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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
26 March 2020 (26.03.2020)

WIPO | PCT

(10) International Publication Number
WO 2020/060625 A1

(51) International Patent Classification:
H04L 29/12 (2006.01)

(21) International Application Number:
PCT/US2019/039097

(22) International Filing Date:
26 June 2019 (26.06.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
201841035632 21 September 2018 (21.09.2018) IN
16/262,626 30 January 2019 (30.01.2019) US

(71) Applicant: **MICROSOFT TECHNOLOGY LICENSING, LLC** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: **MITTAL, Sumeet**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **SHUKLA, Abhishek**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **TEWARI, Rishabh**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **CHEN, Qiming**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **CHANDRAPPA, Harish Kumar**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **SHRIVASTAVA, Pranjal**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **ADUSUMILLI, Anitha**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **SHARMA, Parag**; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US). **SREENATH, Abhishek Ellore**; Microsoft Technol-

(54) **Title:** MAPPING A SERVICE INTO A VIRTUAL NETWORK USING SOURCE NETWORK ADDRESS TRANSLATION

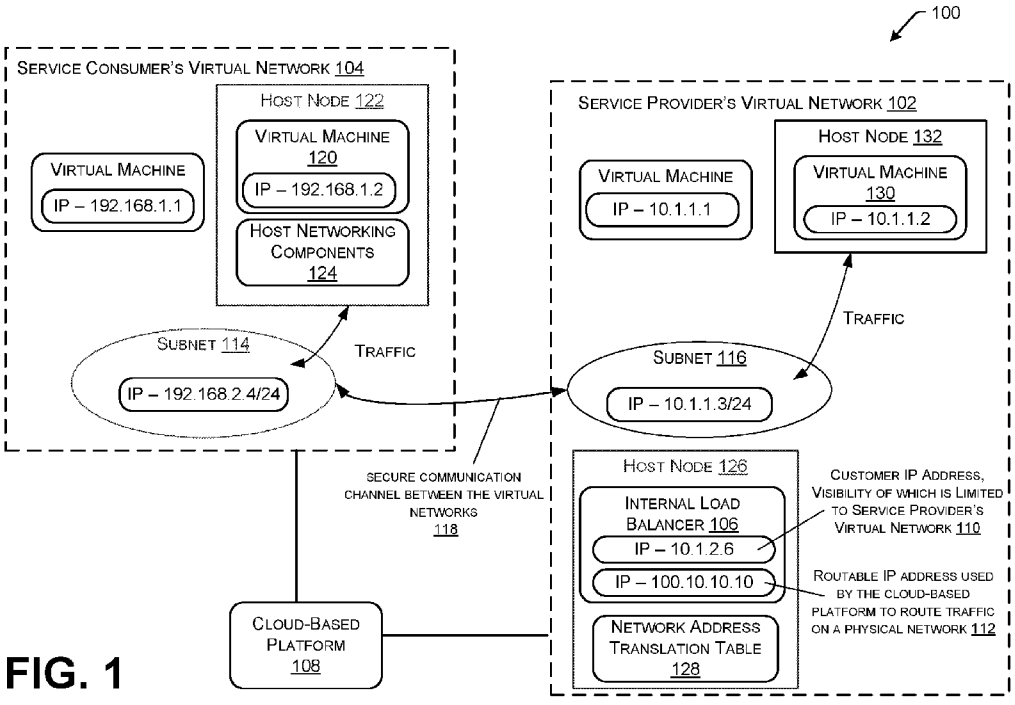


FIG. 1

(57) **Abstract:** The techniques described herein enable a private connectivity solution between a virtual network of a service consumer and a virtual network of a service provider in a cloud-based platform. The techniques map a service (e.g., one or more workloads or containers) executing in the virtual network of the service provider into the virtual network of the service consumer. The mapping uses network address translation (NAT) that is performed by the cloud-based infrastructure. As a result of the techniques described herein, a public Internet Protocol (IP) address does not need to be used to establish a connection thereby alleviating privacy and/or security concerns for the virtual networks of the service provider and/or the service consumer that are hosted by the cloud-based platform.

WO 2020/060625 A1

[Continued on next page]

ogy Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(74) **Agent: MINHAS, Sandip S.** et al.; Microsoft Technology Licensing, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

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

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

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

MAPPING A SERVICE INTO A VIRTUAL NETWORK USING SOURCE NETWORK ADDRESS TRANSLATION

BACKGROUND

[0001] Cloud-based platforms (e.g., AMAZON WEB SERVICES, MICROSOFT AZURE, etc.) enable tenants (e.g., customers) to offer and/or to consume services via their virtual networks. To do this, a virtual network of a service provider, that is hosted by a cloud-based platform, needs to connect to a virtual network of a service consumer that is also hosted by the cloud-based platform. Typically, this connection is established using a public Internet Protocol (IP) address, which causes privacy and/or security concerns for the virtual networks of the service provider and/or the service consumer that are hosted by the cloud-based platform.

[0002] It is with respect to these and other technical challenges that the disclosure made herein is presented.

SUMMARY

[0003] The techniques described herein enable a private connectivity solution between a virtual network of a service consumer and a virtual network of a service provider in a cloud-based platform (e.g., AMAZON WEB SERVICES, MICROSOFT AZURE, etc.). The techniques map a service (e.g., one or more workloads or containers) executing in the virtual network of the service provider into a virtual network of the service consumer. The mapping uses network address translation (NAT) that is performed by the cloud-based infrastructure. As a result of the techniques described herein, a public Internet Protocol (IP) address does not need to be used to establish a connection thereby alleviating the privacy and/or the security concerns for the virtual networks of the service provider and/or the service consumer that are hosted by the cloud-based platform.

[0004] An internal load balancer is configured in association with the virtual network of the service provider. For instance, an entity (e.g., a customer) that provides the service can configure an internal load balancer to scale a service across multiple virtual machines operating within its virtual network and/or to distribute workloads across multiple resources. Thus, the load balancing optimizes resource use, maximizes throughput, minimizes response time, and avoids overload of any single resource. The internal load balancer described herein is assigned a routable IP address, which is a physical IP address. The routable IP address is used by networking components of the cloud-based platform to route traffic on a network of the cloud-based platform. The internal load balancer is also assigned

a customer IP address. The customer IP address is a private virtual IP address that is only visible to virtual machines within the virtual network of the service provider.

[0005] In order to map the service executing in the virtual network of the service provider into the virtual network of the service consumer, subnets (e.g., a collection of IP addresses) are created. That is, the service provider creates and allocates a subnet in its virtual network and the service consumer creates and allocates a subnet in its virtual network. The subnets represent endpoints to a secure communication channel between the virtual networks. Consequently, the virtual machines in the virtual network of the service consumer believe they are sending data packets to, and/or receiving data packets from, an endpoint in their own virtual network when consuming the service. Moreover, the virtual machines in the virtual network of the service provider believe they are sending data packets to, and/or receiving data packets from, an endpoint in their own virtual network when providing the service.

[0006] After the subnets are created, the techniques described herein determine that a data packet is being communicated from a source IP address associated with a virtual machine in the virtual network of the service consumer to a destination IP address associated with the subnet that has been created in the virtual network of the service consumer. Accordingly, it appears, to the virtual machine, that the data packet is being communicated within the same virtual network (i.e., the virtual network of the service consumer). This determination that a data packet is being communicated can be made by networking components of a physical machine (e.g., a server) of the cloud-based platform that is supporting the virtual network of the service consumer (e.g., the virtual machine where the data packet originates). A physical machine that is part of the cloud-based platform may be referred to herein as a "host node".

[0007] The host node that determines that a data packet is being communicated is aware of the routable IP address of the internal load balancer configured in association with the virtual network of the service provider. For instance, networking components of the host node are informed of a policy (e.g., a condition match/action rule) to translate an IP address associated with the subnet in the virtual network of the service consumer to the routable IP address of the internal load balancer configured in association with the virtual network of the service provider. In various examples, the policy, and therefore, the routable IP address, can be accessed via a publication subscription component of the cloud-based platform. Example networking components of a host node that implement the techniques described herein can include one or more of a host networking agent (e.g., NMAgent), a

virtual switch (e.g., vSwitch), the Virtual Filtering Platform ("VFP"), etc.

[0008] The host node is configured to replace the destination IP address in the data packet with the routable IP address of the internal load balancer. The host node can then communicate the data packet over the physical network of the cloud-based platform via the routable IP address of the internal load balancer. That is, the host node can communicate the data packet to another host node of the cloud-based platform that supports the internal load balancer configured in association with the virtual network of the service provider. To this end, the other host node that supports the internal load balancer receives the data packet.

[0009] The communication of the data packet using the routable IP address allows the internal load balancer to implement its load balancing functionality. For instance, the internal load balancer can select a virtual machine in the virtual network of the service provider to receive the data packet (e.g., to handle a request). The host node that supports the internal load balancer is configured to translate, using a network address translation (NAT) table, the IP addresses in the communicated data packet into IP addresses associated with the virtual network of the service provider. For example, the host node translates the original source IP address in the data packet into a source IP address associated with the subnet created in the virtual network of the service provider. In this way, when the data packet is communicated to a destination within the virtual network of the service provider, it appears that the data packet originates within the same virtual network (i.e., the virtual network of the service provider). Furthermore, the host node translates the original source port number to a source port number selected from a pool of port numbers useable in the virtual network of the service provider. This addresses a situation where two different virtual machines of one or multiple service consumers initiate a connection with a same source port number. Given this situation, the same source port number used by the two different virtual machines is translated into two unique port numbers in the virtual network of the service provider.

[0010] The host node supporting the internal load balancer also translates the routable IP address of the internal load balancer in the data packet into the customer IP address of the internal load balancer. Then, the customer IP address is translated into a destination IP address (e.g., a "final" destination IP address) associated with the virtual machine in the virtual network of the service provider that has been selected to receive the data packet. In some instances, the internal load balancer selects a destination port number associated with the virtual machine in the virtual network of the service provider and the host node supporting the internal load balancer translates a an original destination port number in the

data packet into the destination port number associated with the virtual machine in the virtual network of the service provider. The data packet can then be communicated to the virtual machine in the virtual network of the service provider using the final destination IP address and/or the destination port number.

