




(86) Date de dépôt PCT/PCT Filing Date: 2015/04/30
(87) Date publication PCT/PCT Publication Date: 2015/11/05
(45) Date de délivrance/Issue Date: 2023/02/28
(85) Entrée phase nationale/National Entry: 2016/10/17
(86) N° demande PCT/PCT Application No.: US 2015/028393
(87) N° publication PCT/PCT Publication No.: 2015/168350
(30) Priorité/Priority: 2014/04/30 (US14/265,930)

(51) Cl.Int./Int.Cl. *H04L 67/1396* (2022.01),
G06F 21/00 (2013.01), *H04L 9/40* (2022.01)
(72) Inventeurs/Inventors:
CABRERA, LUIS FELIPE, US;
LIETZ, M. SHANNON, US
(73) Propriétaire/Owner:
INTUIT INC., US
(74) Agent: OSLER, HOSKIN & HARCOURT LLP

(54) Titre : PROCEDE ET APPAREIL DE GESTION DE SECRETS MULTI-LOCATAIRES
(54) Title: METHOD AND APPARATUS FOR MULTI-TENANCY SECRETS MANAGEMENT



(57) **Abrégé/Abstract:**

A service provider computing environment includes a service provider secrets policy. A service provider computing device receives tenant secrets policies from tenants. The tenants are tenants of multi-tenant assets of a service provider. The service provider computing environment determines if the tenant secrets policies satisfy the requirements of the service provider secrets policy. If the tenant secrets policies satisfy the requirements of the service provider secrets policy, the service provider computing environment allows the tenant secrets policies to be applied to tenant data or information in the multi-tenant assets.

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

(19) World Intellectual Property
Organization
International Bureau



(10) International Publication Number
WO 2015/168350 A1

(43) International Publication Date
5 November 2015 (05.11.2015)

- (51) **International Patent Classification:**
G06F 21/31 (2013.01) *H04L 9/14* (2006.01)
- (21) **International Application Number:**
PCT/US2015/028393
- (22) **International Filing Date:**
30 April 2015 (30.04.2015)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
14/265,930 30 April 2014 (30.04.2014) US
- (71) **Applicant:** INTUIT INC. [US/US]; 2700 Coast Avenue, Mountain View, California 94043 (US).
- (72) **Inventors:** CABRERA, Luis Felipe; c/o Intuit Inc., 2700 Coast Avenue, Mountain View, California 94043 (US). LI-ETZ, M. Shannon; c/o Intuit Inc., 2700 Coast Avenue, Mountain View, California 94043 (US).
- (74) **Agents:** MCKAY, Philip et al.; Hawley Troxell, P.O. Box 1617, Boise, Idaho 83701-1617 (US).
- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, 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))

Published:

- with international search report (Art. 21(3))

(54) **Title:** METHOD AND APPARATUS FOR MULTI-TENANCY SECRETS MANAGEMENT




FIG. 1

(57) **Abstract:** A service provider computing environment includes a service provider secrets policy. A service provider computing device receives tenant secrets policies from tenants. The tenants are tenants of multi-tenant assets of a service provider. The service provider computing environment determines of the tenant secrets policies satisfy the requirements of the service provider secrets policy. If the tenant secrets policies satisfy the requirements of the service provider secrets policy, the service provider computing environment allows the tenant secrets policies to be applied to tenant data or information in the multi-tenant assets.

METHOD AND APPARATUS FOR MULTI-TENANCY SECRETS MANAGEMENT

BACKGROUND

[0001] A service provider of computing systems or computing environments or computing system assets can allocate portions of a computing system asset to multiple tenants. This type of resource sharing is called multi-tenant or multi-tenancy sharing. By allocating portions of the computing system asset to multiple tenants to establish multi-tenant assets, the service provider can allow the tenants to share the financial burden associated with maintaining computing environments or services while handling much of the logistical burdens associated with maintaining the computing environments.

[0002] To maintain secure multi-tenant assets or multi-tenant computing environments, the service provider can set forth data security rules or policies that must be adhered to by the tenants, for example, as part of a service agreement. However, the service provider's data security rules or policies may differ from the data security rules or policies of the tenants.

[0003] What is needed is a method and system that allows for the management of tenant secrets policies in a multi-tenant computing environment.

SUMMARY

[0004] In accordance with one embodiment, a method and system for managing secrets of tenants in a multi-tenant computing environment includes maintaining a service provider secrets policy, receiving a request to apply a tenant secrets policy to the multi-tenant computing environment, comparing the tenant secrets policy to the service provider secrets policy, and authorizing the application of the tenant secrets policy to the multi-tenant computing environment if the tenant secrets policy satisfies the requirements of the service provider secrets policy. As used herein, an information secrets policy or secrets policy may interchangeably be referred to as an information security policy or a security policy.

[0005] In one embodiment, a service provider computing system receives secrets from tenants and maintains logs of the use, copying, deletion, receipt, transfer, and transmission of the

tenants' secrets. In one embodiment, the service provider computing system provides the logs to each of the tenants for each tenants' secrets.

[0006] In one embodiment, the service provider computing system maintains a cross-tenant secrets policy which includes rules that are associated with sharing one tenant's secrets with another tenant to support multi-tenant operations of multi-tenant assets, such as shared applications, shared servers, shared virtual machines, shared memory devices, or the like.

[0007] In one embodiment, the service provider computing system applies authorized tenant secrets policies to the multi-tenant assets. In other embodiments, a tenant computing system applies authorized tenant secrets policies to multi-tenant assets.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] FIG. 1 is a block diagram of a hardware architecture for implementing multi-tenant secret management, in accordance with one embodiment;

[0009] FIG. 2 is a flow diagram depicting a method for managing multi-tenant secrets policies in a multi-tenant computing environment, in accordance with one embodiment;

[0010] FIG. 3 is a flow diagram depicting a method for managing cross-tenant policies in a multi-tenant computing environment, in accordance with one embodiment; and

[0011] FIG. 4 is a flow diagram depicting a method for providing logging services to tenants in a multi-tenant computing environment, in accordance with one embodiment.

[0012] Common reference numerals are used throughout the FIG.s and the detailed description to indicate like elements. One skilled in the art will readily recognize that the above FIG.s are examples and that other architectures, modes of operation, orders of operation, and elements/functions can be provided and implemented without departing from the characteristics and features of the invention, as set forth in the claims.

DETAILED DESCRIPTION

[0013] Embodiments will now be discussed with reference to the accompanying FIG.s, which depict one or more exemplary embodiments. Embodiments may be implemented in many different forms and should not be construed as limited to the embodiments set forth herein, shown in the FIG.s, and/or described below. Rather, these exemplary embodiments are provided to allow a complete disclosure that conveys the principles of the invention, as set forth in the claims, to those of skill in the art.

[0014] The INTRODUCTORY SYSTEM, HARDWARE ARCHITECTURE, and PROCESS sections herein include systems and processes suitable for incorporating various embodiments of methods and systems for managing tenant secrets in multi-tenant computing environments, according to various embodiments.

INTRODUCTORY SYSTEM

[0015] Herein, the term “production environment” includes the various components, or assets, used to deploy, implement, access, and use, a given application as that application is intended to be used. In various embodiments, production environments include multiple assets that are combined, communicatively coupled, virtually and/or physically connected, and/or associated with one another, to provide the production environment implementing the application.

[0016] As specific illustrative examples, the assets making up a given production environment can include, but are not limited to, one or more computing environments used to implement the application in the production environment such as a data center, a cloud computing environment, a dedicated hosting environment, and/or one or more other computing environments in which one or more assets used by the application in the production environment are implemented; one or more computing systems or computing entities used to implement the application in the production environment; one or more virtual assets used to implement the application in the production environment; one or more supervisory or control systems, such as hypervisors, or other monitoring and management systems, used to monitor and control assets and/or components of the production environment; one or more communications channels for sending and receiving data used to implement the application in the production environment; one or more access control systems for limiting access to various components of the production environment, such as firewalls and gateways; one or more traffic and/or routing systems used to direct, control, and/or buffer, data traffic to components of the production environment, such as routers and switches; one or more communications endpoint proxy systems used to buffer, process, and/or direct data traffic, such as load balancers or buffers; one or more secure communication protocols and/or endpoints used to encrypt/decrypt data, such as Secure Sockets Layer (SSL) protocols, used to implement the application in the production environment; one or more databases used to store data in the production environment; one or more internal or external services used to implement the application in the production environment; one or more backend systems, such as backend servers or other hardware used to process data and implement

the application in the production environment; one or more software systems used to implement the application in the production environment; and/or any other assets/components making up an actual production environment in which an application is deployed, implemented, accessed, and run, e.g., operated, as discussed herein, and/or as known in the art at the time of filing, and/or as developed after the time of filing.

