

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).

Published:

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

TRANSPARENT MOUNTING OF EXTERNAL ENDPOINTS BETWEEN PRIVATE NETWORKS

CROSS-REFERENCE TO RELATED APPLICATION

- 5 [0001] The present application is a PCT application of, and claims the benefit and priority under 35 U.S.C. 119(e) of U.S. Non-Provisional Application 17/515,087, filed October 29, 2021, entitled "Transparent Mounting of External Endpoints between Private Networks," the contents of which are incorporated by reference in its entirety for all purposes.
- [0002] This application is related to Application Serial No 17/515,093, filed October 29,
10 2021, entitled "Secure Bi-Directional Network Connectivity System between Private Networks," the entire contents of which are incorporated herein by reference for all purposes.

TECHNICAL FIELD

- [0003] This disclosure relates generally to cloud-based services. More specifically, but not by way of limitation, the present disclosure describes a secure private network connectivity
15 service within a cloud infrastructure that includes improved capabilities to enable secure private network connectivity to be established between external resources residing in a customer's on-premise environment and the customer's resources residing in the cloud.

BACKGROUND

- [0004] The demand for cloud-based services continues to increase rapidly. The term cloud
20 service is generally used to refer to a service that is made available to users or customers on demand (e.g., via a subscription model) using systems and infrastructure (cloud infrastructure) provided by a cloud services provider. Typically, the servers and systems that make up the cloud service provider's infrastructure are separate from the customer's own on-premise servers and systems. Customers can thus avail themselves of cloud services provided
25 by a cloud service provider without having to purchase separate hardware and software resources for the services. There are various different types of cloud services including Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), Infrastructure-as-a-Service (IaaS), and others.

[0005] To take advantage of the numerous benefits provided by cloud services, an enterprise is often required to migrate on-premise applications and data from their local data center to the public cloud infrastructure. This process typically requires the enterprise to set up a site-to-site network connection to establish secure connectivity between their on-premise data center and the cloud infrastructure. Configuring a high performant, scalable, and highly available site-to-site network connection to process network traffic between different networks can be a complex and time-consuming task for an enterprise especially when the enterprise's on-premises applications and data scale across multiple different networks.

BRIEF SUMMARY

[0006] The present disclosure relates generally to cloud-based services. More specifically, but not by way of limitation, the present disclosure describes a secure private network connectivity service within a cloud infrastructure that includes improved capabilities to enable secure private network connectivity to be established between external resources residing in a customer's on-premise environment and the customer's resources residing in the cloud.

[0007] In certain embodiments, a secure network connectivity system (SNCS) in a cloud service provider is provided. The SNCS executes a secure network connectivity service for enabling secure private network connectivity between an on-premise network associated with a customer of the cloud service provider and a virtual cloud network (VCN) hosted by the cloud service provider for the customer. The SNCS comprises a virtual overlay network comprising a set of computing nodes.

[0008] In a certain implementation, and as part of the secure connectivity services provided by the SNCS, the SNCS registers an external resource residing in the on-premise network as an external endpoint in the virtual cloud network. The external endpoint is identified by an Internet Protocol (IP) address in the virtual cloud network. A computing node in the SNCS then creates an external resource representation for the external endpoint in the virtual cloud network. In a certain implementation, creating the external resource representation comprises creating a virtual network interface card (VNIC) and assigning the Internet Protocol (IP) address associated with the external endpoint to the VNIC. The computing node receives a request for querying information stored in the external resource residing in the on-premise network associated with the customer and establishes a connection between the VNIC created for the external resource representation in the virtual cloud network and the external resource

residing in the on-premise network. The computing node transmits the request to the external resource via the established connection and obtains a result corresponding to the request via the established connection.

[0009] In certain examples, the SNCS provides an external gateway appliance. The external gateway appliance is configured to establish secure private network connectivity between the external resource residing in the on-premise network and the set of one or more computing nodes comprising the secure network connectivity system. In certain examples, the external gateway appliance is configured in the on-premise network associated with the customer by a user of the secure network connectivity system. In certain examples, the external gateway appliance is configured to establish a secure Virtual Private Network (VPN) connection between the external resource residing in the on-premise network and the set of one or more computing nodes comprising the secure network connectivity system.

[0010] In certain examples, the computing node in the SNCS is configured to translate the IP address associated with the VNIC to a physical IP address associated with the external resource in the on-premise network and transmit the request to the physical IP address associated with the external resource. In certain examples, the computing node is configured to establish the connection to the external gateway appliance via the VNIC.

[0011] In certain examples, the SNCS is configured to create an external site representation of the on-premise network associated with the customer. The external site representation is a logical representation of the on-premise network and is identified by an external site identifier and a customer identifier. In certain examples, the external resource is registered as an external endpoint in the external site representation. In certain examples, the computing node is configured to establish the connection between the VNIC created for the external resource representation in the virtual cloud network and the external resource residing in the external site representation.

[0012] In certain examples, the VNIC is identified by a fully qualified domain name and a cloud identifier in the virtual cloud network associated with the customer. In certain examples, the external resource is a database, an application, or a compute instance residing in the on-premise network.

[0013] Various embodiments are described herein, including methods, systems, non-transitory computer-readable storage media storing programs, code, or instructions executable by one or more processors, and the like. These illustrative embodiments are

mentioned not to limit or define the disclosure, but to provide examples to aid understanding thereof. Additional embodiments are discussed in the Detailed Description, and further description is provided there.

BRIEF DESCRIPTION

5 **[0014]** FIG. 1 depicts a distributed environment that includes a secure private network connectivity service within a cloud service provider infrastructure (CSPI), according to certain embodiments.

10 **[0015]** FIG. 2 depicts additional details of the operations performed by the systems and subsystems shown in FIG. 1 for providing secure private network connectivity to a customer's external resources residing in a customer's on-premise network from the customer's virtual cloud network hosted by the CSPI, according to certain embodiments.

15 **[0016]** FIG. 3 depicts an example in which multiple external resources in the customer's external site representation can be represented as multiple Service Virtual Network Interface Cards in the customer's virtual cloud network hosted by the CSPI, according to certain embodiments,

20 **[0017]** FIG. 4 depicts an example of a process performed by the systems and subsystems shown in FIG. 1 for providing secure private network connectivity, according to certain embodiments.

25 **[0018]** FIG. 5 is a flowchart depicting the flow of network packets between an external resource representation in the customer's virtual cloud network and an external resource residing in a customer's on-premise network, according to certain embodiments.

30 **[0019]** FIG. 6 is a high level diagram of a distributed environment showing a virtual or overlay cloud network hosted by a cloud service provider infrastructure according to certain embodiments.

35 **[0020]** FIG. 7 depicts a simplified architectural diagram of the physical components in the physical network within CSPI according to certain embodiments.

40 **[0021]** FIG. 8 shows an example arrangement within CSPI where a host machine is connected to multiple network virtualization devices (NVDs) according to certain embodiments.

[0022] FIG. 9 depicts connectivity between a host machine and an NVD for providing I/O virtualization for supporting multitenancy according to certain embodiments.

[0023] FIG. 10 depicts a simplified block diagram of a physical network provided by a CSPI according to certain embodiments.

5 [0024] FIG. 11 is a block diagram illustrating one pattern for implementing a cloud infrastructure as a service system, according to at least one embodiment.

[0025] FIG. 12 is a block diagram illustrating another pattern for implementing a cloud infrastructure as a service system, according to at least one embodiment.

10 [0026] FIG. 13 is a block diagram illustrating another pattern for implementing a cloud infrastructure as a service system, according to at least one embodiment.

[0027] FIG. 14 is a block diagram illustrating another pattern for implementing a cloud infrastructure as a service system, according to at least one embodiment.

[0028] FIG. 15 is a block diagram illustrating an example computer system, according to at least one embodiment.

15 DETAILED DESCRIPTION

[0029] In the following description, for the purposes of explanation, specific details are set forth in order to provide a thorough understanding of certain embodiments. However, it will be apparent that various embodiments may be practiced without these specific details. The figures and description are not intended to be restrictive. The word “exemplary” is used
20 herein to mean “serving as an example, instance, or illustration.” Any embodiment or design described herein as “exemplary” is not necessarily to be construed as preferred or advantageous over other embodiments or designs.

[0030] The present disclosure relates generally to cloud-based services. More specifically, but not by way of limitation, a secure private network connectivity service within a cloud
25 service provider infrastructure (CSPI) is described that includes improved capabilities to provide secure private network connectivity between external resources residing in a customer’s on-premise environment and the customer’s resources residing in the cloud.

[0031] A cloud infrastructure can offer high-performance compute, storage, and network capabilities in a flexible overlay virtual network that runs on top of the physical underlay
30 network and that is securely accessible from an enterprise’s on-premises network. The cloud

infrastructure allows enterprises to manage their cloud-based workloads in the same way they manage their on-premises workloads. Thus, enterprises can get all benefits of the cloud with the same control, isolation, security, and predictable performance as their on-premises network. An enterprise can build their own networks using compute, memory, and

5 networking resources provided by the cloud. For example, a customer can use resources provided by the cloud to build one or multiple customizable and private network(s) referred to as virtual cloud networks (VCNs). A customer can deploy one or more customer resources, such as compute instances, on these customer VCNs. Compute instances can take the form of virtual machines, bare metal instances, and the like. The cloud thus provides infrastructure
10 and a set of complementary cloud services that enable enterprises (customers) to build and run a wide range of applications and services in a highly available hosted environment.

[0032] To take advantage of the numerous benefits provided by the cloud infrastructure, many enterprises perform the migration of on-premise applications and data from their local data center to the public cloud infrastructure. Migrating workloads from a local (i.e., on-
15 premise) data center of an enterprise to the cloud can be a complex and challenging process. Common challenges encountered during cloud migration involve identifying, by the enterprise, the type of migration to perform, the type of resources that need to be moved, and data dependencies between the resources. While migrating workloads, some resources (e.g., databases, applications and the like) may need to reside at the on-premise data center for a
20 certain period of time until they are able to be successfully migrated into the cloud. Certain resources (e.g., business-critical resources, resources that have data portability restrictions or resources that have strict geographic requirements) may be identified as resources that are not migrated to the cloud due to security reasons and remain at the on-premise data center. For an enterprise to be able to securely access such on-premise resources from their VCN in the
25 cloud, secure private network connectivity needs to be established between the customer's (enterprise's) on-premise data center and the customer's VCN. Setting up a secure site-to-site network connection can be a complex and time-consuming task for an enterprise. This typically requires network policy level management, by a user (e.g., an administrator) of the enterprise to set up the site-to-site network (e.g., VPN) connection, setting up multiple
30 configuration parameters, setting up VPN components (e.g., a customer gateway device, a target gateway device) for the site-to-site network connectivity and so on.

[0033] Additionally, in order to access a remote asset (e.g., an on-premise resource such as a database or application) residing in their on-premise data center from their VCN using the

site-to-site network connection, the user of an enterprise has to perform additional tasks like manually configure the gateway devices to perform route advertisements and network address translations so that the remote asset in the customer's external environment can be reached from the customer's VCN. The user also has to manually configure the remote asset in order for traffic (e.g., network packets) to reach the remote asset from the customer's VCN, configure route tables to include the routes used by the site-to-site VPN connection, enable route propagation for the route table to automatically propagate site-to-site VPN routes, update security rules and so on.

[0034] In certain embodiments, a secure private network connectivity service within a cloud service provider infrastructure (CSPI) is described that includes improved capabilities to provide secure private network connectivity between external resources residing in an enterprise's on-premise environment and the enterprise's resources residing in the cloud. The secure private network connectivity service is implemented using a secure network connectivity system (SNCS) within a cloud service provider infrastructure (CSPI). The SNCS described in the present disclosure provides several technical advancements and/or improvements over conventional cloud-based network connectivity services. The SNCS provides a high performant, scalable, and highly available site-to-site network connection for processing network traffic between a customer's on-premise environment and the CSPI by implementing a robust infrastructure of network elements and computing nodes that are configured to enable secure private network connectivity to be established between external resources residing in a customer's on-premise environment and the customer's resources residing in the cloud. By using the services provided by the SNCS, a user of an enterprise can securely access their on-premise external resources from the cloud without requiring to set an elaborate site-to-site network between their on-premise network and the cloud, without making any changes to their external resources or without configuring routes to be used by the site-to-site connection.

[0035] Referring now to the drawings, FIG. 1 depicts a distributed environment 100 that includes a secure private network connectivity service within a cloud service provider infrastructure (CSPI), according to certain embodiments. The distributed environment 100 includes multiple systems communicatively coupled to each other via one or more communication networks. These communication networks may include public and private networks. The distributed environment 100 depicted in FIG. 1 is merely an example and is not intended to unduly limit the scope of claimed embodiments. Many variations,

alternatives, and modifications are possible. For example, in some implementations, the distributed environment depicted in FIG. 1 may have more or fewer systems or components than those shown in FIG. 1, may combine two or more systems, or may have a different configuration or arrangement of systems.

5 **[0036]** As shown in the example depicted in FIG. 1, the distributed environment 100 comprises a CSPI 102 that provides services and resources that customers can subscribe to. In certain embodiments, the CSPI 102 provides a secure private network connectivity service that includes capabilities to provide secure private network connectivity between a customer's on-premise network and the customer's VCN hosted by the CSPI 102. In the
10 example shown in FIG. 1, the secure private network connectivity service may be implemented by a secure network connectivity system (SNCS) 104 within the CSPI 102. The SNCS 104 additionally includes capabilities to enable a customer of the CSPI 102 to securely access an external resource residing in the customer's on-premise network from within the customer's VCN. The secure access to the external resource is enabled by the SNCS 104 by
15 creating an external resource representation for the external resource in the customer's VCN. In a certain implementation, creating an external resource representation for the external resource comprises creating a virtual network interface card (VNIC) for the external resource in the customer's VCN. The VNIC is identified by an IP address, a host name (e.g., a fully qualified domain name (FQDN)) and a cloud identifier in the VCN. A customer of the CSPI
20 can securely access the external resource residing in their on-premise network from within their VCN by connecting to the IP address assigned to the VNIC. The SNCS 104 enables a customer of the CSPI 102 to securely access an external resource residing in their on-premise network without requiring the customer to set up elaborate site-to-site networking, without making any changes to their on-premise routing configuration or without making any changes
25 to the configuration of the external resource.

[0037] In certain approaches, the SNCS 104 enables secure access to external resources residing in a customer's on-premise network from the customer's VCN using a multi-stage connection process. In a first stage of the connection process, a user (e.g., an administrator) associated with the customer can create an "external site representation" 106 of the
30 customer's on-premise network. For instance, the user may interact with the SNCS 104 via a console user interface (UI) 108 of an application executed by a user device, via APIs or via a command line interface (CLI) executed by the user's device to create the external site representation 106. An external site representation (e.g., 106) may represent a logical or

virtual representation of the customer's external site (e.g., on-premise network/on-premise data center) and is identified by an external site identifier and a tenant (customer) identifier. By way of example, an external site representation 106 may represent a high level container resource that logically represents a portion of the customer's on-premise network and a subset of one or more external resources residing in the on-premise network.

[0038] Upon successful registration of the external site representation 106 with the CSPI 102, in a second stage of the connection process, the user can configure an external gateway appliance 112 in the external site representation 106. The external gateway appliance 112 enables secure private network connectivity to be established from the customer's external site representation 106 to the SNCS 104 across a public network (e.g., the internet 110). In certain embodiments, the external gateway appliance 112 may be a software application that is downloaded by the user as part of a download package provided by the SNCS 104 when the user subscribes to the secure private connectivity services provided by the SNCS 104. In certain implementations, the external gateway appliance 112 may be a virtual machine (VM) image that is installed by the user in the customer's external site representation 106 via the console UI 108. After installing the external gateway appliance 112, the user may register it with the SNCS 104 by providing the SNCS 104 the external site identifier and a tenant identifier (customer account) against which it is trying to register the appliance.

[0039] Upon successful installation and registration of the external gateway appliance 112, the SNCS 104 can authenticate the external gateway appliance 112 and orchestrate the setting up of a tenant-specific overlay network 128 for the customer. The tenant-specific overlay network 128 may represent a virtual overlay network that is built on top of a physical network by the SNCS 104 for each tenant (customer) who subscribes to the services provided by the SNCS. The tenant-specific overlay network 128 is used to establish secure private network connectivity between the customer's external site representation 106 and the customer's VCN 148 in the CSPI. As shown in the embodiment depicted in FIG. 1, the tenant-specific overlay network 128 may comprise a distributed and horizontally scalable fleet of computing nodes that include a set of one or more tunnel hosts (also referred to herein as tunnel virtual machines), tunnel VM-1 116 and tunnel VM-2 122 and a set of one or more resource hosts (also referred to herein as resource virtual machines), resource VM-1 130 and resource VM-2 132. A host (e.g., a tunnel host or a resource host) may be composed of a set of containers (also referred to herein as shards) that are inter-connected with each other in the tenant-specific overlay network 128. The tunnel VMs 116 and 122 are used for running per-

tenant tunnel shards. For instance, as shown in FIG. 1, the tunnel VM-1 116 is used to run a tunnel shard 120 and the tunnel VM-2 122 is used to run a tunnel shard 126 for a specific tenant/customer of the CSPI. Each tunnel shard 120 or 126 is responsible for providing secure connectivity to the customer's external site representation 106. The set of resource VMs 130 and 132 may be used for running per tenant resource shards. For instance, as shown in FIG. 1, the resource VM-1 130 is used to run a resource shard 136 and the resource VM-2 132 is used to run a resource shard 140 for the tenant/customer. A resource shard may be used to receive traffic from the customer's VCN and forward it to the customer's external resource in the external site representation. Additional details of the operations performed by the tunnel shards and the resource shards shown in FIG. 1 for providing secure connectivity between the customer's external site representation 106 and the customer's VCN 148 is described in detail in FIG. 1.

[0040] Each tunnel VM (116,122) and resource VM (130, 132) is additionally configured with a host manager (118, 134) respectively. The host managers (118, 134) represent processes executing on the tunnel VMs and the resource VMs. The host managers (118 or 134) may implement API's that are used to create the tunnel and resource shards. In certain implementations, the host managers (118, 134) may be stateless and operate in an imperative mode (i.e., as a sequence of commands for the host manager to perform) by receiving instructions from a user (via APIs 108) regarding the type of shard (tunnel or resource shard) to be created and the specific configuration of the shard. The host managers may additionally be responsible for collecting and monitoring the status of the tunnel shards and the resource shards.

[0041] After the SNCS 104 sets up the tenant-specific overlay network 128 for the customer as described above, in a third stage of the connection process, the user registers a subset of their on-premise/external resources 114A-114N in the external site representation 106. An "external resource" (e.g., 114A) may represent an on-premise resource (e.g., a database, a computing instance, an application and the like) residing in the customer's on-premise network that the customer intends to enable secure private network connectivity from within their VCN. For instance, the user (via the console UI 108) identifies the external resource (e.g., 114A) that is to be enabled secure private network connectivity from their VCN 148 and using the console UI 108 (or via APIs), registers the external resource as an external endpoint in their VCN. As part of registering the external resource, the user provides configuration information related to the external resource such as the on-premise physical IP

address associated with the external resource, the port number that the external resource is accessible at, and a hostname (or a fully qualified domain name (FQDN)) of the external resource via the console UI or APIs provided by the SNCS. The user also selects a subnet in the customer's VCN where the external endpoint for the external resource is to be created.

- 5 The SNCS receives the configuration information and creates an external endpoint for the external resource in the customer's VCN. The external endpoint is identified by an IP address, a port number and a FQDN (hostname) in the customer's VCN.

[0042] The SNCS 104 (via control plane APIs) then creates an external resource representation for the external endpoint in the customer's VCN. In a certain implementation, 10 the creation of the external resource representation comprises creating a VNIC and assigning the IP address associated with the external endpoint to the VNIC. The SNCS (via control plane APIs) then creates resource shards on the resource VMs which are capable of logically attaching the VNIC (via a worker interface) to the resource shards.

[0043] A user associated with the customer can via, a client application 144 or a private 15 endpoint 146 in the customer's VCN 148 access the on-premise external resource 114A located in the customer's remote site via the IP address of the VNIC 142 as if they were connecting to any other native resource from within their VCN. For instance, if the external resource 114A represents a database in the external site representation 106, the user may submit a query via the client application 144 (which may be a data catalog service provided 20 by the CSPI) to obtain information about metadata stored in the database. The client application 144 receives the query and transmits network packets corresponding to the query to the IP address assigned to the VNIC 142 in the customer's VCN. The network packets are then received by the resource shards (136 or 140) running in the resource VMs 130 and 132 respectively. The resource shard initiates a connection to the registered external resource 25 114A in the external site representation 106. In a certain implementation, the resource shards (136, 140) include capabilities to translate the IP of the external endpoint associated with the external resource to its real/physical on-premise IP address and route the traffic (network packets) via the tunnel shards to the external gateway appliance 112. The external gateway appliance 112 receives the network packets from the tunnel shards and, in turn, routes the 30 packets to the external resource 114A in the external site representation 106. Additional details of the operations performed by the subsystems of the resource shards and the tunnel shards in the tenant-specific overlay network 128 implemented by the SNCS 104 to establish

secure connectivity from the customer's VCN to the registered external resource 114A is discussed in detail in FIG. 2.

[0044] FIG. 2 depicts additional details of the operations performed by the systems and subsystems shown in FIG. 1 for providing secure private network connectivity to a customer's external resources residing in a customer's on-premise network from the customer's VCN hosted by the CSPI, according to certain embodiments. The systems and subsystems depicted in FIG. 2 may be implemented using software (e.g., code, instructions, program) executed by one or more processing units (e.g., processors, cores) of a computing system, hardware, or combinations thereof. The software may be stored on a non-transitory storage medium (e.g., on a memory device). The distributed environment 200 depicted in FIG. 2 is merely an example and is not intended to unduly limit the scope of claimed embodiments. Many variations, alternatives, and modifications are possible. For example, in some implementations, the distributed environment depicted in FIG. 2 may have more or fewer systems or components than those shown in FIG. 2, may combine two or more systems, or may have a different configuration or arrangement of systems.

