



- (51) International Patent Classification:  
G06F 9/50 (2006.01)
- (21) International Application Number:  
PCT/US2020/045516
- (22) International Filing Date:  
07 August 2020 (07.08.2020)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
 

62/884,936	09 August 2019 (09.08.2019)	US
62/884,994	09 August 2019 (09.08.2019)	US
62/884,998	09 August 2019 (09.08.2019)	US
16/986,162	05 August 2020 (05.08.2020)	US
16/986,163	05 August 2020 (05.08.2020)	US
16/986,164	05 August 2020 (05.08.2020)	US

(74) Agent: KENNA, Karl, F. et al.; Tucker Ellis LLP, 201 Mission Street, Suite 2310, San Francisco, California 94105 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, IT, JO, JP, KE, KG, KH, KN, KP, KR, KW, 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, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, 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).

- (71) Applicant: ORACLE INTERNATIONAL CORPORATION [US/US]; 500 Oracle Parkway, Mail Stop 50P7, Redwood Shores, California 94065 (US).
- (72) Inventor: RUSHTON, Matthew; Oracle International Corporation, 500 Oracle Parkway, Mail Stop 50P7, Redwood Shores, California 94065 (US).

(54) Title: SYSTEM AND METHOD FOR SUPPORTING A QUOTA POLICY LANGUAGE IN A CLOUD INFRASTRUCTURE ENVIRONMENT

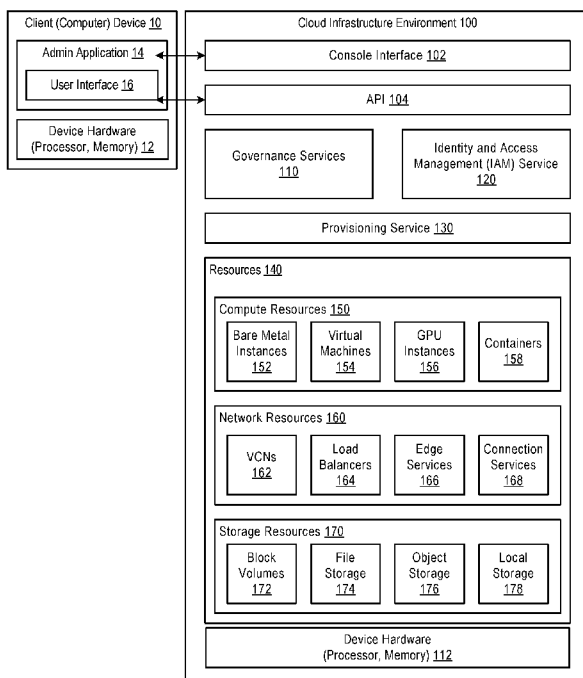


FIGURE 1

(57) Abstract: Systems and methods described herein support compartment quotas in a cloud infrastructure environment. Cloud administrators do not generally have the ability to restrict resource usage in existing clouds. Granting a user permission to create resources allows them to create any number of resources up to a predefined account limit. Compartment quotas allow admins to restrict a user's resource usage to the appropriate level allowing fine-tuned cost control.



**Published:**

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

**SYSTEM AND METHOD FOR SUPPORTING A QUOTA POLICY LANGUAGE  
IN A CLOUD INFRASTRUCTURE ENVIRONMENT**

COPYRIGHT NOTICE

5           A portion of the disclosure of this patent document contains  
material which is subject to copyright protection. The copyright  
owner has no objection to the facsimile reproduction by anyone  
of the patent document or the patent disclosure, as it appears in  
the Patent and Trademark Office patent file or records, but  
10           otherwise reserves all copyright rights whatsoever.

**Claim of Priority and Cross Reference to Related Applications:**

**[0001]**       This application claims the benefit of priority to U.S. Provisional Patent  
Application titled "SYSTEM AND METHOD FOR COMPARTMENT QUOTAS IN A CLOUD  
15   INFRASTRUCTURE ENVIRONMENT", Application No. 62/884,936, filed August 9, 2019; U.S.  
Provisional Patent Application titled "SYSTEM AND METHOD FOR SUPPORTING A QUOTA  
POLICY LANGUAGE IN A CLOUD INFRASTRUCTURE ENVIRONMENT", Application No.  
62/884,994, filed August 9, 2019; and U.S. Provisional Patent Application titled "SYSTEM AND  
METHOD FOR SUPPORTING A USAGE CALCULATION PROCESS IN A CLOUD  
20   INFRASTRUCTURE ENVIRONMENT", Application No. 62/884,998, filed August 9, 2019; each  
of which above applications is herein incorporated by reference in its entirety.

**[0002]**       This application also claims the benefit of priority to the following applications:  
U.S. Patent Application titled "SYSTEM AND METHOD FOR COMPARTMENT QUOTAS IN  
A CLOUD INFRASTRUCTURE ENVIRONMENT", Application No. 16/986,162, filed August 5,  
25   2020; U.S. Patent Application titled "SYSTEM AND METHOD FOR SUPPORTING A QUOTA  
POLICY LANGUAGE IN A CLOUD INFRASTRUCTURE ENVIRONMENT", Application No.  
16/986,163, filed August 5, 2020; and U.S. Patent Application titled "SYSTEM AND METHOD  
FOR SUPPORTING A USAGE CALCULATION PROCESS IN A CLOUD INFRASTRUCTURE  
ENVIRONMENT", Application No. 16/986,164, filed August 5, 2020; each of which above  
30   applications is herein incorporated by reference in its entirety.

**Technical Field**

**[0003]**       Embodiments described herein are generally related to cloud infrastructure  
environments, such as Infrastructure as a Service (IaaS), and are particularly related to  
35   systems and methods for providing systems and methods for providing resource constraints  
within such cloud infrastructure environments.

**Background:**

[0004] Cloud infrastructure environments can comprise sets of complementary cloud services that enable users and clients (throughout the specification, the terms “clients” and “customers” can be used interchangeably) to build and run a wide range of applications and services in a highly available hosted environment.

[0005] Year to year, more and more businesses and organizations are migrating mission critical applications and systems to a cloud infrastructure environment. There are various reasons for this shift. For example, many businesses are moving to the cloud in order to reduce the cost and complexity of operating, maintaining, and building out on-premise infrastructure. As well, cloud infrastructure also allows for a more rapid information technology (IT) delivery mechanism. Some businesses and organizations additionally see the cloud infrastructure environment as a means to gain a leg up on competition by adapting to a nimbler system.

[0006] Within IaaS (Infrastructure as a Service) models, a cloud provider can provide, host, and manage infrastructure components that would, in traditionally settings, be on-premise at each customer’s/client’s location. Such traditionally on-premise components can include hardware, for example, data warehouses and data centers, servers, storage, networking hardware, as well as software, such as virtualization software.

[0007] IaaS providers can, in addition to providing hardware and software that would traditionally be on-premise, also provide services to their clients and customers. As an example, clients and customers can be allowed to tailor their IaaS subscription to fit their needs, which then in turn allows for detailed and broken-down billing and invoicing. IaaS can also support features such as load balancing, redundancy, replication and recover. Because such services are offered and supported by the IaaS provider (and not the customer), this leaves clients and customers to be more focused on improving their business by pushing more into automation and orchestration for their services.

[0008] Cloud infrastructures enable users and clients to seamlessly run traditional enterprise applications along with cloud-native apps, all on the same platform, reducing operational overhead and enabling direct connectivity between both types of workloads.

**Summary:**

[0009] In accordance with an embodiment, described herein are systems and methods for providing compartment quotas for a cloud infrastructure environment. In accordance with an embodiment, cloud administrators do not generally have the ability to restrict resource usage in existing clouds. Granting a user permission to create resources allows them to create any number of resources up to a predefined account limit. Compartment quotas allow admins to restrict a user’s resource usage to the appropriate level allowing fine-tuned cost control.

[00010] In accordance with an embodiment, customers can be assigned service level limits defined by the cloud infrastructure environment at account creation time. These service level limits restrict the total number of resources a customer can create across the entire tenancy (e.g., across multiple regions with multiple compartments). Tenancy and compartment administrators can utilize compartment quotas to set resource-specific limits. Without such compartment limits, a user that is authorized to launch instances can consume all available capacity in the entire tenancy. Compartment resource limits solve this problem and, unlike service limits, are set and customized by the clients and customers via, e.g., a console, SDK, or API. Compartment limits can be applied on top of the service limits and inherited through the nested compartment hierarchy. This allows compartment administrators to limit resource consumption and set boundaries around acceptable resource use.

[00011] In accordance with an embodiment, described herein are systems and methods for supporting a quota policy language in a cloud infrastructure environment. A quota policy can be configured by setting a statement syntax. In order to set a quota, the quota can establish a number of factors. Each quota can be unique within a service family (e.g. compute), and as such, each quota can define a target service along with a quota name. Next, a quota can define the value to set the quota to, as well as a compartment the quota targets. There can be a set of conditions that can be included determine when the quota is applied.

[00012] In accordance with an embodiment, described herein are systems and methods for supporting a usage calculation process or algorithm in a cloud infrastructure environment. The usage calculation process can be used to determine whether a requested transaction that targets a compartment within a tree-structure of compartments violates any compartment quota or limit within parent compartments within the tree-structure.

#### Brief Description of the Drawings:

[00013] Figure 1 illustrates a system for providing a cloud infrastructure environment, in accordance with an embodiment.

[00014] Figure 2 illustrates a system for providing cloud infrastructure regions within a cloud infrastructure environment, in accordance with an embodiment.

[00015] Figure 3 illustrates a system for providing compartment policies within compartments spanning cloud infrastructure regions within a cloud infrastructure environment, in accordance with an embodiment.

[00016] Figure 4 illustrates a system for compartment migration within a cloud infrastructure environment, in accordance with an embodiment.

[00017] Figure 5 shows a system for policy management and enforcement during compartment migration within a cloud infrastructure environment.

[00018] Figure 6 shows a system for supporting compartment quotas in a cloud

infrastructure environment, in accordance with an embodiment.

[00019] Figure 7 shows a system for supporting compartment quotas in a cloud infrastructure environment, in accordance with an embodiment.

5 [00020] Figure 8 is a flowchart of a method for creating a compartment quota, in accordance with an embodiment.

[00021] Figure 9 is a flowchart of a method for modifying a compartment quota, in accordance with an embodiment.

[00022] Figure 10 shows an architecture for a system to enforcing compartment quotas in a cloud infrastructure environment, in accordance with an embodiment.

10 [00023] Figure 11 is a flowchart of a method supporting compartment quotas in a cloud infrastructure environment, in accordance with an embodiment.

[00024] Figure 12a, 12b, and 12c are example compartment quota policy statements, in accordance with an embodiment.

15 [00025] Figure 13 illustrates an example of a quota policy statement, in accordance with an embodiment.

[00026] Figure 14 illustrates an example of a quota policy statement, in accordance with an embodiment.

[00027] Figure 15 illustrates an example of a quota policy statement, in accordance with an embodiment.

20 [00028] Figure 16 illustrates an example of a quota policy statement, in accordance with an embodiment.

[00029] Figure 17 illustrates an example of a quota policy statement, in accordance with an embodiment.

25 [00030] Figure 18 is a flowchart of a method for supporting a quota policy language in a cloud infrastructure environment, in accordance with an embodiment.

[00031] Figure 19 shows an architecture for a system to enforcing compartment quotas in a cloud infrastructure environment, in accordance with an embodiment.

[00032] Figure 20 is a flowchart of a method for supporting a usage calculation process in a cloud infrastructure environment, in accordance with an embodiment.

30 [00033] Figure 21 is a flowchart of a method for supporting a usage calculation process in a cloud infrastructure environment, in accordance with an embodiment.

#### **Detailed Description:**

35 [00034] In accordance with embodiments herein, the terms “bare metal host”, “bare metal instance”, or “bare metal compute instance” can refer to a physical host (“bare metal”) machine that is provided for use by a cloud infrastructure environment. Bare metal compute instances run directly on bare metal servers without a hypervisor. When a bare metal instance

is provisioned, a user/client can maintain sole control of the physical CPU, memory, and network interface card (NIC). The user/client can configure and utilize the full capabilities of each physical machine as if it were hardware running in their own data center. User/clients do not share the physical machine with any other tenants.

5 **[00035]** In accordance with embodiments herein, a cloud infrastructure environment can be physically hosted in regions and availability domains. The term “region” can refer to a localized geographic area. The term “availability domain” can mean one or more data centers located within a region. A region is composed of one or more availability domains. Cloud infrastructures can be region-specific, such as a virtual cloud network, or availability domain-  
10 specific, such as a compute instance. Availability domains are isolated from each other, fault tolerant, and unlikely to fail simultaneously or be impacted by the failure of another availability domain. When a user/client configures their cloud services, the user/client can use multiple availability domains to ensure high availability and to protect against resource failure. Some resources must be created within the same availability domain, such as an instance and the  
15 storage volume attached to it.

**[00036]** In accordance with embodiments herein, the term “realm” can refer to a logical collection of regions. Realms are isolated from each other and do not share any data. A client’s tenancy can exist in a single realm and have access to each of the regions that belong to that realm. Realms can be provided, for example, for commercial regions and government  
20 cloud regions.

**[00037]** In accordance with embodiments herein, the term “console” can refer to a user-interface (e.g., web-based) that users/clients can used to access and manage a cloud infrastructure environment.

**[00038]** In accordance with embodiments herein, the term “tenancy” can refer to a  
25 secure and isolated partition within a cloud infrastructure environment. When a client/user signs up for a cloud infrastructure environment, the cloud infrastructure environment can create a tenancy for the user/client (or their company). A tenancy can be where a user/client creates, organizes, and administers their cloud resources.

**[00039]** In accordance with embodiments herein, the term “compartment” or  
30 “compartments” can refer to a collection of related resources (such as instances, virtual cloud networks, block volumes) that can be accessed only by certain groups that have been given permission by an administrator. Compartments allow a client/user to organize and control access to their cloud resources. When a client/user begins working with resources in the console, the compartment can provide a filter for what the client/user is viewing.

35 **[00040]** In accordance with embodiments herein, the term “tenancy” can refer to a root compartment that comprises all of a tenant’s cloud resources. Additional compartments can be created within a tenancy (root compartment) and corresponding policies to control access

to the resources in each compartment. When a user/client creates a cloud resource, such as an instance, block volume, or cloud network, the user/client can specify to which compartment the resource should belong to.

**[00041]** In accordance with embodiments herein, the term “virtual cloud network” can refer to a virtual version of a traditional network—including subnets, route tables, and gateways—on which instances run. A cloud network resides within a single region but includes all the region's availability domains. Each subnet that is defined in the cloud network can either be in a single availability domain or span all the availability domains in the region. At least one least one cloud network can be set up before instances can be launched.

**[00042]** In accordance with embodiments herein, the term “instance” can refer to a compute host running in the cloud. A cloud infrastructure environment instance can allow a user/client to utilize hosted physical hardware, as opposed to traditional software-based virtual machines, thereby providing increased security and performance.

**[00043]** In accordance with embodiments herein, the term “image” can refer to a template of a virtual hard drive that defines the operating system and other software for an instance, for example, LINUX. When an instance is launched, its characteristics can be defined by choosing its image. Sets of images can be provided, and images can be used from existing instances that have already been configured (e.g., as a template) in order to use instances with the same software and customizations.

**[00044]** In accordance with embodiments herein, the term “shape” can refer to a specification of a number of CPUs and an amount of memory allocated to an instance. In load balancing, the shape can determine the load balancer's total pre-provisioned maximum capacity (bandwidth) for ingress plus egress traffic. Available shapes can include, for example, 100 Mbps, 400 Mbps, and 8000 Mbps.

**[00045]** In accordance with an embodiment, the term “key pair” can refer to an authentication mechanism used a cloud infrastructure environment. A key pair consists of a private key file and a public key file. A public key can be uploaded, while a private key is kept on a user's/client's device. The private key is private to a user/client, like a password. Key pairs can be generated according to different specifications, such as Instance SSH Key pair, and API signing key pair. The Instance SSH key pair is used to establish secure shell (SSH) connection to an instance. When an instance is provisioned, a public key can be provide, which is saved to the instance's authorized key file. To log on to the instance, a private key is provided, which is verified with the public key. The API signing key pair is a PEM (privacy enhanced mail) format and is used to authenticate a user/client when submitting API requests.

**[00046]** In accordance with embodiments herein, the term “block volume” can refer to a virtual disk that provides persistent block storage for a cloud infrastructure environment instance. A block volume can be used in a similar manner as, e.g., a physical hard drive on a

computer to, e.g., store data and applications. A volume can be detached from one instance and attached to another instance without loss of data.

**[00047]** In accordance with embodiments here, the term “object storage” can refer to a storage architecture can be used to store and manage data as objects. Data files can be of any type. Once data is uploaded to object storage, it can be accessed from anywhere. Object storage can be used when a large amount of data needs to be stored, and the data does not change very frequently. Typical uses for object storage include, but are not limited to, data backup, file sharing, storing unstructured data (e.g., logs and sensor-generated data).

**[00048]** In accordance with embodiments herein, the term “bucket” can refer to a logical container for storing data and files. A bucket can contain an unlimited number of objects.

**[00049]** As described above, cloud infrastructure environments can comprise sets of complementary cloud services that enable users and clients to build and run a wide range of applications and services in a highly available hosted environment.

**[00050]** **Figure 1** shows a system for providing a cloud infrastructure environment, in accordance with an embodiment.

**[00051]** In accordance with an embodiment, a cloud infrastructure environment 100, which can be run on a number of hardware and software resources 112, can comprise a console interface 102 and an API 104. In addition, the cloud infrastructure environment 100 can support a number of governance services 110, an identity and access management (IAM) service 120, and a provisioning service 130. The cloud infrastructure environment 100 can also support a number of resources 140, e.g., in layers, such as a computer resource layer 150, a network resource layer 160, and a storage resource layer 170.

**[00052]** In accordance with an embodiment, a client device, such as a computing device 10 having device hardware (processor, memory...etc.) 12, can communicate with the cloud infrastructure environment via a network, such as a wide area network (WAN), a local area network (LAN), or the internet, for example. The client device can comprise an administrator application 14, which can comprise a user interface 16.

**[00053]** In accordance with an embodiment, within the cloud infrastructure environment, tenancy can be supported. On registration and deployment, a tenancy can be created for each client/customer, which can comprise a secure and isolated partition within the cloud infrastructure in which the client can create, organize, and administer their cloud resources.

**[00054]** In accordance with an embodiment, the console interface 102 and the API 104 can provide clients with access to, and control over respective portions of the cloud infrastructure environment. In accordance with an embodiment, the console interface can comprise an intuitive, graphical interface that lets clients create and manage resources, instances, cloud networks, and storage volumes, as well as manage users associated with the client, and set permissions within the client scope. As well, the API 104 can compromise, for

example, a REST API that utilizes HTTPS (hypertext transfer protocol secure).

**[00055]** In accordance with an embodiment, one example of a console interface or API can be a configuration management tool (e.g., Ansible). The configuration management tool can be used for cloud infrastructure provisioning, orchestration, and configuration management. Configuration management tools can allow clients to automate configuring and provisioning of the cloud infrastructure, deploying and updating software assets, and orchestrating complex operational processes.

**[00056]** In accordance with an embodiment, the governance services 110 of the cloud infrastructure environment provides clients tools to help clients enable simple resource governance, manage costs, and control access to the cloud infrastructure. As an example, the governance services provide for tagging which can allow for clients to apply tags to their resources for informational or operational reasons. Defined tags can be controlled to avoid incorrect tags from being applied to resources. Tags can also provide a flexible targeting mechanism for administrative scripts. As well, the governance services can allow for managed budgets, and track actual and forecasted spend all from one place. This allows clients to stay on top of usage with a cost analysis dashboard, and filter by compartments and tags to analyze spending by departments, teams, and projects. Such data can as well be exported for detailed resource utilization reporting and integration with an existing cloud management and business intelligence tools. The governance services can also log events that can later be retrieved, stored, and analyzed for security, compliance, and resource optimization across the cloud infrastructure entitlements and compartments.

**[00057]** In accordance with an embodiment, the identity and access management (IAM) service 120 can create a user profile for each client/customer/user in the IAM service with associated with user credential (e.g., username and password). Clients can be granted administrator privileges in the cloud infrastructure as well via the IAM service.

**[00058]** In accordance with an embodiment, the identity and access management service can be integrated with the cloud infrastructure environment. Upon a client registering. The IAM service can create a separate user credential in an identity service, which can then allow for single sign on to the cloud infrastructure service as well as access to additional cloud services.