[0011] In various embodiments, a host node of the cloud-based platform can perform network functions in association with communicating the data packet from the virtual network of the service consumer to the virtual network of the service provider, and vice versa. For example, the host nodes can perform encapsulation/de-encapsulation ("encap/decap") functionality. That is, the host node supporting the virtual network of the service consumer (e.g., the virtual machine where a data packet originates) can encapsulate the data packet with a header and the host node supporting the virtual network of the service provider (e.g., the internal load balancer) can de-encapsulate the data packet.

[0012] In another example, the IP addresses of the original data packet generated by the virtual machine of the service consumer may be Internet Protocol version 4 (IPv4) addresses, and a host node supporting the virtual network of the service consumer can convert the IPv4 addresses into Internet Protocol version 6 (IPv6) addresses.

[0013] In a further example, the host node supporting the virtual network of the service consumer can embed a general routing encapsulation key (e.g., a GRE Key) into the data packet. The general routing encapsulation key uniquely identifies the virtual network of the service consumer. Consequently, the general routing encapsulation key can be used to ensure a reply data packet is sent back to the correct virtual network. In some instances, the embedding of the general routing encapsulation key can occur with the conversion of the IPv4 addresses into the IPv6 addresses.

[0014] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter. The term "techniques," for instance, may refer to system(s), method(s), computer-readable instructions, module(s), component(s), algorithm(s), hardware logic, and/or operation(s) as permitted by the context described above and throughout the document.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] The detailed description is described with reference to the accompanying figures. In the figures, the left-most digit(s) of a reference number identifies the figure in which the reference number first appears. The same reference numbers in different figures

indicate similar or identical items.

[0016] FIG. 1 illustrates an example environment in which a service (e.g., one or more workloads or containers) executing in a virtual network of a service provider is mapped into a virtual network of a service consumer.

[0017] FIG. 2 illustrates example translations that occur in a data packet to implement the techniques described herein.

[0018] FIG. 3 is a flowchart that illustrates operations for configuring subnets in virtual networks of a service provider and a service consumer so that a service can be mapped from the virtual network of the service provider to the virtual network of the service consumer.

[0019] FIG. 4 is a flowchart that illustrates operations for mapping a service executing in a virtual network of a service provider into a virtual network of a service consumer.

[0020] FIG. 5 is a block diagram of an example computing system usable to implement various aspects of the present disclosure.

DETAILED DESCRIPTION

[0021] The techniques described herein enable a private connectivity solution between a virtual network (e.g., a VNET) of a service consumer and a virtual network of a service provider in a cloud-based platform (e.g., AMAZON WEB SERVICES, MICROSOFT AZURE, etc.). The techniques map a service (e.g., one or more workloads or containers) executing in the virtual network of the service provider into a virtual network of the service consumer. The mapping uses network address translation (NAT) that is performed by the cloud-based infrastructure. As a result of the techniques described herein, a public Internet Protocol (IP) address does not need to be used to establish a connection thereby alleviating the privacy and/or the security concerns for the virtual networks of the service provider and/or the service consumer that are hosted by the cloud-based platform.

[0022] A cloud-based platform may host virtual networks for a number of different tenants, or customers. In various examples, a virtual network is a representation of a tenant's own network in a cloud-based platform. In some instances, the virtual network can be an extension of the tenant's own private network (e.g., an on-premises datacenter network). The virtual network provides an isolated environment for a tenant to run a workload. A virtual network can implement one or more virtual machines to run the workload. In various examples, a virtual machine provides functionality to execute an operating system. A hypervisor in the cloud-based platform can use the functionality to share and manage hardware, allowing for multiple different isolated environments to exist on a same physical machine (e.g., a server that is part of the cloud-based platform).

[0023] A container is an isolated and resource-controlled operating environment. A software application (e.g., a workload) can run in a container without affecting the rest of the system and without the system affecting the application. A container may be specific to one application, or may provide the runtime for multiple applications, providing additional operating system functions and services (e.g. input/output functions, security, etc.). Moreover, a container can be portable and/or mobile. Containers introduce additional levels of abstraction to a network (e.g., a datacenter network operated by a cloud-based platform), thereby increasing the complexity of the networking infrastructure. In a specific example, a host node may have a physical IP address and a virtual machine executing within the host node may have a virtual IP address. Network Virtualization Functions (NVF) in the host node may be used to encapsulate and/or de-encapsulate ("encap/decap") packets and/or to translate addresses from one layer of abstraction to another and/or between physical and virtual layers.

[0024] Accordingly, the techniques described herein use networking aspects of a cloud-based platform to map a service (e.g., a workload, a container, etc.) executing in a virtual network of a tenant that provides the service into a virtual network of a tenant that consumes the service. The mapping uses network address translation (NAT) that is performed by the cloud-based infrastructure that supports the virtual network of the service provider. Stated another way, the virtual network of the service provider is the "source" of the network address translation.

[0025] FIG. 1 illustrates an example environment 100 in which a service (e.g., one or more workloads or containers) executing in a virtual network of a service provider 102 is mapped into a virtual network of a service consumer 104. As described above, an internal load balancer 106 is configured in association with the virtual network of the service provider 102. The internal load balancer 106 is configured to scale the service across a number of virtual machines (e.g., two, five, ten, twenty, etc.) that are capable of executing the service within the virtual network of the service provider 102, such that the service can be exposed to multiple different service consumers. The infrastructure of a cloud-based platform 108 isolates different virtual networks of the multiple different service consumers thereby enabling the different service consumers to use the same or overlapping IP address space.

[0026] The internal load balancer 106 is assigned a customer IP address (CA) (e.g., "10.1.2.6" as shown in FIG. 1). The load balancer is "internal" because the customer IP address is only visible within its own virtual network, as referenced by 110. Further, the

internal load balancer 106 is assigned a routable IP address (e.g., "100.10.10.10" as shown in FIG. 1). The routable IP address is an actual address known to, and used by, the cloud-based platform 108 to route data traffic on a physical network, as referenced by 112.

[0027] To implement the mapping, a subnet 114 (e.g., "192.168.2.4/24" as shown in FIG. 1) is created in the virtual network of the service consumer 104 and a corresponding subnet 116 (e.g., "10.1.1.3/24" as shown in FIG. 1) is created in the virtual network of the service provider 102. These subnets are created so that the virtual machines executing in the respective virtual networks think they are communicating within their own virtual network when providing or consuming the service. Consequently, the subnets 114 and 116 represent endpoints to a secure communication channel between the virtual networks, as referenced by 118.

[0028] When one of multiple virtual machines (e.g., virtual machine 120) in the virtual network of the service consumer 104 sends a data packet, referred to as "traffic" in FIG. 1, via a source IP address (e.g., "192.168.1.2" as shown in FIG. 1) to a destination IP address (e.g., "192.168.2.4") within the subnet 114 allocated to the service being consumed by the virtual network 104, the underlying infrastructure implemented in a host node 122 of the cloud-based platform 108 associates the destination IP address (e.g., "192.168.2.4") with the internal load balancer 106. The underlying infrastructure comprises host networking components 124 such as a host networking agent (e.g., NMAgent), a virtual switch (e.g., vSwitch), the Virtual Filtering Platform ("VFP"), etc. As described herein, operations performed "by a host node" may be operations performed by the host networking components of a host node.

[0029] The host node 122 replaces the destination IP address (e.g., "192.168.2.4") in the data packet with the routable IP address of the internal load balancer (e.g., "100.10.10.10"). In various examples, the host node 122 is aware of the routable IP address of the internal load balancer 106 because the host node 122 has been instructed to apply a policy (e.g., a condition match/action rule) to translate an IP address associated with the subnet 114 in the virtual network of the service consumer 104 into the routable IP address of the internal load balancer 106 configured within the virtual network of the service provider 102. In various examples, the policy, and therefore, the routable IP address, can be accessed via a publication subscription component of the cloud-based platform 108.

[0030] The host node 122 then communicates the data packet to the internal load balancer 106 using the routable IP address. That is, the host node 122 communicates the data packet to a host node 126 supporting the internal load balancer 106. Consequently, the

host node 126 receives the data packet and the internal load balancer 106 is enabled to perform its load balancing functionality (e.g., select a destination for the data packet). Subsequently, the networking components of the host node 126 are configured to use a network address translation (NAT) table 128 to translate the source IP address (e.g., "192.168.1.2") in the data packet to a source IP address associated with the subnet 116 created in the virtual network of the service provider 102 (e.g., "10.1.1.3" as shown in FIG. 1). In this way, when the data packet is communicated to a destination within the virtual network of the service provider 102, it appears that the data packet originates within the same virtual network (i.e., the virtual network of the service provider).

[0031] In various examples, the host node 126 translates an original source port number to a source port number selected from a pool of port numbers useable in the virtual network of the service provider 102. This addresses a situation where two different virtual machines of one or multiple service consumers initiate a connection using a same source port number. Given this situation, the same source port number used by the two different virtual machines is translated into two unique "source" port numbers useable in the virtual network of the service provider 102.

[0032] The host node 126 also uses the NAT table 128 to translate the routable IP address (e.g., "100.10.10.10") in the data packet into an IP address within the virtual network of the service provider 102. For example, the routable IP address (e.g., "100.10.10.10") of the internal load balancer 106 may first be translated into the customer IP address (e.g., "10.1.2.6") of the internal load balancer 106. In this example, the customer IP address (e.g., "10.1.2.6") can subsequently be translated into a destination IP address (e.g., "10.1.1.2" as shown in FIG. 1) associated with the virtual machine 130 selected by the internal load balancer. The host node 126 can then communicate the data packet to the virtual machine 130 via a supporting host node 132. In some instances, the internal load balancer selects a destination port number associated with the virtual machine 130 in the virtual network of the service provider 102 and translates an original destination port number in the data packet into the destination port number associated with the virtual machine 130 in the virtual network of the service provider 102. In association with the virtualization techniques implemented in the cloud-based platforms, any two or more of host nodes 122, 126, and 132 may be the same.