[0045] As previously described in FIG. 1, as part of the connection process implemented by the SNCS 104 to provide secure private network connectivity between a customer's on-premise network and the customer's VCN hosted by the CSPI 102, a user (e.g., an administrator) associated with the customer creates an "external site representation" 106 of the customer's on-premise network and then configures an external gateway appliance 112 in the external site representation 106. Upon successful registration of the external gateway appliance 112, the external gateway appliance 112 establishes a Virtual Private Network (VPN) connection (also referred to herein as a VPN tunnel) to the SNCS 104 via a public network (e.g., Internet 110). The VPN connection is an encrypted connection between the customer's external site representation 106 and the customer's VCN 148. In a certain implementation, the VPN connection utilizes a secure tunneling protocol (e.g., the openVPN protocol) to establish secure private network connectivity to the SNCS 104 via the Internet 110. When the external gateway appliance 112 is installed in the external site representation 106, an agent 204 executing in the external gateway appliance 112 starts a bootstrap process to activate itself with the SNCS 104 by passing information such as a compartment identifier associated with the external gateway appliance 112 and an external site representation identifier in a configuration file to the SNCS 104. When the agent 204 bootstraps itself, it communicates with the control plane of the SNCS which starts the registration and activation

process. During the activation process, the control plane sends back all the required information (e.g., certificates, Public IP and the like) that is needed by the gateway appliance to establish a secure tunnel connection to the VPN server. The agent 204 executing in the external gateway appliance then establishes secure VPN connectivity to the SNCS 104 by
5 executing a VPN client program which opens a secure VPN tunnel connection to a VPN server installed at the SNCS. The secure tunnel terminates on a tunnel shard (e.g., 120) that is placed on a tunnel VM (e.g., 116) in the SNCS 104.

[0046] In the specific implementation depicted in FIG. 1, the external gateway appliance 112 is configured with two VPN clients, VPN client-1 206 and VPN client-2 208 that are
10 each configured to establish tunnels that terminate on two different VPN servers, VPN server-1 224 and VPN server-2 226 executing in tunnel shards, tunnel shard-1 120 and tunnel shard-2 126 respectively. In the implementation depicted in FIG. 2, every physical installation of an external gateway appliance 112 results in the establishment of two tunnels and hence if one tunnel goes down, traffic (i.e., network packets) can automatically be routed
15 across the second tunnel. While the specific implementation shown in FIG. 2 depicts two tunnels, in other implementations, the SNCS 104 may be configured to implement more than two redundant tunnels for every physical installation of the external gateway appliance 112 or may just implement a single tunnel upon installation of the external gateway appliance 112.

[0047] In certain embodiments, the external gateway appliance 112 uses a standard exterior
20 gateway protocol such as the Border Gateway Protocol (BGP) to establish BGP peering sessions with the tunnel shards 120 and 126. Using BGP, the external gateway appliance 112 exchanges routing and reachability information with the tunnel shards via a public interface 220 implemented within the external gateway appliance 112. As part of the configuration information required for the BGP peering sessions, the external gateway appliance 112
25 injects its on-premise IP address into its local route table 212 (also referred to herein as a routing information base (RIB)) which is received by the tunnel shards. The tunnel shards import the routing information into their local route tables (228 and 236) after applying appropriate route filtering policies. The route filtering is preferred to ensure that a compromised external gateway appliance 112 does not inject arbitrary routes to the tunnel
30 shards. In certain examples, the external gateway appliance 112 additionally comprises a routing manager 214. The routing manager 214 may be implemented using an open source routing manager (e.g., Zebra) that is part of a routing suite (e.g., Quagga). When the BGP peering sessions learn routes and import them into its route table 212, the BGP peering

sessions perform the best path calculation and use the routing manager 214 to add the best routes to the local kernel..

[0048] A tunnel shard (e.g., 120, 126) may be composed of a set of one or more containers. In certain implementations, a tunnel shard (120 or 126) may comprise a shell container that may be used to set up various network interfaces that enable the tunnel shard to communicate with both the external gateway appliance 112 and other shards (e.g., resource shards 136 and 140) that are part of the tenant-specific overlay network 128. In the embodiment depicted in FIG. 2, the network interfaces implemented within a shell container in a tunnel shard may include an external site interface (esi), a tunnel interface and a shard backend interface. For instance, the network interfaces implemented in tunnel shard-1 120 include a tunnel interface 216, an external site interface 222 and a shard backend interface 234. Similarly, the network interfaces implemented in tunnel shard-2 126 include a tunnel interface 218, an external site interface 223 and a shard backend interface 235. A tunnel shard (e.g., 120 or 126) is additionally composed of a VPN server (224, 226) that is used to establish a VPN tunnel to a VPN client (206, 208) executing in the external gateway appliance 112. An external site interface (222, 223) may be identified by a public IP that is known to the external gateway appliance 112 on which the VPN client runs and is used to establish a tunnel to the tunnel shard. When the VPN client (e.g., 206, 208) connects to the VPN server (224, 226), a tunnel interface (216, 218) is created and placed in a pre-configured VPN subnet of the tunnel shard.

[0049] In certain implementations, each tunnel shard (e.g., 120, 126) may utilize the Border Gateway Protocol (BGP) to establish BGP peering sessions between the external gateway appliance and the tunnel shards and BGP peering sessions between the tunnel shards and the resource shards. The BGP peering sessions are used to exchange routing and reachability information with the resource shards via the shard backend interfaces (234, 235) and to exchange routing and reachability information with the external site representation 106 via the external site interfaces (222, 223). As part of the configuration information required for the BGP peering sessions, an IP address that identifies the tenant-specific overlay network 128 is added to the route tables 228, 236 (i.e., routing information bases (RIBs)) implemented in the tunnel shards 120, 126 respectively. In a certain implementation, the Classless Inter-Domain Routing (CIDR) technique may be used for allocating the IP address to the tenant-specific overlay network. The route tables (228, 236) list the routes to particular network destinations such as to the resource shards (136, 140) and to the external site representation 106. In some cases, the route tables (228, 236) also list metrics (distances) associated with

those routes. When the BGP peering sessions with the resource shards are established, the route in the route tables is propagated to the resource shards. As shown in FIG. 2, each tunnel shard (120, 126) additionally comprises a routing manager (228, 240). The routing manager (228 or 240) may be implemented in a similar manner to the routing manager (214)

5 implemented in the external gateway appliance 112. When BGP learns routes and imports them into the route tables in the tunnel shards, it performs the best path calculation and uses the routing manager to push the best routes to the external site representation 106 and to the resource shards 136, 140.

[0050] A resource shard (136, 140) may be composed of a set of containers. In certain
10 implementations, a resource shard (136, 140) may comprise a shell container that may be used to set up a virtual tunnel endpoint VTEP (242, 243) with the tunnel shards (116, 126). Each resource shard (136, 140) may use BGP to establish peering sessions with the tunnel shards and to exchange routing and reachability information with the tunnel shards via the virtual tunnel endpoints (242, 243). A resource shard (136, 140) additionally comprises a
15 routing manager (250, 260) that is configured to perform the same functionality as the routing manager (232 or 240) implemented in the tunnel shards. Each resource shard (136, 140) additionally includes a proxy server (244, 254). The proxy servers (244 or 254) may be configured to accept connections from a client application 144 or a private endpoint 146 in the customer's VCN 148 and initiate new connections to the external resource in the external
20 site representation.

[0051] In certain implementations, for every registered external endpoint (i.e.,
corresponding to an external resource in the external site representation), a unique proxy
server container is launched and attached to a resource shard. When registering the external
resource (for e.g., as part of the third stage in the connection process described in FIG. 1), a
25 user may provide configuration information such as the IP address of the external resource,
the port number of the external resource, and the name of the external resource to the SNCS
104. This configuration information is provisioned in the proxy servers and used by a client
application 144 or a private endpoint 146 (e.g., a service) in the customer's VCN to connect
to the external resource. When the external resource is successfully registered as an external
30 endpoint in the customer's VCN by the SNCS, the SNCS 104 (via control plane APIs) creates
a VNIC and assigns the IP address associated with the external endpoint to the VNIC.

[0052] In the example depicted in FIG. 2, the registered external endpoint (i.e., external resource in the external site representation) represents an external database 202 residing in the external site representation 106. The proxy servers (244, 254) listen on the IP address assigned to the VNIC that is created for the registered external endpoint 202. This ensures that the same VNIC IP cannot be used to access other registered external resources in the external site representation. When a client application 144 or a private endpoint 146 running in the customer's VCN receives a query to obtain information stored in the external database 202, it transmits network packets to the VNIC IP in the customer's VCN. The network packets are received by worker VNICs (251, 252) that are attached to the resource shards (136, 140). The worker VNICs are configured with the IP address of the VNIC and in turn, initiate a connection to the registered external resource in the external site representation via the proxy servers (244, 254). The proxy servers (244, 254) perform a network address translation (NAT) to translate the IP address assigned to the VNIC 142 to the real IP address of the external resource 114A in the external site representation 106. The resource shards (136, 140) using the proxy servers (244, 254) initiate a connection to the external gateway appliance 112 via the tunnel shards (120, 126). The external gateway appliance 112 receives the network packets from the tunnel shards (120, 126) and, in turn, routes the packets to the registered external resource 114A in the external site representation 106. The proxy servers (244, 254) additionally include capabilities to load-balance the network traffic to the external resource across the tunnel shards. In certain embodiments, the proxy servers (244, 254) may be configured to multiplex connections from multiple client applications onto a single connection towards the external site representation.

[0053] In the implementation depicted in FIG. 2, each registered external endpoint that requires private connectivity from the customer's VCN 148 may be implemented as a Service VNIC, also referred to herein as an "SVNIC" in the customer's VCN 148. The SVNIC 254 may be implemented by the CSPI 102 using a VNIC as a service (VNICaaS) system (not shown in FIG. 2) which may represent a horizontally scalable service implemented by the CSPI 102 that is capable of hosting multiple VNICs (e.g., service VNIC 254) to process and transmit traffic between virtual cloud networks. Specifically, the VNICaaS is a virtual networking feature that enables a VNIC to be represented or used as a service (i.e., SVNIC 254). The VNICaaS provides functionality of a VNIC without requiring a specific SmartNIC or host of a compute instance within a virtual network to host the VNIC. Techniques used to represent registered endpoints as SVNICs have been described in detail US Patent

Application No. 17/175,573, entitled "Techniques for high performant virtual routing capabilities." The techniques described in US Patent Application No. 17/175,573 are merely meant as examples and are not intended to be limiting. Various other techniques may also be used to represent registered endpoints and to process and transmit traffic between virtual cloud networks in alternative embodiments. Since an SVNIC (e.g., 254) is created for every registered external endpoint (e.g., database 202), a flexible number of workers can be configured per SVNIC. In a specific implementation, every SVNIC (e.g., 254) may be associated with two worker VNICs (e.g., 251, 252) that are attached to the different resource shards (e.g., 136, 140). The worker VNICs are passed into the resource shards and configured with the SVNIC IP.

[0054] In the specific implementation depicted in FIG. 2, a single external gateway appliance 112 is installed in the external site representation 106 and is configured to establish two tunnels resulting in a total of two tunnel shards that need to be placed on the tunnel fleet implemented by the SNCS 104. For a fixed number of tunnel shards as described in this implementation, every resource shard may be pre-configured with a fixed set of BGP peers along with static Address Resolution Protocol (ARP) entries for those peers. In certain approaches, the BGP container on the tunnel shard may be configured with a "Dynamic BGP peer" in which a CIDR block of IP addresses can be specified from where BGP can accept incoming connections. This CIDR block may be set to the CIDR of the tenant-specific overlay network and hence as the resource shards grow/shrink the BGP peers on the tunnel shard also change appropriately. From the perspective of the resource shards, every resource shard now has two BGP peering connections and since each of the BGP peers propagate the same CIDR route to the on-premise network that it learns from the external site representation, every resource shard has two equal-cost paths towards the on-premise network. Similarly from the external gateway's perspective, the external gateway appliance 112 establishes two tunnels towards the SNCS 104 and the tunnel shards advertise the tenant-specific overlay network routes to the external gateway appliance 112 which then installs the appropriate route to the tenant-specific overlay network.

[0055] Using the disclosed new and improved architecture implemented by the SNCS 104, secure private network connectivity to an external resource in a customer's external site representation from the customer's VCN in the cloud can be achieved without a user (e.g., an administrator) of the enterprise having to explicitly configure the external resource, advertise routes or set up site-to-site network connectivity. As described, the secure private network

connectivity is established by the SNCS using an encrypted connection between the customer's external site representation and the customer's VCN and by authenticating the external gateway appliance prior to setting up of a tenant-specific overlay network for the customer. The SNCS 104 provides a high performant, scalable, and highly available site-to-site network connection for processing network traffic between a customer's on-premise environment and the CSPI by implementing a robust infrastructure of network elements and computing nodes (i.e., a tenant-specific overlay network 128) for each tenant/customer that uses the services provided by the SNCS 104. The tunnel hosts in the tenant-specific overlay network 128 are used to provide secure connectivity to the customer's external site representation 106 and the resource hosts are used to receive traffic from the customer's VCN and forward it to the customer's external site representation. By using the robust infrastructure of network elements and computing nodes implemented by the tenant-specific overlay network, a user of an enterprise can securely access its external resources from the cloud as if they were connecting to any other native resource within their VCN. A user of the enterprise can access its external resources without setting an elaborate site-to-site network between their on-premise network and the cloud, without making any changes to their external resources and without configuring routes to be used by the site-to-site connection.

[0056] In the embodiment depicted in FIG. 2, the SNCS 104 was used to establish secure private network connectivity between a single external resource (i.e., database 202) residing in the customer's on-premise network and the representation of the external resource (i.e., SVNIC 254) residing in the customer's VCN. In alternate approaches, a customer may utilize the services provided by the SNCS 104 to establish secure private network connectivity between multiple external resources residing in the customer's on-premise network and the representation of these resources residing in the customer's VCN. FIG. 3 depicts an example in which multiple external resources in the customer's external site representation can be represented as multiple SVNICS in the customer's VCN hosted by the CSPI, according to certain embodiments. In the implementation depicted in FIG. 3, SVNIC-A 306 and SVNIC-B 308 are associated with two external resource representations in the customer's VCN 302. The two external resource representations SVNIC-A 306 and SVNIC-B 308 correspond to two registered external endpoints (i.e., external resources 114A, 114B in the customer's external site representation) in the customer's on-premise network. In this example, SVNIC-A 306 is associated with worker VNIC, worker-A 310 and SVNIC-B 308 is associated with worker VNIC, worker-B 312. The worker VNIC, worker-A 310 is configured with the IP

address (50.0.0.10) corresponding to SVNIC-A 306 and resides in two different resource shards 318 and 320. The worker VNIC, worker-B 310 is configured with the IP address (50.0.0.20) corresponding to SVNIC-B 308 and also resides in different resource shards 318 and 320. The resource shards 318 and 320 may reside in a tenant-specific overlay network similar to the tenant-specific overlay network 128 described in FIG. 2.

[0057] In this example, if a single route table is implemented in a resource shard (318 or 320), the resource shard may encounter a route lookup failure when the resource shard attempts to send network packets out of a specific worker VNIC (e.g., worker-A 310 or worker-B 312). This is because if a client application 304 in the customer's VCN 302 tries to ping SVNIC-B 308 for a return path when the resource shard (318 or 320) does a route lookup it picks up the first matching route and attempts to send the packet out of worker-A 310. This will ultimately fail due to reverse path filtering. In order to avoid this, in certain approaches, each resource shard (318 and 320) is implemented with per SVNIC routing tables. For instance, resource shard 318 may be implemented with a route table 314A for SVNIC-A and a route table 316A for SVNIC-B. Similarly, resource shard 320 may be implemented with a route table 314B for SVNIC-A and a route table 316B for SVNIC-B. The routing tables can use a policy-based scheme in which based on the source IP address (50.0.0.20 or 50.0.0.10) of a network packet, the resource shards can perform a route lookup in the appropriate table.

[0058] FIG. 4 depicts an example of a process performed by the systems and subsystems shown in FIG. 1 for providing secure private network connectivity, according to certain embodiments. The processing depicted in FIG. 4 may be implemented in software (e.g., code, instructions, program) executed by one or more processing units (e.g., processors, cores) of the respective systems, hardware, or combinations thereof. The software may be stored on a non-transitory storage medium (e.g., on a memory device). The process 400 presented in FIG. 4 and described below is intended to be illustrative and non-limiting. Although FIG. 4 depicts the various processing steps occurring in a particular sequence or order, this is not intended to be limiting. In certain alternative embodiments, the steps may be performed in some different order or some steps may also be performed in parallel. In certain embodiments, such as in the embodiment depicted in FIG. 1, the processing depicted in FIG. 4 may be performed by the computing nodes (e.g., 116, 122, 130 and 132) comprising the tenant-specific overlay network 128. In certain embodiments, within the tenant-specific overlay network, the

processing may be performed by one or more computing nodes such as the resource virtual machine (130 or 132) in the tenant-specific overlay network.

[0059] The processing depicted in FIG. 4 assumes that a user (e.g., an administrator) associated with the customer has created an external site representation (e.g., 106) of the customer's on-premise network and has configured an external gateway appliance (e.g., 112) in the external site representation 106. The processing depicted in FIG. 4 further assumes that the SNCS 104 has authenticated the external gateway appliance 112 and has configured/established a tenant-specific overlay network (e.g., 128) comprising a distributed and horizontally scalable fleet of computing nodes for the customer. As described in FIG. 1, the set of one or more computing nodes include resource hosts (i.e., resource virtual machines, 130, 132) and tunnel hosts (i.e., tunnel virtual machines, 116, 122).

[0060] The processing depicted in FIG. 4 may be initiated at block 402 when the SNCS 104 executes the tenant-specific overlay network (e.g., 128) comprising a distributed and horizontally scalable fleet of computing nodes. The tenant-specific overlay network is used to establish secure private network connectivity between the customer's external site representation and the customer's VCN (e.g., 148) in the CSPI.

[0061] At block 404, the SNCS (via control plane APIs) registers an external resource residing in the customer's on-premise network as an external endpoint in the customer's VCN. The external endpoint is identified by an IP address in the customer's VCN. As previously described, the external resource may represent a database, a computing instance, or an application in the external site representation 106 that the customer intends to enable secure private network connectivity from within their VCN. As part of registering the external resource, a user of the SNCS provides configuration information related to the external resource such as the on-premise physical IP address associated with the external resource, the port number that the external resource is accessible at, and a hostname (or a fully qualified domain name (FQDN)) of the external resource via the console UI or APIs. The user also selects a subnet in the customer's VCN where the external endpoint for the external resource is to be created. Based on the configuration information, the SCNS creates an external endpoint for the external resource in the customer's VCN. The external endpoint is identified by an IP address, a port number and a FQDN (hostname) in the customer's VCN.

[0062] At block 406, a computing node (e.g., a resource VM 130 or resource VM 132) in the SNCS creates an external resource representation for the external endpoint in the

customer's VCN. In a certain implementation, the creation of an external resource representation involves creating, by the computing node, a VNIC at block 408 and assigning the IP address of the external endpoint to the VNIC at block 410.

[0063] At block 412, the computing node receives a request for querying information stored in the external resource residing in the on-premise network. For instance, as described in FIG. 1, a user may submit a query via the client application (e.g., 144) to obtain information about metadata stored in the external resource (e.g., 114A). The client application receives the query and transmits network packets corresponding to the query to the IP address assigned to the VNIC 142 in the customer's VCN. The network packets are then received by the resource shards (136 or 140) running in the resource VMs.

[0064] At block 414, the computing node (i.e., the resource shard 136 or 140) establishes a connection to the registered external resource 114A in the external site representation 106 from the VNIC via a tunnel shard (e.g., 120 or 126). In a specific implementation, a resource shard (136 or 140) includes capabilities to translate the IP address assigned to the VNIC 142 to the real IP address of the external resource 114A residing in the external site representation. Specifically, as described in relation to FIG. 2, a proxy server (244 or 254) in the resource shard (136 or 140) is configured with capabilities to perform a network address translation (NAT) to translate the IP address assigned to the VNIC 142 to the real IP address of the external resource 114A installed in the external site representation 106. Using the proxy servers (244 or 254), the resource shard (136 or 140) initiates a connection to the external gateway appliance 112 via the tunnel shards (120 or 126) which in turn, initiates a connection to the registered external resource 114A in the external site representation 106. Additional details of the operations performed by the proxy servers to perform a network address translation (NAT) to translate the IP address assigned to the VNIC to the real IP address of the external resource is described in detail in FIG. 5.

[0065] At block 416, the computing node transmits the request to the external resource residing in the external site representation via the connection established in block 410. At block 418, the computing node obtains a result corresponding to the request via the established connection. The result is then transmitted to the client application in the customer's VCN via the VNIC associated with the external resource representation in the customer's VCN. For example, if the external resource is a database executing in the external

site representation, the result may include information stored in one or more tables in the database.

[0066] FIG. 5 is a flowchart depicting the flow of network packets between an external resource representation in the customer's virtual cloud network and an external resource residing in a customer's on-premise network, according to certain embodiments. The processing depicted in FIG. 5 may be implemented in software (e.g., code, instructions, program) executed by one or more processing units (e.g., processors, cores) of the respective systems, hardware, or combinations thereof. The software may be stored on a non-transitory storage medium (e.g., on a memory device). The process 500 presented in FIG. 5 and described below is intended to be illustrative and non-limiting. Although FIG. 5 depicts the various processing steps occurring in a particular sequence or order, this is not intended to be limiting. In certain alternative embodiments, the steps may be performed in some different order or some steps may also be performed in parallel. In certain embodiments, such as in the embodiment depicted in FIG. 2, the processing depicted in FIG. 5 may be performed by the computing nodes (e.g., 116, 122, 130 and 132) comprising the tenant-specific overlay network 128.

