



US 20220303366A1

(19) **United States**

(12) **Patent Application Publication**  
**Brissette et al.**

(10) **Pub. No.: US 2022/0303366 A1**  
(43) **Pub. Date: Sep. 22, 2022**

(54) **ETHERNET VIRTUAL PRIVATE NETWORK USING SEGMENT ROUTING**

*H04L 61/5007* (2006.01)  
*H04L 67/63* (2006.01)  
*H04L 12/46* (2006.01)  
*H04L 45/74* (2006.01)  
*H04L 49/35* (2006.01)  
*H04L 45/741* (2006.01)

(71) Applicant: **Cisco Technology, Inc.**, San Jose, CA (US)

(72) Inventors: **Patrice Brissette**, Gatineau (CA); **Clarence Filsfils**, Brussels (BE); **Darren Dukes**, Ottawa (CA); **Gaurav Dawra**, San Jose, CA (US); **Francois Clad**, Strasbourg (FR); **Pablo Camarillo Garvia**, Madrid (ES)

(52) **U.S. Cl.**  
**CPC** ..... *H04L 69/22* (2013.01); *H04L 69/324* (2013.01); *H04L 45/34* (2013.01); *H04L 67/10* (2013.01); *H04L 45/04* (2013.01); *H04L 45/50* (2013.01); *H04L 61/5007* (2022.05); *H04L 67/63* (2022.05); *H04L 12/4633* (2013.01); *H04L 45/74* (2013.01); *H04L 45/14* (2013.01); *H04L 49/35* (2013.01); *H04L 45/741* (2013.01); *H04L 12/4641* (2013.01); *H04L 61/2503* (2013.01)

(21) Appl. No.: **17/836,069**

(22) Filed: **Jun. 9, 2022**

**Related U.S. Application Data**


(63) Continuation of application No. 17/671,188, filed on Feb. 14, 2022, which is a continuation of application No. 16/697,134, filed on Nov. 26, 2019, now Pat. No. 11,283,906, which is a continuation of application No. 15/783,753, filed on Oct. 13, 2017, now Pat. No. 10,491,720.  
(60) Provisional application No. 62/525,439, filed on Jun. 27, 2017.

**Publication Classification**

(51) **Int. Cl.**  
*H04L 69/22* (2006.01)  
*H04L 69/324* (2006.01)  
*H04L 45/00* (2006.01)  
*H04L 67/10* (2006.01)  
*H04L 45/02* (2006.01)  
*H04L 45/50* (2006.01)

(57) **ABSTRACT**

In one embodiment, Ethernet Virtual Private Network (EVPN) is implemented using Internet Protocol Version 6 (IPv6) Segment Routing (SRv6) underlay network and SRv6-enhanced Border Gateway Protocol (BGP) signaling. A particular route associated with a particular Internet Protocol Version 6 (IPv6) Segment Routing (SRv6) Segment Identifier (SID) is advertised in a particular route advertisement message of a routing protocol (e.g., BGP). The SID includes encoding representing a particular Ethernet Virtual Private Network (EVPN) Layer 2 (L2) flooding Segment Routing end function of the particular router and a particular Ethernet Segment Identifier (ESI), with the particular SID including a routable prefix to the particular router. The particular router receives a particular packet including the particular SID; and in response, the particular router performs the particular EVPN end function on the particular packet.