[0017] As used herein, the terms “computing system”, “computing device”, and “computing entity”, include, but are not limited to, a virtual asset; a server computing system; a workstation; a desktop computing system; a mobile computing system, including, but not limited to, smart phones, portable devices, and/or devices worn or carried by a user; a database system or storage cluster; a switching system; a router; any hardware system; any communications system; any form of proxy system; a gateway system; a firewall system; a load balancing system; or any device, subsystem, or mechanism that includes components that can execute all, or part, of any one of the processes and/or operations as described herein.

[0018] In addition, as used herein, the terms computing system and computing entity, can denote, but are not limited to, systems made up of multiple: virtual assets; server computing systems; workstations; desktop computing systems; mobile computing systems; database systems or storage clusters; switching systems; routers; hardware systems; communications systems; proxy systems; gateway systems; firewall systems; load balancing systems; or any devices that can be used to perform the processes and/or operations as described herein.

[0019] As used herein, the term “computing environment” includes, but is not limited to, a logical or physical grouping of connected or networked computing systems and/or virtual assets using the same infrastructure and systems such as, but not limited to, hardware systems, software systems, and networking/communications systems. Typically, computing environments are either known environments, e.g., “trusted” environments, or unknown, e.g., “untrusted” environments. Typically, trusted computing environments are those where the assets, infrastructure, communication and networking systems, and security systems associated with the computing systems and/or virtual assets making up the trusted computing environment, are either under the control of, or known to, a party. Examples of trusted computing environments include the assets and components making up data centers associated with, and/or controlled by, an application and/or any computing systems and/or virtual assets, and/or networks of computing systems and/or virtual assets, associated with, known by, and/or controlled by, an application.

[0020] In contrast, unknown, or untrusted computing environments are environments and systems where the assets, components, infrastructure, communication and networking systems, and security systems implemented and associated with the computing systems and/or virtual assets making up the untrusted computing environment, are not under the control of, and/or are not known by, a party, and/or are dynamically configured with new elements capable of being added that are unknown to the party. Examples of untrusted computing environments include, but are not limited to, public networks, such as the Internet, various cloud-based computing environments, and various other forms of distributed computing systems.

[0021] In various embodiments, each computing environment includes allocated assets and virtual assets associated with, and controlled or used to create, and/or deploy, and/or operate an application.

[0022] It is often the case that to create, and/or deploy, and/or operate, an application data must be transferred between a first computing environment that is an untrusted computing environment and a trusted computing environment. However, in other situations a party may wish to transfer data between two trusted computing environments, and/or two untrusted computing environments.

[0023] In various embodiments, one or more cloud computing environments are used to create, and/or deploy, and/or operate an application that can be any form of cloud computing environment, such as, but not limited to, a public cloud; a private cloud; a virtual private network (VPN); a subnet; a Virtual Private Cloud (VPC); a sub-net or any security/communications grouping; or any other cloud-based infrastructure, sub-structure, or architecture, as discussed herein, and/or as known in the art at the time of filing, and/or as developed after the time of filing.

[0024] In many cases, a given application or service may utilize, and interface with, multiple cloud computing environments, such as multiple VPCs, in the course of being created, and/or deployed, and/or operated.

[0025] As used herein, the term “virtual asset” includes any virtualized entity or resource, and/or virtualized part of an actual, or “bare metal” entity. In various embodiments, the virtual assets can be, but are not limited to, virtual machines, virtual servers, and instances implemented in a cloud computing environment; databases associated with a cloud computing environment, and/or implemented in a cloud computing environment; services associated with, and/or delivered through, a cloud computing environment; communications systems used with, part of, or provided through, a cloud computing environment; and/or any other virtualized assets

and/or sub-systems of “bare metal” physical devices such as mobile devices, remote sensors, laptops, desktops, point-of-sale devices, ATMs, electronic voting machines, etc., located within a data center, within a cloud computing environment, and/or any other physical or logical location, as discussed herein, and/or as known/available in the art at the time of filing, and/or as developed/made available after the time of filing.

[0026] In various embodiments, any, or all, of the assets making up a given production environment discussed herein, and/or as known in the art at the time of filing, and/or as developed after the time of filing, can be implemented as virtual assets.

[0027] Typically, virtual assets are created, or instantiated, using steps, instructions, processes, code, or “recipes” referred to herein as “virtual asset creation templates.” Typically, virtual assets that have the same, or similar, operational parameters are created using the same or similar “virtual asset creation templates.”

[0028] Examples of virtual asset creation templates include, but are not limited to, any tool and/or system for creating and managing a collection of related cloud resources. Illustrative examples of such a virtual asset creation template are any of the cloud formation templates/tools provided by Amazon Web Service (AWS), Rack Space, Joyent, and/or any other of the numerous cloud based infrastructure providers.

[0029] Other examples of virtual asset creation templates include, but are not limited to, any configuration management tool associated with, and/or used to create, virtual assets. One specific illustrative example of such a virtual asset creation template is a cookbook or recipe tool such as a Chef Recipe or system or any other fundamental element, or set of elements, used to override the default settings on a node within an infrastructure or architecture.

[0030] Other examples of virtual asset creation templates include, but are not limited to, any virtual appliance used to instantiate virtual assets. One specific illustrative example of such a virtual asset creation template is an Amazon Machine Image (AMI), and/or similar functionality provided by Amazon Web Service (AWS), Rack Space, Joyent, and/or any other of the numerous cloud based infrastructure providers.

[0031] Other examples of virtual asset creation templates include, but are not limited to, any appliance, or tool, or system, or framework, used to instantiate virtual assets as discussed herein, and/or as known/available in the art at the time of filing, and/or as developed/made available after the time of filing.

[0032] Herein virtual assets that have the same, or similar, operational parameters and are created by the same or similar virtual asset creation template are generically referred to as

virtual assets of the same “class.” Examples of virtual asset classes include, but are not limited to, virtual machine classes; virtual server classes; virtual database or data store classes; self-monitoring virtual assets including specific types of instances instantiated in a cloud environment; application development process classes; and application classes.

[0033] In one embodiment, two or more assets, such as computing systems and/or virtual assets, and/or two or more computing environments, are connected by one or more communications channels including but not limited to, Secure Sockets Layer communications channels and various other secure communications channels, and/or distributed computing system networks, such as, but not limited to: a public cloud; a private cloud; a virtual private network (VPN); a subnet; any general network, communications network, or general network/communications network system; a combination of different network types; a public network; a private network; a satellite network; a cable network; or any other network capable of allowing communication between two or more assets, computing systems, and/or virtual assets, as discussed herein, and/or available or known at the time of filing, and/or as developed after the time of filing.

[0034] As used herein, the term “network” includes, but is not limited to, any network or network system such as, but not limited to, a peer-to-peer network, a hybrid peer-to-peer network, a Local Area Network (LAN), a Wide Area Network (WAN), a public network, such as the Internet, a private network, a cellular network, any general network, communications network, or general network/communications network system; a wireless network; a wired network; a wireless and wired combination network; a satellite network; a cable network; any combination of different network types; or any other system capable of allowing communication between two or more assets, virtual assets, and/or computing systems, whether available or known at the time of filing or as later developed.

[0035] As used herein, the term “user” includes, but is not limited to, any party, parties, entity, and/or entities using, or otherwise interacting with any of the methods or systems discussed herein. For instance, in various embodiments, a user can be, but is not limited to, a person, a commercial entity, an application, a service, and/or a computing system.

HARDWARE ARCHITECTURE

[0036] FIG. 1 illustrates a functional diagram of a production environment 1 that is configured to enable and perform multi-tenant secrets management for one or more multi-tenant assets in a multi-tenant (or multi-tenancy) computing environment, according to one

embodiment. Multi-tenant computing environments enable service providers to deliver and/or provide services and resources to tenants or customers in quantities, speeds, and with resources that are flexibly and dynamically catered or customized to the specifications of the tenant/customer. The terms “tenant” and “customer” are used interchangeably herein and represent, according to various embodiments, a purchaser or recipient of computing infrastructure or other computing services from a service provider. A multi-tenant computing environment might include multi-tenant assets, such as an entire data center, one or more servers, one or more storage devices, or one or more applications that are shared by multiple customers. According to the particular needs or intended uses of the customer, it may be financially and/or logistically more feasible or profitable for a customer to rent, lease, or otherwise temporarily receive control over merely some, in contrast to all, assets available in a computing environment that is made available by a service provider. In exchange for lower cost and the dynamic availability of resources offered by the service provider, a customer or tenant accepts security risks associated with at least partially sharing the multi-tenant assets with other tenants. In some implementations, one tenant’s information may be physically separated from another tenant’s information, e.g., by being hosted from different servers or by being stored on different storage devices. However, in other implementations one tenant’s information may merely be logically separated from another tenant’s information and may therefore be susceptible to being intentionally or inadvertently compromised to other tenants. Service providers therefore allow tenants to apply tenant secrets or security measures, such as password protection and/or data encryption, to the tenants’ information and/or to portions of the multi-tenant assets that are allocated to the tenants.