[0067] The processing depicted in FIG. 5 may be initiated at block 502 when a client application (e.g., 144) in the customer's VCN receives a query for information associated with an external resource (e.g., 202) residing in the customer's on-premise network (i.e., the customer's external site representation 106). For instance, the query may be received from a user associated with the customer via the client application. As part of the processing performed in block 502, the client application 144 initiates a connection to the VNIC (e.g., 206) associated with the external resource representation by transmitting network packets corresponding to the query to the IP address assigned to the VNIC in the customer's VCN.

[0068] At block 504, the client application transmits the network packets corresponding to the query to the resource shards (136 or 140) running in the resource VMs in the SNCS. Specifically, and as shown in FIG. 2, the network packets may be received by worker VNICs (251, 252) that are attached to the resource shards (136, 140). The proxy servers (244, 254) in the resource shards (136, 140) perform a network address translation (NAT) to translate the IP address assigned to the VNIC to the real IP address of the external resource 114A in the external site representation 106 and initiate a connection to the external gateway appliance 112 via the tunnel shards (120, 126).

[0069] At block 506, the resource shards transmit the network packets to the tunnel shards (120, 126). As part of the processing performed in block 506, the tunnel shards may encrypt the network packets received from the resource shards prior to transmitting the packets to the external gateway appliance 112 in the external site representation 106.

5 [0070] At block 508, the tunnel shards transmit the network packets to the external gateway appliance 112 residing in the customer's external site representation. As part of the processing performed in block 508, the external gateway appliance 112 decrypts the network packets prior to transmitting the network packets to the external resource (e.g., 202) in the external site representation (e.g., 106).

10 [0071] At block 510, the external gateway appliance 112 transmits/routes the network packets to the external resource (e.g., 202) residing in the customer's external site representation. The external resource receives the network packets corresponding to the query and generates response network packets corresponding to the query. The response network packets may then be transmitted by the external resource back to the client application. For
 15 example, as part of a response network packet flow, network packets corresponding to a response to the query are transmitted from the external resource to the external gateway appliance. The external gateway appliance encrypts the network packets and transmits the encrypted network packets to the tunnel shards. The tunnel shards receive the encrypted packets and transmit the packets to the resource shards. The resource shards decrypt the
 20 network packets and perform a reverse network address translation (NAT) to translate the realIP address assigned to the external resource to the IP address of the VNIC that is associated with the external resource representation in the customer's VCN. The resource shards then transmit the response network packets to the VNIC which in turn transmits the response packets to the requesting client application 144 in the customer's VCN 148.

25 Example Virtual Networking Architecture

[0072] The term cloud service is generally used to refer to a service that is made available by a cloud services provider (CSP) to users or customers on demand (e.g., via a subscription model) using systems and infrastructure (cloud infrastructure) provided by the CSP.

Typically, the servers and systems that make up the CSP's infrastructure are separate from the
 30 customer's own on-premise servers and systems. Customers can thus avail themselves of cloud services provided by the CSP without having to purchase separate hardware and software resources for the services. Cloud services are designed to provide a subscribing

customer easy, scalable access to applications and computing resources without the customer having to invest in procuring the infrastructure that is used for providing the services.

[0073] There are several cloud service providers that offer various types of cloud services. There are various different types or models of cloud services including Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), Infrastructure-as-a-Service (IaaS), and others.

[0074] A customer can subscribe to one or more cloud services provided by a CSP. The customer can be any entity such as an individual, an organization, an enterprise, and the like. When a customer subscribes to or registers for a service provided by a CSP, a tenancy or an account is created for that customer. The customer can then, via this account, access the subscribed-to one or more cloud resources associated with the account.

[0075] As noted above, infrastructure as a service (IaaS) is one particular type of cloud computing service. In an IaaS model, the CSP provides infrastructure (referred to as cloud services provider infrastructure or CSPI) that can be used by customers to build their own customizable networks and deploy customer resources. The customer's resources and networks are thus hosted in a distributed environment by infrastructure provided by a CSP. This is different from traditional computing, where the customer's resources and networks are hosted by infrastructure provided by the customer.

[0076] The CSPI may comprise interconnected high-performance compute resources including various host machines, memory resources, and network resources that form a physical network, which is also referred to as a substrate network or an underlay network. The resources in CSPI may be spread across one or more data centers that may be geographically spread across one or more geographical regions. Virtualization software may be executed by these physical resources to provide a virtualized distributed environment. The virtualization creates an overlay network (also known as a software-based network, a software-defined network, or a virtual network) over the physical network. The CSPI physical network provides the underlying basis for creating one or more overlay or virtual networks on top of the physical network. The physical network (or substrate network or underlay network) comprises physical network devices such as physical switches, routers, computers and host machines, and the like. An overlay network is a logical (or virtual) network that runs on top of a physical substrate network. A given physical network can support one or multiple overlay networks. Overlay networks typically use encapsulation techniques to differentiate between traffic belonging to different overlay networks. A virtual

or overlay network is also referred to as a virtual cloud network (VCN). The virtual networks are implemented using software virtualization technologies (e.g., hypervisors, virtualization functions implemented by network virtualization devices (NVDs) (e.g., smartNICs), top-of-rack (TOR) switches, smart TORs that implement one or more functions performed by an NVD, and other mechanisms) to create layers of network abstraction that can be run on top of the physical network. Virtual networks can take on many forms, including peer-to-peer networks, IP networks, and others. Virtual networks are typically either Layer-3 IP networks or Layer-2 VLANs. This method of virtual or overlay networking is often referred to as virtual or overlay Layer-3 networking. Examples of protocols developed for virtual networks include IP-in-IP (or Generic Routing Encapsulation (GRE)), Virtual Extensible LAN (VXLAN — IETF RFC 7348), Virtual Private Networks (VPNs) (e.g., MPLS Layer-3 Virtual Private Networks (RFC 4364)), VMware's NSX, GENEVE (Generic Network Virtualization Encapsulation), and others.

[0077] For IaaS, the infrastructure (CSPI) provided by a CSP can be configured to provide virtualized computing resources over a public network (e.g., the Internet). In an IaaS model, a cloud computing services provider can host the infrastructure components (e.g., servers, storage devices, network nodes (e.g., hardware), deployment software, platform virtualization (e.g., a hypervisor layer), or the like). In some cases, an IaaS provider may also supply a variety of services to accompany those infrastructure components (e.g., billing, monitoring, logging, security, load balancing and clustering, etc.). Thus, as these services may be policy-driven, IaaS users may be able to implement policies to drive load balancing to maintain application availability and performance. CSPI provides infrastructure and a set of complementary cloud services that enable customers to build and run a wide range of applications and services in a highly available hosted distributed environment. CSPI offers high-performance compute resources and capabilities and storage capacity in a flexible virtual network that is securely accessible from various networked locations such as from a customer's on-premises network. When a customer subscribes to or registers for an IaaS service provided by a CSP, the tenancy created for that customer is a secure and isolated partition within the CSPI where the customer can create, organize, and administer their cloud resources.

[0078] Customers can build their own virtual networks using compute, memory, and networking resources provided by CSPI. One or more customer resources or workloads, such as compute instances, can be deployed on these virtual networks. For example, a customer

can use resources provided by CSPI to build one or multiple customizable and private virtual network(s) referred to as virtual cloud networks (VCNs). A customer can deploy one or more customer resources, such as compute instances, on a customer VCN. Compute instances can take the form of virtual machines, bare metal instances, and the like. The CSPI thus provides
5 infrastructure and a set of complementary cloud services that enable customers to build and run a wide range of applications and services in a highly available virtual hosted environment. The customer does not manage or control the underlying physical resources provided by CSPI but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., firewalls).

10 **[0079]** The CSP may provide a console that enables customers and network administrators to configure, access, and manage resources deployed in the cloud using CSPI resources. In certain embodiments, the console provides a web-based user interface that can be used to access and manage CSPI. In some implementations, the console is a web-based application provided by the CSP.

15 **[0080]** CSPI may support single-tenancy or multi-tenancy architectures. In a single tenancy architecture, a software (e.g., an application, a database) or a hardware component (e.g., a host machine or a server) serves a single customer or tenant. In a multi-tenancy architecture, a software or a hardware component serves multiple customers or tenants. Thus, in a multi-tenancy architecture, CSPI resources are shared between multiple customers or
20 tenants. In a multi-tenancy situation, precautions are taken and safeguards put in place within CSPI to ensure that each tenant's data is isolated and remains invisible to other tenants.

[0081] In a physical network, a network endpoint ("endpoint") refers to a computing device or system that is connected to a physical network and communicates back and forth with the network to which it is connected. A network endpoint in the physical network may be
25 connected to a Local Area Network (LAN), a Wide Area Network (WAN), or other type of physical network. Examples of traditional endpoints in a physical network include modems, hubs, bridges, switches, routers, and other networking devices, physical computers (or host machines), and the like. Each physical device in the physical network has a fixed network address that can be used to communicate with the device. This fixed network address can be
30 a Layer-2 address (e.g., a MAC address), a fixed Layer-3 address (e.g., an IP address), and the like. In a virtualized environment or in a virtual network, the endpoints can include various virtual endpoints such as virtual machines that are hosted by components of the

physical network (e.g., hosted by physical host machines). These endpoints in the virtual network are addressed by overlay addresses such as overlay Layer-2 addresses (e.g., overlay MAC addresses) and overlay Layer-3 addresses (e.g., overlay IP addresses). Network overlays enable flexibility by allowing network managers to move around the overlay addresses associated with network endpoints using software management (e.g., via software implementing a control plane for the virtual network). Accordingly, unlike in a physical network, in a virtual network, an overlay address (e.g., an overlay IP address) can be moved from one endpoint to another using network management software. Since the virtual network is built on top of a physical network, communications between components in the virtual network involves both the virtual network and the underlying physical network. In order to facilitate such communications, the components of CSPI are configured to learn and store mappings that map overlay addresses in the virtual network to actual physical addresses in the substrate network, and vice versa. These mappings are then used to facilitate the communications. Customer traffic is encapsulated to facilitate routing in the virtual network.

[0082] Accordingly, physical addresses (e.g., physical IP addresses) are associated with components in physical networks and overlay addresses (e.g., overlay IP addresses) are associated with entities in virtual or overlay networks. A physical IP address is an IP address associated with a physical device (e.g., a network device) in the substrate or physical network. For example, each NVD has an associated physical IP address. An overlay IP address is an overlay address associated with an entity in an overlay network, such as with a compute instance in a customer's virtual cloud network (VCN). Two different customers or tenants, each with their own private VCNs can potentially use the same overlay IP address in their VCNs without any knowledge of each other. Both the physical IP addresses and overlay IP addresses are types of real IP addresses. These are separate from virtual IP addresses. A virtual IP address is typically a single IP address that represents or maps to multiple real IP addresses. A virtual IP address provides a 1-to-many mapping between the virtual IP address and multiple real IP addresses. For example, a load balancer may use a VIP to map to or represent multiple servers, each server having its own real IP address.

[0083] The cloud infrastructure or CSPI is physically hosted in one or more data centers in one or more regions around the world. The CSPI may include components in the physical or substrate network and virtualized components (e.g., virtual networks, compute instances, virtual machines, etc.) that are in an virtual network built on top of the physical network components. In certain embodiments, the CSPI is organized and hosted in realms, regions

and availability domains. A region is typically a localized geographic area that contains one or more data centers. Regions are generally independent of each other and can be separated by vast distances, for example, across countries or even continents. For example, a first region may be in Australia, another one in Japan, yet another one in India, and the like. CSPI resources are divided among regions such that each region has its own independent subset of CSPI resources. Each region may provide a set of core infrastructure services and resources, such as, compute resources (e.g., bare metal servers, virtual machine, containers and related infrastructure, etc.); storage resources (e.g., block volume storage, file storage, object storage, archive storage); networking resources (e.g., virtual cloud networks (VCNs), load balancing resources, connections to on-premise networks), database resources; edge networking resources (e.g., DNS); and access management and monitoring resources, and others. Each region generally has multiple paths connecting it to other regions in the realm.

[0084] Generally, an application is deployed in a region (i.e., deployed on infrastructure associated with that region) where it is most heavily used, because using nearby resources is faster than using distant resources. Applications can also be deployed in different regions for various reasons, such as redundancy to mitigate the risk of region-wide events such as large weather systems or earthquakes, to meet varying requirements for legal jurisdictions, tax domains, and other business or social criteria, and the like.

[0085] The data centers within a region can be further organized and subdivided into availability domains (ADs). An availability domain may correspond to one or more data centers located within a region. A region can be composed of one or more availability domains. In such a distributed environment, CSPI resources are either region-specific, such as a virtual cloud network (VCN), or availability domain-specific, such as a compute instance.

[0086] ADs within a region are isolated from each other, fault tolerant, and are configured such that they are very unlikely to fail simultaneously. This is achieved by the ADs not sharing critical infrastructure resources such as networking, physical cables, cable paths, cable entry points, etc., such that a failure at one AD within a region is unlikely to impact the availability of the other ADs within the same region. The ADs within the same region may be connected to each other by a low latency, high bandwidth network, which makes it possible to provide high-availability connectivity to other networks (e.g., the Internet, customers' on-premise networks, etc.) and to build replicated systems in multiple ADs for

both high-availability and disaster recovery. Cloud services use multiple ADs to ensure high availability and to protect against resource failure. As the infrastructure provided by the IaaS provider grows, more regions and ADs may be added with additional capacity. Traffic between availability domains is usually encrypted.

5 **[0087]** In certain embodiments, regions are grouped into realms. A realm is a logical collection of regions. Realms are isolated from each other and do not share any data. Regions in the same realm may communicate with each other, but regions in different realms cannot. A customer's tenancy or account with the CSP exists in a single realm and can be spread across one or more regions that belong to that realm. Typically, when a customer
10 subscribes to an IaaS service, a tenancy or account is created for that customer in the customer-specified region (referred to as the "home" region) within a realm. A customer can extend the customer's tenancy across one or more other regions within the realm. A customer cannot access regions that are not in the realm where the customer's tenancy exists.

[0088] An IaaS provider can provide multiple realms, each realm catered to a particular set
15 of customers or users. For example, a commercial realm may be provided for commercial customers. As another example, a realm may be provided for a specific country for customers within that country. As yet another example, a government realm may be provided for a government, and the like. For example, the government realm may be catered for a specific government and may have a heightened level of security than a commercial realm.
20 For example, Oracle Cloud Infrastructure (OCI) currently offers a realm for commercial regions and two realms (e.g., FedRAMP authorized and IL5 authorized) for government cloud regions.

[0089] In certain embodiments, an AD can be subdivided into one or more fault domains. A fault domain is a grouping of infrastructure resources within an AD to provide anti-affinity.
25 Fault domains allow for the distribution of compute instances such that the instances are not on the same physical hardware within a single AD. This is known as anti-affinity. A fault domain refers to a set of hardware components (computers, switches, and more) that share a single point of failure. A compute pool is logically divided up into fault domains. Due to this, a hardware failure or compute hardware maintenance event that affects one fault domain
30 does not affect instances in other fault domains. Depending on the embodiment, the number of fault domains for each AD may vary. For instance, in certain embodiments each AD contains three fault domains. A fault domain acts as a logical data center within an AD.

[0090] When a customer subscribes to an IaaS service, resources from CSPI are provisioned for the customer and associated with the customer's tenancy. The customer can use these provisioned resources to build private networks and deploy resources on these networks. The customer networks that are hosted in the cloud by the CSPI are referred to as virtual cloud networks (VCNs). A customer can set up one or more virtual cloud networks (VCNs) using CSPI resources allocated for the customer. A VCN is a virtual or software defined private network. The customer resources that are deployed in the customer's VCN can include compute instances (e.g., virtual machines, bare-metal instances) and other resources. These compute instances may represent various customer workloads such as applications, load balancers, databases, and the like. A compute instance deployed on a VCN can communicate with public accessible endpoints ("public endpoints") over a public network such as the Internet, with other instances in the same VCN or other VCNs (e.g., the customer's other VCNs, or VCNs not belonging to the customer), with the customer's on-premise data centers or networks, and with service endpoints, and other types of endpoints.

[0091] The CSP may provide various services using the CSPI. In some instances, customers of CSPI may themselves act like service providers and provide services using CSPI resources. A service provider may expose a service endpoint, which is characterized by identification information (e.g., an IP Address, a DNS name and port). A customer's resource (e.g., a compute instance) can consume a particular service by accessing a service endpoint exposed by the service for that particular service. These service endpoints are generally endpoints that are publicly accessible by users using public IP addresses associated with the endpoints via a public communication network such as the Internet. Network endpoints that are publicly accessible are also sometimes referred to as public endpoints.

[0092] In certain embodiments, a service provider may expose a service via an endpoint (sometimes referred to as a service endpoint) for the service. Customers of the service can then use this service endpoint to access the service. In certain implementations, a service endpoint provided for a service can be accessed by multiple customers that intend to consume that service. In other implementations, a dedicated service endpoint may be provided for a customer such that only that customer can access the service using that dedicated service endpoint.

[0093] In certain embodiments, when a VCN is created, it is associated with a private overlay Classless Inter-Domain Routing (CIDR) address space, which is a range of private

overlay IP addresses that are assigned to the VCN (e.g., 10.0/16). A VCN includes associated subnets, route tables, and gateways. A VCN resides within a single region but can span one or more or all of the region's availability domains. A gateway is a virtual interface that is configured for a VCN and enables communication of traffic to and from the VCN to one or more endpoints outside the VCN. One or more different types of gateways may be configured for a VCN to enable communication to and from different types of endpoints.

[0094] A VCN can be subdivided into one or more sub-networks such as one or more subnets. A subnet is thus a unit of configuration or a subdivision that can be created within a VCN. A VCN can have one or multiple subnets. Each subnet within a VCN is associated with a contiguous range of overlay IP addresses (e.g., 10.0.0.0/24 and 10.0.1.0/24) that do not overlap with other subnets in that VCN and which represent an address space subset within the address space of the VCN.

[0095] Each compute instance is associated with a virtual network interface card (VNIC), that enables the compute instance to participate in a subnet of a VCN. A VNIC is a logical representation of physical Network Interface Card (NIC). In general, a VNIC is an interface between an entity (e.g., a compute instance, a service) and a virtual network. A VNIC exists in a subnet, has one or more associated IP addresses, and associated security rules or policies. A VNIC is equivalent to a Layer-2 port on a switch. A VNIC is attached to a compute instance and to a subnet within a VCN. A VNIC associated with a compute instance enables the compute instance to be a part of a subnet of a VCN and enables the compute instance to communicate (e.g., send and receive packets) with endpoints that are on the same subnet as the compute instance, with endpoints in different subnets in the VCN, or with endpoints outside the VCN. The VNIC associated with a compute instance thus determines how the compute instance connects with endpoints inside and outside the VCN. A VNIC for a compute instance is created and associated with that compute instance when the compute instance is created and added to a subnet within a VCN. For a subnet comprising a set of compute instances, the subnet contains the VNICs corresponding to the set of compute instances, each VNIC attached to a compute instance within the set of computer instances.

[0096] Each compute instance is assigned a private overlay IP address via the VNIC associated with the compute instance. This private overlay IP address is assigned to the VNIC that is associated with the compute instance when the compute instance is created and used for routing traffic to and from the compute instance. All VNICs in a given subnet use

the same route table, security lists, and DHCP options. As described above, each subnet within a VCN is associated with a contiguous range of overlay IP addresses (e.g., 10.0.0.0/24 and 10.0.1.0/24) that do not overlap with other subnets in that VCN and which represent an address space subset within the address space of the VCN. For a VNIC on a particular
5 subnet of a VCN, the private overlay IP address that is assigned to the VNIC is an address from the contiguous range of overlay IP addresses allocated for the subnet.

[0097] In certain embodiments, a compute instance may optionally be assigned additional overlay IP addresses in addition to the private overlay IP address, such as, for example, one or more public IP addresses if in a public subnet. These multiple addresses are assigned
10 either on the same VNIC or over multiple VNICs that are associated with the compute instance. Each instance however has a primary VNIC that is created during instance launch and is associated with the overlay private IP address assigned to the instance—this primary VNIC cannot be removed. Additional VNICs, referred to as secondary VNICs, can be added to an existing instance in the same availability domain as the primary VNIC. All the VNICs
15 are in the same availability domain as the instance. A secondary VNIC can be in a subnet in the same VCN as the primary VNIC, or in a different subnet that is either in the same VCN or a different one.

[0098] A compute instance may optionally be assigned a public IP address if it is in a public subnet. A subnet can be designated as either a public subnet or a private subnet at the
20 time the subnet is created. A private subnet means that the resources (e.g., compute instances) and associated VNICs in the subnet cannot have public overlay IP addresses. A public subnet means that the resources and associated VNICs in the subnet can have public IP addresses. A customer can designate a subnet to exist either in a single availability domain or across multiple availability domains in a region or realm.