FIGURE 1A

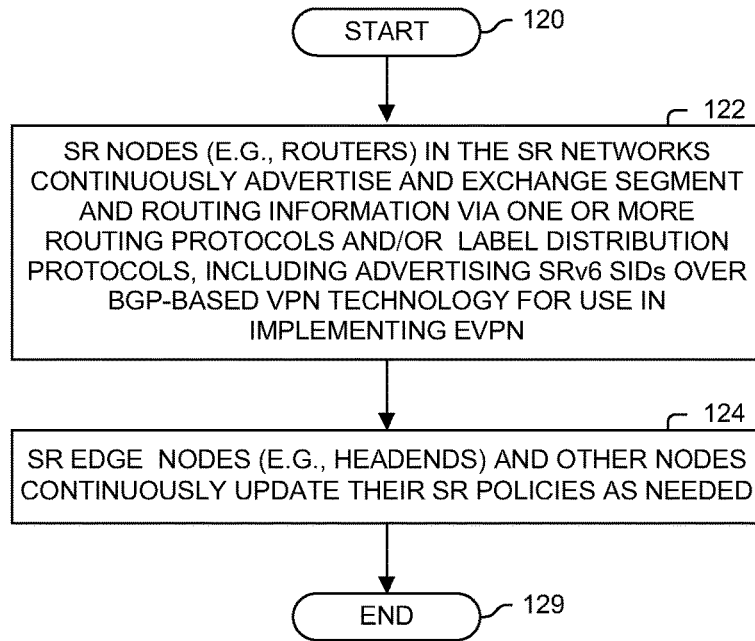


FIGURE 1B

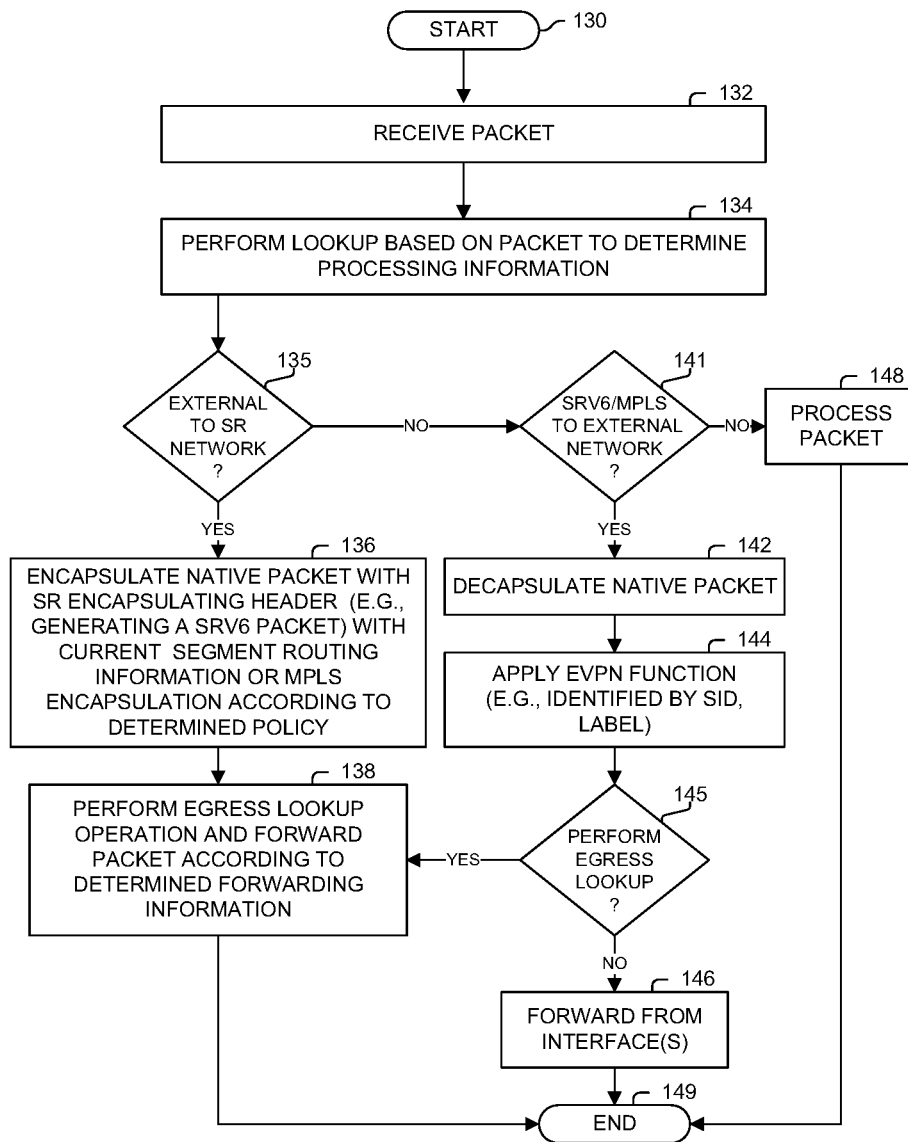


FIGURE 1C

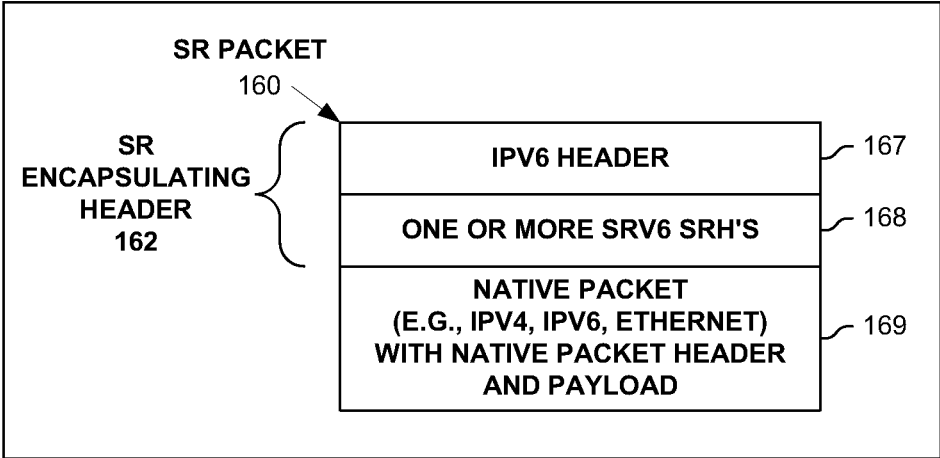


FIGURE 1D

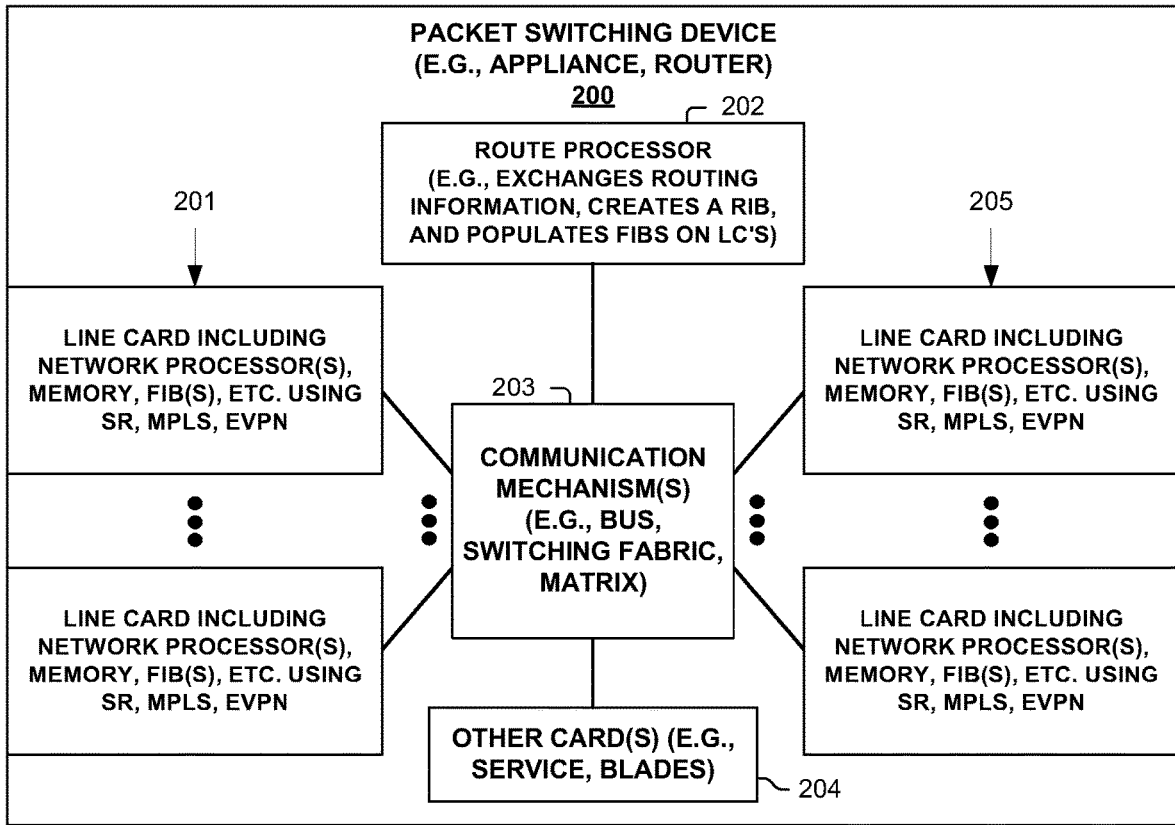


FIGURE 2A

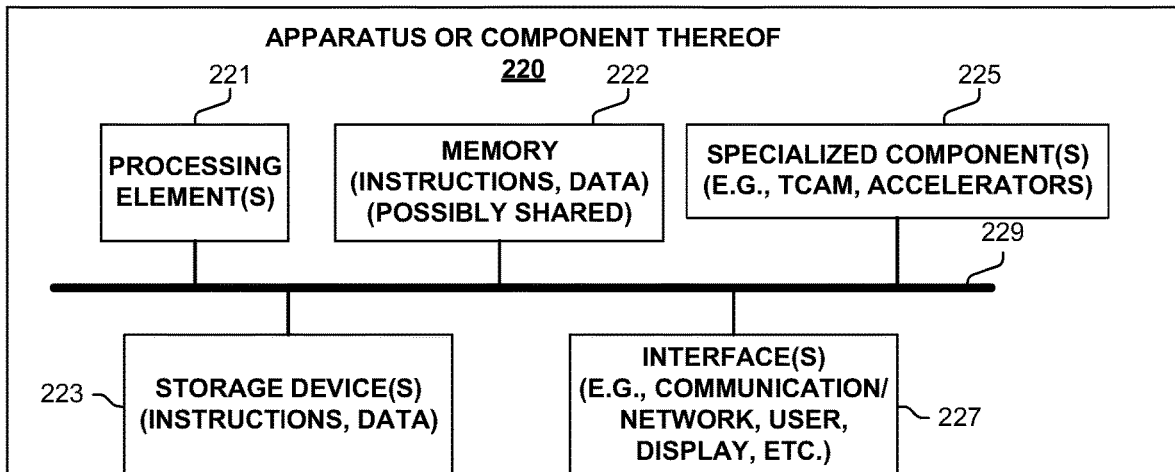


FIGURE 2B

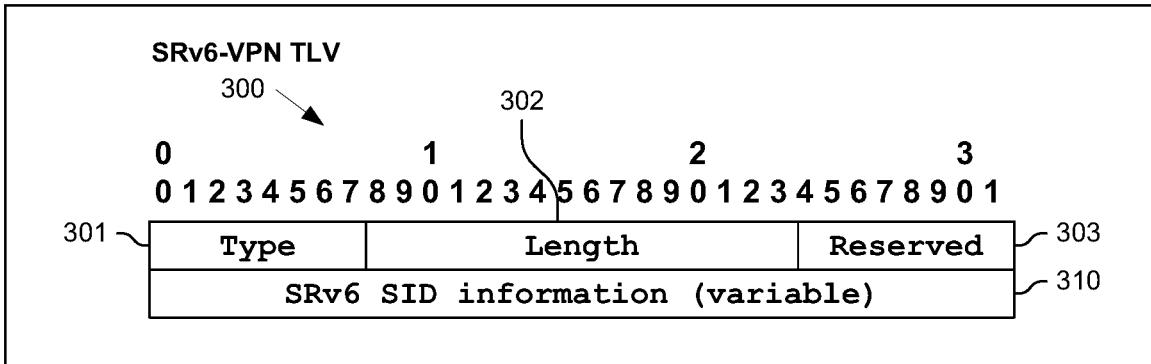


FIG. 3A

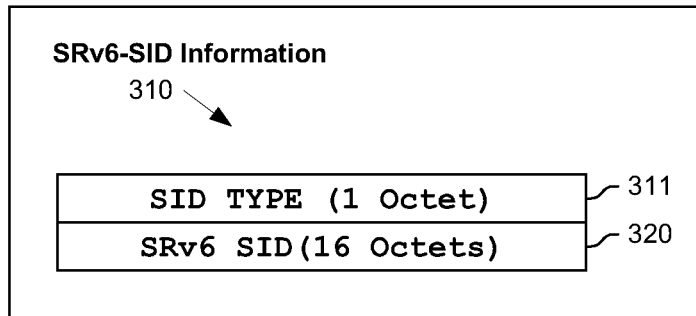


FIG. 3B


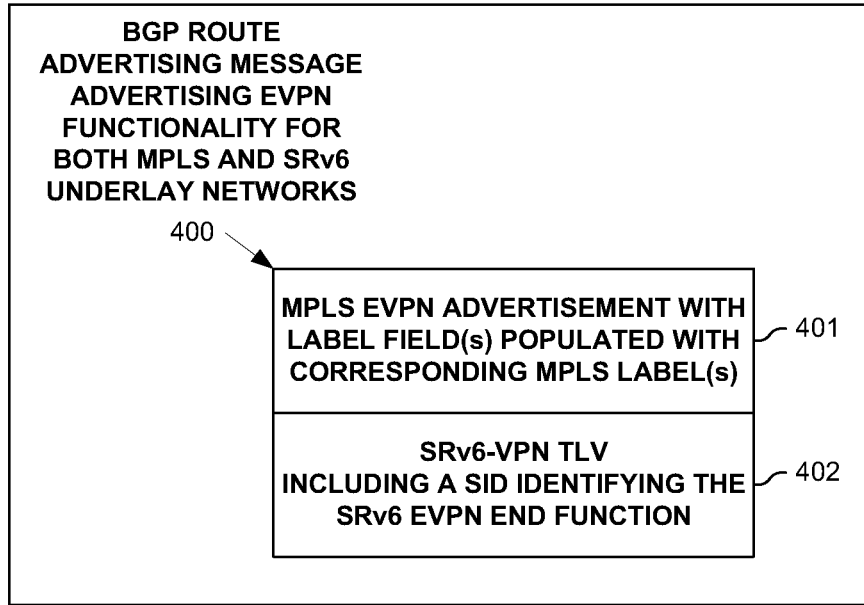
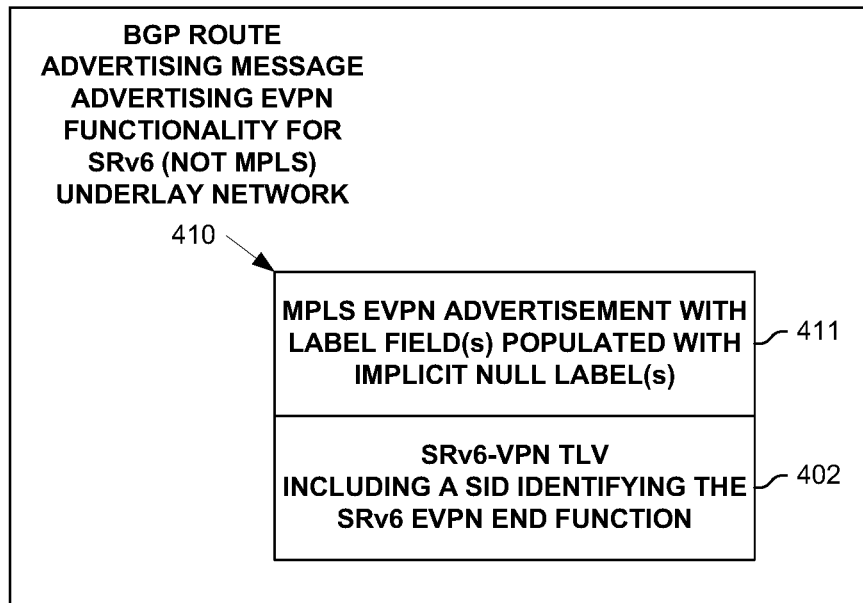


FIG. 3C



**FIGURE 4A**



**FIGURE 4B**

## ETHERNET VIRTUAL PRIVATE NETWORK USING SEGMENT ROUTING

### CROSS-REFERENCE TO RELATED APPLICATIONS

**[0001]** This application is a continuation of U.S. patent application Ser. No. 17/671,188, filed on Feb. 14, 2022, which in turn, is a continuation of U.S. patent application Ser. No. 16/697,134, filed Nov. 26, 2019, which in turn, is a continuation of U.S. patent application Ser. No. 15/783,753, filed Oct. 13, 2017, now U.S. Pat. No. 10,491,720, which claims the benefit of priority to U.S. Provisional Application No. 62/525,439, filed Jun. 27, 2017, all of which are hereby incorporated by reference in their entirety.

### TECHNICAL FIELD

**[0002]** The present disclosure relates generally to sending of packets through a packet network, such as, but not limited to, according to Segment Routing of packets through a packet network.

### BACKGROUND

**[0003]** The communications industry is rapidly changing to adjust to emerging technologies and ever increasing customer demand. This customer demand for new applications and increased performance of existing applications is driving communications network and system providers to employ networks and systems having greater speed and capacity (e.g., greater bandwidth). In trying to achieve these goals, a common approach taken by many communications providers is to use packet switching technology. Packets are typically forwarded in a network forwarded based on one or more values representing network nodes or paths.