[0033] In various embodiments, a host node of the cloud-based platform can perform network functions in association with communicating the data packet from the virtual network of the service consumer 104 to the virtual network of the service provider 102, and

vice versa. For example, the host nodes can perform encapsulation/de-encapsulation ("encap/decap") functionality. That is, the host node 122 supporting the virtual network of the service consumer 104 can encapsulate the data packet with a header and the host node 126 supporting the virtual network of the service provider 102 can de-encapsulate the data packet.

[0034] In another example, the IP addresses of the original data packet generated by the virtual machine 120 of the service consumer 104 may be Internet Protocol version 4 (IPv4) addresses, and a host node 122 supporting the virtual network of the service consumer 104 can convert the IPv4 addresses into Internet Protocol version 6 (IPv6) addresses prior to communicating the data packet over a network of the cloud-based platform 108.

[0035] In a further example, the host node 122 supporting the virtual network of the service consumer 104 can embed a general routing encapsulation key (e.g., a GRE Key) into the data packet. The general routing encapsulation key uniquely identifies the virtual network of the service consumer 104, thereby enabling the different service consumers to use the same or overlapping IP address space. Consequently, the general routing encapsulation key can be used to ensure a reply data packet is sent back to the correct virtual network. In some instances, the embedding of the general routing encapsulation key can occur with the conversion of the IPv4 addresses into the IPv6 addresses.

[0036] Once the data packet has been communicated from the virtual machine 120 in the virtual network of the service consumer 104 to the virtual machine 130 in the virtual network of the service provider 102, as described above, similar translations can be used by the cloud-based infrastructure to communicate a reply data packet from the virtual machine 130 in the virtual network of the service provider 102 back to the virtual machine 120 in the virtual network of the service consumer 104. For example, the final destination IP address (e.g., "10.1.1.2") of the originally communicated data packet becomes the initial source IP address of a reply data packet and the original source IP address (e.g., "192.168.1.2") of the originally communicated data packet becomes the final destination IP address of the reply data packet.

[0037] A host node supporting the virtual network of the service provider 102 can use the same NAT table 128 to implement reply translations. For example, the source IP address in the reply data packet (e.g., "10.1.1.2") can be translated into a source IP address (e.g., "192.168.2.4") associated with the subnet 114 created in the virtual network of the service consumer 104. Moreover, the destination IP address (e.g., "10.1.1.3") in the reply data

packet that is associated with the subnet 116 of the virtual network of the service provider 102 can be translated into a destination IP address (e.g., "192.168.1.2") associated with the virtual machine 120 in the virtual network of the service consumer 104. The host node can then communicate the reply data packet, with the translated IP addresses, to the virtual machine 120 in the virtual network of the service consumer 104. In various embodiments, the general routing encapsulation key is embedded in the reply data packet to ensure the reply data packet is routed back to the correct virtual network (i.e., the virtual network of the service consumer 104 from which the originally communicated data packet was received). Because the general routing encapsulation key uniquely identifies the virtual network of the service consumer 104, traffic can be routed back and a private communication channel between the two subnets 114, 116 is established by the cloud-based platform 108 without involving a public IP address.

[0038] FIG. 2 illustrates example translations that occur in a data packet to implement the techniques described herein. Reference 202 illustrates an example of at least part of a data packet that leaves the virtual machine 120 in the virtual network of the service consumer 104. The data packet includes a virtual source IP address in the virtual network of the service consumer 104 (e.g., "192.168.1.2") and a source port number (e.g., "1729"). The data packet further includes a virtual destination IP address in the virtual network of the service consumer 104 (e.g., "192.168.2.4) and a destination port number (e.g., "443").

[0039] Reference 204 illustrates an example of converting IP addresses in the data packet from IPv4 addresses into IPv6 addresses. Moreover, a general routing encapsulation key (e.g., a GRE Key) that uniquely identifies the virtual network of the service consumer 104 is embedded into the IPv6 address.

[0040] Reference 206 illustrates an example of how the host networking components 124 of a host node 122 supporting the virtual machine 120 in the virtual network of the service consumer 104 use a policy (e.g., a condition match/action rule) to replace, as the destination for the data packet, the subnet IP address (e.g., "192.168.2.4") in the virtual network of the service consumer 104 with the routable IP address of the internal load balancer 106 (e.g., "100.10.10.10"). Then, the host networking components 124 can encapsulate the data packet, that includes the IPv6 addresses, with a header (e.g., an IPv4 header), as shown.

[0041] The encapsulated data packet can then be communicated, via the cloud-based platform 108 (e.g., dropped to the wire), to the host node 126 supporting the internal load balancer 106 where the data packet can be de-encapsulated, as referenced by 208. For

instance, the header (e.g., the outer layer) is removed and network address translation is performed, using a NAT table 128, within the virtual network of the service provider 102. In various examples, the NAT table 128 translates the IPv6 source address of the virtual network of the service consumer 104 into an IPv4 source address in the subnet 116 of the virtual network of the service provider 102 (e.g., "10.1.1.3" as shown in FIG. 1). Moreover, the NAT table 128 translates the routable IP address of the internal load balancer 106 (e.g., "100.10.10.10") into the customer IP address of the internal load balancer 106 and subsequently translates the customer IP address of the internal load balancer 106 into a destination IP address associated with a virtual machine 130 (e.g., "10.1.1.2") as the final destination for the data packet in the virtual network of the service provider 102.

[0042] FIGS. 3 and 4 each illustrates an example flowchart. It should be understood by those of ordinary skill in the art that the operations of the method(s) disclosed herein are not necessarily presented in any particular order and that performance of some or all of the operations in alternative order(s) is possible and is contemplated. The operations have been presented in the demonstrated order for ease of description and illustration. Operations may be added, omitted, performed together, and/or performed simultaneously, without departing from the scope of the appended claims.

[0043] It also should be understood that the illustrated method(s) can end at any time and need not be performed in their entirety. Some or all operations of the method(s), and/or substantially equivalent operations, can be performed by execution of computer-readable instructions included on a computer-storage media, as defined herein. The term "computer-readable instructions," and variants thereof, as used in the description and claims, is used expansively herein to include routines, applications, application modules, program modules, programs, components, data structures, algorithms, and the like.

[0044] Thus, it should be appreciated that the logical operations described herein are implemented (1) as a sequence of computer implemented acts or program modules running on a computing system and/or (2) as interconnected machine logic circuits or circuit modules within the computing system. The implementation is a matter of choice dependent on the performance and other requirements of the computing system. Accordingly, the logical operations may be implemented in software, in firmware, in special purpose digital logic, and any combination thereof.

[0045] FIG. 3 is a flowchart 300 that illustrates operations for configuring subnets in virtual networks of a service provider and a service consumer so that a service can be mapped from the virtual network of the service provider into the virtual network of the

service consumer.

[0046] At operation 302, an internal load balancer is configured in the virtual network of the service provider. As described above, a tenant or another party that operates the virtual network may deploy the internal load balancer and assign both a routable IP address and a customer IP address to the internal load balancer. The routable IP address is used by a cloud-based platform to route traffic on a network and the customer IP address is an internal address that is only visible within the virtual network of the service provider.

[0047] At operation 304, a subnet in the virtual network of the service provider is created. The subnet is created so that the service executing in the virtual network of the service provider can be mapped into the virtual network of the service consumer.

[0048] At operation 306, a subnet in the virtual network of the service consumer is created. This subnet is also created so that the service executing in the virtual network of the service provider can be mapped into the virtual network of the service consumer. This subnet can be created by a tenant that wants to consume a service executing in another virtual network.

[0049] As described above, the subnets represent endpoints to a secure communication channel between the virtual networks. Consequently, the virtual machines in the virtual network of the service consumer believe they are sending data packets to, and/or receiving data packets from, an endpoint in their own virtual network when consuming the service. Moreover, the virtual machines in the virtual network of the service provider believe they are sending data packets to, and/or receiving data packets from, an endpoint in their own virtual network when providing the service.

[0050] FIG. 4 is a flowchart 400 that illustrates operations for mapping a service executing in a virtual network of a service provider into a virtual network of a service consumer.

[0051] At operation 402, it is determined that a data packet is to be communicated from a first IP address (e.g., a consumer-side source virtual IP address) associated with a virtual machine in the virtual network of the service consumer to a second IP address (e.g., a consumer-side destination virtual IP address) associated with the subnet created in the virtual network of the service consumer. This determination can be made by a host node of the cloud-based platform that supports the virtual machine.

[0052] At operation 404, the second IP address in the data packet is replaced with the routable IP address (e.g., a physical IP address) of the internal load balancer configured in association with the virtual network of the service provider. This operation can be

performed by the host node that supports the virtual machine. In various examples, the host node is made aware of the routable IP address of the internal load balancer because the host node has been instructed to apply a policy (e.g., a condition match/action rule) to translate an IP address associated with the subnet in the virtual network of the service consumer into the routable IP address of the internal load balancer configured in association with the virtual network of the service provider 102. In various examples, the policy, and therefore, the routable IP address, is pushed down to the host node via a publication service (e.g., a publication subscription component) of the cloud-based platform.

[0053] At operation 406, the data packet is communicated, using the routable IP address, over the physical network of the cloud-based platform to a host node that supports the internal load balancer.

[0054] At operation 408, the data packet is received by the host node that supports the internal load balancer.

[0055] At operation 410, the first IP address in the data packet is translated into a third IP address (e.g., a provider-side source virtual IP address) associated with the subnet in the virtual network of the service provider that has been created to map the service executing in the virtual network of the service provider into the virtual network of the service consumer.

[0056] At operation 412, the routable IP address of the internal load balancer in the data packet is translated into the customer IP address of the internal load balancer.

[0057] At operation 414, the customer IP address of the internal load balancer is translated into a fourth IP address (e.g., a server-side destination virtual IP address) associated with a virtual machine in the virtual network of the service provider. As described above, the virtual machine may be selected, amongst multiple virtual machines, by the internal load balancer in association with a scaling and/or load balancing function.

[0058] At operation 416, the data packet is communicated, using the fourth IP address, to the virtual machine in the virtual network of the service provider.