[0099] As described above, a VCN may be subdivided into one or more subnets. In certain embodiments, a Virtual Router (VR) configured for the VCN (referred to as the VCN VR or just VR) enables communications between the subnets of the VCN. For a subnet within a VCN, the VR represents a logical gateway for that subnet that enables the subnet (i.e., the compute instances on that subnet) to communicate with endpoints on other subnets within the
30 VCN, and with other endpoints outside the VCN. The VCN VR is a logical entity that is configured to route traffic between VNICs in the VCN and virtual gateways ("gateways") associated with the VCN. Gateways are further described below with respect to FIG. 6. A

VCN VR is a Layer-3/IP Layer concept. In one embodiment, there is one VCN VR for a VCN where the VCN VR has potentially an unlimited number of ports addressed by IP addresses, with one port for each subnet of the VCN. In this manner, the VCN VR has a different IP address for each subnet in the VCN that the VCN VR is attached to. The VR is also connected to the various gateways configured for a VCN. In certain embodiments, a particular overlay IP address from the overlay IP address range for a subnet is reserved for a port of the VCN VR for that subnet. For example, consider a VCN having two subnets with associated address ranges 10.0/16 and 10.1/16, respectively. For the first subnet within the VCN with address range 10.0/16, an address from this range is reserved for a port of the VCN VR for that subnet. In some instances, the first IP address from the range may be reserved for the VCN VR. For example, for the subnet with overlay IP address range 10.0/16, IP address 10.0.0.1 may be reserved for a port of the VCN VR for that subnet. For the second subnet within the same VCN with address range 10.1/16, the VCN VR may have a port for that second subnet with IP address 10.1.0.1. The VCN VR has a different IP address for each of the subnets in the VCN.

[0100] In some other embodiments, each subnet within a VCN may have its own associated VR that is addressable by the subnet using a reserved or default IP address associated with the VR. The reserved or default IP address may, for example, be the first IP address from the range of IP addresses associated with that subnet. The VNICs in the subnet can communicate (e.g., send and receive packets) with the VR associated with the subnet using this default or reserved IP address. In such an embodiment, the VR is the ingress/egress point for that subnet. The VR associated with a subnet within the VCN can communicate with other VRs associated with other subnets within the VCN. The VRs can also communicate with gateways associated with the VCN. The VR function for a subnet is running on or executed by one or more NVDs executing VNICs functionality for VNICs in the subnet.

[0101] Route tables, security rules, and DHCP options may be configured for a VCN. Route tables are virtual route tables for the VCN and include rules to route traffic from subnets within the VCN to destinations outside the VCN by way of gateways or specially configured instances. A VCN's route tables can be customized to control how packets are forwarded/routed to and from the VCN. DHCP options refers to configuration information that is automatically provided to the instances when they boot up.

[0102] Security rules configured for a VCN represent overlay firewall rules for the VCN. The security rules can include ingress and egress rules, and specify the types of traffic (e.g., based upon protocol and port) that is allowed in and out of the instances within the VCN.

The customer can choose whether a given rule is stateful or stateless. For instance, the customer can allow incoming SSH traffic from anywhere to a set of instances by setting up a stateful ingress rule with source CIDR 0.0.0.0/0, and destination TCP port 22. Security rules can be implemented using network security groups or security lists. A network security group consists of a set of security rules that apply only to the resources in that group. A security list, on the other hand, includes rules that apply to all the resources in any subnet that uses the security list. A VCN may be provided with a default security list with default security rules. DHCP options configured for a VCN provide configuration information that is automatically provided to the instances in the VCN when the instances boot up.

[0103] In certain embodiments, the configuration information for a VCN is determined and stored by a VCN Control Plane. The configuration information for a VCN may include, for example, information about: the address range associated with the VCN, subnets within the VCN and associated information, one or more VRs associated with the VCN, compute instances in the VCN and associated VNICs, NVDs executing the various virtualization network functions (e.g., VNICs, VRs, gateways) associated with the VCN, state information for the VCN, and other VCN-related information. In certain embodiments, a VCN Distribution Service publishes the configuration information stored by the VCN Control Plane, or portions thereof, to the NVDs. The distributed information may be used to update information (e.g., forwarding tables, routing tables, etc.) stored and used by the NVDs to forward packets to and from the compute instances in the VCN.

[0104] In certain embodiments, the creation of VCNs and subnets are handled by a VCN Control Plane (CP) and the launching of compute instances is handled by a Compute Control Plane. The Compute Control Plane is responsible for allocating the physical resources for the compute instance and then calls the VCN Control Plane to create and attach VNICs to the compute instance. The VCN CP also sends VCN data mappings to the VCN data plane that is configured to perform packet forwarding and routing functions. In certain embodiments, the VCN CP provides a distribution service that is responsible for providing updates to the VCN data plane. Examples of a VCN Control Plane are also depicted in FIGs. 11, 12, 13, and 14 (see references 1116, 1216, 1316, and 1416) and described below.

[0105] A customer may create one or more VCNs using resources hosted by CSPI. A compute instance deployed on a customer VCN may communicate with different endpoints. These endpoints can include endpoints that are hosted by CSPI and endpoints outside CSPI.

[0106] Various different architectures for implementing cloud-based service using CSPI are depicted in FIGs. 6, 7, 8, 9, 10, 11, 12, 13, and 15, and are described below. FIG. 6 is a high level diagram of a distributed environment 600 showing an overlay or customer VCN hosted by CSPI according to certain embodiments. The distributed environment depicted in FIG. 6 includes multiple components in the overlay network. Distributed environment 600 depicted in FIG. 6 is merely an example and is not intended to unduly limit the scope of claimed embodiments. Many variations, alternatives, and modifications are possible. For example, in some implementations, the distributed environment depicted in FIG. 6 may have more or fewer systems or components than those shown in FIG. 1, may combine two or more systems, or may have a different configuration or arrangement of systems.

[0107] As shown in the example depicted in FIG. 6, distributed environment 600 comprises CSPI 601 that provides services and resources that customers can subscribe to and use to build their virtual cloud networks (VCNs). In certain embodiments, CSPI 601 offers IaaS services to subscribing customers. The data centers within CSPI 601 may be organized into one or more regions. One example region "Region US" 602 is shown in FIG. 6. A customer has configured a customer VCN 604 for region 602. The customer may deploy various compute instances on VCN 604, where the compute instances may include virtual machines or bare metal instances. Examples of instances include applications, database, load balancers, and the like.

[0108] In the embodiment depicted in FIG. 6, customer VCN 604 comprises two subnets, namely, "Subnet-1" and "Subnet-2", each subnet with its own CIDR IP address range. In FIG. 6, the overlay IP address range for Subnet-1 is 10.0/16 and the address range for Subnet-2 is 10.1/16. A VCN Virtual Router 605 represents a logical gateway for the VCN that enables communications between subnets of the VCN 604, and with other endpoints outside the VCN. VCN VR 605 is configured to route traffic between VNICs in VCN 604 and gateways associated with VCN 604. VCN VR 605 provides a port for each subnet of VCN 604. For example, VR 605 may provide a port with IP address 10.0.0.1 for Subnet-1 and a port with IP address 10.1.0.1 for Subnet-2.

[0109] Multiple compute instances may be deployed on each subnet, where the compute instances can be virtual machine instances, and/or bare metal instances. The compute instances in a subnet may be hosted by one or more host machines within CSPI 601. A compute instance participates in a subnet via a VNIC associated with the compute instance.

5 For example, as shown in FIG. 6, a compute instance C1 is part of Subnet-1 via a VNIC associated with the compute instance. Likewise, compute instance C2 is part of Subnet-1 via a VNIC associated with C2. In a similar manner, multiple compute instances, which may be virtual machine instances or bare metal instances, may be part of Subnet-1. Via its associated VNIC, each compute instance is assigned a private overlay IP address and a MAC address.

10 For example, in FIG. 6, compute instance C1 has an overlay IP address of 10.0.0.2 and a MAC address of M1, while compute instance C2 has an private overlay IP address of 10.0.0.3 and a MAC address of M2. Each compute instance in Subnet-1, including compute instances C1 and C2, has a default route to VCN VR 605 using IP address 10.0.0.1, which is the IP address for a port of VCN VR 605 for Subnet-1.

15 [0110] Subnet-2 can have multiple compute instances deployed on it, including virtual machine instances and/or bare metal instances. For example, as shown in FIG. 6, compute instances D1 and D2 are part of Subnet-2 via VNICs associated with the respective compute instances. In the embodiment depicted in FIG. 6, compute instance D1 has an overlay IP address of 10.1.0.2 and a MAC address of MM1, while compute instance D2 has an private
20 overlay IP address of 10.1.0.3 and a MAC address of MM2. Each compute instance in Subnet-2, including compute instances D1 and D2, has a default route to VCN VR 605 using IP address 10.1.0.1, which is the IP address for a port of VCN VR 605 for Subnet-2.

[0111] VCN A 604 may also include one or more load balancers. For example, a load balancer may be provided for a subnet and may be configured to load balance traffic across
25 multiple compute instances on the subnet. A load balancer may also be provided to load balance traffic across subnets in the VCN.

[0112] A particular compute instance deployed on VCN 604 can communicate with various different endpoints. These endpoints may include endpoints that are hosted by CSPI 700 and endpoints outside CSPI 700. Endpoints that are hosted by CSPI 601 may include: an
30 endpoint on the same subnet as the particular compute instance (e.g., communications between two compute instances in Subnet-1); an endpoint on a different subnet but within the same VCN (e.g., communication between a compute instance in Subnet-1 and a compute

instance in Subnet-2); an endpoint in a different VCN in the same region (e.g., communications between a compute instance in Subnet-1 and an endpoint in a VCN in the same region 606 or 610, communications between a compute instance in Subnet-1 and an endpoint in service network 610 in the same region); or an endpoint in a VCN in a different region (e.g., communications between a compute instance in Subnet-1 and an endpoint in a VCN in a different region 608). A compute instance in a subnet hosted by CSPI 601 may also communicate with endpoints that are not hosted by CSPI 601 (i.e., are outside CSPI 601). These outside endpoints include endpoints in the customer's on-premise network 616, endpoints within other remote cloud hosted networks 618, public endpoints 614 accessible via a public network such as the Internet, and other endpoints.

[0113] Communications between compute instances on the same subnet are facilitated using VNICs associated with the source compute instance and the destination compute instance. For example, compute instance C1 in Subnet-1 may want to send packets to compute instance C2 in Subnet-1. For a packet originating at a source compute instance and whose destination is another compute instance in the same subnet, the packet is first processed by the VNIC associated with the source compute instance. Processing performed by the VNIC associated with the source compute instance can include determining destination information for the packet from the packet headers, identifying any policies (e.g., security lists) configured for the VNIC associated with the source compute instance, determining a next hop for the packet, performing any packet encapsulation/decapsulation functions as needed, and then forwarding/routing the packet to the next hop with the goal of facilitating communication of the packet to its intended destination. When the destination compute instance is in the same subnet as the source compute instance, the VNIC associated with the source compute instance is configured to identify the VNIC associated with the destination compute instance and forward the packet to that VNIC for processing. The VNIC associated with the destination compute instance is then executed and forwards the packet to the destination compute instance.

[0114] For a packet to be communicated from a compute instance in a subnet to an endpoint in a different subnet in the same VCN, the communication is facilitated by the VNICs associated with the source and destination compute instances and the VCN VR. For example, if compute instance C1 in Subnet-1 in FIG. 6 wants to send a packet to compute instance D1 in Subnet-2, the packet is first processed by the VNIC associated with compute instance C1. The VNIC associated with compute instance C1 is configured to route the

packet to the VCN VR 605 using default route or port 10.0.0.1 of the VCN VR. VCN VR 605 is configured to route the packet to Subnet-2 using port 10.1.0.1. The packet is then received and processed by the VNIC associated with D1 and the VNIC forwards the packet to compute instance D1.

5 [0115] For a packet to be communicated from a compute instance in VCN 604 to an endpoint that is outside VCN 604, the communication is facilitated by the VNIC associated with the source compute instance, VCN VR 605, and gateways associated with VCN 604. One or more types of gateways may be associated with VCN 604. A gateway is an interface between a VCN and another endpoint, where the another endpoint is outside the VCN. A
10 gateway is a Layer-3/IP layer concept and enables a VCN to communicate with endpoints outside the VCN. A gateway thus facilitates traffic flow between a VCN and other VCNs or networks. Various different types of gateways may be configured for a VCN to facilitate different types of communications with different types of endpoints. Depending upon the gateway, the communications may be over public networks (e.g., the Internet) or over private
15 networks. Various communication protocols may be used for these communications.

[0116] For example, compute instance C1 may want to communicate with an endpoint outside VCN 604. The packet may be first processed by the VNIC associated with source compute instance C1. The VNIC processing determines that the destination for the packet is outside the Subnet-1 of C1. The VNIC associated with C1 may forward the packet to VCN
20 VR 605 for VCN 604. VCN VR 605 then processes the packet and as part of the processing, based upon the destination for the packet, determines a particular gateway associated with VCN 604 as the next hop for the packet. VCN VR 605 may then forward the packet to the particular identified gateway. For example, if the destination is an endpoint within the customer's on-premise network, then the packet may be forwarded by VCN VR 605 to
25 Dynamic Routing Gateway (DRG) gateway 622 configured for VCN 604. The packet may then be forwarded from the gateway to a next hop to facilitate communication of the packet to its final intended destination.

[0117] Various different types of gateways may be configured for a VCN. Examples of gateways that may be configured for a VCN are depicted in FIG. 6 and described below.

30 Examples of gateways associated with a VCN are also depicted in FIGS. 11, 12, 13, and 14 (for example, gateways referenced by reference numbers 1134, 1136, 1138, 1234, 1236, 1238, 1334, 1336, 1338, 1434, 1436, and 1438) and described below. As shown in the

embodiment depicted in FIG. 6, a Dynamic Routing Gateway (DRG) 622 may be added to or be associated with customer VCN 604 and provides a path for private network traffic communication between customer VCN 604 and another endpoint, where the another endpoint can be the customer's on-premise network 616, a VCN 608 in a different region of CSPI 601, or other remote cloud networks 618 not hosted by CSPI 601. Customer on-premise network 616 may be a customer network or a customer data center built using the customer's resources. Access to customer on-premise network 616 is generally very restricted. For a customer that has both a customer on-premise network 616 and one or more VCNs 604 deployed or hosted in the cloud by CSPI 601, the customer may want their on-premise network 616 and their cloud-based VCN 604 to be able to communicate with each other. This enables a customer to build an extended hybrid environment encompassing the customer's VCN 604 hosted by CSPI 601 and their on-premises network 616. DRG 622 enables this communication. To enable such communications, a communication channel 624 is set up where one endpoint of the channel is in customer on-premise network 616 and the other endpoint is in CSPI 601 and connected to customer VCN 604. Communication channel 624 can be over public communication networks such as the Internet or private communication networks. Various different communication protocols may be used such as IPsec VPN technology over a public communication network such as the Internet, Oracle's FastConnect technology that uses a private network instead of a public network, and others. The device or equipment in customer on-premise network 616 that forms one end point for communication channel 624 is referred to as the customer premise equipment (CPE), such as CPE 626 depicted in FIG. 6. On the CSPI 601 side, the endpoint may be a host machine executing DRG 622.

[0118] In certain embodiments, a Remote Peering Connection (RPC) can be added to a DRG, which allows a customer to peer one VCN with another VCN in a different region. Using such an RPC, customer VCN 604 can use DRG 622 to connect with a VCN 608 in another region. DRG 622 may also be used to communicate with other remote cloud networks 618, not hosted by CSPI 601 such as a Microsoft Azure cloud, Amazon AWS cloud, and others.

[0119] As shown in FIG. 6, an Internet Gateway (IGW) 620 may be configured for customer VCN 604 the enables a compute instance on VCN 604 to communicate with public endpoints 614 accessible over a public network such as the Internet. IGW 620 is a gateway that connects a VCN to a public network such as the Internet. IGW 620 enables a public

subnet (where the resources in the public subnet have public overlay IP addresses) within a VCN, such as VCN 604, direct access to public endpoints 612 on a public network 614 such as the Internet. Using IGW 620, connections can be initiated from a subnet within VCN 604 or from the Internet.

- 5 **[0120]** A Network Address Translation (NAT) gateway 628 can be configured for customer's VCN 604 and enables cloud resources in the customer's VCN, which do not have dedicated public overlay IP addresses, access to the Internet and it does so without exposing those resources to direct incoming Internet connections (e.g., L4-L7 connections). This enables a private subnet within a VCN, such as private Subnet-1 in VCN 604, with private
10 access to public endpoints on the Internet. In NAT gateways, connections can be initiated only from the private subnet to the public Internet and not from the Internet to the private subnet.

- [0121]** In certain embodiments, a Service Gateway (SGW) 626 can be configured for customer VCN 604 and provides a path for private network traffic between VCN 604 and
15 supported services endpoints in a service network 610. In certain embodiments, service network 610 may be provided by the CSP and may provide various services. An example of such a service network is Oracle's Services Network, which provides various services that can be used by customers. For example, a compute instance (e.g., a database system) in a private subnet of customer VCN 604 can back up data to a service endpoint (e.g., Object
20 Storage) without needing public IP addresses or access to the Internet. In certain embodiments, a VCN can have only one SGW, and connections can only be initiated from a subnet within the VCN and not from service network 610. If a VCN is peered with another, resources in the other VCN typically cannot access the SGW. Resources in on-premises networks that are connected to a VCN with FastConnect or VPN Connect can also use the
25 service gateway configured for that VCN.

- [0122]** In certain implementations, SGW 626 uses the concept of a service Classless Inter-Domain Routing (CIDR) label, which is a string that represents all the regional public IP address ranges for the service or group of services of interest. The customer uses the service CIDR label when they configure the SGW and related route rules to control traffic to the
30 service. The customer can optionally utilize it when configuring security rules without needing to adjust them if the service's public IP addresses change in the future.

[0123] A Local Peering Gateway (LPG) 632 is a gateway that can be added to customer VCN 604 and enables VCN 604 to peer with another VCN in the same region. Peering means that the VCNs communicate using private IP addresses, without the traffic traversing a public network such as the Internet or without routing the traffic through the customer's on-premises network 616. In preferred embodiments, a VCN has a separate LPG for each peering it establishes. Local Peering or VCN Peering is a common practice used to establish network connectivity between different applications or infrastructure management functions.

[0124] Service providers, such as providers of services in service network 610, may provide access to services using different access models. According to a public access model, services may be exposed as public endpoints that are publicly accessible by compute instance in a customer VCN via a public network such as the Internet and or may be privately accessible via SGW 626. According to a specific private access model, services are made accessible as private IP endpoints in a private subnet in the customer's VCN. This is referred to as a Private Endpoint (PE) access and enables a service provider to expose their service as an instance in the customer's private network. A Private Endpoint resource represents a service within the customer's VCN. Each PE manifests as a VNIC (referred to as a PE-VNIC, with one or more private IPs) in a subnet chosen by the customer in the customer's VCN. A PE thus provides a way to present a service within a private customer VCN subnet using a VNIC. Since the endpoint is exposed as a VNIC, all the features associates with a VNIC such as routing rules, security lists, etc., are now available for the PE VNIC.

[0125] A service provider can register their service to enable access through a PE. The provider can associate policies with the service that restricts the service's visibility to the customer tenancies. A provider can register multiple services under a single virtual IP address (VIP), especially for multi-tenant services. There may be multiple such private endpoints (in multiple VCNs) that represent the same service.

[0126] Compute instances in the private subnet can then use the PE VNIC's private IP address or the service DNS name to access the service. Compute instances in the customer VCN can access the service by sending traffic to the private IP address of the PE in the customer VCN. A Private Access Gateway (PAGW) 630 is a gateway resource that can be attached to a service provider VCN (e.g., a VCN in service network 610) that acts as an ingress/egress point for all traffic from/to customer subnet private endpoints. PAGW 630 enables a provider to scale the number of PE connections without utilizing its internal IP

address resources. A provider needs only configure one PAGW for any number of services registered in a single VCN. Providers can represent a service as a private endpoint in multiple VCNs of one or more customers. From the customer's perspective, the PE VNIC, which, instead of being attached to a customer's instance, appears attached to the service with which the customer wishes to interact. The traffic destined to the private endpoint is routed via PAGW 630 to the service. These are referred to as customer-to-service private connections (C2S connections).

[0127] The PE concept can also be used to extend the private access for the service to customer's on-premises networks and data centers, by allowing the traffic to flow through FastConnect/IPsec links and the private endpoint in the customer VCN. Private access for the service can also be extended to the customer's peered VCNs, by allowing the traffic to flow between LPG 632 and the PE in the customer's VCN.

[0128] A customer can control routing in a VCN at the subnet level, so the customer can specify which subnets in the customer's VCN, such as VCN 604, use each gateway. A VCN's route tables are used to decide if traffic is allowed out of a VCN through a particular gateway. For example, in a particular instance, a route table for a public subnet within customer VCN 604 may send non-local traffic through IGW 620. The route table for a private subnet within the same customer VCN 604 may send traffic destined for CSP services through SGW 626. All remaining traffic may be sent via the NAT gateway 628. Route tables only control traffic going out of a VCN.

[0129] Security lists associated with a VCN are used to control traffic that comes into a VCN via a gateway via inbound connections. All resources in a subnet use the same route table and security lists. Security lists may be used to control specific types of traffic allowed in and out of instances in a subnet of a VCN. Security list rules may comprise ingress (inbound) and egress (outbound) rules. For example, an ingress rule may specify an allowed source address range, while an egress rule may specify an allowed destination address range. Security rules may specify a particular protocol (e.g., TCP, ICMP), a particular port (e.g., 22 for SSH, 3389 for Windows RDP), etc. In certain implementations, an instance's operating system may enforce its own firewall rules that are aligned with the security list rules. Rules may be stateful (e.g., a connection is tracked and the response is automatically allowed without an explicit security list rule for the response traffic) or stateless.

[0130] Access from a customer VCN (i.e., by a resource or compute instance deployed on VCN 604) can be categorized as public access, private access, or dedicated access. Public access refers to an access model where a public IP address or a NAT is used to access a public endpoint. Private access enables customer workloads in VCN 604 with private IP addresses (e.g., resources in a private subnet) to access services without traversing a public network such as the Internet. In certain embodiments, CSPI 601 enables customer VCN workloads with private IP addresses to access the (public service endpoints of) services using a service gateway. A service gateway thus offers a private access model by establishing a virtual link between the customer's VCN and the service's public endpoint residing outside the customer's private network.