[0037] At times, a service provider’s minimum requirements for secrets policies, a.k.a., security policies, may differ from the tenants’ secrets policies, and a tenant’s secrets policies may differ from another tenant’s secrets policies. The system of production environment 1, disclosed herein, enables and/or provides management of multi-tenant secrets and multi-tenant secrets policies in, for example, multi-tenant computing environments, according to one embodiment.

[0038] The production environment 1 provides multi-tenant secrets management using one or more computing environments described hereinafter, according to various embodiments. Although each of the computing environments of production environment 1 are illustrated as distinct components, one or more of the computing environments may be physically and/or functionally combined. Alternatively, as will be appreciated by one of ordinary skill in the art,

in some embodiments the functions that are grouped together by the illustrated computing environments may be physically and/or functionally separated into additional computing environments that are not presently shown. According to one embodiment, the production environment 1 includes a tenant computing environment 10, a tenant computing environment 11, a computing environment 12, a multi-tenant computing environment 13, and a service provider computing environment 14. Although two tenant computing environments 10 and 11 are included, more tenant computing environments may be incorporated into the production environment 1, according to other embodiments.

[0039] The tenant computing environment 10 represents a computing environment of a tenant or customer of the service provider, according to one embodiment. The tenant computing environment 10 may be owned, operated, or be under the direction or control of any one of a number of types of customers. According to various embodiments, a tenant or customer includes, but is not limited to, one or more individuals, families, small/medium/large businesses, educational institutions, or other for-profit or non-profit organizations. In one embodiment, the tenant computing environment 10 includes a tenant secrets database 100, a tenant secrets policy 101, and a policy execution engine 102.

[0040] The tenant secrets database 100 stores, maintains, organizes, sorts, and/or hosts the secrets 103 of a tenant, e.g., a first tenant, according to one embodiment. The secrets 103 of a tenant include, but are not limited to, passwords, codes, encryption keys, digital certificates, and other security measures that are currently available or that will become available to enable a tenant to prevent unauthorized access, viewing, copying, distribution, or other use of a tenant's data or information. According to one embodiment, the secrets 103 are, at least partially, means for securing tenant information or data. According to another embodiment, the tenant secrets database 100 also stores one or more copies of secrets data, which is the information or data of the tenant to which the secrets 103 are applied in order to protect the information or data from unauthorized access or use.

[0041] The tenant secrets policy 101 includes rules for governing the use, distribution, creation, and/or destruction of the secrets 103 of the tenant, according to one embodiment. The tenant secrets policy 101 can define the length, e.g., 800 bit encryption keys, and type, e.g., symmetric public-key infrastructure encryption keys, of secrets or security measures used. The tenant secrets policy 101 can also define the frequency with which secrets are rotated, cycled, or renewed for any particular group or set of information to which the secrets are applied. In one particular example, the tenant secrets policy 101 can include a 3 month rule/requirement to

rotate encryption keys used for protecting tenant data stored or used in the multi-tenant computing environment 13. In other implementations, the rotation rule for the encryption keys can be monthly, bi-monthly, quarterly, semi-annually, annually, or the like, and can apply to other or all secrets 103 maintained by the tenant secrets database 100. In one embodiment, the tenant secrets policy 101 can specify different lengths and/or types of secret rules, for use in various scenarios. For example, the tenant secrets policy 101 may specify a first rule for multi-tenancy assets, a second rule for dedicated or non-shared assets, a third rule for cloud storage, a fourth rule for private network storage, and so forth.

[0042] In yet other embodiments, the tenant secrets policy 101 can provide exceptions or limited provisions for sharing secrets 103 of a tenant with other tenants. For example, if a first tenant and a second tenant mutually provide software, an application, analytical services, or other services to one or more users, the tenant secrets policy 101 can include rules for authorizing or rejecting requests for access to the secrets 103 by other tenants. In a more specific example, the tenant secrets policy 101 may include a rule that authorizes distribution of a password or encryption key of a first tenant to a second tenant if the second tenant obtains permission from a service provider that provides services to both the first and second tenants. As another specific example, the tenant secrets policy 101 may include a rule that authorizes distribution of the password or encryption key of a first tenant to a second tenant if the second tenant mutually subscribes to one or more particular services, e.g. cloud storage or other web-based services, from a particular service provider. As can be appreciated by those of ordinary skill in the art, other examples of rules may be included in the tenant secrets policy 101 to define or specify when a first tenant's secrets 103 may be temporarily or permanently shared with a second tenant. In one embodiment, the tenant secrets policy 101 includes a rule for changing or rotating encryption keys after a predetermined time, e.g., 48 hours, after a first tenant shares one or more encryption keys or other secrets with a second tenant, to limit the duration with which the first tenant's protected data is shared.

[0043] The policy execution engine 102 applies the secrets 103 to the tenant's data or a tenant's users' data, in accordance with the tenant secrets policy 101, according to one embodiment. In one embodiment, the policy execution engine 102 is a daemon that runs in the background of the tenant computing environment 10 to automatically apply, update or rotate the secrets 103 continuously, periodically, or in response to events/triggers, in accordance with the tenant secrets policy 101. Alternatively, the policy execution engine 102 may be a program that executes the tenant secrets policy 101, in direct response to instructions from the tenant of the

tenant computing environment 10. The policy execution engine 102 may be blocked from applying the tenant secrets policy 101 to assets within the multi-tenant computing environment 13, unless or until the policy execution engine 102 receives permission/authorization from the service provider, e.g., via the service provider computing environment 14, according to one embodiment.

[0044] The tenant computing environment 11 is owned, operated, or under the direction/control of a service provider customer or tenant that is different than the customer or tenant of the tenant computing environment 10, according to one embodiment. The tenant computing environment 11 may operate in a similar manner to the tenant computing environment 10. The tenant computing environment 11 includes a tenant secrets database 110, a tenant secrets policy 111, and a policy execution engine 112, according to one embodiment. The tenant secrets database 110 also includes secrets 113, e.g., one or more passwords, encryption keys, and digital certificates that are different than the secrets 103, according to one embodiment. The tenant secrets policy 111 may include similar secrets or security rules as the tenant secrets policy 101, or the tenant secrets policy 111 may include secrets or security rules that are different than the tenant secrets policy 101, according to various embodiments. The policy execution engine 112 is configured to apply the tenant secrets policy 111 to the tenant's information (including the tenant's users' information) that is stored or hosted by the multi-tenant computing environment 13, according to one embodiment. In one embodiment, the tenant of the tenant computing environment 10 is a first tenant or customer of the service provider, and the tenant of the tenant computing environment 11 is a second tenant or customer of the service provider. The first tenant and the second tenant may constitute just two of multiple tenants that are allocated resources or multi-tenant assets within the multi-tenant computing environment 13, according to one embodiment.

[0045] The computing environment 12 includes multiple users that may have contractual or other relationships with the tenants of the tenant computing environments 10 and 11. The computing environment 12 includes a user 120, a user 121, and a user 122, according to one embodiment. As will be appreciated by one of ordinary skill in the art, the computing environment 12 may include more or less users than illustrated in the production environment 1. In one embodiment, the user 120 receives information or services from the tenant of the tenant computing environment 10, e.g., the first tenant, and not from the tenant of the tenant computing environment 11, e.g., the second tenant. In one example embodiment, the user 120 may receive services from the first tenant that are associated with managing financial assets of the user 120.

The user 121 may receive services from the second tenant that are different than the services provided by the first tenant. The user 122 may receive services from both the first tenant and the second tenant either separately or concurrently, according to one embodiment.

[0046] A network communication channel 123 provides communicative coupling between the tenant computing environment 10, the tenant computing environment 11, the computing environment 12, the multi-tenant computing environment 13, and the service provider computing environment 14, according to one embodiment. As described above, the network communication channel 123 within the production environment 1 can include, but is not be limited to, a local area network, a wide area network, a wireless network, the Internet, and intranet, or other network or communications protocols or environments.

