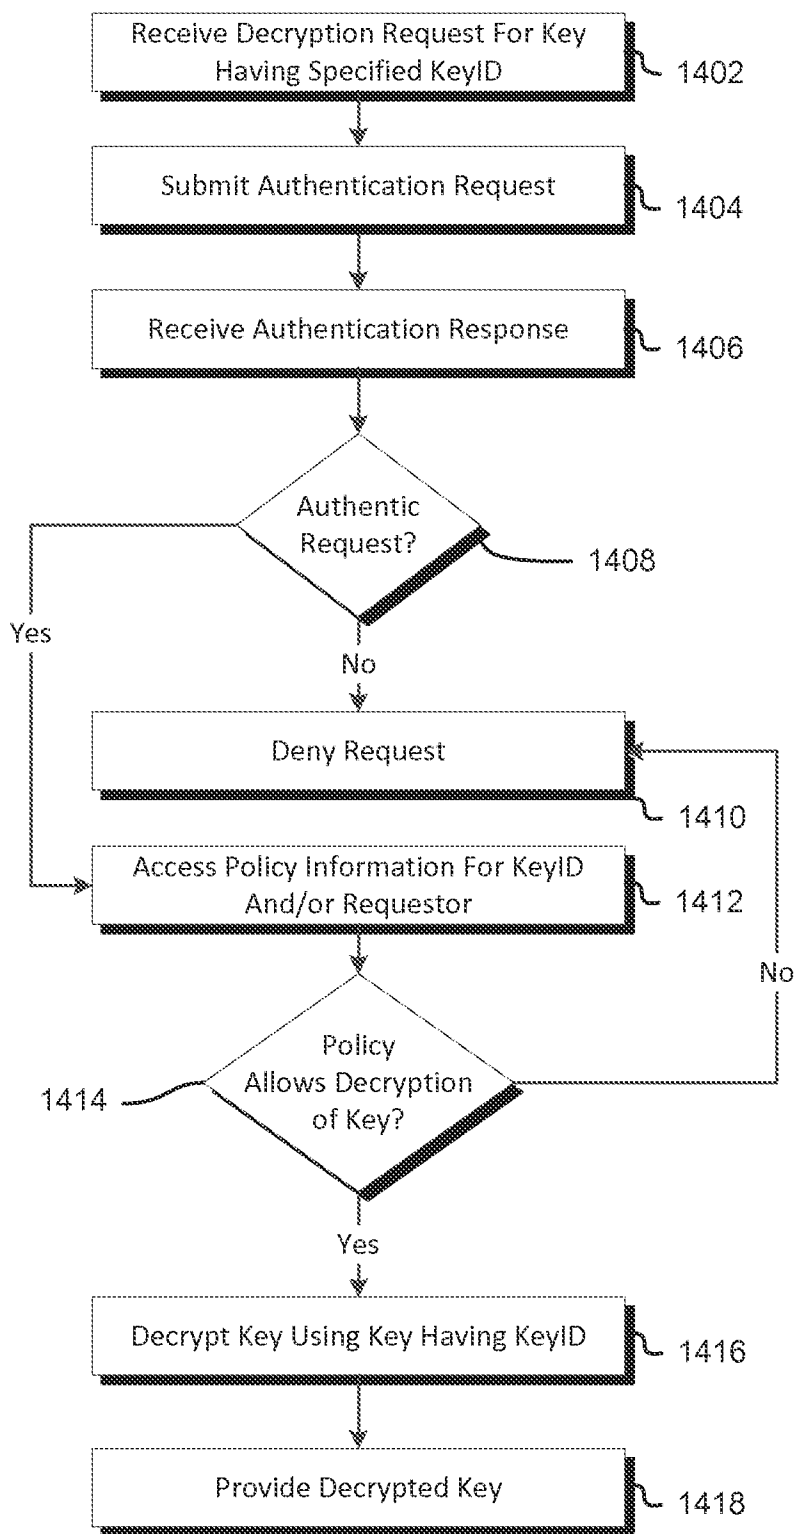


FIG. 13



1400

FIG. 14

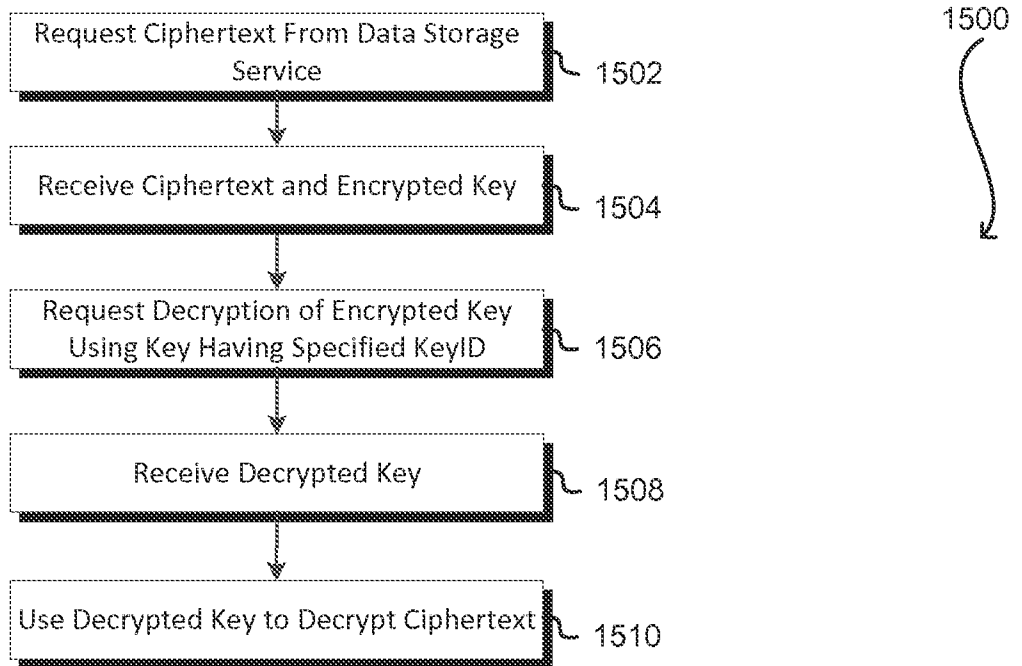


FIG. 15

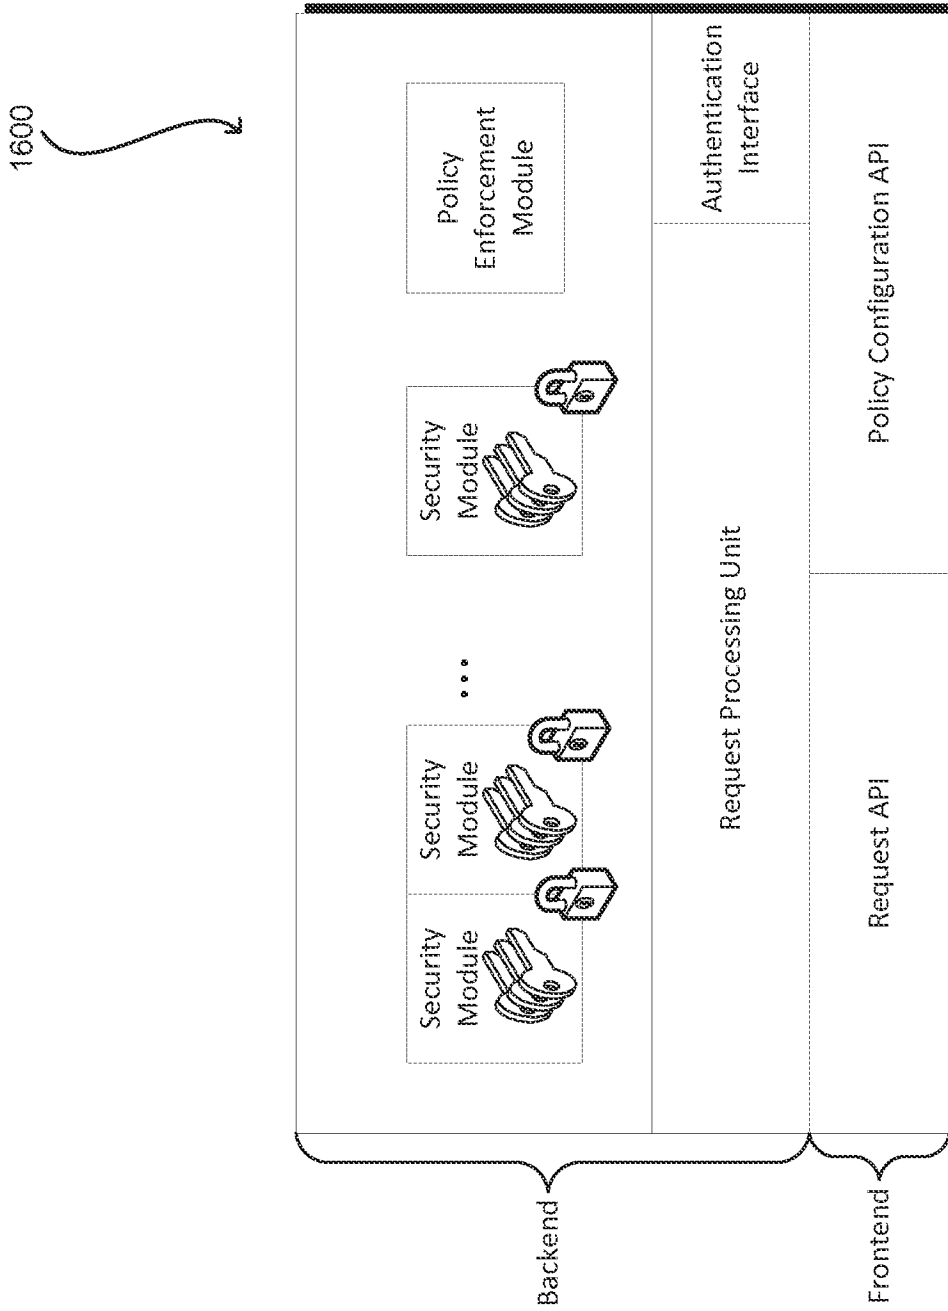


FIG. 16



FIG. 17