[0131] Additionally, CSPI may offer dedicated public access using technologies such as FastConnect public peering where customer on-premises instances can access one or more services in a customer VCN using a FastConnect connection and without traversing a public network such as the Internet. CSPI also may also offer dedicated private access using FastConnect private peering where customer on-premises instances with private IP addresses can access the customer's VCN workloads using a FastConnect connection. FastConnect is a network connectivity alternative to using the public Internet to connect a customer's on-premise network to CSPI and its services. FastConnect provides an easy, elastic, and economical way to create a dedicated and private connection with higher bandwidth options and a more reliable and consistent networking experience when compared to Internet-based connections.

[0132] FIG. 6 and the accompanying description above describes various virtualized components in an example virtual network. As described above, the virtual network is built on the underlying physical or substrate network. FIG. 7 depicts a simplified architectural diagram of the physical components in the physical network within CSPI 700 that provide the underlay for the virtual network according to certain embodiments. As shown, CSPI 700 provides a distributed environment comprising components and resources (e.g., compute, memory, and networking resources) provided by a cloud service provider (CSP). These components and resources are used to provide cloud services (e.g., IaaS services) to subscribing customers, i.e., customers that have subscribed to one or more services provided by the CSP. Based upon the services subscribed to by a customer, a subset of resources (e.g., compute, memory, and networking resources) of CSPI 700 are provisioned for the customer. Customers can then build their own cloud-based (i.e., CSPI-hosted) customizable and private

virtual networks using physical compute, memory, and networking resources provided by CSPI 700. As previously indicated, these customer networks are referred to as virtual cloud networks (VCNs). A customer can deploy one or more customer resources, such as compute instances, on these customer VCNs. Compute instances can be in the form of virtual machines, bare metal instances, and the like. CSPI 700 provides infrastructure and a set of complementary cloud services that enable customers to build and run a wide range of applications and services in a highly available hosted environment.

[0133] In the example embodiment depicted in FIG. 7, the physical components of CSPI 700 include one or more physical host machines or physical servers (e.g., 702, 706, 708), network virtualization devices (NVDs) (e.g., 710, 712), top-of-rack (TOR) switches (e.g., 714, 716), and a physical network (e.g., 718), and switches in physical network 718. The physical host machines or servers may host and execute various compute instances that participate in one or more subnets of a VCN. The compute instances may include virtual machine instances, and bare metal instances. For example, the various compute instances depicted in FIG. 6 may be hosted by the physical host machines depicted in FIG. 7. The virtual machine compute instances in a VCN may be executed by one host machine or by multiple different host machines. The physical host machines may also host virtual host machines, container-based hosts or functions, and the like. The VNICs and VCN VR depicted in FIG. 6 may be executed by the NVDs depicted in FIG. 7. The gateways depicted in FIG. 6 may be executed by the host machines and/or by the NVDs depicted in FIG. 7.

[0134] The host machines or servers may execute a hypervisor (also referred to as a virtual machine monitor or VMM) that creates and enables a virtualized environment on the host machines. The virtualization or virtualized environment facilitates cloud-based computing. One or more compute instances may be created, executed, and managed on a host machine by a hypervisor on that host machine. The hypervisor on a host machine enables the physical computing resources of the host machine (e.g., compute, memory, and networking resources) to be shared between the various compute instances executed by the host machine.

[0135] For example, as depicted in FIG. 7, host machines 702 and 708 execute hypervisors 760 and 766, respectively. These hypervisors may be implemented using software, firmware, or hardware, or combinations thereof. Typically, a hypervisor is a process or a software layer that sits on top of the host machine's operating system (OS), which in turn executes on the hardware processors of the host machine. The hypervisor provides a virtualized environment

by enabling the physical computing resources (e.g., processing resources such as processors/cores, memory resources, networking resources) of the host machine to be shared among the various virtual machine compute instances executed by the host machine. For example, in FIG. 7, hypervisor 760 may sit on top of the OS of host machine 702 and enables the computing resources (e.g., processing, memory, and networking resources) of host machine 702 to be shared between compute instances (e.g., virtual machines) executed by host machine 702. A virtual machine can have its own operating system (referred to as a guest operating system), which may be the same as or different from the OS of the host machine. The operating system of a virtual machine executed by a host machine may be the same as or different from the operating system of another virtual machine executed by the same host machine. A hypervisor thus enables multiple operating systems to be executed alongside each other while sharing the same computing resources of the host machine. The host machines depicted in FIG. 7 may have the same or different types of hypervisors.

[0136] A compute instance can be a virtual machine instance or a bare metal instance. In FIG. 7, compute instances 768 on host machine 702 and 774 on host machine 708 are examples of virtual machine instances. Host machine 706 is an example of a bare metal instance that is provided to a customer.

[0137] In certain instances, an entire host machine may be provisioned to a single customer, and all of the one or more compute instances (either virtual machines or bare metal instance) hosted by that host machine belong to that same customer. In other instances, a host machine may be shared between multiple customers (i.e., multiple tenants). In such a multi-tenancy scenario, a host machine may host virtual machine compute instances belonging to different customers. These compute instances may be members of different VCNs of different customers. In certain embodiments, a bare metal compute instance is hosted by a bare metal server without a hypervisor. When a bare metal compute instance is provisioned, a single customer or tenant maintains control of the physical CPU, memory, and network interfaces of the host machine hosting the bare metal instance and the host machine is not shared with other customers or tenants.

[0138] As previously described, each compute instance that is part of a VCN is associated with a VNIC that enables the compute instance to become a member of a subnet of the VCN. The VNIC associated with a compute instance facilitates the communication of packets or frames to and from the compute instance. A VNIC is associated with a compute instance

when the compute instance is created. In certain embodiments, for a compute instance executed by a host machine, the VNIC associated with that compute instance is executed by an NVD connected to the host machine. For example, in FIG. 7, host machine 702 executes a virtual machine compute instance 768 that is associated with VNIC 776, and VNIC 776 is executed by NVD 710 connected to host machine 702. As another example, bare metal instance 772 hosted by host machine 706 is associated with VNIC 780 that is executed by NVD 712 connected to host machine 706. As yet another example, VNIC 784 is associated with compute instance 774 executed by host machine 708, and VNIC 784 is executed by NVD 712 connected to host machine 708.

[0139] For compute instances hosted by a host machine, an NVD connected to that host machine also executes VCN VRs corresponding to VCNs of which the compute instances are members. For example, in the embodiment depicted in FIG. 7, NVD 710 executes VCN VR 777 corresponding to the VCN of which compute instance 768 is a member. NVD 712 may also execute one or more VCN VRs 783 corresponding to VCNs corresponding to the compute instances hosted by host machines 706 and 708.

[0140] A host machine may include one or more network interface cards (NIC) that enable the host machine to be connected to other devices. A NIC on a host machine may provide one or more ports (or interfaces) that enable the host machine to be communicatively connected to another device. For example, a host machine may be connected to an NVD using one or more ports (or interfaces) provided on the host machine and on the NVD. A host machine may also be connected to other devices such as another host machine.

[0141] For example, in FIG. 7, host machine 702 is connected to NVD 710 using link 720 that extends between a port 734 provided by a NIC 732 of host machine 702 and between a port 736 of NVD 710. Host machine 706 is connected to NVD 712 using link 724 that extends between a port 746 provided by a NIC 744 of host machine 706 and between a port 748 of NVD 712. Host machine 708 is connected to NVD 712 using link 726 that extends between a port 752 provided by a NIC 750 of host machine 708 and between a port 754 of NVD 712.

[0142] The NVDs are in turn connected via communication links to top-of-the-rack (TOR) switches, which are connected to physical network 718 (also referred to as the switch fabric). In certain embodiments, the links between a host machine and an NVD, and between an NVD and a TOR switch are Ethernet links. For example, in FIG. 7, NVDs 710 and 712 are

connected to 1126 TOR switches 714 and 716, respectively, using links 728 and 730. In certain embodiments, the links 720, 724, 726, 728, and 730 are Ethernet links. The collection of host machines and NVDs that are connected to a TOR is sometimes referred to as a rack.

[0143] Physical network 718 provides a communication fabric that enables TOR switches to communicate with each other. Physical network 718 can be a multi-tiered network. In certain implementations, physical network 718 is a multi-tiered Clos network of switches, with TOR switches 714 and 716 representing the leaf level nodes of the multi-tiered and multi-node physical switching network 718. Different Clos network configurations are possible including but not limited to a 2-tier network, a 3-tier network, a 4-tier network, a 5-tier network, and in general a "n"-tiered network. An example of a Clos network is depicted in FIG. 10 and described below.

[0144] Various different connection configurations are possible between host machines and NVDs such as one-to-one configuration, many-to-one configuration, one-to-many configuration, and others. In a one-to-one configuration implementation, each host machine is connected to its own separate NVD. For example, in FIG. 7, host machine 702 is connected to NVD 710 via NIC 732 of host machine 702. In a many-to-one configuration, multiple host machines are connected to one NVD. For example, in FIG. 7, host machines 706 and 708 are connected to the same NVD 712 via NICs 744 and 750, respectively.

[0145] In a one-to-many configuration, one host machine is connected to multiple NVDs. FIG. 8 shows an example within CSPI 800 where a host machine is connected to multiple NVDs. As shown in FIG. 8, host machine 802 comprises a network interface card (NIC) 804 that includes multiple ports 806 and 808. Host machine 800 is connected to a first NVD 810 via port 806 and link 820, and connected to a second NVD 812 via port 808 and link 822. Ports 806 and 808 may be Ethernet ports and the links 820 and 822 between host machine 802 and NVDs 810 and 812 may be Ethernet links. NVD 810 is in turn connected to a first TOR switch 814 and NVD 812 is connected to a second TOR switch 816. The links between NVDs 810 and 812, and TOR switches 814 and 816 may be Ethernet links. TOR switches 814 and 816 represent the Tier-0 switching devices in multi-tiered physical network 818.

[0146] The arrangement depicted in FIG. 8 provides two separate physical network paths to and from physical switch network 818 to host machine 802: a first path traversing TOR switch 814 to NVD 810 to host machine 802, and a second path traversing TOR switch 816 to NVD 812 to host machine 802. The separate paths provide for enhanced availability

(referred to as high availability) of host machine 802. If there are problems in one of the paths (e.g., a link in one of the paths goes down) or devices (e.g., a particular NVD is not functioning), then the other path may be used for communications to/from host machine 802.

[0147] In the configuration depicted in FIG. 8, the host machine is connected to two different NVDs using two different ports provided by a NIC of the host machine. In other embodiments, a host machine may include multiple NICs that enable connectivity of the host machine to multiple NVDs.

[0148] Referring back to FIG. 7, an NVD is a physical device or component that performs one or more network and/or storage virtualization functions. An NVD may be any device with one or more processing units (e.g., CPUs, Network Processing Units (NPU), FPGAs, packet processing pipelines, etc.), memory including cache, and ports. The various virtualization functions may be performed by software/firmware executed by the one or more processing units of the NVD.

[0149] An NVD may be implemented in various different forms. For example, in certain embodiments, an NVD is implemented as an interface card referred to as a smartNIC or an intelligent NIC with an embedded processor onboard. A smartNIC is a separate device from the NICs on the host machines. In FIG. 7, the NVDs 710 and 712 may be implemented as smartNICs that are connected to host machines 702, and host machines 706 and 708, respectively.

[0150] A smartNIC is however just one example of an NVD implementation. Various other implementations are possible. For example, in some other implementations, an NVD or one or more functions performed by the NVD may be incorporated into or performed by one or more host machines, one or more TOR switches, and other components of CSPI 700. For example, an NVD may be embodied in a host machine where the functions performed by an NVD are performed by the host machine. As another example, an NVD may be part of a TOR switch or a TOR switch may be configured to perform functions performed by an NVD that enables the TOR switch to perform various complex packet transformations that are used for a public cloud. A TOR that performs the functions of an NVD is sometimes referred to as a smart TOR. In yet other implementations, where virtual machines (VMs) instances, but not bare metal (BM) instances, are offered to customers, functions performed by an NVD may be implemented inside a hypervisor of the host machine. In some other implementations, some

of the functions of the NVD may be offloaded to a centralized service running on a fleet of host machines.

[0151] In certain embodiments, such as when implemented as a smartNIC as shown in FIG. 7, an NVD may comprise multiple physical ports that enable it to be connected to one or more host machines and to one or more TOR switches. A port on an NVD can be classified as a host-facing port (also referred to as a "south port") or a network-facing or TOR-facing port (also referred to as a "north port"). A host-facing port of an NVD is a port that is used to connect the NVD to a host machine. Examples of host-facing ports in FIG. 7 include port 736 on NVD 710, and ports 748 and 754 on NVD 712. A network-facing port of an NVD is a port that is used to connect the NVD to a TOR switch. Examples of network-facing ports in FIG. 7 include port 756 on NVD 710, and port 758 on NVD 712. As shown in FIG. 7, NVD 710 is connected to TOR switch 714 using link 728 that extends from port 756 of NVD 710 to the TOR switch 714. Likewise, NVD 712 is connected to TOR switch 716 using link 730 that extends from port 758 of NVD 712 to the TOR switch 716.

[0152] An NVD receives packets and frames from a host machine (e.g., packets and frames generated by a compute instance hosted by the host machine) via a host-facing port and, after performing the necessary packet processing, may forward the packets and frames to a TOR switch via a network-facing port of the NVD. An NVD may receive packets and frames from a TOR switch via a network-facing port of the NVD and, after performing the necessary packet processing, may forward the packets and frames to a host machine via a host-facing port of the NVD.

[0153] In certain embodiments, there may be multiple ports and associated links between an NVD and a TOR switch. These ports and links may be aggregated to form a link aggregator group of multiple ports or links (referred to as a LAG). Link aggregation allows multiple physical links between two end-points (e.g., between an NVD and a TOR switch) to be treated as a single logical link. All the physical links in a given LAG may operate in full-duplex mode at the same speed. LAGs help increase the bandwidth and reliability of the connection between two endpoints. If one of the physical links in the LAG goes down, traffic is dynamically and transparently reassigned to one of the other physical links in the LAG. The aggregated physical links deliver higher bandwidth than each individual link. The multiple ports associated with a LAG are treated as a single logical port. Traffic can be load-balanced across the multiple physical links of a LAG. One or more LAGs may be configured

between two endpoints. The two endpoints may be between an NVD and a TOR switch, between a host machine and an NVD, and the like.

[0154] An NVD implements or performs network virtualization functions. These functions are performed by software/firmware executed by the NVD. Examples of network virtualization functions include without limitation: packet encapsulation and de-capsulation functions; functions for creating a VCN network; functions for implementing network policies such as VCN security list (firewall) functionality; functions that facilitate the routing and forwarding of packets to and from compute instances in a VCN; and the like. In certain embodiments, upon receiving a packet, an NVD is configured to execute a packet processing pipeline for processing the packet and determining how the packet is to be forwarded or routed. As part of this packet processing pipeline, the NVD may execute one or more virtual functions associated with the overlay network such as executing VNICs associated with compute instances in the VCN, executing a Virtual Router (VR) associated with the VCN, the encapsulation and decapsulation of packets to facilitate forwarding or routing in the virtual network, execution of certain gateways (e.g., the Local Peering Gateway), the implementation of Security Lists, Network Security Groups, network address translation (NAT) functionality (e.g., the translation of Public IP to Private IP on a host by host basis), throttling functions, and other functions.

[0155] In certain embodiments, the packet processing data path in an NVD may comprise multiple packet pipelines, each composed of a series of packet transformation stages. In certain implementations, upon receiving a packet, the packet is parsed and classified to a single pipeline. The packet is then processed in a linear fashion, one stage after another, until the packet is either dropped or sent out over an interface of the NVD. These stages provide basic functional packet processing building blocks (e.g., validating headers, enforcing throttle, inserting new Layer-2 headers, enforcing L4 firewall, VCN encapsulation/decapsulation, etc.) so that new pipelines can be constructed by composing existing stages, and new functionality can be added by creating new stages and inserting them into existing pipelines.

[0156] An NVD may perform both control plane and data plane functions corresponding to a control plane and a data plane of a VCN. Examples of a VCN Control Plane are also depicted in FIGs. 11, 12, 13, and 14 (see references 1116, 1216, 1316, and 1416) and described below. Examples of a VCN Data Plane are depicted in FIGs. 11, 12, 13, and 14

(see references 1118, 1218, 1318, and 1418) and described below. The control plane functions include functions used for configuring a network (e.g., setting up routes and route tables, configuring VNICs, etc.) that controls how data is to be forwarded. In certain embodiments, a VCN Control Plane is provided that computes all the overlay-to-substrate mappings centrally and publishes them to the NVDs and to the virtual network edge devices such as various gateways such as the DRG, the SGW, the IGW, etc. Firewall rules may also be published using the same mechanism. In certain embodiments, an NVD only gets the mappings that are relevant for that NVD. The data plane functions include functions for the actual routing/forwarding of a packet based upon configuration set up using control plane. A VCN data plane is implemented by encapsulating the customer's network packets before they traverse the substrate network. The encapsulation/decapsulation functionality is implemented on the NVDs. In certain embodiments, an NVD is configured to intercept all network packets in and out of host machines and perform network virtualization functions.

[0157] As indicated above, an NVD executes various virtualization functions including VNICs and VCN VRs. An NVD may execute VNICs associated with the compute instances hosted by one or more host machines connected to the VNIC. For example, as depicted in FIG. 7, NVD 710 executes the functionality for VNIC 776 that is associated with compute instance 768 hosted by host machine 702 connected to NVD 710. As another example, NVD 712 executes VNIC 780 that is associated with bare metal compute instance 772 hosted by host machine 706, and executes VNIC 784 that is associated with compute instance 774 hosted by host machine 708. A host machine may host compute instances belonging to different VCNs, which belong to different customers, and the NVD connected to the host machine may execute the VNICs (i.e., execute VNICs-relate functionality) corresponding to the compute instances.

[0158] An NVD also executes VCN Virtual Routers corresponding to the VCNs of the compute instances. For example, in the embodiment depicted in FIG. 7, NVD 710 executes VCN VR 777 corresponding to the VCN to which compute instance 768 belongs. NVD 712 executes one or more VCN VRs 783 corresponding to one or more VCNs to which compute instances hosted by host machines 706 and 708 belong. In certain embodiments, the VCN VR corresponding to that VCN is executed by all the NVDs connected to host machines that host at least one compute instance belonging to that VCN. If a host machine hosts compute instances belonging to different VCNs, an NVD connected to that host machine may execute VCN VRs corresponding to those different VCNs.

[0159] In addition to VNICs and VCN VRs, an NVD may execute various software (e.g., daemons) and include one or more hardware components that facilitate the various network virtualization functions performed by the NVD. For purposes of simplicity, these various components are grouped together as "packet processing components" shown in FIG. 7. For example, NVD 710 comprises packet processing components 786 and NVD 712 comprises packet processing components 788. For example, the packet processing components for an NVD may include a packet processor that is configured to interact with the NVD's ports and hardware interfaces to monitor all packets received by and communicated using the NVD and store network information. The network information may, for example, include network flow information identifying different network flows handled by the NVD and per flow information (e.g., per flow statistics). In certain embodiments, network flows information may be stored on a per VNIC basis. The packet processor may perform packet-by-packet manipulations as well as implement stateful NAT and L4 firewall (FW). As another example, the packet processing components may include a replication agent that is configured to replicate information stored by the NVD to one or more different replication target stores. As yet another example, the packet processing components may include a logging agent that is configured to perform logging functions for the NVD. The packet processing components may also include software for monitoring the performance and health of the NVD and, also possibly of monitoring the state and health of other components connected to the NVD.

[0160] FIG. 6 shows the components of an example virtual or overlay network including a VCN, subnets within the VCN, compute instances deployed on subnets, VNICs associated with the compute instances, a VR for a VCN, and a set of gateways configured for the VCN. The overlay components depicted in FIG. 6 may be executed or hosted by one or more of the physical components depicted in FIG. 7. For example, the compute instances in a VCN may be executed or hosted by one or more host machines depicted in FIG. 7. For a compute instance hosted by a host machine, the VNIC associated with that compute instance is typically executed by an NVD connected to that host machine (i.e., the VNIC functionality is provided by the NVD connected to that host machine). The VCN VR function for a VCN is executed by all the NVDs that are connected to host machines hosting or executing the compute instances that are part of that VCN. The gateways associated with a VCN may be executed by one or more different types of NVDs. For example, certain gateways may be

executed by smartNICs, while others may be executed by one or more host machines or other implementations of NVDs.

[0161] As described above, a compute instance in a customer VCN may communicate with various different endpoints, where the endpoints can be within the same subnet as the source compute instance, in a different subnet but within the same VCN as the source compute instance, or with an endpoint that is outside the VCN of the source compute instance. These communications are facilitated using VNICs associated with the compute instances, the VCN VRs, and the gateways associated with the VCNs.

[0162] For communications between two compute instances on the same subnet in a VCN, the communication is facilitated using VNICs associated with the source and destination compute instances. The source and destination compute instances may be hosted by the same host machine or by different host machines. A packet originating from a source compute instance may be forwarded from a host machine hosting the source compute instance to an NVD connected to that host machine. On the NVD, the packet is processed using a packet processing pipeline, which can include execution of the VNIC associated with the source compute instance. Since the destination endpoint for the packet is within the same subnet, execution of the VNIC associated with the source compute instance results in the packet being forwarded to an NVD executing the VNIC associated with the destination compute instance, which then processes and forwards the packet to the destination compute instance. The VNICs associated with the source and destination compute instances may be executed on the same NVD (e.g., when both the source and destination compute instances are hosted by the same host machine) or on different NVDs (e.g., when the source and destination compute instances are hosted by different host machines connected to different NVDs). The VNICs may use routing/forwarding tables stored by the NVD to determine the next hop for the packet.