[0047] The multi-tenant computing environment 13 provides hardware, software, and/or virtual resources and assets to customers of the service provider, according to one embodiment. The assets or resources of the multi-tenant computing environment 13 may be assigned or allocated to the tenants associated with the tenant computing environments 10 and 11, in response to one or more of the tenants entering an agreement, contract, or other relationship with the service provider, according to one embodiment. As described above in the INTRODUCTORY SYSTEM section, assets can include virtual and non-virtual assets and can be, but are not limited to, virtual machines, virtual servers, databases, services associated with cloud computing, communications systems, mobile devices, sensors, various computing systems and other physical or logical portions of a data center, according to various embodiments. The multi-tenant computing environment 13 includes a multi-tenant asset 130 hosting an application 131, a multi-tenant asset 132 hosting tenant applications 133 and 134, and a multi-tenant asset 135 hosting an application 136, according to various embodiments. Although three multi-tenant assets 130, 132, and 135 are included in the production environment 1, one of ordinary skill in the art will understand that more multi-tenant assets or less multi-tenant assets than those shown may be included in the multi-tenant computing environment 13, according to other embodiments.

[0048] The service provider computing environment 14 assigns or allocates control over one or more assets or parts of assets in the multi-tenant computing environment 13 to the tenants of, or to the tenant computing environments 10 and 11, according to one embodiment. The service provider computing environment 14 may allocate control over assets to the tenant computing environments 10 and 11 by creating one or more user accounts having, for example, permission levels associated with a 'tenant' or 'customer'. In one embodiment, the tenant or

customer permission levels may enable the tenant computing environments 10 and 11 to add, remove, modify, protect, or otherwise manipulate information within the multi-tenant computing environment 13 that is associated with the particular tenant computing environment 10 or 11. In other words, the service provider may enable a first tenant to add, remove, view, or otherwise manipulate the first tenant's data stored in the multi-tenant computing environment 13 and may enable a second tenant to add, remove, view, or otherwise manipulate the second tenant's data stored in the multi-tenant computing environment 13. In some implementations, a tenant associated with the tenant computing environment 10 or 11 may use authentication information received from the service provider to "log in" to assets allocated to the tenant from any computing system or environment, and is not limited to accessing the multi-tenant computing environment 13 solely from the tenant computing environment 10 or 11.

[0049] According to various embodiments, each of the assets and applications of the multi-tenant computing environment 13 can have a particular function and/or relationship with the tenant computing environments 10 and 11. In one embodiment, the multi-tenant asset 130 and the application 131 provide a service to both of the tenant computing environments 10 and 11. For example, the multi-tenant asset 130 and the application 131 can provide one or more services such as database hosting, information management/security, file transfer services, or the like. In one embodiment, the multi-tenant asset 132 allows tenants to execute, host, or run applications from the resources of the multi-tenant asset 132. For example, the multi-tenant asset 132 may allow one tenant to host the tenant application 133 while concurrently allowing another tenant to host the tenant application 134. The tenant applications 133 and 134 may be independent of one another but may share storage space, network resources, computer processors, random access memory, server rack storage space, or other aspects of the multi-tenant asset 132. In one embodiment, the multi-tenant asset 135 provides resources that enable the tenants to jointly provide a user services through the application 136. For example, the application 136 may receive, store, and organize financial data for the user 122 using functions/features/operations supplied by a first tenant, and the application 136 may analyze the financial data for the user 122 using functions/features/operations supplied by a second tenant. The application 136 can include, for example, a financial data engine supplied by the first tenant and an analytics engine supplied by the second tenant in such a manner that the tenants do not share source code while sharing clientele, e.g., the user 122, and the resulting revenue from the clientele, according to one embodiment.

[0050] The service provider computing environment 14 manages whether and how the tenant secrets policies are applied to the assets of the multi-tenant computing environment 13, according to one embodiment. The service provider computing environment 14 manages the application of tenant secrets policies by authorizing or rejecting requests for implementing tenant secrets policies, according to one embodiment. For example, the service provider computing environment 14 can receive a request from the tenant computing environment 10 or 11 to apply a tenant secrets policy 101 or 111. The service provider computing environment 14 then compares the tenant secrets policy to the service provider secrets policy and authorizes the implementation of the tenant secrets policy if the tenant secrets policy complies with or is more restrictive than the service provider secrets policy, according to one embodiment.

[0051] The service provider computing environment 14 is communicatively coupled to the other computing environments of the production environment 1 via the network communication channel 123. In one embodiment, the service provider computing environment 14 may be communicatively coupled to the multi-tenant computing environment 13 by a back-end or trusted communication channel 140.

[0052] The service provider computing environment 14 includes a service provider secrets policy manager 141, a policy analytic engine 142, a policy execution engine 143, and a policy manager log engine 144. The service provider policy manager 141, in addition to the other modules or components of the service provider computing environment 14, may be used to implement, orchestrate, or apply several tenant secrets policies to the multi-tenant assets of the multi-tenant computing environment 13.

[0053] The service provider policy manager 141 may be configured for receiving, maintaining, analyzing, and/or applying one or more tenant secrets policies 101 and 111 associated with the computing environments of the production environment 1. The service provider policy manager 141 includes a service provider secrets policy 145, tenant secrets policies 146, and cross-tenant policies 147. The service provider policy manager 141 receives requests from tenant computing environments 10 and 11 regarding tenant secrets policies 146. In one embodiment, the requests from the tenant computing environments 10 and 11 are seeking permission or authorization to change or apply tenant secrets policies to assets in the multi-tenant computing environment 13. In response to the request from the tenant computing environments, the service provider policy manager 141 may be configured to compare rules, factors, or characteristics of the tenant secrets policies 146 to corresponding rules, factors, or characteristics of the service provider secrets policy 145. For example, the service provider

policy manager 141 may compare a frequency of encryption key rotation, e.g., every three months, of the tenant secrets policy 101 to a required frequency of encryption key rotation, e.g., annually, of the service provider secrets policy 145. If the service provider policy manager 141 determines that the tenant secrets policy 101 satisfies (or is more restrictive than) the requirements or standards set forth by the service provider secrets policy 145, then the service provider policy manager 141 may authorize the request to, for example, apply, execute, or implement the tenant secrets policy 101 in the multi-tenant computing environment 13, according to one embodiment.

[0054] The service provider secrets policy 145 includes rules, requirements, and/or standards for maintaining information security within the assets of the multi-tenant computing environment 13, according to one embodiment. The service provider secrets policy 145 may specify: the types of secrets that tenants may use; minimum password or encryption key lengths; and a frequency by which the secrets are to be rotated, changed, or updated, according to one embodiment.

[0055] The tenant secrets policies 146 includes the tenant secrets policies of the tenants of the multi-tenant computing environment 13, according to one embodiment. In one embodiment, the tenant secrets policies 146 include the tenant secrets policy 101 and the tenant secrets policy 111. The service provider policy manager 141 may receive and store copies of tenant secrets policies for each of the tenants that are allocated resources/assets in the multi-tenant computing environment 13, according to one embodiment. In other embodiments, the tenant secrets policies 146 receive portions of the tenant secrets policies 101 and 111 with requests for authorization from the tenants or the tenant computing environments 10, 11 to apply or implement the tenant secrets policies. In some embodiments, the tenant secrets policies 146 also include temporary or permanent copies of the secrets 103 and/or 113 to enable the service provider policy manager 141 to verify that the tenant's secrets meet minimum requirements for secrets, as established or set forth in the service provider secrets policy 145. In one embodiment, the tenant secrets policies 146 also include temporary or permanent copies of the secrets 103 and/or 113 to enable the service provider policy manager 141 to selectively distribute copies of the secrets 103 and/or 113 to the applications and/or multi-tenant assets of the multi-computing environment 13.

[0056] The cross-tenant policies 147 include rules for authorizing one tenant to access or receive the secrets of another tenant. For example, the application 136 can include functions/operations from a first tenant and functions/operations from the second tenant to

provide multi-tenant services to data of tenant users, e.g., the users 120, 121, 122. To apply the functions of the first tenant and of the second tenant to the users' data, the second tenant's functions may need to decrypt the users' data from the first tenant's secrets (e.g., the secrets 103) prior to executing the second tenant's functions. The second tenant's functions may be configured to request, from the service provider policy manager 141, some of the first tenant's secrets, to complete execution of the application 136. In response, the service provider policy manager 141 may verify that the request from the second tenant's functions comply with the cross-tenant policies 147. For example, the cross-tenant policies 147 may include a list, a database, a table, or other data structure that identifies which tenants are authorized to receive secrets of other tenants for collaborative use within the multi-tenant computing environment 13. In some implementations, the cross-tenant policies 147 include copies of the secrets 103 and/or 113, to support selective distribution of the secrets 103 and/or 113 to tenants that are collaboratively using multi-tenant assets, e.g., the multi-tenant assets 130, 132, 135.