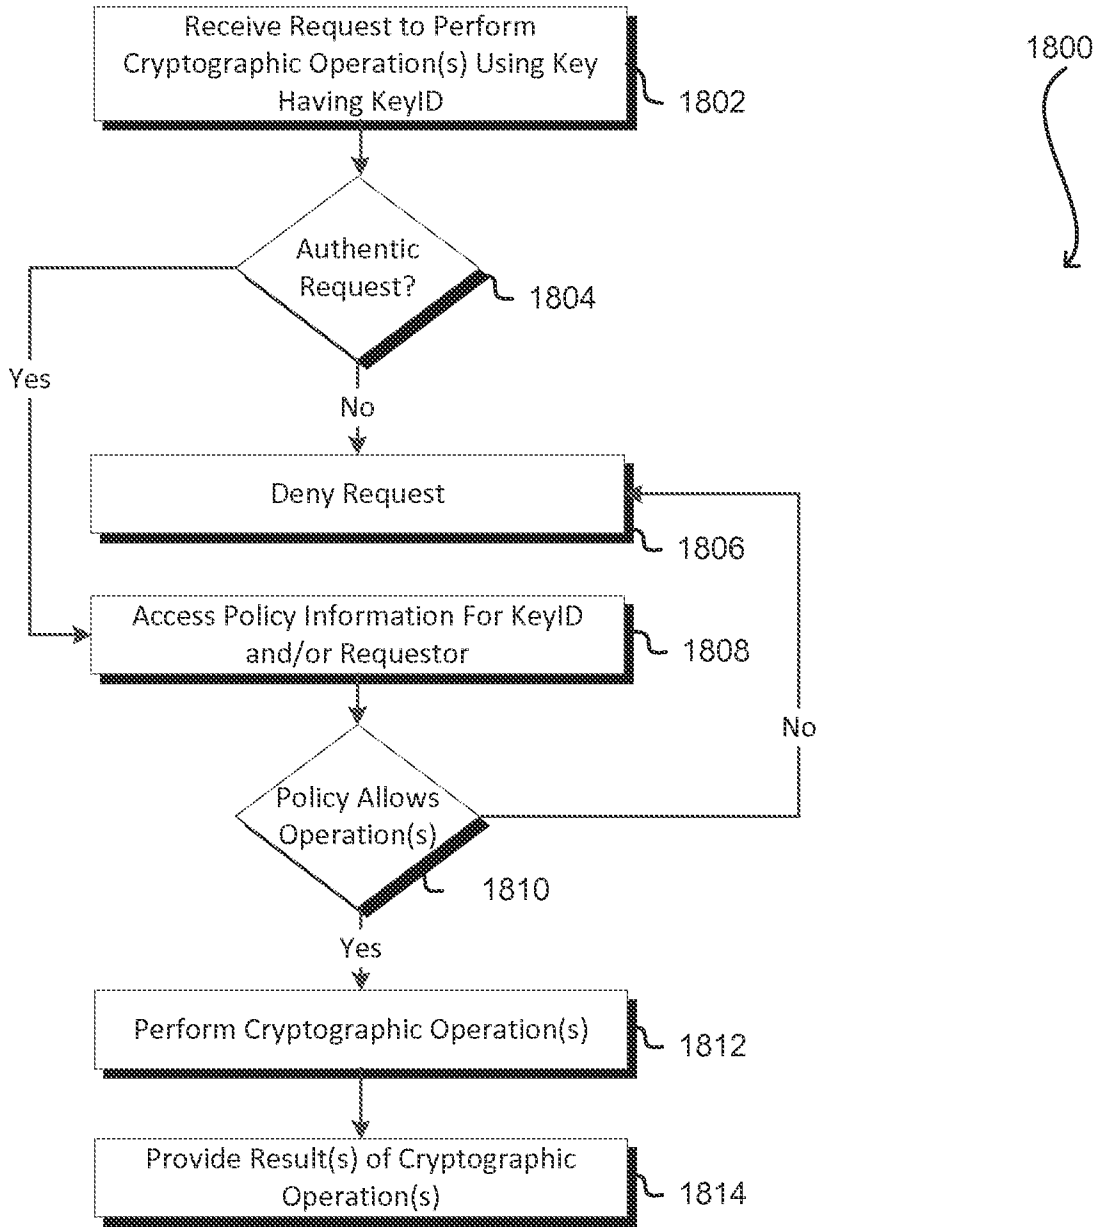


FIG. 18

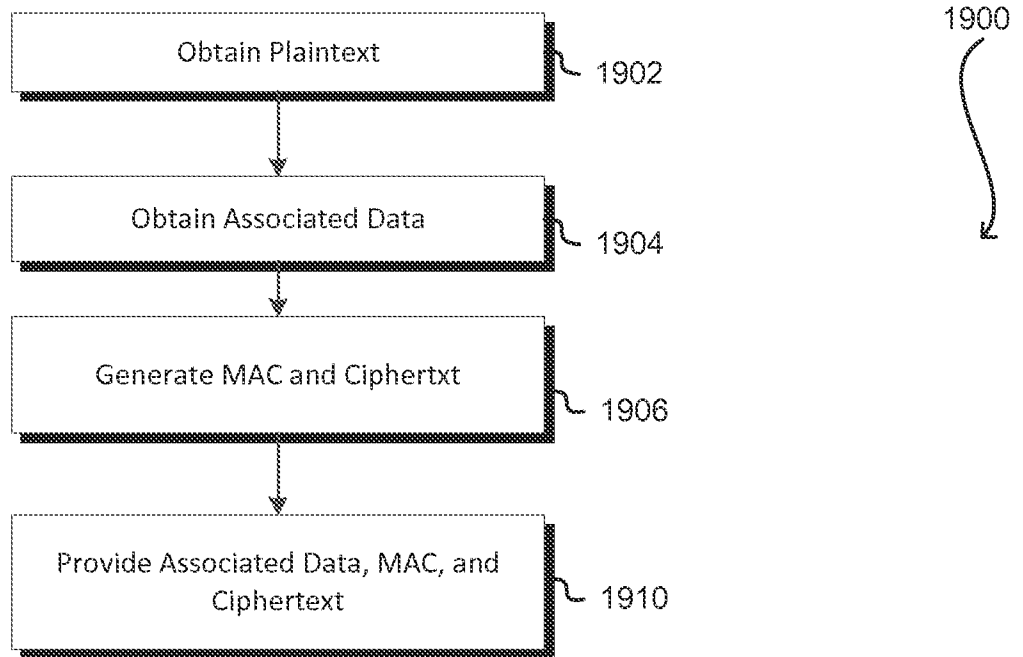


FIG. 19

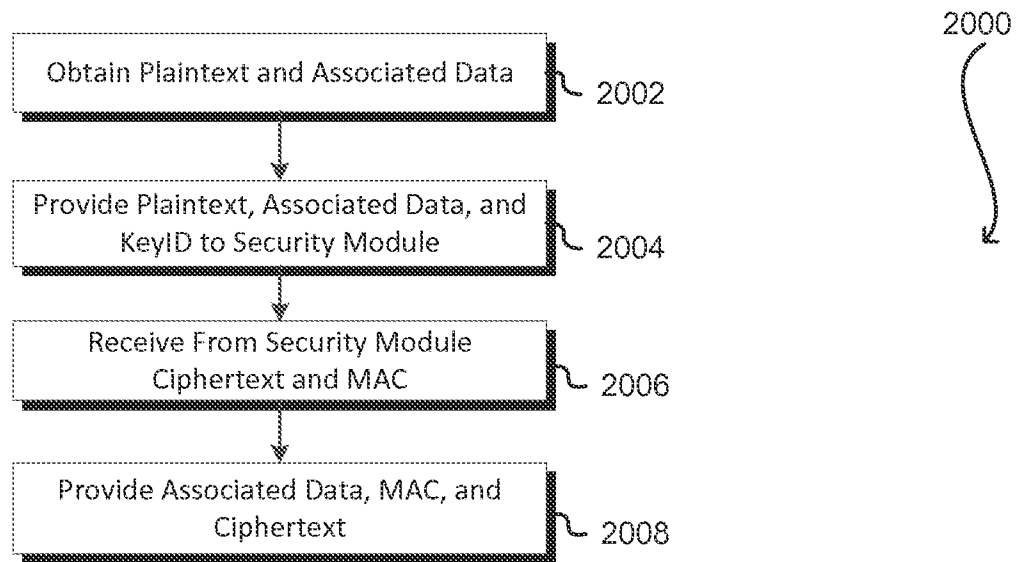


FIG. 20