### BRIEF DESCRIPTION OF THE DRAWINGS

**[0004]** The appended claims set forth the features of one or more embodiments with particularity. The embodiment (s), together with its advantages, may be understood from the following detailed description taken in conjunction with the accompanying drawings of which:

**[0005]** FIG. 1A illustrates a network operating according to one embodiment;

**[0006]** FIG. 1B illustrates a process according to one embodiment;

**[0007]** FIG. 1C illustrates a process according to one embodiment;

**[0008]** FIG. 1D illustrates different Segment Routing (SR) packet formats according to one embodiment;

**[0009]** FIG. 2A illustrates a packet switching device according to one embodiment;

**[0010]** FIG. 2B illustrates an apparatus according to one embodiment;

**[0011]** FIG. 3A illustrates a SRv6 Type-Length-Value (TLV) according to one embodiment;

**[0012]** FIG. 3B illustrates SRv6-SID information according to one embodiment;

**[0013]** FIG. 3C illustrates SRv6-VPN encoding according to one embodiment;

**[0014]** FIG. 4A illustrates a SRv6-enhanced Border Gateway Protocol (BGP) route advertisement message according to one embodiment; and

**[0015]** FIG. 4B illustrates a SRv6-enhanced Border Gateway Protocol (BGP) route advertisement message according to one embodiment.

### DESCRIPTION OF EXAMPLE EMBODIMENTS

#### 1. Overview

**[0016]** Disclosed are, inter alia, methods, apparatus, computer-storage media, mechanisms, and means associated with Ethernet Virtual Private Network (EVPN) using Internet Protocol Version 6 (IPv6) Segment Routing (SRv6) underlay network and SRv6-enhanced Border Gateway Protocol (BGP) signaling.

