



(51) International Patent Classification:

H04L 12/751 (2013.01) H04L 12/713 (2013.01)
H04L 12/709 (2013.01) H04L 12/741 (2013.01)

(21) International Application Number:

PCT/US2020/058896

(22) International Filing Date:

04 November 2020 (04.11.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/742,633	14 January 2020 (14.01.2020)	US
16/742,685	14 January 2020 (14.01.2020)	US
16/742,663	14 January 2020 (14.01.2020)	US

(71) Applicant: **VMWARE, INC.** [US/US]; 3401 Hillview Avenue, Palo Alto, CA 94304 (US).

(72) Inventors: **JAIN, Jayant**; 3401 Hillview Avenue, Palo Alto, CA 94304 (US). **PARSA, Mike**; 3401 Hillview Avenue, Palo Alto, CA 94304 (US). **HONG, Xinhua**; 3401 Hillview Avenue, Palo Alto, CA 94304 (US). **MANUGURI, Subrahmanyam**; 3401 Hillview Avenue, Palo Alto, CA 94304 (US). **SENGUPTA, Anirban**; 20031 Karn Circle, Saratoga, CA 95070 (US).

(74) Agent: **ADELI, Mani** et al.; Adeli LLP, 11859 Wilshire Blvd., Suite 408, Los Angeles, CA 90025 (US).

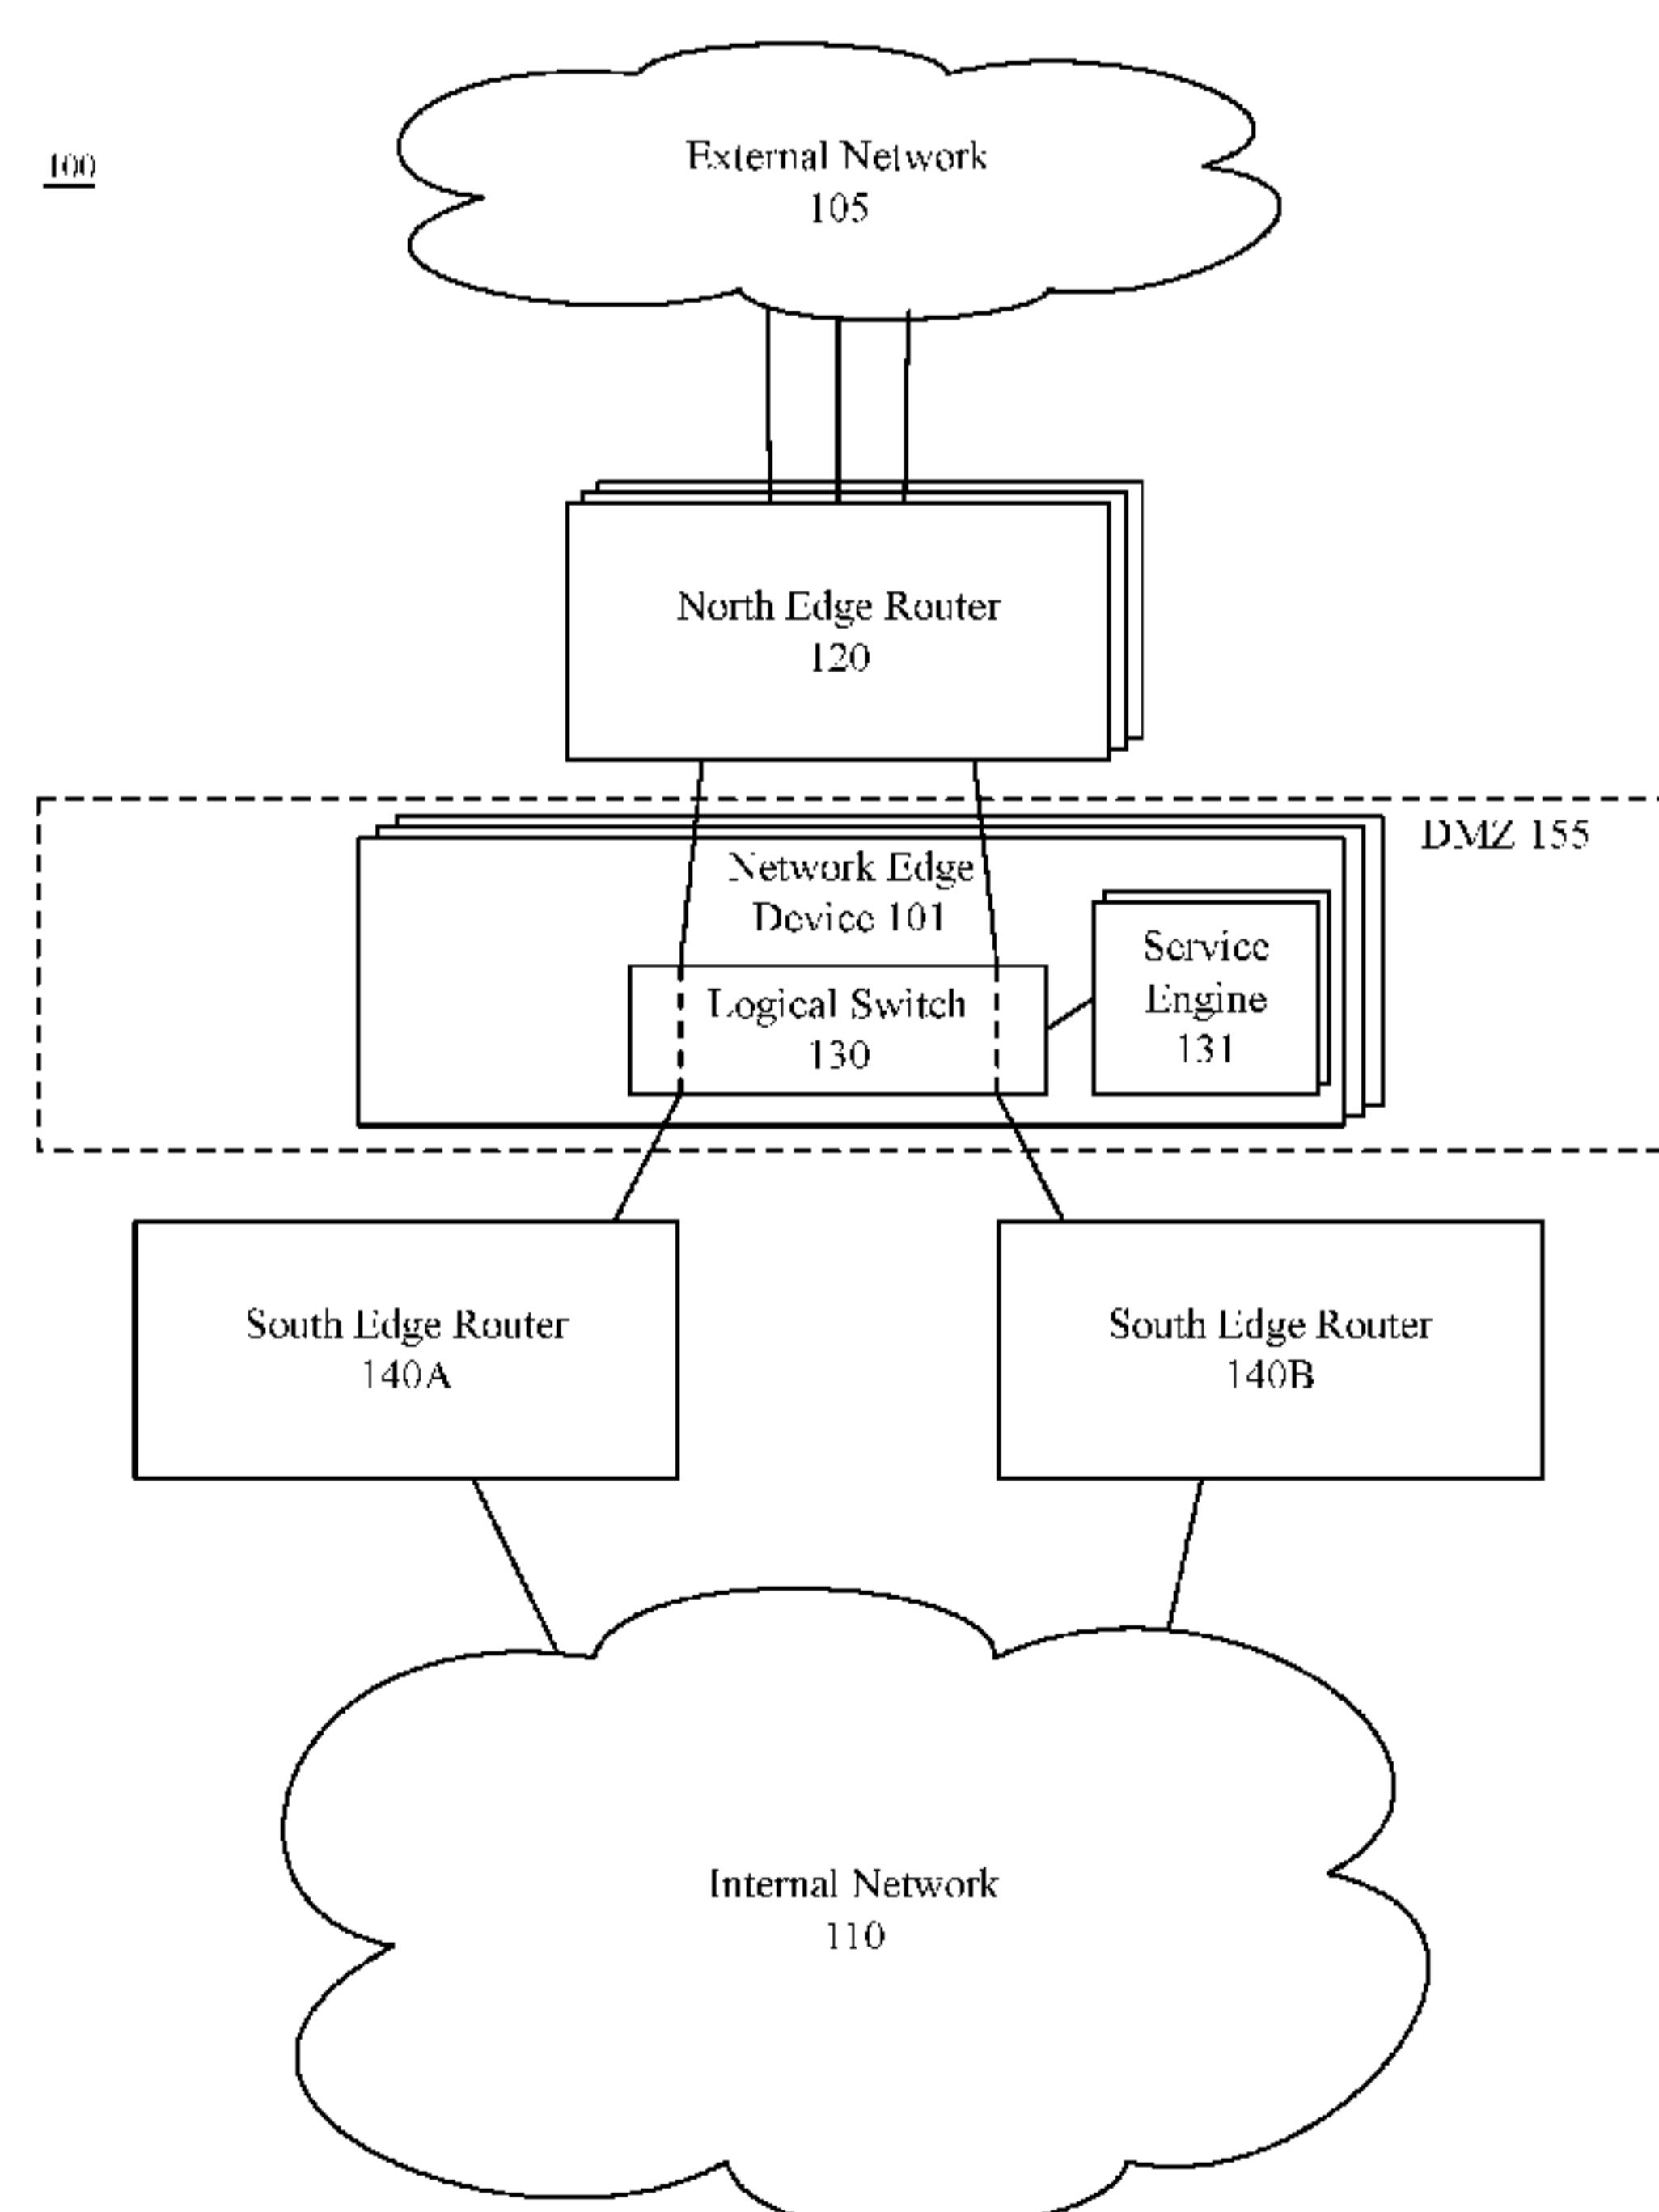
(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS,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

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

Published:

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

(54) Title: TRANSPARENT DEMILITARIZED ZONE PROVIDING STATEFUL SERVICE BETWEEN PHYSICAL AND LOGICAL NETWORKS



(57) Abstract: Some embodiments of the invention provide novel methods for providing a stateful service at a network edge device (e.g., an NSX edge) that has a plurality of north-facing interfaces (e.g., interfaces to an external network) and a plurality of corresponding south-facing interfaces (e.g., interfaces to a logical network). In some embodiments, the network edge device receives data messages from a first gateway device from a logical network, provides the stateful network service to the data message, and forwards the data message towards the destination through a corresponding interface connected to a physical network.

Figure 1

TRANSPARENT DEMILITARIZED ZONE PROVIDING STATEFUL SERVICE BETWEEN PHYSICAL AND LOGICAL NETWORKS

Jayant Jain, Mike Parsa, Xinhua Hong, Subrahmanyam Manuguri, Anirban Sengupta

BACKGROUND

[0001] Providing a stateful service across two or more equal cost multipathing (ECMP) paths presents certain challenges. If different paths are selected for forward and return data messages of a data message flow, a service engine providing the stateful service on one path may not be able to track a state of the data message flow as the service engine will not receive all the data messages of the data message flow (i.e., will not receive data messages for which an alternative path is selected). If the different paths are combined to pass through a same service engine, the path selection is not respected as the combined paths will redistribute incoming data messages among the combined outgoing interfaces such that a particular data message will be forwarded along a different path than the path selected by the ECMP operation. Accordingly, a solution that respects ECMP path selection and maintains correct state information at a stateful service engine is required.

BRIEF SUMMARY

[0002] Some embodiments of the invention provide novel methods for providing a stateful service at a network edge device (e.g., an NSX edge) that has a plurality of north-facing interfaces (e.g., interfaces to an external network) and a plurality of corresponding south-facing interfaces (e.g., interfaces to an internal logical network). North-facing interfaces of the network edge device are associated with south-facing interfaces. In some embodiments, the association is one-to-one, while in other embodiments some interfaces are grouped (e.g., aggregated) and the groups are associated in a one-to-one fashion. The network edge device receives data messages at the interfaces from forwarding elements that select particular paths (e.g., interfaces of forwarding elements on the other side of the network edge device) for the data messages. The selection of paths, in some embodiments, is based on equal cost multipathing (ECMP) or a similar selection mechanism. The network edge device is transparent to the north and south forwarding elements such that they are not aware that the network edge device is interposed between north and south forwarding elements in some embodiments.

[0003] A set of interfaces on the north (south) side of the network edge device for a set of equal cost paths (e.g., an ECMP group), in some embodiments, are bonded together in the network edge device to correspond to a single interface on the north (south) side of a logical bridge including at least one logical switch providing a stateful service implemented by the network edge device. The bond is implemented, in some embodiments, using a bonding module executing on the network edge device that maintains a mapping between ingress and egress interfaces (in addition to standard bonding techniques) to allow deterministic forwarding through the network edge device in the presence of bonded interfaces. A bonding module on the ingress side for a particular data message, in some embodiments, associates an identifier of an ingress interface on which a data message is received with the data message. In some embodiments, the association is stored in metadata of the data message and is removed after a bonding module on the egress side determines the egress interface based on the mapping between ingress and egress interfaces.

[0004] In some embodiments, the at least one logical switch calls a service engine to provide the stateful service. The bonded interfaces force all traffic for the ECMP group to go through the same set of logical switches and call the same service engine to provide the stateful service and maintain state information for each flow crossing the logical switches. If the interfaces were not bonded and separate logical switches (and service engines) processed data messages on

the separate paths, a return data message could use a different path than an original message and state information would not be kept by the service engine and the stateful service would not function properly.

[0005] The network edge device, in some embodiments, receives a data message at a first interface and provides the data message to the logical switch interface associated with the bonding module. The bonding module associates the data message with an identifier of the ingress interface. In some embodiments, the association is inserted in the data message, while in other embodiments the identifier is appended to the data message. A service engine is called by a logical switch, in some embodiments, to provide a stateful service. The service engine provides the stateful service and provides the data message to an egress interface associated with a second bonding module. The second bonding module uses a mapping table, in some embodiments, to determine an egress interface associated with the ingress interface. In some embodiments, the logical processing removes the ingress interface identifier and provides the data message to the egress module.

[0006] The network edge device, in some embodiments, is a primary network edge device and a secondary network edge device is implemented as a standby network edge device. The secondary network edge device replicates the interfaces and associations of the primary network edge device and is available to assume the functions of the primary network edge device in the case of a failure of the primary network edge device. In some embodiments, failure of the primary network edge device is defined by a user. The definition of failure, in some embodiments, is defined by a user-supplied policy (e.g., a threshold for a given metric). In other embodiments, failover occurs when all paths through the primary network edge device fail.

[0007] In some embodiments, instead of forcing all data messages to cross a single bridge, each interface associated with a different bridge calls a service engine based on identifiers included in data messages received at the interface. Each data message flow is associated with a particular identifier that is associated with a particular service engine instance that provides the stateful service. In some embodiments, the interface that receives a data message identifies a service engine to provide the stateful service and provides the data message to the identified service engine. After processing the data message, the service engine provides the data message to the egress interface associated with the ingress interface.

[0008] The preceding Summary is intended to serve as a brief introduction to some embodiments of the invention. It is not meant to be an introduction or overview of all inventive subject matter disclosed in this document. The Detailed Description that follows and the Drawings that are referred to in the Detailed Description will further describe the embodiments described in the Summary as well as other embodiments. Accordingly, to understand all the embodiments described by this document, a full review of the Summary, Detailed Description, the Drawings and the Claims is needed. Moreover, the claimed subject matters are not to be limited by the illustrative details in the Summary, Detailed Description and the Drawing.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The novel features of the invention are set forth in the appended claims. However, for purposes of explanation, several embodiments of the invention are set forth in the following figures.

[0010] **Figure 1** illustrates an exemplary network topology of a system implementing an embodiment of the invention.

[0011] **Figure 2** illustrates a particular embodiment of a system topology that includes a north edge device implementing a logical router and a logical switch (and DMZ).

[0012] **Figure 3** illustrates a particular embodiment of a network topology of a system implementing an embodiment of the invention using a common connection tracker.

[0013] **Figure 4** illustrates a particular embodiment of a network topology of a system implementing an embodiment of the invention using different service engines for different logical networks.

[0014] **Figure 5** conceptually illustrates a process for configuring a network edge device to process data messages using a set of bonded interfaces with deterministic mapping between bonded interfaces.

[0015] **Figure 6** conceptually illustrates a process for configuring a network edge device to implement the bonded interfaces and bonding modules that, in some embodiments, maintain the mapping tables generated.

[0016] **Figure 7** illustrates two sets of mapping tables that are used in different embodiments of the invention.

[0017] **Figure 8** conceptually illustrates a process for using the configured network edge device to provide a stateful service.

[0018] **Figure 9** illustrates a particular network edge device in a set of network edge devices connecting a first north edge routing element and a south edge routing element.

[0019] **Figure 10** illustrates a second embodiment of a network in which the invention is carried out.

[0020] **Figure 11** illustrates an embodiment including active and standby network edge devices.

[0021] **Figure 12** conceptually illustrates a process for failure handling.

[0022] **Figure 13** conceptually illustrates a process for configuring a network edge device to implement an embodiment of the invention that allows for separate logical switches for different equal cost paths.

[0023] **Figure 14** conceptually illustrates a process for processing data messages at a network edge device configured according to an embodiment in which a set of shared service engines are implemented for the set of logical switches.

[0024] **Figure 15** illustrates an embodiment in which the process of **Figure 14** is carried out.

[0025] **Figure 16** conceptually illustrates a process for processing data messages at a network edge device configured according to an embodiment in which a common connection tracker is shared by each logical switch and the service engine associated with each logical switch.

[0026] **Figure 17** illustrates an embodiment in which the process of **Figure 16** is carried out.

[0027] **Figure 18** illustrates the system of some embodiments.

[0028] **Figure 19** conceptually illustrates a computer system with which some embodiments of the invention are implemented.

DETAILED DESCRIPTION

[0029] In the following detailed description of the invention, numerous details, examples, and embodiments of the invention are set forth and described. However, it will be clear and apparent to one skilled in the art that the invention is not limited to the embodiments set forth and that the invention may be practiced without some of the specific details and examples discussed.

[0030] As used in this document, the term data packet, packet, data message, or message refers to a collection of bits in a particular format sent across a network. It should be understood that the term data packet, packet, data message, or message may be used herein to refer to various formatted collections of bits that may be sent across a network, such as Ethernet frames, IP packets, TCP segments, UDP datagrams, etc. While the examples below refer to data packets, packets, data messages, or messages, it should be understood that the invention should not be limited to any specific format or type of data message. Also, as used in this document, references to L2, L3, L4, and L7 layers (or layer 2, layer 3, layer 4, layer 7) are references to the second data link layer, the third network layer, the fourth transport layer, and the seventh application layer of the OSI (Open System Interconnection) layer model, respectively.

[0031] Data messages, in some embodiments, are organized into data message flows based on being part of a same communication session between a source and destination compute node. A unidirectional data message flow is often defined by a set of header values (e.g., a 5-tuple for a TCP flow, or, more generally, an n-tuple) of a data message. A bidirectional data message flow includes both the forward unidirectional flow and the return unidirectional data message flow (e.g., a return flow that has a set of source and destination addresses reversed in relation to the forward flow). For the purposes of this discussion, a data message flow will refer to the bidirectional data message flow unless otherwise indicated.

[0032] A user-defined logical network as used in this application, refers to a particular logical abstraction of a network. In some embodiments, the logical abstraction includes logical counterparts to network elements of a physical network such as forwarding elements (e.g., switches, hubs, routers, bridges, etc.), load balancers, and firewalls. The logical forwarding elements (e.g., a logical switch or logical router) in some embodiments are implemented by a set of MFEs (e.g., physical or virtual/software switches, or routers) executing on host machines. A particular host machine may host data compute nodes (DCNs) (e.g., containers or virtual machines (VMs)) connected to multiple different logical networks and the set of MFEs implements all the

logical networks to which the DCNs logically connect. Additional details of the structure and function of logical networks are described in U.S. Patent Publication 2016/0226754, which is hereby incorporated by reference.

[0033] A network edge device is implemented in some logical networks. A network edge device connects a first set of forwarding elements on a first side of the network edge device to a second set of forwarding elements on a second side of the network edge device. In some embodiments, the network edge device executes in a 'bare-metal' environment (i.e., directly on a physical device). The first set of forwarding elements, in some embodiments, are MFEs implementing a distributed logical router in a particular logical network and the second set of forwarding elements is a service router in the logical network that provides connectivity to an external network. In other embodiments, the first set of forwarding elements is a set of logical forwarding elements (e.g., routers, switches, etc.) in a logical network and the second set of forwarding elements is a set of forwarding elements external to the logical network.

[0034] In other embodiments, the network edge device operates between two networks (e.g., two logical networks, two physical networks, or between one logical network and one physical network). In some embodiments, the two gateway devices (e.g., network edge routers) associated with the two networks are equivalent to a single logical gateway that is separated into two logical gateways (e.g., one associated with each network) for additional security. The network edge device implements a set of logical switches in some embodiments. Logical switches, in some embodiments, are distributed logical switches implemented by a plurality of network edge devices. The logical switches are associated with a distributed stateful service, or set of stateful services, in some embodiments.

[0035] Some embodiments of the invention provide novel methods for providing a stateful service at a network edge device (e.g., an NSX edge) that has a plurality of north-facing interfaces (e.g., interfaces to an external network) and a plurality of corresponding south-facing interfaces (e.g., interfaces to a logical network). North-facing interfaces of the network edge device are associated with south-facing interfaces. In some embodiments, the association is one-to-one, while in other embodiments some interfaces are grouped (e.g., aggregated) and the groups are associated in a one-to-one fashion. The network edge device receives data messages at the interfaces from forwarding elements that select particular paths (e.g., interfaces of forwarding elements on the other side of the network edge device) for the data messages. The selection of paths, in some

embodiments, is based on equal cost multipathing (ECMP) or a similar selection mechanism. The network edge device is transparent to the north and south forwarding elements such that they are not aware that the network edge device is interposed between north and south forwarding elements in some embodiments. For example, a time to live parameter is not decremented, and, conjunctively or alternatively, a source IP address is not changed.

[0036] **Figure 1** illustrates an exemplary network topology of a system 100 implementing an embodiment of the invention. System 100 includes a set of network forwarding elements connecting external network 105 and internal network 110. External network 105 and internal network 110 connect to a set of edge routers, i.e., north edge routers 120 and south edge routers 140 respectively, that connect the networks 105 and 110. The edge routers 120 and 140, in some embodiments, are any combination of physical routers and software routers implemented on a physical network edge device. The software routers, in some embodiments, are implemented as a plurality of forwarding elements such as a distributed router, a transit logical switch, and a service router as described in U.S. Patent Publication 2016/0226754. In some embodiments, the edge routers 120 and 140 are specifically the service routers of an edge device. The south edge routers 140, in some embodiments, are active-active edge devices for a same logical network (or set of logical networks) in internal network 110. In some other embodiments, the south edge routers 140 are in an active-standby configuration.

[0037] Interposed between north edge routers 120 and south edge routers 140 is a set of network edge devices 101 that each execute a set of logical switches 130 and a set of service engines 131 that provide a set of gateway services in a transparent bridge between interfaces of the north and south edge routers. A network edge device 101, in some embodiments, is a bare metal edge device. The provision of gateway services will be described in further detail below in relation to **Figures 5-17**. The set of network edge devices 101 implements a demilitarized zone (DMZ) 155 that provides an extra layer of protection between the external and internal networks 105 and 110. In some embodiments, each network edge device provides a connection between a specific internal (e.g., logical network) and the external network.

[0038] **Figure 2** illustrates a particular embodiment of a system topology that includes a set of north edge devices 225 each implementing logical router 220 (corresponding to a service router of north edge router 120) and the logical switch 130 (and DMZ 155). As opposed to **Figure 1**, the external network is an external logical network 205. Additionally, the logical router 220, in

some embodiments, is a component of external logical network 205. **Figure 2** illustrates a set of connections for a single north edge device 225 with a set of edge routers (e.g., in an active-active or active-standby configuration) of an internal network (e.g., a logical network) 110. However, in this and other embodiments, additional similar devices are employed in the system 200 for connecting to additional internal elements (e.g., edge routers of additional logical networks in the internal network 110).

[0039] **Figure 3** illustrates a particular embodiment of a network topology of a system 300 implementing an embodiment of the invention using a common connection tracker. **Figure 3** includes a set of logical switches 330 in place of the single logical switch 130. One embodiment using a set of logical switches 330 is described in **Figure 17**. The multiple logical switches 330, in some embodiments, are each implemented by a set of network edge devices making up the DMZ 155. Each logical switch in some embodiments includes an associated service engine 331 to provide a service for data messages traversing the DMZ 155. In other embodiments, each of a set of network edge devices implements a subset of logical switches (e.g., anywhere from one to all of the logical switches) in the set of logical switches 330. The service engines 331, in some embodiments, share a connection tracker 346 that maintains connection states for data flows passing through the DMZ using any of the logical switches 330.

[0040] **Figure 3** also includes a set of south edge routers 340 and a set of north edge routers 120. In some embodiments, each set of edge routers 120 and 340 represents a set of edge devices in one of an active-active or active-standby configuration. As discussed above, some embodiments implement the south edge router as a component of an edge device that also includes a transit logical switch and a distributed router for a logical network (e.g., a logical network in internal network 110). In some embodiments, a single physical device implements multiple south (north) edge routers in the set of south (north) edge routers 340 (120). In other embodiments, each edge router in the set of edge routers 120 and 340 are implemented by a different edge device. Some embodiments using an active-active configuration assign different traffic to different edge devices (e.g., using a sharding technique) to ensure that each data message in a particular data message flow (e.g., a bidirectional data message flow between a compute node in the internal network 110 and a compute node in the external network 105) is processed at a same edge device so as to maintain accurate state information without having to synchronize data-message-flow state information across edge devices.