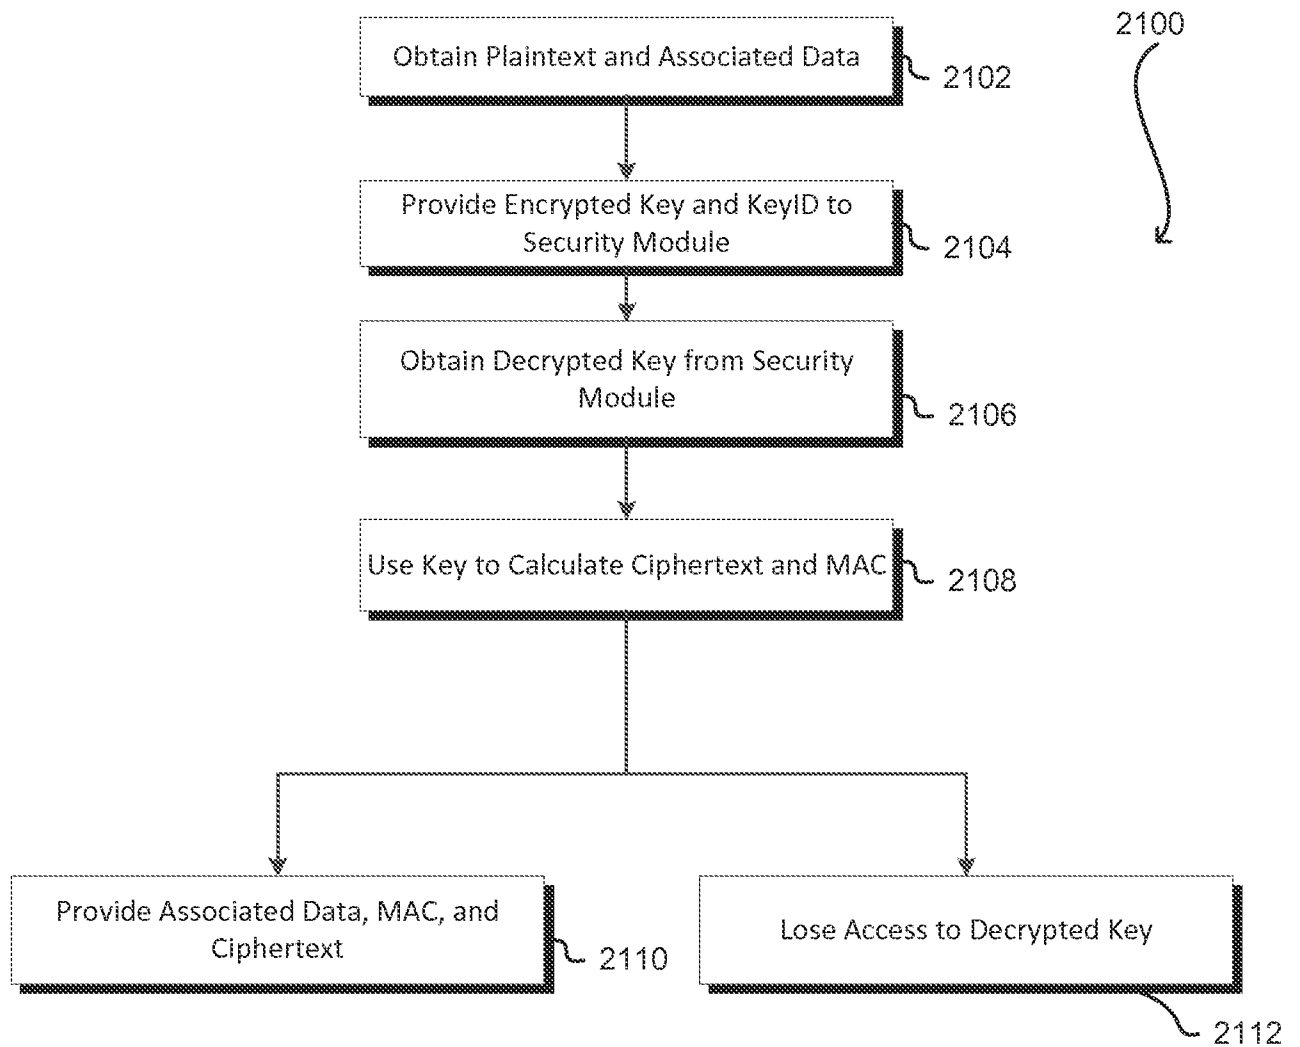


FIG. 21

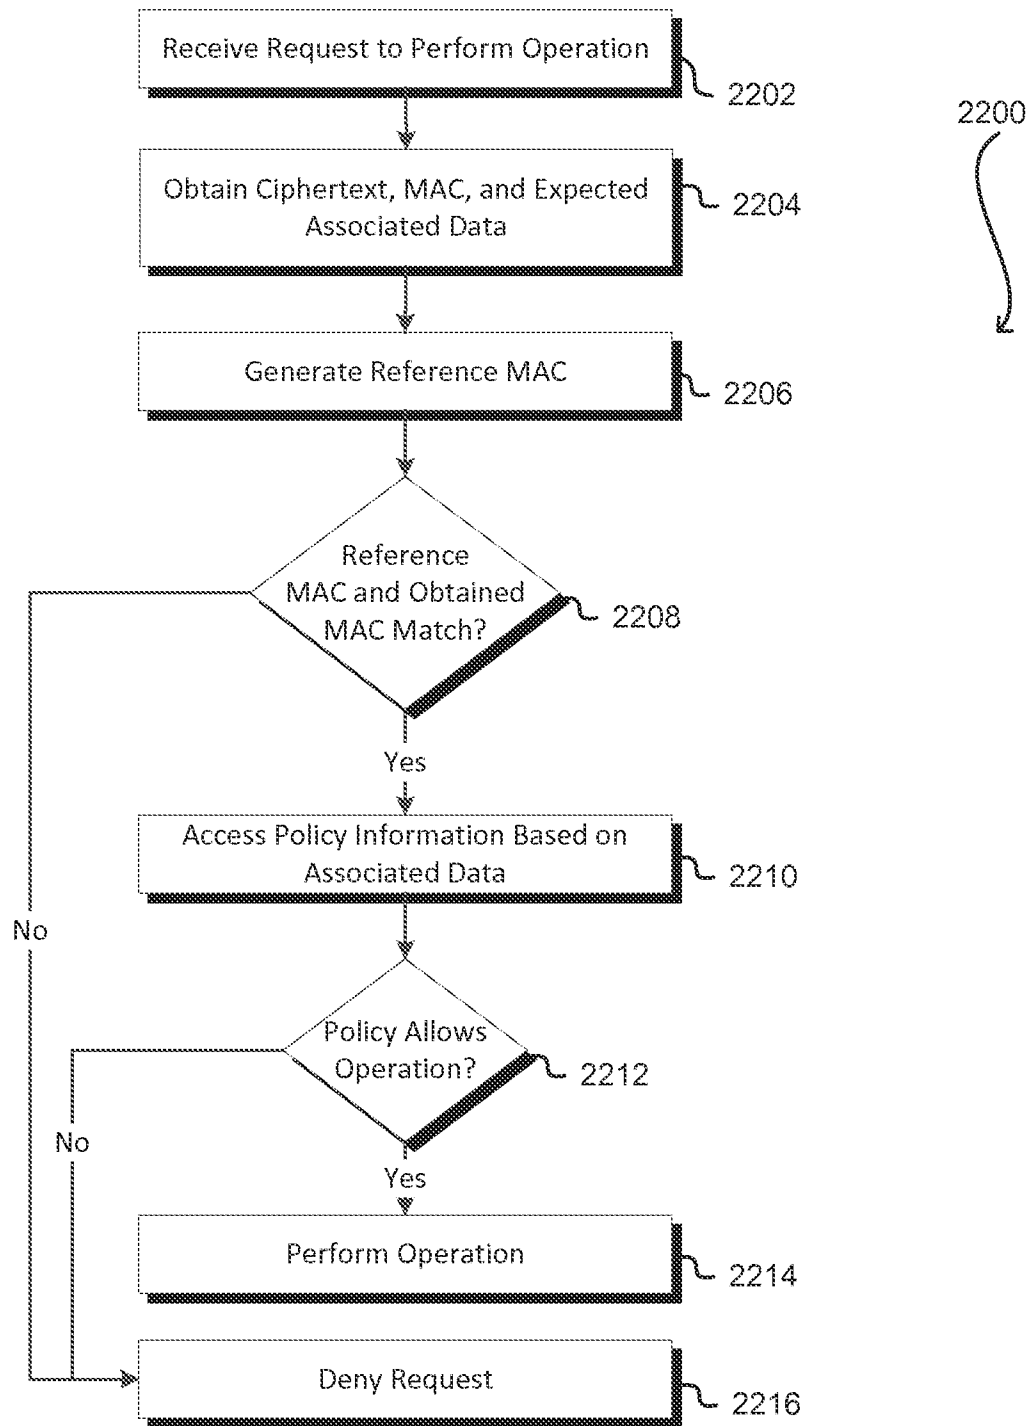


FIG. 22

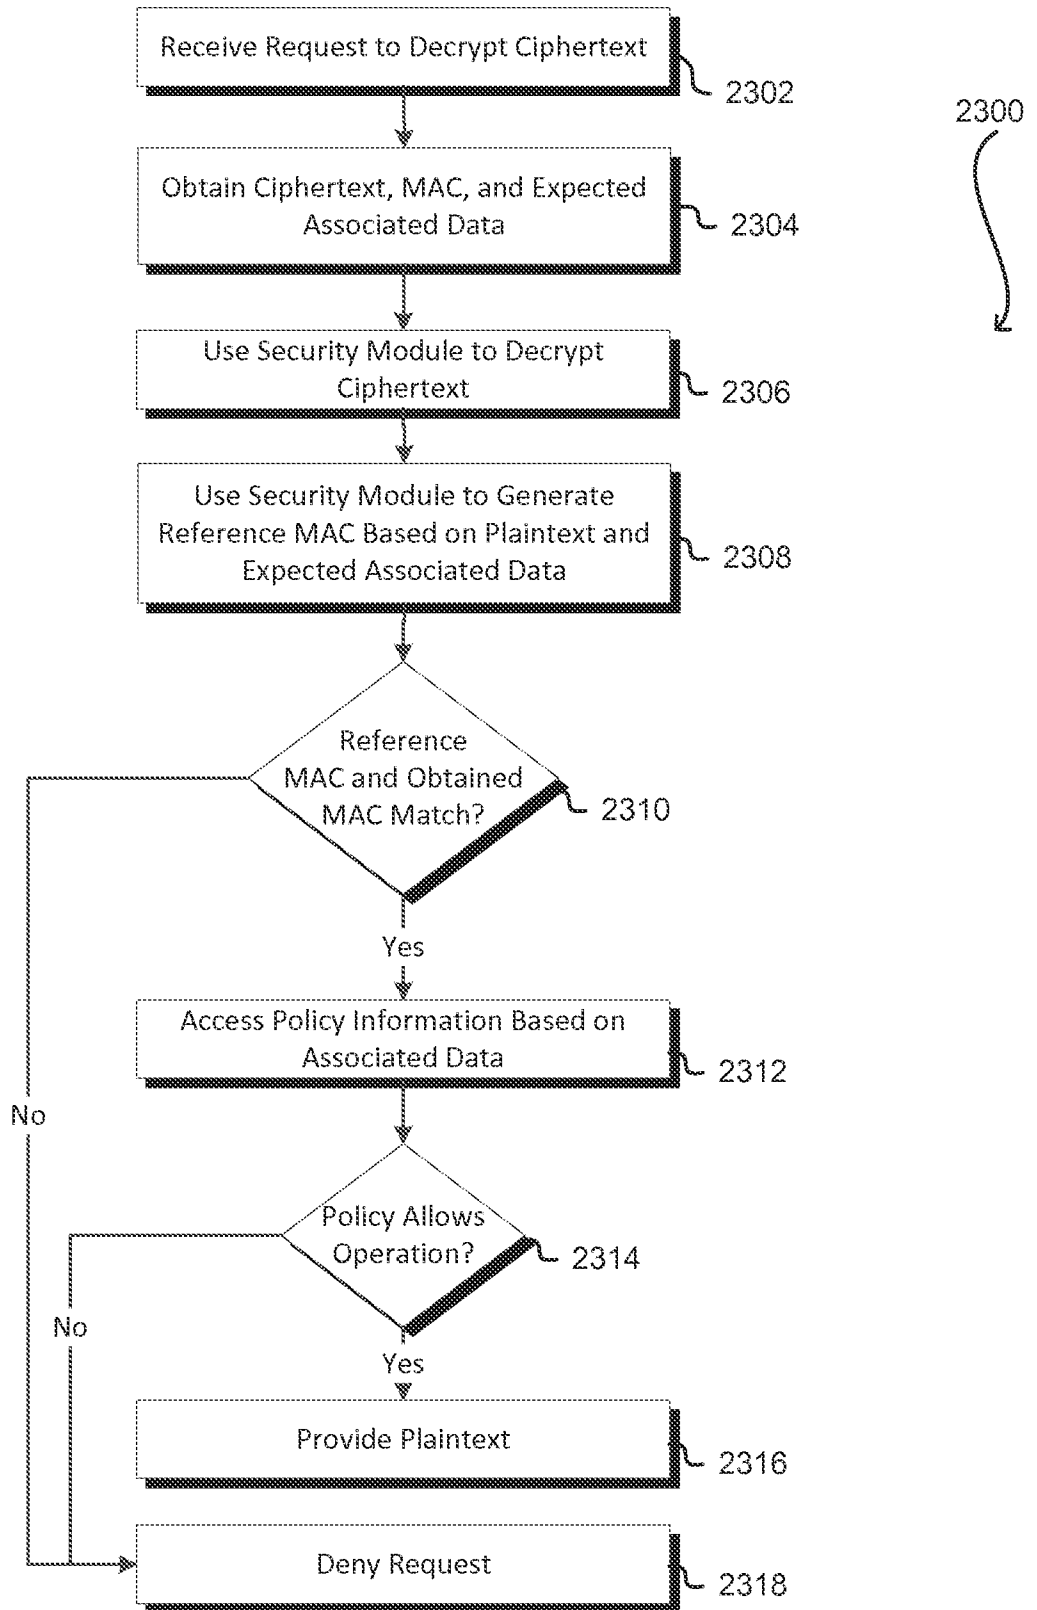


FIG. 23

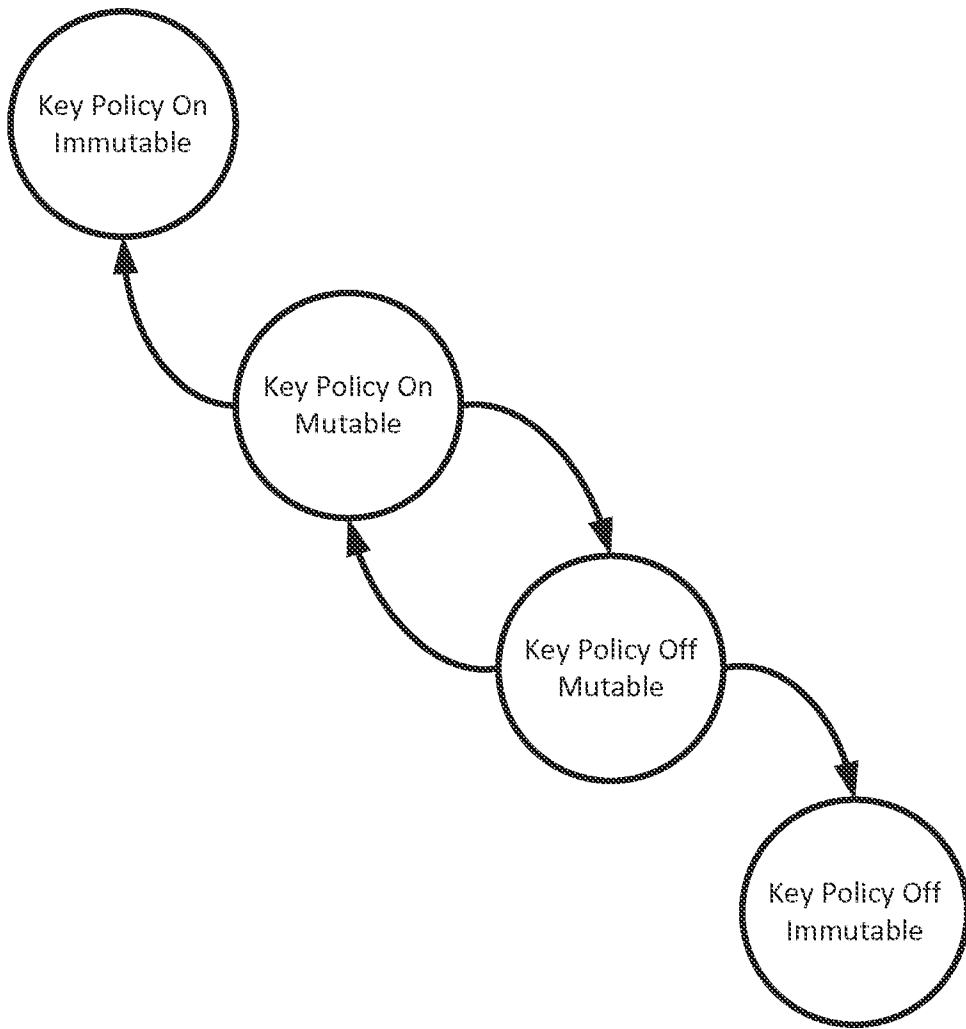