**[00059]** In accordance with an embodiment, the provisioning service 130 can provision, for example, a tenancy within cloud infrastructure service, such as within the resources 140. The provisioning service can be accessed and controller through, for example, the console interface or via one or more APIs, such as API 104. The provisioning service can allow for lets clients to provision and manage compute hosts, which can be referred to as instances. Clients can launch instances as needed to meet compute and application requirements. After a client launches an instance, the provisioned instance can be accessed from, for example, a client

device. The provisioning service can also provide for restarting an instance, attaching and detaching volumes from an instance, and terminating an instance.

**[00060]** In accordance with an embodiment, resources 140 provided by an cloud infrastructure environment can be broken down into a plurality of layers, such as a compute resources layer 150, a network resources layer 160, and a storage resource layer 170.

**[00061]** In accordance with an embodiment, the compute resources layer 150 can comprise a number of resources, such as, for example, bare metal instances 152, virtual machines 154, edge services 156, and containers 158. The compute resources layer can be used to, for example, provision and manage bare metal compute instances, provision instances as needed to deploy and run applications, just as in an on-premises data center.

**[00062]** In accordance with an embodiment, the cloud infrastructure environment can provide control of one or more physical host (“bare metal”) machines within the compute resources layer. Bare metal compute instances run directly on bare metal servers without a hypervisor. When a bare metal compute instance is provisioned, the client can maintain sole control of the physical CPU, memory, and network interface card (NIC). The bare metal compute instance can be configured and utilize the full capabilities of each physical machine as if it were hardware running in an on-premise own data center. As such, bare metal compute instances are generally not shared between tenants.

**[00063]** In accordance with an embodiment, bare metal compute instances can provide, via the associated physical hardware as opposed to a software-based virtual environment, a high level of security and performance.

**[00064]** In accordance with an embodiment, the cloud infrastructure environment can provide control of a number of virtual machines within the compute resources layer. A virtual machine compute host can be launched, for example, from an image that can determine the virtual machines operation system as well as other software. The types and quantities of resources available to a virtual machine instance can be determined, for example, based upon the image that the virtual machine was launched from.

**[00065]** In accordance with an embodiment, a virtual machine (VM) compute instance can comprise an independent computing environment that runs on top of physical bare metal hardware. The virtualization makes it possible to run multiple VMs that are isolated from each other. VMs can be used, for example, for running applications that do not require the performance and resources (CPU, memory, network bandwidth, storage) of an entire physical machine.

**[00066]** In some embodiments, virtual machine instances can run on the same hardware as a bare metal instance, which can provide leverage over using the same cloud-optimized hardware, firmware, software stack, and networking infrastructure

**[00067]** In accordance with an embodiment, the cloud infrastructure environment can

provide a number of graphical processing unit (GPU) compute instances within the compute resources layer. Accelerated computing requires consistently-fast infrastructure across every service. With GPU instances, clients can process and analyze massive data sets more efficiently, making them useful for complex machine learning (ML), artificial intelligence (AI) algorithms, and many industrial HPC applications. GPU compute instances can be provisioned as either virtualized compute instances (where multiple GPU compute instances share the same bare metal hardware), or as bare metal instances which provide dedicate hardware for each GPU compute instance.

**[00068]** In accordance with an embodiment, the cloud infrastructure environment can provide a number of containerized compute instances within the compute resources layer. A standalone container engine service can be used to build and launch containerized applications to the cloud. The container service can be used, for example, to build, deploy, and manage cloud-native applications. The container service can specify the compute resources that the containerized applications require, and the container engine can then provision, via the provisioning service, the required compute resources for use within the cloud infrastructure environment (e.g., in the context of a tenancy).

**[00069]** In accordance with an embodiment, one such container service engine that can be used is Kubernetes, an open-source system for automating deployment, scaling, and management of containerized applications across clusters of hosts. Such container services can group the containers that make up an application into logical units for easy management and discovery.

**[00070]** In accordance with an embodiment, the network resources layer 160 can comprise a number of resources, such as, for example, virtual cloud networks (VCNs) 162, load balancers 164, edge services 166, and connection services 168.

**[00071]** In accordance with an embodiment, the cloud infrastructure environment can provide a number of virtual cloud networks 162 at the networking resources layer. A virtual cloud network can comprise a virtual version of a traditional network—including subnets, route tables, and gateways—on which client instances can run. A cloud network resides within a single region but includes all the region's availability domains. Each subnet defined in the cloud network can either be in a single availability domain or span all the availability domains in the region (recommended). At least one cloud network can be configured before launching instances. In certain embodiments, VCNs can be configured via an internet gateway to handle public traffic, a VPN connection, or a fast connect service to securely extend on-premises network.

**[00072]** In accordance with an embodiment, the cloud infrastructure environment can provide a number of load balancers 164 at the networking resources layer. A load balancing service can provide automated traffic distribution from one entry point to multiple servers

reachable from a virtual cloud network (VCN). Various load balances can provide a public or private IP address, and provisioned bandwidth.

**[00073]** In accordance with an embodiment, a load balancer can improve resource utilization, scaling, and help ensure high availability. Multipole load balancing policies can be configured, and application-specific health checks can be provided to ensure that the load balancer directs traffic only to healthy instances. The load balancer can reduce maintenance window by draining traffic from an unhealthy application server before it is removed from service for maintenance.

**[00074]** In accordance with an embodiment, a load balancing service enables creation of a public or private load balancer in conjunction with a VCN. A public load balancer has a public IP address that is accessible from the internet. A private load balancer has an IP address from the hosting subnet, which is visible only within the VCN. Multiple listeners can be configured for an IP address to load balance transport different layers of traffic (e.g., Layer 4 and Layer 7 (TCP and HTTP) traffic). Both public and private load balancers can route data traffic to any backend server that is reachable from the VCN.

**[00075]** In accordance with an embodiment, a public load balancer can accept traffic from the internet, a public load balance can be created that is assigned a public address, which serves as the entry point for incoming traffic.

**[00076]** In accordance with an embodiment, a public load balancer is regional in scope. If a region includes multiple availability domains, a public load balancer can have, for example, a regional subnet, or two availability domain-specific (AD-specific) subnets, each in a separate availability domain. With a regional subnet, the load balancer can create a primary load balancer and a standby load balancer, each in a different availability domain, to ensure accessibility even during an availability domain outage. If a load balance is created in multiple AD-specific subnets, one subnet can host the primary load balancer and the other hosts a standby load balancer. If the primary load balancer fails, the public IP address can switch to the secondary load balancer. The service treats the two load balancers as equivalent.

**[00077]** In accordance with an embodiment, if a region includes only one availability domain, the service requires just one subnet, either regional or AD-specific, to host both the primary and standby load balancers. The primary and standby load balancers can each have a private IP address from the host subnet, in addition to the assigned floating public IP address. If there is an availability domain outage, the load balancer has no failover.

**[00078]** In accordance with an embodiment, private load balances can also be provided so as to isolate the load balancer from the internet and simplify security posture. The load balancer service can assign a private address to the load balancer that serves as the entry point for incoming traffic.

**[00079]** In accordance with an embodiment, a private load balancer can be created by

a service to service only one subnet to host both the primary and standby load balancers. The load balancer can be regional or AD-specific, depending on the scope of the host subnet. The load balancer is accessible only from within the VCN that contains the host subnet, or as further restricted by security rules.

5 [00080] In accordance with an embodiment, the assigned floating private IP address is local to the host subnet. The primary and standby load balancers each require an extra private IP address from the host subnet.

[00081] In accordance with an embodiment, if there is an availability domain outage, a private load balancer created in a regional subnet within a multi-AD region provides failover  
10 capability. A private load balancer created in an AD-specific subnet, or in a regional subnet within a single availability domain region, has no failover capability in response to an availability domain outage.

[00082] In accordance with an embodiment, the cloud infrastructure environment can provide a number of edge services 166 at the networking resources layer. In general, edge  
15 services comprise a number of services that allow clients to manage, secure, and maintain domains and endpoints. These include, for example, DNS (domain name system), DDoS (distributed denial of service) protection, and email delivery. These services enable clients to optimize performance, thwart cyberattacks, and scale communication.

[00083] In accordance with an embodiment, the cloud infrastructure environment can  
20 provide a number of connection services 168 at the networking resources layer. Such connection services can provide an easy way to create a dedicated, private connection between a client data center or existing network and the cloud infrastructure environment. The connection service can provide high bandwidth, and a reliable and consistent network.

[00084] In accordance with an embodiment, the storage resources layer 170 can  
25 comprise a number of resources, such as, for example, block volumes 172, file storage 174, object storage 176, and local storage 178.

[00085] In accordance with an embodiment, block volumes 172 provide high-  
performance network storage capacity that supports a broad range of I/O intensive workloads. Clients can use block volumes to expand the storage capacity of compute instances, to provide  
30 durable and persistent data storage that can be migrated across compute instances, and to host large databases.

[00086] In accordance with an embodiment, file storage 174 allows clients to create a  
scalable, distributed, enterprise-grade network file system. File storage supports semantics, snapshots capabilities, and data at-rest encryption.

35 [00087] In accordance with an embodiment, object storage provides high throughput storage for unstructured data. Object storage service enables near limitless storage capacity for large amounts of analytic data, or rich content like images and videos. Block volumes can

be backed up to object storage for added durability.

**[00088]** In accordance with an embodiment, local storage 178 can provide, for example, high speed and reliable storage in the form of solid state drives for I/O intensive applications. These can be provided, for example, within bare metal instances. Local storage provides high storage performance for VM's and bare metal compute instances. Some examples include relational databases, data warehousing, big data, analytics, AI and HPC application.

**[00089]** **Figure 2** illustrates a system for providing cloud infrastructure regions within a cloud infrastructure environment, in accordance with an embodiment.

**[00090]** In accordance with an embodiment, instances of the cloud infrastructure environment described above in Figure 1 can be hosted in different regions, called cloud infrastructure regions 220. These can be accessed, as described above, via a console, SDK, or APIs, by customer networks 200 via a network, such as the internet 210. Each cloud infrastructure region can comprise management services 222, compute services 224, storage services 226, edge serves 228, network services 230, and physical infrastructure 232.

**[00091]** In accordance with an embodiment, a cloud infrastructure can be hosted in regions and availability domains. A region can be a localized geographic area, and an availability domain can be one or more data centers located within a region. A region is composed of one or more availability domains. Most cloud infrastructure resources can be either region-specific, such as a virtual cloud network, or availability domain-specific, such as a compute instance. Traffic between availability domains and between regions is encrypted.

**[00092]** In accordance with an embodiment, availability domains are isolated from each other, fault tolerant, and very unlikely to fail simultaneously. Because availability domains do not share infrastructure such as power or cooling, or the internal availability domain network, a failure at one availability domain within a region is unlikely to impact the availability of the others within the same region.

**[00093]** In accordance with an embodiment, availability domains within the same region can be connected to each other by a low latency, high bandwidth network, which can provide high-availability connectivity to the internet and on-premises, and to build replicated systems in multiple availability domains for both high-availability and disaster recovery.

**[00094]** In accordance with an embodiment, regions are independent of other regions and can be separated geographically (e.g., across countries or continents). This then leads to the deployment of an application within a region where the application would most likely be utilized the most frequently.

**[00095]** In accordance with an embodiment, however, applications can also be deployed in different regions for various reasons. This can include, for example, risk mitigation when events, such as weather systems, take a region offline. In addition, applications can be deployed in other regions for strategic reasons, such as tax domains or other business or social

criteria.

**[00096]** In accordance with an embodiment, there are several services that are available across regions. These include, for example, management services 222, compute services 224, storage services 226, edge services 228, and network services 230.

5 **[00097]** In accordance with an embodiment, compartments allow clients to organize and control access to cloud resources. A compartment is a collection of related resources (such as instances, virtual cloud networks, block volumes) that can be accessed only by certain groups that have been given permission by an administrator. A compartment can be thought of as a logical group and not a physical container. When working within a console, a compartment can  
10 act as a filter for what is allowed to be viewed.

**[00098]** In accordance with an embodiment, compartments can have several layers. For example, a tenancy can be considered a root compartment that holds all of a client's cloud resources. Additional compartments can be created within that root compartment (tenancy) and corresponding policies to control access to the resources in each compartment. When  
15 clients create a cloud resource such as an instance, block volume, or cloud network, such a resource can be directed to a specific compartment or compartments. Compartments can span regions.

### **Fault Domains**

20 **[00099]** In accordance with an embodiment, a fault domain can comprise a grouping of hardware and infrastructure within an availability domain. Each availability domain can comprise three fault domains. Fault domains allow instances to be distributed so that they are not on the same physical hardware within a single availability domain. A hardware failure or Compute hardware maintenance that affects one fault domain does not affect instances in  
25 other fault domains.

**[000100]** In accordance with an embodiment, placement of resources, such as compute, bare metal DB system, or virtual machine DB system instances, can optionally specify a fault domain or a new instance at launch time. The resources can additionally change fault domains after placement by terminating the resource at the current fault domain and launching a new  
30 instance of the resource at another fault domain.

**[000101]** In accordance with an embodiment, fault domains can be utilized for a number of reasons, such as protecting against unexpected hardware failures and protecting against planned outages due to maintenance.

### 35 **Availability**

**[000102]** In accordance with an embodiment, service availability can be provided. Regions within cloud infrastructure environments can provide core infrastructure services and

resources, including the following:

- Compute: Compute (Bare Metal & VM, DenseIO & Standard), Container Engine for Kubernetes, Registry
- Storage: Block Volume, File Storage, Object Storage, Archive Storage
- 5 • Networking: Virtual Cloud Network, Load Balancing, FastConnect
- Database: Database, Exadata Cloud Service, Autonomous Data Warehouse, Autonomous Transaction
- Processing
- Edge: DNS
- 10 • Platform: Identity and Access Management, Tagging, Audit

**[000103]** In accordance with an embodiment, the above services and resources can be generally available, while other services and resources can additionally be available as well (e.g., based upon regional demand or customer request). As an example, new cloud services can be made available in regions as quickly based on a variety of considerations including regional customer demand, ability to achieve regulatory compliance where applicable, resource availability, and other factors. Because of low latency interconnect backbone, customers can use cloud services in other geographic regions with effective results when they are not available in their home region, provided that data residency requirements do not prevent them from doing so.

**[000104]** In accordance with an embodiment, resource availability can be considered in the context of global availability, regional availability, single region availability, and domain availability. Generally speaking, IAM resources are globally available, DB systems, instances, and volumes are specific to a viability domain. Most other resources are regional.

**[000105]** In accordance with an embodiment, examples of globally available resources can include API signing keys, compartments, dynamic groups, federation resources, groups, policies, tag namespaces, tag keys, and users.

**[000106]** In accordance with an embodiment, examples of regionally available resources can include, alarms, applications, buckets (although buckets are regional resources, they can be accessed from any location when the correct region-specific Object Storage URL for the API calls is used), clusters, cloudevents-rules, customer-premises equipment (CPE), DHCP options sets, dynamic routing gateways (DRGs), encryption keys, functions, images, internet gateways, jobs, key vaults, load balancers, local peering gateways (LPGs), metrics, NAT gateways, network security groups, node pools, ons-subscriptions, ons-topics, repositories, reserved public Ips, route tables, security lists, service gateways, stacks, subnets (when a subnet is created, it can be declared to be a regional or specific to an availability domain), virtual cloud networks (VCNs), and volume backups (volume backups can be restored as new

volumes to any availability domain within the same region in which they are stored).

**[000107]** In accordance with an embodiment, examples of availability domain-specific resources can include DB Systems, ephemeral public Ips, instances (instances can be attached to volumes in the same availability domain), subnets (when a subnet is created, it can be declared to be a regional or specific to an availability domain), and volumes (volumes can be attached to an instance in a same availability domain).

## Compartments

**[000108]** In accordance with an embodiment, administrators can manage compartments within a cloud infrastructure environment.

**[000109]** In accordance with an embodiment, tags can be applied to resources within a compartment. Tags can be used to, for example, organize resources according a schema, such as a business needs schema. Tags can be applied to resources at the time of creation of a resource, or a tag can be updated on an existing resource.

**[000110]** In accordance with an embodiment, compartments are important to the construction and organization of a cloud infrastructure. Resources can be moved between compartments, and resources can be displayed (e.g., via a user interface) organized by compartment within a current region. When working with and managing resources, a compartment can first be selected.

**[000111]** In accordance with an embodiment, compartments are tenancy-wide, and can span across regions. When a compartment is created, the compartment can be made available within every region that a tenancy is subscribed to.

**[000112]** In accordance with an embodiment, compartments can be deleted. In order for a compartment to be deleted, the compartment can have all resources therein removed prior to deletion.

**[000113]** In accordance with an embodiment, the action to delete a compartment can be asynchronous and initiates a work request. The state of the compartment changes to "Deleting" while the work request is executing. If the work request fails, the compartment is not deleted and it returns to the active state.

**[000114]** In accordance with an embodiment, each compartment created within the cloud infrastructure environment can have certain properties. For example, each compartment can be assigned a unique identifier (ID), and can additionally, and optionally, be provided with a modifiable description, as well as a name. In accordance with an embodiment, sub-compartments (or subcompartments) can be defined in a hierarchical manner under a base compartment.

**[000115]** In accordance with an embodiment, access and control over compartments and subcompartments can be limited to administrators or other users with sufficient credentials.

Credentials can be associated with differing levels of compartment access. For example, an administrator can have permission to view and access all compartments and work with resources within any compartment of a tenancy, but a user with more limited access will not have such a level of access and control.

5 [000116] **Figure 3** illustrates a system for providing compartment policies within compartments spanning cloud infrastructure regions within a cloud infrastructure environment, in accordance with an embodiment.

[000117] In accordance with an embodiment, as described above, instances of the cloud infrastructure environment 300 described above in Figure 1 can be hosted in different regions, such as cloud infrastructure regions 320, 330, 340, 350. These can be accessed, as described above, via a console, SDK, or APIs, by customer networks 301 via a network 302, such as the internet.

[000118] In accordance with an embodiment, a customer network 301 can comprise, for example, a single computer, a network of customer computers, or other such networks.

15 [000119] In accordance with an embodiment, although not shown in the Figure, each cloud infrastructure region can comprise a number of services, each comprising a number of resources, such as management services, compute services, storage services, edge serves, network services, and physical infrastructure.

[000120] In accordance with an embodiment, a cloud infrastructure can be hosted in regions and availability domains. A region can be a localized geographic area, and an availability domain can be one or more data centers located within a region. A region is composed of one or more availability domains. Most cloud infrastructure resources can be either region-specific, such as a virtual cloud network, or availability domain-specific, such as a compute instance. Traffic between availability domains and between regions is encrypted.

25 [000121] In accordance with an embodiment, availability domains are isolated from each other, fault tolerant, and very unlikely to fail simultaneously. Because availability domains do not share infrastructure such as power or cooling, or the internal availability domain network, a failure at one availability domain within a region is unlikely to impact the availability of the others within the same region.

30 [000122] In accordance with an embodiment, availability domains within the same region can be connected to each other by a low latency, high bandwidth network, which can provide high-availability connectivity to the internet and on-premises, and to build replicated systems in multiple availability domains for both high-availability and disaster recovery.

[000123] In accordance with an embodiment, regions are independent of other regions and can be separated geographically (e.g., across countries or continents). This then leads to the deployment of an application within a region where the application would most likely be utilized the most frequently.

**[000124]** In accordance with an embodiment, however, applications can also be deployed in different regions for various reasons. This can include, for example, risk mitigation when events, such as weather systems, take a region offline. In addition, applications can be deployed in other regions for strategic reasons, such as tax domains or other business or social criteria.

**[000125]** In accordance with an embodiment, there are several services that are available across regions. These include, for example, management services, compute services, storage services, edge services, and network services.

**[000126]** In accordance with an embodiment, compartments allow clients to organize and control access to cloud resources. A compartment is a collection of related resources (such as instances, virtual cloud networks, block volumes) that can be accessed only by certain groups that have been given permission by an administrator. A compartment can be thought of as a logical group and not a physical container. When working within a console, a compartment can act as a filter for what is allowed to be viewed.

**[000127]** In accordance with an embodiment, compartments can have several layers. For example, a tenancy 305 can be considered a root compartment that holds all of a client's cloud resources. Compartments can be organized in a hierarchical manner, such as compartment 360 being a level below the tenancy compartment, with sub compartment 370 being an additional layer below the compartment 360. In accordance with an embodiment, each compartment can be associated with one or more compartment policies, such as compartment policy 365, and sub compartment policy 375. Tenant compartment policy is not shown in the Figure.

**[000128]** In accordance with an embodiment, during, upon, or after creation of a compartment, or sub compartment, such as compartment 360 and sub compartment 370, a policy, such as compartment policy 365 and sub compartment policy 375 can be written/created for each compartment and sub compartment. Compartment policies can, for example, restrict which users are allowed to access/utilize resources and services within a compartment. Without a policy in place, access to the compartments and/or sub compartments can be restricted to users having permissions at the tenancy 305 level.