[0163] For a packet to be communicated from a compute instance in a subnet to an endpoint in a different subnet in the same VCN, the packet originating from the source compute instance is communicated from the host machine hosting the source compute instance to the NVD connected to that host machine. On the NVD, the packet is processed using a packet processing pipeline, which can include execution of one or more VNICs, and the VR associated with the VCN. For example, as part of the packet processing pipeline, the NVD executes or invokes functionality corresponding to the VNIC (also referred to as

executes the VNIC) associated with source compute instance. The functionality performed by the VNIC may include looking at the VLAN tag on the packet. Since the packet's destination is outside the subnet, the VCN VR functionality is next invoked and executed by the NVD. The VCN VR then routes the packet to the NVD executing the VNIC associated with the destination compute instance. The VNIC associated with the destination compute instance then processes the packet and forwards the packet to the destination compute instance. The VNICs associated with the source and destination compute instances may be executed on the same NVD (e.g., when both the source and destination compute instances are hosted by the same host machine) or on different NVDs (e.g., when the source and destination compute instances are hosted by different host machines connected to different NVDs).

[0164] If the destination for the packet is outside the VCN of the source compute instance, then the packet originating from the source compute instance is communicated from the host machine hosting the source compute instance to the NVD connected to that host machine.

The NVD executes the VNIC associated with the source compute instance. Since the destination end point of the packet is outside the VCN, the packet is then processed by the VCN VR for that VCN. The NVD invokes the VCN VR functionality, which may result in the packet being forwarded to an NVD executing the appropriate gateway associated with the VCN. For example, if the destination is an endpoint within the customer's on-premise network, then the packet may be forwarded by the VCN VR to the NVD executing the DRG gateway configured for the VCN. The VCN VR may be executed on the same NVD as the NVD executing the VNIC associated with the source compute instance or by a different NVD. The gateway may be executed by an NVD, which may be a smartNIC, a host machine, or other NVD implementation. The packet is then processed by the gateway and forwarded to a next hop that facilitates communication of the packet to its intended destination endpoint. For example, in the embodiment depicted in FIG. 7, a packet originating from compute instance 768 may be communicated from host machine 702 to NVD 710 over link 720 (using NIC 732). On NVD 710, VNIC 776 is invoked since it is the VNIC associated with source compute instance 768. VNIC 776 is configured to examine the encapsulated information in the packet, and determine a next hop for forwarding the packet with the goal of facilitating communication of the packet to its intended destination endpoint, and then forward the packet to the determined next hop.

[0165] A compute instance deployed on a VCN can communicate with various different endpoints. These endpoints may include endpoints that are hosted by CSPI 700 and endpoints outside CSPI 700. Endpoints hosted by CSPI 700 may include instances in the same VCN or other VCNs, which may be the customer's VCNs, or VCNs not belonging to the customer. Communications between endpoints hosted by CSPI 700 may be performed over physical network 718. A compute instance may also communicate with endpoints that are not hosted by CSPI 700, or are outside CSPI 700. Examples of these endpoints include endpoints within a customer's on-premise network or data center, or public endpoints accessible over a public network such as the Internet. Communications with endpoints outside CSPI 700 may be performed over public networks (e.g., the Internet) (not shown in FIG. 7) or private networks (not shown in FIG. 7) using various communication protocols.

[0166] The architecture of CSPI 700 depicted in FIG. 7 is merely an example and is not intended to be limiting. Variations, alternatives, and modifications are possible in alternative embodiments. For example, in some implementations, CSPI 700 may have more or fewer systems or components than those shown in FIG. 7, may combine two or more systems, or may have a different configuration or arrangement of systems. The systems, subsystems, and other components depicted in FIG. 7 may be implemented in software (e.g., code, instructions, program) executed by one or more processing units (e.g., processors, cores) of the respective systems, using hardware, or combinations thereof. The software may be stored on a non-transitory storage medium (e.g., on a memory device).

[0167] FIG. 9 depicts connectivity between a host machine and an NVD for providing I/O virtualization for supporting multitenancy according to certain embodiments. As depicted in FIG. 9, host machine 902 executes a hypervisor 904 that provides a virtualized environment. Host machine 902 executes two virtual machine instances, VM1 906 belonging to customer/tenant #1 and VM2 908 belonging to customer/tenant #2. Host machine 902 comprises a physical NIC 910 that is connected to an NVD 912 via link 914. Each of the compute instances is attached to a VNIC that is executed by NVD 912. In the embodiment in FIG. 9, VM1 906 is attached to VNIC-VM1 920 and VM2 908 is attached to VNIC-VM2 922.

[0168] As shown in FIG. 9, NIC 910 comprises two logical NICs, logical NIC A 916 and logical NIC B 918. Each virtual machine is attached to and configured to work with its own logical NIC. For example, VM1 906 is attached to logical NIC A 916 and VM2 908 is

attached to logical NIC B 918. Even though host machine 902 comprises only one physical NIC 910 that is shared by the multiple tenants, due to the logical NICs, each tenant's virtual machine believes they have their own host machine and NIC.

[0169] In certain embodiments, each logical NIC is assigned its own VLAN ID. Thus, a specific VLAN ID is assigned to logical NIC A 916 for Tenant #1 and a separate VLAN ID is assigned to logical NIC B 918 for Tenant #2. When a packet is communicated from VM1 906, a tag assigned to Tenant #1 is attached to the packet by the hypervisor and the packet is then communicated from host machine 902 to NVD 912 over link 914. In a similar manner, when a packet is communicated from VM2 908, a tag assigned to Tenant #2 is attached to the packet by the hypervisor and the packet is then communicated from host machine 902 to NVD 912 over link 914. Accordingly, a packet 924 communicated from host machine 902 to NVD 912 has an associated tag 926 that identifies a specific tenant and associated VM. On the NVD, for a packet 924 received from host machine 902, the tag 926 associated with the packet is used to determine whether the packet is to be processed by VNIC-VM1 920 or by VNIC-VM2 922. The packet is then processed by the corresponding VNIC. The configuration depicted in FIG. 9 enables each tenant's compute instance to believe that they own their own host machine and NIC. The setup depicted in FIG. 9 provides for I/O virtualization for supporting multi-tenancy.

[0170] FIG. 10 depicts a simplified block diagram of a physical network 1000 according to certain embodiments. The embodiment depicted in FIG. 10 is structured as a Clos network. A Clos network is a particular type of network topology designed to provide connection redundancy while maintaining high bisection bandwidth and maximum resource utilization. A Clos network is a type of non-blocking, multistage or multi-tiered switching network, where the number of stages or tiers can be two, three, four, five, etc. The embodiment depicted in FIG. 10 is a 3-tiered network comprising tiers 1, 2, and 3. The TOR switches 1004 represent Tier-0 switches in the Clos network. One or more NVDs are connected to the TOR switches. Tier-0 switches are also referred to as edge devices of the physical network. The Tier-0 switches are connected to Tier-1 switches, which are also referred to as leaf switches. In the embodiment depicted in FIG. 10, a set of "n" Tier-0 TOR switches are connected to a set of "n" Tier-1 switches and together form a pod. Each Tier-0 switch in a pod is interconnected to all the Tier-1 switches in the pod, but there is no connectivity of switches between pods. In certain implementations, two pods are referred to as a block. Each block is served by or connected to a set of "n" Tier-2 switches (sometimes referred to as

spine switches). There can be several blocks in the physical network topology. The Tier-2 switches are in turn connected to "n" Tier-3 switches (sometimes referred to as super-spine switches). Communication of packets over physical network 1000 is typically performed using one or more Layer-3 communication protocols. Typically, all the layers of the physical network, except for the TORs layer are n-ways redundant thus allowing for high availability. Policies may be specified for pods and blocks to control the visibility of switches to each other in the physical network so as to enable scaling of the physical network.

[0171] A feature of a Clos network is that the maximum hop count to reach from one Tier-0 switch to another Tier-0 switch (or from an NVD connected to a Tier-0 switch to another NVD connected to a Tier-0 switch) is fixed. For example, in a 3-Tiered Clos network at most seven hops are needed for a packet to reach from one NVD to another NVD, where the source and target NVDs are connected to the leaf tier of the Clos network. Likewise, in a 4-tiered Clos network, at most nine hops are needed for a packet to reach from one NVD to another NVD, where the source and target NVDs are connected to the leaf tier of the Clos network. Thus, a Clos network architecture maintains consistent latency throughout the network, which is important for communication within and between data centers. A Clos topology scales horizontally and is cost effective. The bandwidth/throughput capacity of the network can be easily increased by adding more switches at the various tiers (e.g., more leaf and spine switches) and by increasing the number of links between the switches at adjacent tiers.

[0172] In certain embodiments, each resource within CSPI is assigned a unique identifier called a Cloud Identifier (CID). This identifier is included as part of the resource's information and can be used to manage the resource, for example, via a Console or through APIs. An example syntax for a CID is:

ocid1.<RESOURCE TYPE>.<REALM> [REGION][.FUTURE USE].<UNIQUE ID>

where,

ocid1: The literal string indicating the version of the CID;

resource type: The type of resource (for example, instance, volume, VCN, subnet, user, group, and so on);

realm: The realm the resource is in. Example values are "c1" for the commercial realm, "c2" for the Government Cloud realm, or "c3" for the Federal Government Cloud realm, etc. Each realm may have its own domain name;

region: The region the resource is in. If the region is not applicable to the resource, this part might be blank;

future use: Reserved for future use.

unique ID: The unique portion of the ID. The format may vary depending on the type of resource or service.

[0173] FIG. 11 is a block diagram 1100 illustrating an example pattern of an IaaS architecture, according to at least one embodiment. Service operators 1102 can be communicatively coupled to a secure host tenancy 1104 that can include a virtual cloud network (VCN) 1106 and a secure host subnet 1108. In some examples, the service operators 1102 may be using one or more client computing devices, which may be portable handheld devices (e.g., an iPhone®, cellular telephone, an iPad®, computing tablet, a personal digital assistant (PDA)) or wearable devices (e.g., a Google Glass® head mounted display), running software such as Microsoft Windows Mobile®, and/or a variety of mobile operating systems such as iOS, Windows Phone, Android, BlackBerry 8, Palm OS, and the like, and being Internet, e-mail, short message service (SMS), Blackberry®, or other communication protocol enabled. Alternatively, the client computing devices can be general purpose personal computers including, by way of example, personal computers and/or laptop computers running various versions of Microsoft Windows®, Apple Macintosh®, and/or Linux operating systems. The client computing devices can be workstation computers running any of a variety of commercially-available UNIX® or UNIX-like operating systems, including without limitation the variety of GNU/Linux operating systems, such as for example, Google Chrome OS. Alternatively, or in addition, client computing devices may be any other electronic device, such as a thin-client computer, an Internet-enabled gaming system (e.g., a Microsoft Xbox gaming console with or without a Kinect® gesture input device), and/or a personal messaging device, capable of communicating over a network that can access the VCN 1106 and/or the Internet.

[0174] The VCN 1106 can include a local peering gateway (LPG) 1110 that can be communicatively coupled to a secure shell (SSH) VCN 1112 via an LPG 1110 contained in the SSH VCN 1112. The SSH VCN 1112 can include an SSH subnet 1114, and the SSH VCN 1112 can be communicatively coupled to a control plane VCN 1116 via the LPG 1110 contained in the control plane VCN 1116. Also, the SSH VCN 1112 can be communicatively coupled to a data plane VCN 1118 via an LPG 1110. The control plane VCN 1116 and the

data plane VCN 1118 can be contained in a service tenancy 1119 that can be owned and/or operated by the IaaS provider.

[0175] The control plane VCN 1116 can include a control plane demilitarized zone (DMZ) tier 1120 that acts as a perimeter network (e.g., portions of a corporate network between the corporate intranet and external networks). The DMZ-based servers may have restricted responsibilities and help keep breaches contained. Additionally, the DMZ tier 1120 can include one or more load balancer (LB) subnet(s) 1122, a control plane app tier 1124 that can include app subnet(s) 1126, a control plane data tier 1128 that can include database (DB) subnet(s) 1130 (e.g., frontend DB subnet(s) and/or backend DB subnet(s)). The LB subnet(s) 1122 contained in the control plane DMZ tier 1120 can be communicatively coupled to the app subnet(s) 1126 contained in the control plane app tier 1124 and an Internet gateway 1134 that can be contained in the control plane VCN 1116, and the app subnet(s) 1126 can be communicatively coupled to the DB subnet(s) 1130 contained in the control plane data tier 1128 and a service gateway 1136 and a network address translation (NAT) gateway 1138. The control plane VCN 1116 can include the service gateway 1136 and the NAT gateway 1138.

[0176] The control plane VCN 1116 can include a data plane mirror app tier 1140 that can include app subnet(s) 1126. The app subnet(s) 1126 contained in the data plane mirror app tier 1140 can include a virtual network interface controller (VNIC) 1142 that can execute a compute instance 1144. The compute instance 1144 can communicatively couple the app subnet(s) 1126 of the data plane mirror app tier 1140 to app subnet(s) 1126 that can be contained in a data plane app tier 1146.

[0177] The data plane VCN 1118 can include the data plane app tier 1146, a data plane DMZ tier 1148, and a data plane data tier 1150. The data plane DMZ tier 1148 can include LB subnet(s) 1122 that can be communicatively coupled to the app subnet(s) 1126 of the data plane app tier 1146 and the Internet gateway 1134 of the data plane VCN 1118. The app subnet(s) 1126 can be communicatively coupled to the service gateway 1136 of the data plane VCN 1118 and the NAT gateway 1138 of the data plane VCN 1118. The data plane data tier 1150 can also include the DB subnet(s) 1130 that can be communicatively coupled to the app subnet(s) 1126 of the data plane app tier 1146.

[0178] The Internet gateway 1134 of the control plane VCN 1116 and of the data plane VCN 1118 can be communicatively coupled to a metadata management service 1152 that can

be communicatively coupled to public Internet 1154. Public Internet 1154 can be communicatively coupled to the NAT gateway 1138 of the control plane VCN 1116 and of the data plane VCN 1118. The service gateway 1136 of the control plane VCN 1116 and of the data plane VCN 1118 can be communicatively couple to cloud services 1156.

5 [0179] In some examples, the service gateway 1136 of the control plane VCN 1116 or of the data plane VCN 1118 can make application programming interface (API) calls to cloud services 1156 without going through public Internet 1154. The API calls to cloud services 1156 from the service gateway 1136 can be one-way: the service gateway 1136 can make API calls to cloud services 1156, and cloud services 1156 can send requested data to the
10 service gateway 1136. But, cloud services 1156 may not initiate API calls to the service gateway 1136.

[0180] In some examples, the secure host tenancy 1104 can be directly connected to the service tenancy 1119, which may be otherwise isolated. The secure host subnet 1108 can communicate with the SSH subnet 1114 through an LPG 1110 that may enable two-way
15 communication over an otherwise isolated system. Connecting the secure host subnet 1108 to the SSH subnet 1114 may give the secure host subnet 1108 access to other entities within the service tenancy 1119.

[0181] The control plane VCN 1116 may allow users of the service tenancy 1119 to set up or otherwise provision desired resources. Desired resources provisioned in the control plane
20 VCN 1116 may be deployed or otherwise used in the data plane VCN 1118. In some examples, the control plane VCN 1116 can be isolated from the data plane VCN 1118, and the data plane mirror app tier 1140 of the control plane VCN 1116 can communicate with the data plane app tier 1146 of the data plane VCN 1118 via VNICs 1142 that can be contained in the data plane mirror app tier 1140 and the data plane app tier 1146.

25 [0182] In some examples, users of the system, or customers, can make requests, for example create, read, update, or delete (CRUD) operations, through public Internet 1154 that can communicate the requests to the metadata management service 1152. The metadata management service 1152 can communicate the request to the control plane VCN 1116 through the Internet gateway 1134. The request can be received by the LB subnet(s) 1122
30 contained in the control plane DMZ tier 1120. The LB subnet(s) 1122 may determine that the request is valid, and in response to this determination, the LB subnet(s) 1122 can transmit the request to app subnet(s) 1126 contained in the control plane app tier 1124. If the request is

validated and requires a call to public Internet 1154, the call to public Internet 1154 may be transmitted to the NAT gateway 1138 that can make the call to public Internet 1154. Memory that may be desired to be stored by the request can be stored in the DB subnet(s) 1130.

[0183] In some examples, the data plane mirror app tier 1140 can facilitate direct communication between the control plane VCN 1116 and the data plane VCN 1118. For example, changes, updates, or other suitable modifications to configuration may be desired to be applied to the resources contained in the data plane VCN 1118. Via a VNIC 1142, the control plane VCN 1116 can directly communicate with, and can thereby execute the changes, updates, or other suitable modifications to configuration to, resources contained in the data plane VCN 1118.

[0184] In some embodiments, the control plane VCN 1116 and the data plane VCN 1118 can be contained in the service tenancy 1119. In this case, the user, or the customer, of the system may not own or operate either the control plane VCN 1116 or the data plane VCN 1118. Instead, the IaaS provider may own or operate the control plane VCN 1116 and the data plane VCN 1118, both of which may be contained in the service tenancy 1119. This embodiment can enable isolation of networks that may prevent users or customers from interacting with other users', or other customers', resources. Also, this embodiment may allow users or customers of the system to store databases privately without needing to rely on public Internet 1154, which may not have a desired level of threat prevention, for storage.

[0185] In other embodiments, the LB subnet(s) 1122 contained in the control plane VCN 1116 can be configured to receive a signal from the service gateway 1136. In this embodiment, the control plane VCN 1116 and the data plane VCN 1118 may be configured to be called by a customer of the IaaS provider without calling public Internet 1154. Customers of the IaaS provider may desire this embodiment since database(s) that the customers use may be controlled by the IaaS provider and may be stored on the service tenancy 1119, which may be isolated from public Internet 1154.

[0186] FIG. 12 is a block diagram 1200 illustrating another example pattern of an IaaS architecture, according to at least one embodiment. Service operators 1202 (e.g. service operators 1102 of FIG. 11) can be communicatively coupled to a secure host tenancy 1204 (e.g. the secure host tenancy 1104 of FIG. 11) that can include a virtual cloud network (VCN) 1206 (e.g. the VCN 1106 of FIG. 11) and a secure host subnet 1208 (e.g. the secure host subnet 1108 of FIG. 11). The VCN 1206 can include a local peering gateway (LPG) 1210

(e.g. the LPG 1110 of FIG. 11) that can be communicatively coupled to a secure shell (SSH) VCN 1212 (e.g. the SSH VCN 1112 of FIG. 11) via an LPG 1110 contained in the SSH VCN 1212. The SSH VCN 1212 can include an SSH subnet 1214 (e.g. the SSH subnet 1114 of FIG. 11), and the SSH VCN 1212 can be communicatively coupled to a control plane VCN 1216 (e.g. the control plane VCN 1116 of FIG. 11) via an LPG 1210 contained in the control plane VCN 1216. The control plane VCN 1216 can be contained in a service tenancy 1219 (e.g. the service tenancy 1119 of FIG. 11), and the data plane VCN 1218 (e.g. the data plane VCN 1118 of FIG. 11) can be contained in a customer tenancy 1221 that may be owned or operated by users, or customers, of the system.

[0187] The control plane VCN 1216 can include a control plane DMZ tier 1220 (e.g. the control plane DMZ tier 1120 of FIG. 11) that can include LB subnet(s) 1222 (e.g. LB subnet(s) 1122 of FIG. 11), a control plane app tier 1224 (e.g. the control plane app tier 1124 of FIG. 11) that can include app subnet(s) 1226 (e.g. app subnet(s) 1126 of FIG. 11), a control plane data tier 1228 (e.g. the control plane data tier 1128 of FIG. 11) that can include database (DB) subnet(s) 1230 (e.g. similar to DB subnet(s) 1130 of FIG. 11). The LB subnet(s) 1222 contained in the control plane DMZ tier 1220 can be communicatively coupled to the app subnet(s) 1226 contained in the control plane app tier 1224 and an Internet gateway 1234 (e.g. the Internet gateway 1134 of FIG. 11) that can be contained in the control plane VCN 1216, and the app subnet(s) 1226 can be communicatively coupled to the DB subnet(s) 1230 contained in the control plane data tier 1228 and a service gateway 1236 (e.g. the service gateway of FIG. 11) and a network address translation (NAT) gateway 1238 (e.g. the NAT gateway 1138 of FIG. 11). The control plane VCN 1216 can include the service gateway 1236 and the NAT gateway 1238.

[0188] The control plane VCN 1216 can include a data plane mirror app tier 1240 (e.g. the data plane mirror app tier 1140 of FIG. 11) that can include app subnet(s) 1226. The app subnet(s) 1226 contained in the data plane mirror app tier 1240 can include a virtual network interface controller (VNIC) 1242 (e.g. the VNIC of 1142) that can execute a compute instance 1244 (e.g. similar to the compute instance 1144 of FIG. 11). The compute instance 1244 can facilitate communication between the app subnet(s) 1226 of the data plane mirror app tier 1240 and the app subnet(s) 1226 that can be contained in a data plane app tier 1246 (e.g. the data plane app tier 1146 of FIG. 11) via the VNIC 1242 contained in the data plane mirror app tier 1240 and the VNIC 1242 contained in the data plane app tier 1246.

[0189] The Internet gateway 1234 contained in the control plane VCN 1216 can be communicatively coupled to a metadata management service 1252 (e.g. the metadata management service 1152 of FIG. 11) that can be communicatively coupled to public Internet 1254 (e.g. public Internet 1154 of FIG. 11). Public Internet 1254 can be communicatively coupled to the NAT gateway 1238 contained in the control plane VCN 1216. The service gateway 1236 contained in the control plane VCN 1216 can be communicatively couple to cloud services 1256 (e.g. cloud services 1156 of FIG. 11).