[0057] The service provider policy manager 141 uses the policy analytic engine 142 and the policy execution engine 143 to apply tenant secrets policies to the assets of the multi-tenant computing environment 13, according to one embodiment. The policy analytic engine 142 can be used or configured to compare aspects of the tenant secrets policies 146 to the service provider secrets policy 145 to assist the service provider policy manager 141 in authorizing or rejecting requests to apply tenant secrets policies to multi-tenant assets.

[0058] The policy execution engine 143 may be used or configured to apply tenant secrets policies 146, if the policies satisfy the rules or requirements of the service provider secrets policy 145. According to one embodiment, the function of the policy execution engine 143 resides in and is executed by tenant computing environments 10 and 11 and the policy execution engine 102 and 112, respectively. In some embodiments, the policy analytic engine 142 and/or the policy execution engine 143 are modules that are separate from the service provider policy manager 141. In other embodiments, the policy analytic engine 142 and/or the policy execution engine 143 are modules that are integrated into the service provider policy manager 141.

[0059] The service provider computing environment 14 includes a policy manager log engine 144 to monitor, track, record, and/or report any handling of secrets 103 and/or 113 in addition to any handling of tenant secrets policies 101 and/or 111, according to one embodiment. Mishandling or mismanagement of secrets can result in direct and indirect damage to a user. For example, when private information, such as names and credit card information,

are hacked or stolen from a user such as a large consumer goods retailer, the reputation, investor interest, and consumer confidence related to the consumer goods retailer are negatively impacted. Thus, tenants have a strong interest in protecting user data and in safeguarding secrets, e.g., the secrets 103, 113. The policy manager log engine 144 monitors, tracks, records, and/or reports the receipt, copy, transfer, distribution, destruction, backup, and/or restoration of the secrets 103 and/or 113, according to various embodiments. The policy manager log engine 144 can report this information to the tenants, upon request.

[0060] In one embodiment, the policy manager log engine 144 enables tenants to view and/or receive reports regarding the handling of their secrets through, for example, the tenant interface 148. The tenant interface 148 may provide a graphical user interface or a text-based user interface to enable a tenant to manually retrieve or automatically receive reports concerning the handling and/or management of the tenant's secrets.

PROCESS

[0061] FIG. 2 illustrates a diagram of a process 200 for managing security policies, e.g., tenant secrets policies, for multi-tenant assets within the production environment 1, according to various embodiments.

[0062] At block 202, the tenant computing environment 10 or 11 sends a request to apply a secrets policy to a multi-tenant asset, according to one embodiment. The tenant computing environment 10 or 11 sends the request to the service provider policy manager 141. According to one embodiment, the tenant computing environment 10 or 11 also sends a copy of the tenant secrets policy 101 or 111 and/or a copy of the secrets 103 or 113 with the request. According to another embodiment, the service provider policy manager 141 already maintains, retains, hosts, or manages copies of the tenant secrets policies 101 and 111, in addition to a copy of the secrets 103 and 113.

[0063] At block 204, the service provider policy manager 141 receives the request from the tenant computing environment 10 or 11, according to one embodiment. As discussed above, according to one embodiment, the service provider policy manager 141 also receives a copy of the tenant secrets policies 101 or 111 along with the request to apply the tenant secrets policy. Although the service provider policy manager 141 is specifically identified as receiving the request to apply a tenant secrets policy, in another embodiment, the service provider computing environment 14 or another component or modules within the service provider computing environment 14 receives the request.

[0064] At block 206, the service provider policy manager 141 sends a request to the policy analytic engine 142 to perform a comparison of the tenant secrets policy 101 or 111 to the service provider secrets policy 145, according to one embodiment. In alternative embodiments, the service provider policy manager 141 performs the comparison itself rather than delegating the operation to another module.

[0065] At block 208, the policy analytic engine 142 compares the tenant secrets policy 101 or 111 to the service provider secrets policy 145, according to one embodiment. The policy analytic engine 142 can be configured to compare the tenant secrets policy 101 or 111 to the service provider secrets policy 145 by comparing corresponding rules, factors, or characteristics. For example, each of the secrets policies can include rules relating to the frequency with which secrets, e.g., encryption keys, are rotated in the multi-tenant assets. As another example, each of the secrets policies can include rules relating to the types of secrets allowable, e.g., passwords and/or encryption keys, and can include rules specifying allowable lengths of secrets, e.g., 400 bit encryption keys. The policy analytic engine 142 compares the rules of the tenant secrets policy 101 or 111 to the service provider secrets policy 145 to determine whether the tenant secrets policy satisfies the service provider secrets policy. In one embodiment, the tenant secrets policy satisfies the service provider secrets policy if the tenant secrets policy is at least as restrictive as the service provider secrets policy. In other words, if the tenant secrets policy is as restrictive as the service provider secrets policy, or if the tenant secrets policy is more restrictive than the service provider secrets policy, then the policy analytic engine 142 determines that the tenant secrets policy satisfies the requirements of the service provider secrets policy. As a specific example, if the tenant secrets policy 101 specifies that encryption keys are required to be rotated once every six months and the service provider secrets policy 145 only requires that encryption keys be rotated once every nine months then the policy analytic engine 142 determines that the tenant secrets policy 101 is more restrictive than the service provider secrets policy 145 and that the tenant secrets policy 101 satisfies the service provider secrets policy 145.

[0066] The policy analytic engine 142 notifies the service provider policy manager 141 of the results of the comparison, according to one embodiment. If the tenant secrets policy does not satisfy the service provider secrets policy, the process 200 proceeds to block 210 to notify the service provider policy manager 141 that the tenant secrets policy did not satisfy the service provider secrets policy. If the tenant secrets policy satisfies the service provider secrets policy, the process 200 proceeds to block 214 to notify the service provider policy manager that the tenant secrets policy does satisfy the service provider secrets policy.

[0067] At block 210, the service provider policy manager 141 transmits a notification of request rejection to the tenant computing environment 10 or 11, from which the request originated. In one embodiment, the service provider policy manager 141 includes reasons for the rejection and/or includes instructions for overcoming the rejection.

[0068] At block 212, the tenant computing environment 10 or 11 receives the notification rejection of the request and the process 200 terminates, according to one embodiment.

[0069] At block 214, the service provider policy manager 141 transmits a request approval notification to the tenant computing environment 10 or 11. The service provider policy manager 141 also proceeds to execute the approved tenant secrets policy, according to one embodiment. From the block 214, the process 200 proceeds concurrently to block 216 and block 218, according to one embodiment.

[0070] At block 216, the tenant computing environment 10 or 11 receives the notification of the approval from the service provider policy manager 141, according to one embodiment. In other embodiments, the tenant computing environment 10 or 11 receives the notification of approval directly from the policy analytic engine 142. In response to receiving notification of the approval of the request, the tenant computing environment 10 or 11 authorizes a policy execution engine 102 to apply the tenant secrets policy, e.g., as shown in block 220. In other embodiments, as illustrated in the process 200, the service provider policy manager 141 authorizes the execution of the tenant secrets policy that satisfies the service provider secrets policy, e.g., as shown in block 218.

[0071] At block 218, the service provider policy manager 141 authorizes execution of the tenant secrets policy, in response to receiving notification from the policy analytic engine 142 that the tenant secrets policy satisfies the service provider secrets policy, according to one embodiment. The service provider policy manager 141 may then apply the tenant secrets policy to the multi-tenant assets 130, 132, and/or 135, in one embodiment. In another embodiment, the process 200 proceeds to block 220.

[0072] At block 220, the policy execution engine 102 or 143 applies the tenant secrets policy, in response to approval or instruction from the service provider policy manager 141, according to one embodiment. The policy execution engine 102 or 143 may apply the tenant secrets policy by parsing the request, identifying portions of the tenant secrets policy to be applied, establishing a connection to one or more multi-tenant assets 130, 132, and 135, and

providing instructions to the multi-tenant asset(s) in accordance with the tenant secrets policy associated with the request.

[0073] At block 222, one or more multi-tenant assets 130, 132, and 135 apply the authorized tenant secrets policy to assets allocated to the tenant associated with the tenant secrets policy, according to one embodiment. In one embodiment, the one or more multi-tenant assets 130, 132, and 135 apply the authorized tenant secrets policy by updating one or more tables, databases, or other data structures with the new rules, or with updates to pre-existing rules, that are being modified by the tenant secrets policy.

[0074] At block 224, the tenant computing environment 10 or 11 receives notification from the one or more multi-tenant assets 130, 132, and 135 that the tenant secrets policy has been applied to the multi-tenant asset(s), according to one embodiment. Upon receipt of the notification, the tenant computing environment 10 or 11 may update its corresponding tenant secrets database 100 or 110 to reflect the current tenant secrets policy that is associated with and applied to the multi-tenant asset(s). At block 224, the process 200 terminates.