[0059] FIG. 5 is a block diagram of an example computing system 500 usable to implement various aspects of the present disclosure. The computing system 500 may be deployed in a shared network environment, including in a datacenter, a cloud computing environment, or other network of computing devices. In one example configuration, the computing system 500 comprises at least one processor 502 and computer-readable media 504. The computing system 500 also contains communication connection(s) 506 that allow communications with various other systems. The computing system 500 may also include one or more input devices 508, such as a keyboard, mouse, pen, voice input device, touch

input device, etc., and/or one or more output devices 510, such as a display (including a touch-screen display), speakers, printer, etc. coupled communicatively to the processor(s) 502 and the computer-readable media 504 via connections 512.

[0060] In the illustrated example, computer-readable media 504 stores operating system(s) 514, which provide basic system functionality to virtual machines 516 (which may be the same as or similar to one or more of virtual machines described above) and/or other cloud-based platform components 518 (e.g., networking resource provider, a publication service, a networking switch, a Virtual Filtering Platform, APIs, etc.).

[0061] The operations described herein can be implemented via computer-executable instructions stored on one or more computer storage media that, when executed by one or more processors, enable the one or more processors to perform the operations. Generally, computer-executable instructions include routines, programs, objects, modules, components, data structures, and the like that perform particular functions or implement particular abstract data types.

[0062] According to various examples, the computing systems described herein includes one or more devices, such as servers, storage devices, and other cloud-based platform devices. Computer-readable media 504 stores computer-executable instructions that are loadable and executable by one or more processor(s), as well as data generated during execution of, and/or usable in conjunction with, these programs. In the illustrated example, computer-readable media stores operating system instances, which provide basic system functionality to applications stored thereon. One or more of these components, including the operating systems, may be instantiated as virtual machines, containers, or as some other type of virtualized instantiation.

[0063] Processor(s) 502 may include one or more single-core processing unit(s), multi-core processing unit(s), central processing units (CPUs), graphics processing units (GPUs), general-purpose graphics processing units (GPGPUs), or hardware logic components configured, e.g., via specialized programming from modules or application program interfaces (APIs), to perform functions described herein. In alternative examples one or more functions of the present disclosure may be performed or executed by, and without limitation, hardware logic components including Field-Programmable Gate Arrays (FPGAs), Application-Specific Integrated Circuits (ASICs), Application-Specific Standard Products (ASSPs), System-on-a-Chip Systems (SOCs), Complex Programmable Logic Devices (CPLDs), Digital Signal Processing unit(s) (DSPs), and other types of customized processing unit(s). For example, a processing unit configured to perform one or more of the

functions described herein may represent a hybrid device that includes a CPU core embedded in an FPGA fabric. These or other hardware logic components may operate independently or, in some instances, may be driven by a CPU. In some examples, examples of the computing systems may include a plurality of processing units of multiple types. Different processing units may have different execution models, e.g., as is the case for graphics processing units (GPUs) and central processing units (CPUs).

[0064] Depending on the configuration and type of computing device used, computer-readable media include volatile memory (such as random access memory (RAM)) and/or non-volatile memory (such as read-only memory (ROM), flash memory, 3D XPoint, resistive RAM, etc.). The computer-readable media can also include additional removable storage and/or non-removable storage including, but not limited to, SSD (e.g., flash memory), HDD (Hard Disk Drive) storage or other type of magnetic storage, optical storage, and/or other storage that can provide non-volatile storage of computer-executable instructions, data structures, program modules, and other data for computing systems.

[0065] Computer-readable media can, for example, represent computer memory, which is a form of computer storage media. Computer-readable media includes at least two types of computer-readable media, namely computer storage media and communications media. Computer storage media includes volatile and non-volatile, removable and non-removable media implemented in any process or technology for storage of information such as computer-executable instructions, data structures, programming modules, or other data. Computer storage media includes, but is not limited to, phase change memory (PRAM), resistive RAM, 3D Xpoint non-volatile memory, static random-access memory (SRAM), dynamic random-access memory (DRAM), other types of random-access memory (RAM), read-only memory (ROM), electrically erasable programmable read-only memory (EEPROM), flash memory or other memory technology, compact disk read-only memory (CD-ROM), digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium that can be used to store information for access and retrieval by a computing device. In contrast, communication media can embody computer-executable instructions, data structures, program modules, or other data in a modulated data signal, such as a carrier wave, or other transmission mechanism. As defined herein, computer storage media does not include communication media.

ILLUSTRATIVE CONFIGURATIONS

[0066] The following clauses described multiple possible configurations for

implementing the features described in this disclosure. The various configurations described herein are not limiting nor is every feature from any given configuration required to be present in another configuration. Any two or more of the configurations may be combined together unless the context clearly indicates otherwise. As used herein in this document "or" means and/or. For example, "A or B" means A without B, B without A, or A and B. As used herein, "comprising" means including listed all features and potentially including addition of other features that are not listed.

[0067] The disclosure presented herein also encompasses the subject matter set forth in the following clauses.

[0068] Example Clause A, a method for mapping a service executing in a virtual network of a service provider into a virtual network of a service consumer, the method comprising: creating a subnet, in the virtual network of the service provider, for mapping the service executing in the virtual network of the service provider into the virtual network of the service consumer; receiving a data packet at a host node of a cloud-based platform that supports an internal load balancer associated with the virtual network of the service provider, wherein the data packet is routed to the host node using a routable Internet Protocol (IP) address assigned to the internal load balancer; translating, by the host node, a first source IP address in the data packet into a second source IP address, wherein the first source IP address is associated with a virtual machine in the virtual network of the service consumer and the second source IP address is associated with the subnet; translating, by the host node, the routable IP address in the data packet into a destination IP address associated with a virtual machine in the virtual network of the service provider; and communicating, by the host node using the destination IP address, the data packet to the virtual machine in the virtual network of the service provider, wherein the data packet indicates the second source IP address as a source.

[0069] Example Clause B, the method of Example Clause A, further comprising: creating a subnet, in the virtual network of the service consumer, for mapping the service executing in the virtual network of the service provider into the virtual network of the service consumer; determining, by another host node of the cloud-based platform supporting the virtual network of the service consumer, that the data packet is being communicated from the first source IP address to another destination IP address associated with the subnet in the virtual network of the service consumer; replacing, by the other host node, the other destination IP address in the data packet with the routable IP address; and communicating, by the other host node using the routable IP address, the data packet over a network to the

host node.

[0070] Example Clause C, the method of Example Clause B, wherein the first source IP address and the other destination IP address comprise Internet Protocol version 4 (IPv4) addresses, and the method further comprises converting, by the other host node, the IPv4 addresses into Internet Protocol version 6 (IPv6) addresses.

[0071] Example Clause D, the method of Example Clause B or Example Clause C, further comprising: encapsulating, by the other host node, the data packet; embedding, by the other host node, a general routing encapsulation key that uniquely identifies the virtual network of the service consumer in the data packet; and de-encapsulating, by the host node, the data packet.

[0072] Example Clause E, the method of any one of Example Clauses B through D, further comprising: determining that a reply data packet is communicated from the destination IP address to the second source IP address; translating, by the host node, the destination IP address in the reply data packet into the other destination IP address associated with the subnet in the virtual network of the service consumer; translating, by the host node, the second source IP address in the reply data packet into the first source IP address associated with the virtual machine in the virtual network of the service consumer; and communicating, by the host node using the first source IP address as a destination, the reply data packet to the virtual machine in the virtual network of the service consumer.

[0073] Example Clause F, the method of Example Clause E, wherein translating the second source IP address into the first source IP address uses a general routing encapsulation key to identify the virtual network of the service consumer.

[0074] Example Clause G, the method of any one of Example Clauses A through F, wherein the internal load balancer selects the virtual machine in the virtual network of the service provider as a final destination for the data packet.

[0075] Example Clause H, the method of any one of Example Clauses A through G, wherein translating the routable IP address into the destination IP address associated with the virtual machine in the virtual network of the service provider comprises: translating the routable IP address into a customer IP address associated with the internal load balancer; and subsequently translating the customer IP address into the destination IP address.

[0076] Example Clause I, the method of any one of Example Clauses A through H, further comprising translating, by the host node, an original destination port number in the data packet into a selected destination port number that is selected by the internal load balancer, the selected destination port number being associated with the virtual machine in

the virtual network of the service provider.

[0077] Example Clause J, a system comprising: one or more processors; and computer-readable media storing instructions that, when executed by the one or more processors, cause the system to: create a subnet, in a virtual network of a service provider, for mapping a service into a virtual network of a service consumer; receive a data packet at a host node of a cloud-based platform that supports an internal load balancer associated with the virtual network of the service provider, wherein the data packet is routed to the host node using a routable Internet Protocol (IP) address assigned to the internal load balancer; translate, by the host node, a first source IP address in the data packet into a second source IP address, wherein the first source IP address is associated with a virtual machine in the virtual network of the service consumer and the second source IP address is associated with the subnet; translate, by the host node, the routable IP address in the data packet into a destination IP address associated with a virtual machine in the virtual network of the service provider; and communicate, by the host node using the destination IP address, the data packet to the virtual machine in the virtual network of the service provider, wherein the data packet indicates the second source IP address as a source.

[0078] Example Clause K, the system of Example Clause J, wherein the instructions further cause the system to: create a corresponding subnet, in the virtual network of the service consumer, for mapping the service into the virtual network of the service consumer; determine, by another host node of the cloud-based platform supporting the virtual network of the service consumer, that the data packet is being communicated from the first source IP address to another destination IP address associated with the corresponding subnet; replace, by the other host node, the other destination IP address in the data packet with the routable IP address; and communicate, by the other host node using the routable IP address, the data packet over a network to the host node.

[0079] Example Clause L, the system of Example Clause K, wherein the first source IP address and the other destination IP address comprise Internet Protocol version 4 (IPv4) addresses, and the method further comprises converting, by the other host node, the IPv4 addresses into Internet Protocol version 6 (IPv6) addresses.

[0080] Example Clause M, the system of Example Clause K or Example Clause L, wherein the instructions further cause the system to: encapsulate, by the other host node, the data packet; embed, by the other host node, a general routing encapsulation key that uniquely identifies the virtual network of the service consumer in the data packet; and de-encapsulate, by the host node, the data packet.

[0081] Example Clause N, the system of any one of Example Clauses J through M, wherein the internal load balancer selects the virtual machine in the virtual network of the service provider as a final destination for the data packet.