FIG. 24

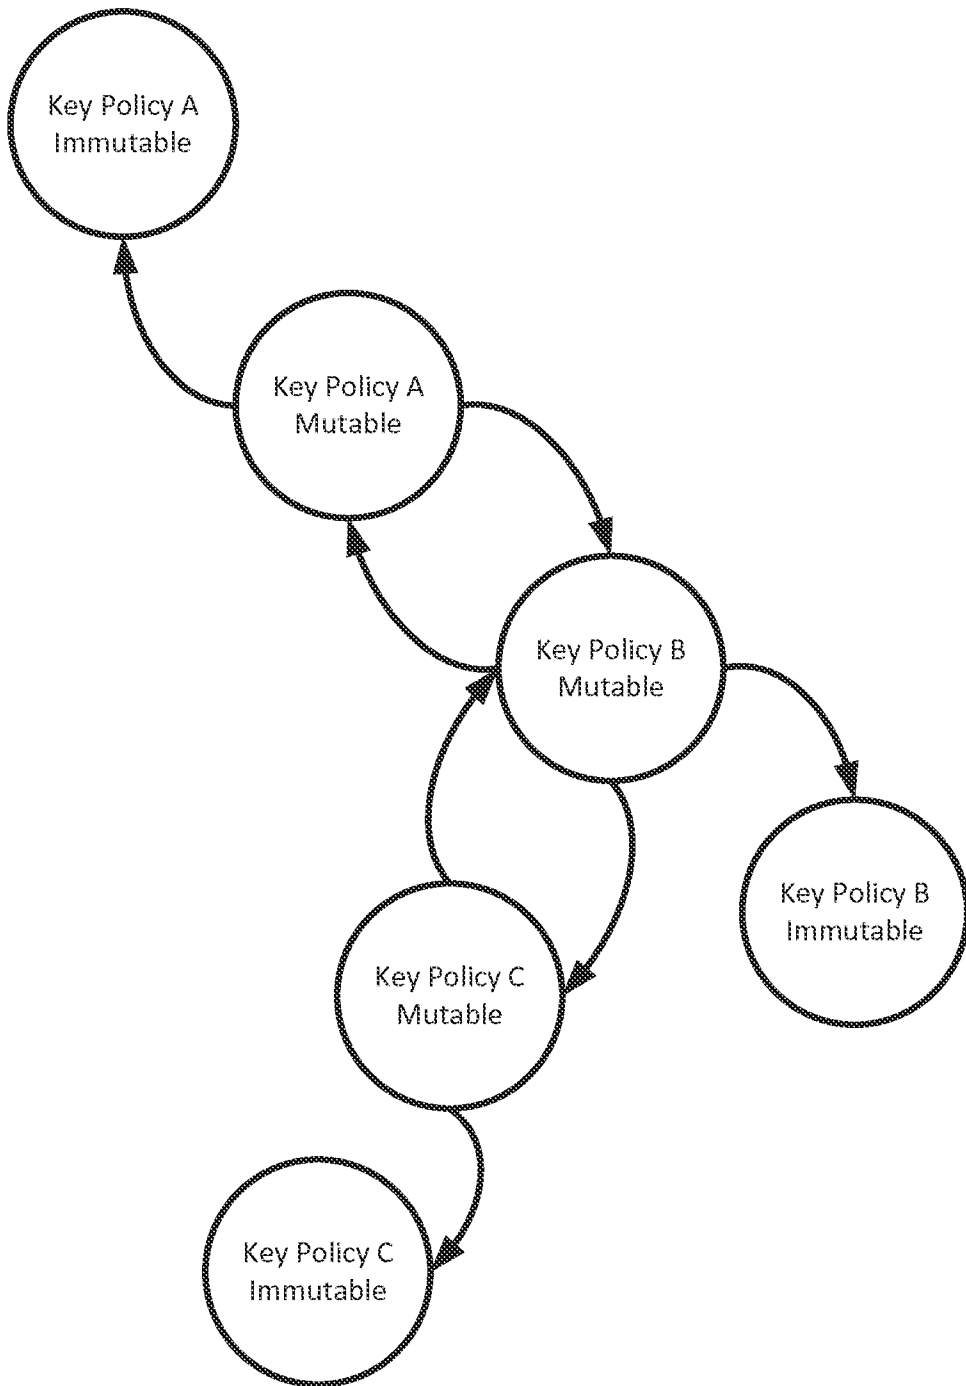


FIG. 25

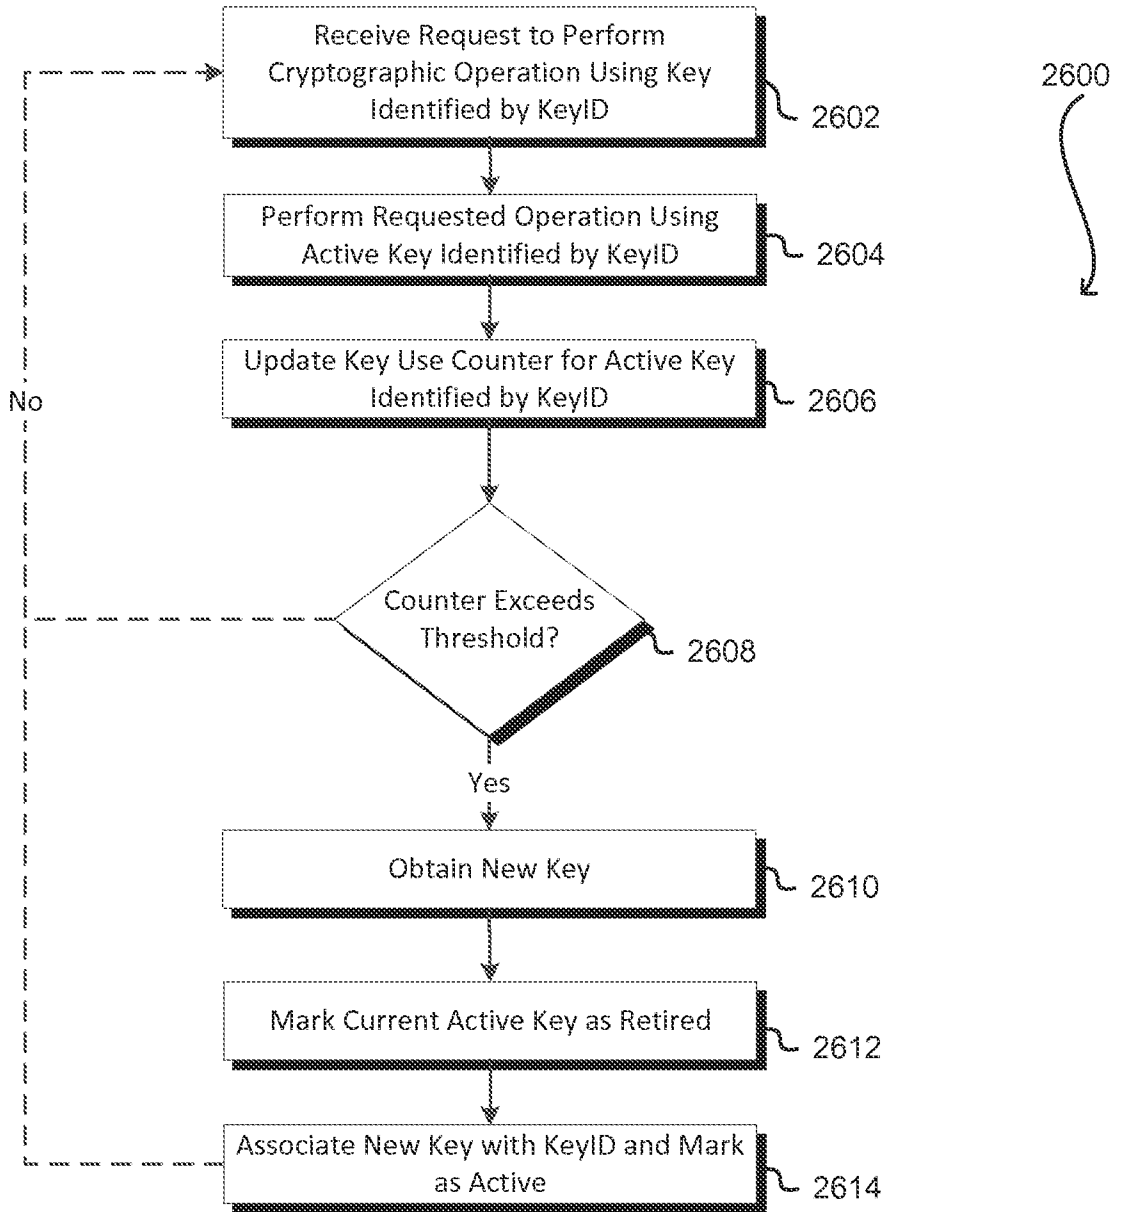


FIG. 26

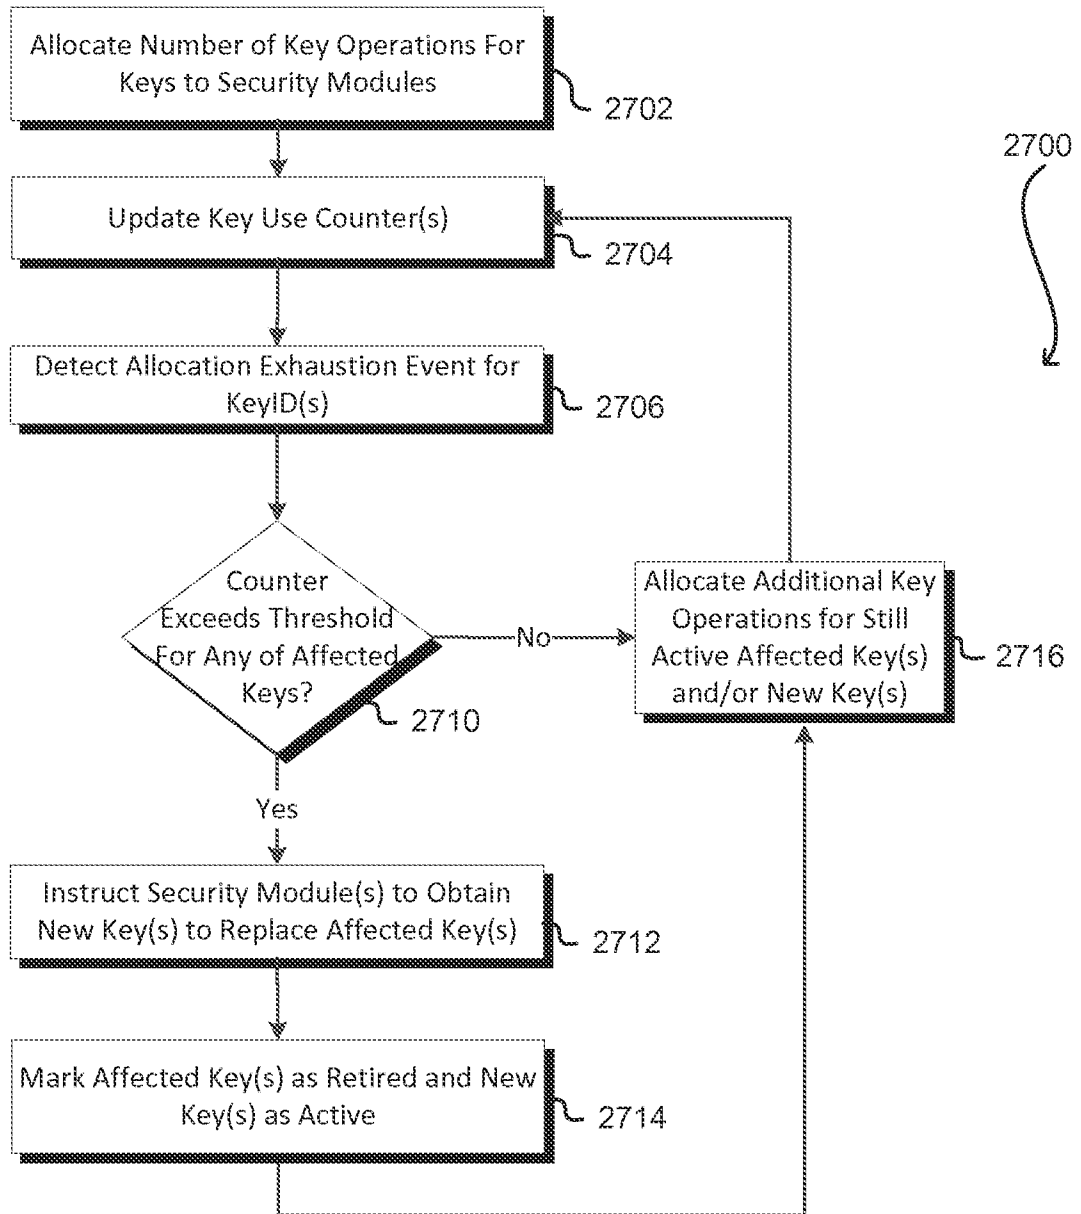


FIG. 27

KeyID	Key Version	Usability	Counter
⋮	⋮	⋮	⋮
31415926	1	Retired	4294967296
31415926	2	Retired	4294967296
31415926	3	Retired	4294967296
31415926	4	Active	1048576
31415927	1	Active	2097152
⋮	⋮	⋮	⋮