**[000129]** In accordance with an embodiment, upon creation of a compartment within a compartment (i.e., a sub compartment), the sub compartment inherits access permissions from compartments higher up its hierarchy.

**[000130]** In accordance with an embodiment, upon creation of a compartment or sub compartment policy, the policy can comprise a specification indicating which compartment the policy attaches to. Such a specification can contain controls limiting access for subsequent control, modification, or deletion of the policy. In some embodiments, the policies can be attached to a tenancy, a parent compartment, or the specific compartment to which the policy

is directed.

[000131] In accordance with an embodiment, resources can be placed into a compartment. This can be accomplished by specifying the targeted compartment upon creation of a resource (the compartment is one of the required pieces of information to create  
5 a resource). This can be accomplished via a console interface.

[000132] In accordance with an embodiment, existing resources can also be moved to different compartments. Most resources can be moved after they are created.

[000133] In accordance with an embodiment, some resources have attached resource dependencies and some do not. Not all attached dependencies behave the same way when  
10 the parent resource moves.

[000134] In accordance with an embodiment, for some resources, the attached dependencies move with the parent resource to the new compartment. The parent resource moves immediately, but in some cases attached dependencies move asynchronously and are not visible in the new compartment until the move is complete.

[000135] In accordance with an embodiment, for other resources, the attached resource dependencies do not move to the new compartment. Such attached resources can be moved independently.  
15

[000136] In accordance with an embodiment, after a resource is moved to a new compartment, the policies that govern the new compartment apply immediately and affect  
20 access to the resource. Depending on the structure of the compartment organization, metering, billing, and alarms can also be affected.

[000137] In accordance with an embodiment, after creation, a compartment can be moved to, e.g., a different parent compartment within a same tenancy. Upon moving a compartment, all of the compartment's contents (including sub compartments and resources)  
25 are moved along with the compartment.

[000138] **Figure 4** illustrates a system for compartment migration within a cloud infrastructure environment, in accordance with an embodiment.

[000139] In accordance with an embodiment, as described above, instances of the cloud infrastructure environment 400 described above in Figure 1 can be hosted in different regions.  
30 Compartments, such as tenancy 405, compartment A 460 and compartment B 470, can be defined within the cloud infrastructure environment, and these compartments can span across regions. Such compartments can be accessed, as described above, via a console, SDK, or APIs, by customer networks 401 via a network 402, such as the internet.

[000140] In accordance with an embodiment, a customer network 401 can comprise, for  
35 example, a single computer, a network of customer computers, or other such networks.

[000141] In accordance with an embodiment, compartments allow clients to organize and control access to cloud resources. A compartment is a collection of related resources (such as

instances, virtual cloud networks, block volumes) that can be accessed only by certain groups that have been given permission by an administrator. A compartment can be thought of as a logical group and not a physical container. When working within a console, a compartment can act as a filter for what is allowed to be viewed.

5 **[000142]** In accordance with an embodiment, compartments can have several layers. For example, a tenancy 405 can be considered a root compartment that holds all of a client's cloud resources. Compartments can be organized in a hierarchical manner, such as compartment A 460 and compartment B 470 being a level below the tenancy compartment, with sub compartment A 465 being defined below compartment A, and sub compartment B  
10 475 being defined below compartment B. In accordance with an embodiment, each compartment can be associated with one or more compartment policies (not shown).

**[000143]** In accordance with an embodiment, compartments defined within a tenancy, for example, can be moved by, for example, re-defining a compartment or sub-compartment.

**[000144]** In accordance with an embodiment, in order to move a compartment, a request  
15 with sufficient permissions can be received. That is, a request from a user belonging to a group that has, for example, a "manage all-resources" permissions on the lowest shared parent compartment to the current compartment and the destination compartment of the moving compartment.

**[000145]** That is, for example, a request to move sub-compartment A 465 from (move  
20 from 465 to 476) compartment A 460 to compartment B 470 must be received from a user with sufficient permissions. Because the tenancy 405 is the lowest shared parent compartment of both the source compartment, compartment A 460, and the destination compartment, compartment B 470, then the request to move sub-compartment A, as shown in the Figure, must be received from a user having "manage all-resources" permissions within the tenancy  
25 405 compartment.

**[000146]** In accordance with an embodiment, in another example, if the request to move sub-compartment A 465 from compartment A to compartment B was received from a user having "manage all-resources" permissions within compartment A only, then the request may fail as the request from the user cannot manage resources within the destination compartment,  
30 namely compartment B.

**[000147]** In accordance with an embodiment, upon moving a compartment to a new parent compartment, the access policies of the new parent take effect and the policies of the previous parent compartment no longer apply. In some cases, when moving nested compartments with policies that specify the hierarchy, the policies can be  
35 automatically updated to ensure consistency.

**[000148]** In accordance with an embodiment, therefore, a compartment policy of compartment A 460 which was previously applied to sub-compartment A would no longer apply

on migration of the sub-compartment A to compartment B. Then, a compartment policy of compartment B would apply to sub-compartment A instead. This is explained more in the description following Figure.

**[000149]** Figure 5 shows a system for policy management and enforcement during  
5 compartment migration within a cloud infrastructure environment.

**[000150]** In accordance with an embodiment, and more specifically, Figure 5 shows a compartment hierarchy in which a compartment is moved, and the consequences for different policies.

**[000151]** In accordance with an embodiment, as described above, instances of the cloud  
10 infrastructure environment 500 described above in Figure 1 can be hosted in different regions. Compartments, such as tenancy 505, compartment A 560 and compartment B 565, and compartment D 566 can be defined within the cloud infrastructure environment, and these compartments can span across regions. Such compartments can be accessed, as described above, via a console, SDK, or APIs, by customer networks 501 via a network 502, such as the  
15 internet.

**[000152]** In accordance with an embodiment, a customer network 501 can comprise, for example, a single computer, a network of customer computers, or other such networks.

**[000153]** In accordance with an embodiment, compartments can have several layers. For example, a tenancy 505 can be considered a root compartment that holds all of a client's  
20 cloud resources. Compartments can be organized in a hierarchical manner, such as compartment A 560 being a level below the tenancy. Compartments B 565 and compartment D 566 are then organized as being yet another level below compartment A 560, while compartment C 570 is shown as being originally a level below compartment B. In accordance with an embodiment, each compartment can be associated with one or more compartment  
25 policies, such as compartment B policy 582, compartment A policy 580, and compartment D policy 583. Such policies can govern, for example, user/client permissions for access to the compartments, as well as permissions for access to and control of resources within any given compartment. As described above, compartment policies can add to each other (i.e., "stack"), such that a user accessing compartment B 565 would have their interactions with compartment  
30 B 565 being governed by/limited by compartment B policy 582 in addition to compartment A policy 580.

**[000154]** In accordance with an embodiment, for example, suppose that compartment B policy 582 allows a group, group 1, to manage the instance-family in compartment A-B (the compartment hierarchy comprising compartment B being a sub compartment of compartment  
35 A).

**[000155]** In accordance with an embodiment, suppose also that compartment D policy 583 allows another group, group 2, to manage the instance family in compartment A-D (the

compartment hierarchy comprising compartment D being a sub compartment of compartment A).

[000156] In accordance with an embodiment, upon compartment C being moved from compartment B to compartment D, members of group 1 can no longer manage instance families in compartment C, while members of group 2 can now manage instance families in compartment C.

[000157] In accordance with an embodiment, in certain situations, upon moving a compartment, certain policies can be automatically updated. Policies, for example, that specify the compartment hierarchy down to the compartment being moved can be automatically be updated when the policy is attached to a shared ancestor of the current and target parent.

[000158] Referring back to Figure 5, for example, in accordance with an embodiment, suppose that compartment A policy allows members of a group, group X, to manage buckets in compartment B:C. On moving compartment C to compartment D, because of the shared ancestor (compartment A) between compartments B and D, then the compartment A policy can be automatically updated to allow members of group X to manage buckets in compartment D:C.

[000159] In accordance with an embodiment, polices attached to tenancies can be likewise automatically updated upon a compartment moving within the tenancy.

[000160] In accordance with an embodiment, however, not all policies are automatically updated upon a compartment moving. For example, in referring to Figure 5, in the situation where compartment C is moved from compartment B to compartment D. Suppose that the compartment B policy allows management of buckets in compartment C (prior to moving). When compartment C is moved, then, compartment B policy is not automatically updated. Instead, the policy is no longer valid and can be removed (e.g., manually or automatically).

### **Compartment Quotas**

[000161] In accordance with an embodiment, cloud administrators do not generally have the ability to restrict resource usage in existing clouds. Granting a user permission to create resources allows them to create any number of resources up to a predefined account limit. Compartment quotas allow restrictions to be placed on the ability to create or use resources within a compartment to the appropriate level allowing fine-tuned cost control.

[000162] In accordance with an embodiment, customers can be assigned service level limits defined by the cloud infrastructure environment at account creation time. These service level limits restrict the total number of resources a customer can create across the entire tenancy (e.g., across multiple regions with multiple compartments). Tenancy and compartment administrators can utilize compartment quotas to set resource-specific hard limits. Without such compartment limits, a user that is authorized to launch instances can consume all

available capacity in the entire tenancy. Compartment resource limits solve this problem and, unlike service limits, are set and customized by the clients and customers via, e.g., a console, SDK, or API. Compartment limits can be applied on top of the service limits and inherited through the nested compartment hierarchy. This allows compartment administrators to limit resource consumption and set boundaries around acceptable resource use.

5 [000163] In accordance with an embodiment, compartment quotas give tenant and compartment administrators better control over how resources are consumed in a cloud infrastructure environment, enabling administrators to easily allocate resources to compartments using, for example, a console, SDK, or API. Compartment quotas are a powerful toolset to manage client spending in within tenancies.

[000164] Figure 6 shows a system for supporting compartment quotas in a cloud infrastructure environment, in accordance with an embodiment.

15 [000165] In accordance with an embodiment, as described above, instances of the cloud infrastructure environment 600 described above in Figure 1 can be hosted in different regions, such as cloud infrastructure regions 620, 630, 640, 650. These can be accessed, as described above, via a console 603, SDK, or APIs, by customer networks 601 via a network 602, such as the internet.

[000166] In accordance with an embodiment, a customer network 601 can comprise, for example, a single computer, a network of customer computers, or other such networks.

20 [000167] In accordance with an embodiment, although not shown in the Figure, each cloud infrastructure region can comprise a number of services, each comprising a number of resources, such as management services, compute services, storage services, edge serves, network services, and physical infrastructure.

[000168] In accordance with an embodiment, a cloud infrastructure can be hosted in regions and availability domains. A region can be a localized geographic area, and an availability domain can be one or more data centers located within a region. A region is composed of one or more availability domains. Most cloud infrastructure resources can be either region-specific, such as a virtual cloud network, or availability domain-specific, such as a compute instance. Traffic between availability domains and between regions is encrypted.

30 [000169] In accordance with an embodiment, availability domains are isolated from each other, fault tolerant, and very unlikely to fail simultaneously. Because availability domains do not share infrastructure such as power or cooling, or the internal availability domain network, a failure at one availability domain within a region is unlikely to impact the availability of the others within the same region.

35 [000170] In accordance with an embodiment, availability domains within the same region can be connected to each other by a low latency, high bandwidth network, which can provide high-availability connectivity to the internet and on-premises, and to build replicated systems

in multiple availability domains for both high-availability and disaster recovery.

**[000171]** In accordance with an embodiment, regions are independent of other regions and can be separated geographically (e.g., across countries or continents). This then leads to the deployment of an application within a region where the application would most likely be  
5 utilized the most frequently.

**[000172]** In accordance with an embodiment, however, applications can also be deployed in different regions for various reasons. This can include, for example, risk mitigation when events, such as weather systems, take a region offline. In addition, applications can be deployed in other regions for strategic reasons, such as tax domains or other business or social  
10 criteria.

**[000173]** In accordance with an embodiment, compartments allow clients to organize and control access to cloud resources. A compartment is a collection of related resources (such as instances, virtual cloud networks, block volumes) that can be accessed only by certain groups that have been given permission by an administrator. A compartment can be thought of as a  
15 logical group and not a physical container. When working within a console, a compartment can act as a filter for what is allowed to be viewed.

**[000174]** In accordance with an embodiment, compartments can have several layers, or levels, arranged in a hierarchical manner. For example, a tenancy 605 can be considered a root compartment that holds all of a client's cloud resources. Compartments can be organized  
20 in a hierarchical manner, such as compartment 660 being a level below the tenancy compartment, with sub compartment 670 being an additional layer below the compartment 660. In accordance with an embodiment, each compartment can be associated with one or more compartment resource quotas, such as tenancy resource quota 680 (i.e., the service limit discussed above), compartment resource quota 682, and sub compartment resource quota  
25 684.

**[000175]** In accordance with an embodiment, during, upon, or after creation of a compartment, or sub compartment, such as compartment 660 and sub compartment 670, each corresponding compartment resource quota can be generated via, e.g., the console 603.

**[000176]** In accordance with an embodiment, because of the hierarchical nature of  
30 compartments, a compartment resource quota set can be set on a root/base level compartment. An administrator (having a level of access sufficient to manage and create resource quotas within the root compartment) can set quotas on the root compartment, as well as any nested or "child" compartments. Compartment resource quotas set on a parent/root compartment override any compartment resource quotas set on child compartments. In this  
35 way, an administrator of a parent compartment can create a quota on a child compartment that cannot be overridden by the child.

**[000177]** In accordance with an embodiment, quotas can have different scopes, and work

at the availability domain, the region, or globally. There are a few important things to understand about scope when working with compartment quotas. When setting a quota at the availability domain (AD) level, the quota is allocated to each AD. So, for example, setting a quota of 2 X7 VMs on a compartment actually sets a limit of 2 VMs per AD. Quotas can be written in such a way to target a specific AD. Likewise, regional quotas can apply to each region. For example, if a quota of 10 functions is set on a compartment, 10 functions will be allocated per region. To target a specific region, a quota can use a region parameter.

**[000178]** In accordance with an embodiment, due to the hierarchical nature of compartment quotas, usage (e.g., of resources), resource usage of a sub-compartment counts as resource usage towards a main compartment. That is, for example, in a three-level compartment nest (parent compartment, child compartment, grandchild compartment), where each compartment and sub-compartment has a compartment quota policy attached to it, any resource usage by the grandchild compartment counts against a compartment quota of both the child compartment as well as the parent compartment.

**[000179]** In accordance for example, for example, suppose compartment resource quota 682 specifies that, compartment wide, compartment 660 may only utilize 10 virtual machines of type X. On creation of sub-compartment 670, sub compartment resource quota 684 specifies that, sub-compartment wide, sub compartment 670 may only utilize 8 virtual machines of type X. If, then, sub-compartment 670 utilizes its maximum capacity of virtual machines of type X (namely 8), then compartment 660, and any other sub-compartments thereof (e.g., sub-compartments not shown in the figure), can only utilize an addition 2 virtual machines of type X. The above example additionally applies to the tenancy resource quota 680 (i.e., the service limit) in a similar manner.

**[000180]** In accordance with an embodiment, compartment quotas can be evaluated following these guidelines. Within a policy, quota statements can be evaluated in order, and later statements can supersede previous statements that target the same resource. In cases where more than one policy is set for the same resource, the most restrictive policy can be applied. Service limits always take precedence over quotas. Although it is possible to specify a quota for a resource that exceeds the service limit for that resource, the service limit will still be enforced.

**[000181]** **Figure 7** shows a system for supporting compartment quotas in a cloud infrastructure environment, in accordance with an embodiment.

**[000182]** More specifically, **Figure 7** shows the replication of compartment quotas to different regions after creation/modification/deletion of a compartment quota within a region.

**[000183]** In accordance with an embodiment, as described above, instances of the cloud infrastructure environment 700 described above in **Figure 1** can be hosted in different regions, such as cloud infrastructure regions 720, 730, 740, 750. These can be accessed, as described

above, via a console 703, SDK, or APIs, by customer networks 701 via a network 702, such as the internet.

**[000184]** In accordance with an embodiment, a customer network 701 can comprise, for example, a single computer, a network of customer computers, or other such networks.

5 **[000185]** In accordance with an embodiment, although not shown in the Figure, each cloud infrastructure region can comprise a number of services, each comprising a number of resources, such as management services, compute services, storage services, edge serves, network services, and physical infrastructure.

10 **[000186]** In accordance with an embodiment, a cloud infrastructure can be hosted in regions and availability domains. A region can be a localized geographic area, and an availability domain can be one or more data centers located within a region. A region is composed of one or more availability domains. Most cloud infrastructure resources can be either region-specific, such as a virtual cloud network, or availability domain-specific, such as a compute instance. Traffic between availability domains and between regions is encrypted.

15 **[000187]** In accordance with an embodiment, availability domains are isolated from each other, fault tolerant, and very unlikely to fail simultaneously. Because availability domains do not share infrastructure such as power or cooling, or the internal availability domain network, a failure at one availability domain within a region is unlikely to impact the availability of the others within the same region.

20 **[000188]** In accordance with an embodiment, availability domains within the same region can be connected to each other by a low latency, high bandwidth network, which can provide high-availability connectivity to the internet and on-premises, and to build replicated systems in multiple availability domains for both high-availability and disaster recovery.

25 **[000189]** In accordance with an embodiment, regions are independent of other regions and can be separated geographically (e.g., across countries or continents). This then leads to the deployment of an application within a region where the application would most likely be utilized the most frequently.

30 **[000190]** In accordance with an embodiment, however, applications can also be deployed in different regions for various reasons. This can include, for example, risk mitigation when events, such as weather systems, take a region offline. In addition, applications can be deployed in other regions for strategic reasons, such as tax domains or other business or social criteria.

35 **[000191]** In accordance with an embodiment, compartments allow clients to organize and control access to cloud resources. A compartment is a collection of related resources (such as instances, virtual cloud networks, block volumes) that can be accessed only by certain groups that have been given permission by an administrator. A compartment can be thought of as a logical group and not a physical container. When working within a console, a compartment can

act as a filter for what is allowed to be viewed.

**[000192]** In accordance with an embodiment, compartments can have several layers, or levels, arranged in a hierarchical manner. For example, a tenancy 705 can be considered a root compartment that holds all of a client's cloud resources. Compartments can be organized in a hierarchical manner, such as compartment 760 being a level below the tenancy compartment, with sub compartment 770 being an additional layer below the compartment 760. In accordance with an embodiment, each compartment can be associated with one or more compartment resource quotas, such as tenancy resource quota 780 (i.e., the service limit discussed above), compartment resource quota 782, and sub compartment resource quota 784.

**[000193]** In accordance with an embodiment, during, upon, or after creation of a compartment, or sub compartment, such as compartment 760 and sub compartment 770, each corresponding compartment resource quota can be generated via, e.g., the console 703.

**[000194]** In accordance with an embodiment, because of the hierarchical nature of compartments, a compartment resource quota set can be set on a root/base level compartment. An administrator (having a level of access sufficient to manage and create resource quotas within the root compartment) can set quotas on the root compartment, as well as any nested or "child" compartments. Compartment resource quotas set on a parent/root compartment override any compartment resource quotas set on child compartments. In this way, an administrator of a parent compartment can create a quota on a child compartment that cannot be overridden by the child.

**[000195]** In accordance with an embodiment, quotas can have different scopes, and work at the availability domain, the region, or globally. There are a few important things to understand about scope when working with compartment quotas. When setting a quota at the availability domain (AD) level, the quota is allocated to each AD. So, for example, setting a quota of 2 X7 VMs on a compartment actually sets a limit of 2 VMs per AD. Quotas can be written in such a way to target a specific AD. Likewise, regional quotas can apply to each region. For example, if a quota of 10 functions is set on a compartment, 10 functions will be allocated per region. To target a specific region, a quota can use a region parameter.

**[000196]** In accordance with an embodiment, due to the hierarchical nature of compartment quotas, usage (e.g., of resources), resource usage of a sub-compartment counts as resource usage towards a main compartment. That is, for example, in a three-level compartment nest (parent compartment, child compartment, grandchild compartment), where each compartment and sub-compartment has a compartment quota policy attached to it, any resource usage by the grandchild compartment counts against a compartment quota of both the child compartment as well as the parent compartment.

**[000197]** In accordance with an embodiment, quotas can be managed centrally in a

single region and replicated globally. This reduces the effort required involved for managing quotas. For example, in looking at Figure 7, a sub-compartment resource quota policy 784 and a compartment resource quota 782 can be initially defined within region 720. Upon creation, each of the resource quotas 782 and 784 can be replicated in to regions 730, 740, and 750. Such replication can be an asynchronous process that globally replicates the quota policies to the quotas data plane in each data center.

**[000198]** In accordance with an embodiment, compartment quotas provide the ability to completely turn off individual services. As well, an algorithm for calculating resource usage over a hierarchy of resource compartments is provided. Quotas can consist of several primary components: quotas user interface; quotas control plane; quotas data plane; and quotas client SDKs.