**[0017]** One embodiment includes: advertising, in a network via a particular route advertisement message of a routing protocol, a particular route associated with a particular Internet Protocol Version 6 (IPv6) Segment Routing (SRv6) Segment Identifier (SID), with the SID including a locator of a particular router and a function encoding representing a particular Ethernet Virtual Private Network (EVPN) end function of the particular router, with the particular SID including a routable prefix to the particular router; receiving, by the particular router, a particular packet including the particular SID; and in response to said received particular packet including the particular SID, the particular router performing the particular EVPN end function on the particular packet.

**[0018]** One embodiment includes: receiving an advertisement of the particular route being associated with the particular SID by an ingress router in the network; receiving an Ethernet frame by the ingress router; encapsulating the Ethernet frame into the particular packet by the ingress router; and sending the particular packet into the network by the ingress router; wherein said performing the particular EVPN end function on the particular packet includes decapsulating the Ethernet frame from the packet and sending the Ethernet frame from the particular router.

**[0019]** In one embodiment, the particular packet includes a Segment Routing Header (SRH) including the particular SID as the currently active SID. In one embodiment, the routing protocol is Border Gateway Protocol, and the particular route advertising message advertises a BGP route type 1, 2, 3, or 5, and associated with a particular SRv6-VPN Type Length Value (TLV) including the particular SID. In one embodiment, the particular route advertising message includes advertising the particular route being associated with one or more BGP Multiprotocol Label Switching-based (MPLS-based) EVPN labels for which to use in an MPLS packet to invoke corresponding BGP MPLS-based EVPN functionality on a corresponding MPLS packet by the particular router.

**[0020]** One embodiment includes receiving, by an ingress router, an advertisement of the particular route being associated with both the particular SID and said one or more MPLS-EVPN labels; receiving an Ethernet frame by the ingress router; encapsulating the Ethernet frame into the particular packet by the ingress router; and sending the particular packet into the network by the ingress router. In one embodiment, performing the particular EVPN end function on the particular packet includes decapsulating the Ethernet frame from the packet and sending the Ethernet frame from the particular router.

**[0021]** One embodiment includes: receiving, by a second ingress router, an advertisement of the particular route being



associated with both the particular SID and said one or more MPLS-EVPN labels; receiving a second Ethernet frame by the second ingress router; encapsulating the second Ethernet frame into a particular MPLS packet by the second ingress router; sending the particular MPLS packet into the network by the second ingress router; and performing said corresponding BGP MPLS-based EVPN functionality on the particular MPLS packet which includes decapsulating the second Ethernet frame from the particular MPLS packet and sending the second Ethernet frame from the particular router.

## 2. Description

**[0022]** Disclosed are, inter alia, methods, apparatus, computer-storage media, mechanisms, and means associated with EVPN using a SRv6 Underlay Network and SRv6-enhanced Border Gateway Protocol (BGP) signaling. As used herein Segment Routing refers to SRv6. Also, a segment identifier (SID) related to SRv6 typically refers to an IPv6 address, however, it also includes other SRv6 encoding (s) (e.g., compact encoding of SIDs) if so used in the embodiment.

**[0023]** As used herein, a node is a device in a network; a router (also referred to herein as a packet switching device) is a node that forwards received packets not explicitly addressed to itself (e.g., an L2 or L3 packet switching device); and a host is any node that is not a router.

**[0024]** The term “route” is used to refer to a fully or partially expanded prefix (e.g., 10.0.0.1, 10.0.\*.\*), BGP Network Layer Reachability Information (NLRI), EVPN NLRI), which is different than a “path” through the network which refers to a nexthop (e.g., next router) or complete path (e.g., traverse router A then router B, and so on). Also, the use of the term “prefix” without a qualifier herein refers to a fully or partially expanded prefix.

**[0025]** Embodiments described herein include various elements and limitations, with no one element or limitation contemplated as being a critical element or limitation. Each of the claims individually recites an aspect of the embodiment in its entirety. Moreover, some embodiments described may include, but are not limited to, inter alia, systems, networks, integrated circuit chips, embedded processors, ASICs, methods, and computer-readable media containing instructions. One or multiple systems, devices, components, etc., may comprise one or more embodiments, which may include some elements or limitations of a claim being performed by the same or different systems, devices, components, etc. A processing element may be a general processor, task-specific processor, a core of one or more processors, or other co-located, resource-sharing implementation for performing the corresponding processing. The embodiments described hereinafter embody various aspects and configurations, with the figures illustrating exemplary and non-limiting configurations. Computer-readable media and means for performing methods and processing block operations (e.g., a processor and memory or other apparatus configured to perform such operations) are disclosed and are in keeping with the extensible scope of the embodiments. The term “apparatus” is used consistently herein with its common definition of an appliance or device.

**[0026]** The steps, connections, and processing of signals and information illustrated in the figures, including, but not limited to, any block and flow diagrams and message sequence charts, may typically be performed in the same or in a different serial or parallel ordering and/or by different

components and/or processes, threads, etc., and/or over different connections and be combined with other functions in other embodiments, unless this disables the embodiment or a sequence is explicitly or implicitly required (e.g., for a sequence of read the value, process said read value—the value must be obtained prior to processing it, although some of the associated processing may be performed prior to, concurrently with, and/or after the read operation). Also, nothing described or referenced in this document is admitted as prior art to this application unless explicitly so stated.

**[0027]** The term “one embodiment” is used herein to reference a particular embodiment, wherein each reference to “one embodiment” may refer to a different embodiment, and the use of the term repeatedly herein in describing associated features, elements and/or limitations does not establish a cumulative set of associated features, elements and/or limitations that each and every embodiment must include, although an embodiment typically may include all these features, elements and/or limitations. In addition, the terms “first,” “second,” etc., as well as “particular” and “specific” are typically used herein to denote different units (e.g., a first widget or operation, a second widget or operation, a particular widget or operation, a specific widget or operation). The use of these terms herein does not necessarily connote an ordering such as one unit, operation or event occurring or coming before another or another characterization, but rather provides a mechanism to distinguish between elements units. Moreover, the phrases “based on x” and “in response to x” are used to indicate a minimum set of items “x” from which something is derived or caused, wherein “x” is extensible and does not necessarily describe a complete list of items on which the operation is performed, etc. Additionally, the phrase “coupled to” is used to indicate some level of direct or indirect connection between two elements or devices, with the coupling device or devices modifying or not modifying the coupled signal or communicated information. Moreover, the term “or” is used herein to identify a selection of one or more, including all, of the conjunctive items. Additionally, the transitional term “comprising,” which is synonymous with “including,” “containing,” or “characterized by,” is inclusive or open-ended and does not exclude additional, unrecited elements or method steps. Finally, the term “particular machine,” when recited in a method claim for performing steps, refers to a particular machine within the 35 USC § 101 machine statutory class.

**[0028]** Disclosed are, inter alia, methods, apparatus, computer-storage media, mechanisms, and means associated with Ethernet Virtual Private Network (EVPN) using Internet Protocol Version 6 (IPv6) Segment Routing (SRv6) underlay network and SRv6-enhanced Border Gateway Protocol (BGP) signaling.

**[0029]** In one embodiment, to support SRv6 based EVPN overlays, a SID is advertised in route-type 1, 2, 3 and 5. The SID (or SIDS), per route-type, are advertised in a new SRv6-VPN SID TLV which is optional transitive BGP Prefix SID attribute. This attribute serves two purposes; first it indicates that the BGP egress device is reachable via an SRv6 underlay and the BGP ingress device receiving this route MAY choose to encapsulate or insert an SRv6 SRH. Second, it indicates the value of the SID or SIDS to include in the SRH encapsulation. A BGP speaker supporting an SRv6 underlay MAY distribute SIDs per route via the BGP SRv6 Attribute. If the BGP speaker supports MPLS based EVPN simultaneously, then it MAY also populate the Label

values in EVPN route types and allow the BGP ingress device to decide which encapsulation to use. If the BGP speaker does not support MPLS based EVPN services the MPLS Labels in EVPN route types MUST be set to IMPLICIT-NULL.