[0075] FIG. 3 illustrates a diagram of a process 300 for managing cross-tenant secrets sharing within the production environment 1, according to one embodiment.

[0076] At block 302, the application 136 includes a first tenant function that requests access to second tenant secrets, according to one embodiment. As discussed above, the application 136 may be an application that includes functions/operations that are provided by more than one tenant, e.g., a first tenant and a second tenant. In other words, a first tenant may provide some functions of the application 136 and a second tenant may provide other functions of the application 136. According to one embodiment, the first tenant may apply secrets 103 to data of first tenant users associated with the application 136, and the second tenant may apply the secrets 113 to data of second tenant users associated with the application 136. For the application to apply the features of both tenants to user data, the first tenant may need access to second tenant secrets, e.g., the secrets 113, and the second tenant may need access to first tenant secrets, e.g., the secrets 103, according to one embodiment. However, to maintain security over each tenants' users' data, each tenant may maintain encryption of its users' data until and unless access to the tenant's secrets is requested by another tenant.

[0077] At block 304, the multi-tenant asset 135 requests authorization to share second tenant secrets with the first tenant function of the application 136, in response to a request received from the application 136, according to one embodiment. In one embodiment, the multi-tenant asset 135 includes copies of cross-tenant policies and copies of tenants' secrets, and

the multi-tenant asset 135 authorizes or rejects the request to share secrets that was received from the application 136. In another embodiment, the multi-tenant asset 135 requests authorization to share tenant secrets with other tenants from, for example, the service provider policy manager 141.

[0078] At block 306, the service provider policy manager 141 determines if the request to share tenant secrets violates cross-tenant policies, e.g., cross-tenant policies 147, in one embodiment. The service provider policy manager 141 may itself analyze the cross-tenant policies 147 to determine if the requesting tenant is authorized to receive permanent or temporary access to the secrets of the other tenant of the request. The service provider policy manager 141 uses a spreadsheet, a table, a database, or other data structure to determine if the cross-tenant policies 147 permit secrets sharing between the first tenant and the second tenant, according to one embodiment. According to another embodiment, the service provider policy manager 141 analyzes the cross-tenant policies 147 to determine if cross-tenant secrets sharing is permitted for application 136 or for the users of the application 136, e.g., users 120, 121, 122. The service provider policy manager 141 may then directly reply to the request of the multi-tenant asset. In one embodiment, the service provider policy manager 141 delegates or assigns responsibility to the policy analytic engine 142 to determine if the cross-tenant secrets sharing request is permissible.

[0079] At block 308, the policy analytic engine 142 determines if secret sharing is allowed between the first and second tenants of the application 136, according to one embodiment. As discussed above, the policy analytic engine 142 may use a spreadsheet, a table, a database, or other data structure to determine if the cross-tenant policies 147 permit or allow secrets sharing between the first tenant and the second tenant, according to one embodiment. If the cross-tenant policies 147 do not permit the requested secrets sharing, the process 300 proceeds to: block 310, where the service provider policy manager 141 issues a request rejection; block 312, where the multi-tenant asset 135 issues a request rejection; and block 314, where the application 136 receives notification of the request rejection and notifies the first tenant, e.g., the tenant computing environment 10, according to one embodiment.

[0080] If the policy analytic engine 142 determines that the cross-tenant policies 147 permit or allow the first tenant function to receive the requested access to the second tenant secrets, e.g., the secrets 113, the process 300 proceeds to block 316.

[0081] At block 316, the service provider policy manager 141 approves the request from the first tenant function to receive at least partial access to secrets of the second tenant. The

service provider policy manager 141 may provide one or more copies of second tenant secrets, e.g., the secrets 113, in addition to the copy of the cross-tenant policies 147 to the multi-tenant asset 135 for distribution to the application 136, according to one embodiment. In another embodiment, the service provider policy manager 141 transmits notification of authorization or approval of the request to the multi-tenant asset 135, so that the multi-tenant asset 135 can authorize the application 136 to use the second tenant secrets.

[0082] At block 318, the multi-tenant asset 135 provides the first tenant function with access to the second tenant secrets for use in the application 136, according to one embodiment. As described above, the multi-tenant asset 135 may receive copies of the second tenant secrets from the service provider policy manager 141, or the multi-tenant asset 135 may already store or maintain local copies of the second tenant secrets, for selective distribution to, for example, the application 136.

[0083] At block 320, the application 136 receives approval for the access request of the first tenant function to the second tenant secrets. The first tenant function may use the second tenant secrets to decrypt user data associated with the application 136 prior to performing one or more operations or functions on the unencrypted data, according to one embodiment. In one embodiment, the first tenant function is configured to decrypt the second tenant secrets, perform the first tenant function on the unencrypted data, and re-encrypt the data with the second tenant secrets, in accordance with the cross-tenant policies 147. At block 320, the process 300 ends.

[0084] FIG. 4 illustrates a diagram of a process 400 for logging and reporting use of a tenant's secrets by a service provider, according to various embodiments.

[0085] At block 402, the process 400 begins.

[0086] At block 404, the service provider maintains a first log of one or more of the use, copying, deletion, receipt, transfer, transmission, backup, and restoration of a first tenant's secrets, according to one embodiment. The service provider may create and/or maintain the first log with a service provider computing environment, with a service provider policy manager, with a policy manager log engine, or with one or more other components or modules. The service provider may update the first log with each use or manipulation of one or more of the first tenant's secrets. In one embodiment, the service provider logs the creation of any backup copies of tenant secrets in addition to logging any data restoration operations that include one or more of the tenant's secrets. By allowing the tenant to view any or all operations or manipulations of the tenant's secrets, the tenant can determine if the service provider is taking appropriate actions and/or measurements to protect and manage the secrets of the tenant. In the

event that the tenant determines that the service provider's management of the tenant's secrets is inadequate, the tenant may seek to establish relationships with a different, more secure or competent, service provider, according to one embodiment.

[0087] At block 406, the service provider maintains a second log of one or more of the use, copying, deletion, receipt, transfer, transmission, backup, and restoration of a second tenant's secrets, according to one embodiment.

[0088] At block 408, the service provider provides the first log to the first tenant, in response to receiving a request for the first log from the first tenant, according to one embodiment. In practice, the service provider may provide the first log and may receive a request for the first log by the service provider computing environment, and the request from the first tenant may be received by the service provider computing environment from a first tenant computing environment, according to one embodiment.

[0089] At block 410, the service provider or service provider computing environment provides the second log to the second tenant or second tenant computing environment, according to one embodiment. The service provider may provide the second log to the second tenant in response to a request for the second log from the second tenant. In one embodiment, the first and second tenants may require the service provider computing environment to automatically and/or periodically transmit copies of the first/second logs to the respective tenant computing environments, absent a request for the log from the tenant computing environments. In another embodiment, the service provider computing environment is configured to transmit a copy of one or more logs to the one or more tenant computing environments in response to a trigger or event. According to one embodiment, the trigger or events may be determined by the tenant computing environment and may include triggers or events such as a security breach, suspicious or excessive requests for access to tenant secrets, or any predetermined use or manipulation of one or more tenant secrets, e.g., including copying, transmitting, receiving, deleting, backing up, and restoring a copy of the tenant secret.

[0090] At block 412, the process 400 ends.

[0091] As noted above, the specific illustrative examples discussed above are but illustrative examples of implementations of embodiments of the method or process for providing rotating-key encrypted databases. Those of skill in the art will readily recognize that other implementations and embodiments are possible. Therefore the discussion above should not be construed as a limitation on the claims provided below.

[0092] In the discussion above, certain aspects of one embodiment include process steps and/or operations and/or instructions described herein for illustrative purposes in a particular order and/or grouping. However, the particular order and/or grouping shown and discussed herein are illustrative only and not limiting. Those of skill in the art will recognize that other orders and/or grouping of the process steps and/or operations and/or instructions are possible and, in some embodiments, one or more of the process steps and/or operations and/or instructions discussed above can be combined and/or deleted. In addition, portions of one or more of the process steps and/or operations and/or instructions can be re-grouped as portions of one or more other of the process steps and/or operations and/or instructions discussed herein. Consequently, the particular order and/or grouping of the process steps and/or operations and/or instructions discussed herein do not limit the scope of the invention as claimed below.

[0093] As discussed in more detail above, using the above embodiments, with little or no modification and/or input, there is considerable flexibility, adaptability, and opportunity for customization to meet the specific needs of various parties under numerous circumstances.