**[000199]** In accordance with an embodiment, administrators can manage compartment quota policies through a user interface which behind the scenes calls the quotas control plane where the policies are stored. An asynchronous process globally replicates the data to the quotas data plane in each data center. Other cloud services integrate with the client SDK which can fetch the usage data from the service's persistent storage layer and the quota values from the quotas data plane. It then makes the decision whether the resource creation is permitted.

**[000200]** In accordance with an embodiment, there are a number of conditions that can be set in order for compartment quotas to be initiated. These include, but are not limited to: create/update/delete/list compartment level resource quotas; viewing compartment resource usage; compartment quotas can be set globally to apply to all regions/ADs; region/AD specific compartment quotas can be set and override globally defined quotas; compartment quotas can correspond to existing service quotas where applicable; setting and viewing compartment quotas can be possible through the console; quotas can be strictly enforced by each respective control plane; setting a quota should not fail based on current resource counts or parent compartment values; integrate with core services compute, DBaaS (database as a service) and block storage; achieving strong read after write consistency for get/list operations; ability to enforce zero quotas for free for teams unable to onboard; declarative quota statements via policies; and define compartment quotas to act as either a whitelist or a blacklist (i.e.. deny everything not set by a quota or allow everything).

**[000201]** In accordance with an embodiment, control planes can be extended to support compartment resource limits.

**[000202]** **Figure 8** is a flowchart of a method for creating a compartment quota, in accordance with an embodiment.

**[000203]** In accordance with an embodiment, a console 800, on receiving instructions indicative of a user of a tenant interaction to create a compartment quota, the console can call a limits control plane 805 (e.g., within a home region of the tenant) to create a compartment

quota via an API.

**[000204]** In accordance with an embodiment, the call can comprise a compartment ID to which the compartment quota is directed.

5 **[000205]** In accordance with an embodiment, the limits control plane 805 can communicate with an identity data plane 810 (e.g., within the home region of the tenant). The communication can comprise the compartment ID of the target compartment, and the limits control plane can be requesting a tenant associated with the targeted compartment. The identity control plane can return the tenant.

10 **[000206]** In accordance with an embodiment, on receiving the tenant, the limits control plane can then see if the request to create the compartment quota is sufficiently authorized (e.g., user of tenant has sufficient permissions to instantiate such compartment quota on the targeted compartment). The limits control plane can do this by again communicating with the identity control plane.

15 **[000207]** In accordance with an embodiment, on determining that the request is authorized, the limits control plane can create the quota entity within a database 815 (e.g., Kiev) (e.g., within the home region of the tenant). A successful creation of the compartment quota can then be returned to the console.

20 **[000208]** In accordance with an embodiment, the database 815 can be configured to store, and can store, all quotas associated with a tenant (or multiple tenants). That is, for example, both service level limits (that apply to an entire tenancy), as well as each of a plurality of compartment level quotas can be stored in one database (e.g., Kiev), with complete limits owned by one tenant.

**[000209]** **Figure 9** is a flowchart of a method for modifying a compartment quota, in accordance with an embodiment.

25 **[000210]** In accordance with an embodiment, a console 900, on receiving instructions indicative of a user of a tenant interaction to get, update, or delete a compartment quota, the console can call the database 915 to get the requested targeted compartment quota by the quota ID contained within the request.

30 **[000211]** In accordance with an embodiment, the call can comprise a quota ID to which the get/update/delete operation of the quota is directed.

35 **[000212]** In accordance with an embodiment, the limits control plane 905 can communicate with an identity data plane 910 (e.g., within the home region of the tenant). The communication can comprise the compartment ID of the target compartment, and the limits control plane can be requesting a tenant associated with the targeted compartment. The identity control plane can return the tenant.

**[000213]** In accordance with an embodiment, on receiving the tenant, the limits control plane can then see if the request to create the compartment quota is sufficiently authorized

(e.g., user of tenant has sufficient permissions to instantiate such compartment quota on the targeted compartment). The limits control plane can do this by again communicating with the identity control plane.

**[000214]** In accordance with an embodiment, on determining that the request is authorized, the limits control plane can get/update/delete the compartment quota within a database 915 (e.g., Kiev) (e.g., within the home region of the tenant). A successful creation of the compartment quota can then be returned to the console.

**[000215]** In accordance with an embodiment, the database 815 can be configured to store, and can store, all quotas associated with a tenant (or multiple tenants). That is, for example, both service level limits (that apply to an entire tenancy), as well as each of a plurality of compartment level quotas can be stored in one database (e.g., Kiev), with complete limits owned by one tenant.

**[000216]** **Figure 10** shows an architecture for a system to enforcing compartment quotas in a cloud infrastructure environment, in accordance with an embodiment.

**[000217]** In accordance with an embodiment, Figure 10 shows an architecture flow for a user launching an instance that requires a check against one or more compartment resource quotas.

**[000218]** In accordance with an embodiment, at step 1, instructions can be received from, e.g., a user 1001. The instructions can be indicative of a request to launch an instance in a region 1000, where the request specifies a compartment within the region.

**[000219]** In accordance with an embodiment, at step 2, a load balancer 1020 (e.g., Flamingo) can forward the request to an API, such as a public API proxy.

**[000220]** In accordance with an embodiment, at step 3, the API can perform an authentication at the identity data plane 1035. Such an authentication step can check both that a user has correct credentials (e.g., user credentials such as username and password, one-time password, or other authentication methods), and as well that the user 1001 has sufficient authorization within a region to perform the requested launching of an instance within the targeted compartment. If, for example, the user 1001 does not have sufficient authorization to launch the instance at the targeted compartment (for example, the user 1001 does only has authorization to launch a similar instance at a child compartment of the targeted compartment), then the API 1025 can deny the request. If the user does have sufficient authorization, then the API can allow the request to proceed.

**[000221]** In accordance with an embodiment, at step 4, the request to launch the instance can be forwarded to compute control plate (CCP) 1030. The CCP can then additionally authenticate the request at the identity data plane at step 5.

**[000222]** In accordance with an embodiment, at step 6, the compute control plane can call the identity data plane 1035 to fetch the compartment tree associated with the

compartment to which the request is targeted, or the compartment tree for the entire tenancy to which the request is directed.

**[000223]** In accordance with an embodiment, at step 7, the compute control plane can fetch all compartment quotas for the tenancy or compartment tree from the limit service data plane 1040, which comprises a compartment quota map 1041. This can comprise a number of compartment quotas for the specific compartment to which the request is directed, as well as compartment quotas for the entire tree above the specific compartment.

**[000224]** In accordance with an embodiment, at step 8, the compute control plane can get the resource usage for the compartment from the database 1045, as well as check to see if the requested usage would violate any of the compartment quotas for any compartment along the tree comprising the specific compartment to which the request is directed. If the requested usage does not violate any compartment quota, then the request can be processed. If the requested usage would violate any of the compartment quotas along the compartment tree, then the request can be dropped and a message can be sent to notify the user.

**[000225]** In accordance with an embodiment, compartment quotas are applied to specific compartments within the tenancy. When the compartment tree is fetched, along with all associated compartment quotas, then the CCP can get the usage and update the usage based upon the request. From the database, as part of step 8, CCP can get the actual usage in each compartment. For the entire compartment tenancy, CCP can map each compartment ID to a number of counts (usage). Then CCP can map the entire compartment tree layout and the quotas that exist. The CCP can then ensure that the new request doesn't violate any of the quotas. If the request would exceed a quota, CCP can fail the request.

**[000226]** **Figure 11** is a flowchart of a method supporting compartment quotas in a cloud infrastructure environment, in accordance with an embodiment.

**[000227]** In accordance with an embodiment, at step 1110, the method can provide a computer comprising one or more microprocessors.

**[000228]** In accordance with an embodiment, at step 1120, the method can provide a cloud infrastructure environment at the computer, the cloud infrastructure environment comprising a plurality of regions, wherein a tenancy is defined within the plurality of regions.

**[000229]** In accordance with an embodiment, at step 1130, the method can provide a plurality of compartments of the tenancy.

**[000230]** In accordance with an embodiment, at step 1140, the method can provide a limit service data plane within the cloud infrastructure environment, the limit service data plane having access to a database storing a plurality of compartment resource quotas.

**[000231]** In accordance with an embodiment, at step 1150, the method can associate a compartment resource quota of the plurality of compartment resource quotas with a compartment of the plurality of compartments, the compartment resource quota limiting at least

one type of resource within the compartment.

**[000232]** In accordance with an embodiment, at step 1160, the method can receive a request for a resource to be provisioned in the compartment, the requested resource being of the type of resource limited by the resource quota.

5 **[000233]** In accordance with an embodiment, at step 1170, the method can determine whether the provisioning of the requested resource would violate the compartment resource quota.

**[000234]** In accordance with an embodiment, at step 1180, the method can, based upon a determination that the compartment resource quota would not be violated, provision the  
10 requested resource.

**[000235]** In accordance with an embodiment, at step 1190, the method can, based upon a determination that the compartment resource quota would be violated, provide an objection to the request.

## 15 **Quota Policy Language**

**[000236]** In accordance with an embodiment, as described above, a cloud infrastructure environment can provide for the ability for customers to set resource quotas on a compartment. Quotas act like existing service limits except they are set by customers to restrict their own usage rather than by a cloud infrastructure environment. Service teams can check against  
20 these quotas in their control planes and hard block requests that exceed a quota. Quotas can be stored and accessed from a limits service, and be set by clients through the console or the REST API. In addition, the cloud infrastructure environment can support a declarative, human readable, quota policy language to be used in the creation of such compartment quotas.

**[000237]** In accordance with an embodiment, a compartment resource quota can be set  
25 in a similar manner as writing, for example, an identity policy. Both live in a compartment and can be written to target a different compartment. A resource quota can also be associated with a set of conditions that can be applied, such as, for example, region, availability domain, and overlaps.

**[000238]** In accordance with an embodiment, a resource quota can be configured by  
30 setting a statement syntax. In order to set a quota, the quota can establish a number of factors. Each quota can be unique within a service family (e.g. compute), and as such, each quota can define a target service along with a quota name. Next, a quota can define the value to set the quota to, as well as a compartment the quota targets. Finally, there are a set of conditions that can be included determine when this quota is applied.

35 **[000239]** In accordance with an embodiment, compartment quota policies can be written using an expressive human readable policy language. The policy language allows for rich expressiveness including the ability to both whitelist or blacklist resource usage through

statement overrides. A whitelist allows all resources to be created except those that are explicitly denied. A blacklist allows no resources to be created except those that are explicitly allowed.

**[000240]** In accordance with an embodiment, compartment quotas can be set using policy statements written in a simple declarative language. There are, in general, three types of quota policy statements: 1) set - sets the maximum number of a cloud resource that can be used for a compartment; 2) unset - resets quotas back to the default service limits; and 3) zero - removes access to a cloud resource for a compartment.

**[000241]** Figure 12a, 12b, and 12c are example compartment quota policy statements, in accordance with an embodiment.

**[000242]** Figure 12a shows a set quota policy, in accordance with an embodiment.

**[000243]** Figure 12b shows an unset quota policy, in accordance with an embodiment.

**[000244]** Figure 12c shows a zero quota policy, in accordance with an embodiment.

**[000245]** In accordance with an embodiment, the language components for each quota policy statement comprise the following components: 1) an action keyword – the action keyword can correspond to a type of quota being defined. This can be, for example, set, unset, or zero; 2) the name of the service family. This can be, for example: compute; 3) the quota or quotas keyword; 4) the name of the quota, which varies by service family. For example, a valid quota in the compute family can be "vm-standard2-16-count". In addition, wildcards can be used to specify a range of names. For example, "/vm-\*/" matches all compute shapes that start with the letters "vm"; 5) a "for" set statements that sets a value of the quota; 6) and the compartment that the quota covers; 7) an optional condition can also be provided. For example where request.region = 'us-phoenix-1'. Such conditionals can include request.region and request.ad.

**[000246]** Figure 13 illustrates an example of a quota policy statement, in accordance with an embodiment.

**[000247]** In accordance with an embodiment, the statement of Figure 13 is a set statement, which is applied to a family 1302. The quota 1303 is applied to a name 1304 (e.g., a resource), which quota sets to 1305 a value 1306. The quota can be applied in 1307 a location 1308, and then conditions 1309 can be set as well.

**[000248]** In accordance with an embodiment, provided below are some exemplary quota policy statements, provided below are some examples along with an explanation of each quota policy.

**[000249]** In accordance with an embodiment, the quota policy:

```
Set compute quotas to 0 in compartment My Compartment
```

[000250] In accordance with an embodiment, the above quota policy would set all compute quotas to 0 in "MyCompartment", thereby not allowing any further resources to be created.

[000251] In accordance with an embodiment, the quota policy:

5

```
Set block-storage quota 'volume-count' to 10 in compartment MyCompartment
```

[000252] In accordance with an embodiment, the above quota policy would set block storage quota for volume count to 10 in "MyCompartment".

10

[000253] In accordance with an embodiment, the quota policy:

```
Set compute quota 'VM.Dense101.16' to 10 in compartment MyCompartment where request.region = 'phx'
```

15

[000254] In accordance with an embodiment, the above quota policy would set the quota for VM.DenseIO1.16 Compute shapes to 10 on compartment MyCompartment in PHX.

[000255] In accordance with an embodiment, the quota policy:

20

```
Set compute quota /VM.*/ to 10 in compartment MyCompartment
```

[000256] In accordance with an embodiment, a wildcard expression can also be used to match many names at once. The above quota policy sets the quota for each VM shape to 10, not the combined quota to 10.

25

[000257] In accordance with an embodiment, the quota policy format can require a specification of a quota family along with an optional name or list of names.

[000258] In accordance with an embodiment, the quota location 1008 can comprise a target compartment to which the quota applies. This is often different than the compartment where the quota resides and can be a descendant of it. The location can also be "tenancy" which is equivalent to specifying the root compartment.

30

[000259] In accordance with an embodiment, another example of a quota policy statement is a zero statement. **Figure 14** illustrates an example of a quota policy statement, in accordance with an embodiment.

[000260] In accordance with an embodiment, the statement of Figure 14 is a zero 1401 statement, which is applied to a family 1402. The quota 1403 is applied to a name 1404 (e.g., a resource). The zero quota is applied in 1405 a location 1406, and then conditions 1407 can be set as well.

35

[000261] In accordance with an embodiment, provided below are some exemplary quota policy statements, provided below are some examples along with an explanation of each quota policy.

[000262] In accordance with an embodiment, the quota policy:

5

```
Zero compute quota 'VMStandard2.32' in compartment MyCompartment
```

[000263] In accordance with an embodiment, the above quota policy statement applies to VM Standard 2.32 instance type in MyCompartment. Once this quota is set, a user will not be able to launch any of these instances (e.g., VM Standard 2.32) in that compartment.

10

[000264] In accordance with an embodiment, the quota policy languages can additionally support a whitelist. In other words, setting every quota in a family to 0 and allocating resources explicitly on top of this. This can be implemented by establishing a simple statement precedence rule, namely - when two or more statements evaluate to true within the same policy the last statement in the set will control.

15

[000265] In accordance with an embodiment, consider the below two quota policies to see how precedence works to support a whitelist.

```
Zero compute quotas in tenancy
```

20

```
Set compute quota 'VM.DenseIO1.16' to 10 in tenancy
```

[000266] In accordance with an embodiment, a base statement can be written that reduces all Compute quotas to 0. On top of this, a second statement can be written for VM.DenseIO1.16 which wants to allow up to 10. Importantly, this override only applies to statements written in the same policy. There is no precedence between statements from different policies, they must all be satisfied. This makes it impossible for a child compartment admin to override a quota in a parent compartment. Admins can safely give child compartment admins "manage all-resources" on the child and maintain proper quotas. Child compartment admins can only further reduce quotas that are already present within a parent compartment.

25

[000267] In accordance with an embodiment, another example of a quota policy statement is an unset statement. **Figure 15** illustrates an example of a quota policy statement, in accordance with an embodiment.

30

[000268] In accordance with an embodiment, the statement of Figure 15 is an unset 1501 statement, which is applied to a family 1502. The quota 1503 is unset to a name 1504 (e.g., a resource), in 1505 a location 1506, with optional conditions 1507.

35

[000269] In accordance with an embodiment, re-working the above example with regard to precedence, the second quota policy can be unset so that there is no quota on VM

DenseIO1.16.

Zero compute quotas in tenancy

Unset compute quota 'VM.DenseIO1.16' to 10 in tenancy

5

**[000270]** In accordance with an embodiment, as with the previous example, the unset does not work across different quota policies.

**[000271]** In accordance with an embodiment, within each service's control plane, there is a critical section (e.g., time) where the database transaction stays open. The longer this period of time extends, conflicts and failures start to creep in. In order to minimize this possible problem, quota policies can be considered and applied in a range of approximately 5 milliseconds.

**[000272]** In accordance with an embodiment, as described above, quotas can be enforced at the compartment level. Because of this, quotas can additionally be specified to set quotas on additional entities within compartments, such as users, groups and tags. For example, consider the below quota statement:

Set compute quota 'VM.DenseIO1.16' to 20 in tenancy where request.user =  
 20 ocid1.user.dev..aaaaaaaaan4xh2xpvzcmabc5m7pd7ioaovc2ktt56fvbhfy77y4v

**[000273]** The above allows for a quota to partially get placed on a user, but it doesn't ensure the user will get 20 instances since the quota is on the compartment and not the user. Instead, a quota can be specified that sets a quota for a user regardless of the current compartment usage.

**[000274]** Figure 16 illustrates an example of a quota policy statement, in accordance with an embodiment.

**[000275]** More specifically, Figure 16 shows a quota policy statement that additionally includes a scope.

**[000276]** In accordance with an embodiment, the statement of Figure 16 is a set 1601 statement, which is applied to a family 1602. The quota 1603 is applied to a name 1604 (e.g., a resource), which quota sets to 1605 a value 1606. The quota can be applied for 1607 a scope 1608, in 1609 a location 1610, where 1611 conditions 1612 can be set as well. Take the above compartment quota and re-writing it using the statement language of Figure 16 then, the quota policy becomes:

Set compute quota 'VM.DenseIO1.16' to 10 for user MyUser in tenancy

[000277] In accordance with an embodiment, the quota policy language clearly expresses that this quota is applied to MyUser and not the compartment. This would be an entirely separate quota from the compartment quota, and all quota scopes (both the above  
5 quota policy as well as any compartment quota policy) would need to be satisfied when creating a resource.

[000278] **Figure 17** illustrates an example of a quota policy statement, in accordance with an embodiment.

[000279] More specifically, Figure 17 shows an override quota policy.

10 [000280] In accordance with an embodiment, precedence rules associated with compartment quotas do not allow for a statement in one policy to override a statement from a different policy. That means to implement a whitelist, the systems and methods are limited by the maximum number of statements in a single policy. To get around this, an explicit override statement can be utilized. This statement establishes the precedence of one policy over  
15 another. It's equivalent to having written the statements together as part of one long policy. In order for an override statement to apply, the policy must live in the same compartment

[000281] In accordance with an embodiment, the override statement can comprise an override statement 1700, an override 1701, one or more other quotas 1702, in 1703 policy 1704, and a name 1705.

20 [000282] In accordance with an embodiment, consider the below override statement:

```
Override quotas in policy BasePolicy
```

[000283] In accordance with an embodiment, the above override statement provides that  
25 any statements in this policy will override statements in BasePolicy. Note that overrides are transitive. If a overrides b and b overrides c, then a overrides c.

[000284] **Figure 18** is a flowchart of a method for supporting a quota policy language in a cloud infrastructure environment, in accordance with an embodiment.

[000285] In accordance with an embodiment, at step 1810, the method can provide a  
30 computer comprising one or more microprocessors.

[000286] In accordance with an embodiment, at step 1820, the method can provide a cloud infrastructure environment comprising a plurality of regions, wherein a tenancy is defined within the plurality of regions.

[000287] In accordance with an embodiment, at step 1830, the method can provide a  
35 plurality of compartments of the tenancy.

[000288] In accordance with an embodiment, at step 1840, the method can provide a console interface providing adapted to receive inputs indicative of instructions.

[000289] In accordance with an embodiment, at step 1850, the method can provide a limit service data plane within the cloud infrastructure environment, the limit service data plane having access to a database storing a plurality of compartment resource quotas,

[000290] In accordance with an embodiment, at step 1860, the method can associate a compartment of the tenancy with a compartment quota of the plurality of compartment quotas.

[000291] In accordance with an embodiment, at step 1870, the method can configure the quota policy to be one of a set policy, unset policy, or zero quota policy.

### Usage Calculation Process (Algorithm)

[000292] In accordance with an embodiment, systems and methods can support a usage calculation process or algorithm to determine whether a requested transaction violates any limits or quotas with regard to compartments.