**[0030]** One embodiment includes: advertising, in a network via a particular route advertisement message of a routing protocol, a particular route associated with a particular Internet Protocol Version 6 (IPv6) Segment Routing (SRv6) Segment Identifier (SID), with the SID including a locator of a particular router and a function encoding representing a particular Ethernet Virtual Private Network (EVPN) end function of the particular router, with the particular SID including a routable prefix to the particular router; receiving, by the particular router, a particular packet including the particular SID; and in response to said received particular packet including the particular SID, the particular router performing the particular EVPN end function on the particular packet.

**[0031]** One embodiment includes: receiving an advertisement of the particular route being associated with the particular SID by an ingress router in the network; receiving an Ethernet frame by the ingress router; encapsulating the Ethernet frame into the particular packet by the ingress router; and sending the particular packet into the network by the ingress router; wherein said performing the particular EVPN end function on the particular packet includes decapsulating the Ethernet frame from the packet and sending the Ethernet frame from the particular router.

**[0032]** In one embodiment, the particular packet includes a Segment Routing Header (SRH) including the particular SID as the currently active SID. In one embodiment, the routing protocol is Border Gateway Protocol, and the particular route advertising message advertises a BGP route type 1, 2, 3, or 5, and associated with a particular SRv6-VPN Type Length Value (TLV) including the particular SID. In one embodiment, the particular route advertising message includes advertising the particular route being associated with one or more BGP Multiprotocol Label Switching-based (MPLS-based) EVPN labels for which to use in an MPLS packet to invoke corresponding BGP MPLS-based EVPN functionality on a corresponding MPLS packet by the particular router.

**[0033]** One embodiment includes receiving, by an ingress router, an advertisement of the particular route being associated with both the particular SID and said one or more MPLS-EVPN labels; receiving an Ethernet frame by the ingress router; encapsulating the Ethernet frame into the particular packet by the ingress router; and sending the particular packet into the network by the ingress router. In one embodiment, performing the particular EVPN end function on the particular packet includes decapsulating the Ethernet frame from the packet and sending the Ethernet frame from the particular router.

**[0034]** One embodiment includes: receiving, by a second ingress router, an advertisement of the particular route being associated with both the particular SID and said one or more MPLS-EVPN labels; receiving a second Ethernet frame by the second ingress router; encapsulating the second Ethernet frame into a particular MPLS packet by the second ingress router; sending the particular MPLS packet into the network by the second ingress router; and performing said corresponding BGP MPLS-based EVPN functionality on the particular MPLS packet which includes decapsulating the

second Ethernet frame from the particular MPLS packet and sending the second Ethernet frame from the particular router.

**[0035]** Disclosed are, inter alia, methods, apparatus, computer-storage media, mechanisms, and means associated with EVPN using a SRv6 Underlay Network and SRv6-enhanced Border Gateway Protocol (BGP) signaling.

**[0036]** FIG. 1A illustrates network **100** operating according to one embodiment. As shown, network **100** includes client networks **101** and **103** (which are the same network in one embodiment) external to network **110**, which includes edge nodes (e.g., SRv6 and/or MPLS) **111** and **113** and a network **112** of nodes (e.g., routers, hosts, gateways, and service functions) some, none, or all of which may be SR-capable nodes. In response to receiving a native packet, a SR edge node **111**, **113** identifies a SR policy (e.g., list of one or more segments) through or to which to forward a SR packet encapsulating the native packet. These policies can change in response to network conditions, network programming, route advertisements, etc. SR edge nodes **111** and **113** also decapsulate native packets from SR packets and forward the native packets into network **101** and **103**.

**[0037]** FIG. 1B illustrates a process according to one embodiment. Processing begins with process block **120**. In process block **122**, SR nodes (e.g., routers) in the SR networks continuously advertise and exchange segment and routing information via one or more routing protocols and/or via one or more label distribution protocols. In one embodiment, a BGP route advertisement includes SRv6 EVPN information including one or more SIDs. In one embodiment, a BGP route advertisement includes SRv6 EVPN information (e.g., including one or more SIDs) and MPLS EVPN information. In one embodiment, a BGP route advertisement includes SRv6 EVPN information (e.g., including one or more SIDs) and an implicit null in each label field MPLS EVPN information signaling that the advertising node does not support MPLS EVPN but supports SRv6 EVPN. In process block **124**, SR edge devices (e.g., SRv6 headends) continuously update their SR policies as needed. Processing of the flow diagram of FIG. 1B is complete as indicated by process block **129**.

**[0038]** FIG. 1C illustrates a process performed by a SR edge router according to one embodiment. Processing begins with process block **130**. In process block **132**, a packet is received. In process block **134**, an ingress lookup operation is performed based on the packet to determine processing information. As determined in process block **135**, if the packet was received on an external network interface and will be sent into a SRv6 or MPLS network, then processing proceeds to process block **136**, else processing proceeds to process block **141**. As determined in process block **141**, if the packet was received on a SRv6 or MPLS network interface and will be sent to an external network, then processing proceeds to process block **142**, else processing proceeds to process block **148** wherein the packet is processed normally and processing proceeds to process block **149**.

**[0039]** Continuing in process block **136**, the received native packet is encapsulated into a SRv6 packet with a SRv6 encapsulating header with at least one SID identifying an EVPN End Function according to a corresponding policy, and processing proceeds to process block **138**.

**[0040]** Continuing in process block **142**, the native packet is removed from being encapsulated in the received SRv6 or MPLS packet; in process block **144**, the Segment Routing

EVPN End Function identified from the SID in the received SRv6 packet is performed or the EVPN function identified by an MPLS label is performed; and processing proceeds to process block 145. As determined in process block 145, if an egress lookup operation is to be performed as part of the EVPN function being applied, then processing proceeds to process block 138. Otherwise, in process block 146, the native packet is forwarded from one or more corresponding interfaces as part of the EVPN function being applied, and processing proceeds to process block 149.

[0041] Continuing with process block 138, an egress lookup operation is performed on the outbound packet to determine forwarding information, with the packet being forwarded accordingly. Processing proceeds to process block 149.

[0042] Continuing in process block 149, processing of the flow diagram of FIG. 1C is complete as indicated by process block 149.

[0043] FIG. 1D illustrates a Segment Routing (SR) packet 160 according to one embodiment. As shown, SR packet 160 (e.g. SRv6 packet) includes a SR encapsulating header 162 and native packet 169. SR encapsulating header 162 includes an IPv6 header 167 and one or more SR headers (SRHs) 168.

[0044] FIGS. 2A-B and their discussion herein provide a description of various SRv6 and/or MPLS nodes according to one embodiment.

[0045] FIG. 2A illustrates one embodiment of a packet switching device 200 (e.g., SR-capable router and/or MPLS-capable router) according to one embodiment. As shown, packet switching device 200 includes multiple line cards 201 and 205, each with one or more network interfaces for sending and receiving packets over communications links (e.g., possibly part of a link aggregation group), and with one or more processing elements that are used in one embodiment associated with EVPN using a SRv6 Underlay Network and SRv6-enhanced Border Gateway Protocol (BGP) signaling. Packet switching device 200 also has a control plane with one or more processing elements 202 for managing the control plane and/or control plane processing of packets associated with EVPN using a SRv6 Underlay Network and SRv6-enhanced Border Gateway Protocol (BGP) signaling. Packet switching device 200 also includes other cards 204 (e.g., service cards, blades) which include processing elements that are used in one embodiment to process (e.g., forward/send, drop, manipulate, change, modify, receive, create, duplicate, perform SR gateway functionality possibly with shared memory with one or more service functions, apply a service according to one or more service functions) packets associated with EVPN using a SRv6 Underlay Network and SRv6-enhanced Border Gateway Protocol (BGP) signaling, and some hardware-based communication mechanism 203 (e.g., bus, switching fabric, and/or matrix, etc.) for allowing its different entities 201, 202, 204 and 205 to communicate. Line cards 201 and 205 typically perform the actions of being both an ingress and egress line card, in regards to multiple other particular packets and/or packet streams being received by, or sent from, packet switching device 200. In one embodiment, SRv6 EVPN processing of packets is implemented by packet switching device 200 (e.g., on one or more of line cards 201, 205, on a service card 204).