FIG. 28

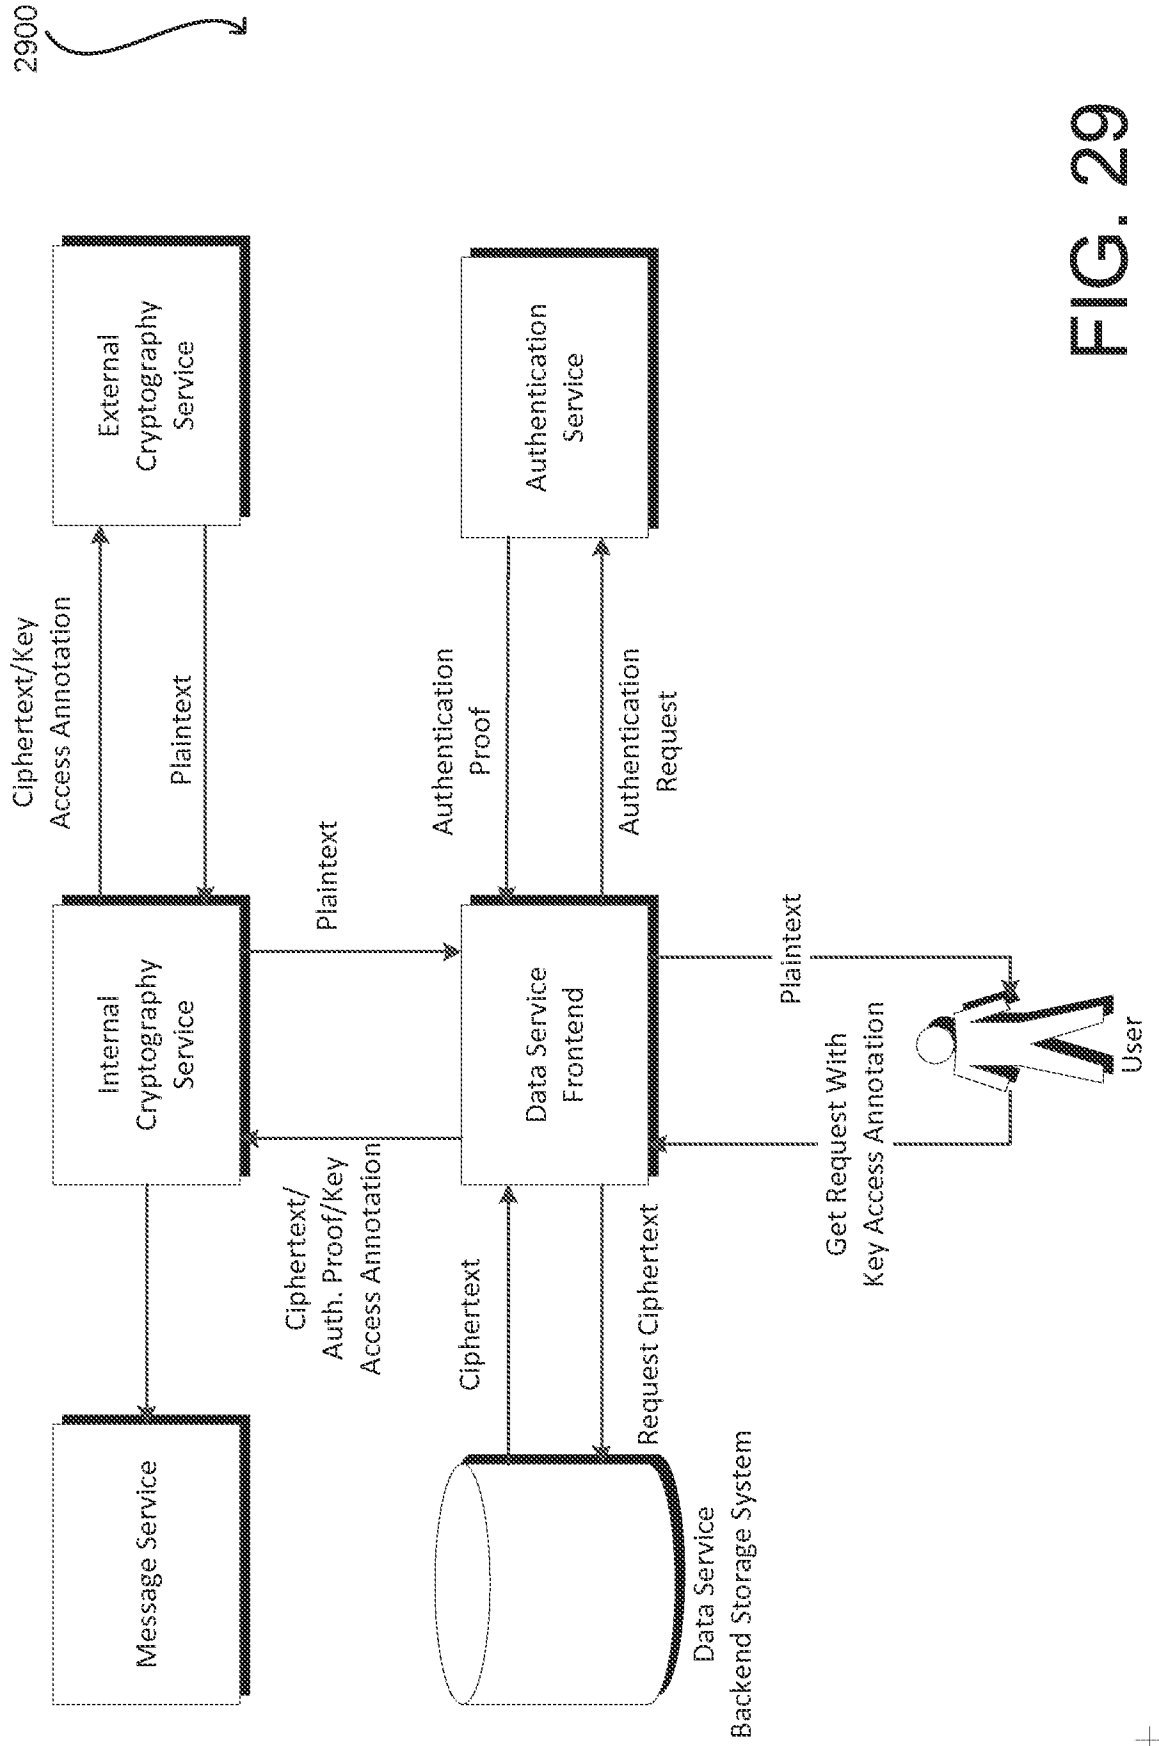


FIG. 29

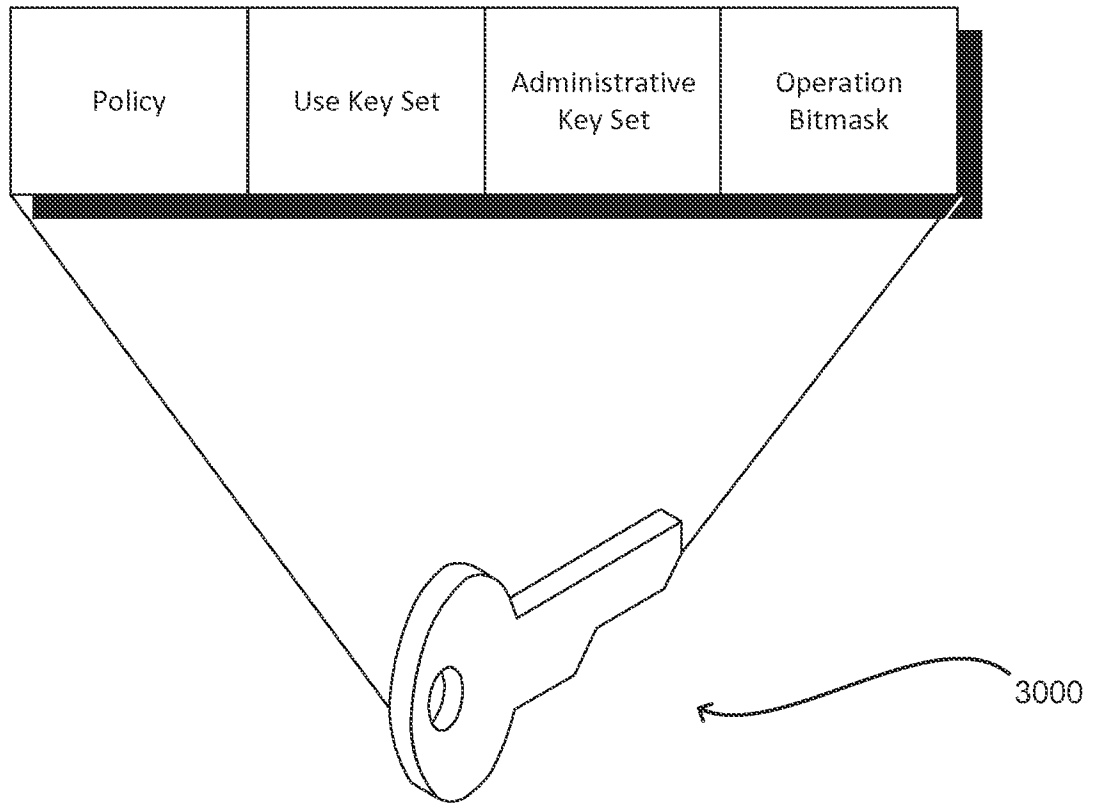


FIG. 30

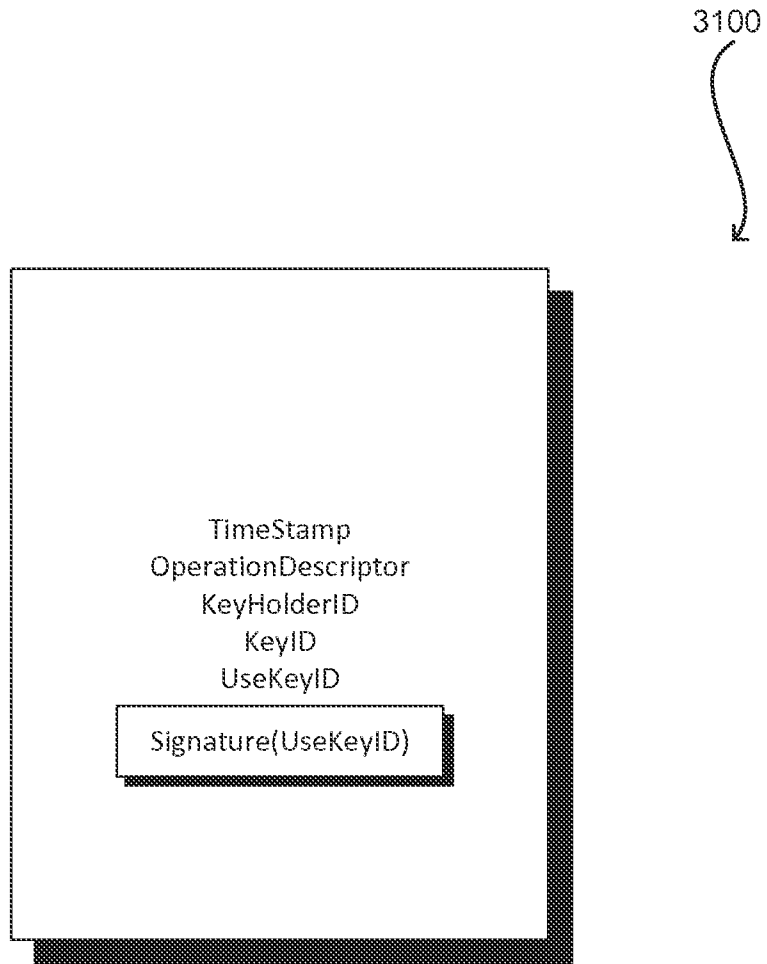


FIG. 31

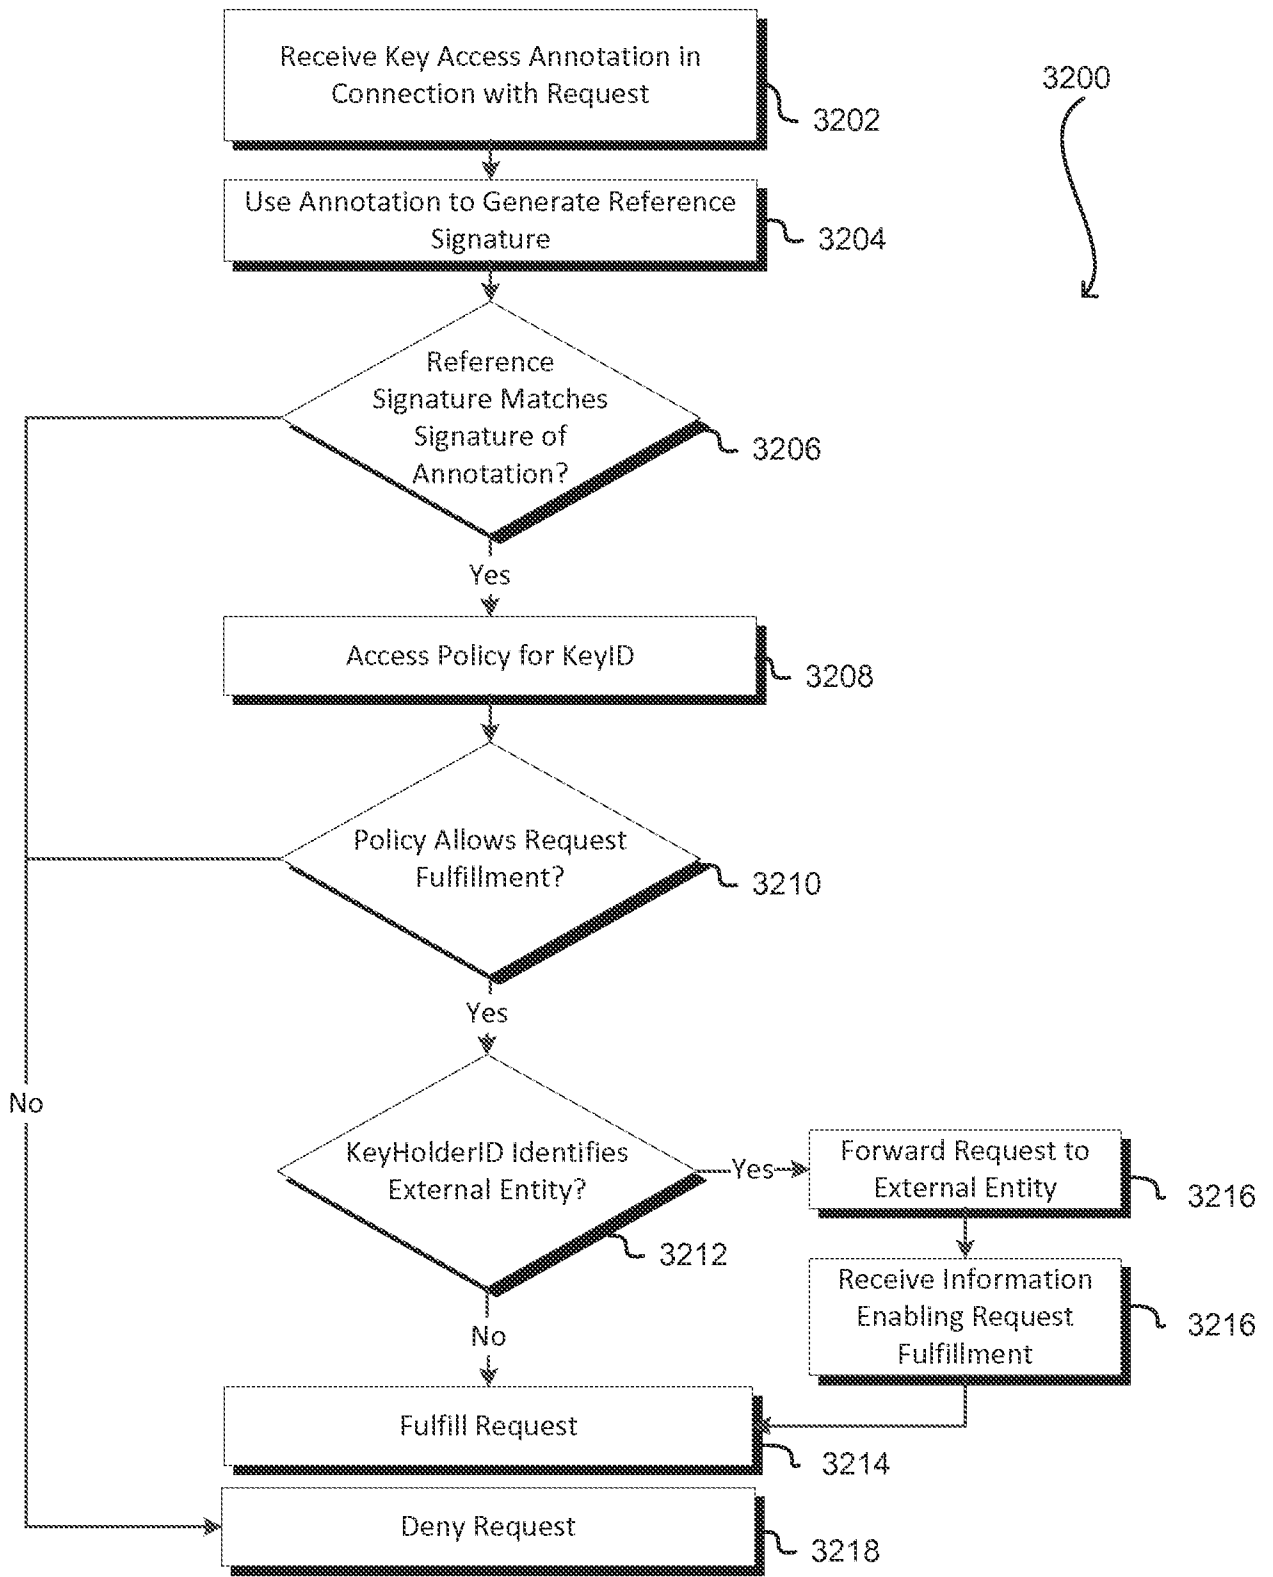


FIG. 32

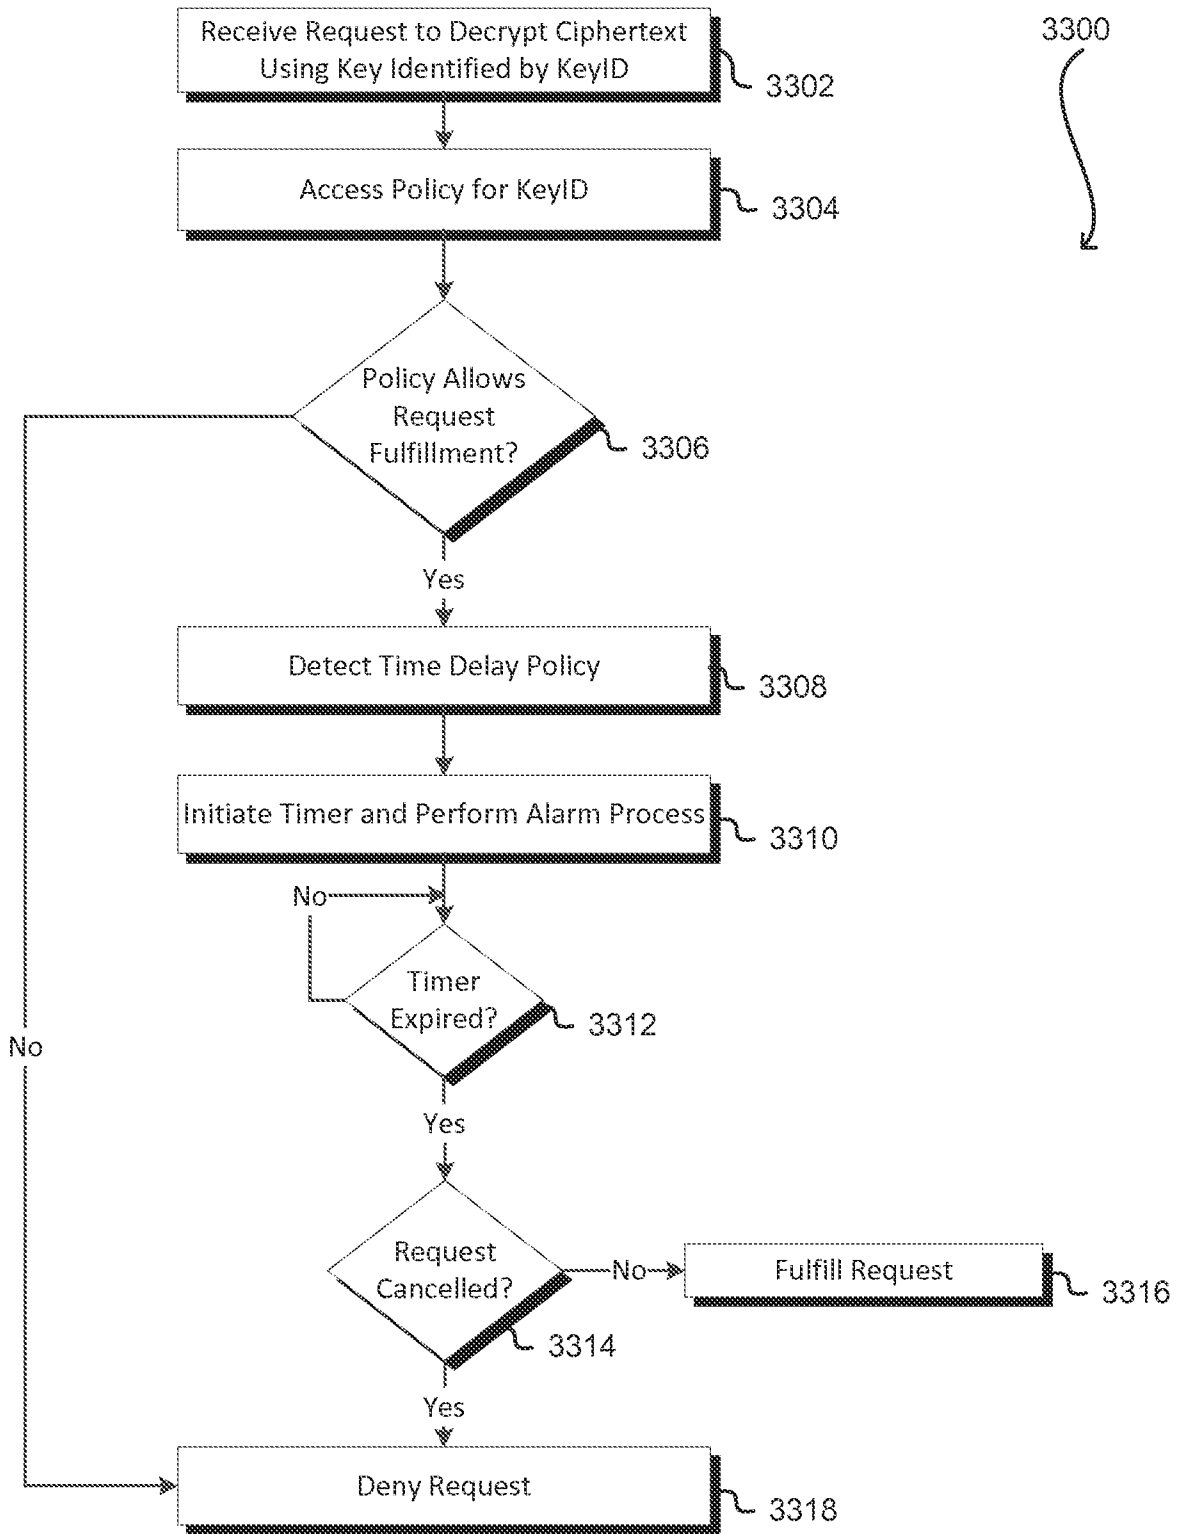


FIG. 33

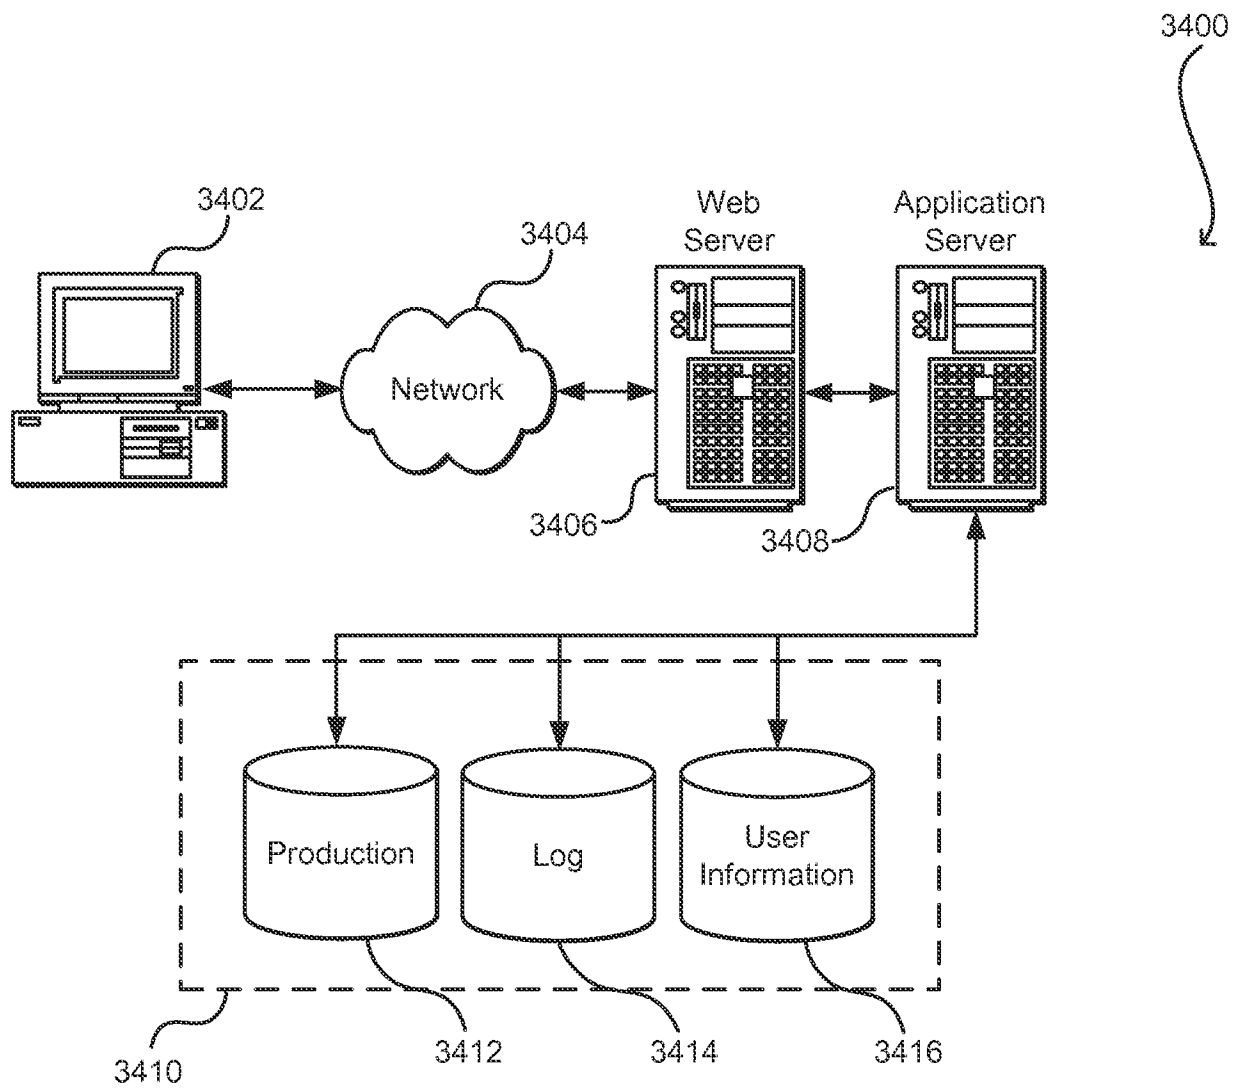


FIG. 34

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2014/015414

<p>A. CLASSIFICATION OF SUBJECT MATTER IPC(8) - G06F 21/00 (2014.01) USPC - 713/189 According to International Patent Classification (IPC) or to both national classification and IPC</p>																							
<p>B. FIELDS SEARCHED</p> <p>Minimum documentation searched (classification system followed by classification symbols) IPC(8) - G06F 12/14, 17/30, 21/00; H04L 9/00 (2014.01) USPC - 713/150, 713/168, 713/189, 713/500</p> <p>Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched CPC - G06F 17/30, 21/00; H04L 9/00 (2014.02)</p> <p>Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) PatBase, Google Patents, ProQuest</p>																							
<p>C. DOCUMENTS CONSIDERED TO BE RELEVANT</p> <table border="1"> <thead> <tr> <th>Category*</th> <th>Citation of document, with indication, where appropriate, of the relevant passages</th> <th>Relevant to claim No.</th> </tr> </thead> <tbody> <tr> <td>Y</td> <td>US 8,261,320 B1 (SERENYI et al) 04 September 2012 (04.09.2012) entire document.</td> <td>1-3, 9-11</td> </tr> <tr> <td>Y</td> <td>US 2007/0180153 A1 (CORNWELL et al) 02 August 2007 (02.08.2007) entire document.</td> <td>1-3, 9-11</td> </tr> <tr> <td>A</td> <td>US 2010/0008499 A1 (LEE et al) 14 January 2010 (14.01.2010) entire document.</td> <td>1-3, 9-11</td> </tr> <tr> <td>A</td> <td>US 2003/0172269 A1 (NEWCOMBE) 11 September 2003 (11.09.2003) entire document.</td> <td>1-3, 9-11</td> </tr> <tr> <td>A</td> <td>US 2004/0107345 A1 (BRANDT et al) 03 June 2004 (03.06.2004) entire document.</td> <td>1-3, 9-11</td> </tr> <tr> <td>A</td> <td>US 8,245,037 B1 (DURGIN et al) 14 August 2012 (14.08.2012) entire document.</td> <td>1-3, 9-11</td> </tr> </tbody> </table>			Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.	