[0082] Example Clause O, the system of any one of Example Clauses J through N, wherein translating the routable IP address in the data packet into the destination IP address comprises: translating the routable IP address into a customer IP address of the internal load balancer; and subsequently translating the customer IP address into the destination IP address.

[0083] Example Clause P, the system of any one of Example Clauses J through O, wherein the instructions further cause the system to translate, by the host node, an original destination port number in the data packet into a selected destination port number selected by the internal load balancer, the selected destination port number being associated with the virtual machine in the virtual network of the service provider.

[0084] Example Clause Q, a method comprising: receiving a data packet at a host node of a cloud-based platform that supports an internal load balancer associated with a virtual network of a service provider, wherein the data packet is routed to the host node using a routable Internet Protocol (IP) address assigned to the internal load balancer; translating, by the host node, a first source IP address in the data packet into a second source IP address associated with a subnet in the virtual network of the service provider, the subnet configured to map a service executing in the virtual network of the service provider into a virtual network of a service consumer; translating, by the host node, the routable address in the data packet into a destination IP address associated with a virtual machine in the virtual network of the service provider; and communicating, by the host node using the destination IP address, the data packet to the virtual machine in the virtual network of the service provider, wherein the data packet indicates the second source IP address as a source.

[0085] Example Clause R, the method of Example Clause Q, wherein the internal load balancer selects the virtual machine in the virtual network of the service provider as a final destination for the data packet.

[0086] Example Clause S, the method of Example Clause Q or Example Clause R, wherein translating the routable IP address in the data packet into the destination IP address comprises: translating the routable IP address into a customer IP address of the internal load balancer; and subsequently translating the customer IP address into the destination IP address.

[0087] Example Clause T, the method of any one of Example Clauses Q through S,

further comprising translating, by the host node, an original destination port number in the data packet into a selected destination port number selected by the internal load balancer, the selected destination port number being associated with the virtual machine in the virtual network of the service provider.

CONCLUSION

[0088] For ease of understanding, the processes discussed in this disclosure are delineated as separate operations represented as independent blocks. However, these separately delineated operations should not be construed as necessarily order dependent in their performance. The order in which the process is described is not intended to be construed as a limitation, and any number of the described process blocks may be combined in any order to implement the process or an alternate process. Moreover, it is also possible that one or more of the provided operations is modified or omitted.

[0089] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts are disclosed as example forms of implementing the claims.

[0090] The terms "a," "an," "the" and similar referents used in the context of describing the invention (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 "based on," "based upon," and similar referents are to be construed as meaning "based at least in part" which includes being "based in part" and "based in whole" unless otherwise indicated or clearly contradicted by context.

[0091] It should be appreciated that any reference to "first," "second," "third", etc. elements within the Summary and/or Detailed Description is not intended to and should not be construed to necessarily correspond to any reference of "first," "second," etc. elements of the claims. Rather, unless otherwise specified, any use of ""first," "second," and "third" within the Summary and/or Detailed Description may be used to distinguish between different instances of a same or a similar element (e.g., different IP addresses, different virtual machines, etc.).

[0092] Certain configurations are described herein, including the best mode known to the inventors for carrying out the invention. Of course, variations on these described configurations will become apparent to those of ordinary skill in the art upon reading the foregoing description. Skilled artisans will know how to employ such variations as

appropriate, and the configurations disclosed herein may be practiced otherwise than specifically described. Accordingly, all modifications and equivalents of the subject matter recited in the claims appended hereto are included within the scope of this disclosure. Moreover, any combination of the above-described elements in all possible variations thereof is encompassed by the invention unless otherwise indicated herein or otherwise clearly contradicted by context.

CLAIMS

1. A method for mapping a service executing in a virtual network of a service provider into a virtual network of a service consumer, the method comprising:

creating a subnet, in the virtual network of the service provider, for mapping the service executing in the virtual network of the service provider into the virtual network of the service consumer;

receiving a data packet at a host node of a cloud-based platform that supports an internal load balancer associated with the virtual network of the service provider, wherein the data packet is routed to the host node using a routable Internet Protocol (IP) address assigned to the internal load balancer;

translating, by the host node, a first source IP address in the data packet into a second source IP address, wherein the first source IP address is associated with a virtual machine in the virtual network of the service consumer and the second source IP address is associated with the subnet;

translating, by the host node, the routable IP address in the data packet into a destination IP address associated with a virtual machine in the virtual network of the service provider; and

communicating, by the host node using the destination IP address, the data packet to the virtual machine in the virtual network of the service provider, wherein the data packet indicates the second source IP address as a source.

2. The method of claim 1, further comprising:

creating a subnet, in the virtual network of the service consumer, for mapping the service executing in the virtual network of the service provider into the virtual network of the service consumer;

determining, by another host node of the cloud-based platform supporting the virtual network of the service consumer, that the data packet is being communicated from the first source IP address to another destination IP address associated with the subnet in the virtual network of the service consumer;

replacing, by the other host node, the other destination IP address in the data packet with the routable IP address; and

communicating, by the other host node using the routable IP address, the data packet over a network to the host node.

3. The method of claim 2, wherein the first source IP address and the other destination IP address comprise Internet Protocol version 4 (IPv4) addresses, and the

method further comprises converting, by the other host node, the IPv4 addresses into Internet Protocol version 6 (IPv6) addresses.

4. The method of claim 2 or claim 3, further comprising:

encapsulating, by the other host node, the data packet;

embedding, by the other host node, a general routing encapsulation key that uniquely identifies the virtual network of the service consumer in the data packet; and

de-encapsulating, by the host node, the data packet.

5. The method of any one of claims 2 through 4, further comprising:

determining that a reply data packet is communicated from the destination IP address to the second source IP address;

translating, by the host node, the destination IP address in the reply data packet into the other destination IP address associated with the subnet in the virtual network of the service consumer;

translating, by the host node, the second source IP address in the reply data packet into the first source IP address associated with the virtual machine in the virtual network of the service consumer; and

communicating, by the host node using the first source IP address as a destination, the reply data packet to the virtual machine in the virtual network of the service consumer.

6. The method of claim 5, wherein translating the second source IP address into the first source IP address uses a general routing encapsulation key to identify the virtual network of the service consumer.

7. The method of any one of claims 1 through 6, wherein the internal load balancer selects the virtual machine in the virtual network of the service provider as a final destination for the data packet.

8. The method of any one of claims 1 through 7, wherein translating the routable IP address into the destination IP address associated with the virtual machine in the virtual network of the service provider comprises:

translating the routable IP address into a customer IP address associated with the internal load balancer; and

subsequently translating the customer IP address into the destination IP address.

9. A system comprising:

one or more processors; and

computer-readable media storing instructions that, when executed by the one or more processors, cause the system to:

create a subnet, in a virtual network of a service provider, for mapping a service into a virtual network of a service consumer;

receive a data packet at a host node of a cloud-based platform that supports an internal load balancer associated with the virtual network of the service provider, wherein the data packet is routed to the host node using a routable Internet Protocol (IP) address assigned to the internal load balancer;

translate, by the host node, a first source IP address in the data packet into a second source IP address, wherein the first source IP address is associated with a virtual machine in the virtual network of the service consumer and the second source IP address is associated with the subnet;

translate, by the host node, the routable IP address in the data packet into a destination IP address associated with a virtual machine in the virtual network of the service provider; and

communicate, by the host node using the destination IP address, the data packet to the virtual machine in the virtual network of the service provider, wherein the data packet indicates the second source IP address as a source.

10. The system of claim 9, wherein the instructions further cause the system to:

create a corresponding subnet, in the virtual network of the service consumer, for mapping the service into the virtual network of the service consumer;

determine, by another host node of the cloud-based platform supporting the virtual network of the service consumer, that the data packet is being communicated from the first source IP address to another destination IP address associated with the corresponding subnet;

replace, by the other host node, the other destination IP address in the data packet with the routable IP address; and

communicate, by the other host node using the routable IP address, the data packet over a network to the host node.

11. The system of claim 9 or claim 10, wherein the internal load balancer selects the virtual machine in the virtual network of the service provider as a final destination for the data packet.

12. The system of any one of claims 9 through 11, wherein translating the routable IP address in the data packet into the destination IP address comprises:

translating the routable IP address into a customer IP address of the internal load balancer; and

subsequently translating the customer IP address into the destination IP address.

13. The system of any one of claims 9 through 12, wherein the instructions further cause the system to translate, by the host node, an original destination port number in the data packet into a selected destination port number selected by the internal load balancer, the selected destination port number being associated with the virtual machine in the virtual network of the service provider.

14. A method comprising:

receiving a data packet at a host node of a cloud-based platform that supports an internal load balancer associated with a virtual network of a service provider, wherein the data packet is routed to the host node using a routable Internet Protocol (IP) address assigned to the internal load balancer;

translating, by the host node, a first source IP address in the data packet into a second source IP address associated with a subnet in the virtual network of the service provider, the subnet configured to map a service executing in the virtual network of the service provider into a virtual network of a service consumer;

translating, by the host node, the routable address in the data packet into a destination IP address associated with a virtual machine in the virtual network of the service provider; and

communicating, by the host node using the destination IP address, the data packet to the virtual machine in the virtual network of the service provider, wherein the data packet indicates the second source IP address as a source.

15. The method of claim 14, wherein translating the routable IP address in the data packet into the destination IP address comprises:

translating the routable IP address into a customer IP address of the internal load balancer; and

subsequently translating the customer IP address into the destination IP address.