[0046] FIG. 2B is a block diagram of an apparatus 220 used in one embodiment associated with EVPN using a

SRv6 Underlay Network and SRv6-enhanced Border Gateway Protocol (BGP) signaling. In one embodiment, apparatus 220 performs one or more processes, or portions thereof, corresponding to one of the flow diagrams illustrated or otherwise described herein, and/or illustrated in another diagram or otherwise described herein.

[0047] In one embodiment, apparatus 220 includes one or more processor(s) 221 (typically with on-chip memory), memory 222 (possibly shared memory), storage device(s) 223, specialized component(s) 225 (e.g. optimized hardware such as for performing lookup and/or packet processing operations and/or service function, associative memory, binary and/or ternary content-addressable memory, etc.), and interface(s) 227 for communicating information (e.g., sending and receiving packets, user-interfaces, displaying information, etc.), which are typically communicatively coupled via one or more communications mechanisms 229 (e.g., bus, links, switching fabric, matrix), with the communications paths typically tailored to meet the needs of a particular application.

[0048] Various embodiments of apparatus 220 may include more or fewer elements. The operation of apparatus 220 is typically controlled by processor(s) 221 using memory 222 and storage device(s) 223 to perform one or more tasks or processes. Memory 222 is one type of computer-readable/computer-storage medium, and typically comprises random access memory (RAM), read only memory (ROM), flash memory, integrated circuits, and/or other memory components. Memory 222 typically stores computer-executable instructions to be executed by processor(s) 221 and/or data which is manipulated by processor(s) 221 for implementing functionality in accordance with an embodiment. Storage device(s) 223 are another type of computer-readable medium, and typically comprise solid state storage media, disk drives, diskettes, networked services, tape drives, and other storage devices. Storage device(s) 223 typically store computer-executable instructions to be executed by processor(s) 221 and/or data which is manipulated by processor(s) 221 for implementing functionality in accordance with an embodiment.

[0049] One embodiment provides new extensions to BGP to allow advertising certain attributes and functionalities related to SRv6, and end functions to process SRv6 EVPN packets.

[0050] To support various SRv6-based EVPN overlay services, one embodiment advertises SIDs using any BGP route-type. To support various SRv6-based EVPN overlay services, one embodiment advertises SIDs using a BGP route-type 1, 2, 3 and/or 5.

[0051] In one embodiment, these SIDs are advertised in a new SRv6-VPN SID TLV, which is optional transitive BGP Prefix SID attribute. This attribute serves multiple purposes. The advertising of the SRv6 SID in the SRv6-VPN SID TLV identifies that the BGP egress router is reachable via a SRv6 underlay network. In one embodiment, the BGP ingress router receiving this BGP route advertisement then encapsulates the received packet into a SRv6 packet or inserts a SRv6 SRH into a received IPv6 packet. The advertising of the SRv6 SID in the SRv6-VPN SID TLV also identifies the value of the SID or SIDs to include in the SRH of the packet. In one embodiment, a BGP speaker supporting a SRv6 underlay network distributes SIDs per route via the BGP SRv6 Attribute.

**[0052]** In one embodiment, a BGP Speaker (e.g., a SRv6 egress router supporting both MPLS-based EVPN and SRv6-based EVPN) populates the MPLS label fields in the EVPN route types of a route advertisement, thus providing a BGP ingress router receiving this route advertisement two different underlay networks (e.g., MPLS and SRv6) to select among to send to the egress SRv6 router (that is both MPLS and SRv6 capable). In one embodiment, a BGP Speaker (e.g., a SRv6 egress router not supporting MPLS-based EVPN but supporting SRv6-based EVPN) populates the MPLS label fields in the EVPN route types of a route advertisement to IMPLICIT-NULL, thus providing a BGP ingress router receiving this route advertisement the SRv6 routing information and informing it that it does not support MPLS-based EVPN. In this manner, ingress and egress routers no longer have to negotiate underlay network capabilities, nor use BGP Capabilities Advertisements to ensure that they both are capable of properly processing such advertised NLRI.

**[0053]** FIG. 3A illustrates a SRv6 Type-Length-Value (TLV) **300** according to one embodiment. As shown, SRv6 TLV **300** includes a Type field **301** (e.g., eight bits), a Length field **302** (e.g., sixteen bits) indicating the length of the value field, a Reserve field **303** (e.g., eight bits), and a value field containing SRv6 SID information **310**.

**[0054]** FIG. 3B illustrates SRv6-SID information **310** according to one embodiment. As shown, SRv6-SID information **310** includes a SID Type Field **311** (e.g., one octet) and a SRv6 SID **320** (e.g., an 128 bit IPv6 routable address). In one embodiment, SID Type Field **311** is set to type-1 for a Layer-3 End Function. In one embodiment, SID Type Field **311** is set to type-2 for a Layer-2 End Function.

**[0055]** FIG. 3C illustrates SRv6-VPN encoding of a SRv6 SID **320** according to one embodiment. As shown, SRv6 SID **320** includes a Locator **321** (identifying a router to which the SID corresponds), a Function **322** (e.g., identifies processing to be performed by the router), and an optional Argument **323** for use in performing the processing. In one embodiment, the entire SRv6 **320** is used in routing a packet through the network to the BGP egress router. In one embodiment, a non-fully expanded prefix of SID **320** (e.g., locator **321**, locator **321** plus function **322**) is used in routing a packet through the network to the BGP advertising router.

**[0056]** FIG. 4A illustrates a SRv6-enhanced Border Gateway Protocol (BGP) route advertisement message **400** according to one embodiment, advertising a BGP EVPN route for both SRv6 and MPLS underlay networks. As shown, BGP route advertising message **400** includes an MPLS EVPN advertisement **401** including corresponding MPLS label(s), and a SRv6 TLV **402** including a SID identifying the SRv6 End Function.

**[0057]** FIG. 4B illustrates a SRv6-enhanced Border Gateway Protocol (BGP) route advertisement message **410** according to one embodiment, advertising a BGP EVPN route for a SRv6 underlay network but not for an MPLS underlay network. As shown, BGP route advertising message **410** includes an MPLS EVPN advertisement **411** including Implicit Null label(s) in MPLS label field(s), and a SRv6 TLV **402** including a SID identifying the SRv6 End Function (e.g., the same SRv6 TLV **402** in both BGP advertising messages **410** (of FIG. 4B) and **400** (of FIG. 4A)).

Ethernet Auto-Discovery Route Over SRv6 Core

**[0058]** Ethernet Auto-discovery (A-D) routes are Type-1 route type, defined in RFC 7432, and may be used to achieve split horizon filtering, fast convergence and aliasing. EVPN route type-1 is also used in EVPN-VPWS as well as in EVPN flexible cross-connect; mainly used to advertise point-to-point services id.

**[0059]** Multihomed PEs may advertise an Ethernet auto discovery route per Ethernet segment with the introduced ESI MPLS label extended community, defined in RFC 7432. Provider edge routers (PEs) may identify other PEs connected to the same Ethernet segment after the EVPN type-4 ES route exchange. All the multihomed PE and all remote PEs that are part of same EVI may import the auto discovery route.