[000293] In accordance with an embodiment, a usage calculation method can provide an SDK for both evaluating quotas or limits, as well as counting resources per compartment.

[000294] In accordance with an embodiment, the process can be distributed amongst a number of systems and sub-systems within a cloud infrastructure environment.

[000295] In accordance with an embodiment, Figure 10 shows an architecture flow for a user launching an instance that requires a check against one or more compartment resource quotas.

[000296] Figure 19 shows an architecture for a system to enforcing compartment quotas in a cloud infrastructure environment, in accordance with an embodiment.

[000297] In accordance with an embodiment, at step 1, instructions can be received from, e.g., a user 1901. The instructions can be indicative of a request to launch an instance in a region 1900, where the request specifies a compartment within the region.

[000298] In accordance with an embodiment, at step 2, a load balancer 1920 (e.g., Flamingo) can forward the request to an API, such as a public API proxy.

[000299] In accordance with an embodiment, at step 3, the API can perform an authentication at the identity data plane 1935. Such an authentication step can check both that a user has correct credentials (e.g., user credentials such as username and password, one-time password, or other authentication methods), and as well that the user 1901 has sufficient authorization within a region to perform the requested launching of an instance within the targeted compartment. If, for example, the user 1901 does not have sufficient authorization to launch the instance at the targeted compartment (for example, the user 1901 does only has authorization to launch a similar instance at a child compartment of the targeted compartment), then the API 1925 can deny the request. If the user does have sufficient authorization, then the API can allow the request to proceed.

[000300] In accordance with an embodiment, at step 4, the request to launch the instance

can be forwarded to compute control plate (CCP) 1930. The CCP can then additionally authenticate the request at the identity data plane at step 5.

5 [000301] In accordance with an embodiment, at step 6, the compute control plane can call the identity data plane 1935 to fetch the compartment tree associated with the compartment to which the request is targeted, or the compartment tree for the entire tenancy to which the request is directed.

10 [000302] In accordance with an embodiment, at step 7, the compute control plane can fetch all compartment quotas for the tenancy or compartment tree from the limit service data plane 1940, which comprises a compartment quota map 1941. This can comprise a number of compartment quotas for the specific compartment to which the request is directed, as well as compartment quotas for the entire tree above the specific compartment.

15 [000303] In accordance with an embodiment, at step 8, the compute control plane can calculate the total usage for each compartment in the lineage. This total usage is equal to the sum of compartment's usage and the total usage of each child compartment recursively down the hierarchy.

20 [000304] In accordance with an embodiment, the CCP 1930 can compare each limit/quota against the total usage for that compartment. If any limit/quota is violated, the CCP can return an objection indicating which quotas and/or limits the prospective resource would exceed. In addition to the objection, the CCP can also return an amount the prospective resource would exceed a quota by.

[000305] In accordance with an embodiment, if the CCP determines that the prospective resource would be allowed, then at step 9, the transaction can be committed at the database 1945.

25 [000306] **Figure 20** is a flowchart of a method for supporting a usage calculation process in a cloud infrastructure environment, in accordance with an embodiment.

[000307] In accordance with an embodiment, at step 2001, the process can fetch the compartment-tree hierarchy from, for example, an identity data plane. This can comprise single request (e.g., HTTP) with the amount of data getting returned growing linearly with the number of compartments.

30 [000308] In accordance with an embodiment, at step 2002, the process can determine a lineage of a target compartment (i.e., the compartment to which a transaction is targeted). The lineage can grow linearly with the height of the compartment tree.

35 [000309] In accordance with an embodiment, at step 2003, the method can fetch the quota policies / service limits that are enacted on the target compartment by calling the limit service data plane. In doing so, the process can fetch both quotas and service limit requests within a single limit service data plane request.

[000310] In accordance with an embodiment, at step 2004, the database transaction can

be started.

**[000311]** In accordance with an embodiment, at step 2005, resource counts for each affected quota from the database bucket can be fetched. For each quota, the process can check against an object from the database with the entire tenant usage data for that quota.

5 The data returned can grow linearly with the number compartments in that tenant that have usage of that resource/quota.

**[000312]** In accordance with an embodiment, at step 2006, the process can calculate the total usage for each compartment in the lineage. The total usage can be an aggregate of the compartment's usage plus the total usage of all children, recursively down the tree. This operation is linear in respect to the number of compartments in a tree. The process can perform a recursive depth first tree traversal with counts coming up the tree.

10

**[000313]** In accordance with an embodiment, at step 2007, the process can compare each limit / quota against the total usage for that compartment. If a limit / quota is violated, the process can return a result object indicating which quota or limit was violated, as well as by how much.

15

**[000314]** In accordance with an embodiment, at step 2008, the process can commit the transaction to the database.

**[000315]** **Figure 21** is a flowchart of a method for supporting a usage calculation process in a cloud infrastructure environment, in accordance with an embodiment.

20 **[000316]** In accordance with an embodiment, at step 2110, the method can provide a computer comprising one or more microprocessors.

**[000317]** In accordance with an embodiment, at step 2120, the method can provide a cloud infrastructure environment comprising a plurality of regions, wherein a tenancy is defined within the plurality of regions.

25 **[000318]** In accordance with an embodiment, at step 2130, the method can provide a compartment of the tenancy, the compartment belonging to a tree structure of a plurality of compartments, wherein the compartment is associated with a compartment quota policy.

**[000319]** In accordance with an embodiment, at step 2140, the method can receive a requested transaction, the requested transaction targeting the compartment.

30 **[000320]** In accordance with an embodiment, at step 2150, the method can determine whether the requested transaction violates the compartment quota policy.

**[000321]** In accordance with an embodiment, the requested transaction can request a type of resource.

35 **[000322]** In accordance with an embodiment, the determining of described at step 2150 can comprise additional steps, comprising: fetching a compartment-tree hierarchy, the compartment-tree hierarchy comprising the targeted compartment; determining a lineage of the targeted compartment; fetching, by calling a limit service data plane, a plurality of

compartment quotas that are enacted on the targeted compartment, wherein the plurality of compartment quotas comprises the compartment quota policy, and wherein each of the plurality of compartment quotas belong to a compartment within the compartment-tree hierarchy; starting the requested transaction at a database; fetching resource counts for each of the plurality of compartment quotas quota from the database, each resource count comprising a count of the requested resource type; calculating a total usage of the requested resource type for each compartment in the determined lineage, wherein the total usage for each compartment comprises an aggregate of a usage of a compartment and a total usage of each child compartment; and comparing each of the plurality of compartment quotas to the total usage for each compartment.

**[000323]** In accordance with various embodiments, the teachings herein may be conveniently implemented using one or more conventional general purpose or specialized computer, computing device, machine, or microprocessor, including one or more processors, memory and/or computer readable storage media programmed according to the teachings of the present disclosure. Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent to those skilled in the software art.

**[000324]** In some embodiments, the teachings herein can include a computer program product which may use a computer readable medium (or media) to convey (for example store or carry) instructions for one or more programmable processors or computers to perform any of the processes of the present teachings. One example of a computer-readable medium is a non-transitory computer readable storage medium (media) having instructions stored thereon/in which can be used to program a computer to perform any of the processes of the present teachings. Examples of such storage mediums can include, but are not limited to, hard disk drives, hard disks, hard drives, fixed disks, or other electromechanical data storage devices, floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices, magnetic or optical cards, nanosystems, or other types of storage media or devices suitable for non-transitory storage of instructions and/or data. Another example of a computer readable medium is a transient medium which conveys such instructions between computer systems, or between components of a given computer system. Examples of transient media can include, but are not limited to, carried waves and transmissions signals.

**[000325]** Therefore, from one perspective, there have been described systems and methods that support compartment quotas in a cloud infrastructure environment. Cloud administrators do not generally have the ability to restrict resource usage in existing clouds. Granting a user permission to create resources allows them to create any number of resources

up to a predefined account limit. Compartment quotas allow admins to restrict a user's resource usage to the appropriate level allowing fine-tuned cost control.

**[000326]** The foregoing description has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the scope of protection to the precise forms disclosed. Many modifications and variations will be apparent to the practitioner skilled in the art. For example, although several of the examples provided herein illustrate use with enterprise software applications components such as Oracle Fusion Applications; cloud environments such as Oracle Cloud Infrastructure; and cloud services such as Oracle Fusion Analytics; in accordance with various embodiments, the systems and methods described herein can be used with other types of enterprise software applications, cloud environments, cloud services, cloud computing, or other computing environments.

**[000327]** The embodiments were chosen and described in order to best explain the principles of the present teachings and their practical application, thereby enabling others skilled in the art to understand the various embodiments and with various modifications that are suited to the particular use contemplated. It is intended that the scope be defined by the following claims and their equivalents.

**Claims:**

What is claimed is:

- 5 1. A system for supporting a quota policy language in a cloud infrastructure environment, comprising:
- 10 a computer comprising one or more microprocessors;
  - a cloud infrastructure environment comprising a plurality of regions, wherein a tenancy is defined within the plurality of regions;
  - 10 a plurality of compartments of the tenancy;
  - a console interface providing adapted to receive inputs indicative of instructions;
  - a limit service data plane within the cloud infrastructure environment, the limit service data plane having access to a database storing a plurality of compartment resource quotas,
  - 15 a compartment of the tenancy, wherein the compartment is associated with a compartment quota of the plurality of compartment quotas;
  - 15 wherein the compartment quota policy is one of a set policy, unset policy, or zero quota policy.
2. The system of claim 1,
- 20 wherein the compartment quota policy comprises a human-readable format.
3. The system of claim 1 or 2,
- 25 wherein the compartment quota policy defines a family of resources to which the compartment quota policy is targeted.
4. The system of claim 3,
- 30 wherein the compartment quota policy is a set policy; and
- wherein the compartment quota policy sets an upper threshold on the family of resources to which the compartment quota policy is targeted.
5. The system of claim 3,
- 35 wherein the compartment quota policy is an unset policy; and
- wherein the compartment quota policy unsets a previously set upper threshold on the family of resources to which the compartment quota policy is targeted.
6. The system of claim 3,
- wherein the compartment quota policy is a zero quota; and

wherein the compartment quota policy removes access to the family of resources to which the compartment quota policy is targeted.

7. The system of any of claims 3 to 6, wherein the compartment quota policy further  
5 comprises at least one condition.
8. A method for supporting a quota policy language in a cloud infrastructure environment, comprising:
- 10 providing a computer comprising one or more microprocessors;
  - providing a cloud infrastructure environment comprising a plurality of regions, wherein a tenancy is defined within the plurality of regions;
  - providing a plurality of compartments of the tenancy;
  - providing a console interface providing adapted to receive inputs indicative of instructions;
  - 15 providing a limit service data plane within the cloud infrastructure environment, the limit service data plane having access to a database storing a plurality of compartment resource quotas,
  - associating a compartment of the tenancy with a compartment quota of the plurality of compartment quotas;
  - 20 configuring the quota policy to be one of a set policy, unset policy, or zero quota policy.
9. The method of claim 8,  
wherein the compartment quota policy comprises a human-readable format.
- 25 10. The method of claim 8 or 9,  
wherein the compartment quota policy defines a family of resources to which the compartment quota policy is targeted.
11. The method of claim 10,  
30 wherein the compartment quota policy is a set policy; and  
wherein the compartment quota policy sets an upper threshold on the family of resources to which the compartment quota policy is targeted.
12. The method of claim 10,  
35 wherein the compartment quota policy is an unset policy; and  
wherein the compartment quota policy unsets a previously set upper threshold on the family of resources to which the compartment quota policy is targeted.

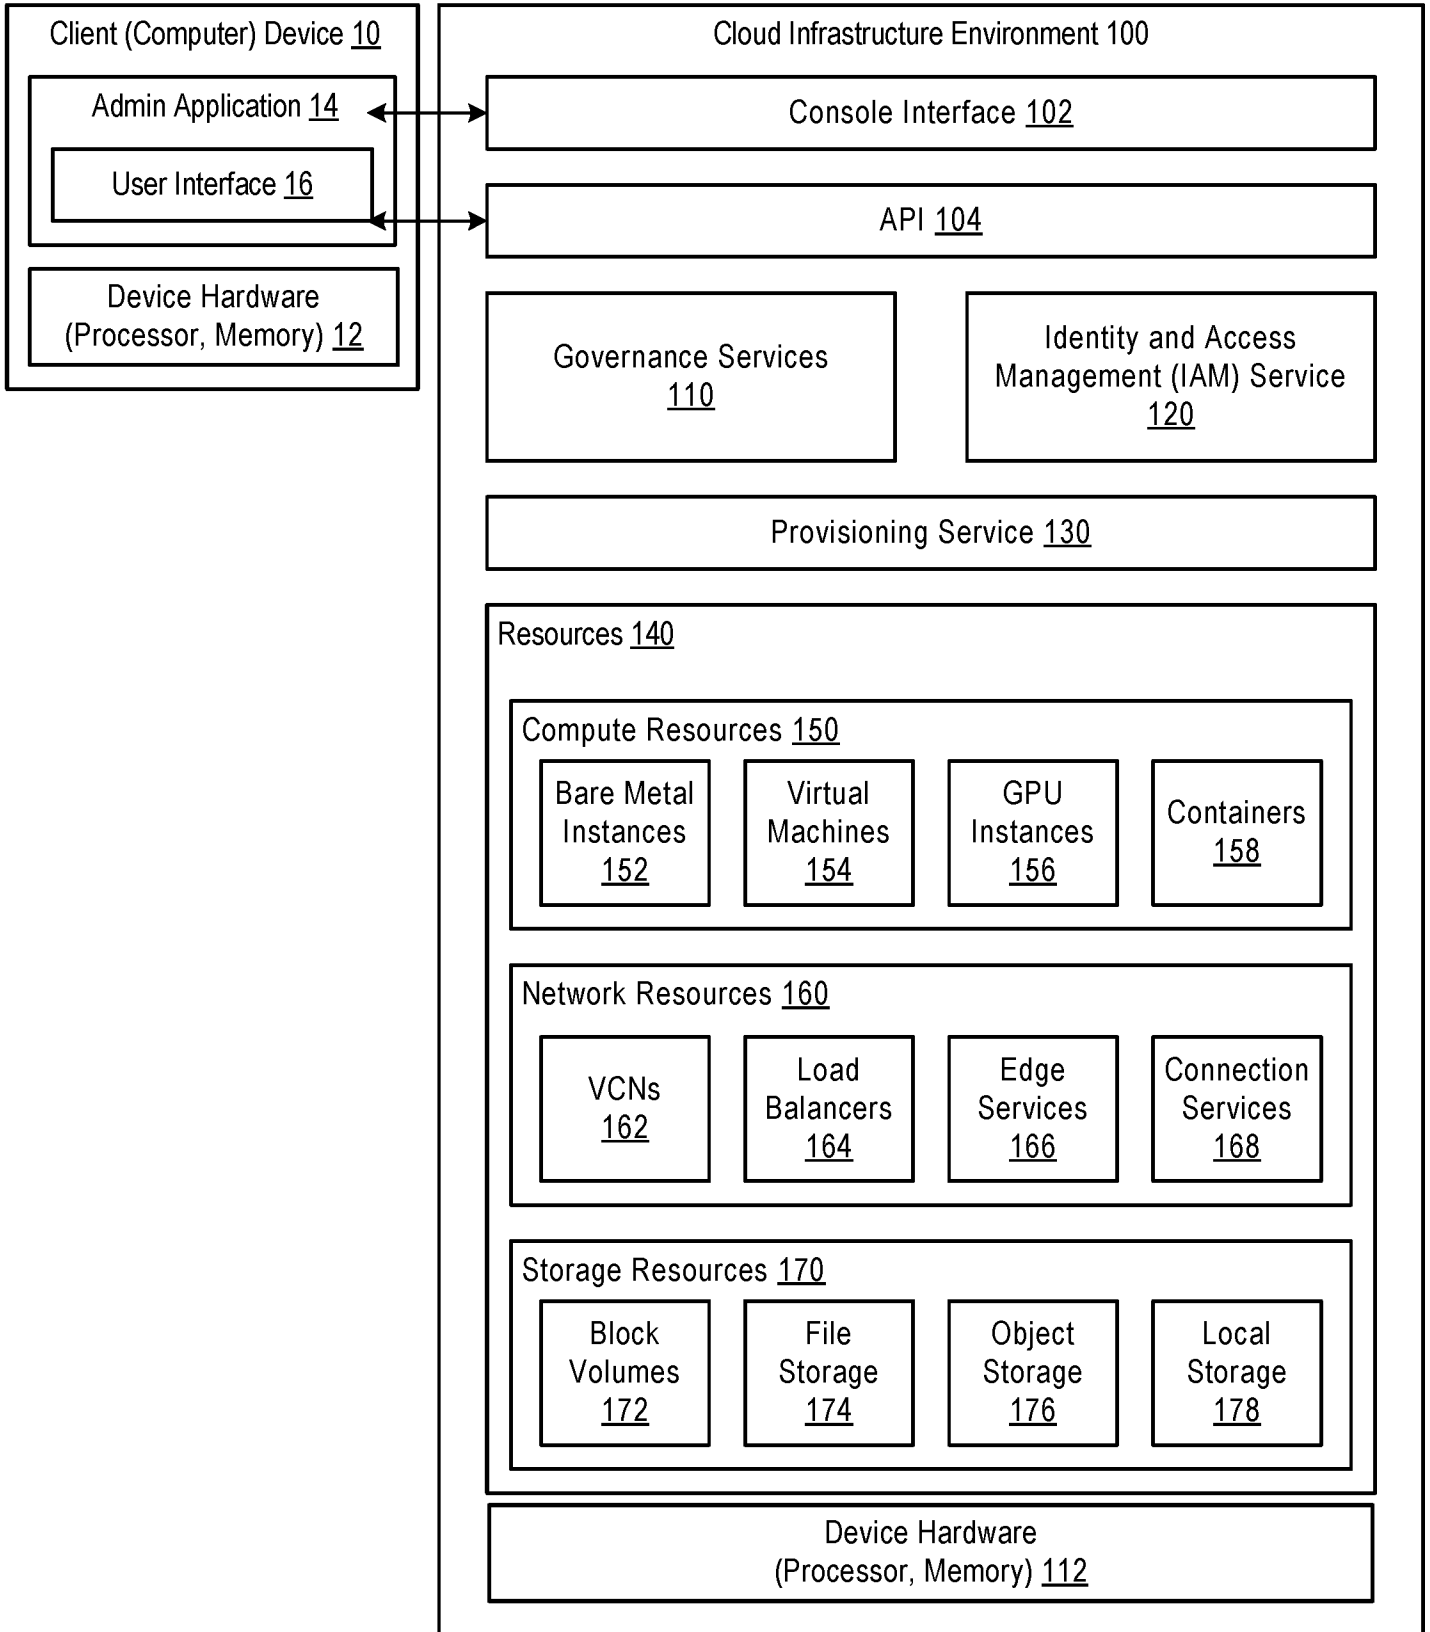
13. The method of claim 10,  
wherein the compartment quota policy is a zero quota; and  
wherein the compartment quota policy removes access to the family of resources to  
5 which the compartment quota policy is targeted.
14. The method of any of claims 10 to 13, wherein the compartment quota policy further  
comprises at least one condition.
- 10 15. A computer readable medium having instructions thereon for supporting a quota policy  
language in a cloud infrastructure environment, which when read and executed by a computer  
cause the computer to perform steps comprising:  
providing a computer comprising one or more microprocessors;  
providing a cloud infrastructure environment comprising a plurality of regions, wherein  
15 a tenancy is defined within the plurality of regions;  
providing a plurality of compartments of the tenancy;  
providing a console interface providing adapted to receive inputs indicative of  
instructions;  
providing a limit service data plane within the cloud infrastructure environment, the limit  
20 service data plane having access to a database storing a plurality of compartment resource  
quotas,  
associating a compartment of the tenancy with a compartment quota of the plurality of  
compartment quotas;  
configuring the quota policy to be one of a set policy, unset policy, or zero quota policy.  
25
16. The medium of claim 15,  
wherein the compartment quota policy comprises a human-readable format.
17. The medium of claim 16,  
30 wherein the compartment quota policy defines a family of resources to which the  
compartment quota policy is targeted.
18. The medium of claim 17,  
wherein the compartment quota policy is a set policy; and  
35 wherein the compartment quota policy sets an upper threshold on the family of  
resources to which the compartment quota policy is targeted.

19. The medium of claim 17,  
wherein the compartment quota policy is an unset policy; and  
wherein the compartment quota policy unsets a previously set upper threshold on the  
family of resources to which the compartment quota policy is targeted.

5

20. The medium of claim 17,  
wherein the compartment quota policy is a zero quota; and  
wherein the compartment quota policy removes access to the family of resources to  
which the compartment quota policy is targeted.