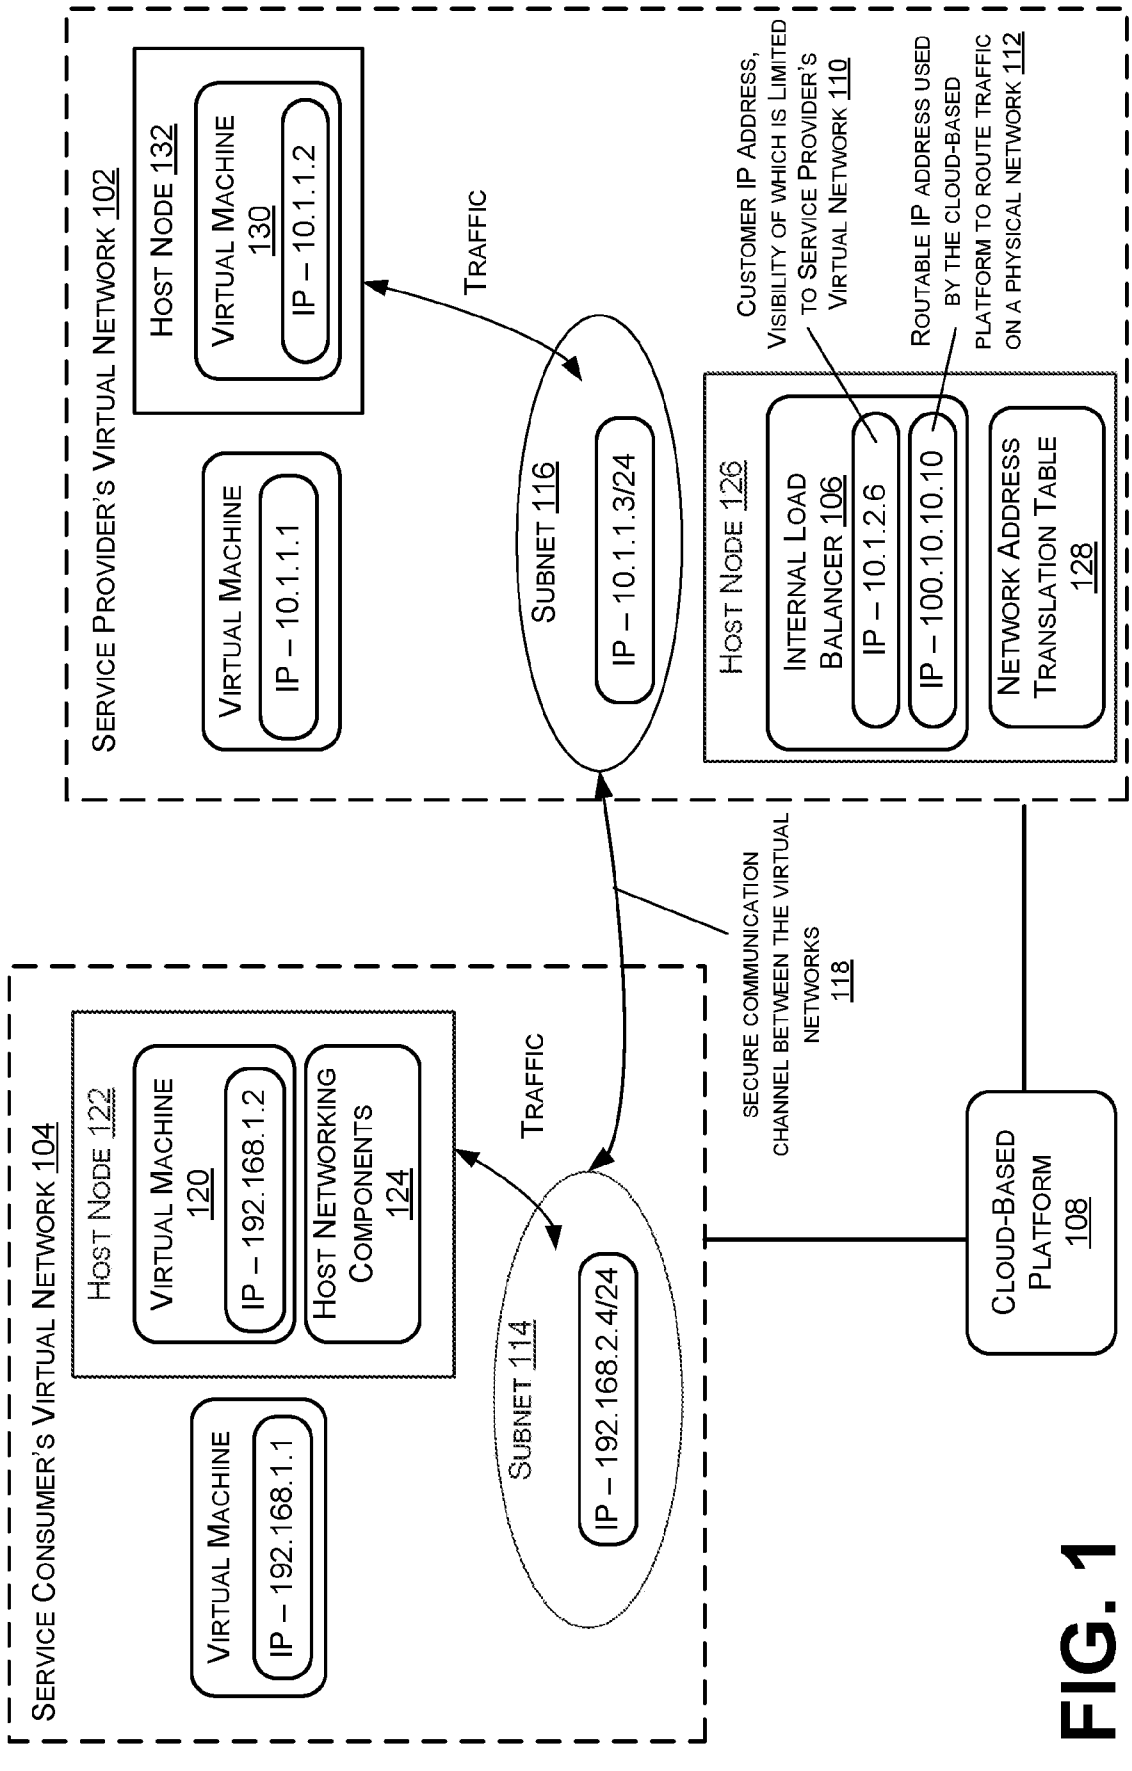
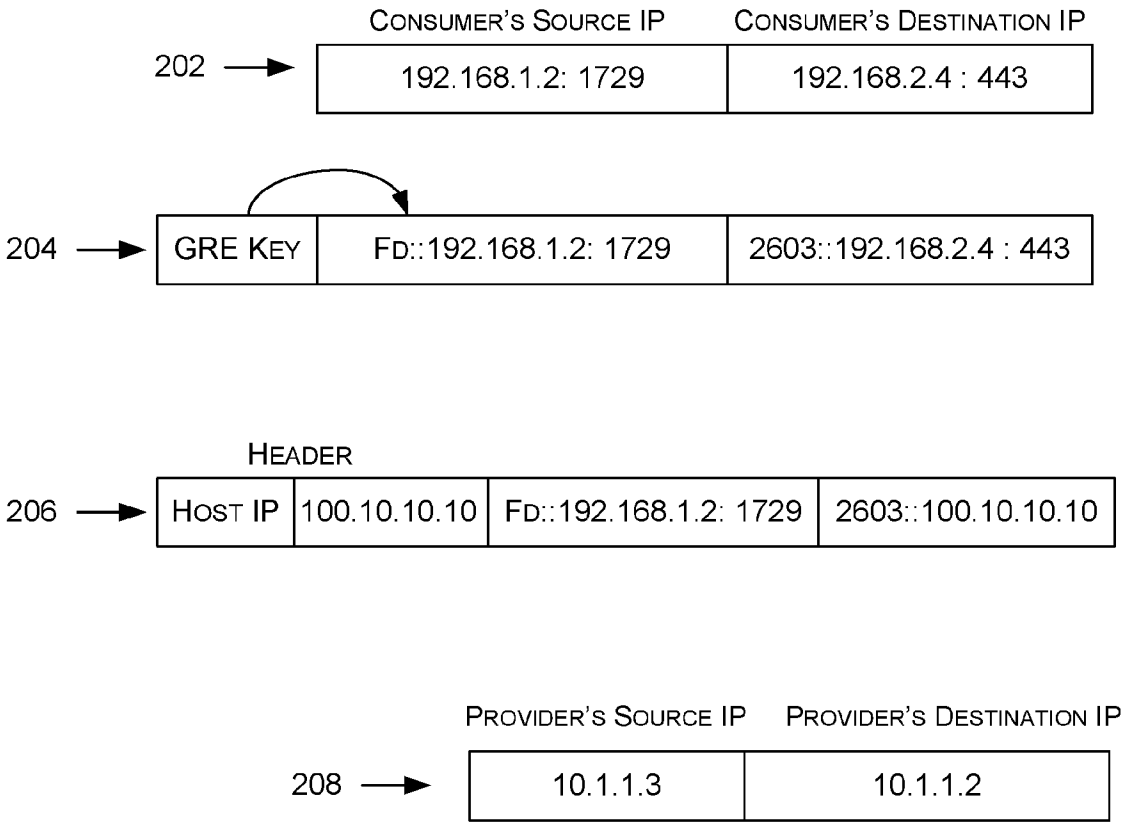

100
SERVICE PROVIDER'S VIRTUAL NETWORK 102
HOST NODE 132
VIRTUAL MACHINE 130
IP – 10.1.1.2
VIRTUAL MACHINE
IP – 10.1.1.1
TRAFFIC
SUBNET 116
IP – 10.1.1.3/24
CUSTOMER IP ADDRESS, VISIBILITY OF WHICH IS LIMITED TO SERVICE PROVIDER'S VIRTUAL NETWORK 110
ROUTABLE IP ADDRESS USED BY THE CLOUD-BASED PLATFORM TO ROUTE TRAFFIC ON A PHYSICAL NETWORK 112
HOST NODE 126
INTERNAL LOAD BALANCER 106
IP – 10.1.2.6
IP – 100.10.10.10
NETWORK ADDRESS TRANSLATION TABLE 128
SECURE COMMUNICATION CHANNEL BETWEEN THE VIRTUAL NETWORKS 118
SERVICE CONSUMER'S VIRTUAL NETWORK 104
HOST NODE 122
VIRTUAL MACHINE 120
IP – 192.168.1.2
HOST NETWORKING COMPONENTS 124
TRAFFIC
VIRTUAL MACHINE
IP – 192.168.1.1
SUBNET 114
IP – 192.168.2.4/24
CLOUD-BASED PLATFORM 108