**[0060]** In one embodiment, EVPN Route Type-1 is encoded as follows for SRv6 Core:

**[0061]** Route Distinguisher (RD) (e.g., eight octets),

**[0062]** Ethernet Segment Identifier (e.g., ten octets),

**[0063]** Ethernet Tag ID (e.g., eight octets), and

**[0064]** MPLS label (e.g., three octets).

**[0065]** For a SRv6 only BGP speaker for a SRv6 Core: SRv6-VPN SID TLV may be advertised with the route.

EVPN Route Type-1 (Per-ES AD)

**[0066]** Where:

**[0067]** BGP next-hop: IPv6 address of an egress PE

**[0068]** Ethernet Tag ID: all FFFF's

**[0069]** MPLS Label: always set to zero value

**[0070]** Extended Community: ESI label extended community

**[0071]** In one embodiment, a SRv6-VPN TLV is (or may be) advertised along with the route advertisement and the behavior of the SRv6-VPN SID is entirely up to the originator of the advertisement. In one embodiment, the behavior is Arg.FE2.

EVPN Route Type-1 (Per-EVI AD)

**[0072]** Where:

**[0073]** BGP next-hop: IPv6 address of an egress PE

**[0074]** Ethernet Tag ID: non-zero for VLAN aware bridging, EVPN VPWS and FXC

**[0075]** MPLS Label: Implicit-Null

**[0076]** In one embodiment, a SRv6-VPN TLV is (or may be) advertised along with the route advertisement and the behavior of the SRv6-VPN SID is entirely up to the originator of the advertisement. In one embodiment, the behavior is End.DX2, End.DXV2 or End.DT2U.

MAC/IP Advertisement Route (Type-2) with SRv6 Core

**[0077]** EVPN route type-2 is used to advertise MAC or MAC+IP address reachability for unicast traffic through MP-BGP to all other PEs in a given EVPN instance. In one embodiment, a MAC/IP Advertisement route type is encoded as follows for SRv6 Core:

**[0078]** Route Distinguisher (RD) (e.g., eight octets),

**[0079]** Ethernet Segment Identifier (e.g., ten octets),

**[0080]** Ethernet Tag ID (e.g., eight octets),

**[0081]** MAC Address Length (e.g., one octet),

**[0082]** MAC Address (e.g., six octets),

**[0083]** IP Address Length (e.g., one octet),

**[0084]** IP Address (e.g., zero, four or sixteen octets),

**[0085]** MPLS Label-1 (e.g., three octets), and

**[0086]** MPLS label-2 (e.g., zero or three octets).

[0087] where:

- [0088] BGP next-hop: IPv6 address of an egress PE,
- [0089] MPLS Label-1: Implicit-null, and
- [0090] MPLS Label-2: Implicit-null.

[0091] In one embodiment, the SRv6-VPN SID TLV is (or may be) advertised. The behavior of the SRv6-VPN SID is entirely up to the originator of the advertisement. In one embodiment, the behavior of the SRv6 SID is as follows:

- [0092] End.DX2, End.DTU2 (Layer 2 portion of the route), and
- [0093] End.DT6/4 or End.DX6/4 (Layer 3 portion of the route).

[0094] Described below are different types of Type-2 advertisements used in one embodiment. The first type is used in Layer-2 Gateway whereas the second type is used in EVPN IRB.

[0095] MAC/IP Advertisement Route (Type-2) with MAC Only:

- [0096] BGP next-hop: IPv6 address of egress PE,
- [0097] MPLS Label1: Implicit-null,
- [0098] MPLS Label2: Implicit-null, and
- [0099] SRv6-VPN SID TLV MAY encode End.DX2 or End.DTU2 behavior.

[0100] MAC/IP Advertisement Route (Type-2) with MAC+IP:

- [0101] BGP next-hop: IPv6 address of egress PE,
- [0102] MPLS Label1: Implicit-Null,
- [0103] MPLS Label2: Implicit-Null, and
- [0104] SRv6-VPN SID TLV MAY encode Layer2 End.DX2 or End.DTU2 behavior and Layer3 End.DT6/4 or End.DX6/4 behavior.

Inclusive Multicast Ethernet Tag Route with SRv6 Core

[0105] EVPN route Type-3 is used to advertise multicast traffic reachability information through MP-BGP to all other PEs in a given EVPN instance.

- [0106] Route Distinguisher (RD) (e.g., eight octets),
- [0107] Ethernet Tag ID (e.g., eight octets),
- [0108] IP Address Length (e.g., one octet), and
- [0109] Originating Router's IP Address (e.g., 4, or 16 octets).

[0110] Where:

- [0111] BGP next-hop: IPv6 address of egress PE, and
- [0112] PMSI Tunnel Attribute (per RFC 6514).

[0113] In one embodiment, a SRv6-VPN TLV is (or may be) advertised along with the route advertisement and the behavior of the SRv6-VPN SID is entirely up to the originator of the advertisement. In one embodiment, the behavior is End.DT2M.

[0114] In one embodiment, the PMSI Tunnel Attribute (see, RFC 6514) is (or may) contain MPLS implicit-null label in the case of Tunnel Type set to ingress-replication.

- [0115] Flag (e.g., 1 octet),
- [0116] Tunnel Type (e.g., 1 octet),
- [0117] MPLS label (e.g., three octets), and
- [0118] Tunnel Identifier (e.g., variable).

[0119] Where:

- [0120] Tunnel Type (e.g., defined per RFC 6514),
- [0121] MPLS label: Implicit-Null, and
- [0122] Tunnel Identifier: IP address of egress PE.

[0123] In one embodiment, an Arg.FE2 argument SID is (or may) get applied to an End.DT2M function SID to create a single SID on the ingress PE. The result End.DT2M(FE2) SID is used as EVPN Split-horizon mechanism (e.g., as described in RFC 7432).

Ethernet Segment Route with SRv6 Core

[0124] In one embodiment, EVPN route type-4 is used to discover PE participating to the same redundancy group. It is also used in Designated Forwarder (DF) election procedure. In one embodiment, an Ethernet Segment route type is encoded as follows for SRv6 Core:

- [0125] Route Distinguisher (RD) (e.g., eight octets),
- [0126] Ethernet Tag ID (e.g., eight octets),
- [0127] IP Address Length (e.g., one octet), and
- [0128] Originating Router's IP Address (e.g., 4, or 16 octets).

[0129] Where:

- [0130] BGP next-hop: IPv6 address of egress PE.

[0131] In one embodiment, SRv6-VPN TLV is advertised along with the route. In one embodiment, SRv6-VPN TLV is not advertised along with the route. The processing of that route has not changed; it remains as described in RFC 7432. IP Prefix Router (Type-5) with SRv6 Core

[0132] In one embodiment, EVPN route Type-5 is used to advertise IP address reachability through MP-BGP to all other PEs in a given EVPN instance. The IP address may include host IP prefix or any specific subnet. EVPN route Type-5 is defined in ietf-bess-evpn-prefix-advertisement-05. In one embodiment, the encoding is as follows:

- [0133] Route Distinguisher (RD) (e.g., eight octets),
- [0134] Ethernet Segment Identifier (e.g., ten octets),
- [0135] Ethernet Tag ID (e.g., eight octets),
- [0136] IP Prefix Length (e.g., one octet),
- [0137] IP Prefix (e.g., six octets),
- [0138] GW Address (e.g., four or sixteen octets), and
- [0139] MPLS Label (e.g., three octets),

[0140] where:

- [0141] BGP next-hop: IPv6 address of an egress PE, and
- [0142] MPLS Label: Implicit-null.

[0143] In one embodiment, SRv6-VPN SID TLV is (or may be) advertised. The behavior of the SRv6-VPN SID is entirely up to the originator of the advertisement. In one embodiment, the behavior of the SRv6 SID is an End.DT6/4 or End.DX6/4.

Multicast Related Routes (EVPN Route Type-6, Type-7, Type-8)

[0144] In one embodiment, these routes do not require any additional SRv6-VPN TLV. In one embodiment, as per EVPN route-type 4, the BGP next-hop is equal to the IPv6 address of egress PE.