Y	US 8,261,320 B1 (SERENYI et al) 04 September 2012 (04.09.2012) entire document.	1-3, 9-11	Y	US 2007/0180153 A1 (CORNWELL et al) 02 August 2007 (02.08.2007) entire document.	1-3, 9-11	A	US 2010/0008499 A1 (LEE et al) 14 January 2010 (14.01.2010) entire document.	1-3, 9-11	A	US 2003/0172269 A1 (NEWCOMBE) 11 September 2003 (11.09.2003) entire document.	1-3, 9-11	A	US 2004/0107345 A1 (BRANDT et al) 03 June 2004 (03.06.2004) entire document.	1-3, 9-11	A	US 8,245,037 B1 (DURGIN et al) 14 August 2012 (14.08.2012) entire document.	1-3, 9-11
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.																					
Y	US 8,261,320 B1 (SERENYI et al) 04 September 2012 (04.09.2012) entire document.	1-3, 9-11																					
Y	US 2007/0180153 A1 (CORNWELL et al) 02 August 2007 (02.08.2007) entire document.	1-3, 9-11																					
A	US 2010/0008499 A1 (LEE et al) 14 January 2010 (14.01.2010) entire document.	1-3, 9-11																					
A	US 2003/0172269 A1 (NEWCOMBE) 11 September 2003 (11.09.2003) entire document.	1-3, 9-11																					
A	US 2004/0107345 A1 (BRANDT et al) 03 June 2004 (03.06.2004) entire document.	1-3, 9-11																					
A	US 8,245,037 B1 (DURGIN et al) 14 August 2012 (14.08.2012) entire document.	1-3, 9-11																					
<p><input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/></p>																							
<p>* Special categories of cited documents:</p> <table border="0"> <tr> <td>“A” document defining the general state of the art which is not considered to be of particular relevance</td> <td>“T” later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</td> </tr> <tr> <td>“E” earlier application or patent but published on or after the international filing date</td> <td>“X” document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone</td> </tr> <tr> <td>“L” document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</td> <td>“Y” document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art</td> </tr> <tr> <td>“O” document referring to an oral disclosure, use, exhibition or other means</td> <td>“&” document member of the same patent family</td> </tr> <tr> <td>“P” document published prior to the international filing date but later than the priority date claimed</td> <td></td> </tr> </table>			“A” document defining the general state of the art which is not considered to be of particular relevance	“T” later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention	“E” earlier application or patent but published on or after the international filing date	“X” document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone	“L” document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	“Y” document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art	“O” document referring to an oral disclosure, use, exhibition or other means	“&” document member of the same patent family	“P” document published prior to the international filing date but later than the priority date claimed												
“A” document defining the general state of the art which is not considered to be of particular relevance	“T” later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention																						
“E” earlier application or patent but published on or after the international filing date	“X” document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone																						
“L” document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	“Y” document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art																						
“O” document referring to an oral disclosure, use, exhibition or other means	“&” document member of the same patent family																						
“P” document published prior to the international filing date but later than the priority date claimed																							
<p>Date of the actual completion of the international search 12 May 2014</p>		<p>Date of mailing of the international search report 30 MAY 2014</p>																					
<p>Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201</p>		<p>Authorized officer: Blaine R. Copenheaver PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774</p>																					

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2014/015414

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. Claims Nos.: 4-8,12-15
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.