[0041] **Figure 4** illustrates a particular embodiment of a network topology of a system 400 implementing an embodiment of the invention using different service engines for different logical networks. **Figure 4** includes a set of logical switches 430 as described in **Figure 15**. The multiple logical switches 430, in some embodiments, are each implemented by a set of network edge devices making up the DMZ 155. Each network edge device, in some embodiments, also executes a set of service engines associated with different tenants or logical networks (e.g., logical networks 450A and 450B) for which it provides stateful services. Each logical switch, in some embodiments, can call any service engine 431 executing on the same network edge device to provide a service for data messages traversing the DMZ 155. The service engine called, in some embodiments, depends on the logical network associated with the data message to be provided the service. The service engines 431, in some embodiments, each have their own connection tracker 446 to maintain state data for data messages associated with the logical network 450 associated with the service engine 431 using the connection tracker 446. Further details of the implementation are discussed in relation to **Figures 14 and 15**.

[0042] **Figure 4** also includes a set of south edge routers 440 and a set of north edge routers 120. In some embodiments, each set of edge routers 120 and 440 represents a set of edge devices in one of an active-active or active-standby configuration. As discussed above, some embodiments implement the south edge router as a component of an edge device that also includes a transit logical switch and a distributed router for a logical network (e.g., a logical network in internal network 110). In some embodiments, the south edge routers 440 include south edge routers for both logical networks 450 that execute on a same edge device. In other embodiments, each logical network's south edge routers 440 execute on a different set of edge devices. In some embodiments, a single physical device implements multiple south (north) edge routers in the set of south (north) edge routers 440 (120). In other embodiments, each edge router in the set of edge routers 120 and 440 are implemented by a different edge device. Some embodiments using an active-active configuration assign different traffic to different edge devices (e.g., using a sharding technique) to ensure that each data message in a particular data message flow (e.g., a bidirectional data message flow between a compute node in the internal network 110 and a compute node in the external network 105) is processed at a same edge device so as to maintain accurate state information without having to synchronize data-message-flow state information across edge devices.

[0043] A set of interfaces on the north (south) side of the network edge device for a set of equal cost paths (e.g., an ECMP group), in some embodiments, are bonded together in the network edge device to correspond to a single interface on the north (south) side of a logical bridge including at least one logical switch providing a stateful service implemented by the network edge device. The bond is implemented, in some embodiments, using a bonding module executing on the network edge device that maintains a mapping between ingress and egress interfaces to allow deterministic forwarding through the network edge device in the presence of bonded interfaces (in addition to standard bonding techniques). A bonding module on the ingress side for a particular data message, in some embodiments, associates an identifier of an ingress interface on which a data message is received with the data message. In some embodiments, the association is stored in metadata of the data message and is removed after a bonding module on the egress side determines the egress interface based on the mapping between ingress and egress interfaces.

[0044] In some embodiments, the at least one logical switch calls a service engine to provide the stateful service. The bonded interfaces force all traffic for the ECMP group to go through the same set of logical switches and call the same service engine to provide the stateful service and maintain state information for each flow crossing the logical switches. If the interfaces were not bonded and separate logical switches (and service engines) processed data messages on the separate paths, a return data message could use a different path than an original message and state information would not be kept by the service engine and the stateful service would not function properly. **Figures 8-7** illustrate aspects of configuring and using an exemplary embodiment using bonded interfaces depicted in **Figure 9**.

[0045] **Figure 5** conceptually illustrates a process 500 for configuring a network edge device to process data messages using a set of bonded interfaces with deterministic mapping between bonded interfaces. Process 500, in some embodiments, is performed by a network manager computer (e.g., a virtualization controller, a network manager, etc.). Process 500 begins (at 510) by determining a mapping between interfaces of the network edge device. In some embodiments, the network edge device is a transparent (e.g., bump-in-the-wire) device that provides a service without changing the data message in a way that is apparent to the forwarding elements. Accordingly, each source forwarding element expects that a data message sent from a first interface of the source forwarding element/network edge device will arrive at a second interface of the network edge device/destination forwarding element and not at a third interface of

the network edge device/destination forwarding element. Because of this expectation, the normal function of a set of bonded interfaces (e.g., selecting an egress interface through a round robin, or other algorithm) would break the transparency of the network edge device (and, in some embodiments, would result in a forwarding failure) if a data message received at the first interface was received at the third interface.

[0046] After determining (at 510) the mapping of the interfaces, the process 500 generates (at 520) configuration information for bonding interfaces on the network edge device. In some embodiments, the configuration information includes an identification of sets of interfaces to bond into a single interface of a set of logical switches that bridge networks on either side of the network edge device. In some embodiments, the identified sets of interfaces include all the interfaces of the network edge device. In other embodiments, the sets of interfaces include all the interfaces that are not management interfaces. The configuration information, in some embodiments, includes configuration for a service engine that provides a stateful service (e.g., firewall, network address translation, virtual private network, etc.) associated with the set of logical switches.

[0047] The configuration information, in some embodiments, includes a set of mapping information (e.g., a mapping table or tables) to maintain deterministic forwarding between interfaces of the network edge device. In some embodiments, the set of mapping tables includes separate mapping tables for a north- and south-side bonding module. The generated configuration data is then provided (at 530) to the network edge device (or devices) to be used by the network edge device to configure the bonded interfaces and bonding modules executing on the network edge device and the process ends.

[0048] **Figure 6** conceptually illustrates a process 600 for configuring a network edge device to implement the bonded interfaces and bonding modules that, in some embodiments, maintain the mapping tables generated (at 520). The process 600 begins (at 610) by receiving configuration information for configuring the bonded interfaces and bonding modules of the network edge device. The configuration data, in some embodiments, is received from a network manager computer. In some embodiments, the configuration information identifies interfaces that are to be bonded into a single interface of a logical switch. In some embodiments, the configuration information identifies a first set of interfaces that are bonded into an interface of a south-side logical switch and a second set of interfaces that are bonded into an interface of a north-side logical

switch. The configuration information also includes, in some embodiments, a set of mapping tables that include mappings between individual interfaces of the network edge device.

[0049] After receiving (at 610) the configuration information, the process 600 configures (at 620) the network edge device to bond the identified interfaces into a single logical interface and maintain the mapping tables that, in some embodiments, are stored in a bonding module. Configuring the interface bonding, in some embodiments, includes configuring the set of logical switches that make up a bridge between the networks on either side of the network edge device (through the bonded interfaces). In some embodiments, a single logical switch is implemented with each bonded interface corresponding to a different interface of the implemented logical switch. In other embodiments, a pair of logical switches (i.e., a north-side logical switch and south side logical switch) is implemented with each set of bonded interfaces connecting to a corresponding logical switch. The logical switch configuration data, in some embodiments, includes configuration information for a service engine that provides a stateful service.

[0050] The network edge device, in some embodiments, also configures the bonding modules to store the received mapping table(s). In some embodiments, the network edge device stores a single mapping table that is accessible to each bonding module, while in other embodiments, each bonding module stores a separate mapping table. The mapping table, in some embodiments, is received as a single mapping table that is split into the two separate tables at the network edge device. Two exemplary sets of mapping tables are depicted in **Figure 7**. One of ordinary skill in the art will appreciate that there are many ways to store the mapping tables that are excluded by the description of the specific examples.

[0051] **Figure 7** illustrates two sets of mapping tables 710 and 720 that are used in different embodiments of the invention. Mapping table 710 is a single mapping table that includes an association between every possible ingress interface of the network edge device and an egress interface of the network edge device. Each bonding module uses table 710 to identify an egress interface for a data message associated with a particular ingress interface.

[0052] Mapping table set 720 includes two different mapping tables (e.g., 720A and 720B) stored by different bonding modules. Mapping table 720A is for a first bonding module bonding interfaces fp-eth0 and fp-eth1 and mapping table 720B is for a second bonding module bonding interfaces fp-eth2 and fp-eth3. Each individual table 720A and 720B stored by a bonding module only stores entries 721 for interfaces that it is not bonding, as data messages received from the

interfaces it bonds are provided to the logical switch and the other bonding module to determine an egress interface based on the ingress interface.

[0053] The network edge device, in some embodiments, receives a data message at a first interface and provides the data message to the logical switch interface associated with the bonding module. The bonding module associates the data message with an identifier of the ingress interface. In some embodiments, the association is inserted in the data message, while in other embodiments the identifier is appended to the data message. In yet other embodiments, the ingress interface is stored as metadata associated with the data message. A service engine is called by a logical switch, in some embodiments, to provide a stateful service. The service engine provides the stateful service and provides the data message to an egress interface associated with a second bonding module. The second bonding module uses a mapping table, in some embodiments, to determine an egress interface associated with the ingress interface. In some embodiments, the logical processing removes the ingress interface identifier and provides the data message to the egress module.

[0054] **Figure 8** conceptually illustrates a process 800 for using the configured network edge device to provide a stateful service. The process 800, is performed by the network edge device, but one of ordinary skill in the art will understand that different modules of the network edge device may be responsible for different operations of process 800. The process 800 begins (at 810) by receiving a data message at a first, ingress interface of the network edge device. The data message is received, in some embodiments, based on a selection of the interface from a plurality of possible interfaces (e.g., paths) by an upstream forwarding element (e.g., a router performing an ECMP operation).

[0055] After receiving (at 810) the data message at the first interface, the process 800, determines (at 820) an ingress interface associated with the data message and stores an identifier of the ingress interface for the data message. In some embodiments, the identified ingress interface is stored in metadata associated with the data message. The identifier, in some embodiments, is stored in the data packet itself (e.g., in an encapsulation header or other unused portion of the data message). One of ordinary skill in the art will appreciate that additional methods of storing the ingress interface identifier are possible. The data message associated with the identifier for the ingress interface is then provided to the logical switch interface associated with the bonded interfaces.

[0056] The logical switch (or service engine associated with the logical switch) provides (at 830) the stateful service to the data message. In some embodiments, the logical switch interface calls the service engine associated with the logical switch. The stateful service, in some embodiments, is at least one of a firewall, network address translation, and a VPN. If the stateful service is a firewall, providing the stateful service can lead to a data message being dropped. However, for the purposes of illustration, the description of **Figure 8** assumes that the data message is allowed. One of ordinary skill in the art will understand that if a data message is dropped the process 800 does not proceed to the next operation. Once the service engine has provided the stateful service the data message, in some embodiments, is returned to the logical switch interface from where it was called, while in other embodiments it is provided to an interface (e.g., the interface associated with the bonded interfaces) of the next logical switch on the egress side (for the particular data message) of the network edge device.

[0057] The bonding module then identifies (at 840) an egress interface for the data message based on the stored ingress interface identifier using the mapping table. The ingress interface identifier is associated with a particular egress interface in the mapping table. In some embodiments, different associations between interfaces can exist for different flows. In such embodiments, a separate mapping table is stored for each flow and is identified by a flow identifier such as a five tuple or a hash of header values of the data messages in the data message flow.

[0058] The bonding module then removes (at 850) the ingress interface identifier and provides the data message to the egress interface to forward to the next hop of the data message. The removal of the identifier, in some embodiments, is optional. For example, if the identifier is stored in an unused portion of a data message header, such that it will not affect the subsequent forwarding operations or the content of the data message, there is no need to remove the identifier and operation 850 will merely provide the data message to the identified egress interface. Once the data message is provided to the egress interface and forwarded to the next hop, the process ends.

[0059] **Figure 9** illustrates a particular network edge device 910 in a set of network edge devices connecting a first north edge routing element 920 (in a set of north edge routing elements) connected to interfaces 905A (e.g., fp-eth0) and 905B (e.g., fp-eth1) and a south edge routing element 930 (in a set of south edge routing elements) connected to interfaces 905C (e.g., fp-eth2) and 905D (e.g., fp-eth3). The paths between routing elements 920 and 930 (i.e., 905A to 905C and

905B to 905D) are equal cost in the illustrated embodiment and either can be selected for transmitting data messages between the routing (e.g., forwarding) elements 920 and 930. The network edge device 910 includes bonded interfaces indicated by the dotted ellipse 960 and bonding modules 912A and 912B that support the bonded interfaces on the north and south sides of the network edge device 910, respectively. The bonding modules 912 are implemented, in some embodiments, as part of a set of logical switches 914 that (1) provide a bridge between the north and south networks and (2) provide a stateful service for data messages crossing the bridge. The bridge in **Figure 9** is depicted as two logical switches 914A and 914B, however, in some other embodiments, the bridge is implemented using a single logical switch with separate interfaces (e.g., ports) connecting to the north and south side bonded interfaces.

[0060] The logical switches 914 provide a stateful service to data messages passing through the set of logical switches 914 between the two networks. The stateful service, in some embodiments is provided by a service engine 913 that provides a firewall service, network address translation, virtual private network, etc. that may require processing all the data messages in a data message flow for proper performance. Additional services, in some embodiments, are provided by additional service engines (not shown for clarity). By bonding the two equal cost paths at the logical switch interface, the invention ensures that all data messages that are sent along either path are processed by the same service engine 913 to ensure that state information maintained at the service engine is complete. As path 940A connects interfaces 905A and 905C and path 940B connects interfaces 905B and 905D, the bonding modules store a mapping table that ensures that traffic received at any particular interface are forwarded to the destination using the corresponding interface. One of ordinary skill in the art would appreciate that in some embodiments there are many more than two equal cost paths and that two paths are shown for clarity in **Figure 9**.

[0061] **Figure 10** illustrates a second embodiment of a network in which the invention is carried out. **Figure 10** illustrates a particular network edge device 1010 connecting a first set of north edge routing elements 1020 connected to interfaces 1005A-C and a second set of south edge routing elements 1030 connected to interfaces 1005D-F. The sets of routing elements 1020 and 1030 help to connect other routing elements 1025 in the networks on either side of the network edge device 1010. The paths between routing (e.g., forwarding) elements in the sets of routing elements 1020 and 1030 (i.e., 1005A to 1005D and 1005B to 1005E and 1005C to 1005F) are equal cost in the illustrated embodiment and any of them can be selected for transmitting data

messages between the routing elements 1025. The network edge device 1010 includes bonding module 1012 that supports the bonded interfaces on the north and south sides of the network edge device 1010, respectively. In some embodiments, the bonding module is a function of the logical switch that implements the interface bonding and stores the mapping tables or otherwise maintains the connections between specific interfaces. The bonding module 1012 supports different interfaces of a single logical switch 1014 that (1) provides a bridge between the north and south networks and (2) provides a stateful service for data messages crossing the bridge using service engine 1013. The bridge in **Figure 10** is depicted as one logical switch 1014, however, in some other embodiments, the bridge is implemented using multiple logical switches as in **Figure 9**.

[0062] The logical switch 1014 provides a stateful service to data messages passing through the logical switch 1014 between the two networks. The stateful service, in some embodiments is provided by the service engine 1013 that provides a firewall service, network address translation, virtual private network, etc. that may require processing all the data messages in a data message flow for proper performance. Additional services, in some embodiments, are provided by additional service engines (not shown for clarity). By bonding the two equal cost paths, the invention ensures that all data messages that are sent along either path are processed by the same service engine 1013 to ensure that state information maintained at the service engine is complete. As path 1040A connects interfaces 1005A and 1005D, path 1040B connects interfaces 1005B and 1005E, and path 1040C connects interfaces 1005C and 1005F, the bonding modules store a mapping table that ensures that traffic received at any particular interface are forwarded to the destination using the corresponding interface. One of ordinary skill in the art would appreciate that in some embodiments there are many more than three equal cost paths and that three paths are shown for clarity in **Figure 10**.

[0063] The network edge device, in some embodiments, is a primary network edge device and a secondary network edge device is implemented as a standby network edge device. The secondary network edge device replicates the interfaces and associations of the primary network edge device and is available to assume the functions of the primary network edge device in the case of a failure of the primary network edge device. In some embodiments, failure of the primary network edge device is defined by a user. The definition of failure, in some embodiments, is defined by a user-supplied policy (e.g., a threshold for a given metric). In other embodiments, failover occurs when all paths through the primary network edge device fail.

[0064] **Figure 11** illustrates an embodiment including active and standby network edge devices 1110A and 1110B. **Figure 11** illustrates pairs of edge routing (forwarding) elements 1120 and 1130 connected by three paths through each of an active network edge device 1110A and a standby network edge device 1110B. The bonding is indicated by 1160 which, in some embodiments, represents a standard bonding operation by each logical switch 1114 for the logical interfaces associated with the bonded interfaces of edge devices 1110. Similar to **Figure 9**, the network edge devices 1110 provide a set of bonded interfaces 1105A-C (on the north side) and 1105D-F (on the south side) and a set of logical switches 1114 that connect the pairs of edge routing (forwarding) elements 1120 and 1130. A service engine that provides stateful services is omitted from the illustrated embodiment for clarity but is understood to be present as in **Figure 9 and 10**. As in **Figures 9 and 10**, the interfaces of the network edge devices 1110 are deterministically paired such that data messages entering interface 1105A exits from 1105D (bonding modules supporting the deterministic mapping are not shown).

[0065] The standby edge device 1110B, in some embodiments, has an equivalent set of interfaces (1105A-C on the north side and 1105D-F on the south side) and advertises the cost of the path through the interfaces of the standby edge device 1110B as having a higher cost than the paths through the equivalent interfaces of the active edge device 1110A such that all data messages destined for the pair of edge routers 1130 passes through active edge device 1110A. An active-standby configuration is often used when stateful services are provided at the edge device. The failover process of similar embodiments employing active and standby network edge devices is described in **Figure 12**.

[0066] Because the edge device is transparent to the routing elements, in some embodiments, when a particular connection to the edge device fails (e.g., a connection between the north edge routing element 1120B and interface 1105B on edge device 1110A) the edge device 1110 is responsible for bringing down a corresponding connection on the other side of the edge device 1110 (e.g., the connection between the south edge routing element 1130A and interface 1105E on edge device 1110A). In the case of a failure of either interface 1105B or 1105E, traffic from within the network on the north side can still reach destinations in the network on the south side and vice versa using edge routing element 1130B. Accordingly, as will be described in relation to **Figure 12** the standby edge device 1110B does not necessarily need to become the active edge device.

[0067] **Figure 12** conceptually illustrates a process 1200 for failure handling (i.e., failover). The process 1200 is performed by a network edge device, or network manager computer that monitors the health of an active network edge device (e.g., network edge device 1110A). The process 1200 begins by receiving (at 1210) a failover policy that defines how different failures are to be handled. In some embodiments, the failover policy is received from a user and specifies threshold values for a set of connection (link) metrics (e.g., dropped data messages, bandwidth, a current load metric, a reliability metric, and a capacity metric) that trigger the standby network edge device to become the active network edge device. The change in active network edge devices, in some embodiments, applies only to new data message flows with existing data message flows directed to the previously active network edge device. In other embodiments, all data message flows are directed to the new active network edge device upon failover.

[0068] In some embodiments, link failure is determined at the level of the logical switch (or switches) implemented by the network edge device (e.g., at the level of a logical switch port or interface). In previous embodiments using different sets of logical switches for each set of associated (paired) interfaces (i.e., not using bonding with deterministic forwarding), failure detection required separate monitoring and failover which introduced additional complexity. In the illustrated embodiment, monitoring a single logical switch failure is simpler and does not require failover until all paths have failed. Additionally, as described above, a user can set a policy for failover other than complete failure by setting up conditional failover based on connection metrics of the combined multiple links between the networks. The network topology is therefore maintained for longer periods than would be the case if a single link failure out of multiple links could trigger a failover event.

[0069] The process 1200 then proceeds to monitor (at 1220) the link (representing the combined links over all the bonded interfaces). In some embodiments, the link monitoring includes monitoring the metrics specified in a default or user-specified policy. Monitoring, in some embodiments, is performed by a network manager computer that initiates a failover process when the conditions for failover are met. In some embodiments, monitoring is performed by the standby network edge device to determine if it needs to initiate a failover process and become the new active network edge device.

[0070] The process 1200 determines (at 1230) whether a failover process has been triggered based on the monitored metrics and the received failover policy. If the policy has not

been met, i.e., the metrics have not crossed the specified threshold from an acceptable to an unacceptable value, the process returns to monitor (at 1220) the link. In some embodiments, the policy specifies a plurality of thresholds such that a first threshold may be crossed for a specified period before triggering a failover process, while a second threshold is specified such that any crossing of the second threshold triggers a failover process.

[0071] If the process 1200 determines (at 1230) that the failover policy triggers a failover based on the monitored metrics, the process initiates (at 1240) a failover process. In some embodiments a failover process causes a formerly-standby network edge device to advertise its availability with a lower cost and a formerly-active network edge device to advertise its availability with a higher cost such that all traffic is now directed to new active (formerly-standby) device. One of ordinary skill in the art will appreciate that there are additional ways of handling failover that are not described here.

[0072] In some embodiments, instead of forcing all data messages to cross a single bridge, each interface associated with a different bridge calls a service engine based on identifiers included in data messages received at the interface. Each data message flow is associated with a particular identifier that is associated with a particular service engine instance that provides the stateful service. In some embodiments, the interface that receives a data message identifies a service engine to provide the stateful service and provides the data message to the identified service engine. After processing the data message, the service engine provides the data message to the egress interface associated with the ingress interface.

[0073] In other embodiments, each logical switch has an associated service engine and uses a single connection tracker to maintain state information accessible to each service engine to ensure that all data messages of a data message flow are processed using current state information. The connection tracker, in some embodiments, is implemented as a lightweight database in the network edge device that is queried by each service engine in the network edge device. In some embodiments, the connection tracker is implemented as a simple table. One of ordinary skill in the art will appreciate that these are just two examples of many ways that a connection tracker may be implemented.

[0074] In some embodiments, each service engine is programmed for rules associated with different logical networks and the set of rules associated with a particular data message is identified using an identifier stored in, or associated with, the data message (e.g., a VLAN tag). The service

engine uses the identified set of rules and the information in the connection tracker to provide the stateful service. After processing the data message, the service engine provides the data message to the logical switch (e.g., to the egress interface of the logical switch) with which the service engine is associated.

[0075] **Figure 13** conceptually illustrates a process 1300 for configuring a network edge device to implement an embodiment of the invention that allows for separate logical switches for different equal cost paths. The process 1300, is described from the perspective of a single network edge device that is being configured, but one of ordinary skill in the art will appreciate that, similarly to the process 500 described in **Figure 5**, a network manager computer, in some embodiments, generates the configuration data and sends it to the network edge device. Additionally, in some embodiments, the process 1300 is performed by each network edge device in a set of multiple edge devices implementing the DMZ. One of ordinary skill in the art will understand that there are additional internal components of a network edge device that may perform different parts of the process 1300 (e.g., an agent for receiving the configuration data and a local controller for configuring the network edge device to configure the network edge device according to the received configuration data).

[0076] The process 1300 begins by receiving (at 1310) configuration information for configuring (implementing) a set of service engine instances that each provide a set of stateful services on the network edge device. In some embodiments, each service engine instance is associated with a different identifier. The identifier, in some embodiments, is an identifier associated with data message flows such that each data message in a particular data message flow will have a same identifier and be associated with a same service engine instance. In some embodiments, the identifier is associated with a network from which the data message flow is received (e.g., a VLAN or VXLAN tag, an IP address prefix, etc.). Thus, a connection tracker of the service engine associated with a particular identifier will process all data messages of a data message flow and maintain current and correct state information.

[0077] In other embodiments, each service engine is associated with a particular logical switch and includes multiple sets of rules for multiple different logical networks serviced by the network edge device. The configuration information also includes, in some embodiments, configuration information for a shared connection tracker that maintains state information for data message flows processed by all the service engines and logical switches implemented on the

network edge device. By using the shared connection tracker that is accessible from each service engine/logical switch, the network edge device allows for data messages of a same data message flow to be processed by different logical switches and service engines while maintaining state data that is current and complete. This is in contrast to a situation in which each logical switch/service engine uses a separate connection tracker, in which case a data message received at a particular logical switch in a data message flow initiated using a different logical switch will not be recognized as belonging to an existing flow and may be dropped or processed differently than the other data messages in the data message flow. In some embodiments, the state information is maintained for each data message flow and includes a network identifier (e.g., a VLAN or VXLAN tag an IP address prefix, etc. associated with the network) along with flow identifiers (e.g., an n-tuple) that uniquely identifies the data message flow.

[0078] The network edge device then receives (at 1320) configuration information to configure a plurality of bridges between particular sets of interfaces of the network edge device. In some embodiments, the configuration information includes a set of logical switches to implement, the identity of the bridged interfaces, and linked interfaces (e.g., linked into a link aggregation group (LAG) using a link aggregation control protocol (LACP)). In some embodiments, the configuration information received also includes configuration information to configure the interfaces (e.g., the physical interfaces of the network edge device or the logical interfaces of the logical switches) to call particular service engine instances based on identifiers included in received data messages. In some embodiments, the configuration information for a particular interface includes a particular service engine called by the interface and the identifier of a received data message is used by the service engine to determine the set of service rules that apply to the received data message. One of ordinary skill in the art will appreciate that the order of receiving the configuration information may be reversed or combined into a single operation in some embodiments.

[0079] The network edge device then configures (at 1330) the bridges and service engine instances specified by the configuration information. The network edge device, in some embodiments, now includes a set of bridged interfaces and a set of service engine instances that can be called from any interface. In other embodiments, the network edge device now includes a common connection tracker shared by the set of bridged interfaces (i.e., by the logical switches) and the set of service engine instances associated with the set of bridged interfaces. As will be

described in relation to **Figures 14-17**, the network edge device is now ready to implement the invention and the process ends.

[0080] **Figure 14** conceptually illustrates a process 1400 for processing data messages at a network edge device configured according to an embodiment of process 1300 in which a set of shared service engines are implemented for the set of logical switches. The process 1400 is performed by the network edge device, in some embodiments, and one of ordinary skill in the art will appreciate that different modules or components of the network edge device perform different operations of process 1400. The process 1400 begins by receiving (at 1410) a data message of a data message flow at a particular interface of the network edge device. The data message includes a particular identifier, e.g., a VLAN tag, that is associated with every data message in the data message flow.

[0081] The process 1400 then identifies (at 1420) the service engine instance that is associated with the particular identifier. In some embodiments, the service engine instances are associated with particular identifiers in a table accessible by the interfaces. The association, in some embodiments, is between a VLAN tag and service engine. In other embodiments, the association is based on other data message flow attributes that do not change for a particular flow.

[0082] After the service engine instance is identified, the process 1400 calls (at 1430) the identified service engine. In some embodiments, the identified service engine instance is called from the interface of the network edge device. In other embodiments, the identified service engine is called from a logical switch interface of a set of logical switches that bridges (connects) two interfaces of the network edge device.

[0083] The identified service engine provides (at 1440) the stateful service to the data message and forwards the data message to its destination. In some embodiments, the data message is returned to the ingress interface to be sent across the bridge to the egress interface. In other embodiments, the data message is forwarded through the logical switch interface to the egress interface and ultimately to the destination. Once the service engine has forwarded the data message to its destination, the process 1400 ends. One of ordinary skill in the art will appreciate that, by using identifiers that are constant for a particular data message flow to determine a service engine instance, all data messages of a data message flow are processed by a same service engine instance and state information maintained at the service engine instance is complete.

[0084] **Figure 15** illustrates an embodiment in which the process 1400 is carried out. **Figure 15** illustrates a network edge device 1510 providing a set of paths between a first set of north edge routing elements 1520 connected to interfaces 1505A-E and a second set of south edge routing elements 1530 connected to interfaces 1505F-J. Each set of associated interfaces (e.g., 1505A and 1505F) is connected by a single logical switch and includes, in the depicted embodiment, sets of interfaces (i.e., 1505D / 1505E and 1505I / 1505J) configured as a link aggregation group 1540 with the aggregated interfaces also connecting through a single logical switch 1514. **Figure 15** also illustrates a set of service engine instances 1545 that execute on the network edge device and can be called from any interface (e.g., any interface configured to call a service engine) to provide a stateful service for data messages received at the interface.