FIG. 1



I) THIS IS PART OF A DATA PACKET THAT LEAVES THE SERVICE CONSUMER'S VIRTUAL MACHINE

II) CONVERT IPV4 ADDRESSES INTO IPV6 ADDRESSES AND EMBED A GENERAL ROUTING ENCAPSULATION KEY INTO IPV6 ADDRESS

III) HOST NODE REPLACES THE DESTINATION IP ADDRESS IN THE SERVICE CONSUMER'S VIRTUAL NETWORK WITH THE ROUTABLE IP ADDRESS OF THE INTERNAL LOAD BALANCER IN THE SERVICE PROVIDER'S VIRTUAL NETWORK, AND ENCAPSULATES THE DATA PACKET USING A HEADER

IV) DE-ENCAPSULATE DATA PACKET TO REMOVE HEADER AND PERFORM NETWORK ADDRESS TRANSLATION ON HOST NODE(S) SUPPORTING THE INTERNAL LOAD BALANCER SO THAT DATA PACKET IS COMMUNICATED TO A VIRTUAL MACHINE IN THE SERVICE PROVIDER'S VIRTUAL NETWORK

FIG. 2

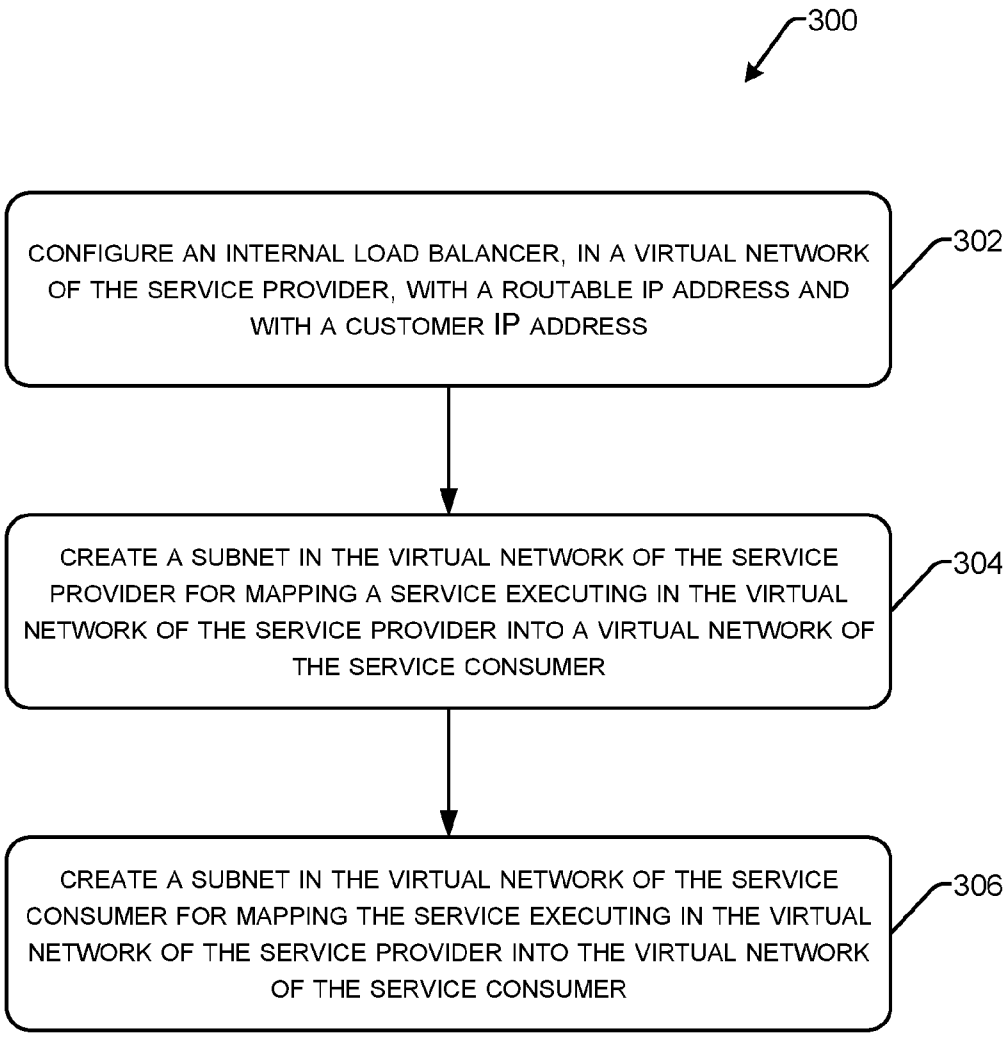

300
CONFIGURE AN INTERNAL LOAD BALANCER, IN A VIRTUAL NETWORK OF THE SERVICE PROVIDER, WITH A ROUTABLE IP ADDRESS AND WITH A CUSTOMER IP ADDRESS
302
CREATE A SUBNET IN THE VIRTUAL NETWORK OF THE SERVICE PROVIDER FOR MAPPING A SERVICE EXECUTING IN THE VIRTUAL NETWORK OF THE SERVICE PROVIDER INTO A VIRTUAL NETWORK OF THE SERVICE CONSUMER
304
CREATE A SUBNET IN THE VIRTUAL NETWORK OF THE SERVICE CONSUMER FOR MAPPING THE SERVICE EXECUTING IN THE VIRTUAL NETWORK OF THE SERVICE PROVIDER INTO THE VIRTUAL NETWORK OF THE SERVICE CONSUMER
306