10



**FIGURE 1**




FIGURE 2

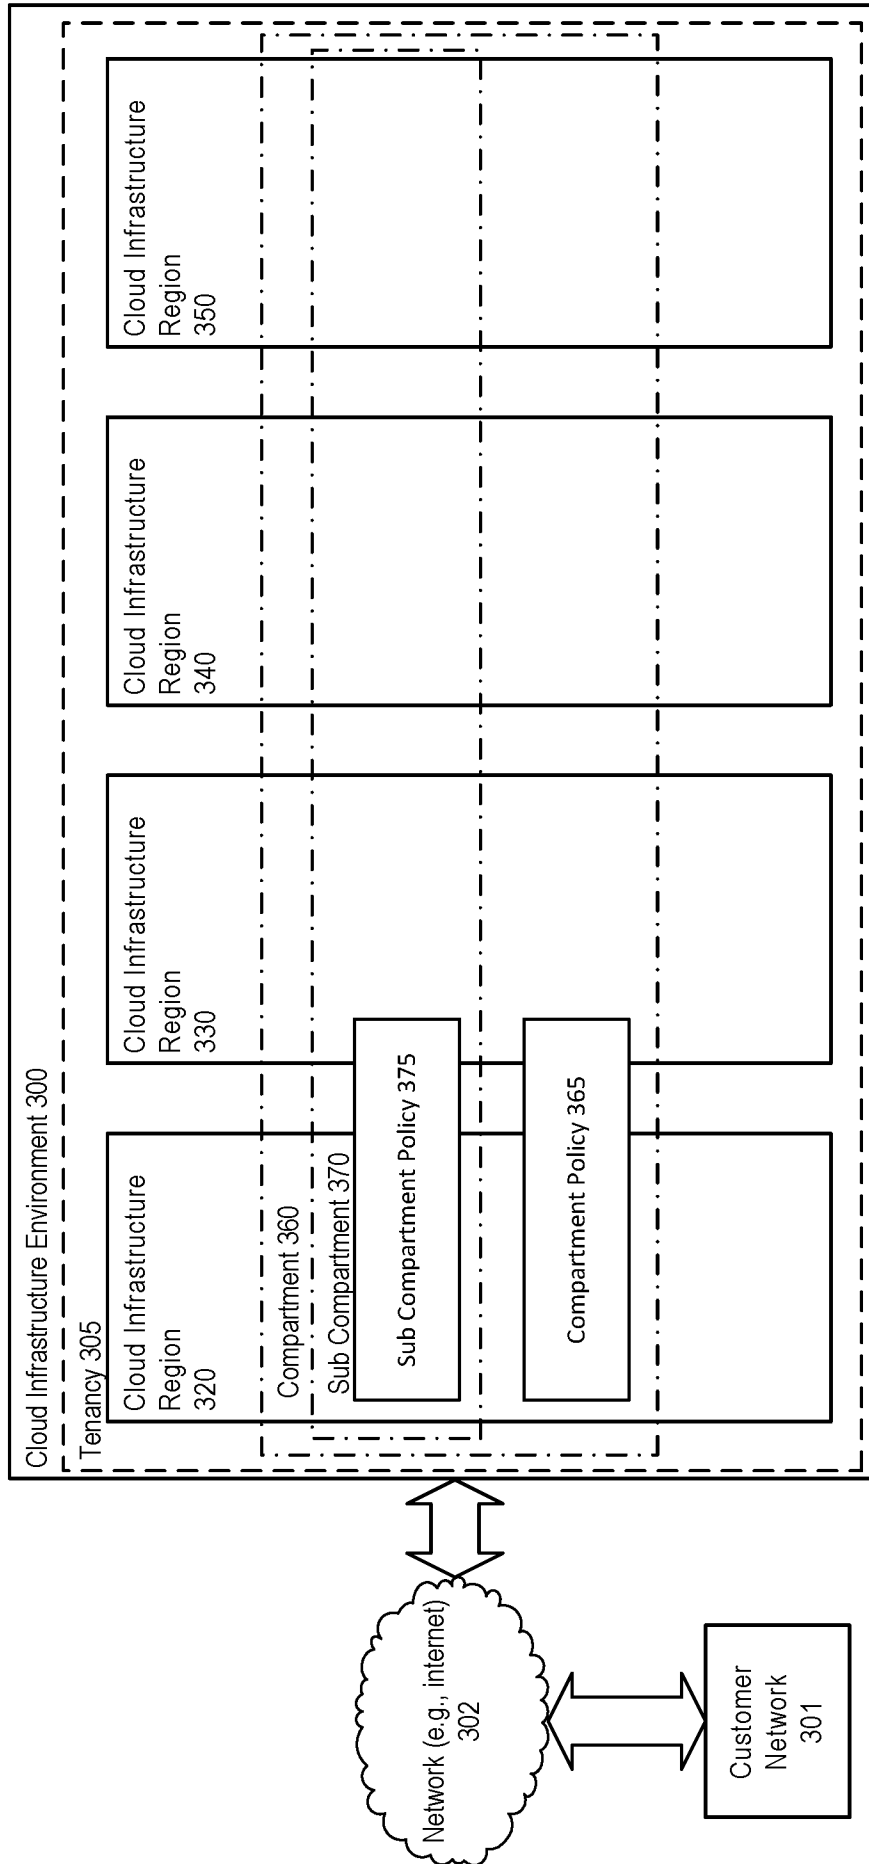
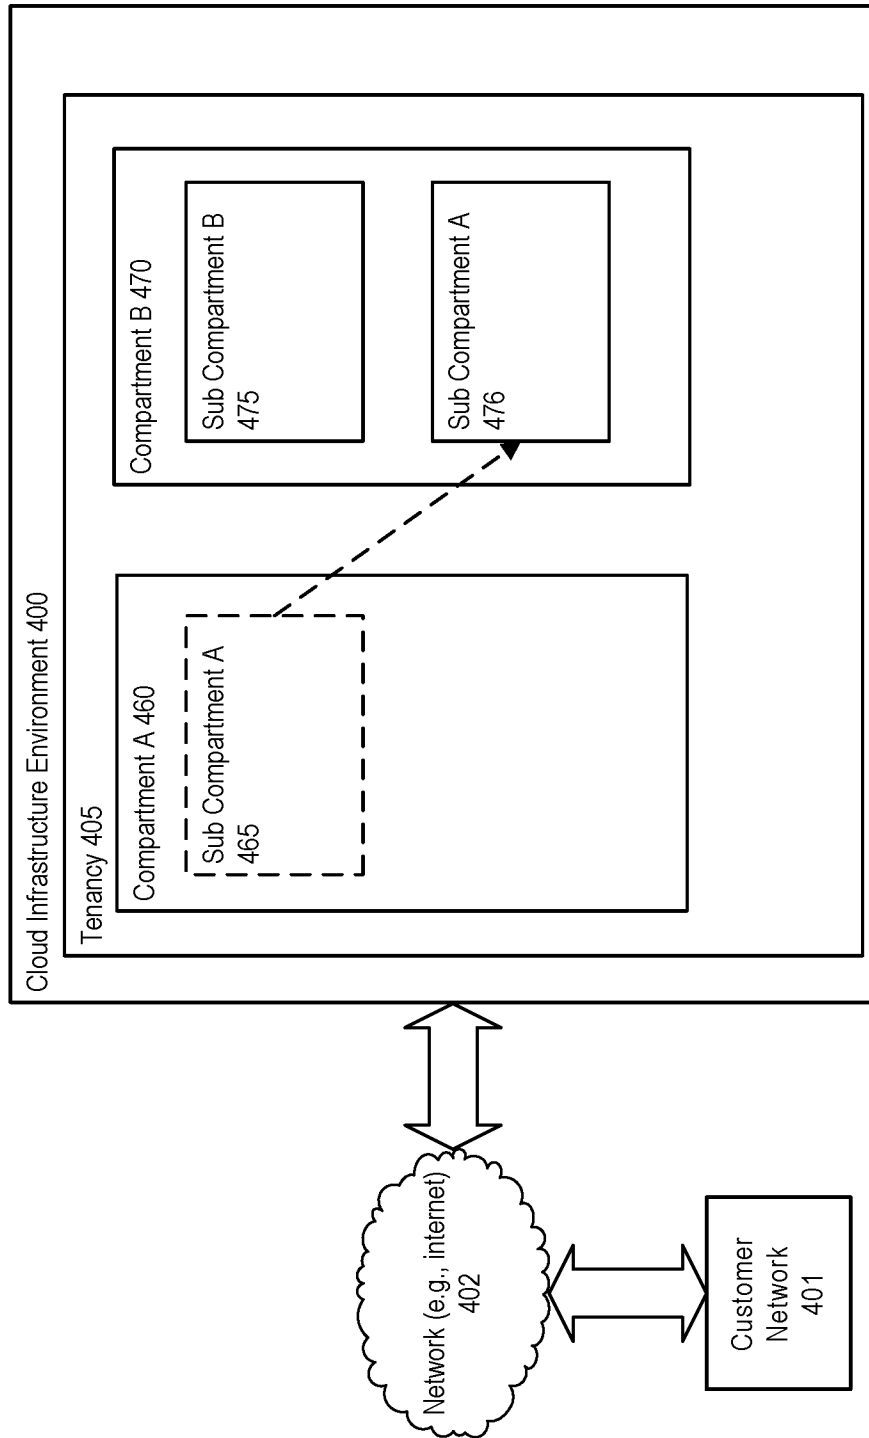