[0085] **Figure 15** also illustrates exemplary routing elements 1525A and 1525B in the external (north-side) and internal (south-side) networks. In some embodiments, external routing element 1525A considers each north edge routing element 1520 to be an equal cost path to a destination behind (or below) the internal routing element 1525B. Accordingly, in trying to reach the destination, external routing element 1525A may select either of north edge routers 1520 as a next hop and the decision may change over the life of a flow based on factors not having to do with any of the illustrated system components. Similarly, the north edge routing elements 1520, in some embodiments, see each path to the internal routing element 1525B to be equal cost and may select either of south edge routers 1530 as a next hop and the decision may change in the middle of a flow based on factors not having to do with any of the illustrated system components. In such embodiments practicing the process 1400, the change in path selection does not affect the provision of the stateful service as the same service engine is called by each logical switch based on the flow identifiers that do not change over the life of the flow.

[0086] In some embodiments, each logical switch has an associated service engine and uses a single connection tracker to maintain state information accessible to each service engine to ensure that all data messages of a data message flow are processed using current state information. In some embodiments, each service engine is programmed for rules associated with different logical networks and the set of rules associated with a particular data message is identified using an identifier stored in, or associated with, the data message (e.g., a VLAN tag). The service engine uses the identified set of rules and the information in the connection tracker to provide the stateful service. After processing the data message, the service engine provides the data message to the

logical switch (e.g., to the egress interface of the logical switch) with which the service engine is associated.

[0087] **Figure 16** conceptually illustrates a process 1600 for processing data messages at a network edge device configured according to an embodiment of process 1300 in which a common connection tracker is shared by each logical switch and the service engine associated with each logical switch. The process 1600 is performed by the network edge device, in some embodiments, and one of ordinary skill in the art will appreciate that different modules or components of the network edge device perform different operations of process 1600. The process 1600 begins by receiving (at 1610) a data message of a data message flow at a particular interface of the network edge device. The data message includes a particular identifier, e.g., a VLAN tag, that is associated with every data message in the data message flow.

[0088] The process 1600 then records (at 1620) the data message state information in the shared connection tracker. The state information, in some embodiments, includes a set of identifiers of the data message (e.g., an n-tuple) and an identifier of the network associated with the data message (e.g., a VLAN or VXLAN tag, an IP address source or destination subnet, etc.) that are used to unambiguously identify the data message flow in the shared connection tracker. One example of state information recorded by the process 1600 occurs upon receiving an ACK data message for a TCP connection. In response to receiving the ACK data message, the process 1600 updates an entry in the connection tracker to identify that the TCP connection has been established. This behavior is distinct from embodiments using separate connection trackers for each logical switch in which an ACK data message received on a different interface than the interface on which a SYN or SYN/ACK data message was received would not have the required state information to identify that the TCP connection was established. One of ordinary skill in the art will understand that additional state information is maintained and updated based on received data messages of different data message flows.

[0089] Once the data message state information is recorded in the shared connection tracker, the process calls (at 1630) the service engine instance that is associated with the particular interface (i.e., logical switch) on which the data message was received. In some embodiments, the identified service engine instance is called from the interface of the network edge device. In other embodiments, the identified service engine is called from a logical switch interface of a set of logical switches that bridges (connects) two interfaces of the network edge device. In some

embodiments, the service engine instances store multiple sets of service rules that apply to different networks (e.g., logical networks, or subnets) that are associated with particular identifiers. One of ordinary skill in the art will appreciate that, in some embodiments, the connection tracker information is accessed (and updated) by the service engine as a part of providing the service.

[0090] Once the data message has been sent (at 1630) to the particular service engine, the process 1600 identifies (at 1640) the set of service rules applicable to data message based on an identifier associated with the data message. The identifier, in some embodiments, is a network identifier (e.g., a VLAN or VXLAN tag, an IP address prefix, etc.) that is associated with each data message in a data message flow and with a set of service rules stored by the service engine. In other embodiments, the association is based on other data message flow attributes that do not change for a particular flow. In some embodiments, the connection tracker is used to identify a specific rule or action previously applied to the data message flow based on the identifier associated with the data message (e.g., an n-tuple, or an n-tuple and network identifier).

[0091] The service engine provides (at 1650) the stateful service to the data message based on the identified set of applicable rules and forwards the data message to its destination. In some embodiments, the data message is returned to the ingress interface to be sent across the bridge to the egress interface. In other embodiments, the data message is forwarded through the logical switch interface to the egress interface and ultimately to the destination. Once the service engine has forwarded the data message to its destination, the process 1600 ends. One of ordinary skill in the art will appreciate that, by using identifiers that are constant for a particular data message flow to determine a set of service engine instance, all data messages of a data message flow are processed by a same service engine instance and state information maintained at the service engine instance is complete.

[0092] **Figure 17** illustrates an embodiment in which the process 1600 is carried out. **Figure 17** illustrates a network edge device 1710 providing a set of paths between a first set of north edge routing elements 1720 connected to interfaces 1705A-E and a second set of south edge routing elements 1730 connected to interfaces 1705F-J. The north and south edge routing elements, in some embodiments, are centralized logical routers implemented by edge devices (e.g., NSX edge) that, in some embodiments, also implement transit logical switches and distributed routers as described in U.S. Patent Publication 2016/0226754. The edge routing elements, in some

embodiments, are configured in an active-active configuration, while in other embodiments, they are configured in an active-standby configuration.

[0093] Each set of associated interfaces (e.g., 1705A and 1705F) is connected by a single logical switch and includes, in the depicted embodiment, sets of interfaces (i.e., 1705D / 1705E and 1705I / 1705J) configured as a link aggregation group 1740 with the aggregated interfaces also connecting through a single logical switch 1714. **Figure 17** also illustrates a set of identical service engine instances 1745 that execute on the network edge device and are associated with interfaces of different logical switches to provide a stateful service for data messages received at the interface. **Figure 17** also depicts a shared connection tracker 1746 used by each logical switch to maintain the state of the data message flows traversing network edge device 1710. As described in relation to **Figure 1**, the service engines 1745, in some embodiments, are replaced by a set of service engines for providing different stateful or stateless services and only one service engine is shown for clarity. In embodiments using additional service engines associated with each logical switch, the service engines providing different services all use the connection tracker 1746. In other embodiments, each service engine providing a particular stateful service is associated with (e.g., calls or queries) a particular connection tracker common to the service engines providing the particular stateful service.

[0094] **Figure 17** also illustrates exemplary routing elements 1725A and 1725B in the external (north-side) and internal (south-side) networks. In some embodiments, external routing element 1725A considers each north edge routing element 1720 to be an equal cost path to a destination behind (or below) the internal routing element 1725B. Accordingly, in trying to reach the destination, external routing element 1725A may select either of north edge routers 1720 as a next hop and the decision may change over the life of a flow based on factors not having to do with any of the illustrated system components. Similarly, the north edge routing elements 1720, in some embodiments, see each path to the internal routing element 1725B to be equal cost and may select either of south edge routers 1730 as a next hop and the decision may change in the middle of a flow based on factors not having to do with any of the illustrated system components. In such embodiments practicing the process 1600, the change in path selection does not affect the provision of the stateful service because using the shared connection tracker 1746 ensures that correct state data is maintained for data messages of a same data message flow that arrive on different interfaces. One of ordinary skill in the art will understand that while **Figure 17** illustrates only one

network edge device 1710 forming a connection between two sets of edge routing elements, in some embodiments each of the depicted elements represents a set of multiple similar elements forming multiple connections between multiple different networks through multiple edge devices.

[0095] **Figure 18** illustrates the system 1800 of some embodiments. As shown, this system includes multiple virtualized hosts 1805 and 1810 (hosting guest virtual machines 1850 and software forwarding elements (SFEs) 1852), a set of network manager computers 1820, and a network edge device 1815. The virtualized hosts 1805 and 1810 host compute nodes that can be sources and destinations of data messages sent through network 1875 and network edge device 1815 to or from a compute node in network 1885. The network edge device 1815 is shown executing a software forwarding element (SFE) implementing a logical edge router for the managed network including hosts 1805-1815 connected by network 1875. The edge device 1840 executes a set of logical switches 1814 and a set of service engines (e.g., service engine instances) 1845 as described above. As shown in **Figure 18**, the hosts 1805 and 1810, the controller set 1820, and the network edge device 1815 communicatively couple through a network 1875, which can include a local area network (LAN), a wide area network (WAN) or a network of networks (e.g., Internet). Network 1885, in some embodiments, is one of a second managed network in a same datacenter, a managed network in a second datacenter, and an external network.

[0096] The set of network manager computers 1820 provide control and management functionality for defining and managing the instantiation of one or more GVMs on each host (for the purposes of this discussion, network controllers 1820 includes both management plane and control plane controllers). These controllers are also responsible, in some embodiments, for configuring the edge device to provide the functionality described above. These controllers, in some embodiments, also provide control and management functionality for defining and managing multiple logical networks that are defined on the common software forwarding elements of the hosts.

[0097] Many of the above-described features and applications are implemented as software processes that are specified as a set of instructions recorded on a computer readable storage medium (also referred to as computer readable medium). When these instructions are executed by one or more processing unit(s) (e.g., one or more processors, cores of processors, or other processing units), they cause the processing unit(s) to perform the actions indicated in the instructions. Examples of computer readable media include, but are not limited to, CD-ROMs,

flash drives, RAM chips, hard drives, EPROMs, etc. The computer readable media does not include carrier waves and electronic signals passing wirelessly or over wired connections.

[0098] In this specification, the term “software” is meant to include firmware residing in read-only memory or applications stored in magnetic storage, which can be read into memory for processing by a processor. Also, in some embodiments, multiple software inventions can be implemented as sub-parts of a larger program while remaining distinct software inventions. In some embodiments, multiple software inventions can also be implemented as separate programs. Finally, any combination of separate programs that together implement a software invention described here is within the scope of the invention. In some embodiments, the software programs, when installed to operate on one or more electronic systems, define one or more specific machine implementations that execute and perform the operations of the software programs.

[0099] **Figure 19** conceptually illustrates a computer system 1900 with which some embodiments of the invention are implemented. The computer system 1900 can be used to implement any of the above-described hosts, controllers, and managers. As such, it can be used to execute any of the above described processes. This computer system includes various types of non-transitory machine readable media and interfaces for various other types of machine readable media. Computer system 1900 includes a bus 1905, processing unit(s) 1910, a system memory 1925, a read-only memory 1930, a permanent storage device 1935, input devices 1940, and output devices 1945.

[0100] The bus 1905 collectively represents all system, peripheral, and chipset buses that communicatively connect the numerous internal devices of the computer system 1900. For instance, the bus 1905 communicatively connects the processing unit(s) 1910 with the read-only memory 1930, the system memory 1925, and the permanent storage device 1935.

[0101] From these various memory units, the processing unit(s) 1910 retrieve instructions to execute and data to process in order to execute the processes of the invention. The processing unit(s) may be a single processor or a multi-core processor in different embodiments. The read-only-memory (ROM) 1930 stores static data and instructions that are needed by the processing unit(s) 1910 and other modules of the computer system. The permanent storage device 1935, on the other hand, is a read-and-write memory device. This device is a non-volatile memory unit that stores instructions and data even when the computer system 1900 is off. Some embodiments of the

invention use a mass-storage device (such as a magnetic or optical disk and its corresponding disk drive) as the permanent storage device 1935.

[0102] Other embodiments use a removable storage device (such as a floppy disk, flash drive, etc.) as the permanent storage device. Like the permanent storage device 1935, the system memory 1925 is a read-and-write memory device. However, unlike storage device 1935, the system memory is a volatile read-and-write memory, such a random access memory. The system memory stores some of the instructions and data that the processor needs at runtime. In some embodiments, the invention's processes are stored in the system memory 1925, the permanent storage device 1935, and/or the read-only memory 1930. From these various memory units, the processing unit(s) 1910 retrieve instructions to execute and data to process in order to execute the processes of some embodiments.

[0103] The bus 1905 also connects to the input and output devices 1940 and 1945. The input devices enable the user to communicate information and select commands to the computer system. The input devices 1940 include alphanumeric keyboards and pointing devices (also called "cursor control devices"). The output devices 1945 display images generated by the computer system. The output devices include printers and display devices, such as cathode ray tubes (CRT) or liquid crystal displays (LCD). Some embodiments include devices such as a touchscreen that function as both input and output devices.

[0104] Finally, as shown in **Figure 19**, bus 1905 also couples computer system 1900 to a network 1965 through a network adapter (not shown). In this manner, the computer can be a part of a network of computers (such as a local area network ("LAN"), a wide area network ("WAN"), or an Intranet, or a network of networks, such as the Internet. Any or all components of computer system 1900 may be used in conjunction with the invention.

[0105] Some embodiments include electronic components, such as microprocessors, storage and memory that store computer program instructions in a machine-readable or computer-readable medium (alternatively referred to as computer-readable storage media, machine-readable media, or machine-readable storage media). Some examples of such computer-readable media include RAM, ROM, read-only compact discs (CD-ROM), recordable compact discs (CD-R), rewritable compact discs (CD-RW), read-only digital versatile discs (e.g., DVD-ROM, dual-layer DVD-ROM), a variety of recordable/rewritable DVDs (e.g., DVD-RAM, DVD-RW, DVD+RW, etc.), flash memory (e.g., SD cards, mini-SD cards, micro-SD cards, etc.), magnetic and/or solid

state hard drives, read-only and recordable Blu-Ray® discs, ultra-density optical discs, any other optical or magnetic media, and floppy disks. The computer-readable media may store a computer program that is executable by at least one processing unit and includes sets of instructions for performing various operations. Examples of computer programs or computer code include machine code, such as is produced by a compiler, and files including higher-level code that are executed by a computer, an electronic component, or a microprocessor using an interpreter.

[0106] While the above discussion primarily refers to microprocessor or multi-core processors that execute software, some embodiments are performed by one or more integrated circuits, such as application specific integrated circuits (ASICs) or field programmable gate arrays (FPGAs). In some embodiments, such integrated circuits execute instructions that are stored on the circuit itself.

[0107] As used in this specification, the terms “computer”, “server”, “processor”, and “memory” all refer to electronic or other technological devices. These terms exclude people or groups of people. For the purposes of the specification, the terms display or displaying means displaying on an electronic device. As used in this specification, the terms “computer readable medium,” “computer readable media,” and “machine readable medium” are entirely restricted to tangible, physical objects that store information in a form that is readable by a computer. These terms exclude any wireless signals, wired download signals, and any other ephemeral or transitory signals.

[0108] While the invention has been described with reference to numerous specific details, one of ordinary skill in the art will recognize that the invention can be embodied in other specific forms without departing from the spirit of the invention. For instance, several figures conceptually illustrate processes. The specific operations of these processes may not be performed in the exact order shown and described. The specific operations may not be performed in one continuous series of operations, and different specific operations may be performed in different embodiments. Furthermore, the process could be implemented using several sub-processes, or as part of a larger macro process. Thus, one of ordinary skill in the art would understand that the invention is not to be limited by the foregoing illustrative details, but rather is to be defined by the appended claims.

CLAIMS

We claim:

1. A method for transparently providing a set of network services at a network edge device providing a demilitarized zone between first and second gateway devices for a first, logical network and a second, physical network respectively, the network edge device, and first and second gateway devices executing in the data center, the method comprising:

receiving, at a first plurality of interfaces of the network edge device, a set of data messages from a first plurality of interfaces of the first gateway device of the logical network destined for a corresponding plurality of interfaces of the second gateway device of the physical network;

providing the set of network services on each data message in the set of data messages; and

forwarding each data message in the set of data messages received at a particular interface of the network edge device to a particular interface of the second gateway device of the received data message, wherein each data message in the set of data messages received at the particular interface of the network edge device is destined to a same interface of the second gateway device, and the first and second gateway devices function as if the network edge device is not interposed between the gateway devices.

2. The method of claim 1, wherein the network edge device implements a logical switch at which the set of network services is provided for all data messages sent between the first and second gateway devices .

3. The method of claim 1, wherein the demilitarized zone provides a layer of protection between the first and second networks.

4. The method of claim 1, wherein the network edge device implements a plurality of logical switches at which at least one network service is provided.

5. The method of claim 4, wherein each logical switch connects a first interface of the network edge device connected to the first gateway device to a second interface of the network edge device connected to the second gateway device.

6. The method of claim 5, wherein the network edge device implements a plurality of service engine instances that are called by the plurality of logical switches.

7. The method of claim 6, wherein the first gateway device is a gateway device for a plurality of logical networks and logical switches implemented by the network edge device that process data messages sent from a particular logical network call a particular service engine.

8. The method of claim 6, wherein the network edge device implements a shared connection tracker that is accessible to each logical network in the plurality of logical networks and maintains state data for each data message flow processed by the network edge device.

9. The method of claim 8, wherein the first gateway device is a gateway device for a plurality of logical networks and service engines implemented by the network edge device maintain sets of service rules for each logical network.

10. The method of claim 9, wherein for a received data message the service engine identifies a set of rules applicable to the received data message based on a logical network identifier associated with the data message.

11. A non-transitory machine readable medium storing a program for execution by a set of processing units of a network edge device, the program for transparently providing a set of network services at a network edge device providing a demilitarized zone between first and second gateway devices for a first, logical network and a second, physical network respectively, the network edge device, and first and second gateway devices executing in the data center, the program comprising sets of instructions for:

receiving, at a first plurality of interfaces of the network edge device, a set of data messages from a first plurality of interfaces of the first gateway device of the logical network destined for a corresponding plurality of interfaces of the second gateway device of the physical network;

providing the set of network services on each data message in the set of data messages; and

forwarding each data message in the set of data messages received at a particular interface of the network edge device to a particular interface of the second gateway device of the received data message, wherein each data message in the set of data messages received at a particular interface of the network edge device is destined to a same interface of the second gateway device, and the first and second gateway devices function as if the network edge device is not interposed between the gateway devices.

12. The non-transitory machine readable medium of claim 11, wherein the network edge device implements a logical switch at which the set of network services is provided for all data messages sent between the first and second gateway devices .

13. The non-transitory machine readable medium of claim 11, wherein the demilitarized zone provides a layer of protection between the first and second networks.

14. The non-transitory machine readable medium of claim 11, wherein the network edge device implements a plurality of logical switches at which at least one network service is provided.

15. The non-transitory machine readable medium of claim 14, wherein each logical switch connects a first interface of the network edge device connected to the first gateway device to a second interface of the network edge device connected to the second gateway device.

16. The non-transitory machine readable medium of claim 15, wherein the network edge device implements a plurality of service engine instances that are called by the plurality of logical switches.

17. The non-transitory machine readable medium of claim 6, wherein the first gateway device is a gateway device for a plurality of logical networks and logical switches implemented by the network edge device that process data messages sent from a particular logical network call a particular service engine.

18. The non-transitory machine readable medium of claim 16, wherein the network edge device implements a shared connection tracker that is accessible to each logical network in the plurality of logical networks and maintains state data for each data message flow processed by the network edge device.

19. The non-transitory machine readable medium of claim 18, wherein the first gateway device is a gateway device for a plurality of logical networks and service engines implemented by the network edge device maintain sets of service rules for each logical network.

20. The non-transitory machine readable medium of claim 19, wherein for a received data message the service engine identifies a set of rules applicable to the received data message based on a logical network identifier associated with the data message.