[0190] In some examples, the data plane VCN 1218 can be contained in the customer tenancy 1221. In this case, the IaaS provider may provide the control plane VCN 1216 for each customer, and the IaaS provider may, for each customer, set up a unique compute instance 1244 that is contained in the service tenancy 1219. Each compute instance 1244 may allow communication between the control plane VCN 1216, contained in the service tenancy 1219, and the data plane VCN 1218 that is contained in the customer tenancy 1221. The compute instance 1244 may allow resources, that are provisioned in the control plane VCN 1216 that is contained in the service tenancy 1219, to be deployed or otherwise used in the data plane VCN 1218 that is contained in the customer tenancy 1221.

[0191] In other examples, the customer of the IaaS provider may have databases that live in the customer tenancy 1221. In this example, the control plane VCN 1216 can include the data plane mirror app tier 1240 that can include app subnet(s) 1226. The data plane mirror app tier 1240 can reside in the data plane VCN 1218, but the data plane mirror app tier 1240 may not live in the data plane VCN 1218. That is, the data plane mirror app tier 1240 may have access to the customer tenancy 1221, but the data plane mirror app tier 1240 may not exist in the data plane VCN 1218 or be owned or operated by the customer of the IaaS provider. The data plane mirror app tier 1240 may be configured to make calls to the data plane VCN 1218 but may not be configured to make calls to any entity contained in the control plane VCN 1216. The customer may desire to deploy or otherwise use resources in the data plane VCN 1218 that are provisioned in the control plane VCN 1216, and the data plane mirror app tier 1240 can facilitate the desired deployment, or other usage of resources, of the customer.

[0192] In some embodiments, the customer of the IaaS provider can apply filters to the data plane VCN 1218. In this embodiment, the customer can determine what the data plane VCN 1218 can access, and the customer may restrict access to public Internet 1254 from the data plane VCN 1218. The IaaS provider may not be able to apply filters or otherwise control

access of the data plane VCN 1218 to any outside networks or databases. Applying filters and controls by the customer onto the data plane VCN 1218, contained in the customer tenancy 1221, can help isolate the data plane VCN 1218 from other customers and from public Internet 1254.

5 [0193] In some embodiments, cloud services 1256 can be called by the service gateway 1236 to access services that may not exist on public Internet 1254, on the control plane VCN 1216, or on the data plane VCN 1218. The connection between cloud services 1256 and the control plane VCN 1216 or the data plane VCN 1218 may not be live or continuous. Cloud services 1256 may exist on a different network owned or operated by the IaaS provider.

10 Cloud services 1256 may be configured to receive calls from the service gateway 1236 and may be configured to not receive calls from public Internet 1254. Some cloud services 1256 may be isolated from other cloud services 1256, and the control plane VCN 1216 may be isolated from cloud services 1256 that may not be in the same region as the control plane VCN 1216. For example, the control plane VCN 1216 may be located in “Region 1,” and
15 cloud service “Deployment 11,” may be located in Region 1 and in “Region 2.” If a call to Deployment 11 is made by the service gateway 1236 contained in the control plane VCN 1216 located in Region 1, the call may be transmitted to Deployment 11 in Region 1. In this example, the control plane VCN 1216, or Deployment 11 in Region 1, may not be communicatively coupled to, or otherwise in communication with, Deployment 11 in Region
20 2.

[0194] FIG. 13 is a block diagram 1300 illustrating another example pattern of an IaaS architecture, according to at least one embodiment. Service operators 1302 (e.g. service operators 1102 of FIG. 11) can be communicatively coupled to a secure host tenancy 1304 (e.g. the secure host tenancy 1104 of FIG. 11) that can include a virtual cloud network (VCN) 1306 (e.g. the VCN 1106 of FIG. 11) and a secure host subnet 1308 (e.g. the secure host subnet 1108 of FIG. 11). The VCN 1306 can include an LPG 1310 (e.g. the LPG 1110 of FIG. 11) that can be communicatively coupled to an SSH VCN 1312 (e.g. the SSH VCN 1112 of FIG. 11) via an LPG 1310 contained in the SSH VCN 1312. The SSH VCN 1312 can include an SSH subnet 1314 (e.g. the SSH subnet 1114 of FIG. 11), and the SSH VCN 1312
25 can be communicatively coupled to a control plane VCN 1316 (e.g. the control plane VCN 1116 of FIG. 11) via an LPG 1310 contained in the control plane VCN 1316 and to a data plane VCN 1318 (e.g. the data plane 1118 of FIG. 11) via an LPG 1310 contained in the data
30

plane VCN 1318. The control plane VCN 1316 and the data plane VCN 1318 can be contained in a service tenancy 1319 (e.g. the service tenancy 1119 of FIG. 11).

[0195] The control plane VCN 1316 can include a control plane DMZ tier 1320 (e.g. the control plane DMZ tier 1120 of FIG. 11) that can include load balancer (LB) subnet(s) 1322 (e.g. LB subnet(s) 1122 of FIG. 11), a control plane app tier 1324 (e.g. the control plane app tier 1124 of FIG. 11) that can include app subnet(s) 1326 (e.g. similar to app subnet(s) 1126 of FIG. 11), a control plane data tier 1328 (e.g. the control plane data tier 1128 of FIG. 11) that can include DB subnet(s) 1330. The LB subnet(s) 1322 contained in the control plane DMZ tier 1320 can be communicatively coupled to the app subnet(s) 1326 contained in the control plane app tier 1324 and to an Internet gateway 1334 (e.g. the Internet gateway 1134 of FIG. 11) that can be contained in the control plane VCN 1316, and the app subnet(s) 1326 can be communicatively coupled to the DB subnet(s) 1330 contained in the control plane data tier 1328 and to a service gateway 1336 (e.g. the service gateway of FIG. 11) and a network address translation (NAT) gateway 1338 (e.g. the NAT gateway 1138 of FIG. 11). The control plane VCN 1316 can include the service gateway 1336 and the NAT gateway 1338.

[0196] The data plane VCN 1318 can include a data plane app tier 1346 (e.g. the data plane app tier 1146 of FIG. 11), a data plane DMZ tier 1348 (e.g. the data plane DMZ tier 1148 of FIG. 11), and a data plane data tier 1350 (e.g. the data plane data tier 1150 of FIG. 11). The data plane DMZ tier 1348 can include LB subnet(s) 1322 that can be communicatively coupled to trusted app subnet(s) 1360 and untrusted app subnet(s) 1362 of the data plane app tier 1346 and the Internet gateway 1334 contained in the data plane VCN 1318. The trusted app subnet(s) 1360 can be communicatively coupled to the service gateway 1336 contained in the data plane VCN 1318, the NAT gateway 1338 contained in the data plane VCN 1318, and DB subnet(s) 1330 contained in the data plane data tier 1350. The untrusted app subnet(s) 1362 can be communicatively coupled to the service gateway 1336 contained in the data plane VCN 1318 and DB subnet(s) 1330 contained in the data plane data tier 1350. The data plane data tier 1350 can include DB subnet(s) 1330 that can be communicatively coupled to the service gateway 1336 contained in the data plane VCN 1318.

[0197] The untrusted app subnet(s) 1362 can include one or more primary VNICs 1364(1)-(N) that can be communicatively coupled to tenant virtual machines (VMs) 1366(1)-(N). Each tenant VM 1366(1)-(N) can be communicatively coupled to a respective app subnet 1367(1)-(N) that can be contained in respective container egress VCNs 1368(1)-(N) that can

be contained in respective customer tenancies 1370(1)-(N). Respective secondary VNICs 1372(1)-(N) can facilitate communication between the untrusted app subnet(s) 1362 contained in the data plane VCN 1318 and the app subnet contained in the container egress VCNs 1368(1)-(N). Each container egress VCNs 1368(1)-(N) can include a NAT gateway 1338 that can be communicatively coupled to public Internet 1354 (e.g. public Internet 1154 of FIG. 11).

[0198] The Internet gateway 1334 contained in the control plane VCN 1316 and contained in the data plane VCN 1318 can be communicatively coupled to a metadata management service 1352 (e.g. the metadata management system 1152 of FIG. 11) that can be communicatively coupled to public Internet 1354. Public Internet 1354 can be communicatively coupled to the NAT gateway 1338 contained in the control plane VCN 1316 and contained in the data plane VCN 1318. The service gateway 1336 contained in the control plane VCN 1316 and contained in the data plane VCN 1318 can be communicatively couple to cloud services 1356.

[0199] In some embodiments, the data plane VCN 1318 can be integrated with customer tenancies 1370. This integration can be useful or desirable for customers of the IaaS provider in some cases such as a case that may desire support when executing code. The customer may provide code to run that may be destructive, may communicate with other customer resources, or may otherwise cause undesirable effects. In response to this, the IaaS provider may determine whether to run code given to the IaaS provider by the customer.

[0200] In some examples, the customer of the IaaS provider may grant temporary network access to the IaaS provider and request a function to be attached to the data plane tier app 1346. Code to run the function may be executed in the VMs 1366(1)-(N), and the code may not be configured to run anywhere else on the data plane VCN 1318. Each VM 1366(1)-(N) may be connected to one customer tenancy 1370. Respective containers 1371(1)-(N) contained in the VMs 1366(1)-(N) may be configured to run the code. In this case, there can be a dual isolation (e.g., the containers 1371(1)-(N) running code, where the containers 1371(1)-(N) may be contained in at least the VM 1366(1)-(N) that are contained in the untrusted app subnet(s) 1362), which may help prevent incorrect or otherwise undesirable code from damaging the network of the IaaS provider or from damaging a network of a different customer. The containers 1371(1)-(N) may be communicatively coupled to the customer tenancy 1370 and may be configured to transmit or receive data from the customer

tenancy 1370. The containers 1371(1)-(N) may not be configured to transmit or receive data from any other entity in the data plane VCN 1318. Upon completion of running the code, the IaaS provider may kill or otherwise dispose of the containers 1371(1)-(N).

[0201] In some embodiments, the trusted app subnet(s) 1360 may run code that may be owned or operated by the IaaS provider. In this embodiment, the trusted app subnet(s) 1360 may be communicatively coupled to the DB subnet(s) 1330 and be configured to execute CRUD operations in the DB subnet(s) 1330. The untrusted app subnet(s) 1362 may be communicatively coupled to the DB subnet(s) 1330, but in this embodiment, the untrusted app subnet(s) may be configured to execute read operations in the DB subnet(s) 1330. The containers 1371(1)-(N) that can be contained in the VM 1366(1)-(N) of each customer and that may run code from the customer may not be communicatively coupled with the DB subnet(s) 1330.

[0202] In other embodiments, the control plane VCN 1316 and the data plane VCN 1318 may not be directly communicatively coupled. In this embodiment, there may be no direct communication between the control plane VCN 1316 and the data plane VCN 1318. However, communication can occur indirectly through at least one method. An LPG 1310 may be established by the IaaS provider that can facilitate communication between the control plane VCN 1316 and the data plane VCN 1318. In another example, the control plane VCN 1316 or the data plane VCN 1318 can make a call to cloud services 1356 via the service gateway 1336. For example, a call to cloud services 1356 from the control plane VCN 1316 can include a request for a service that can communicate with the data plane VCN 1318.

[0203] FIG. 14 is a block diagram 1400 illustrating another example pattern of an IaaS architecture, according to at least one embodiment. Service operators 1402 (e.g. service operators 1102 of FIG. 11) can be communicatively coupled to a secure host tenancy 1404 (e.g. the secure host tenancy 1104 of FIG. 11) that can include a virtual cloud network (VCN) 1406 (e.g. the VCN 1106 of FIG. 11) and a secure host subnet 1408 (e.g. the secure host subnet 1108 of FIG. 11). The VCN 1406 can include an LPG 1410 (e.g. the LPG 1110 of FIG. 11) that can be communicatively coupled to an SSH VCN 1412 (e.g. the SSH VCN 1112 of FIG. 11) via an LPG 1410 contained in the SSH VCN 1412. The SSH VCN 1412 can include an SSH subnet 1414 (e.g. the SSH subnet 1114 of FIG. 11), and the SSH VCN 1412 can be communicatively coupled to a control plane VCN 1416 (e.g. the control plane VCN 1116 of FIG. 11) via an LPG 1410 contained in the control plane VCN 1416 and to a data

plane VCN 1418 (e.g. the data plane 1118 of FIG. 11) via an LPG 1410 contained in the data plane VCN 1418. The control plane VCN 1416 and the data plane VCN 1418 can be contained in a service tenancy 1419 (e.g. the service tenancy 1119 of FIG. 11).

[0204] The control plane VCN 1416 can include a control plane DMZ tier 1420 (e.g. the control plane DMZ tier 1120 of FIG. 11) that can include LB subnet(s) 1422 (e.g. LB subnet(s) 1122 of FIG. 11), a control plane app tier 1424 (e.g. the control plane app tier 1124 of FIG. 11) that can include app subnet(s) 1426 (e.g. app subnet(s) 1126 of FIG. 11), a control plane data tier 1428 (e.g. the control plane data tier 1128 of FIG. 11) that can include DB subnet(s) 1430 (e.g. DB subnet(s) 1330 of FIG. 13). The LB subnet(s) 1422 contained in the control plane DMZ tier 1420 can be communicatively coupled to the app subnet(s) 1426 contained in the control plane app tier 1424 and to an Internet gateway 1434 (e.g. the Internet gateway 1134 of FIG. 11) that can be contained in the control plane VCN 1416, and the app subnet(s) 1426 can be communicatively coupled to the DB subnet(s) 1430 contained in the control plane data tier 1428 and to a service gateway 1436 (e.g. the service gateway of FIG. 11) and a network address translation (NAT) gateway 1438 (e.g. the NAT gateway 1138 of FIG. 11). The control plane VCN 1416 can include the service gateway 1436 and the NAT gateway 1438.

[0205] The data plane VCN 1418 can include a data plane app tier 1446 (e.g. the data plane app tier 1146 of FIG. 11), a data plane DMZ tier 1448 (e.g. the data plane DMZ tier 1148 of FIG. 11), and a data plane data tier 1450 (e.g. the data plane data tier 1150 of FIG. 11). The data plane DMZ tier 1448 can include LB subnet(s) 1422 that can be communicatively coupled to trusted app subnet(s) 1460 (e.g. trusted app subnet(s) 1360 of FIG. 13) and untrusted app subnet(s) 1462 (e.g. untrusted app subnet(s) 1362 of FIG. 13) of the data plane app tier 1446 and the Internet gateway 1434 contained in the data plane VCN 1418. The trusted app subnet(s) 1460 can be communicatively coupled to the service gateway 1436 contained in the data plane VCN 1418, the NAT gateway 1438 contained in the data plane VCN 1418, and DB subnet(s) 1430 contained in the data plane data tier 1450. The untrusted app subnet(s) 1462 can be communicatively coupled to the service gateway 1436 contained in the data plane VCN 1418 and DB subnet(s) 1430 contained in the data plane data tier 1450. The data plane data tier 1450 can include DB subnet(s) 1430 that can be communicatively coupled to the service gateway 1436 contained in the data plane VCN 1418.

[0206] The untrusted app subnet(s) 1462 can include primary VNICs 1464(1)-(N) that can be communicatively coupled to tenant virtual machines (VMs) 1466(1)-(N) residing within the untrusted app subnet(s) 1462. Each tenant VM 1466(1)-(N) can run code in a respective container 1467(1)-(N), and be communicatively coupled to an app subnet 1426 that can be contained in a data plane app tier 1446 that can be contained in a container egress VCN 1468. Respective secondary VNICs 1472(1)-(N) can facilitate communication between the untrusted app subnet(s) 1462 contained in the data plane VCN 1418 and the app subnet contained in the container egress VCN 1468. The container egress VCN can include a NAT gateway 1438 that can be communicatively coupled to public Internet 1454 (e.g. public Internet 1154 of FIG. 11).

[0207] The Internet gateway 1434 contained in the control plane VCN 1416 and contained in the data plane VCN 1418 can be communicatively coupled to a metadata management service 1452 (e.g. the metadata management system 1152 of FIG. 11) that can be communicatively coupled to public Internet 1454. Public Internet 1454 can be communicatively coupled to the NAT gateway 1438 contained in the control plane VCN 1416 and contained in the data plane VCN 1418. The service gateway 1436 contained in the control plane VCN 1416 and contained in the data plane VCN 1418 can be communicatively couple to cloud services 1456.

[0208] In some examples, the pattern illustrated by the architecture of block diagram 1400 of FIG. 14 may be considered an exception to the pattern illustrated by the architecture of block diagram 1300 of FIG. 13 and may be desirable for a customer of the IaaS provider if the IaaS provider cannot directly communicate with the customer (e.g., a disconnected region). The respective containers 1467(1)-(N) that are contained in the VMs 1466(1)-(N) for each customer can be accessed in real-time by the customer. The containers 1467(1)-(N) may be configured to make calls to respective secondary VNICs 1472(1)-(N) contained in app subnet(s) 1426 of the data plane app tier 1446 that can be contained in the container egress VCN 1468. The secondary VNICs 1472(1)-(N) can transmit the calls to the NAT gateway 1438 that may transmit the calls to public Internet 1454. In this example, the containers 1467(1)-(N) that can be accessed in real-time by the customer can be isolated from the control plane VCN 1416 and can be isolated from other entities contained in the data plane VCN 1418. The containers 1467(1)-(N) may also be isolated from resources from other customers.

[0209] In other examples, the customer can use the containers 1467(1)-(N) to call cloud services 1456. In this example, the customer may run code in the containers 1467(1)-(N) that requests a service from cloud services 1456. The containers 1467(1)-(N) can transmit this request to the secondary VNICs 1472(1)-(N) that can transmit the request to the NAT gateway that can transmit the request to public Internet 1454. Public Internet 1454 can transmit the request to LB subnet(s) 1422 contained in the control plane VCN 1416 via the Internet gateway 1434. In response to determining the request is valid, the LB subnet(s) can transmit the request to app subnet(s) 1426 that can transmit the request to cloud services 1456 via the service gateway 1436.

[0210] It should be appreciated that IaaS architectures 1100, 1200, 1300, 1400 depicted in the figures may have other components than those depicted. Further, the embodiments shown in the figures are only some examples of a cloud infrastructure system that may incorporate an embodiment of the disclosure. In some other embodiments, the IaaS systems may have more or fewer components than shown in the figures, may combine two or more components, or may have a different configuration or arrangement of components.

[0211] In certain embodiments, the IaaS systems described herein may include a suite of applications, middleware, and database service offerings that are delivered to a customer in a self-service, subscription-based, elastically scalable, reliable, highly available, and secure manner. An example of such an IaaS system is the Oracle Cloud Infrastructure (OCI) provided by the present assignee.

[0212] FIG. 15 illustrates an example computer system 1500, in which various embodiments may be implemented. The system 1500 may be used to implement any of the computer systems described above. As shown in the figure, computer system 1500 includes a processing unit 1504 that communicates with a number of peripheral subsystems via a bus subsystem 1502. These peripheral subsystems may include a processing acceleration unit 1506, an I/O subsystem 1508, a storage subsystem 1518 and a communications subsystem 1524. Storage subsystem 1518 includes tangible computer-readable storage media 1522 and a system memory 1510.

[0213] Bus subsystem 1502 provides a mechanism for letting the various components and subsystems of computer system 1500 communicate with each other as intended. Although bus subsystem 1502 is shown schematically as a single bus, alternative embodiments of the bus subsystem may utilize multiple buses. Bus subsystem 1502 may be any of several types

of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. For example, such architectures may include an Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus, which can be implemented as a Mezzanine bus manufactured to the IEEE P1386.1 standard.

[0214] Processing unit 1504, which can be implemented as one or more integrated circuits (e.g., a conventional microprocessor or microcontroller), controls the operation of computer system 1500. One or more processors may be included in processing unit 1504. These processors may include single core or multicore processors. In certain embodiments, processing unit 1504 may be implemented as one or more independent processing units 1532 and/or 1534 with single or multicore processors included in each processing unit. In other embodiments, processing unit 1504 may also be implemented as a quad-core processing unit formed by integrating two dual-core processors into a single chip.

[0215] In various embodiments, processing unit 1504 can execute a variety of programs in response to program code and can maintain multiple concurrently executing programs or processes. At any given time, some or all of the program code to be executed can be resident in processor(s) 1504 and/or in storage subsystem 1518. Through suitable programming, processor(s) 1504 can provide various functionalities described above. Computer system 1500 may additionally include a processing acceleration unit 1506, which can include a digital signal processor (DSP), a special-purpose processor, and/or the like.

[0216] I/O subsystem 1508 may include user interface input devices and user interface output devices. User interface input devices may include a keyboard, pointing devices such as a mouse or trackball, a touchpad or touch screen incorporated into a display, a scroll wheel, a click wheel, a dial, a button, a switch, a keypad, audio input devices with voice command recognition systems, microphones, and other types of input devices. User interface input devices may include, for example, motion sensing and/or gesture recognition devices such as the Microsoft Kinect® motion sensor that enables users to control and interact with an input device, such as the Microsoft Xbox® 360 game controller, through a natural user interface using gestures and spoken commands. User interface input devices may also include eye gesture recognition devices such as the Google Glass® blink detector that detects eye activity (e.g., 'blinking' while taking pictures and/or making a menu selection) from users and

transforms the eye gestures as input into an input device (e.g., Google Glass®). Additionally, user interface input devices may include voice recognition sensing devices that enable users to interact with voice recognition systems (e.g., Siri® navigator), through voice commands.

[0217] User interface input devices may also include, without limitation, three dimensional (3D) mice, joysticks or pointing sticks, gamepads and graphic tablets, and audio/visual devices such as speakers, digital cameras, digital camcorders, portable media players, webcams, image scanners, fingerprint scanners, barcode reader 3D scanners, 3D printers, laser rangefinders, and eye gaze tracking devices. Additionally, user interface input devices may include, for example, medical imaging input devices such as computed tomography, magnetic resonance imaging, position emission tomography, medical ultrasonography devices. User interface input devices may also include, for example, audio input devices such as MIDI keyboards, digital musical instruments and the like.