[0094] In the discussion above, certain aspects of one embodiment include process steps and/or operations and/or instructions described herein for illustrative purposes in a particular order and/or grouping. However, the particular order and/or grouping shown and discussed herein are illustrative only and not limiting. Those of skill in the art will recognize that other orders and/or grouping of the process steps and/or operations and/or instructions are possible and, in some embodiments, one or more of the process steps and/or operations and/or instructions discussed above can be combined and/or deleted. In addition, portions of one or more of the process steps and/or operations and/or instructions can be re-grouped as portions of one or more other of the process steps and/or operations and/or instructions discussed herein. Consequently, the particular order and/or grouping of the process steps and/or operations and/or instructions discussed herein do not limit the scope of the invention as claimed below.

[0095] The present invention has been described in particular detail with respect to specific possible embodiments. Those of skill in the art will appreciate that the invention may be practiced in other embodiments. For example, the nomenclature used for components, capitalization of component designations and terms, the attributes, data structures, or any other programming or structural aspect is not significant, mandatory, or limiting, and the mechanisms that implement the invention or its features can have various different names, formats, or protocols. Further, the system or functionality of the invention may be implemented via various combinations of software and hardware, as described, or entirely in hardware elements. Also,

particular divisions of functionality between the various components described herein are merely exemplary, and not mandatory or significant. Consequently, functions performed by a single component may, in other embodiments, be performed by multiple components, and functions performed by multiple components may, in other embodiments, be performed by a single component.

[0096] Some portions of the above description present the features of the present invention in terms of algorithms and symbolic representations of operations, or algorithm-like representations, of operations on information/data. These algorithmic or algorithm-like descriptions and representations are the means used by those of skill in the art to most effectively and efficiently convey the substance of their work to others of skill in the art. These operations, while described functionally or logically, are understood to be implemented by computer programs or computing systems. Furthermore, it has also proven convenient at times to refer to these arrangements of operations as steps or modules or by functional names, without loss of generality.

[0097] Unless specifically stated otherwise, as would be apparent from the above discussion, it is appreciated that throughout the above description, discussions utilizing terms such as, but not limited to, “activating”, “accessing”, “adding”, “aggregating”, “alerting”, “applying”, “analyzing”, “associating”, “calculating”, “capturing”, “categorizing”, “classifying”, “comparing”, “creating”, “defining”, “detecting”, “determining”, “distributing”, “eliminating”, “encrypting”, “extracting”, “filtering”, “forwarding”, “generating”, “identifying”, “implementing”, “informing”, “monitoring”, “obtaining”, “posting”, “processing”, “providing”, “receiving”, “requesting”, “saving”, “sending”, “storing”, “substituting”, “transferring”, “transforming”, “transmitting”, “using”, etc., refer to the action and process of a computing system or similar electronic device that manipulates and operates on data represented as physical (electronic) quantities within the computing system memories, registers, caches or other information storage, transmission or display devices.

[0098] The present invention also relates to an apparatus or system for performing the operations described herein. This apparatus or system may be specifically constructed for the required purposes, or the apparatus or system can comprise a general purpose system selectively activated or configured/reconfigured by a computer program stored on a computer program product as discussed herein that can be accessed by a computing system or other device.

[0099] Those of skill in the art will readily recognize that the algorithms and operations presented herein are not inherently related to any particular computing system, computer

architecture, computer or industry standard, or any other specific apparatus. Various general purpose systems may also be used with programs in accordance with the teaching herein, or it may prove more convenient/efficient to construct more specialized apparatuses to perform the required operations described herein. The required structure for a variety of these systems will be apparent to those of skill in the art, along with equivalent variations. In addition, the present invention is not described with reference to any particular programming language and it is appreciated that a variety of programming languages may be used to implement the teachings of the present invention as described herein, and any references to a specific language or languages are provided for illustrative purposes only and for enablement of the contemplated best mode of the invention at the time of filing.

[0100] The present invention is well suited to a wide variety of computer network systems operating over numerous topologies. Within this field, the configuration and management of large networks comprise storage devices and computers that are communicatively coupled to similar or dissimilar computers and storage devices over a private network, a LAN, a WAN, a private network, or a public network, such as the Internet.

[0101] It should also be noted that the language used in the specification has been principally selected for readability, clarity and instructional purposes, and may not have been selected to delineate or circumscribe the inventive subject matter. Accordingly, the disclosure of the present invention is intended to be illustrative, but not limiting, of the scope of the invention, which is set forth in the claims below.

[0102] In addition, the operations shown in the FIG.s, or as discussed herein, are identified using a particular nomenclature for ease of description and understanding, but other nomenclature is often used in the art to identify equivalent operations.

[0103] Therefore, numerous variations, whether explicitly provided for by the specification or implied by the specification or not, may be implemented by one of skill in the art in view of this disclosure.

The embodiments of the present invention for which an exclusive property or privilege is claimed are defined as follows:

1. A computing system implemented method for managing secrets of tenants of a multi-tenant computing environment, comprising:

maintaining, by a service provider computing system, a service provider secrets policy, wherein the service provider secrets policy includes security requirements associated with the secrets of tenants within the multi-tenant computing environment;

receiving, by the service provider computing system, a first tenant secrets policy from a first tenant computing system for a first tenant of the multi-tenant computing environment;

receiving a request from the first tenant computing system to apply the first tenant secrets policy to at least one multi-tenant asset in the multi-tenant computing environment;

in response to receiving the request, comparing the first tenant secrets policy with the security requirements of the service provider secrets policy;

if the first tenant secrets policy satisfies the security requirements, authorizing, with the service provider computing system, the request from the first tenant computing system to apply the first tenant secrets policy to the at least one multi-tenant asset;

if the first tenant secrets policy fails the security requirements, rejecting the request from the first tenant computing system to apply the first tenant secrets policy to the at least one multi-tenant asset; and

applying the first tenant secrets policy to the at least one multi-tenant asset if the request from the first tenant computing system is authorized.

2. The computing system implemented method of claim 1 further comprising:

receiving, by the service provider computing system, a second tenant secrets policy from a second tenant computing system for a second tenant of the multi-tenant computing environment.

3. The computing system implemented method of claim 2 wherein receiving the first tenant secrets policy includes receiving first tenant secrets,
wherein receiving the second tenant secrets policy includes receiving second tenant secrets, the method further comprising:
maintaining a first log of one or more of use, copying, deletion, receipt, transfer, transmission, backup, and restoration of the first tenant secrets;
maintaining a second log of one or more of use, copying, deletion, receipt, transfer, transmission, backup, and restoration of the second tenant secrets;
providing the first log to the first tenant computing system, in response to receiving a request for the first log from the first tenant computing system;
and
providing the second log to the second tenant computing system, in response to receiving a request for the second log from the second tenant computing system.
4. The computing system implemented method of claim 2, further comprising:
maintaining a cross-tenant policy that includes rules associated with sharing first tenant secrets with the second tenant computing system, and with sharing second tenant secrets with the first tenant computing system.
5. The computing system implemented method of claim 2, further comprising:
maintaining a cross-tenant policy,
wherein the cross-tenant policy includes rules and exceptions for enabling the at least one multi-tenant asset to concurrently use first tenant secrets and second tenant secrets to decrypt encrypted user data,
wherein the at least one multi-tenant asset is shared by the first tenant and the second tenant.
6. The computing system implemented method of claim 2 wherein the first tenant secrets policy includes rules for applying first tenant secrets and the second tenant secrets policy includes rules for applying second tenant secrets, wherein the first tenant secrets and the second tenant secrets include encryption keys,

wherein the security requirements of the service provider secrets policy include a minimum frequency by which the encryption keys are required to be rotated if some of the encryption keys are applied by either the first tenant or the second tenant to data stored by the at least one multi-tenant asset in the multi-tenant computing environment.

7. The computing system implemented method of claim 2 wherein the request is a first request, the method further comprising:

receiving a second request from the second tenant computing system to apply the second tenant secrets policy to the at least one multi-tenant asset in the multi-tenant computing environment;

in response to receiving the second request, comparing the second tenant secrets policy with the security requirements of the service provider secrets policy; and

if the second tenant secrets policy satisfies the security requirements, authorizing the second request and applying the second tenant secrets policy to the at least one multi-tenant asset.

8. The computing system implemented method of claim 1 wherein applying the first tenant secrets policy includes applying, by the service provider, the first tenant secrets policy to the at least one multi-tenant asset.

9. The computing system implemented method of claim 1 wherein applying the first tenant secrets policy includes applying, by the first tenant computing system, the first tenant secrets policy to the at least one multi-tenant asset.

10. The computing system implemented method of claim 2 wherein the at least one multi-tenant asset includes at least one:

application shared by the first tenant and the second tenant;

server computing system shared by the first tenant and the second tenant; and

non-volatile memory device shared by the first tenant and the second tenant.