100




Figure 1

200




Figure 2




Figure 3




Figure 4

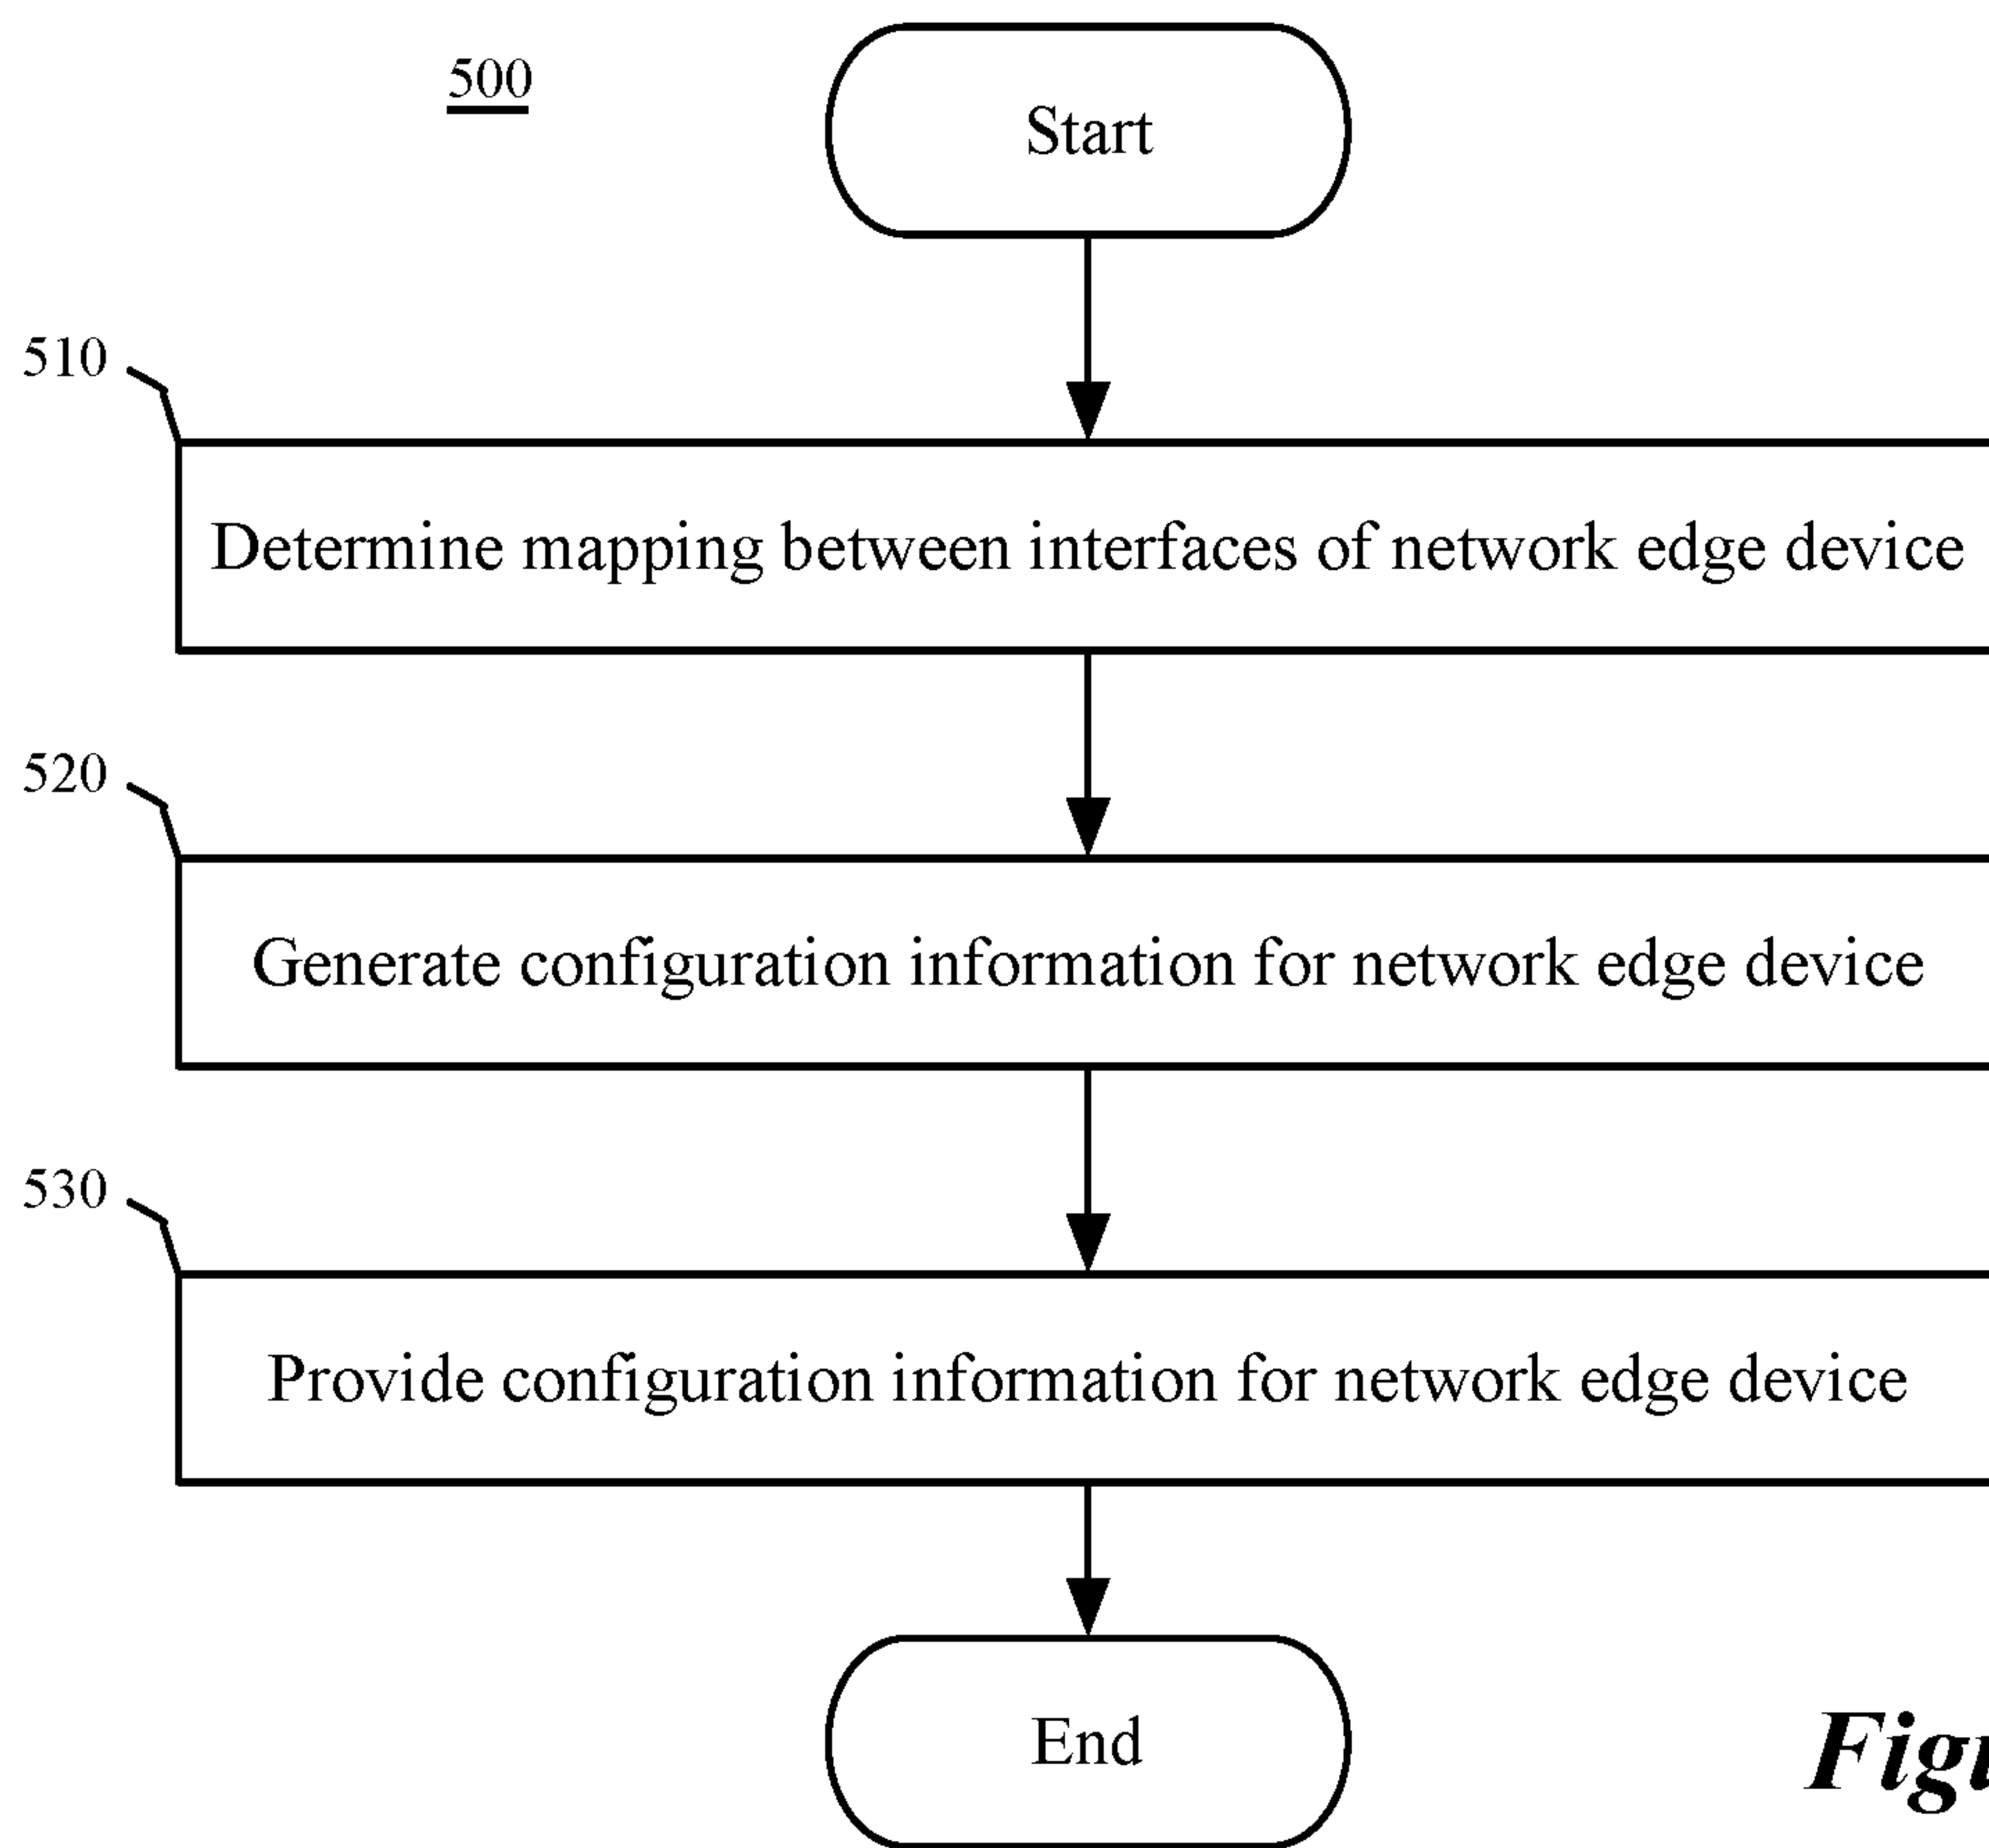


Figure 5

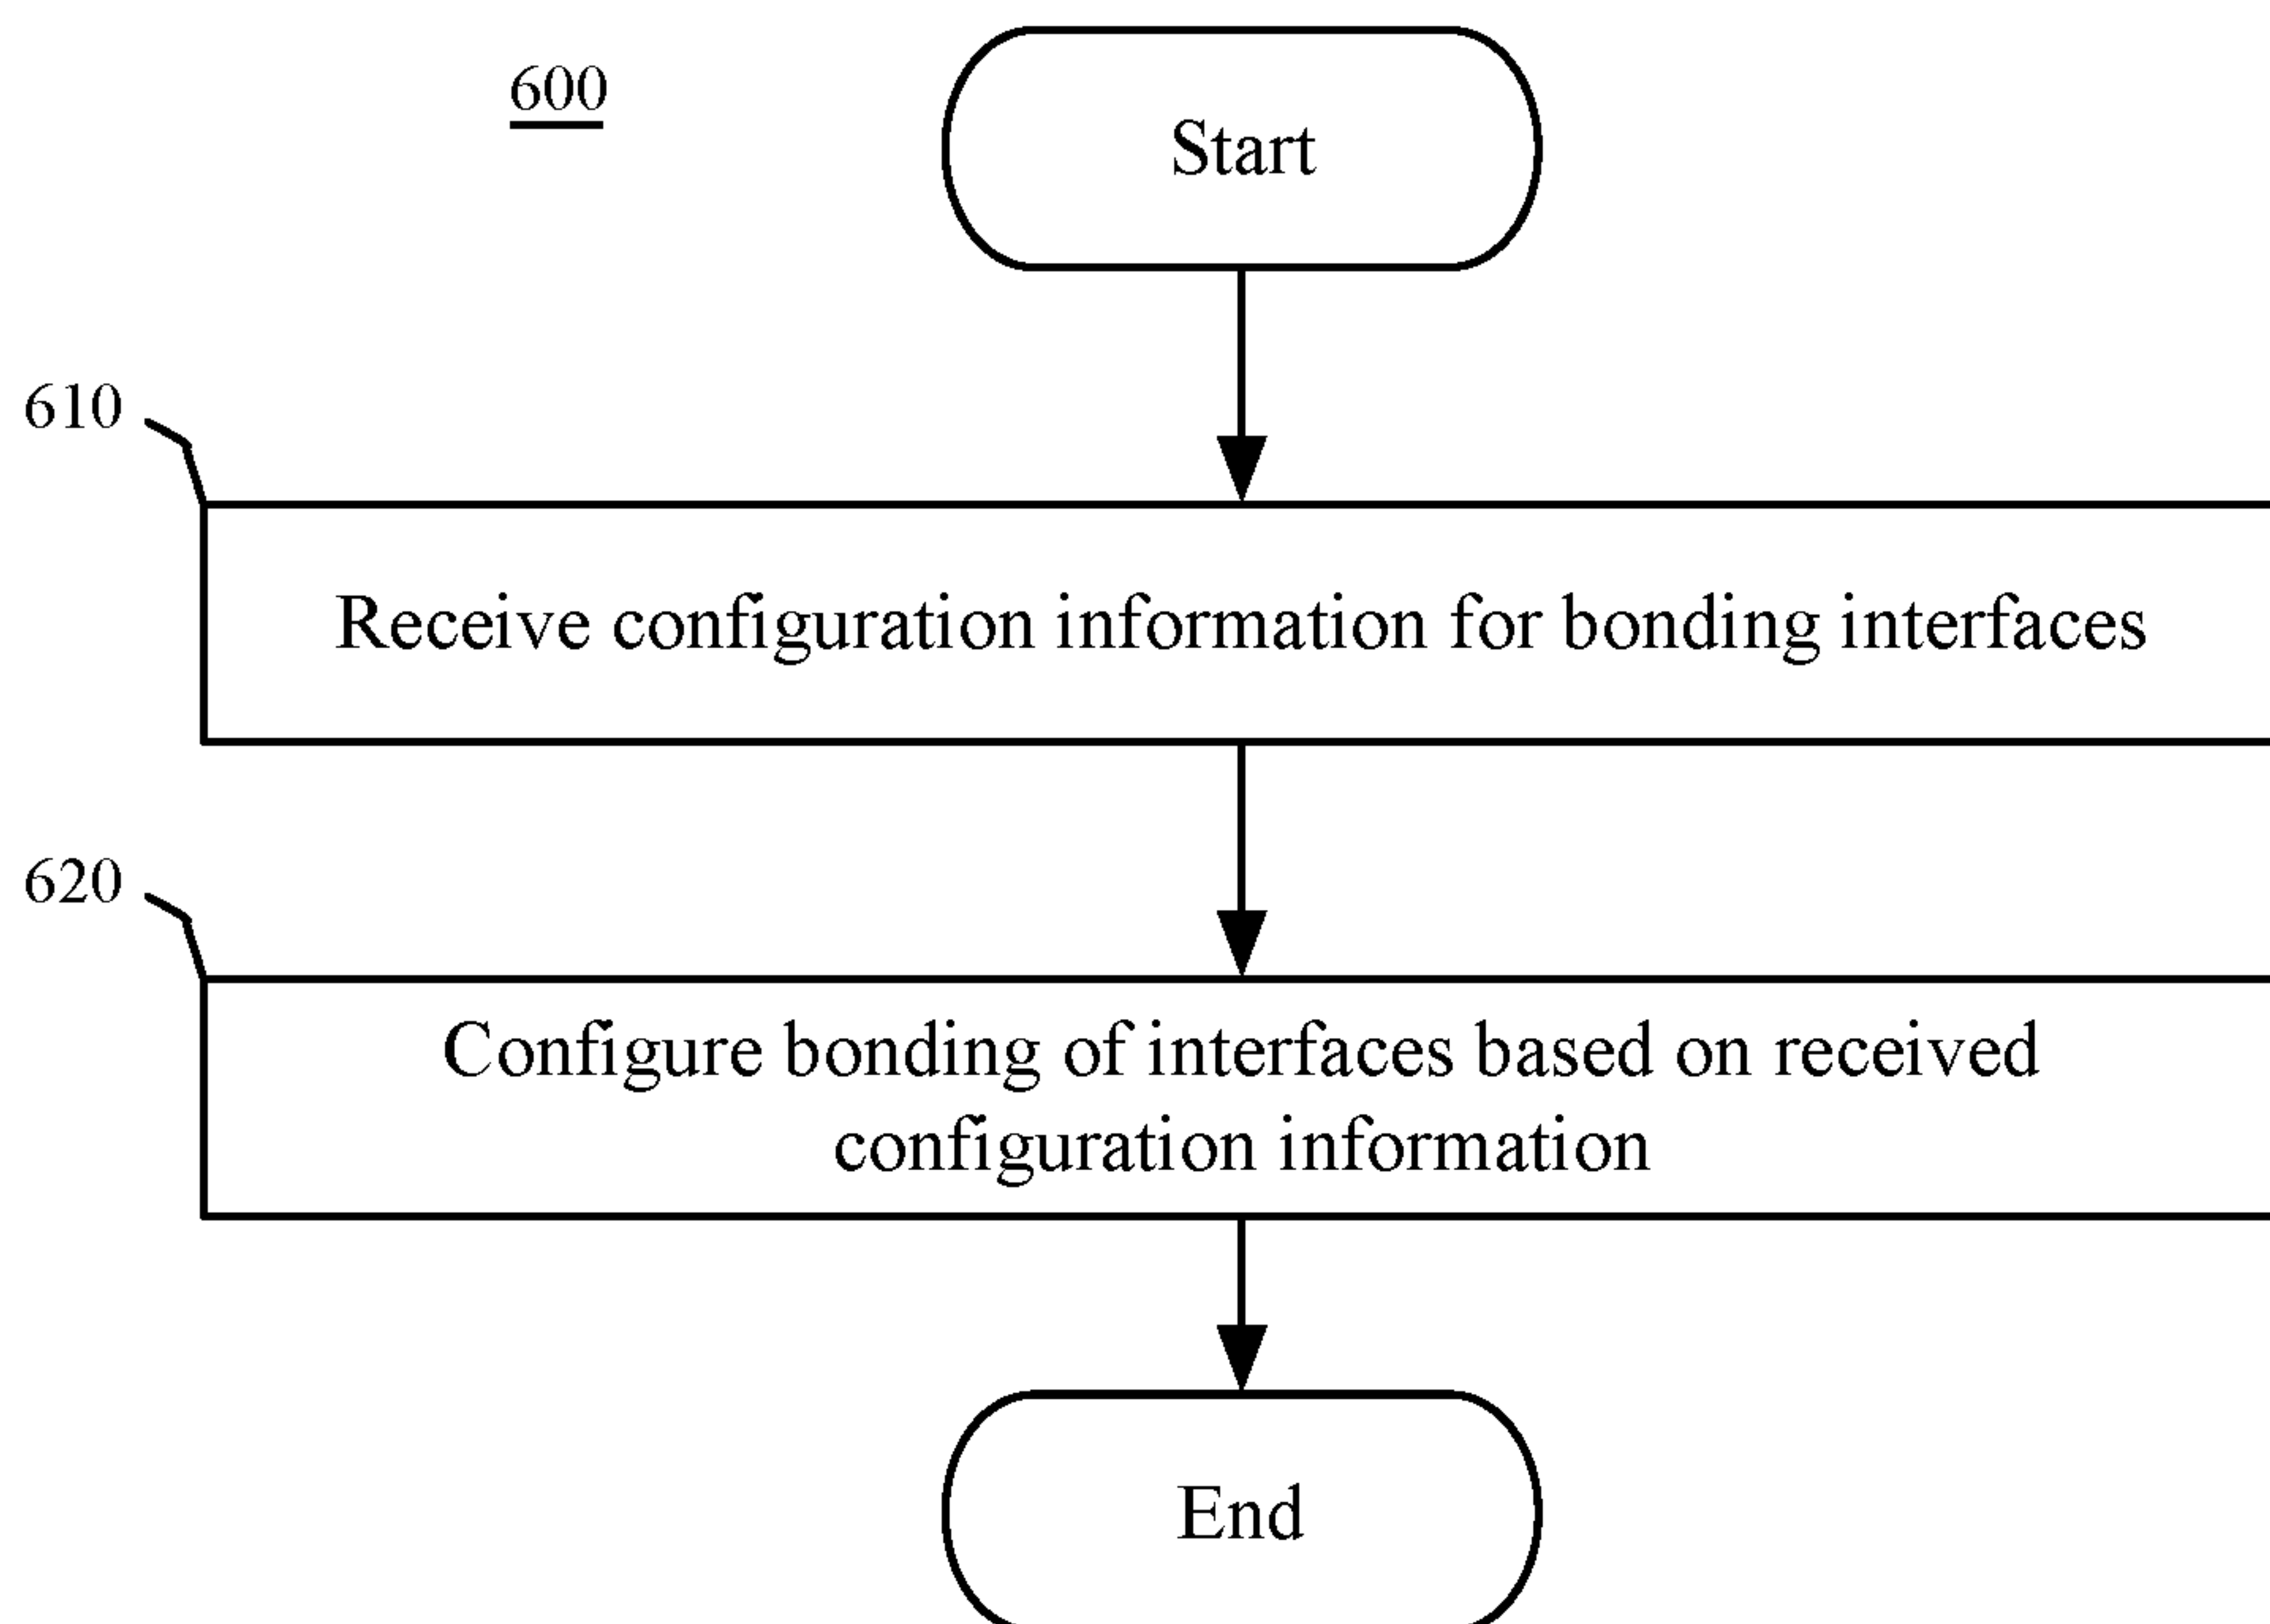


Figure 6

Figure 7

710

	Ingress Interface	Egress Interface
721	fp-eth0	fp-eth2
	fp-eth1	fp-eth3
	fp-eth2	fp-eth0
	fp-eth3	fp-eth1