FIG. 3

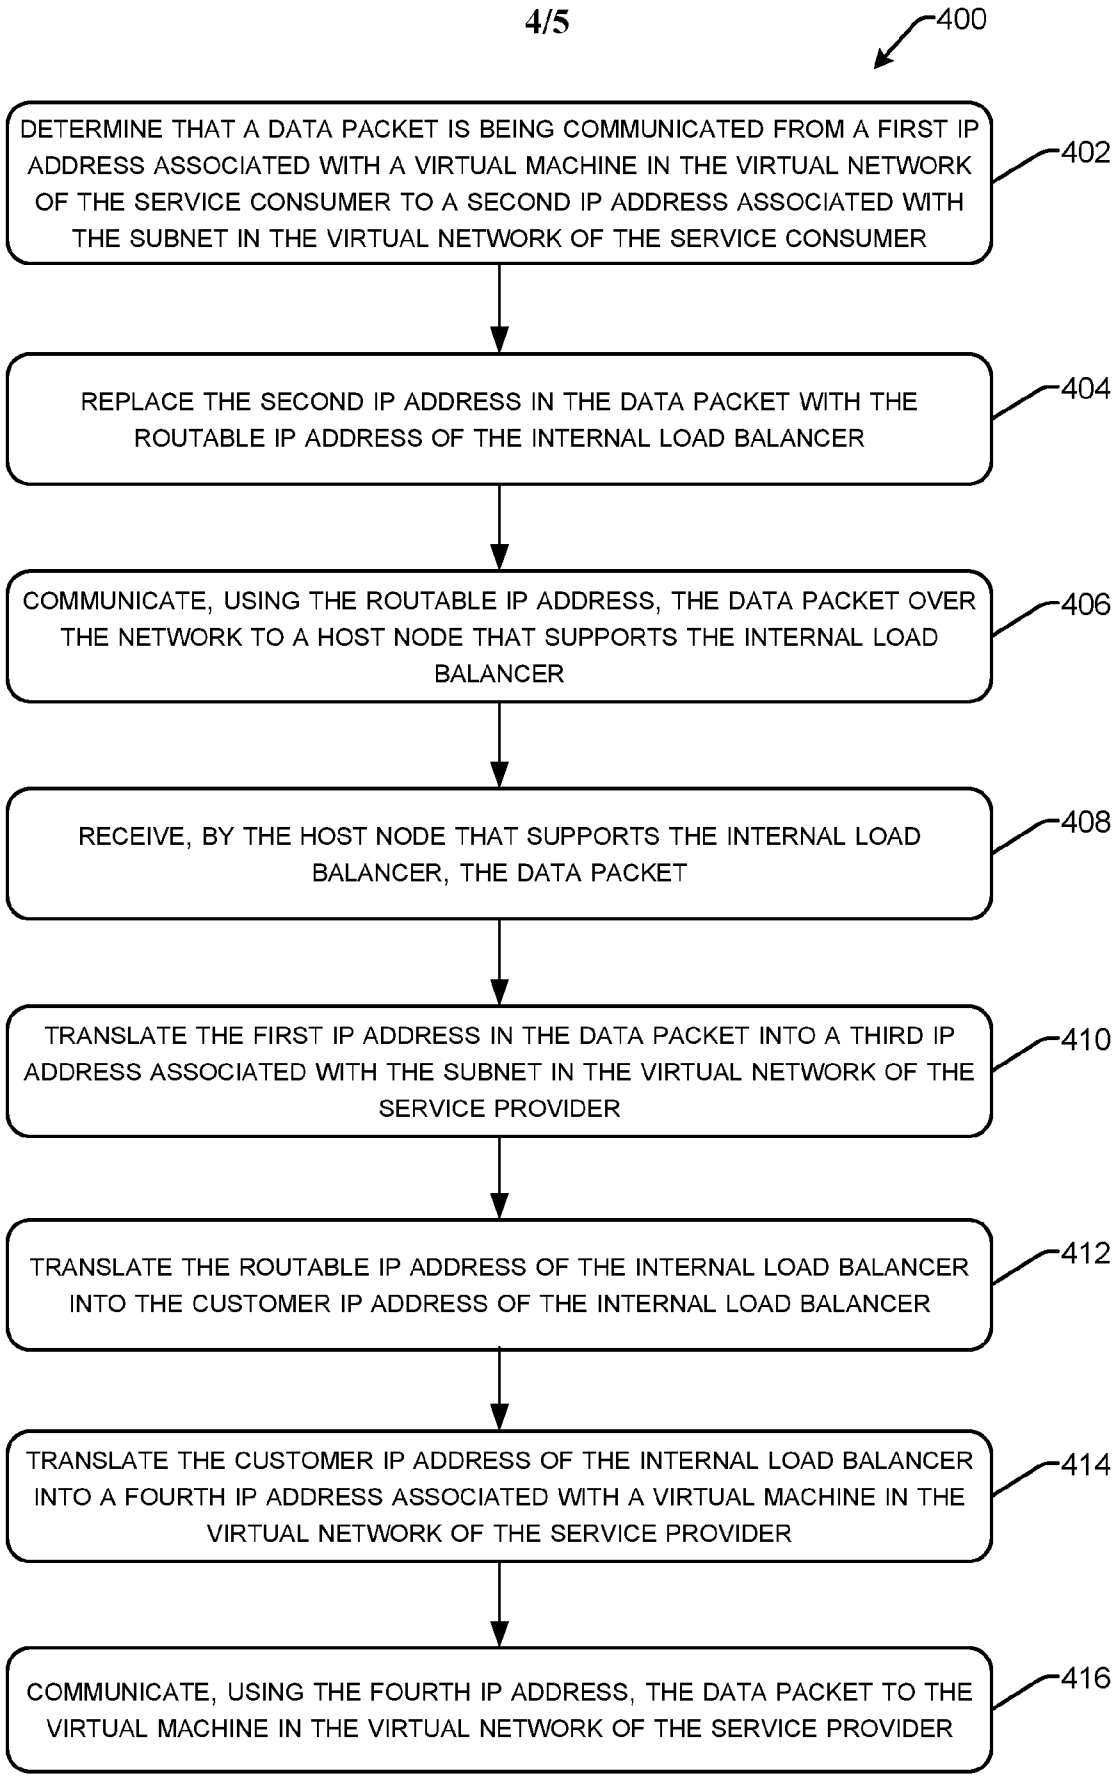

400
DETERMINE THAT A DATA PACKET IS BEING COMMUNICATED FROM A FIRST IP ADDRESS ASSOCIATED WITH A VIRTUAL MACHINE IN THE VIRTUAL NETWORK OF THE SERVICE CONSUMER TO A SECOND IP ADDRESS ASSOCIATED WITH THE SUBNET IN THE VIRTUAL NETWORK OF THE SERVICE CONSUMER
402
REPLACE THE SECOND IP ADDRESS IN THE DATA PACKET WITH THE ROUTABLE IP ADDRESS OF THE INTERNAL LOAD BALANCER
404
COMMUNICATE, USING THE ROUTABLE IP ADDRESS, THE DATA PACKET OVER THE NETWORK TO A HOST NODE THAT SUPPORTS THE INTERNAL LOAD BALANCER
406
RECEIVE, BY THE HOST NODE THAT SUPPORTS THE INTERNAL LOAD BALANCER, THE DATA PACKET
408
TRANSLATE THE FIRST IP ADDRESS IN THE DATA PACKET INTO A THIRD IP ADDRESS ASSOCIATED WITH THE SUBNET IN THE VIRTUAL NETWORK OF THE SERVICE PROVIDER
410
TRANSLATE THE ROUTABLE IP ADDRESS OF THE INTERNAL LOAD BALANCER INTO THE CUSTOMER IP ADDRESS OF THE INTERNAL LOAD BALANCER
412
TRANSLATE THE CUSTOMER IP ADDRESS OF THE INTERNAL LOAD BALANCER INTO A FOURTH IP ADDRESS ASSOCIATED WITH A VIRTUAL MACHINE IN THE VIRTUAL NETWORK OF THE SERVICE PROVIDER
414
COMMUNICATE, USING THE FOURTH IP ADDRESS, THE DATA PACKET TO THE VIRTUAL MACHINE IN THE VIRTUAL NETWORK OF THE SERVICE PROVIDER
416

FIG. 4

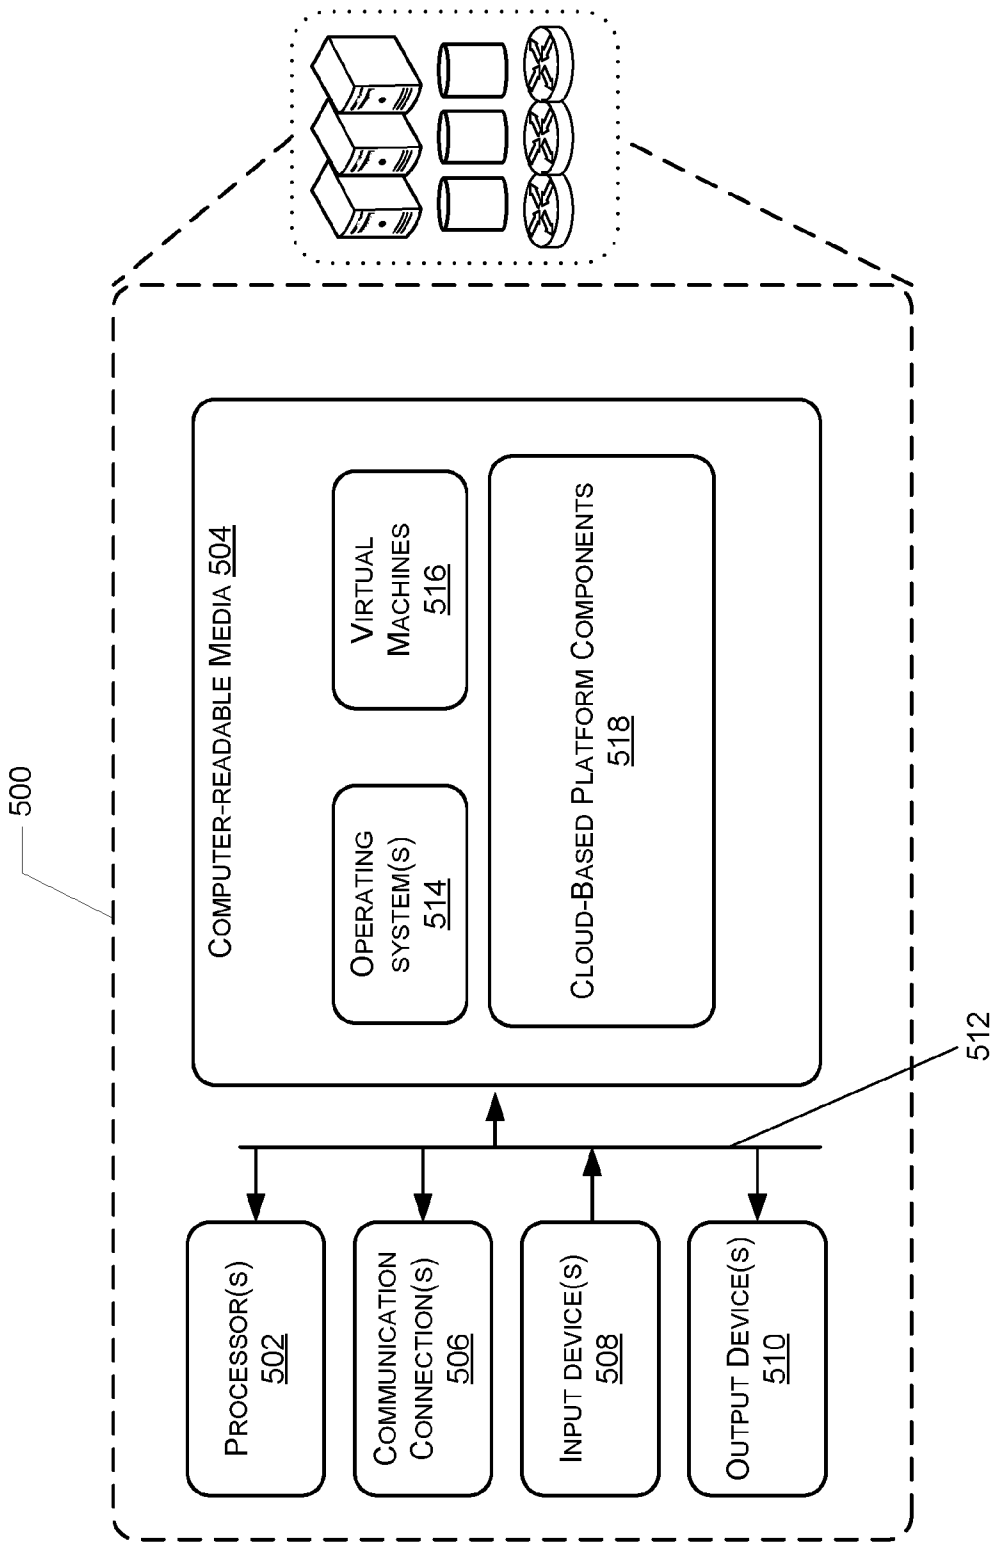

500
COMPUTER-READABLE MEDIA 504
OPERATING SYSTEM(S) 514
VIRTUAL MACHINES 516
CLOUD-BASED PLATFORM COMPONENTS 518
512
PROCESSOR(S) 502
COMMUNICATION CONNECTION(S) 506
INPUT DEVICE(S) 508
OUTPUT DEVICE(S) 510

FIG. 5

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/039097

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L29/12
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 2018/139101 A1 (PURI RANJIT [US] ET AL) 17 May 2018 (2018-05-17) paragraphs [0005], [0023] - [0025] -----	1-15
X	PARVEEN PATEL ET AL: "Ananta", COMPUTER COMMUNICATION REVIEW., vol. 43, no. 4, 27 August 2013 (2013-08-27), pages 207-218, XP055440057, US ISSN: 0146-4833, DOI: 10.1145/2534169.2486026 paragraphs [02.1], [3.2.1], [3.2.2], [03.3], [3.3.1], [3.3.2], [3.4.1] - [3.4.3] ----- -/--	1,7-9, 11-15

[X] Further documents are listed in the continuation of Box C.

[X] See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
5 September 2019	16/09/2019

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Goller, Wolfgang

1

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/039097

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2015/195196 A1 (PATEL PARVEEN [US]) 9 July 2015 (2015-07-09) paragraphs [0005], [0006], [0015], [0016], [0035], [0036], [0039] -----	1-15

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2019/039097

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2018139101 A1	17-05-2018	CN 109937400 A	25-06-2019
		US 2018139101 A1	17-05-2018
		WO 2018093615 A1	24-05-2018
US 2015195196 A1	09-07-2015	CN 105900407 A	24-08-2016
		EP 3092753 A1	16-11-2016
		US 2015195196 A1	09-07-2015
		US 2017295096 A1	12-10-2017
		US 2018212873 A1	26-07-2018
		WO 2015105690 A1	16-07-2015