FIGURE 3



**FIGURE 4**

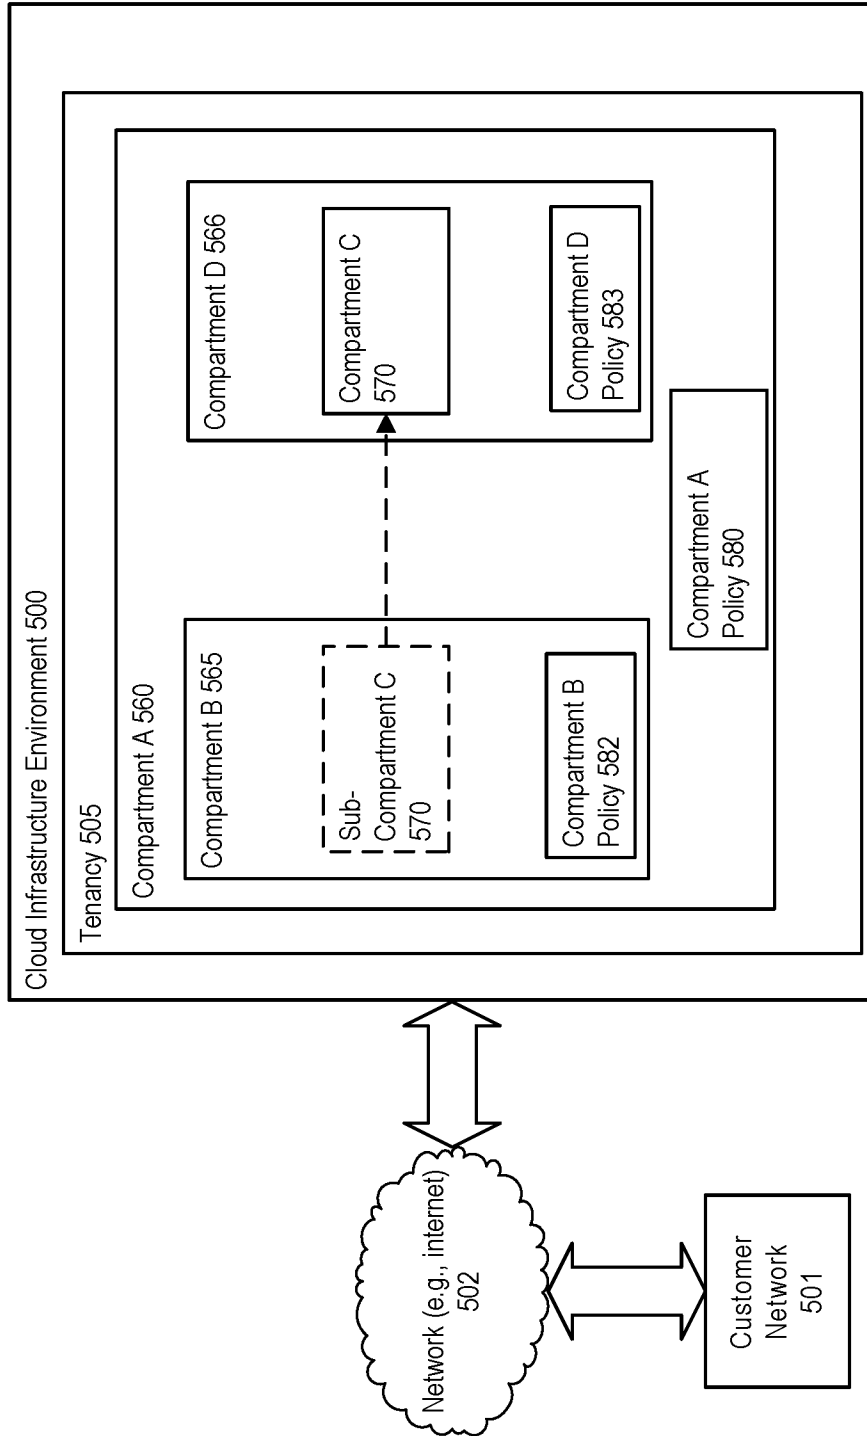


FIGURE 5

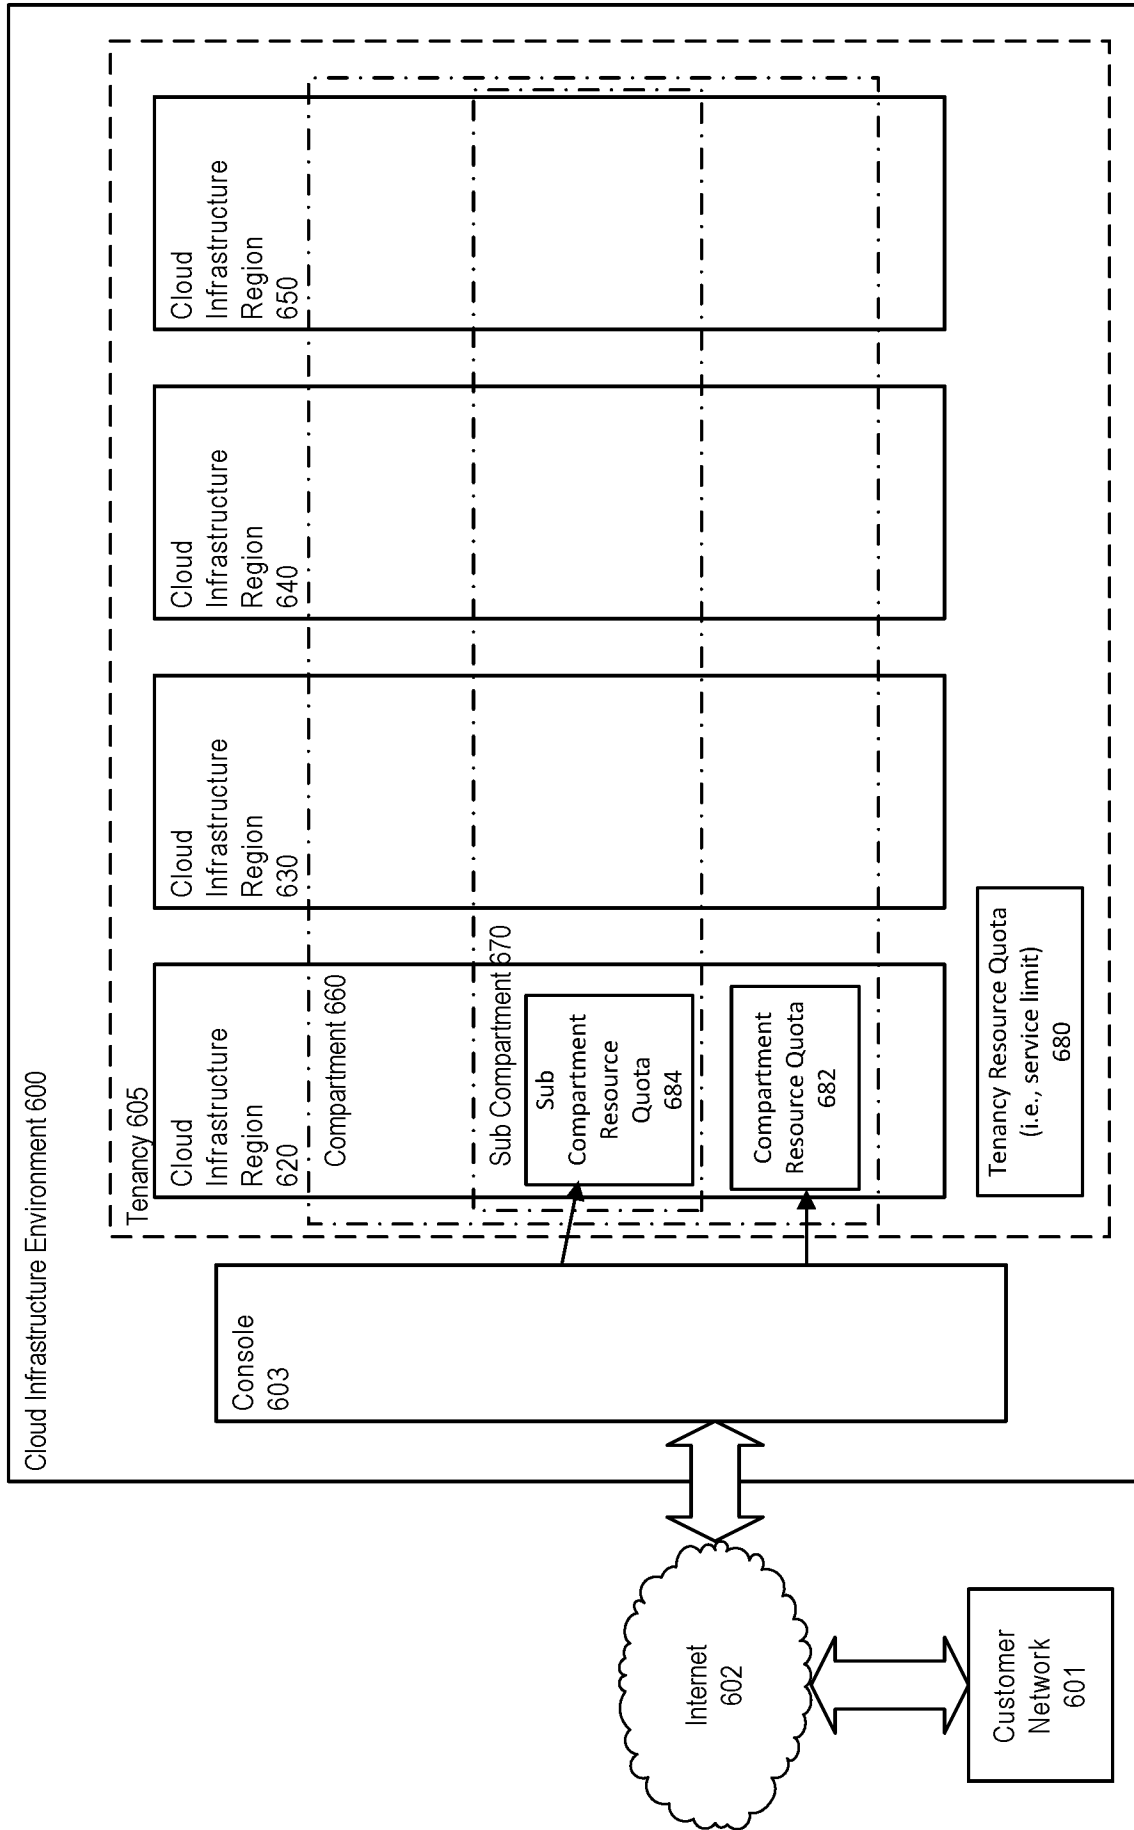


FIGURE 6

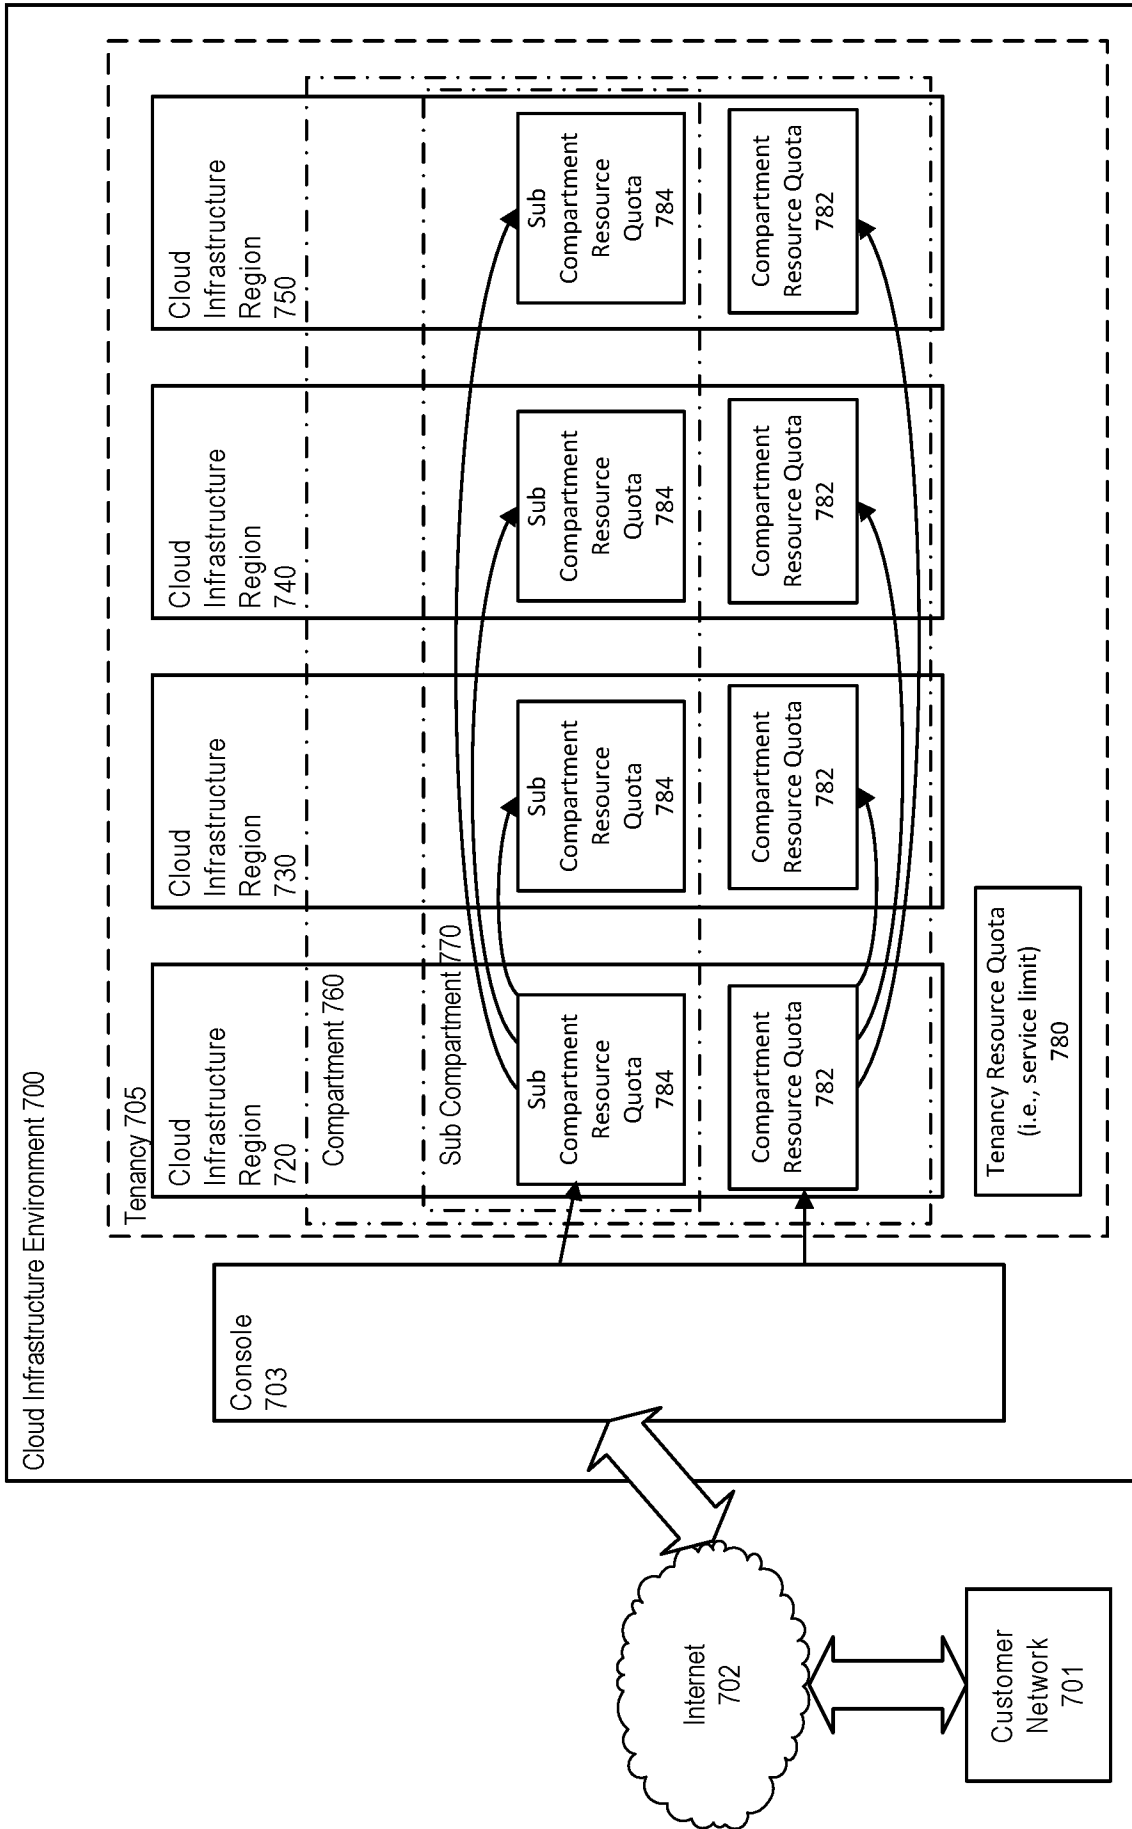


FIGURE 7




FIGURE 8




FIGURE 9

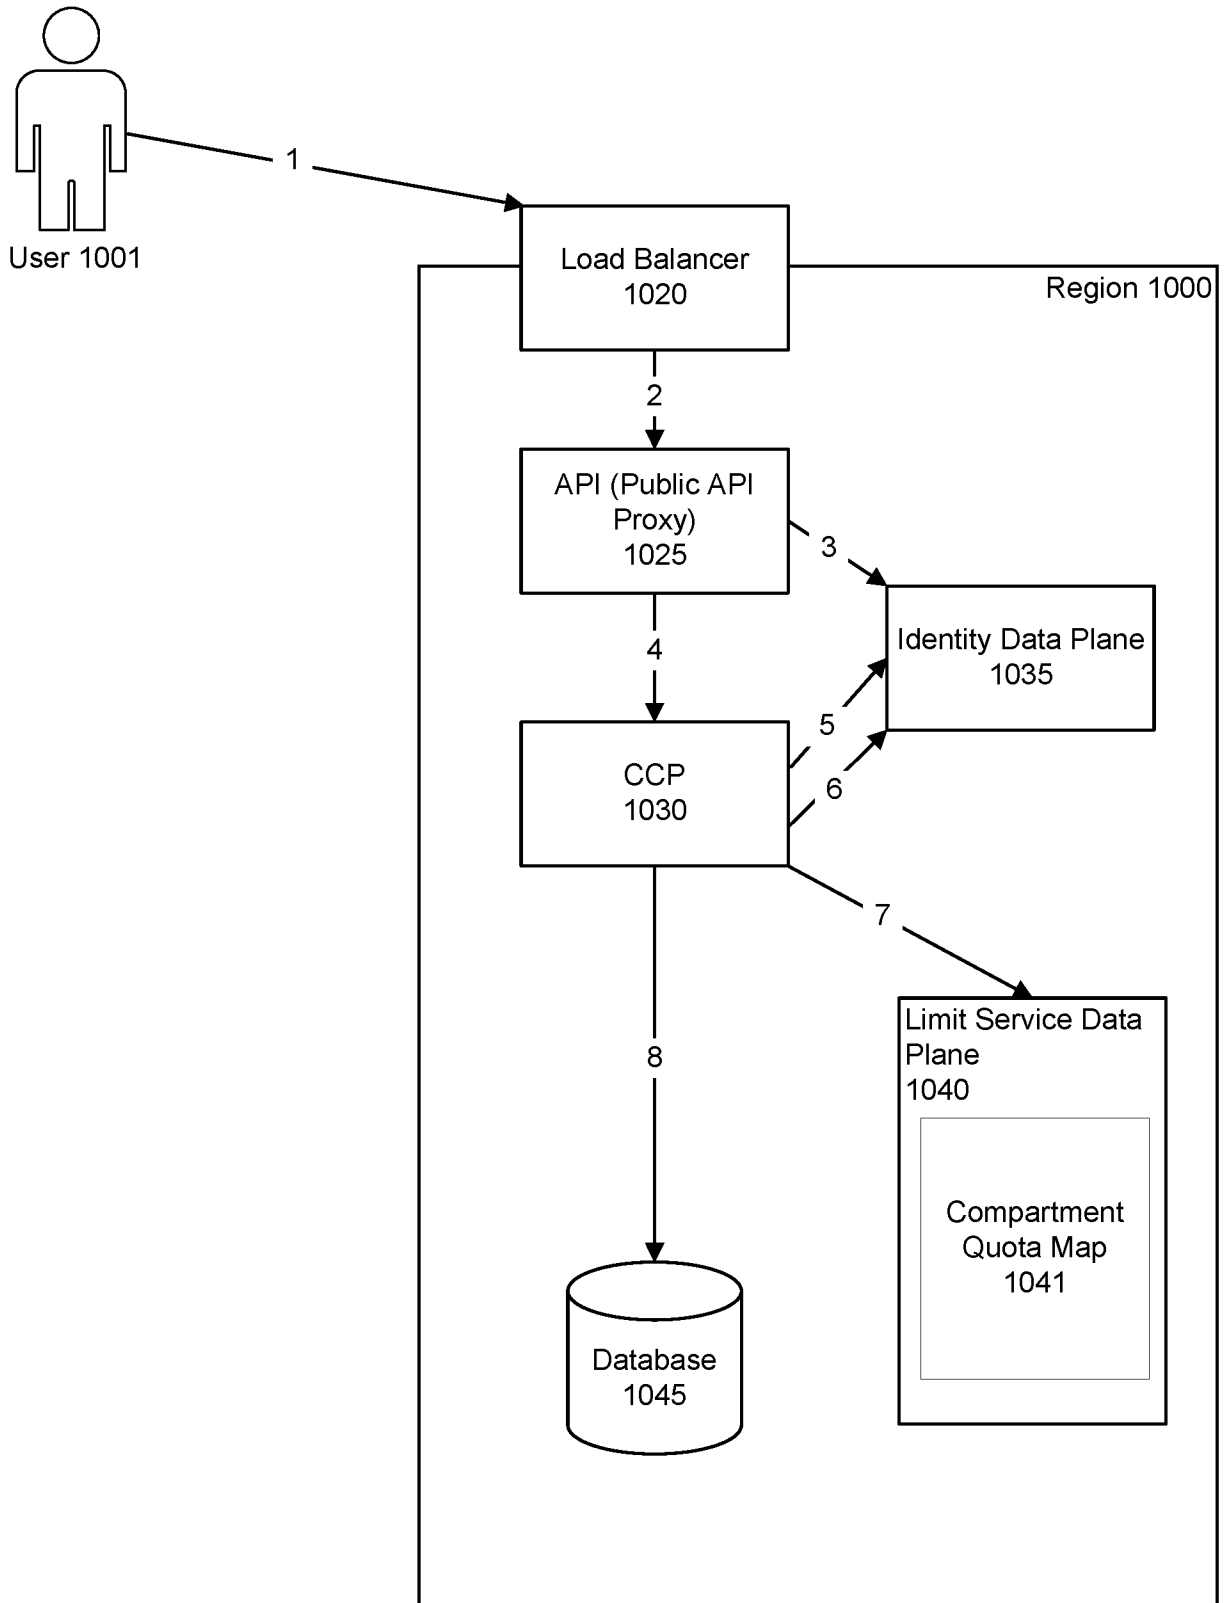
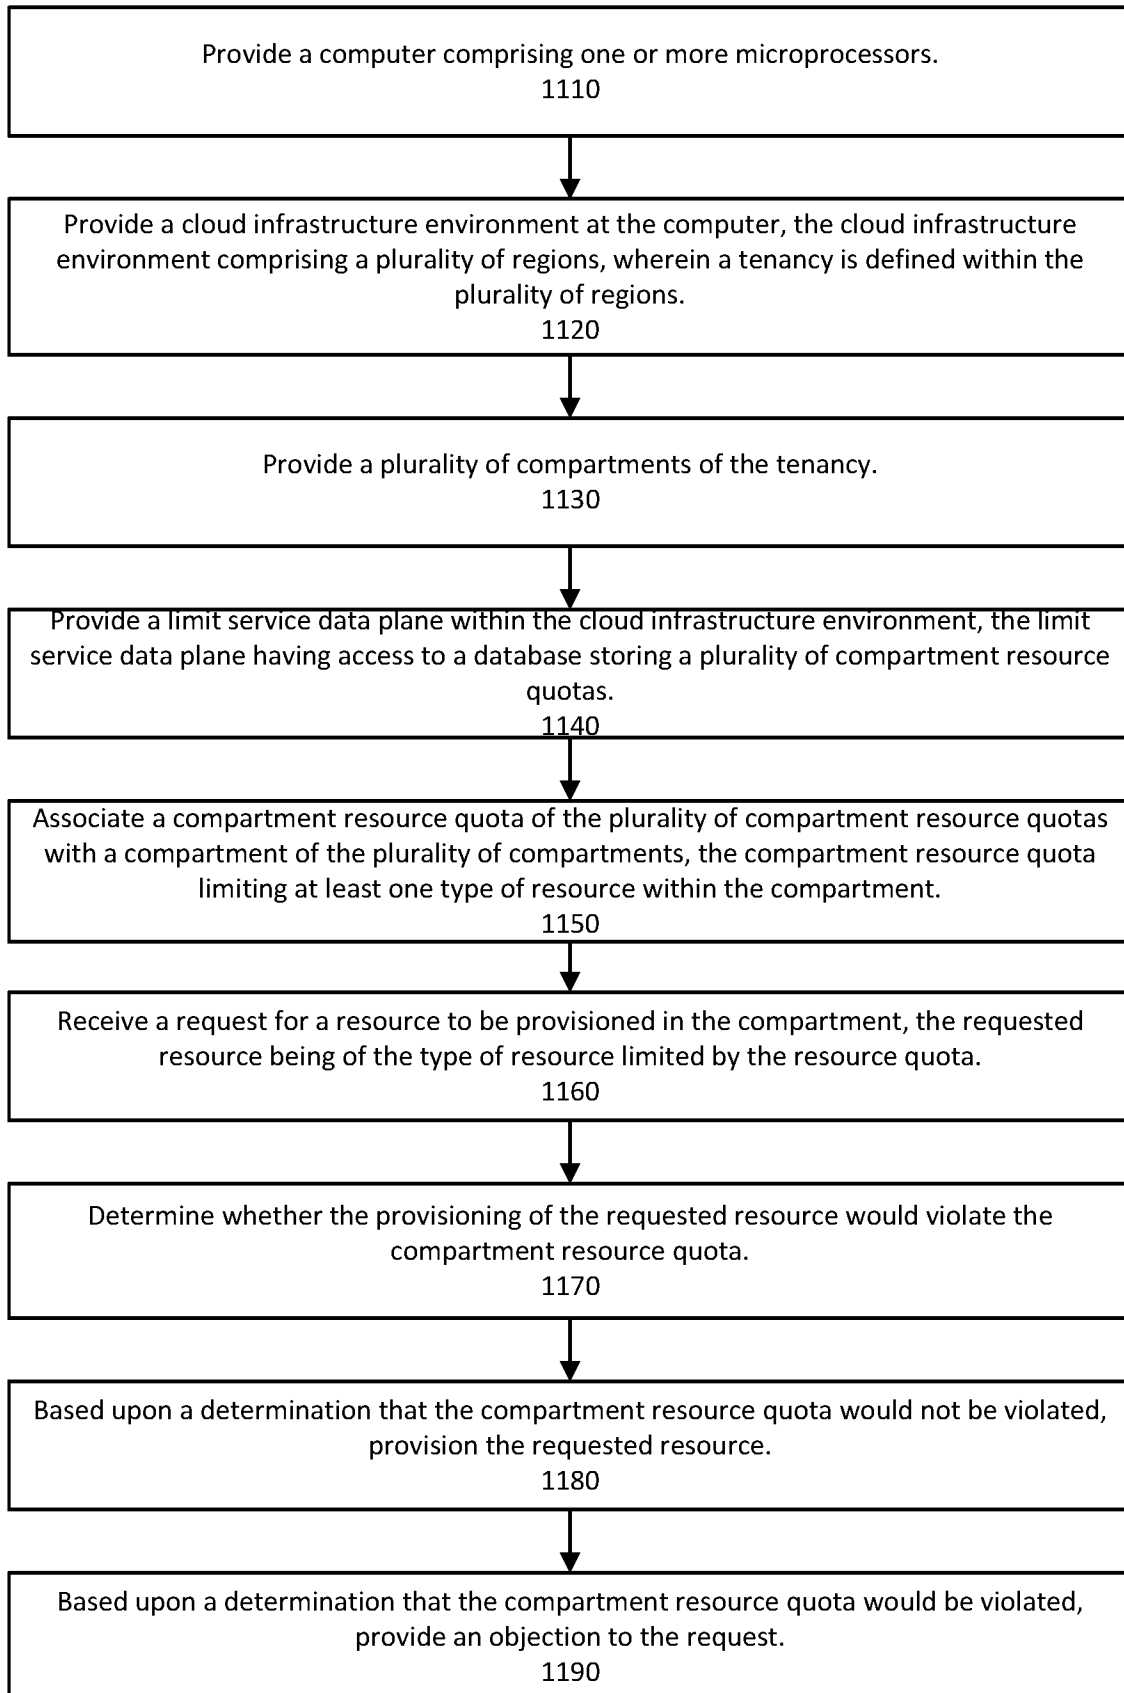