720

720A

	Ingress Interface	Egress Interface
721	fp-eth0	fp-eth2
	fp-eth1	fp-eth3

720B

	Ingress Interface	Egress Interface
721	fp-eth2	fp-eth0
	fp-eth3	fp-eth1

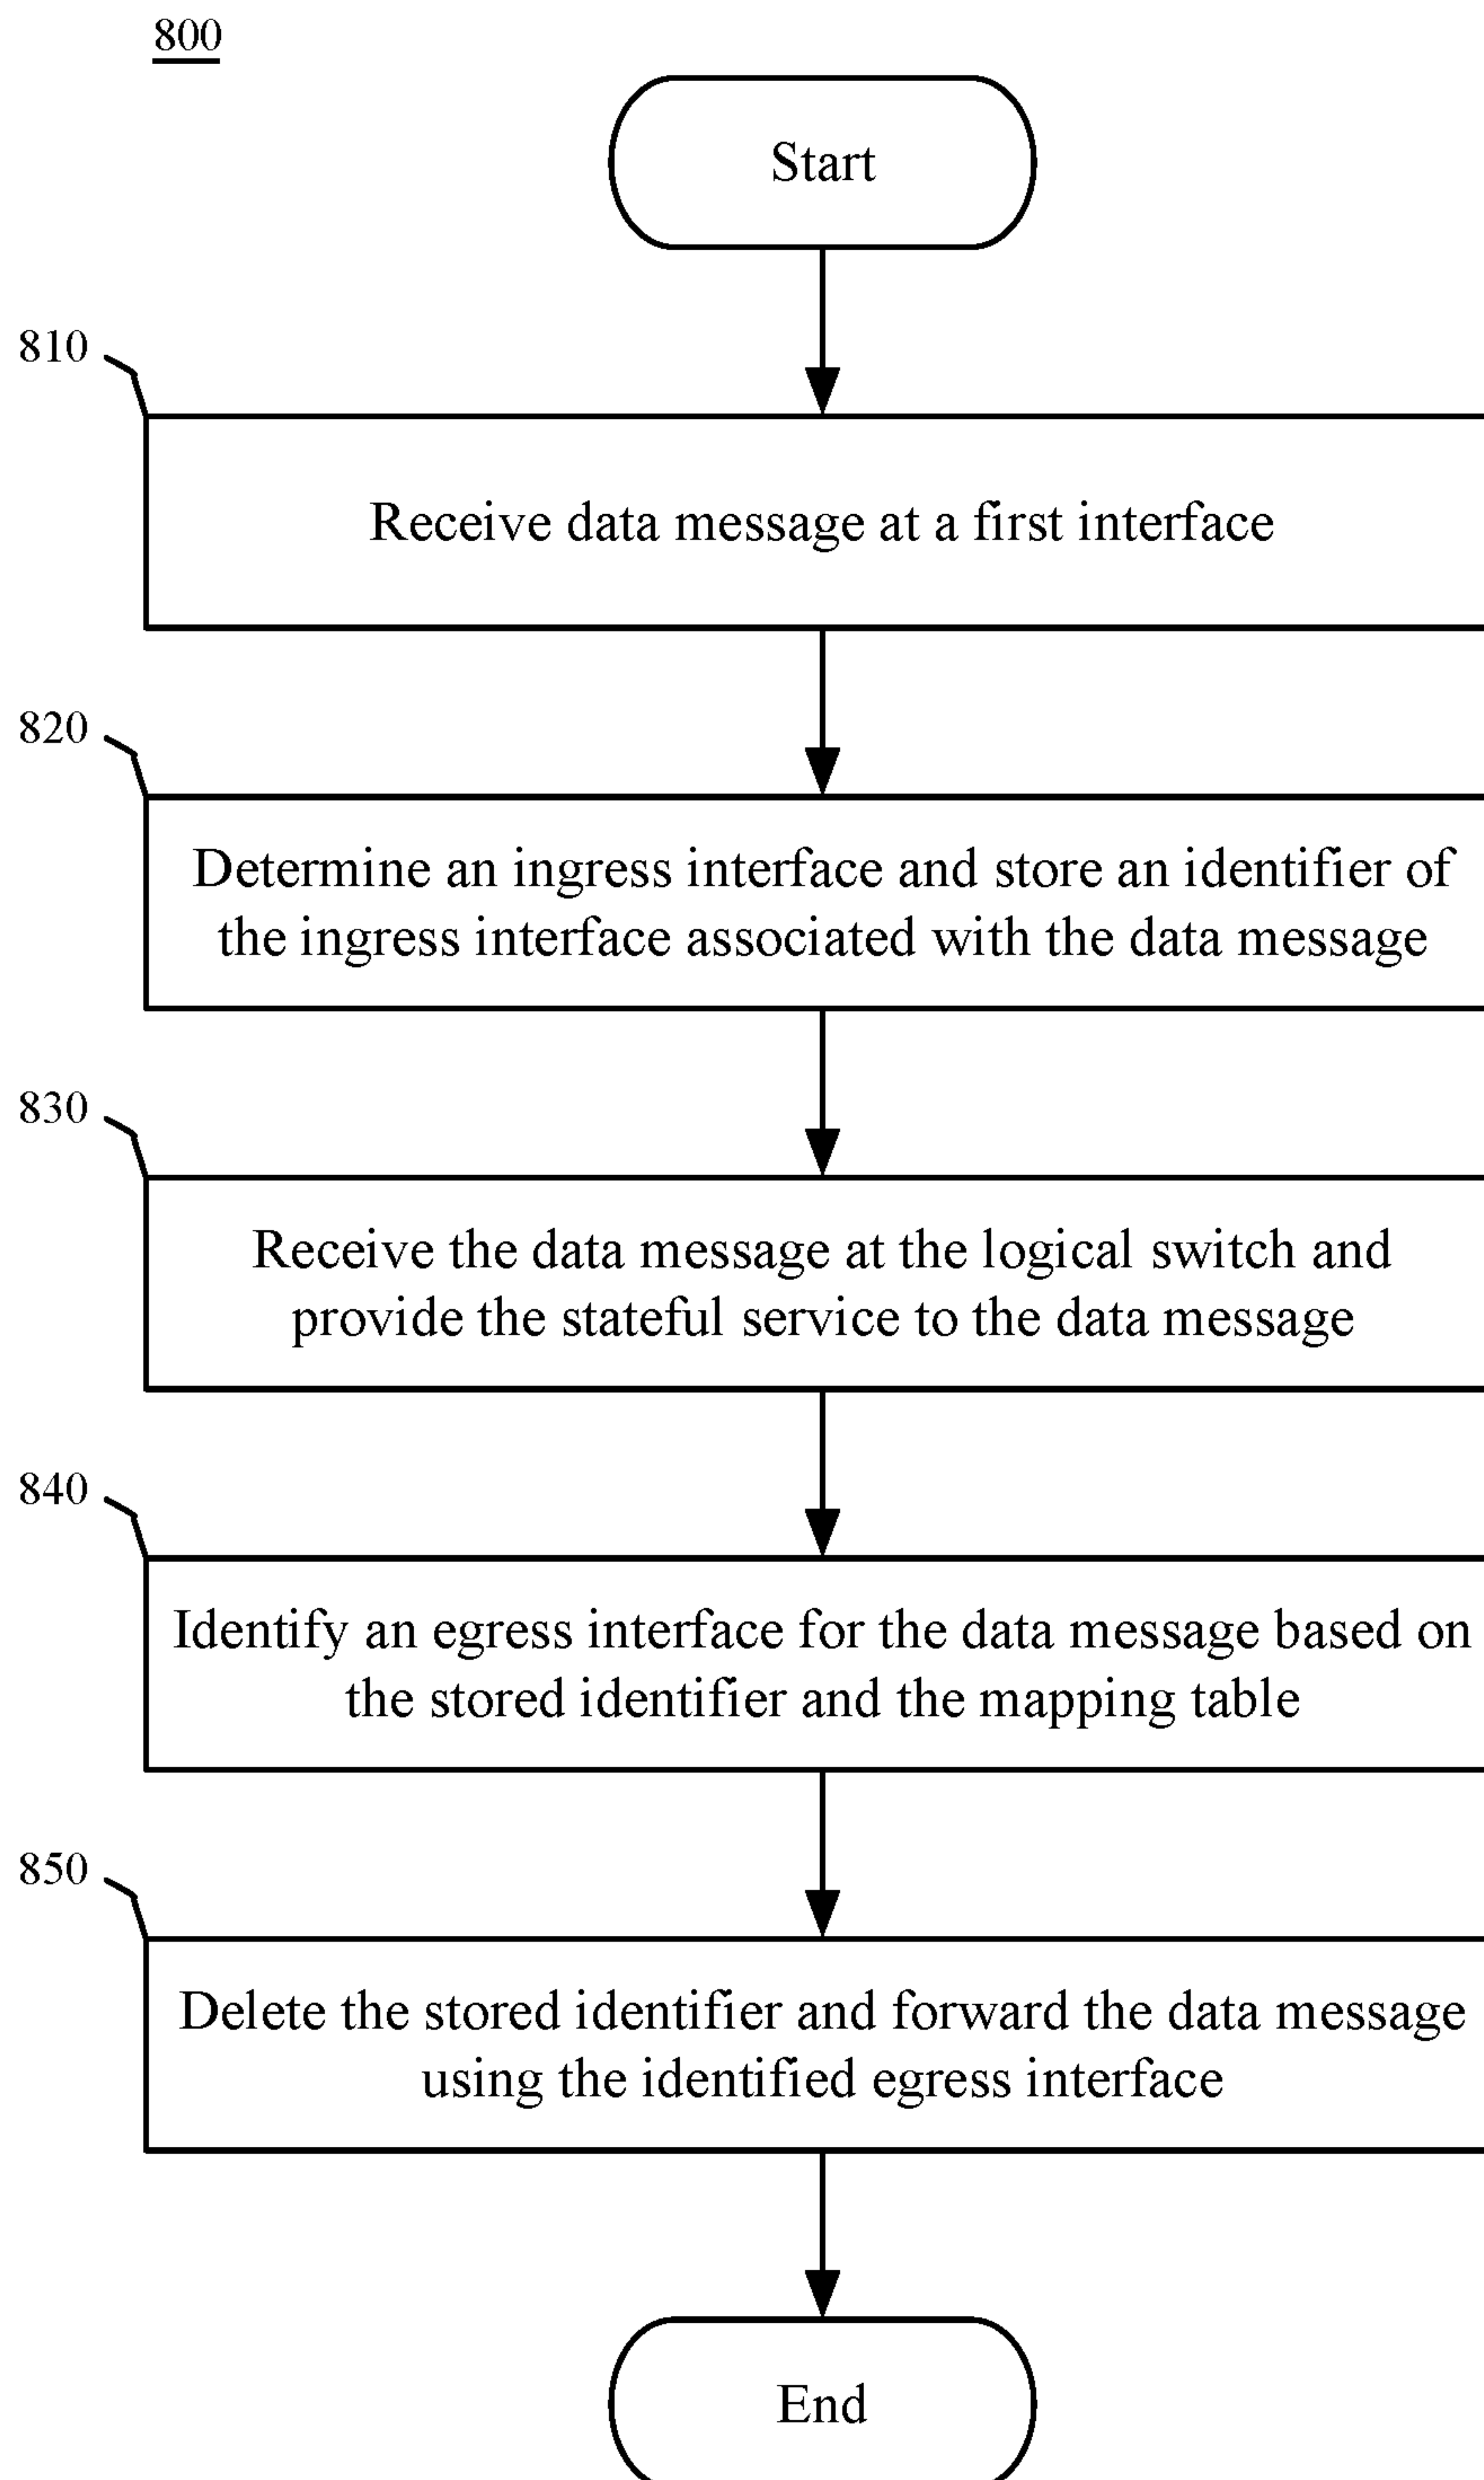
**Figure 8**

Figure 9

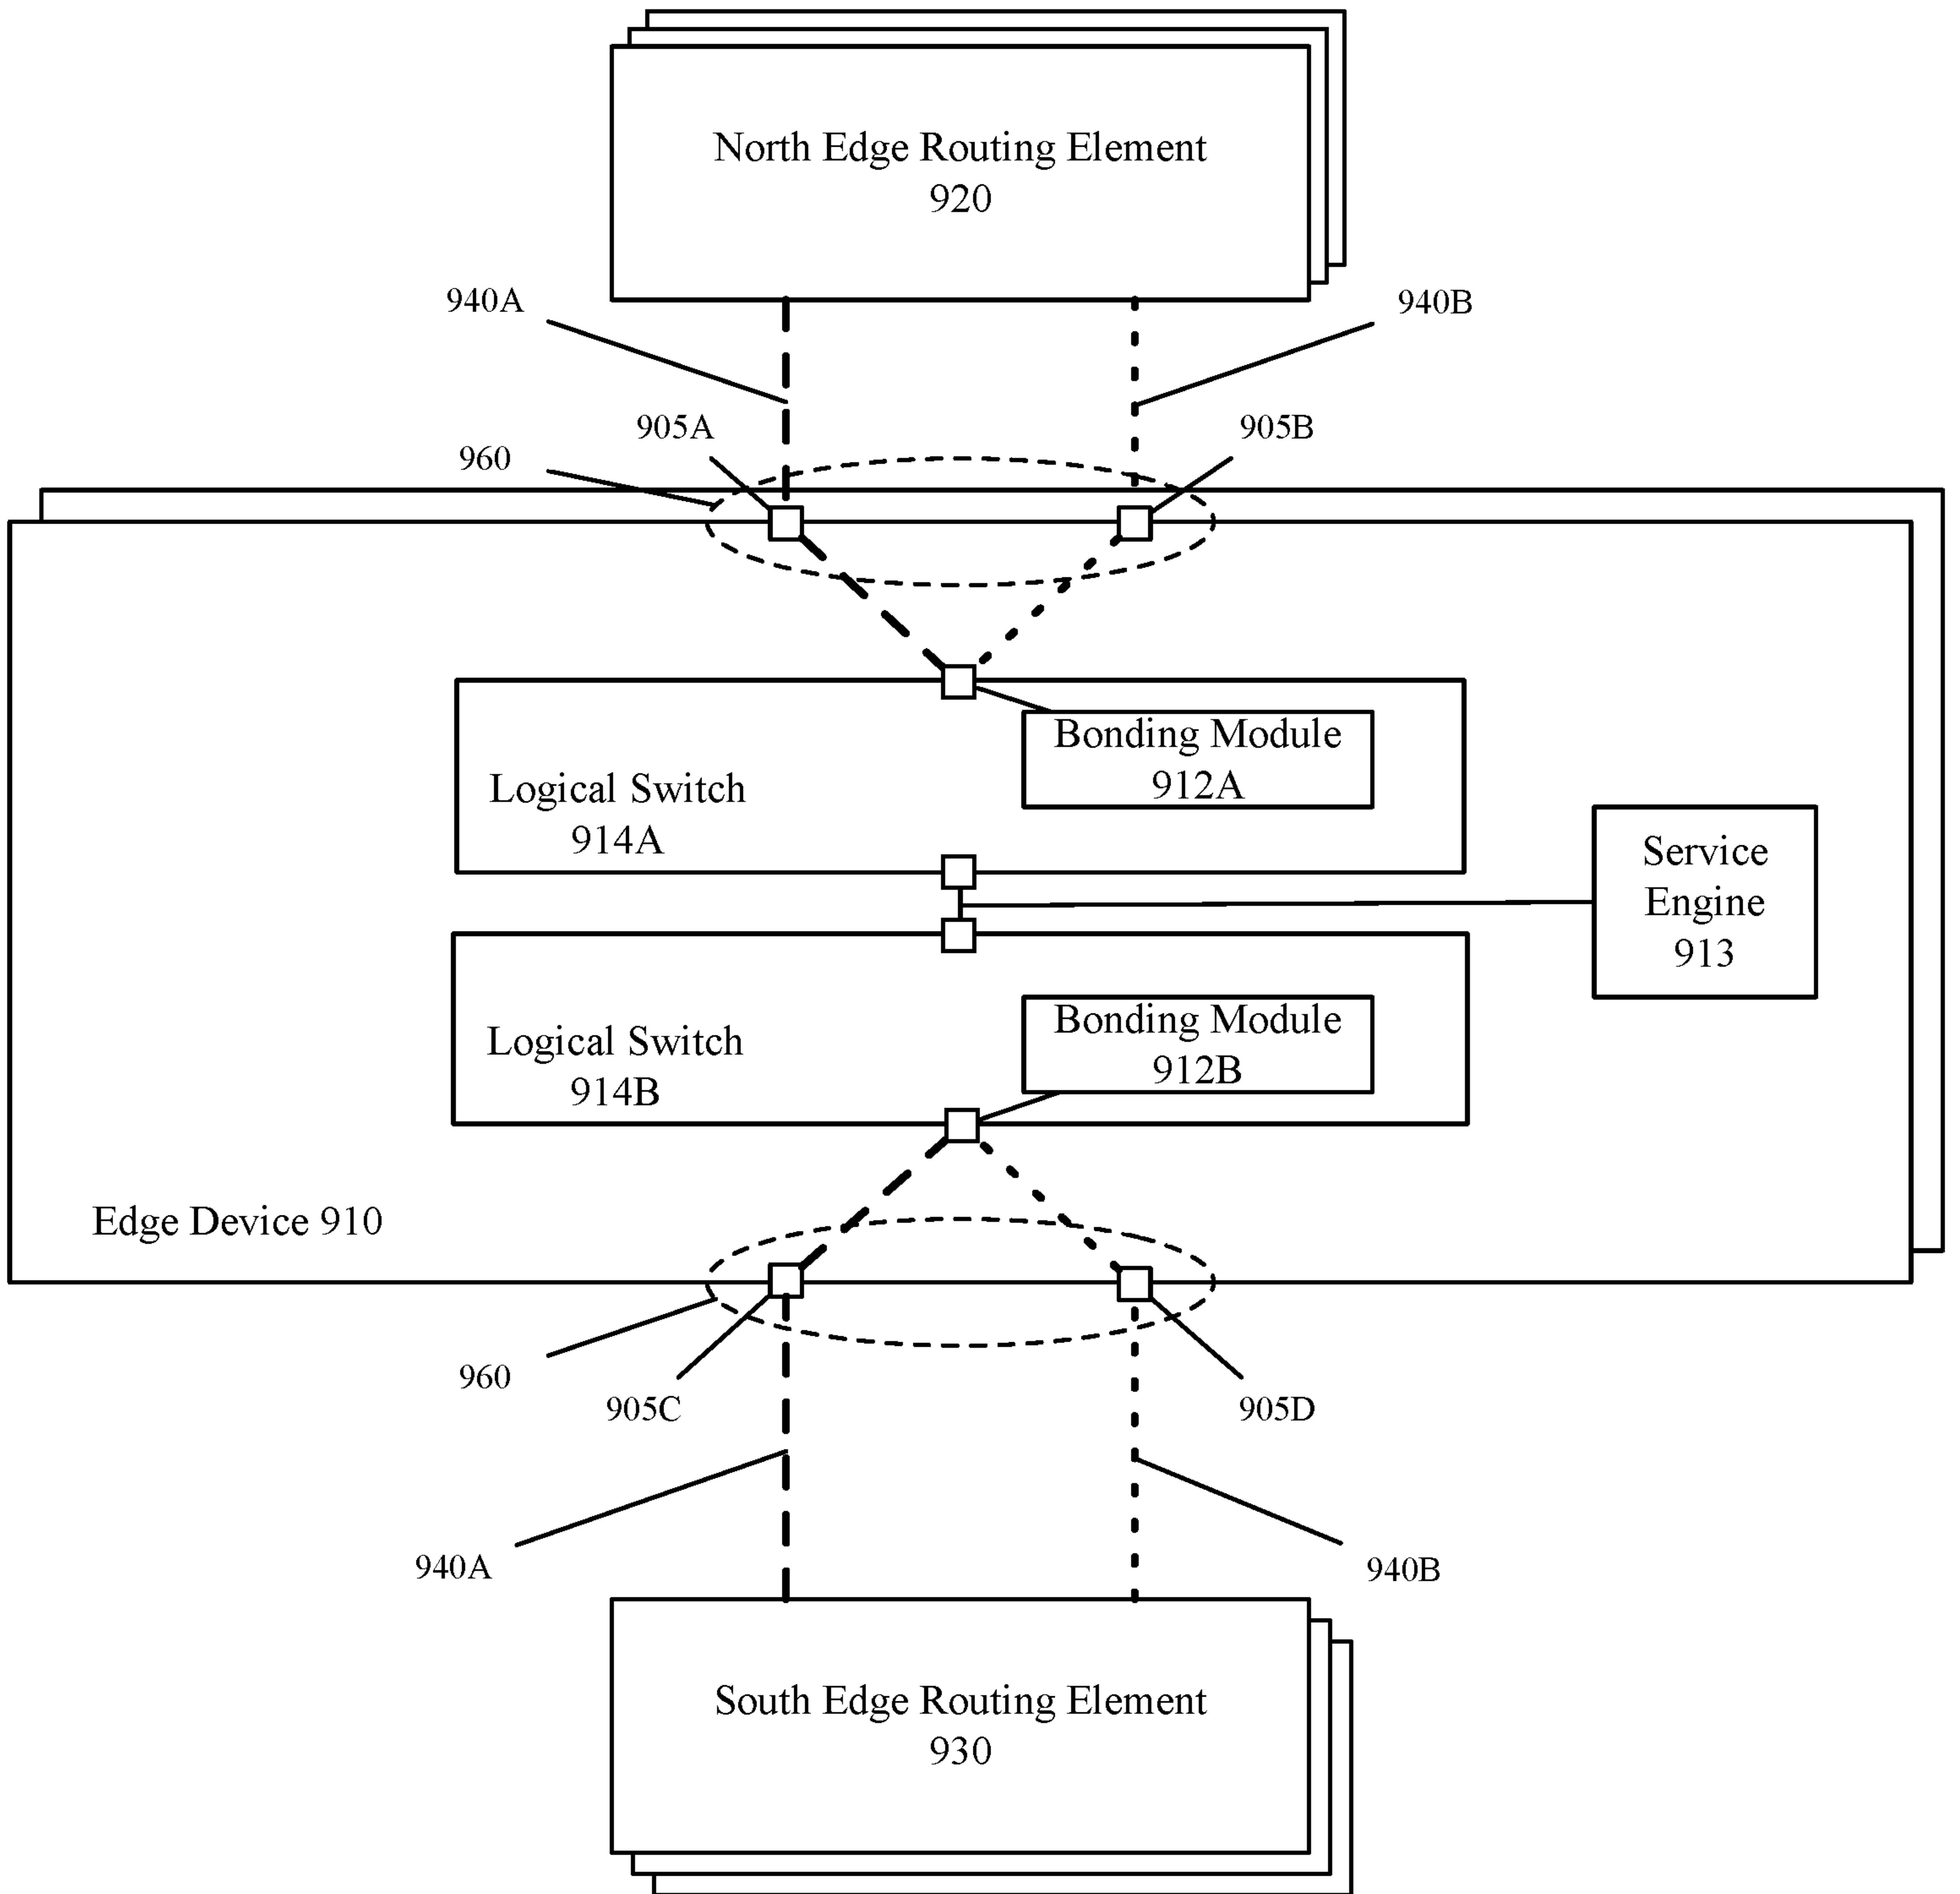


Figure 10

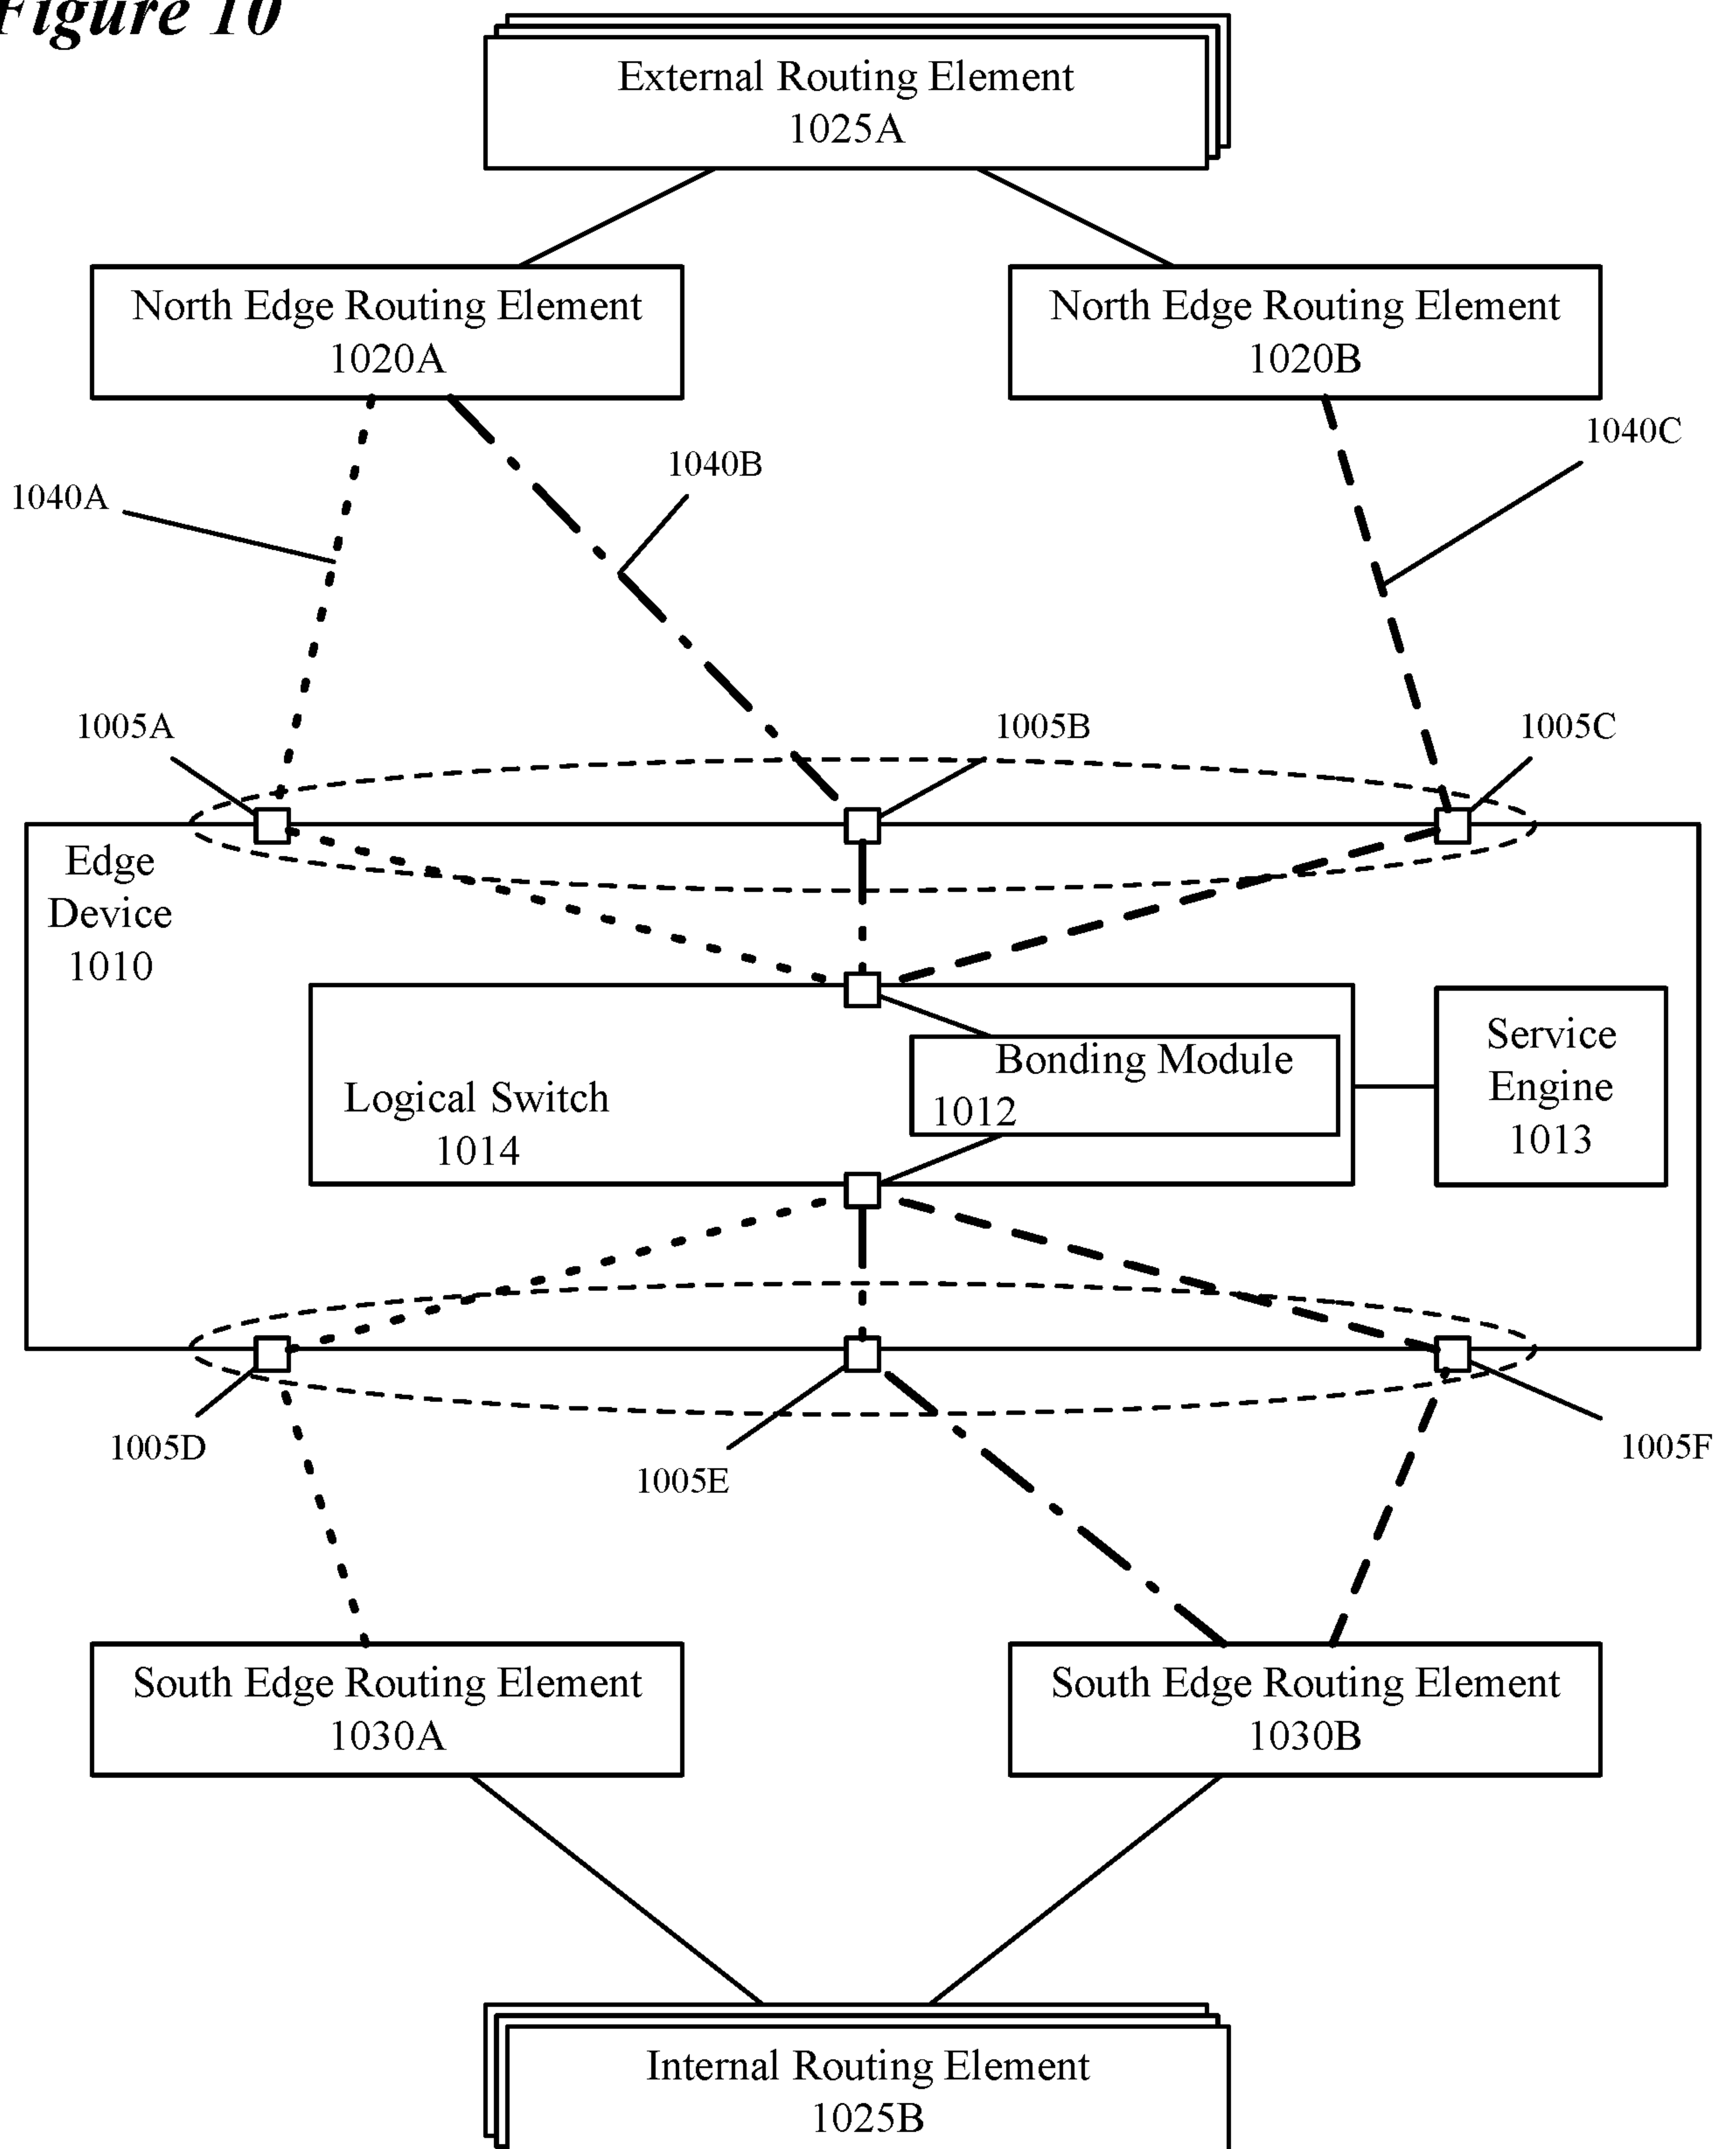
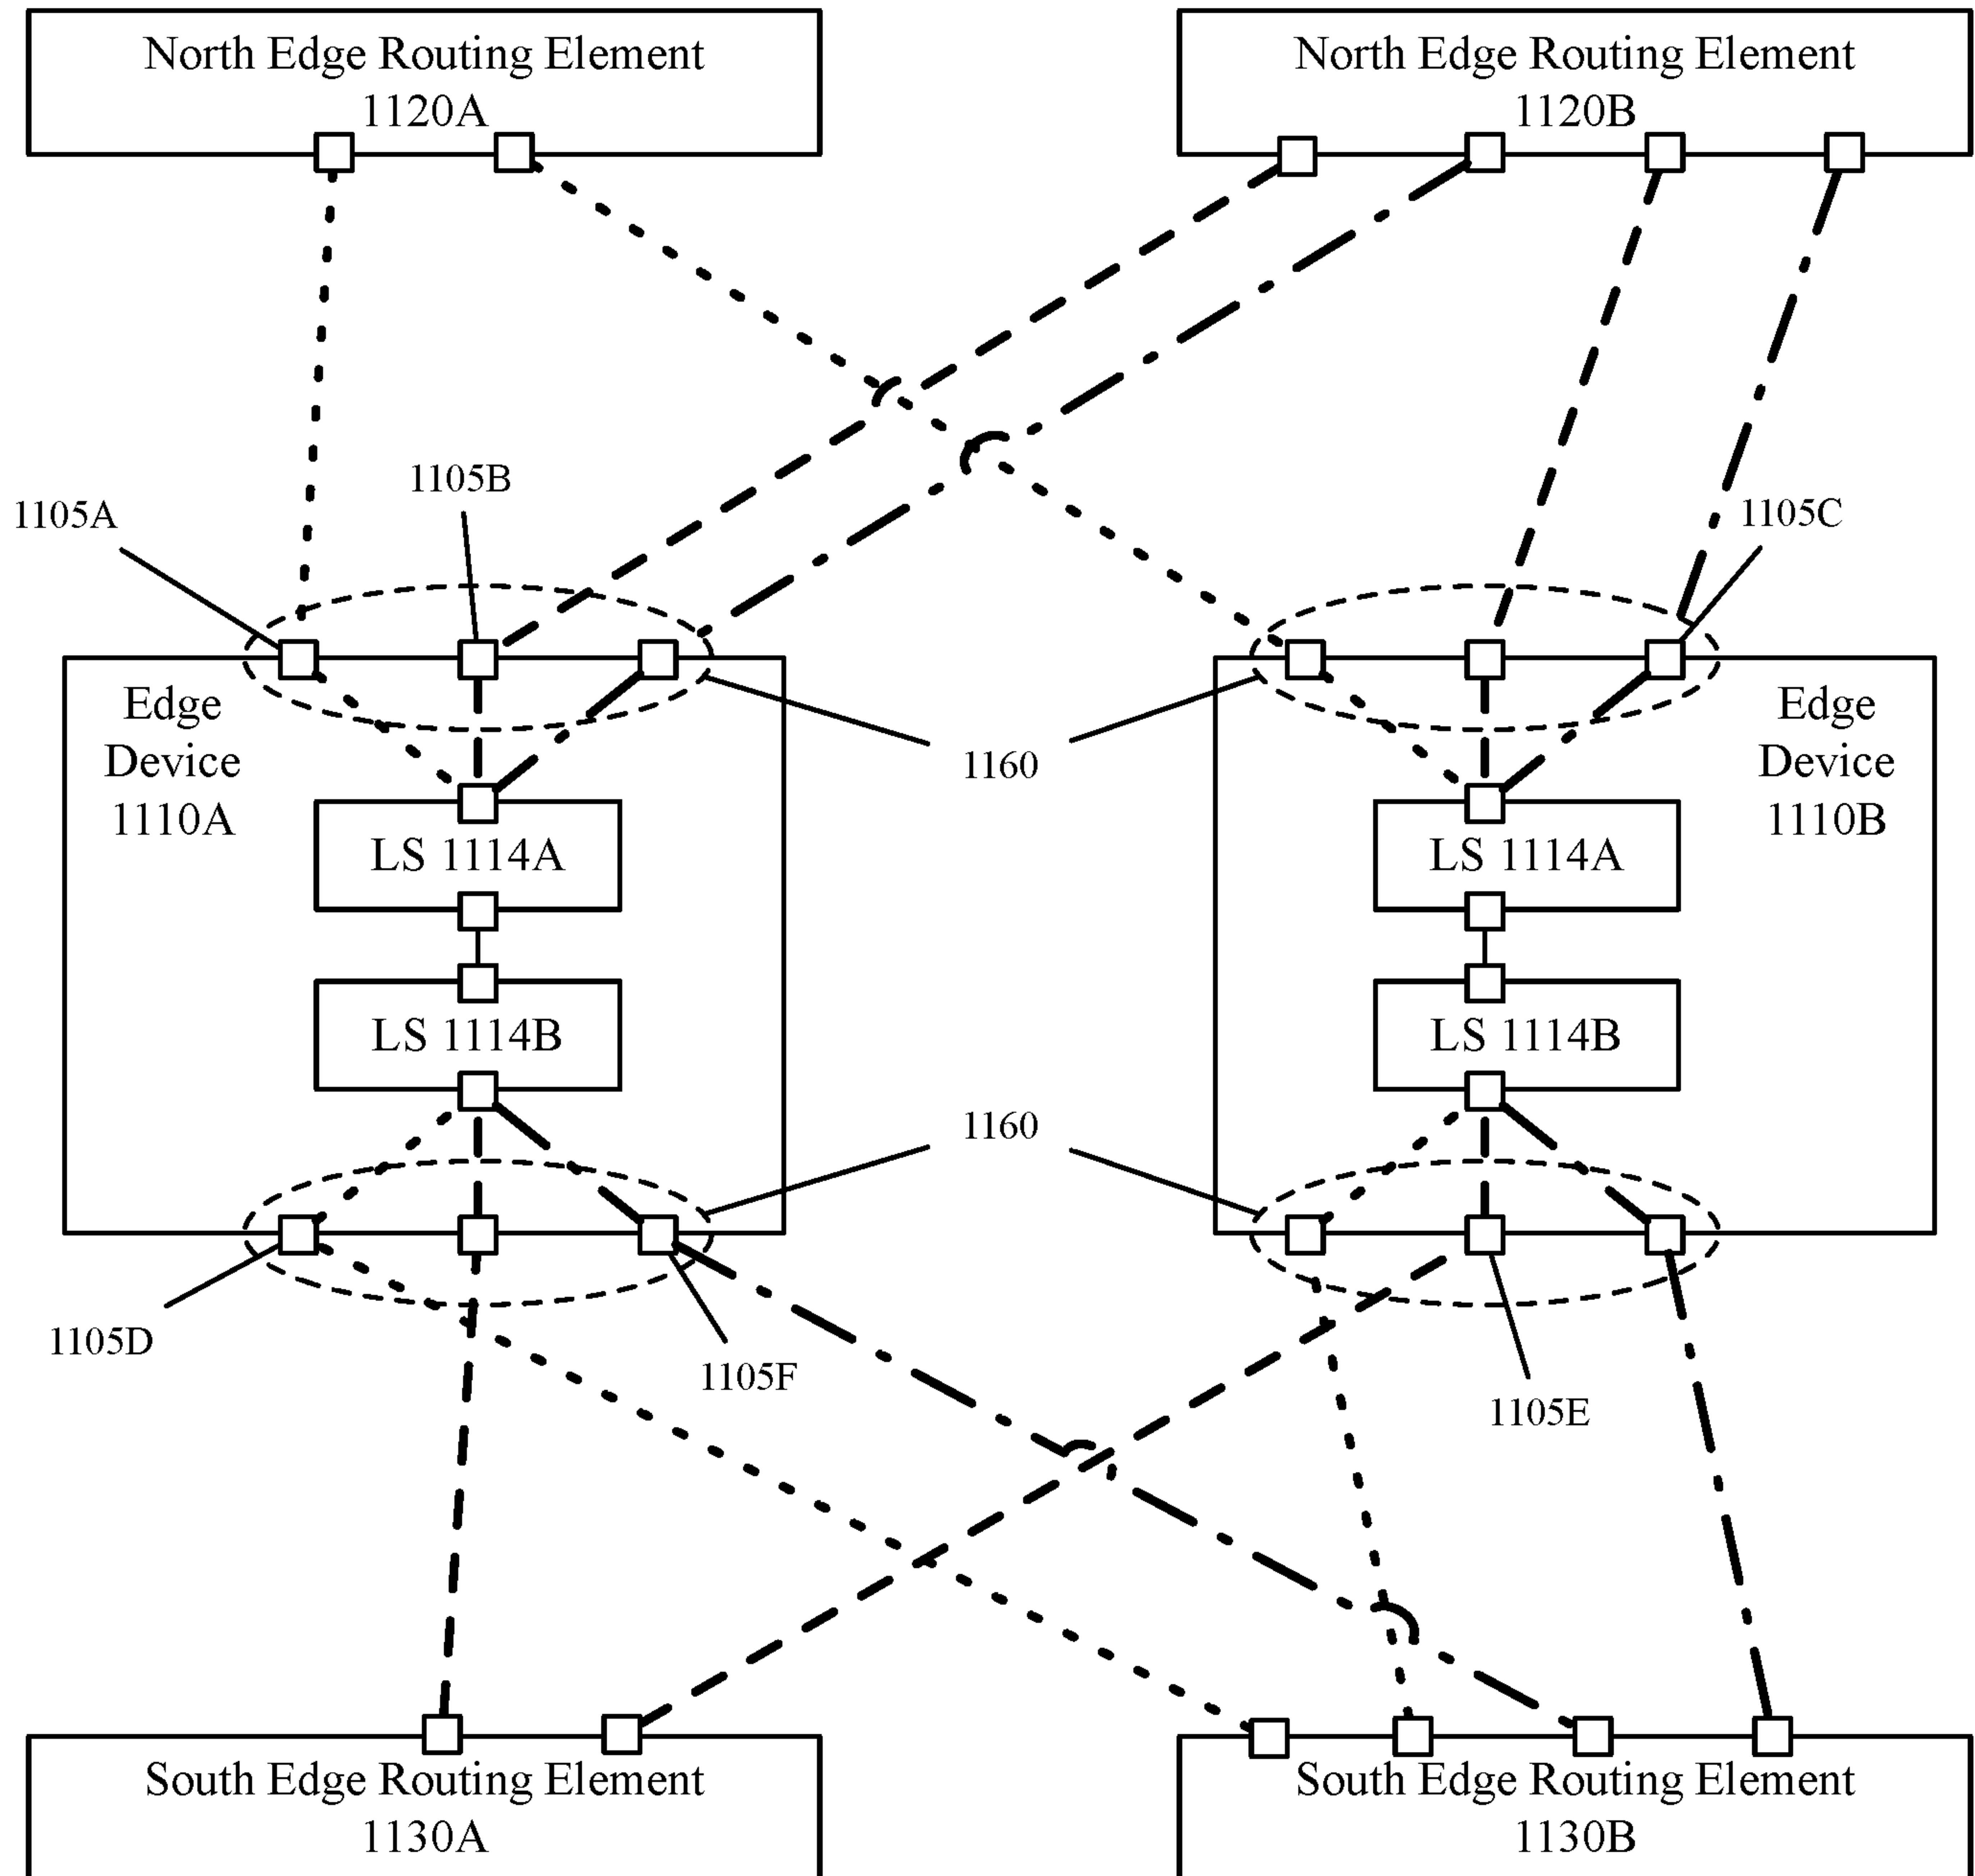