[0218] User interface output devices may include a display subsystem, indicator lights, or non-visual displays such as audio output devices, etc. The display subsystem may be a cathode ray tube (CRT), a flat-panel device, such as that using a liquid crystal display (LCD) or plasma display, a projection device, a touch screen, and the like. In general, use of the term "output device" is intended to include all possible types of devices and mechanisms for outputting information from computer system 1500 to a user or other computer. For example, user interface output devices may include, without limitation, a variety of display devices that visually convey text, graphics and audio/video information such as monitors, printers, speakers, headphones, automotive navigation systems, plotters, voice output devices, and modems.

[0219] Computer system 1500 may comprise a storage subsystem 1518 that comprises software elements, shown as being currently located within a system memory 1510. System memory 1510 may store program instructions that are loadable and executable on processing unit 1504, as well as data generated during the execution of these programs.

[0220] Depending on the configuration and type of computer system 1500, system memory 1510 may be volatile (such as random access memory (RAM)) and/or non-volatile (such as read-only memory (ROM), flash memory, etc.) The RAM typically contains data and/or program modules that are immediately accessible to and/or presently being operated and executed by processing unit 1504. In some implementations, system memory 1510 may include multiple different types of memory, such as static random access memory (SRAM) or

dynamic random access memory (DRAM). In some implementations, a basic input/output system (BIOS), containing the basic routines that help to transfer information between elements within computer system 1500, such as during start-up, may typically be stored in the ROM. By way of example, and not limitation, system memory 1510 also illustrates

5 application programs 1512, which may include client applications, Web browsers, mid-tier applications, relational database management systems (RDBMS), etc., program data 1514, and an operating system 1516. By way of example, operating system 1516 may include various versions of Microsoft Windows®, Apple Macintosh®, and/or Linux operating systems, a variety of commercially-available UNIX® or UNIX-like operating systems
10 (including without limitation the variety of GNU/Linux operating systems, the Google Chrome® OS, and the like) and/or mobile operating systems such as iOS, Windows® Phone, Android® OS, BlackBerry® 15 OS, and Palm® OS operating systems.

[0221] Storage subsystem 1518 may also provide a tangible computer-readable storage medium for storing the basic programming and data constructs that provide the functionality
15 of some embodiments. Software (programs, code modules, instructions) that when executed by a processor provide the functionality described above may be stored in storage subsystem 1518. These software modules or instructions may be executed by processing unit 1504. Storage subsystem 1518 may also provide a repository for storing data used in accordance with the present disclosure.

20 [0222] Storage subsystem 1500 may also include a computer-readable storage media reader 1520 that can further be connected to computer-readable storage media 1522. Together and, optionally, in combination with system memory 1510, computer-readable storage media 1522 may comprehensively represent remote, local, fixed, and/or removable storage devices plus storage media for temporarily and/or more permanently containing, storing, transmitting, and
25 retrieving computer-readable information.

[0223] Computer-readable storage media 1522 containing code, or portions of code, can also include any appropriate media known or used in the art, including storage media and communication media, such as but not limited to, volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or
30 transmission of information. This can include tangible computer-readable storage media such as RAM, ROM, electronically erasable programmable ROM (EEPROM), flash memory or other memory technology, CD-ROM, digital versatile disk (DVD), or other optical storage,

magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or other tangible computer readable media. This can also include nontangible computer-readable media, such as data signals, data transmissions, or any other medium which can be used to transmit the desired information and which can be accessed by computing system

5 1500.

[0224] By way of example, computer-readable storage media 1522 may include a hard disk drive that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive that reads from or writes to a removable, nonvolatile magnetic disk, and an optical disk drive that reads from or writes to a removable, nonvolatile optical disk such as a CD
10 ROM, DVD, and Blu-Ray® disk, or other optical media. Computer-readable storage media 1522 may include, but is not limited to, Zip® drives, flash memory cards, universal serial bus (USB) flash drives, secure digital (SD) cards, DVD disks, digital video tape, and the like. Computer-readable storage media 1522 may also include, solid-state drives (SSD) based on non-volatile memory such as flash-memory based SSDs, enterprise flash drives, solid state
15 ROM, and the like, SSDs based on volatile memory such as solid state RAM, dynamic RAM, static RAM, DRAM-based SSDs, magnetoresistive RAM (MRAM) SSDs, and hybrid SSDs that use a combination of DRAM and flash memory based SSDs. The disk drives and their associated computer-readable media may provide non-volatile storage of computer-readable instructions, data structures, program modules, and other data for computer system 1500.

20 [0225] Communications subsystem 1524 provides an interface to other computer systems and networks. Communications subsystem 1524 serves as an interface for receiving data from and transmitting data to other systems from computer system 1500. For example, communications subsystem 1524 may enable computer system 1500 to connect to one or more devices via the Internet. In some embodiments communications subsystem 1524 can
25 include radio frequency (RF) transceiver components for accessing wireless voice and/or data networks (e.g., using cellular telephone technology, advanced data network technology, such as 3G, 4G or EDGE (enhanced data rates for global evolution), WiFi (IEEE 802.11 family standards, or other mobile communication technologies, or any combination thereof), global positioning system (GPS) receiver components, and/or other components. In some
30 embodiments communications subsystem 1524 can provide wired network connectivity (e.g., Ethernet) in addition to or instead of a wireless interface.

[0226] In some embodiments, communications subsystem 1524 may also receive input communication in the form of structured and/or unstructured data feeds 1526, event streams 1528, event updates 1530, and the like on behalf of one or more users who may use computer system 1500.

5 [0227] By way of example, communications subsystem 1524 may be configured to receive data feeds 1526 in real-time from users of social networks and/or other communication services such as Twitter® feeds, Facebook® updates, web feeds such as Rich Site Summary (RSS) feeds, and/or real-time updates from one or more third party information sources.

10 [0228] Additionally, communications subsystem 1524 may also be configured to receive data in the form of continuous data streams, which may include event streams 1528 of real-time events and/or event updates 1530, that may be continuous or unbounded in nature with no explicit end. Examples of applications that generate continuous data may include, for example, sensor data applications, financial tickers, network performance measuring tools (e.g. network monitoring and traffic management applications), clickstream analysis tools,
15 automobile traffic monitoring, and the like.

[0229] Communications subsystem 1524 may also be configured to output the structured and/or unstructured data feeds 1526, event streams 1528, event updates 1530, and the like to one or more databases that may be in communication with one or more streaming data source computers coupled to computer system 1500.

20 [0230] Computer system 1500 can be one of various types, including a handheld portable device (e.g., an iPhone® cellular phone, an iPad® computing tablet, a PDA), a wearable device (e.g., a Google Glass® head mounted display), a PC, a workstation, a mainframe, a kiosk, a server rack, or any other data processing system.

[0231] Due to the ever-changing nature of computers and networks, the description of
25 computer system 1500 depicted in the figure is intended only as a specific example. Many other configurations having more or fewer components than the system depicted in the figure are possible. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, firmware, software (including applets), or a combination. Further, connection to other computing devices, such as network input/output
30 devices, may be employed. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

[0232] Embodiments may be implemented by using a computer program product, comprising computer program instructions which, when executed by a processor, cause the processor to perform any of the methods described in the disclosure.

[0233] Although specific embodiments have been described, various modifications, alterations, alternative constructions, and equivalents are also encompassed within the scope of the disclosure. Embodiments are not restricted to operation within certain specific data processing environments, but are free to operate within a plurality of data processing environments. Additionally, although embodiments have been described using a particular series of transactions and steps, it should be apparent to those skilled in the art that the scope of the present disclosure is not limited to the described series of transactions and steps. Various features and aspects of the above-described embodiments may be used individually or jointly.

[0234] Further, while embodiments have been described using a particular combination of hardware and software, it should be recognized that other combinations of hardware and software are also within the scope of the present disclosure. Embodiments may be implemented only in hardware, or only in software, or using combinations thereof. The various processes described herein can be implemented on the same processor or different processors in any combination. Accordingly, where components or modules are described as being configured to perform certain operations, such configuration can be accomplished, e.g., by designing electronic circuits to perform the operation, by programming programmable electronic circuits (such as microprocessors) to perform the operation, or any combination thereof. Processes can communicate using a variety of techniques including but not limited to conventional techniques for inter process communication, and different pairs of processes may use different techniques, or the same pair of processes may use different techniques at different times.

[0235] The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that additions, subtractions, deletions, and other modifications and changes may be made thereunto without departing from the broader spirit and scope as set forth in the claims. Thus, although specific disclosure embodiments have been described, these are not intended to be limiting. Various modifications and equivalents are within the scope of the following claims.

[0236] The use of the terms “a” and “an” and “the” and similar referents in the context of describing the disclosed embodiments (especially in the context of the following claims) are to be construed to cover both the singular and the plural, unless otherwise indicated herein or clearly contradicted by context. The terms “comprising,” “having,” “including,” and “containing” are to be construed as open-ended terms (i.e., meaning “including, but not limited to,”) unless otherwise noted. The term “connected” is to be construed as partly or wholly contained within, attached to, or joined together, even if there is something intervening. Recitation of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within the range, unless otherwise indicated herein and each separate value is incorporated into the specification as if it were individually recited herein. All methods described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. The use of any and all examples, or exemplary language (e.g., “such as”) provided herein, is intended merely to better illuminate embodiments and does not pose a limitation on the scope of the disclosure unless otherwise claimed. No language in the specification should be construed as indicating any non-claimed element as essential to the practice of the disclosure.

[0237] Disjunctive language such as the phrase “at least one of X, Y, or Z,” unless specifically stated otherwise, is intended to be understood within the context as used in general to present that an item, term, etc., may be either X, Y, or Z, or any combination thereof (e.g., X, Y, and/or Z). Thus, such disjunctive language is not generally intended to, and should not, imply that certain embodiments require at least one of X, at least one of Y, or at least one of Z to each be present.

[0238] Preferred embodiments of this disclosure are described herein, including the best mode known for carrying out the disclosure. Variations of those preferred embodiments may become apparent to those of ordinary skill in the art upon reading the foregoing description. Those of ordinary skill should be able to employ such variations as appropriate and the disclosure may be practiced otherwise than as specifically described herein. Accordingly, this disclosure includes all modifications and equivalents of the subject matter recited in the claims appended hereto as permitted by applicable law. Moreover, any combination of the above-described elements in all possible variations thereof is encompassed by the disclosure unless otherwise indicated herein.

[0239] All references, including publications, patent applications, and patents, cited herein are hereby incorporated by reference to the same extent as if each reference were individually and specifically indicated to be incorporated by reference and were set forth in its entirety herein.

- 5 [0240] In the foregoing specification, aspects of the disclosure are described with reference to specific embodiments thereof, but those skilled in the art will recognize that the disclosure is not limited thereto. Various features and aspects of the above-described disclosure may be used individually or jointly. Further, embodiments can be utilized in any number of environments and applications beyond those described herein without departing from the
- 10 broader spirit and scope of the specification. The specification and drawings are, accordingly, to be regarded as illustrative rather than restrictive.

WHAT IS CLAIMED IS:

1 1. A method comprising:
2 providing, by a secure network connectivity system implemented in a cloud
3 service provider, a secure network connectivity service for enabling secure private network
4 connectivity between an on-premise network associated with a customer of the cloud service
5 provider and a virtual cloud network (VCN) hosted by the cloud service provider for the
6 customer, the secure network connectivity system comprising a virtual overlay network
7 comprising a set of one or more computing nodes;
8 registering, by the secure network connectivity system, an external resource
9 residing in the on-premise network as an external endpoint in the virtual cloud network, the
10 external endpoint identified by an Internet Protocol (IP) address in the virtual cloud network;
11 creating, by a computing node in the set of one or more computing nodes in
12 the secure network connectivity system, an external resource representation for the external
13 endpoint in the virtual cloud network, creating the external resource representation
14 comprising:
15 creating, by the computing node in the set of one or more computing
16 nodes in the secure network connectivity system, a virtual network interface card
17 (VNIC); and
18 assigning, by the computing node in the set of one or more computing
19 nodes in the secure network connectivity system, the Internet Protocol (IP) address
20 associated with the external endpoint to the virtual network interface card (VNIC);
21 receiving, by the computing node in the secure network connectivity system, a
22 request for querying information stored in the external resource residing in the on-premise
23 network associated with the customer;
24 establishing, by the computing node in the secure network connectivity
25 system, a connection between the virtual network interface card created for the external
26 resource representation in the virtual cloud network and the external resource residing in the
27 on-premise network;
28 transmitting, by the computing node in the secure network connectivity
29 system, the request to the external resource via the established connection; and
30 obtaining, by the computing node in the secure network connectivity system, a
31 result corresponding to the request via the established connection.

1 2. The method of claim 1, further comprising:
2 providing, by the secure network connectivity system, an external gateway
3 appliance, the external gateway appliance configured to establish secure private network
4 connectivity between the external resource residing in the on-premise network and the set of
5 one or more computing nodes comprising the secure network connectivity system.

1 3. The method of claim 2, wherein the external gateway appliance is
2 configured in the on-premise network associated with the customer by a user of the secure
3 network connectivity system.

1 4. The method of claim 2 or claim 3, wherein the external gateway
2 appliance is configured to establish a secure Virtual Private Network (VPN) connection
3 between the external resource residing in the on-premise network and the set of one or more
4 computing nodes comprising the secure network connectivity system.

1 5. The method of any of claims 1 to 4 , wherein transmitting, by the
2 computing node, the request to the external resource residing in the on-premise network via
3 the established connection comprises:

4 translating, by the computing node, the IP address associated with the virtual
5 network interface card to a physical IP address associated with the external resource in the
6 on-premise network; and

7 transmitting, by the computing node, the request to the physical IP address
8 associated with the external resource.

1 6. The method of any of claims 2-4, wherein establishing, by the
2 computing node, the connection to the external resource residing in the on-premise network
3 comprises establishing, by the computing node, the connection to the external gateway
4 appliance via the virtual network interface card.

1 7. The method of any of claims 1 to 6, further comprising enabling, by
2 the secure network connectivity system, a creation of an external site representation of the
3 on-premise network associated with the customer, wherein the external site representation is
4 a logical representation of the on-premise network and identified by an external site identifier
5 and a customer identifier.

1 8. The method of claim 7, wherein the external resource is registered in
2 the external site representation.

1 9. The method of claim 7 or 8, wherein the computing node is configured
2 to establish the connection between the virtual network interface card created for the external
3 resource representation in the virtual cloud network and the external resource residing in the
4 external site representation.

1 10. The method of any of claims 1 to 9, wherein the virtual network
2 interface card is identified by a fully qualified domain name and a cloud identifier in the
3 virtual cloud network associated with the customer.

1 11. The method of any of claims 1 to 10, wherein the external resource is a
2 database, an application, or a compute instance residing in the on-premise network.

1 12. A secure network connectivity system implemented in a cloud service
2 provider for enabling secure private network connectivity between an on-premise network
3 associated with a customer of the cloud service provider and a virtual cloud network (VCN)
4 hosted by the cloud service provider for the customer, the secure network connectivity system
5 comprising a virtual overlay network comprising a set of one or more computing nodes, a
6 computing node in the set of computing nodes comprising:

7 a memory; and

8 one or more processors configured to perform processing, the processing
9 comprising:

10 registering, by the secure network connectivity system, an external resource
11 residing in the on-premise network as an external endpoint in the virtual cloud network, the
12 external endpoint identified by an Internet Protocol (IP) address in the virtual cloud network;

13 creating, by a computing node in the set of one or more computing nodes in
14 the secure network connectivity system, an external resource representation for the external
15 endpoint in the virtual cloud network, creating the external resource representation
16 comprising:

17 creating, by the computing node in the set of one or more computing
18 nodes in the secure network connectivity system, a virtual network interface card
19 (VNIC); and

20 assigning, by the computing node in the set of one or more computing
21 nodes in the secure network connectivity system, the Internet Protocol (IP) address
22 associated with the external endpoint to the virtual network interface card (VNIC);
23 receiving a request for querying information stored in the external resource
24 residing in the on-premise network associated with the customer;
25 establishing a connection between the virtual network interface card created
26 for the external resource representation in the virtual cloud network and the external resource
27 residing in the on-premise network;
28 transmitting the request to the external resource via the established connection;
29 and
30 obtaining a result corresponding to the request, via the established connection.

1 13. The system of claim 12, further comprising:
2 providing an external gateway appliance, the external gateway appliance
3 configured to establish secure private network connectivity between the external resource
4 residing in the on-premise network and the set of one or more computing nodes comprising
5 the secure network connectivity system, wherein the external gateway appliance is configured
6 in the on-premise network associated with the customer by a user of the secure network
7 connectivity system.

1 14. The system of claim 13, wherein the external gateway appliance is
2 configured to establish a secure Virtual Private Network (VPN) connection between the
3 external resource residing in the on-premise network and the set of one or more computing
4 nodes comprising the secure network connectivity system.

1 15. The system of any of claims 12 to 14, wherein transmitting the request
2 to the external resource residing in the on-premise network via the established connection
3 comprises:
4 translating the IP address associated with the virtual network interface card to
5 a physical IP address associated with the external resource in the on-premise network; and
6 transmitting the request to the physical IP address associated with the external
7 resource.

1 16. The system of any of claims 13-14, wherein establishing the
2 connection to the external resource residing in the on-premise network comprises establishing
3 the connection to the external gateway appliance via the virtual network interface card.

1 17. A non-transitory computer-readable medium having program code that
2 is stored thereon, the program code executable by one or more processing devices for
3 performing operations comprising:

4 registering, by a secure network connectivity system, an external resource
5 residing in the on-premise network as an external endpoint in the virtual cloud network, the
6 external endpoint identified by an Internet Protocol (IP) address in the virtual cloud network;

7 creating, by a computing node in the set of one or more computing nodes in
8 the secure network connectivity system, an external resource representation for the external
9 endpoint in the virtual cloud network, creating the external resource representation
10 comprising:

11 creating, by the computing node in the set of one or more computing
12 nodes in the secure network connectivity system, a virtual network interface card
13 (VNIC); and

14 assigning, by the computing node in the set of one or more computing
15 nodes in the secure network connectivity system, the Internet Protocol (IP) address
16 associated with the external endpoint to the virtual network interface card (VNIC);

17 receiving a request for querying information stored in the external resource
18 residing in the on-premise network associated with the customer;

19 establishing a connection between the virtual network interface card created
20 for the external resource representation in the virtual cloud network and the external resource
21 residing in the on-premise network;

22 transmitting the request to the external resource via the established connection;
23 and

24 obtaining a result corresponding to the request, via the established connection.

1 18. The non-transitory computer-readable medium of claim 17 further
2 comprising enabling a creation of an external site representation of the on-premise network
3 associated with the customer, wherein the external site representation is a logical
4 representation of the on-premise network and identified by an external site identifier and a
5 customer identifier.

1 19. The non-transitory computer-readable medium of claim 17 or claim 18,
2 wherein the virtual network interface card is identified by a fully qualified domain name and
3 a cloud identifier in the virtual cloud network associated with the customer.

1 20. The non-transitory computer-readable medium of any of claims 17 to
2 19, wherein the external resource is a database, an application, or a compute instance
3 residing in the on-premise network.

1 21. A computer program product, comprising computer program
2 instructions which, when executed by a processor, cause the processor to:
3 register an external resource residing in an on-premise network as an external
4 endpoint a virtual cloud network, the external endpoint identified by an Internet Protocol (IP)
5 address in the virtual cloud network;

6 create an external resource representation for the external endpoint in the
7 virtual cloud network, the operations to create the external resource representation
8 comprising operations to:

9 create a virtual network interface card (VNIC); and
10 assign the Internet Protocol (IP) address associated with the external
11 endpoint to the virtual network interface card (VNIC);

12 receive a request for querying information stored in the external resource
13 residing in the on-premise network associated with the customer;

14 establish a connection between the virtual network interface card created for
15 the external resource representation in the virtual cloud network and the external resource
16 residing in the on-premise network;

17 transmit the request to the external resource via the established connection;

18 and

19 obtain a result corresponding to the request, via the established connection.





FIG. 1



300




FIG. 3




FIG. 4




FIG. 5






FIG. 6



**FIG. 8**




FIG. 9






FIG. 10









Fig. 14






FIG. 15

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2022/034751

A. CLASSIFICATION OF SUBJECT MATTER INV. H04L9/40 H04L45/00 G06F9/455 ADD.		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED Minimum documentation searched (classification system followed by classification symbols) H04L 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, WPI Data		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2020/382471 A1 (JANAKIRAMAN RAJAGOPALAN [US] ET AL) 3 December 2020 (2020-12-03) paragraph [0001] - paragraph [0075]; figures 1-4 -----	1-21
X	US 2018/062920 A1 (SRINIVASAN ARAVIND [US] ET AL) 1 March 2018 (2018-03-01) paragraph [0002] - paragraph [0046] -----	1-21
X	US 2019/222559 A1 (WANG NAN [CN] ET AL) 18 July 2019 (2019-07-18) paragraph [0001] - paragraph [0031] -----	1, 12, 17, 21
☐ Further documents are listed in the continuation of Box C. ☒ 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 15 September 2022		Date of mailing of the international search report 23/09/2022
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 Veshi, Erzim

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2022/034751

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2020382471 A1	03-12-2020	NONE	
US 2018062920 A1	01-03-2018	US 2018062920 A1	01-03-2018
		US 2018063074 A1	01-03-2018
		US 2018063077 A1	01-03-2018
		US 2018063743 A1	01-03-2018
		US 2018069924 A1	08-03-2018
US 2019222559 A1	18-07-2019	NONE	