FIGURE 10


**FIGURE 11**



**FIGURE 12a**



**FIGURE 12b**



**FIGURE 12c**





FIGURE 13



**FIGURE 14**





FIGURE 15



**FIGURE 16**


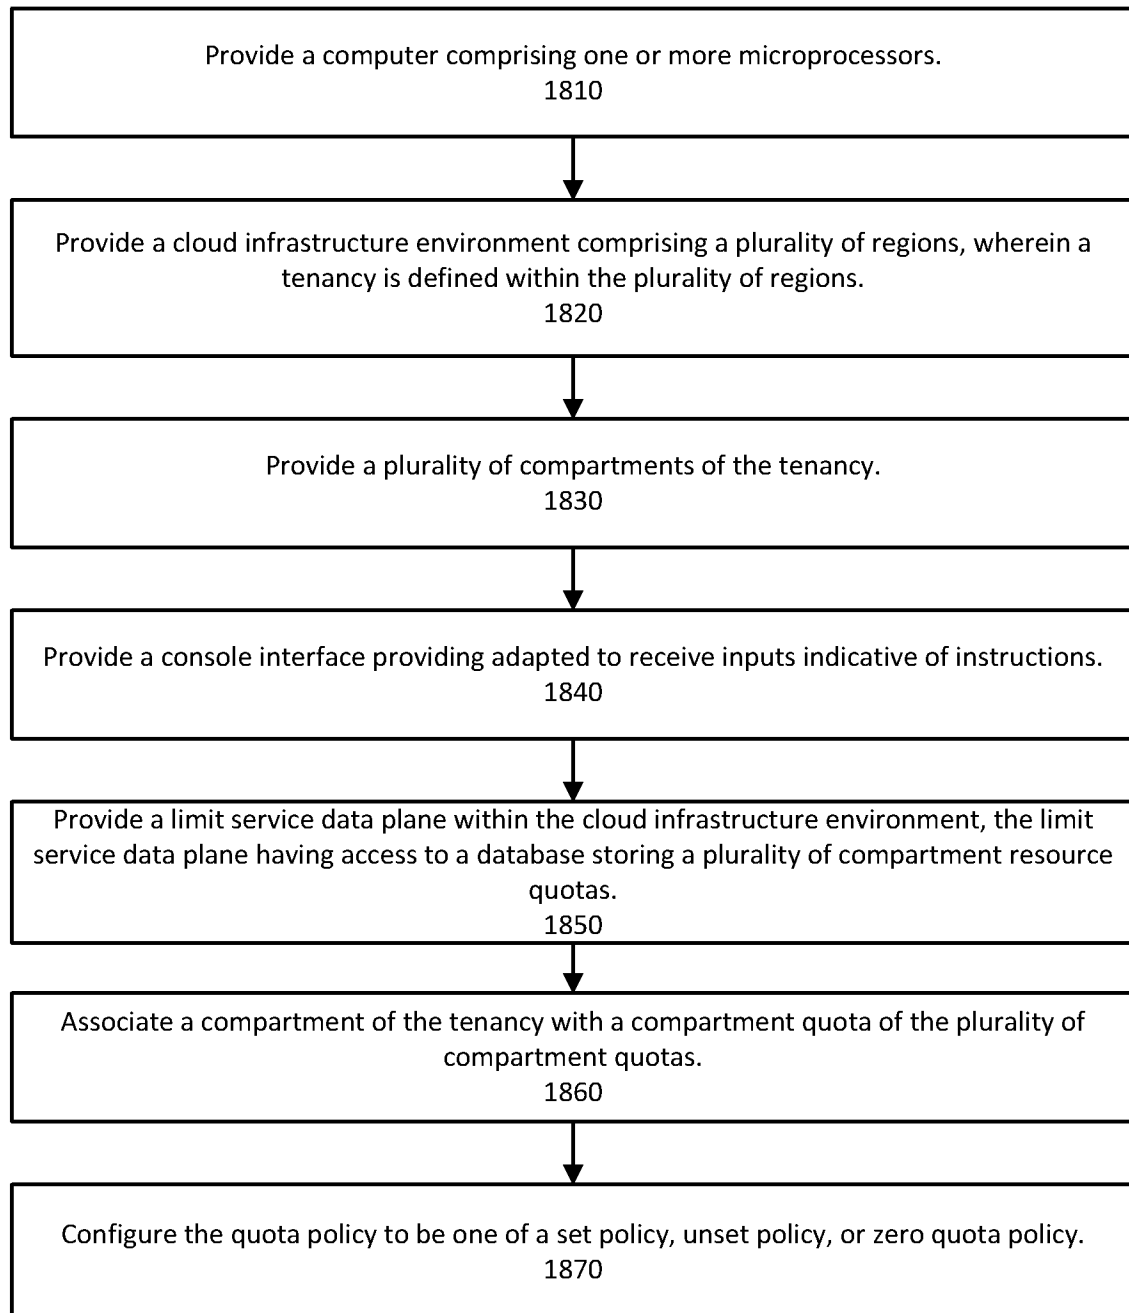


FIGURE 17

**FIGURE 18**

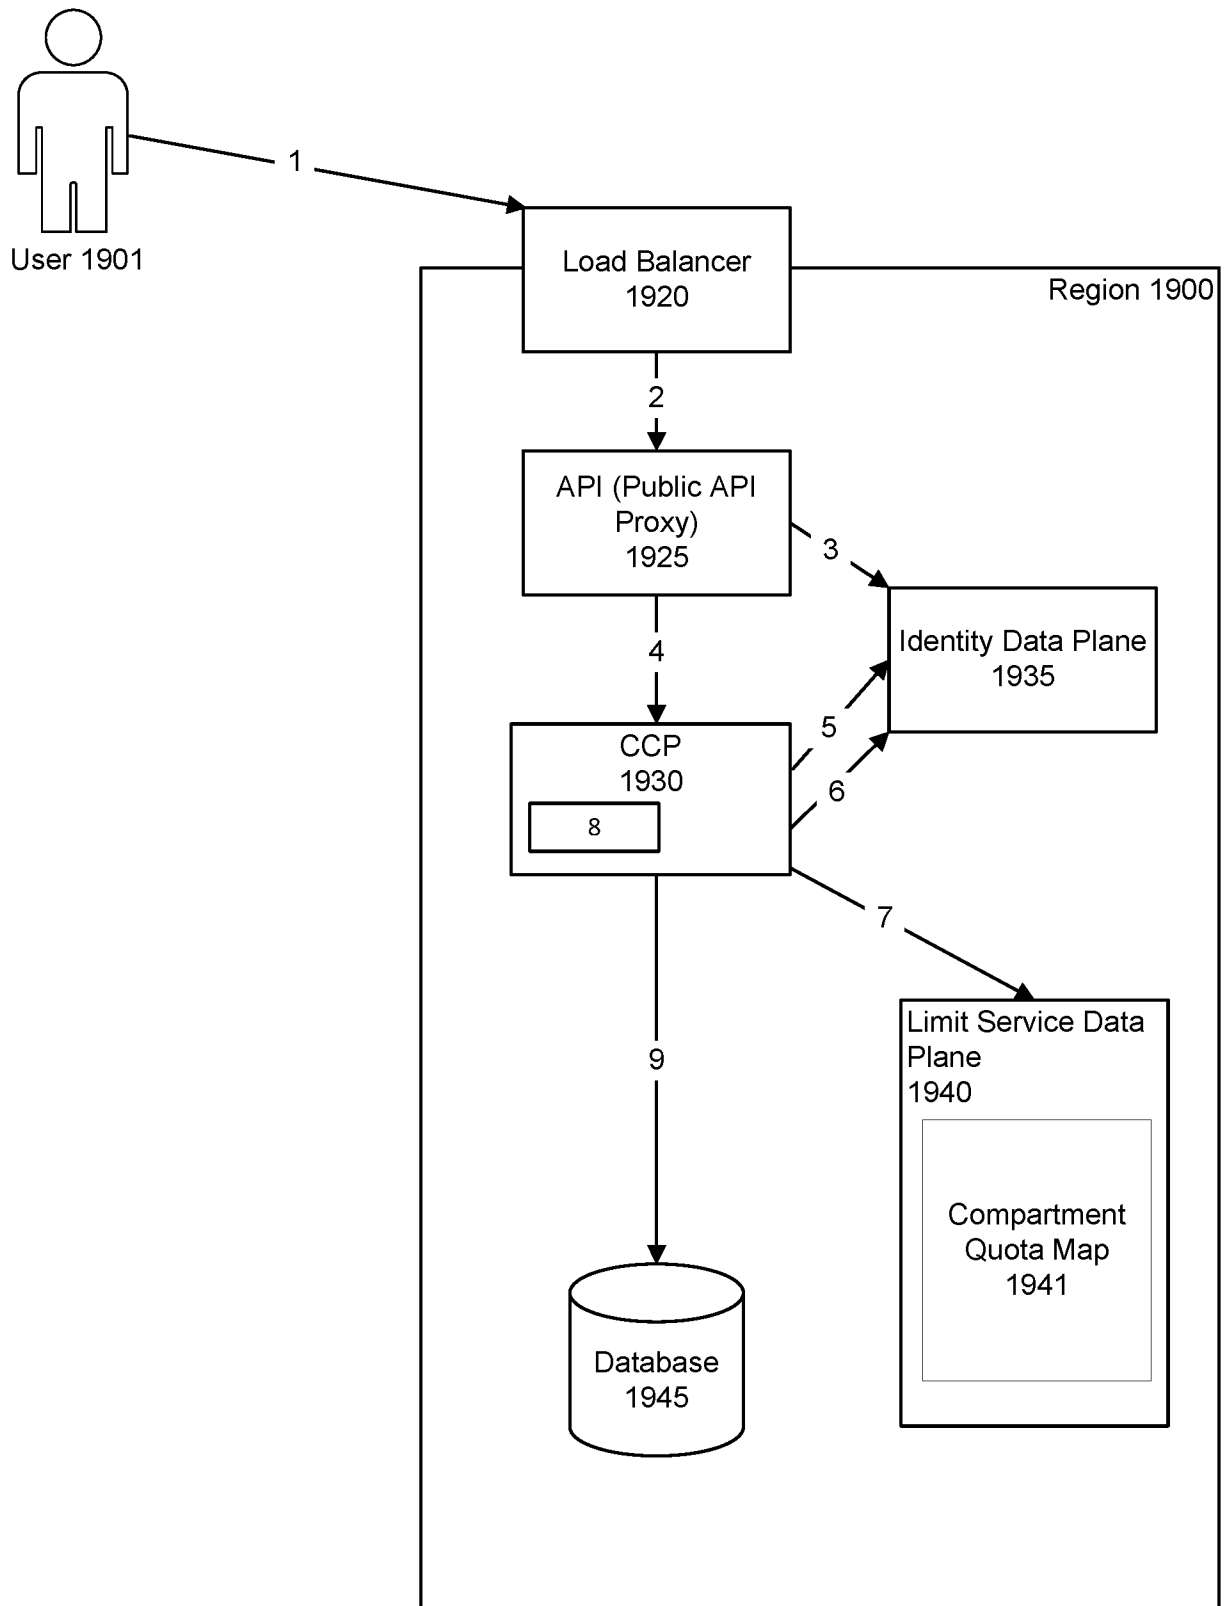
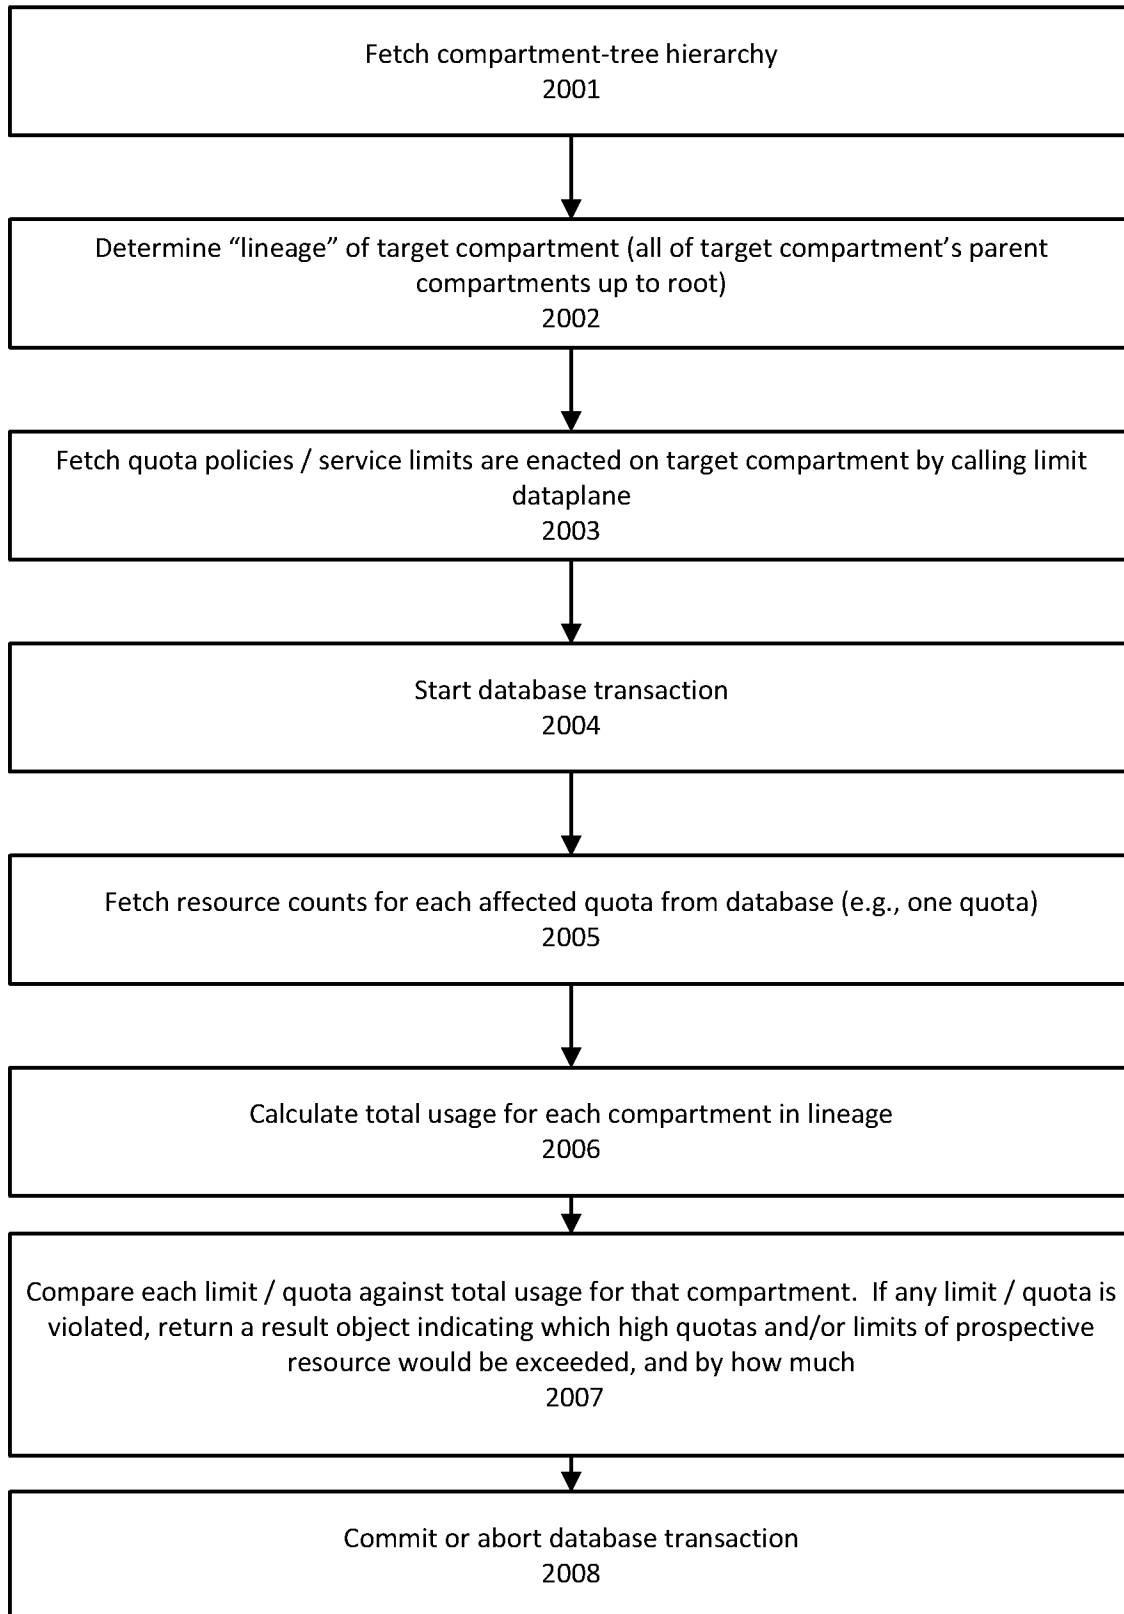
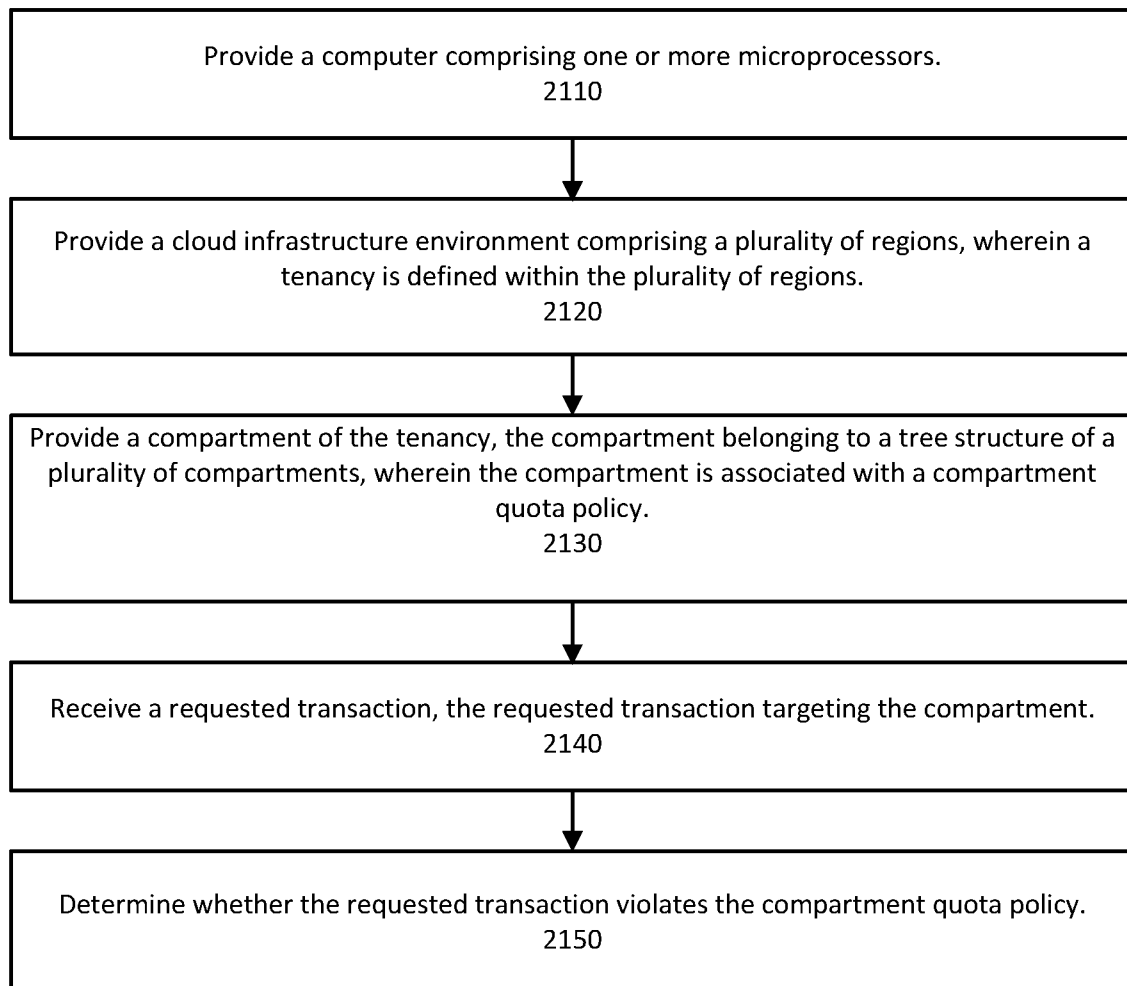


FIGURE 19

**FIGURE 20**

**FIGURE 21**

# INTERNATIONAL SEARCH REPORT

International application No PCT/US2020/045516
---

<b>A. CLASSIFICATION OF SUBJECT MATTER</b> INV. G06F9/50 ADD.				
According to International Patent Classification (IPC) or to both national classification and IPC				
<b>B. FIELDS SEARCHED</b> Minimum documentation searched (classification system followed by classification symbols) G06F				
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched				
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) EPO-Internal				
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>				
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.		
X	Anonymous: "Compartment Quotas", 16 July 2019 (2019-07-16), XP055741917, Retrieved from the Internet: URL:https://docs.cloud.oracle.com/en-us/iaas/Content/General/Concepts/resourcequotas.htm [retrieved on 2020-10-20] page 1 - page 2 page 4 - page 5 page 14	1-20		
X	----- US 2014/297781 A1 (CTERA NETWORKS, LTD.) 2 October 2014 (2014-10-02) abstract paragraphs [0002], [0006], [0011], [0014], [0028] - [0036] ----- -/--	1-20		
<table style="width: 100%; border: none;"> <tr> <td style="width: 50%; border: none;"><input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C.</td> <td style="width: 50%; border: none;"><input checked="" type="checkbox"/> See patent family annex.</td> </tr> </table>			<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C.	<input checked="" type="checkbox"/> See patent family annex.
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C.	<input checked="" type="checkbox"/> See patent family annex.			
* Special categories of cited documents :				
"A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family			
Date of the actual completion of the international search	Date of mailing of the international search report			
22 October 2020	03/11/2020			
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer  Tomàs Blanch, F			

**INTERNATIONAL SEARCH REPORT**

International application No PCT/US2020/045516
---

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2015/089065 A1 (ORACLE INTERNATIONAL CORPORATION) 26 March 2015 (2015-03-26) abstract paragraphs [0005], [0138] - [0142], [0147] - [0148], [0164] - [0165] -----	1-20
A	ANONYMOUS: Oracle / oci-cli  16 July 2019 (2019-07-16), XP002800778, Retrieved from the Internet: URL:https://github.com/oracle/oci-cli/blob/v2.5.19/services/limits/docs/inline-help/limits.txt [retrieved on 2020-10-22] the whole document -----	1-20
A	US 2016/179576 A1 (EMC CORPORATION) 23 June 2016 (2016-06-23) the whole document -----	1-20
A	US 2015/067128 A1 (CONNECTCLOUD, INC.) 5 March 2015 (2015-03-05) the whole document -----	1-20

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/US2020/045516
---

Patent document cited in search report	A1	Publication date	Patent family member(s)	Publication date
US 2014297781	A1	02-10-2014	NONE	
US 2015089065	A1	26-03-2015	US 2015089065 A1	26-03-2015
			US 2015089066 A1	26-03-2015
			US 2015089067 A1	26-03-2015
			US 2015089068 A1	26-03-2015
US 2016179576	A1	23-06-2016	CN 105893138 A	24-08-2016
			US 2016179576 A1	23-06-2016
US 2015067128	A1	05-03-2015	NONE	