Figure 11



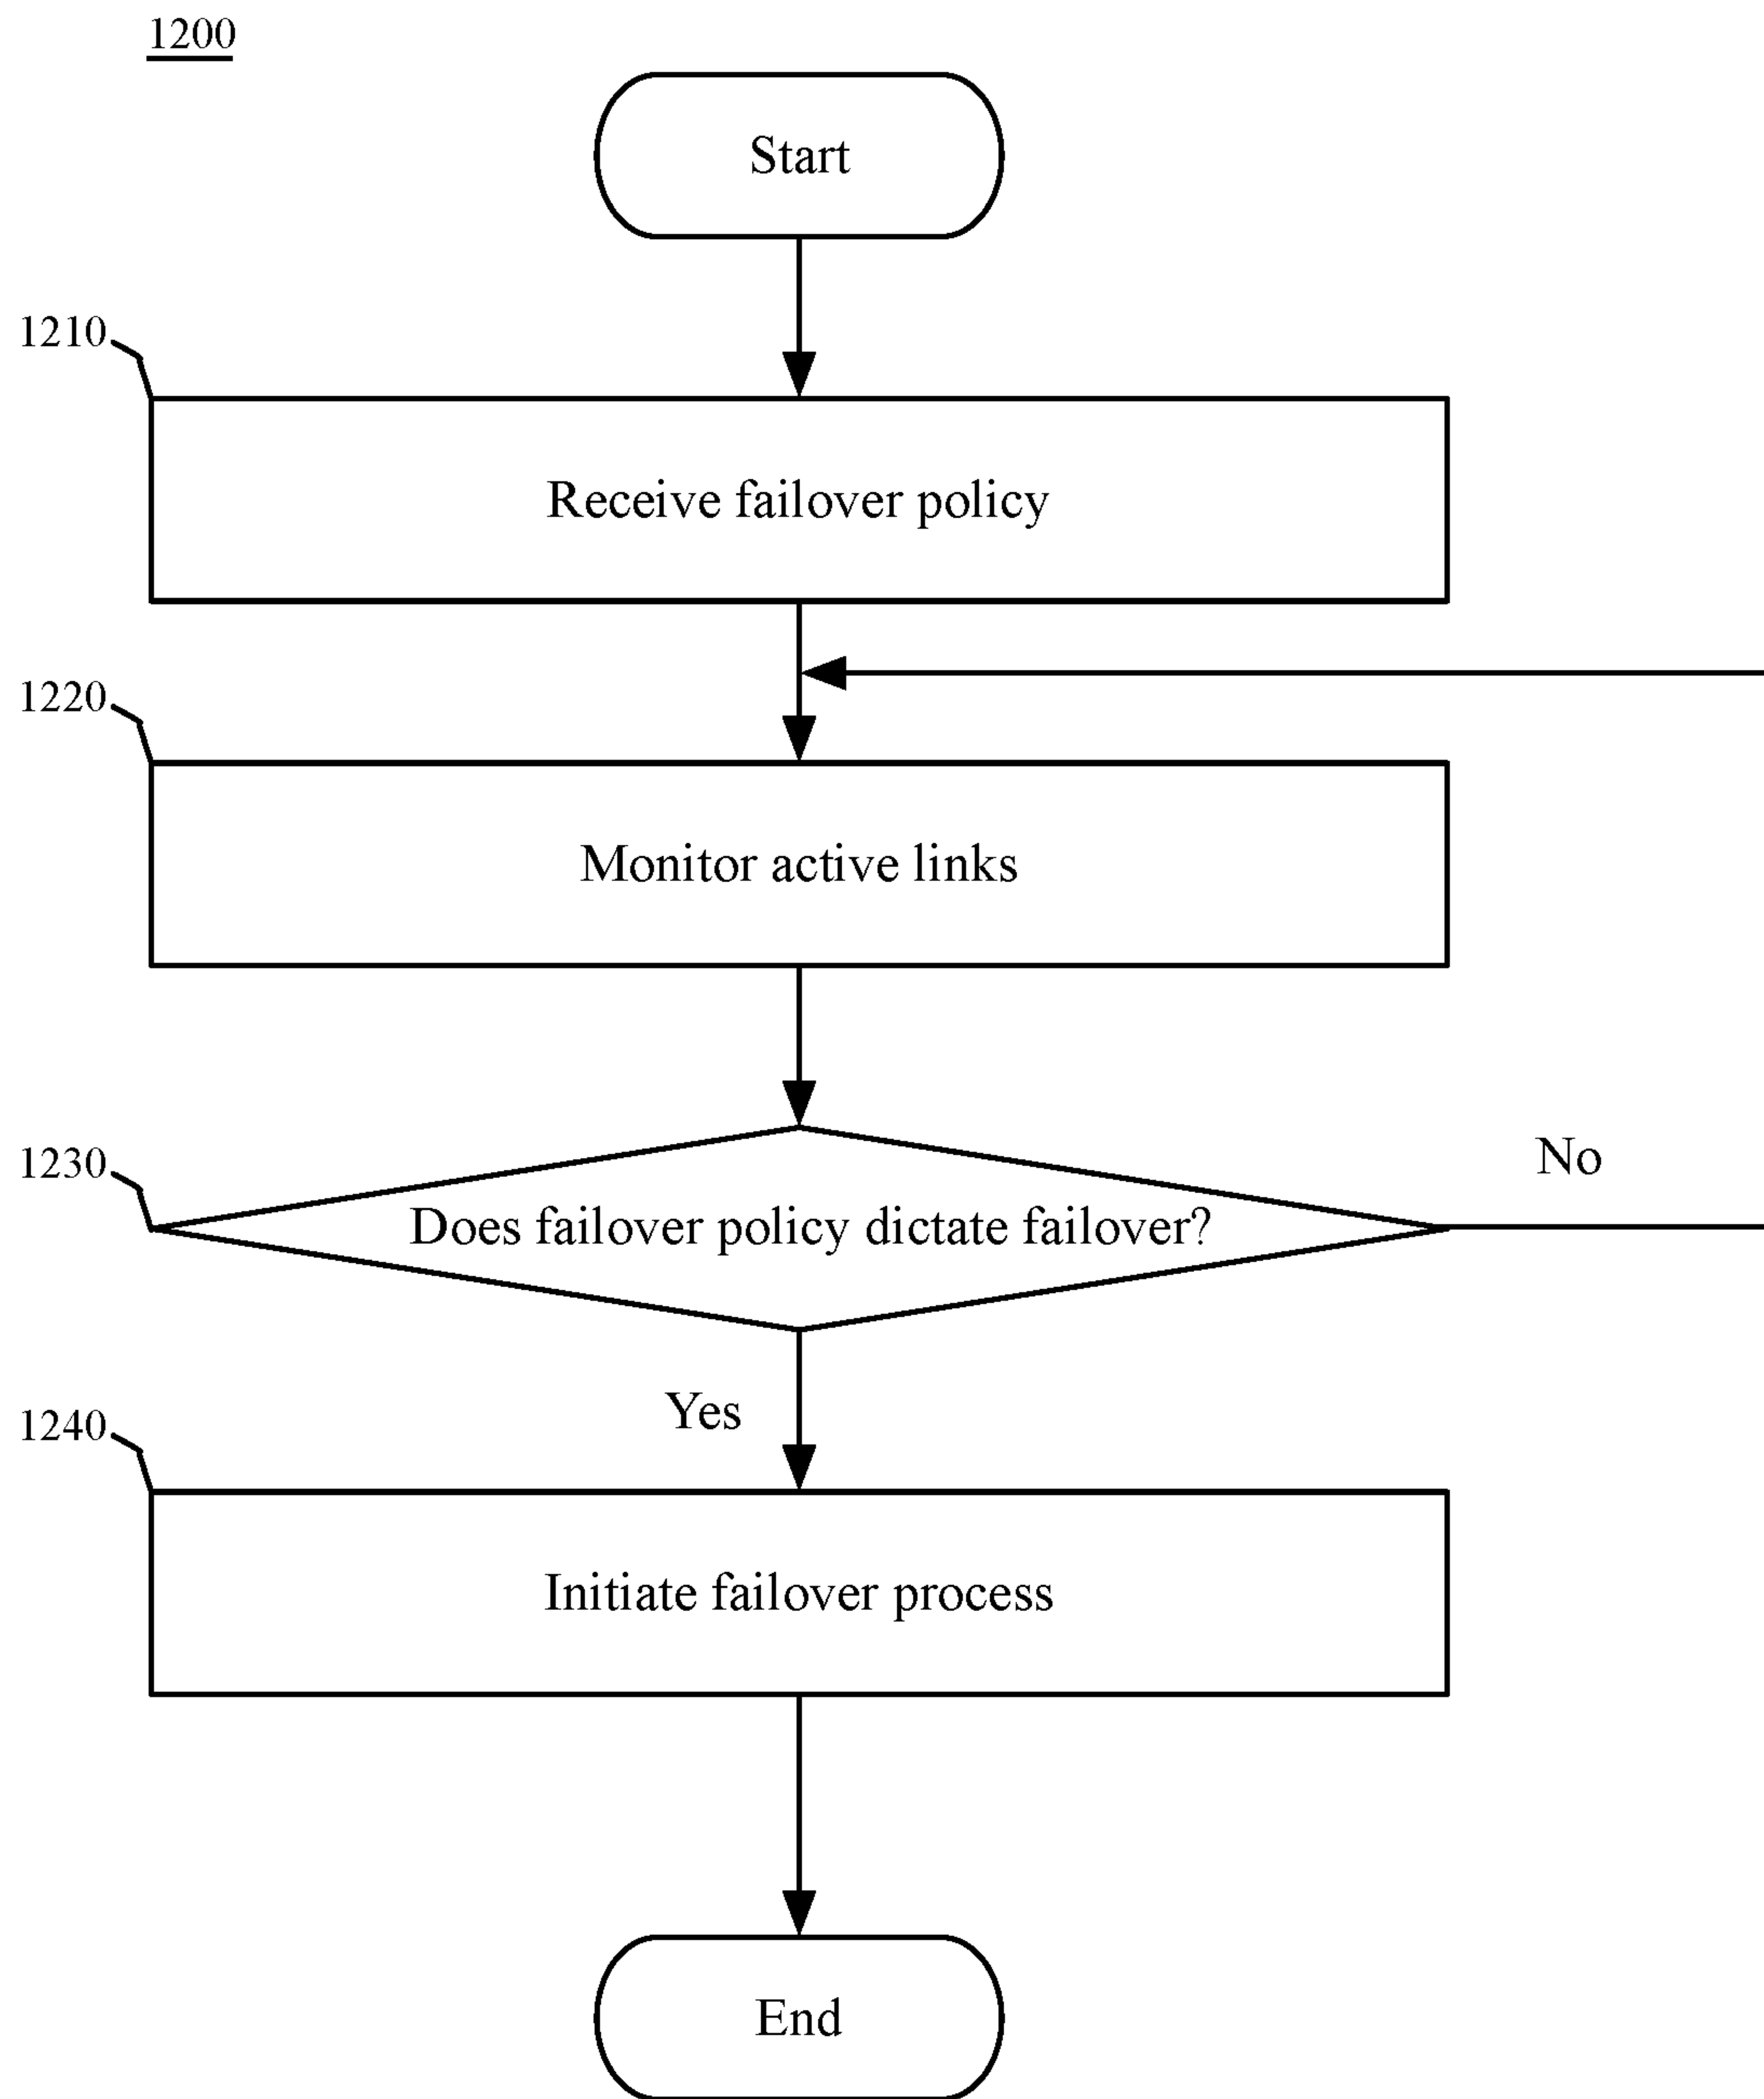


Figure 12

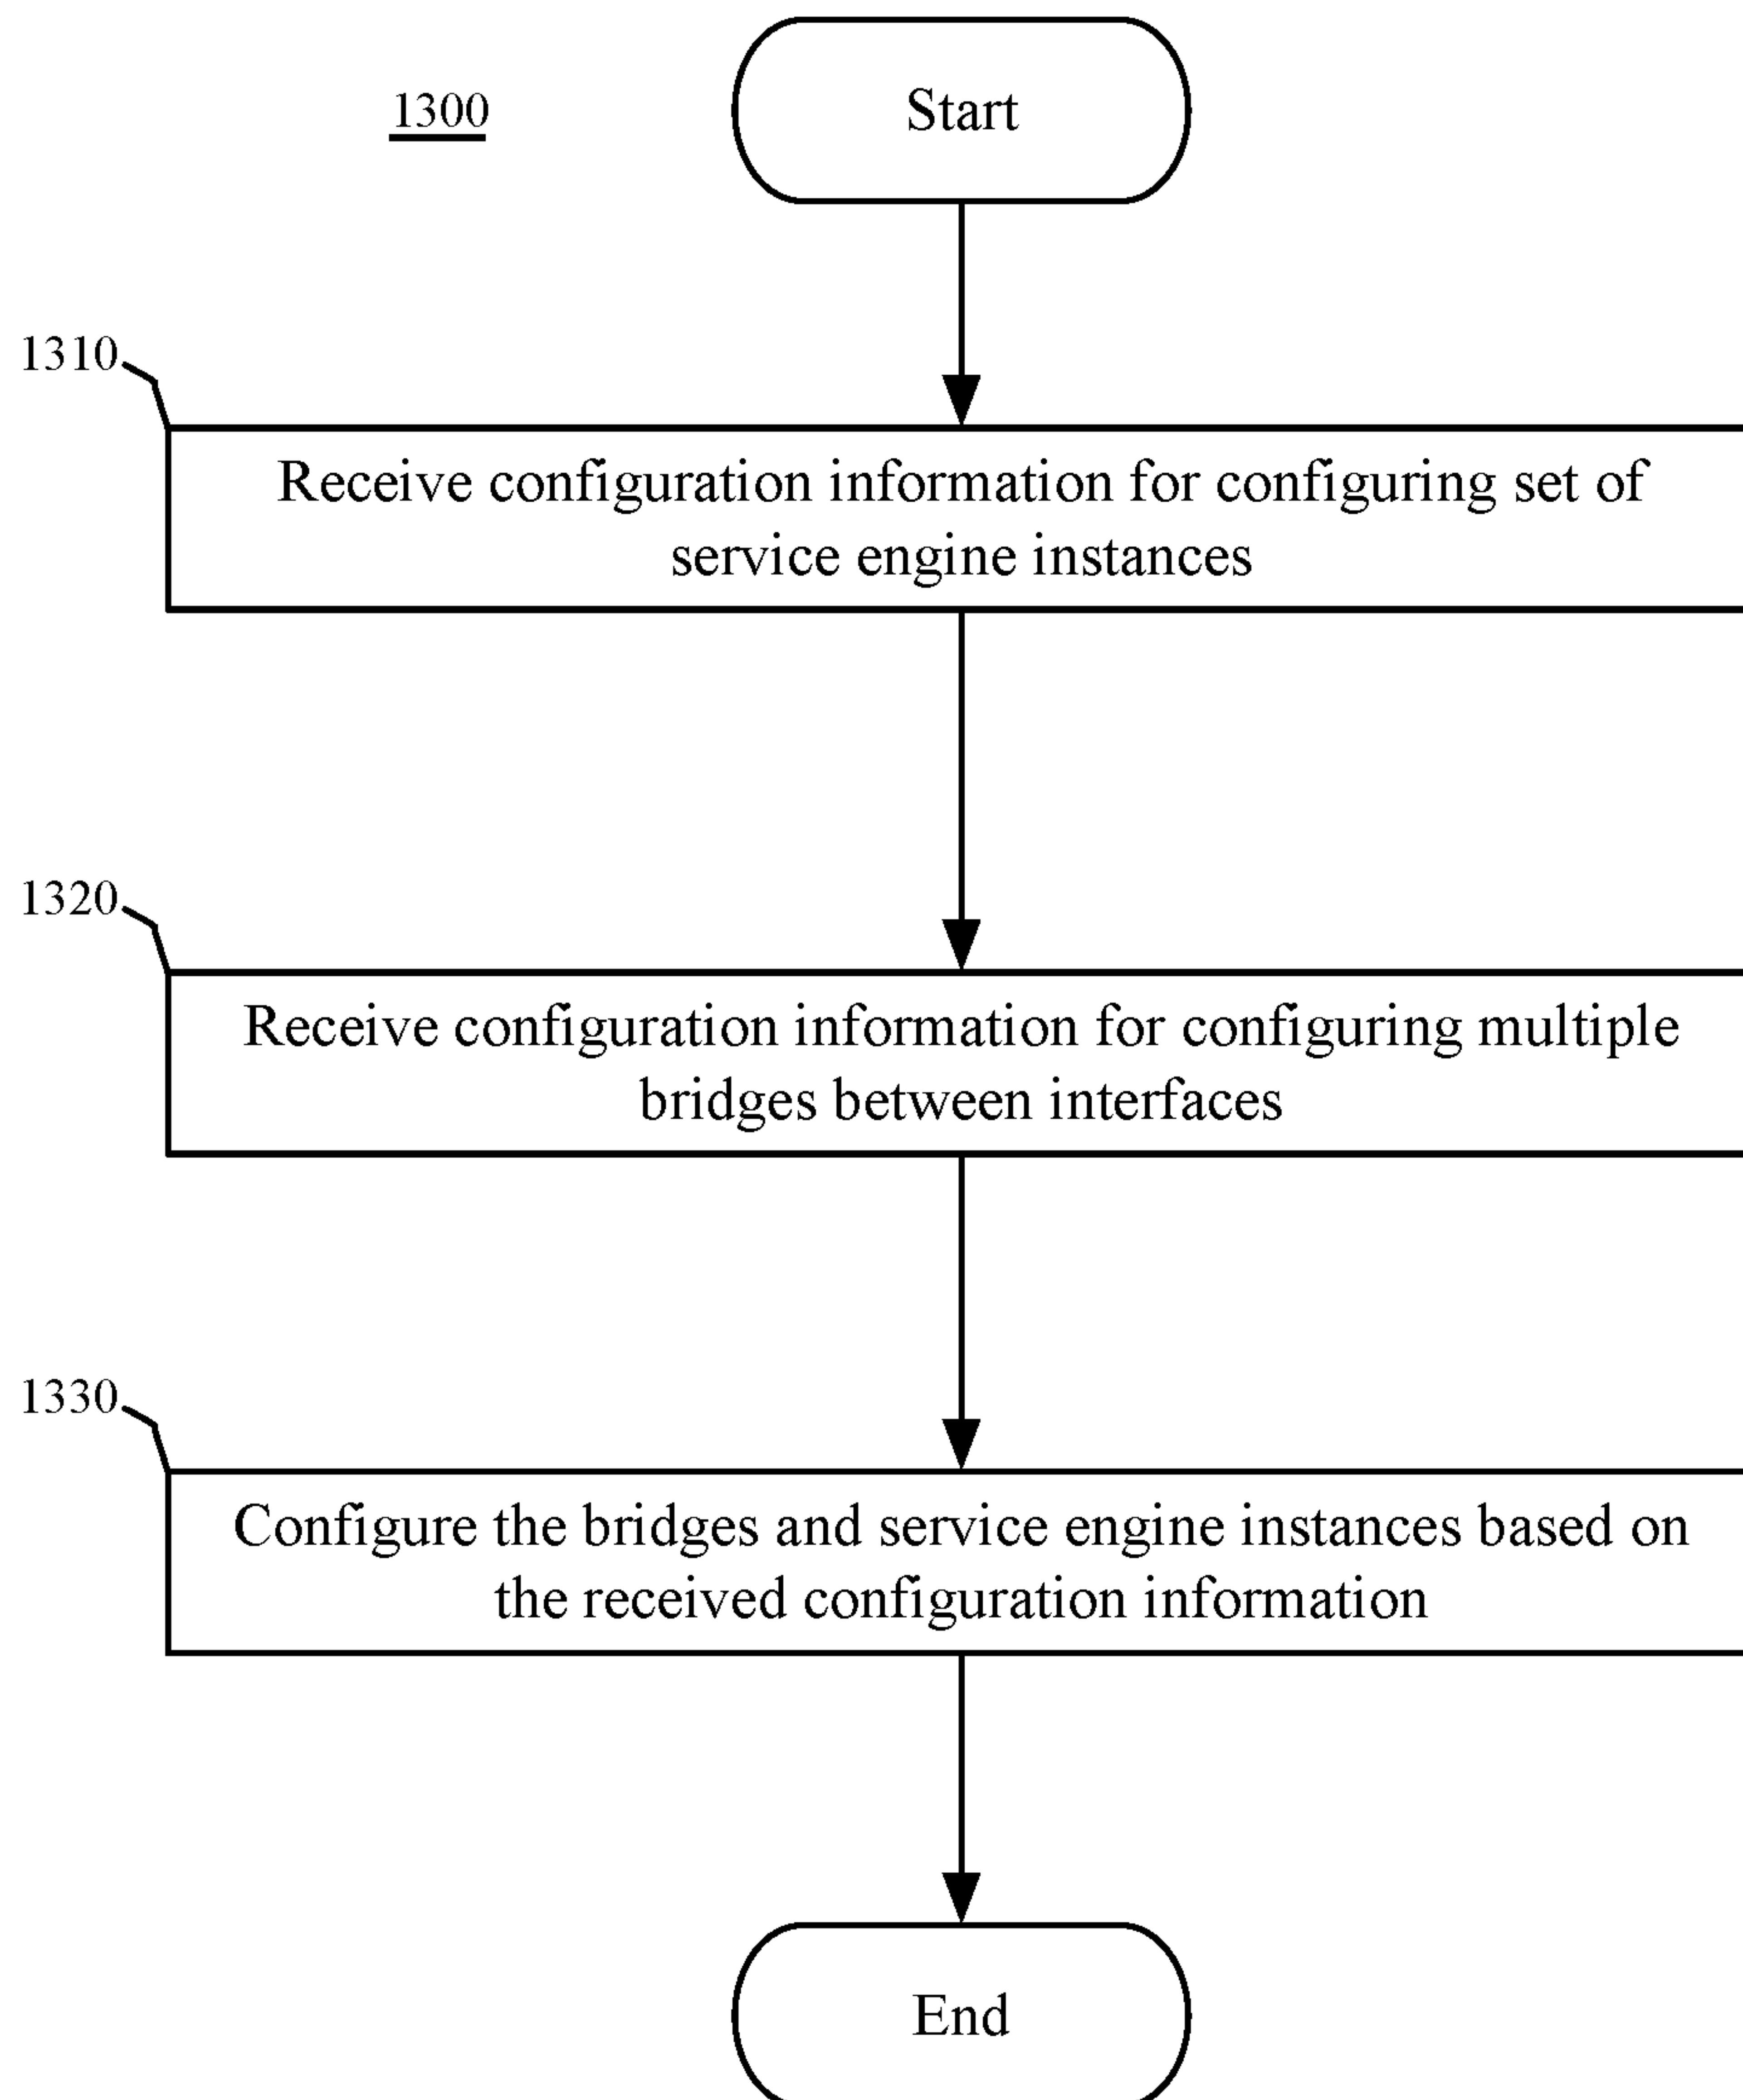


Figure 13

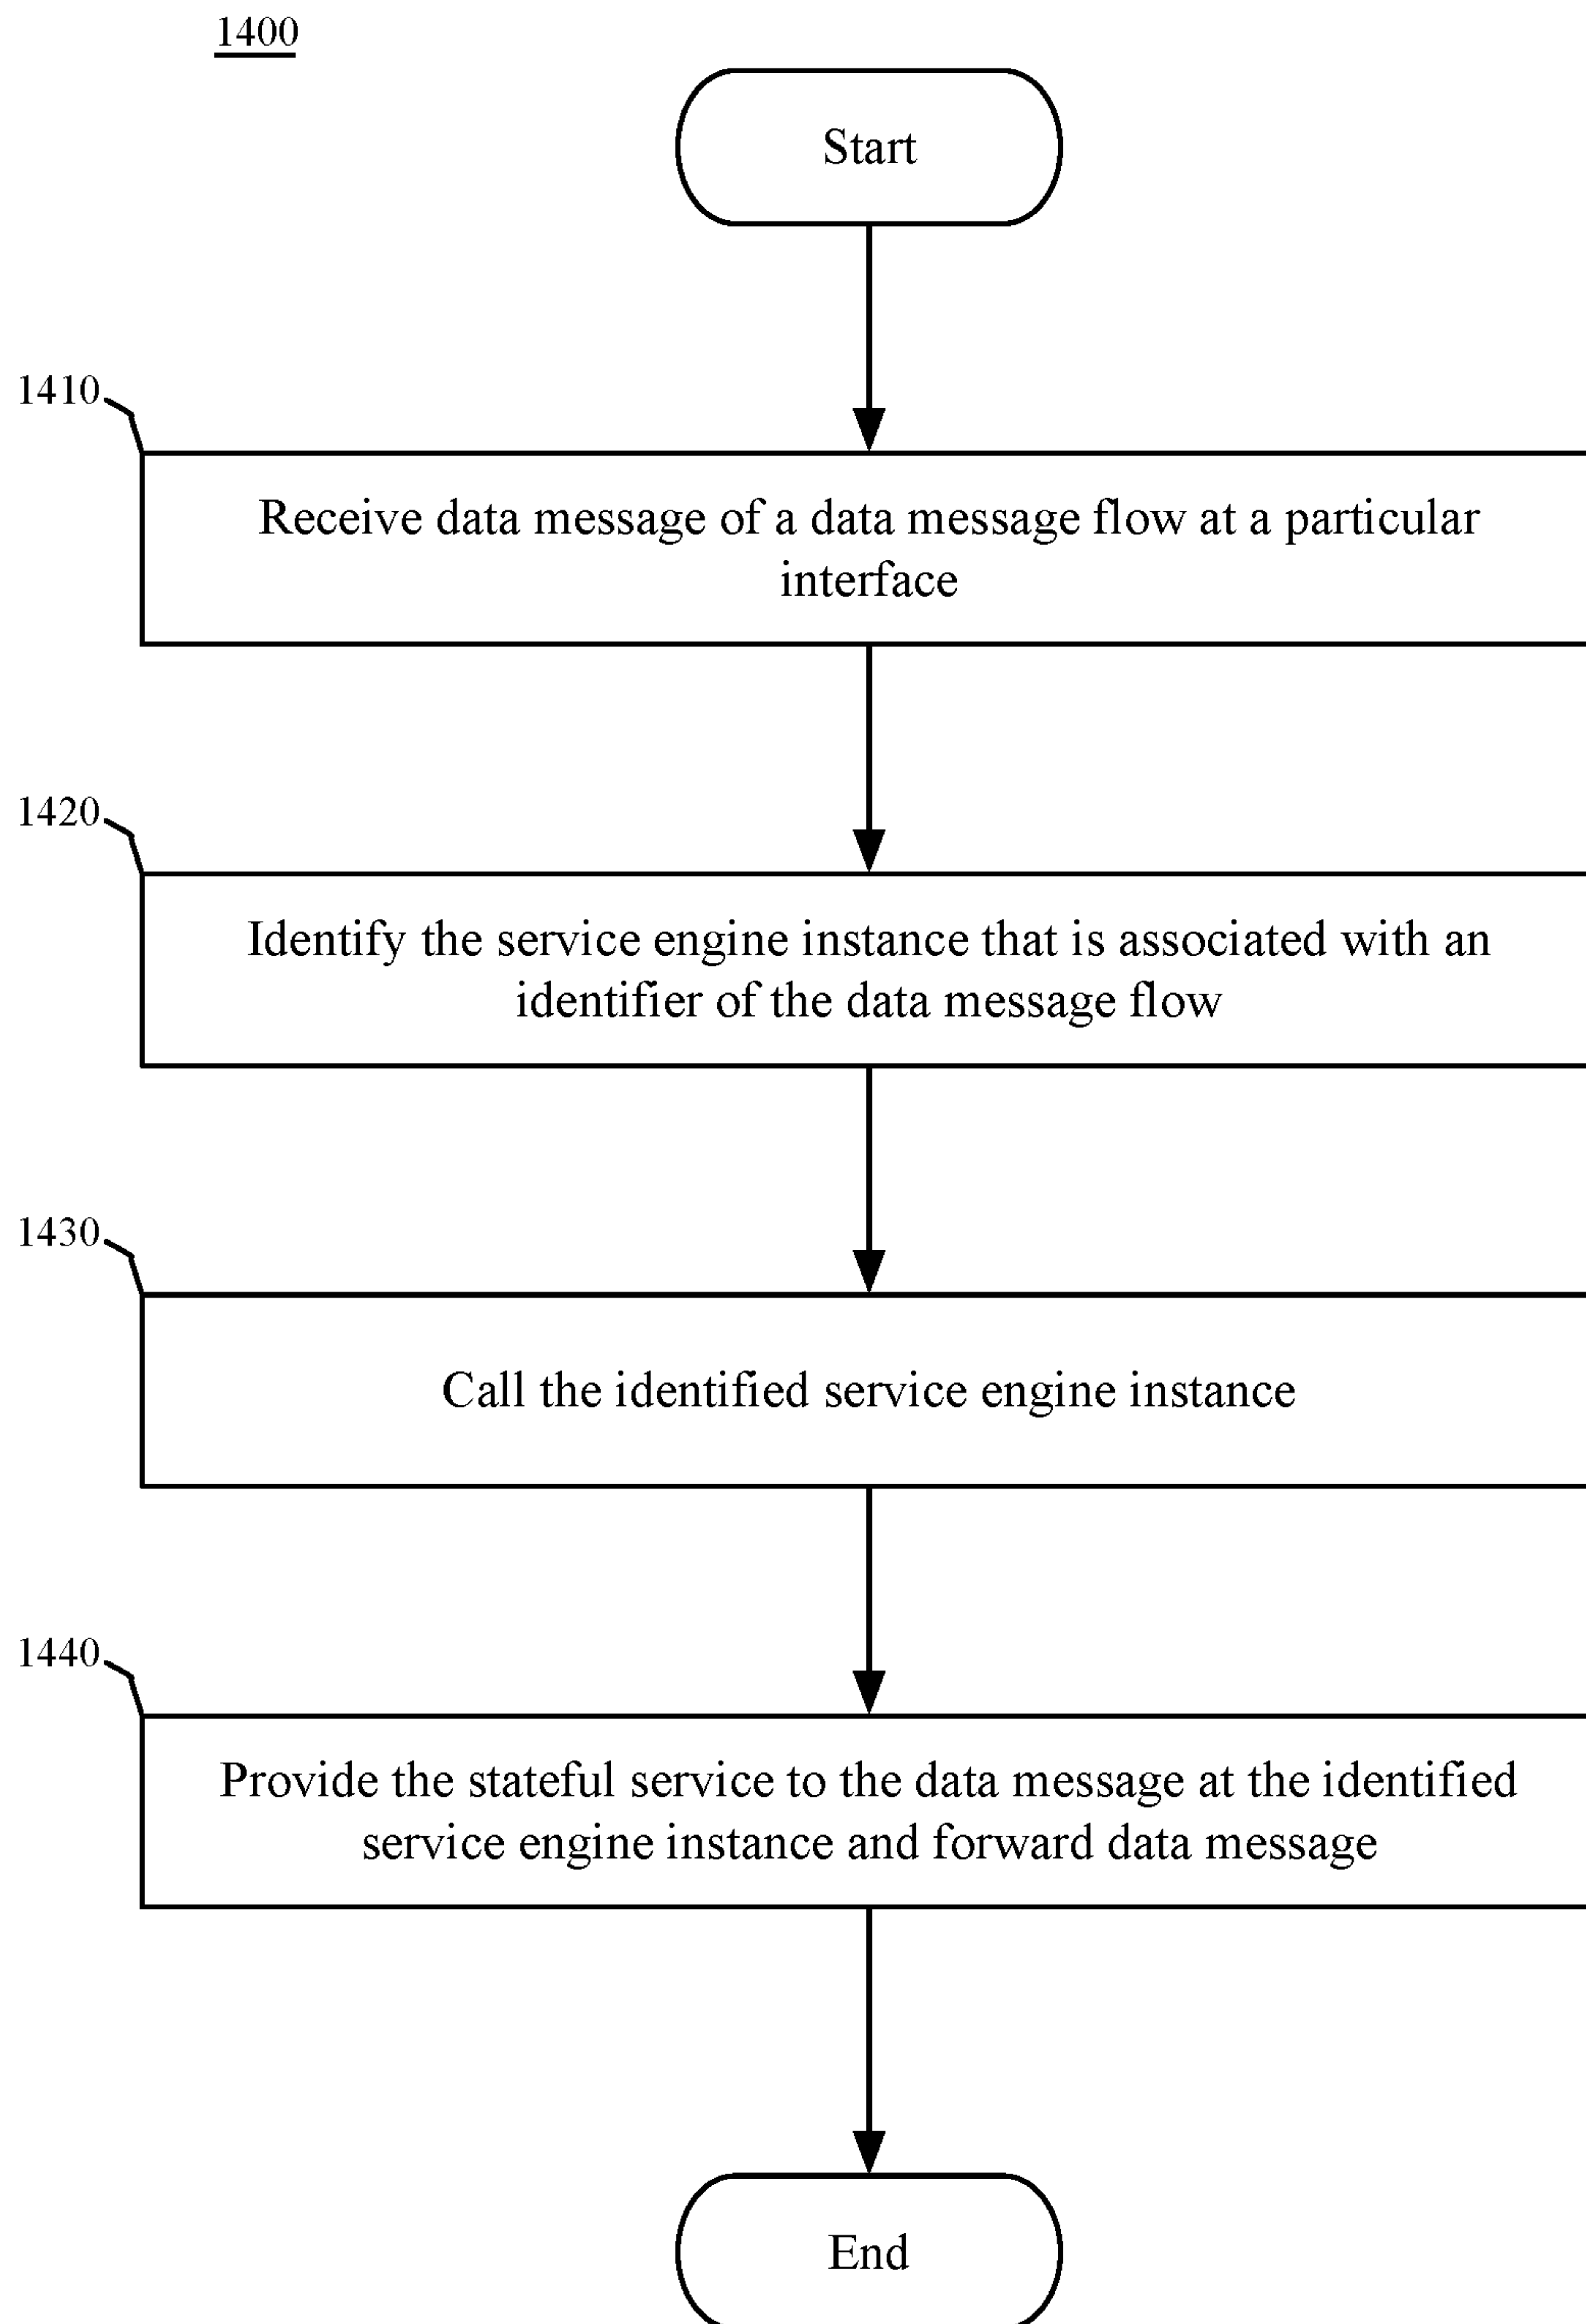
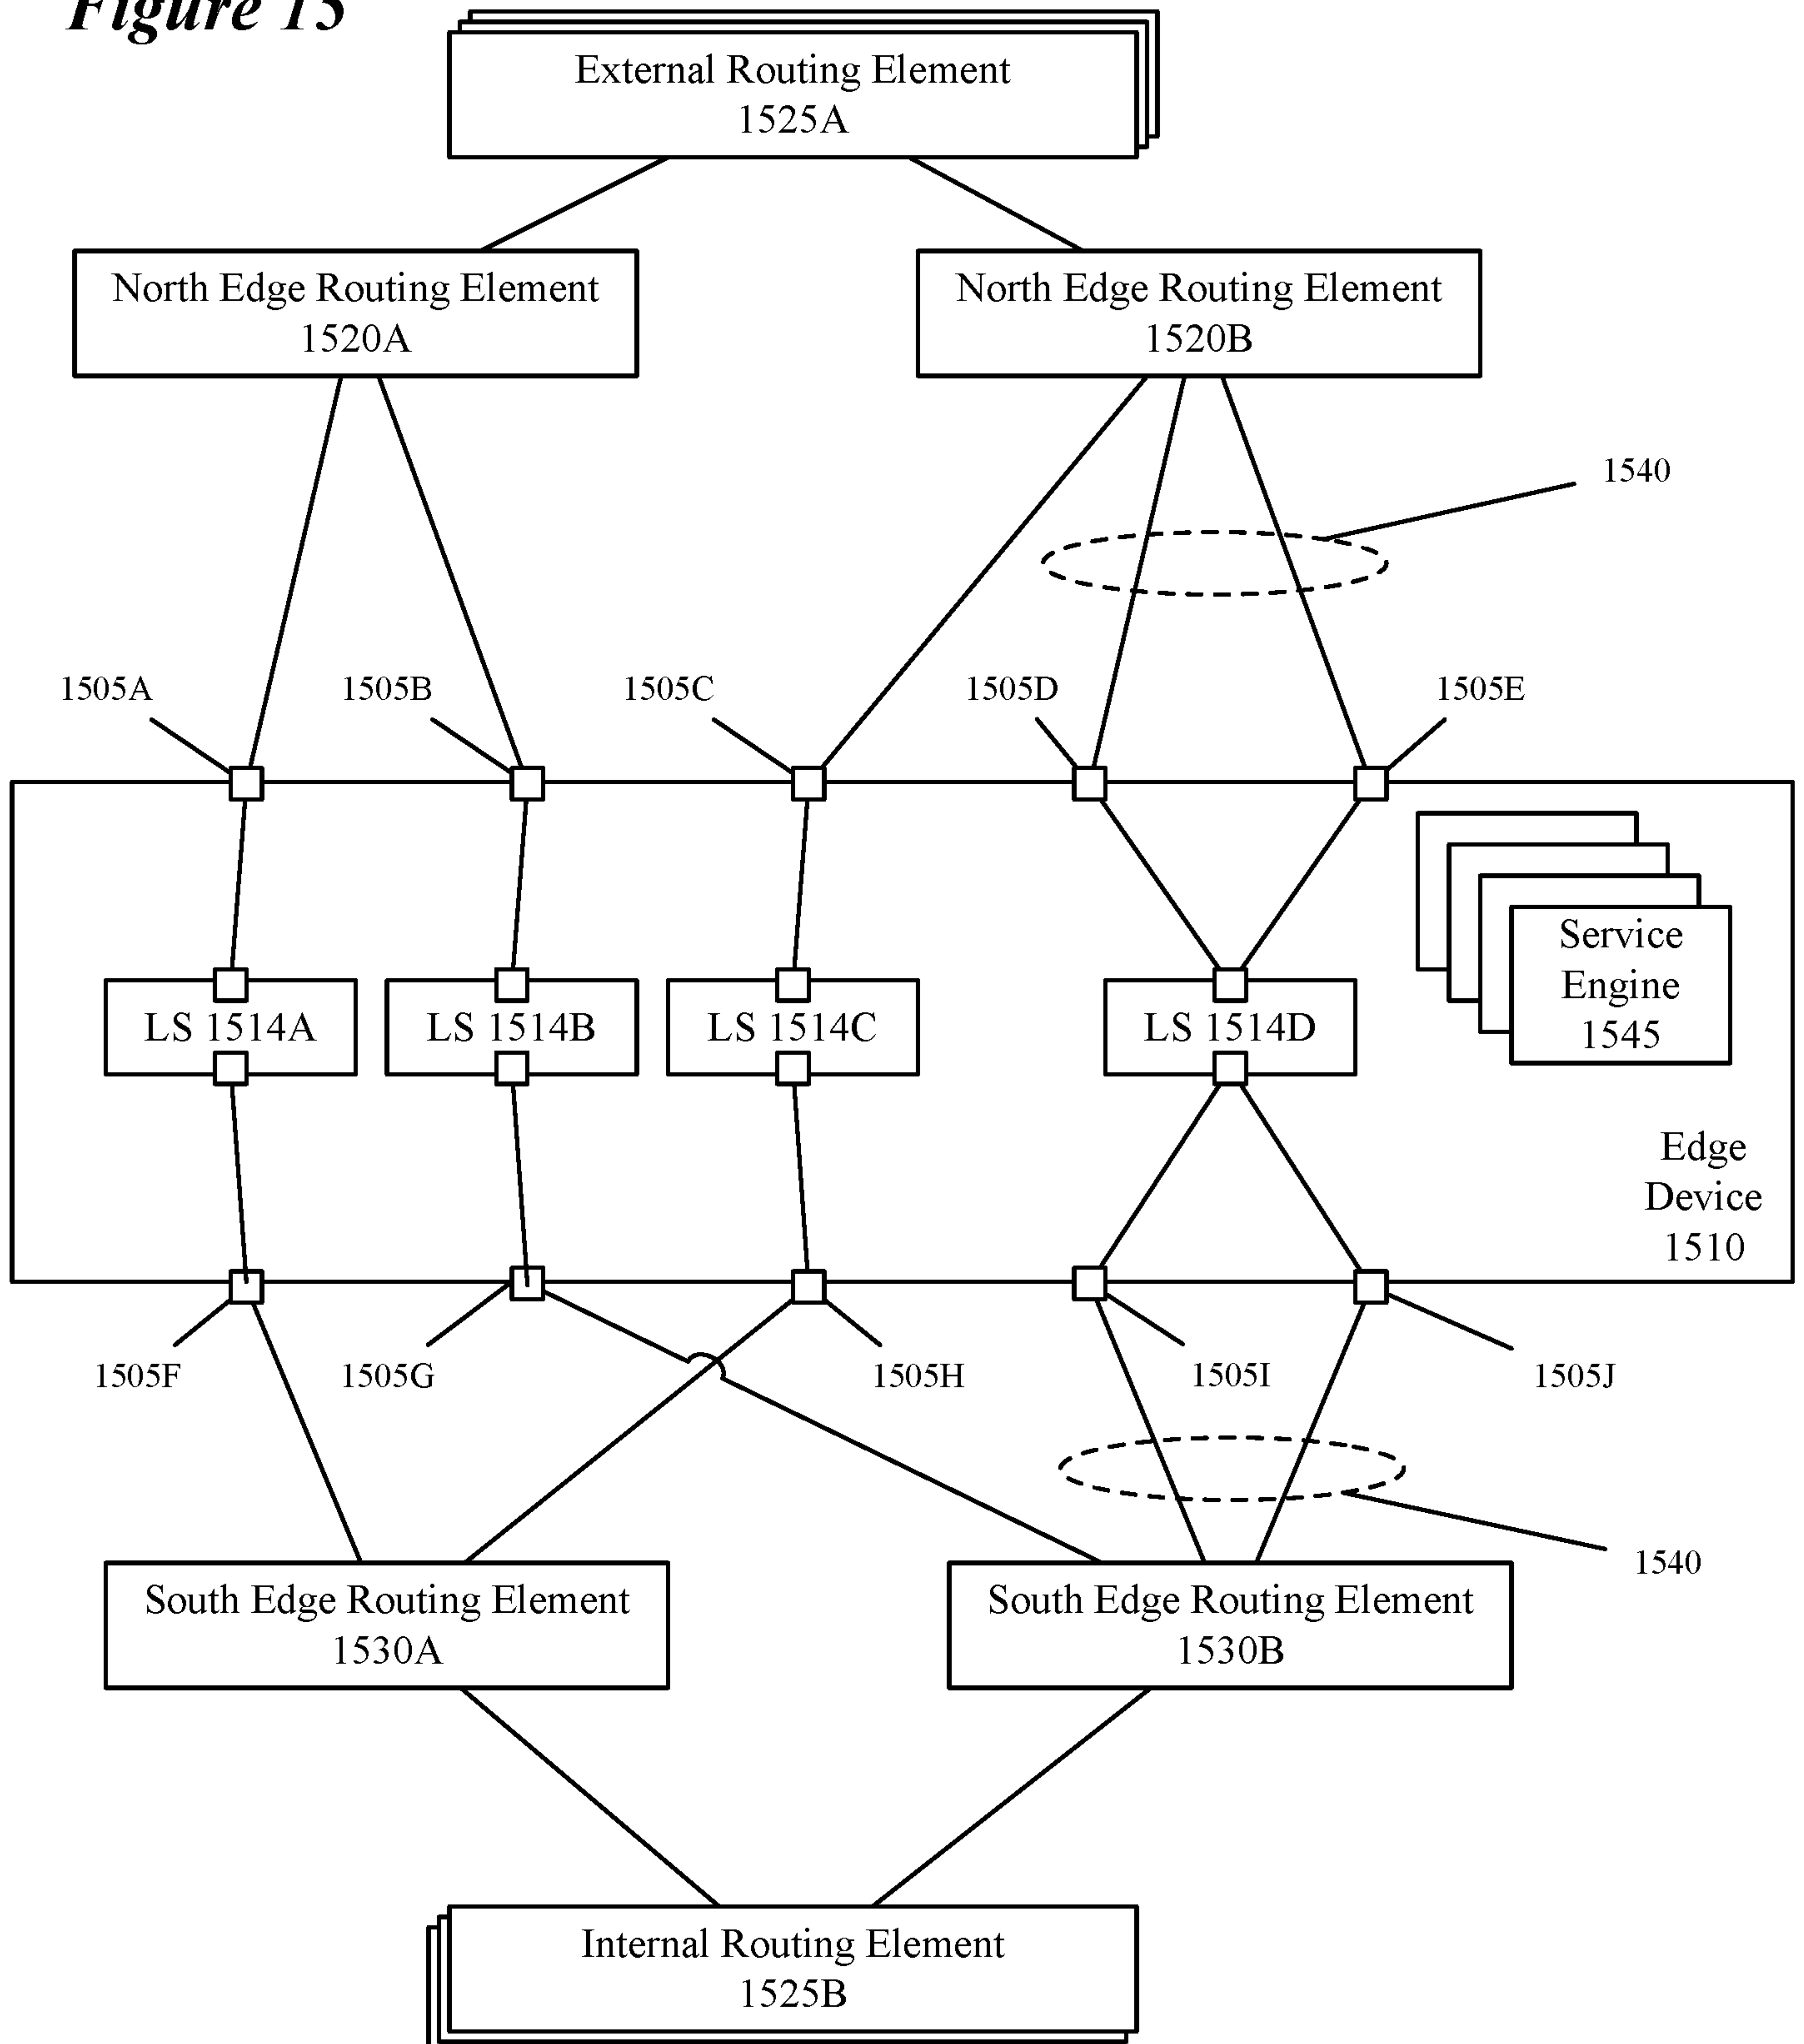
**Figure 14**