11. A computing system implemented method for managing secrets of customers of a service provider for a multi-tenant computing environment, comprising:

maintaining, by a service provider computing system, a service provider security policy that includes security requirements for the customers of a multi-tenant asset, wherein the multi-tenant asset is hosted for the customers by the service provider, wherein the multi-tenant asset includes at least one of:

an application shared by the customers;

a server computing system shared by customers; a virtual machine; and non-volatile memory device logically divided between at least two of

the customers;

receiving a request from a first one of the customers to apply a first customer security policy to a part of the multi-tenant asset allocated to the first one of the customers,

wherein the first customer security policy includes rules for managing first customer secrets with the multi-tenant asset in the multi-tenant computing environment;

comparing the first customer security policy of the request to the service provider security policy; and

if the first customer security policy is at least as restrictive as the service provider security policy, authorizing the request to enable the first one of the customers to apply the first customer security policy to the part of the multi-tenant asset allocated to the first one of the customers.

12. The computing system implemented method of claim 11, further comprising:

receiving, by the service provider computing system, a second request from a second one of the customers for access to the first customer secrets to enable the second one of the customers to execute an operation associated with the multi-tenant asset;

accessing a cross-tenant security policy maintained by the service provider computing system to determine access rights of the second one of the customers to the first customer secrets; and

authorizing the second request if the cross-tenant security policy includes access rights for the second one of the customers to the first customer secrets for the operation associated with the multi-tenant asset.

13. The computing system implemented method of claim 11 wherein the first customer secrets include one or more passwords, encryption keys, and digital certificates that are required to access data of the first one of the customers within the multi-tenant asset.

14. The computing system implemented method of claim 11, further comprising:
receiving, by the multi-tenant asset, the first customer secrets and second customer secrets;

managing, with the service provider computing system, the first customer secrets and the second customer secrets received by the multi-tenant asset;

recording a first log associated with managing the first customer secrets and a second log associated with managing the second customer secrets; and

providing the first log to the first one of the customers and the second log to the second one of the customers.

15. The computing system implemented method of claim 14, wherein providing the first log and the second log includes hosting a user interface for viewing and retrieving the first log and the second log.

16. A system for managing secrets of tenants of a multi-tenant computing environment, the system comprising:

at least one processor; and

at least one memory coupled to the at least one processor, the at least one memory having stored therein instructions which when executed by any set of the one or more processors, perform a process for managing secrets of tenants of the multi-tenant computing environment, the process including:

maintaining, by a service provider computing system, a service provider secrets policy,

wherein the service provider secrets policy includes security requirements associated with the secrets of tenants within the multi-tenant computing environment;

receiving, by the service provider computing system, a first tenant secrets policy from a first tenant computing system for a first tenant of the multi-tenant computing environment;

receiving a request from the first tenant computing system to apply the first tenant secrets policy to at least one multi-tenant asset in the multi-tenant computing environment;

in response to receiving the request, comparing the first tenant secrets policy with the security requirements of the service provider secrets policy;

if the first tenant secrets policy satisfies the security requirements, authorizing, with the service provider computing system, the request from the first tenant computing system to apply the first tenant secrets policy to the at least one multi-tenant asset;

if the first tenant secrets policy fails the security requirements, rejecting the request from the first tenant computing system to apply the first tenant secrets policy to the at least one multi-tenant asset; and

applying the first tenant secrets policy to the at least one multi-tenant asset if the request from the first tenant computing system is authorized.

17. The system of claim 16 further comprising:

receiving, by the service provider computing system, a second tenant secrets policy from a second tenant computing system for a second tenant of the multi-tenant computing environment.

18. The system of claim 17 wherein receiving the first tenant secrets policy includes receiving first tenant secrets,

wherein receiving the second tenant secrets policy includes receiving second tenant secrets, the process further comprising:

maintaining a first log of one or more of use, copying, deletion, receipt, transfer, transmission, backup, and restoration of the first tenant secrets;

maintaining a second log of one or more of use, copying, deletion, receipt, transfer, transmission, backup, and restoration of the second tenant secrets;

providing the first log to the first tenant computing system, in response to receiving a request for the first log from the first tenant computing system; and

providing the second log to the second tenant computing system, in response to receiving a request for the second log from the second tenant computing system.

19. The system of claim 17 wherein the process further comprises:

maintaining a cross-tenant service policy that includes rules associated with sharing first tenant secrets with the second tenant computing system, and with sharing second tenant secrets with the first tenant computing system.

20. The system of claim 17 wherein the process further comprises:

maintaining a cross-tenant policy,

wherein the cross-tenant policy includes rules and exceptions for enabling the at least one multi-tenant asset to concurrently use first tenant secrets and second tenant secrets to decrypt encrypted user data,

wherein the at least one multi-tenant asset is shared by the first tenant and the second tenant.

21. The system of claim 17 wherein the first tenant secrets policy includes rules for applying first tenant secrets and the second tenant secrets policy includes rules for applying second tenant secrets, wherein the first tenant secrets and the second tenant secrets include encryption keys,

wherein the security requirements of the service provider secrets policy include a minimum frequency by which the encryption keys are required to be rotated if some of the encryption keys are applied by either the first tenant or the second tenant to data stored by the at least one multi-tenant asset in the multi-tenant computing environment.

22. The system of claim 17 wherein the request is a first request, the process further comprising:

receiving a second request from the second tenant computing system to apply the second tenant secrets policy to the at least one multi-tenant asset in the multi-tenant computing environment;

in response to receiving the second request, comparing the second tenant secrets policy with the security requirements of the service provider secrets policy;
and

if the second tenant secrets policy satisfies the security requirements, authorizing the second request and applying the second tenant secrets policy to the at least one multi-tenant asset.

23. The system of claim 16 wherein applying the first tenant secrets policy includes applying, by the service provider, the first tenant secrets policy to the at least one multi-tenant asset.

24. The system of claim 16 wherein applying the first tenant secrets policy includes applying, by the first tenant computing system, the first tenant secrets policy to the at least one multi-tenant asset.

25. The system of claim 17 wherein the at least one multi-tenant asset includes at least one:

an application shared by the first tenant and the second tenant;

a server computing system shared by the first tenant and the second tenant;

and

non-volatile memory device shared by the first tenant and the second tenant.

26. A system for managing secrets of customers of a service provider for a multi-tenant computing environment, comprising:

at least one processor; and

at least one memory coupled to the at least one processor, the at least one memory having stored therein instructions which when executed by any set of the one

or more processors, perform a process for managing secrets of customers of a service provider for a multi-tenant computing environment, the process including:

maintaining, by a service provider computing system, a service provider security policy that includes security requirements for customers of a multi-tenant asset,

wherein the multi-tenant asset is hosted for the customers by the service provider,

wherein the multi-tenant asset includes at least one of:

an application shared by the customers;

server computing system shared by customers;

a virtual machine; and

non-volatile memory device logically divided between at least some of the customers;

receiving a request from a first one of the customers to apply a first customer security policy to a part of the multi-tenant asset allocated to the first one of the customers,

wherein the first customer security policy includes rules for managing first customer secrets with the multi-tenant asset in the multi-tenant computing environment;

comparing the first customer security policy of the request to the service provider security policy; and

if the first customer security policy is at least as restrictive as the service provider security policy, authorizing the request to enable the first one of the customers to apply the first customer security policy to the part of the multi-tenant asset allocated to the first one of the customers.

27. The system of claim 26 wherein the process further comprises:

receiving, by the service provider computing system, a second request from a second one of the customers for access to the first customer secrets to enable the second one of the customers to execute an operation associated with the multi-tenant asset;

accessing a cross-tenant security policy maintained by the service provider computing system to determine access rights of the second one of the customers to the first customer secrets; and

authorizing the second request if the cross-tenant security policy includes access rights for the second one of the customers to the first customer secrets for the operation associated with the multi-tenant asset.

28. The system of claim 26 wherein the first customer secrets include one or more passwords, encryption keys, and digital certificates that are required to access data of the first one of the customers within the multi-tenant asset.

29. The system of claim 26 wherein the process further comprises:

receiving, by the multi-tenant asset, the first customer secrets and second customer secrets;

managing, with the service provider computing system, the first customer secrets and the second customer secrets received by the multi-tenant asset;

recording a first log associated with managing the first customer secrets and a second log associated with managing the second customer secrets; and

providing the first log to the first one of the customers and the second log to the second one of the customers.

30. The system of claim 29, wherein providing the first log and the second log includes hosting a user interface for viewing and retrieving the first log and the second log.

1/4


PRODUCTION ENVIRONMENT 1

FIG. 1




FIG. 2

3/4




FIG. 3

4/4




FIG. 4

PRODUCTION ENVIRONMENT 1