Figure 15



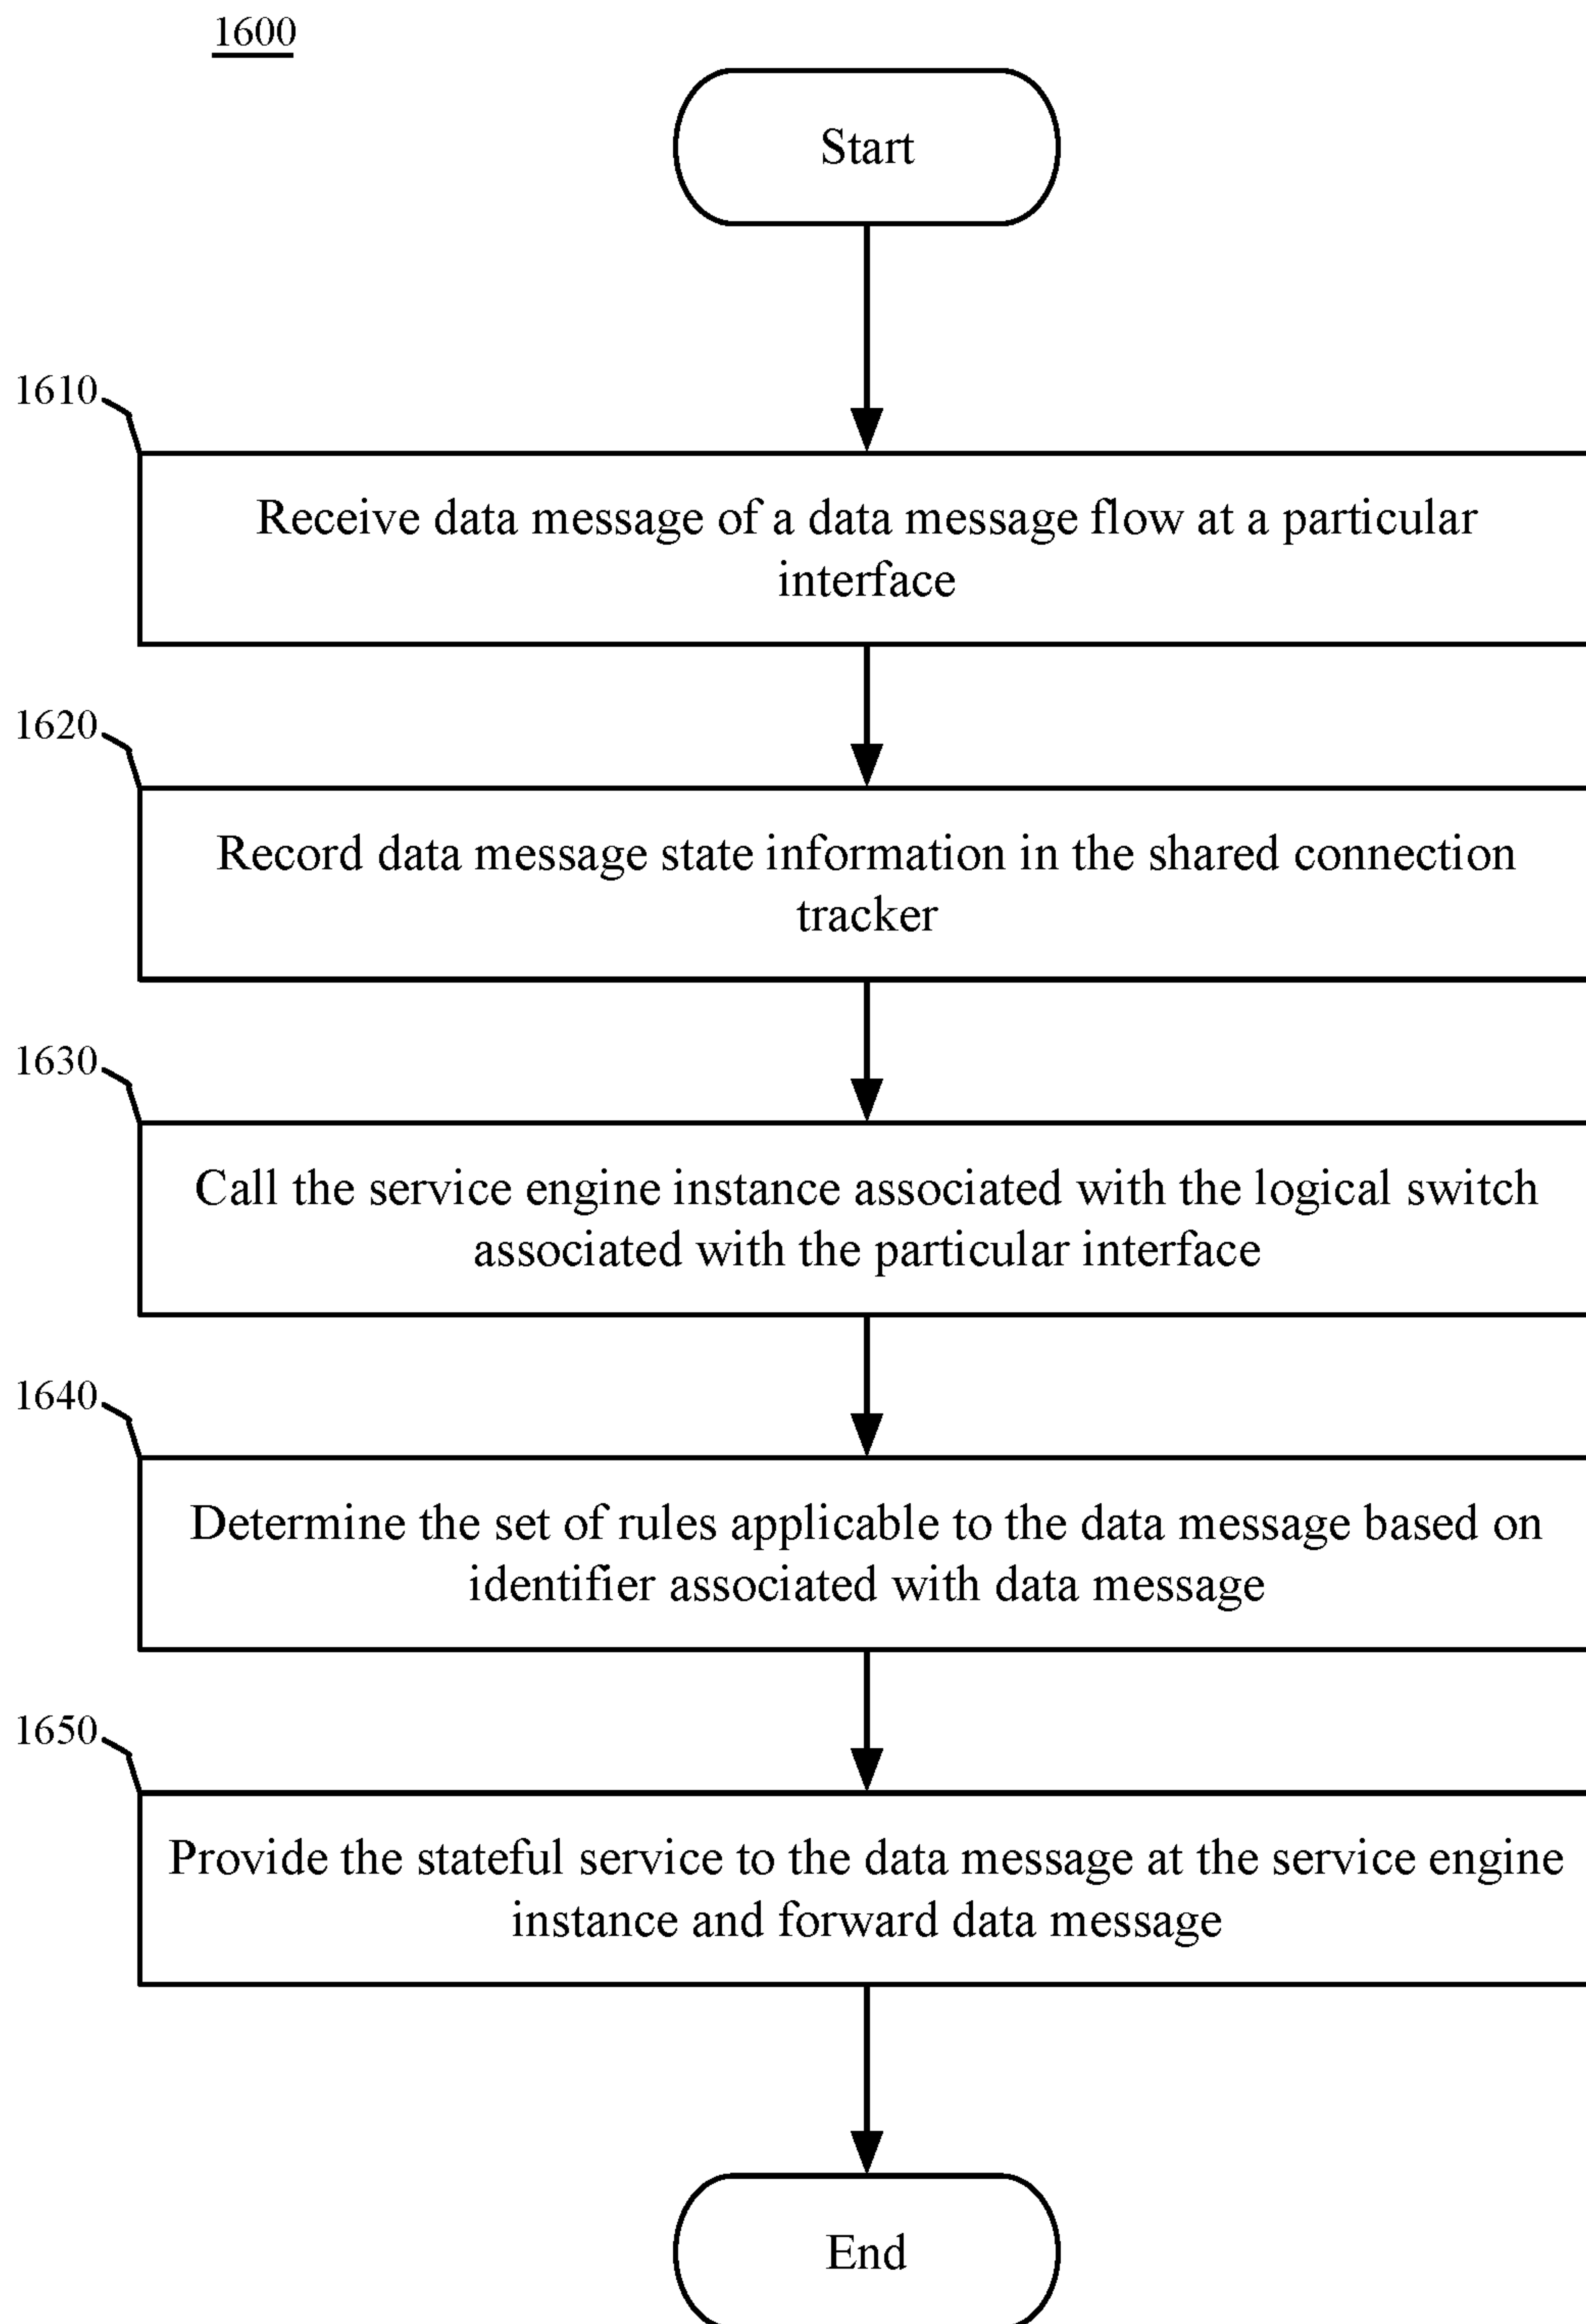
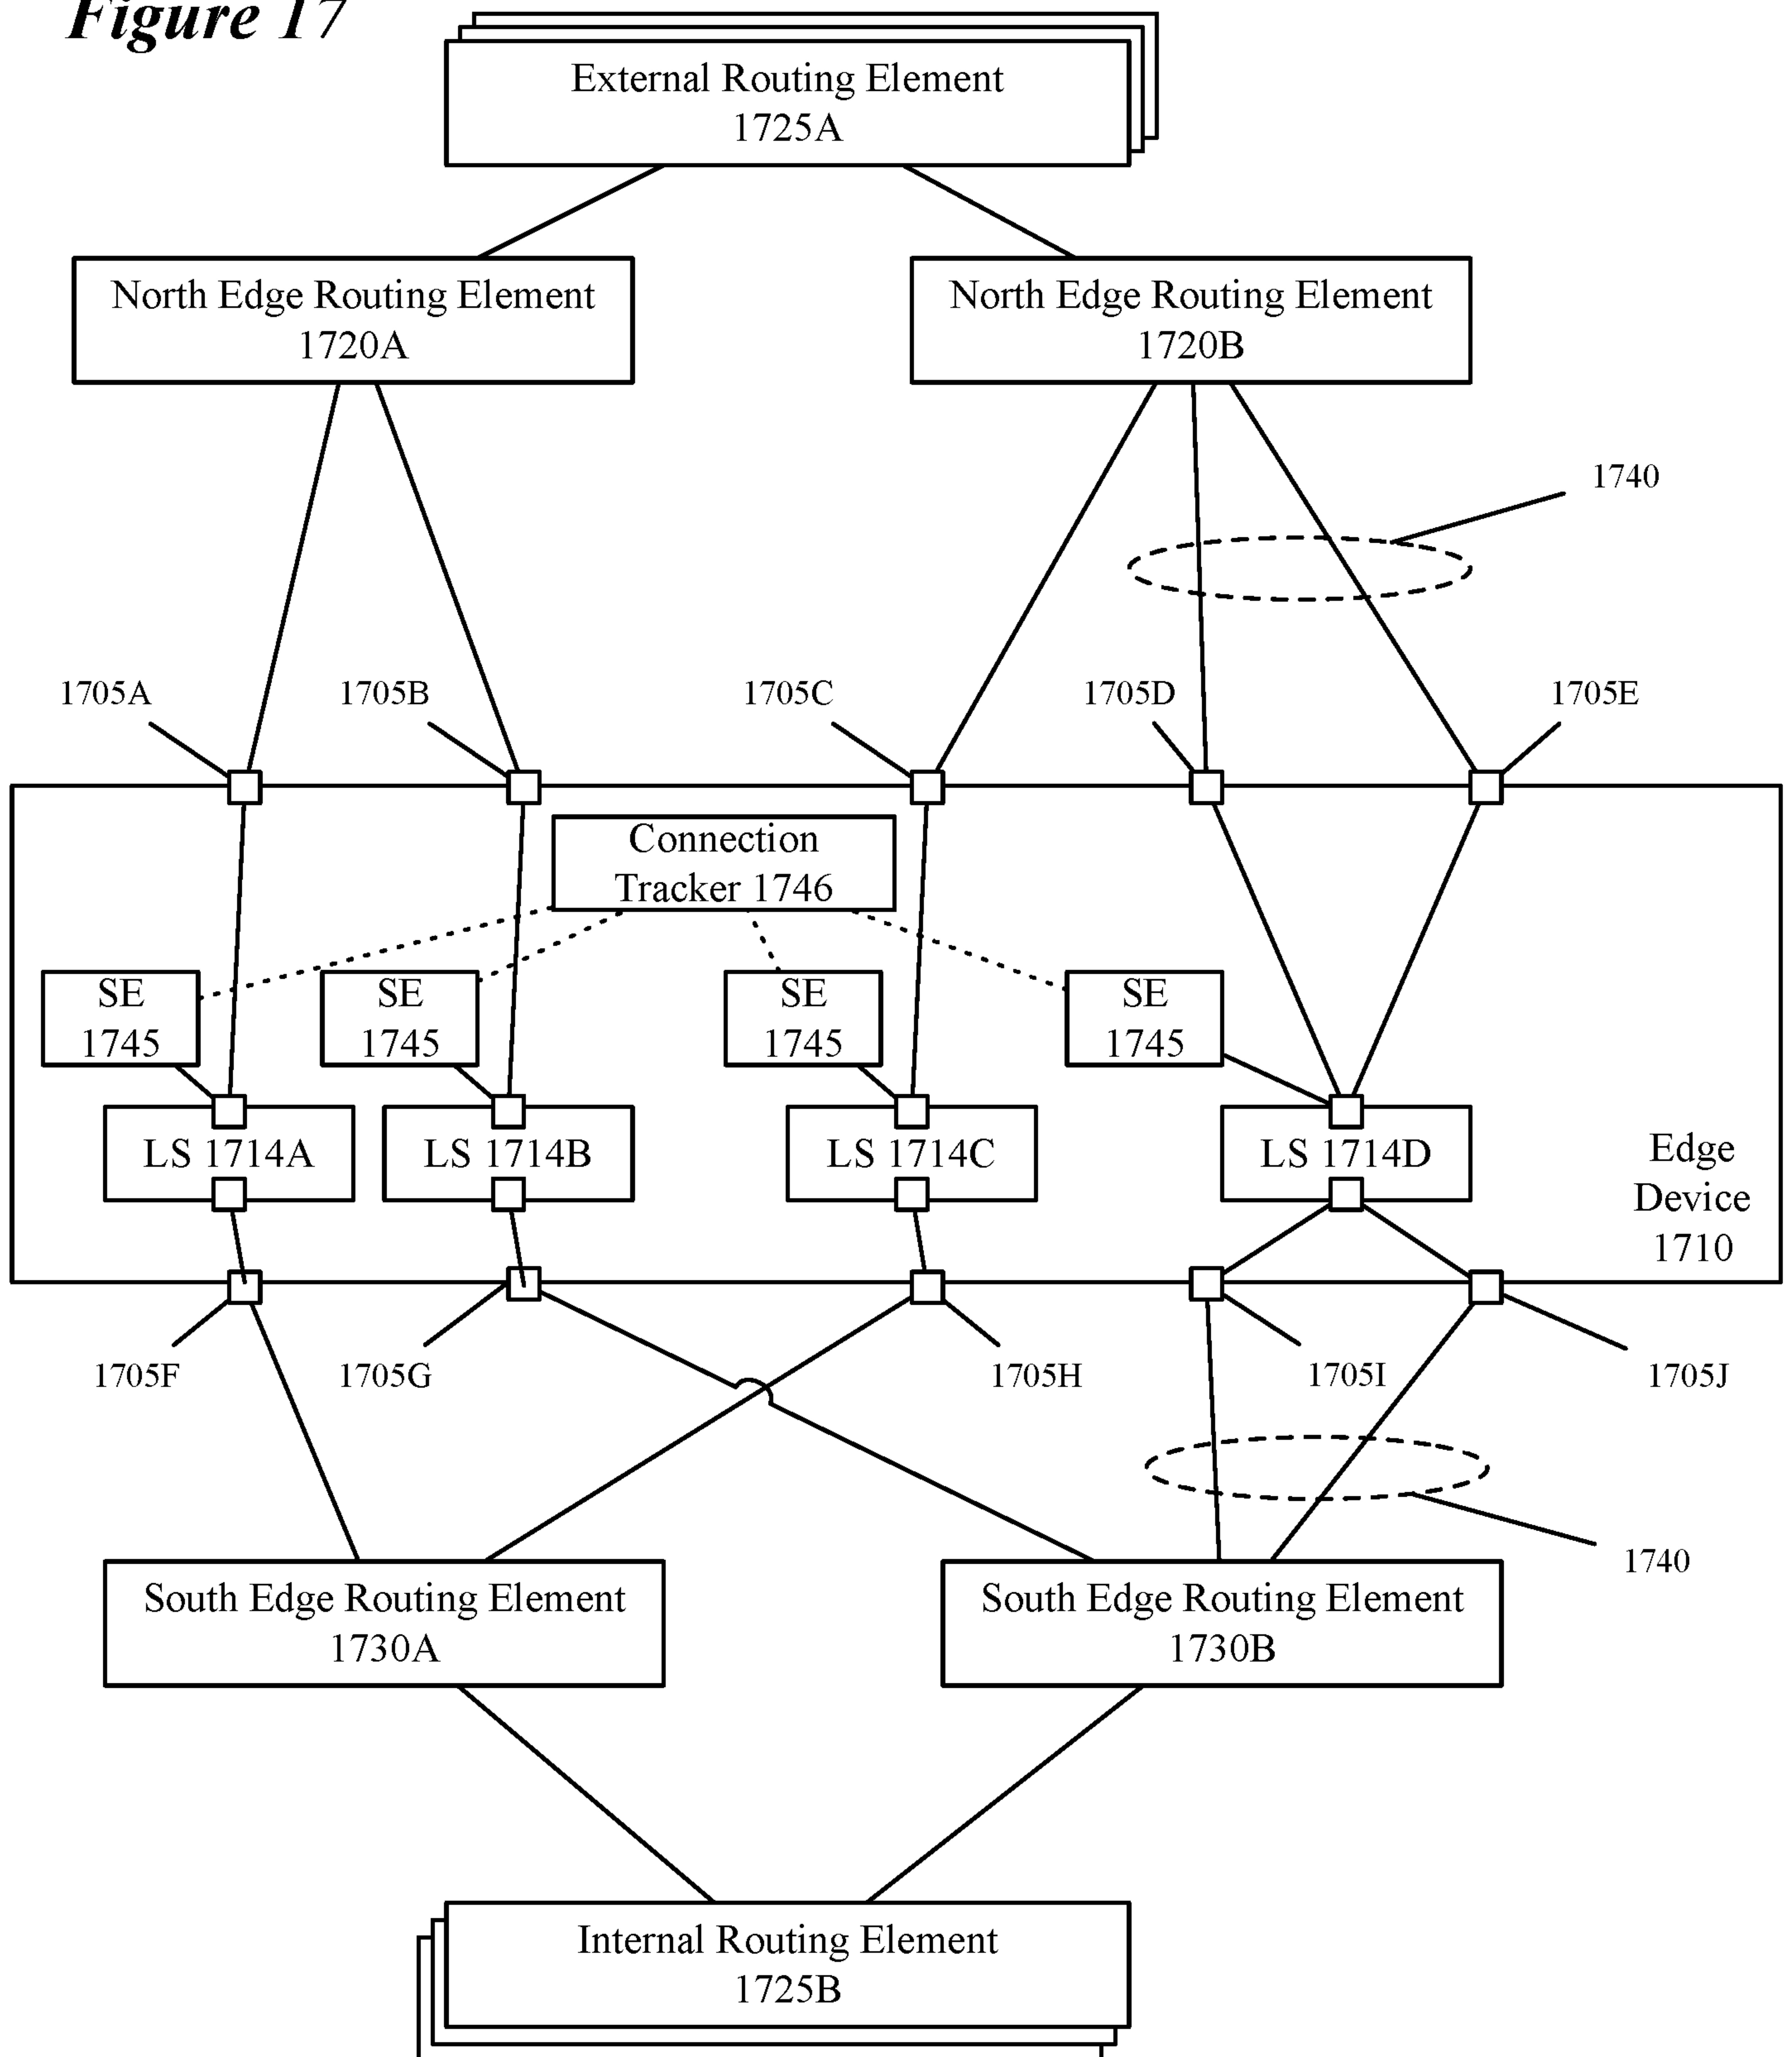


Figure 16

Figure 17



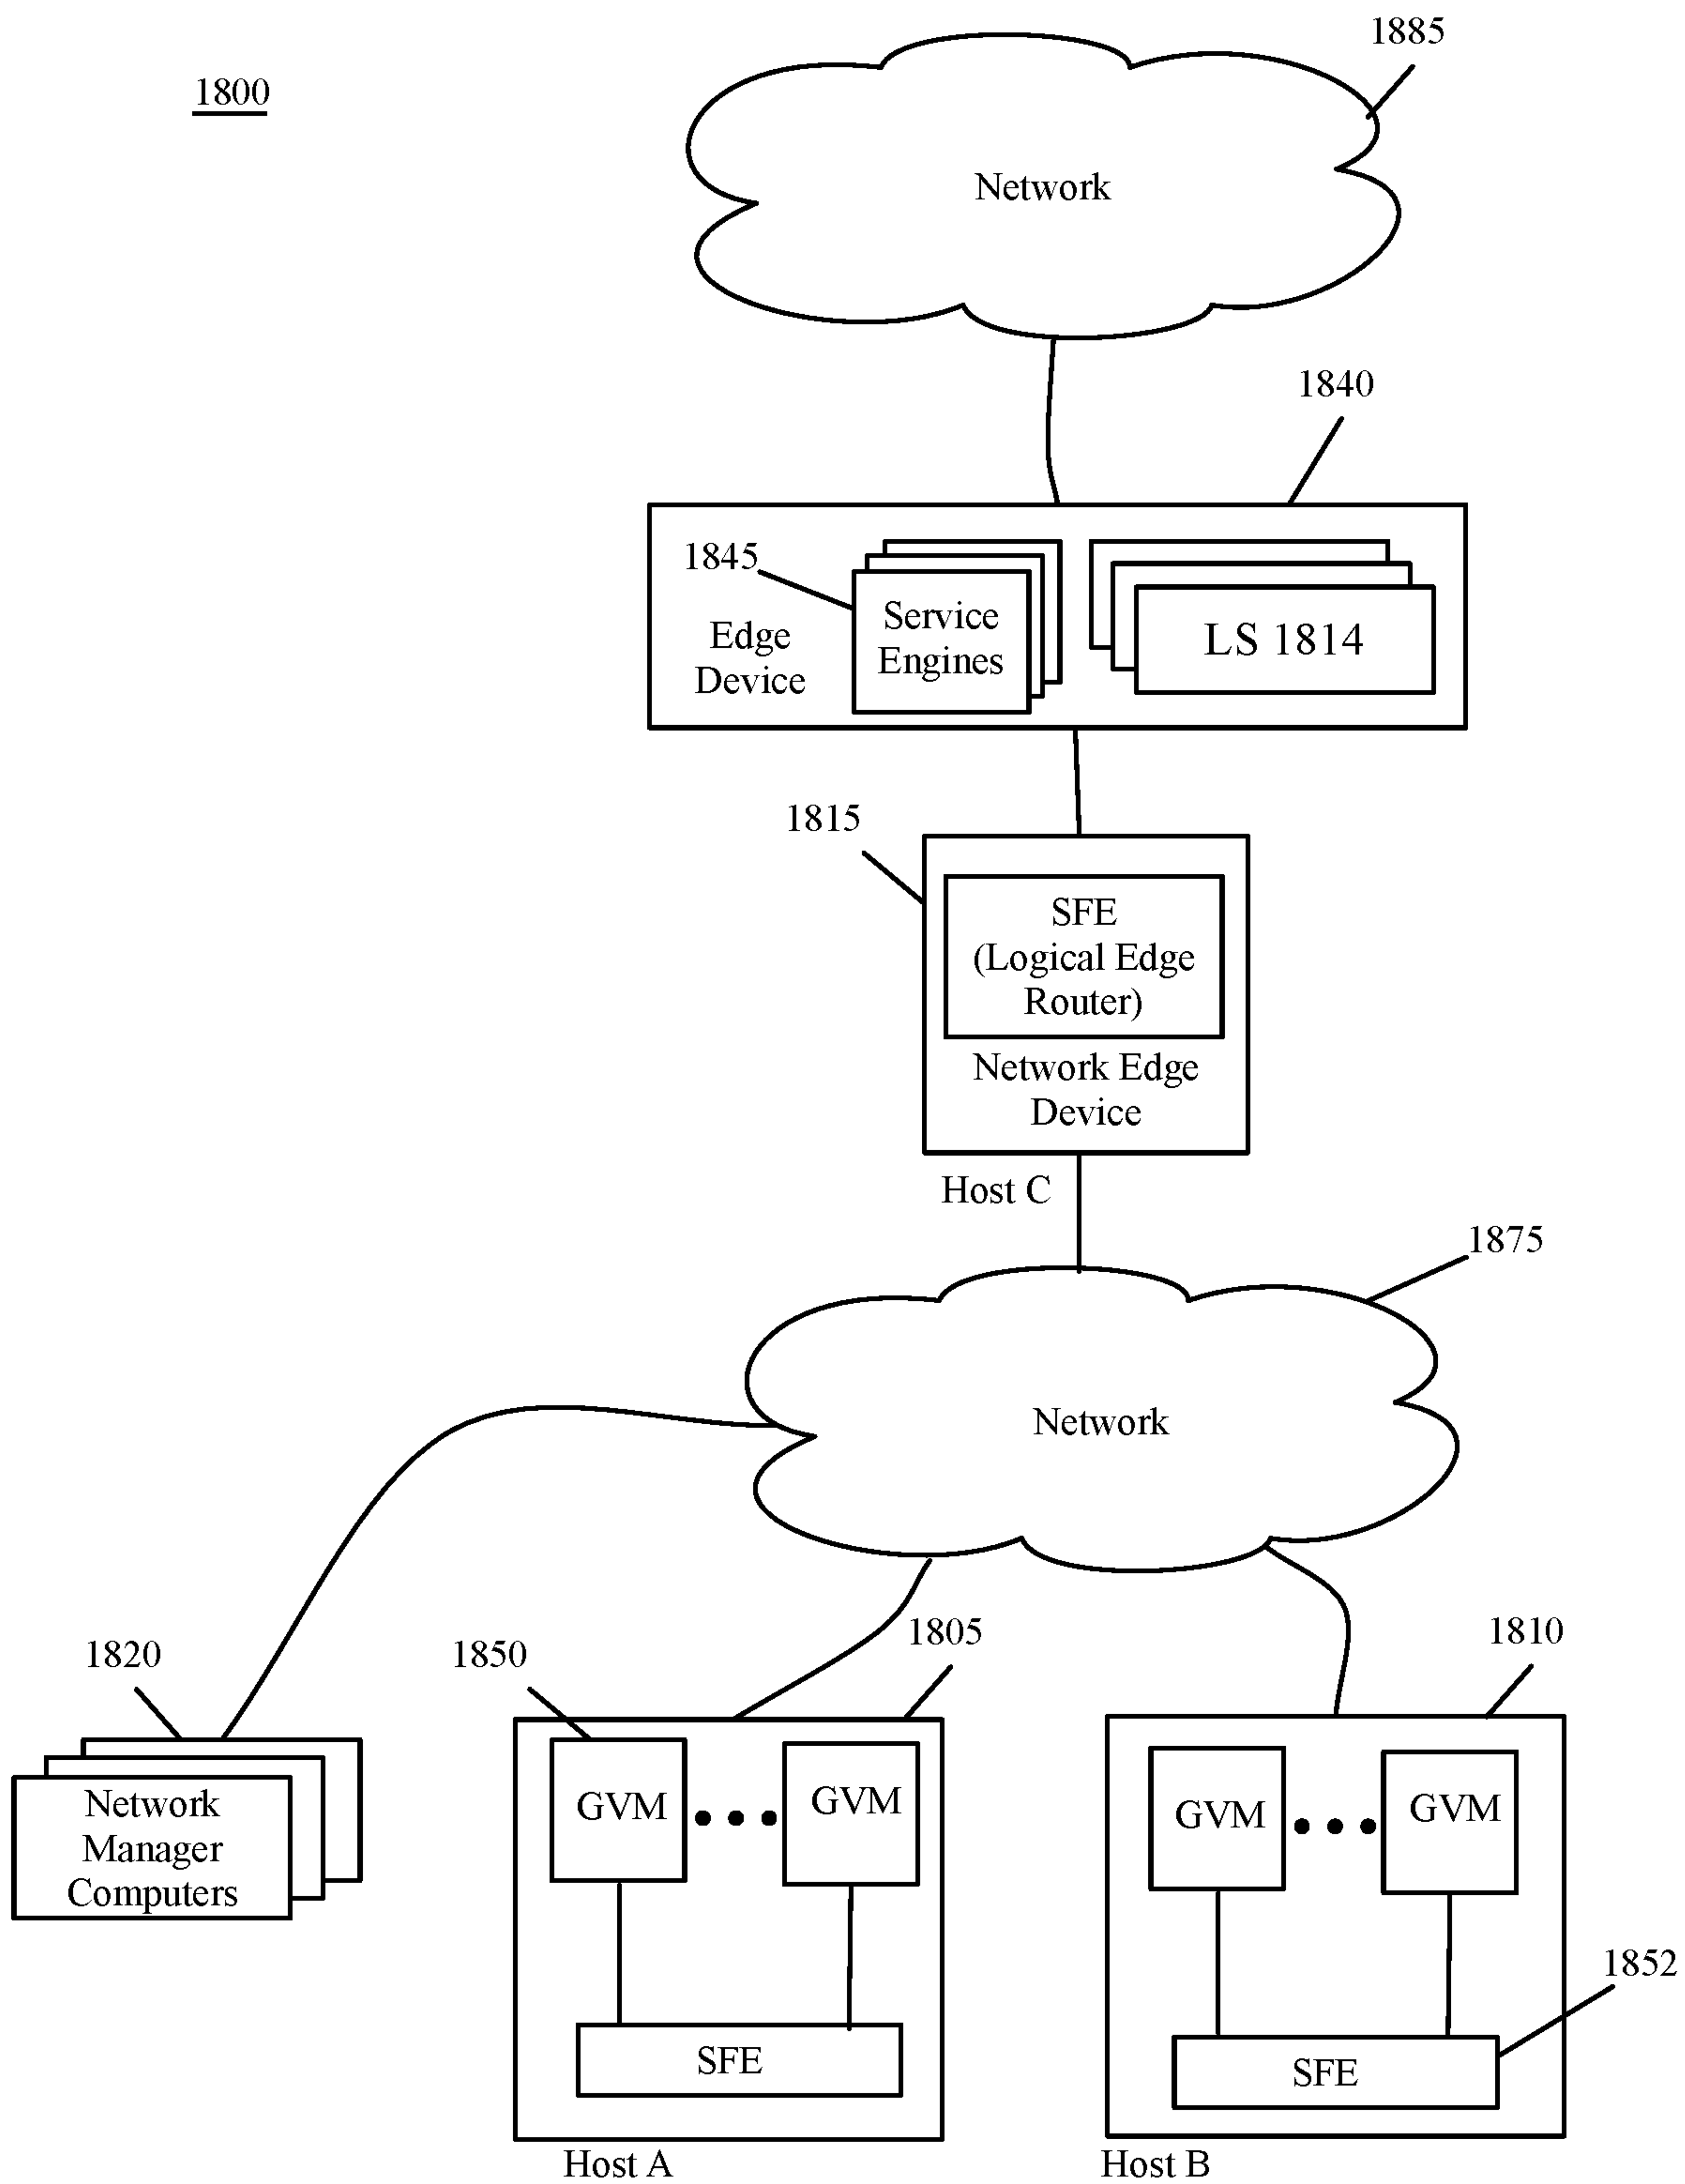


Figure 18




Figure 19

INTERNATIONAL SEARCH REPORT

International application No PCT/US2020/058896

A. CLASSIFICATION OF SUBJECT MATTER
 INV. H04L12/751 H04L12/709 H04L12/713
 ADD. H04L12/741

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

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 2015/341428 A1 (CHAUHAN ABHISHEK [US]) 26 November 2015 (2015-11-26)	1,3-7, 11-17
Y	paragraphs [0111] - [0122], [0187], [0188], [0232] - [0242]; figures 2B,5B,6 -----	8-10, 18-20
Y	US 2017/126812 A1 (SINGHAL PRAVIN [US]) 4 May 2017 (2017-05-04) paragraphs [0241], [0267] - [0273]; figures 5B,6 -----	8-10, 18-20
A	WO 2019/042372 A1 (NEW H3C TECH CO LTD [CN]) 7 March 2019 (2019-03-07) the whole document -----	1-20

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

21 January 2021

Date of mailing of the international search report

29/01/2021

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

Boureau, Vlad

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/US2020/058896

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015341428 A1	26-11-2015	EP 3146696 A1 US 2015341428 A1 US 2019028536 A1 WO 2015179106 A1	29-03-2017 26-11-2015 24-01-2019 26-11-2015

US 2017126812 A1	04-05-2017	CN 108476231 A EP 3369237 A1 JP 6648265 B2 JP 2019502972 A KR 20180054801 A US 2017126812 A1 US 2020382609 A1 WO 2017074766 A1	31-08-2018 05-09-2018 14-02-2020 31-01-2019 24-05-2018 04-05-2017 03-12-2020 04-05-2017

WO 2019042372 A1	07-03-2019	CN 108259328 A EP 3657740 A1 JP 2020532230 A US 2020389388 A1 WO 2019042372 A1	06-07-2018 27-05-2020 05-11-2020 10-12-2020 07-03-2019