SRv6 END Functions Associated with a Local SID

[0145] The following describes end functions performed by an egress router in response to receiving a SRv6 packet including an advertised SID. The end function is typically determined by the function portion of the SID.

End.DX2: Endpoint with Decapsulation and Layer-2 Cross-Connect

[0146] One embodiment uses the "Endpoint with decapsulation and Layer-2 cross-connect to OIF" function (End.DX2 for short).

[0147] When N receives a packet destined to S and S is a local End.DX2 SID, N does:

---

IF NH=SRH and SL > 0  
drop the packet

:: Ref1

-continued

---

```

ELSE IF ENH = 59                ;; Ref2
  pop the (outer)IPv6 header and its extension headers
  forward the resulting frame via OIF associated to the local SID
ELSE
  drop the packet

```

---

**[0148]** Ref1: An End.DX2 SID must always be the last SID, or it can be the Destination Address of an IPv6 packet with no SRH header.

**[0149]** Ref2: The next-header value 59 is allocated to IPv6 No Next Header per RFC 2460. When the SID is of function End.DX2 and the Next-Header is 59, an Ethernet frame is in the payload without any further header.

**[0150]** In one embodiment, an End.DX2 function is customized to expect a specific VLAN format and rewrite the egress VLAN header before forwarding on the outgoing interface. In one embodiment, the End.DX2 is used for L2VPN point-to-point use-cases.

End.DXV2: Endpoint with Decapsulation and Specific VLAN L2 Table Lookup

**[0151]** One embodiment uses the “Endpoint with decapsulation and specific VLAN L2 table lookup” function (End.DXV2 for short).

**[0152]** When N receives a packet destined to S and S is a local End.DXV2 SID, N does:

---

```

IF NH=SRH and SL > 0
  drop the packet                ;; Ref1
ELSE IF ENH = 59                ;; Ref2
  pop the (outer)IPv6 header and its extension headers
  lookup the exposed inner VLANs in L2 table T
  forward via the matched table entry
ELSE
  drop the packet

```

---

**[0153]** Ref1: An End.DXV2 SID must always be the last SID, or it can be the Destination Address of an IPv6 packet with no SRH header.

**[0154]** Ref2: The next-header value 59 is allocated to IPv6 No Next Header per RFC 2460. When the SID is of function End.DXV2 and the Next-Header is 59, an Ethernet frame is in the payload without any further header.

**[0155]** In one embodiment, an End.DXV2 function is customized to expect a specific VLAN format and rewrite the egress VLAN header before forwarding on the outgoing interface. In one embodiment, the End.DXV2 is used for EVPN Flexible cross-connect use-cases. The table lookup support dot1q, dot1ad with single and double tags.

End.DT2U: Endpoint with Decapsulation and Specific Unicast MAC L2 Table Lookup

**[0156]** One embodiment uses the “Endpoint with decapsulation and specific unicast MAC L2 table lookup” function (End.DT2U for short).

**[0157]** When N receives a packet destined to S and S is a local End.DT2U SID, N does:

---

```

IF NH=SRH and SL > 0
  drop the packet                ;; Ref1
ELSE IF ENH = 59                ;; Ref2
  pop the (outer)IPv6 header and its extension headers
  learn the exposed inner MAC SA in L2 table T ;; Ref3
  lookup the exposed inner MAC DA in L2 table T

```

---

-continued

---

```

forward via the matched table entry
ELSE
  drop the packet

```

---

**[0158]** Ref1: An End. DT2U SID must always be the last SID, or it can be the Destination Address of an IPv6 packet with no SRH header.

**[0159]** Ref2: The next-header value 59 is allocated to IPv6 No Next Header per RFC 2460. When the SID is of function End. DT2U and the Next-Header is 59, an Ethernet frame is in the payload without any further header.

**[0160]** Ref3: In EVPN, the learning of the exposed inner MAC SA is done via control plane.

**[0161]** In one embodiment, the End.DT2U is used for EVPN Bridging unicast use cases.

End.DT2M: Endpoint with Decapsulation and Specific L2 Table Flooding

**[0162]** One embodiment uses the “Endpoint with decapsulation and specific L2 table flooding” function (End. DT2M for short).

**[0163]** When N receives a packet destined to S and S is a local End.DT2M SID, N does:

---

```

IF NH=SRH and SL > 0
  drop the packet                ;; Ref1
ELSE IF ENH = 59                ;; Ref2
  pop the (outer)IPv6 header and its extension headers
  learn the exposed inner MAC SA in L2 table T ;; Ref3
  forward based on all OIF in L2 table T
ELSE
  drop the packet

```

---

**[0164]** Ref1: An End. DT2M SID must always be the last SID, or it can be the Destination Address of an IPv6 packet with no SRH header.

**[0165]** Ref2: The next-header value 59 is allocated to IPv6 No Next Header per RFC 2460. When the SID is of function End. DT2U and the Next-Header is 59, an Ethernet frame is in the payload without any further header.

**[0166]** Ref3: In EVPN, the learning of the exposed inner MAC SA is done via control plane.

**[0167]** The End.DT2M is used for EVPN Bridging BUM use cases.

Arg.FE2: L2 Argument Specific to EVPN ESI Filtering

**[0168]** One embodiment uses the “L2 argument specific to EVPN ESI filtering” (Arg.FE2 for short) which is an argument of End.DT2M endpoint function. The Arg.FE2 argument is used to exclude a specific OIF from L2 table T flooding. The Arg.FE2 SID is merged with an End.DT2M function by bit ORing operation to form an End.DT2M (FE2) single SID. In one embodiment, the End.DT2M is used for EVPN Bridging with ESI filtering use cases.

End.DT2M(FE2): Endpoint with Decapsulation and Specific L2 Table Flooding with Argument

**[0169]** One embodiment uses the “Endpoint with decapsulation and specific L2 table flooding with argument” function (End.DT2M(FE2) for short).

**[0170]** When N receives a packet destined to S and S is a local End.DT2M SID, N does:

---

```

IF NH=SRH and SL > 0
  drop the packet                ;; Ref1
ELSE IF ENH = 59                 ;; Ref2
  pop the (outer)IPv6 header and its extension headers
  learn the exposed inner MAC SA in L2 table T ;; Ref3
  forward based on all OIF excluding the one specified by
  ARGUMENT in L2 table T
ELSE
  drop the packet

```

---

**[0171]** Ref1: An End. DT2M SID must always be the last SID, or it can be the Destination Address of an IPv6 packet with no SRH header.

**[0172]** Ref2: The next-header value 59 is allocated to IPv6 No Next Header per RFC 2460. When the SID is of function End. DT2U and the Next-Header is 59, an Ethernet frame is in the payload without any further header.

**[0173]** Ref3: In EVPN, the learning of the exposed inner MAC SA is done via control plane.

**[0174]** The End.DT2M(FE2) is used for EVPN Bridging with ESI filtering use cases.

End.DX6: Endpoint with Decapsulation and IPv6 Cross-Connect

**[0175]** One embodiment uses the “Endpoint with decapsulation and cross-connect to an array of IPv6 adjacencies” function (End.DX6 for short).

**[0176]** When N receives a packet destined to S and S is a local End.DX6 SID, N does:

---

```

IF NH=SRH and SL > 0
  drop the packet                ;; Ref1
ELSE IF ENH = 41                ;; Ref2
  pop the (outer)IPv6 header and its extension headers
  forward to layer-3 adjacency bound to the SID S ;; Ref3
ELSE
  drop the packet

```

---

**[0177]** Ref1: An End. DX6 SID must always be the last SID, or it can be the Destination Address of an IPv6 packet with no SRH header.

**[0178]** Ref2: 41 refers to IPv6 encapsulation as defined by IANA allocation for Internet Protocol Numbers.

**[0179]** Ref3: Selected based on a hash of the packet’s header (at least SA, DA, Flow Label in one embodiment).

**[0180]** One of the applications of the End.DX6 function of one embodiment is the L3VPN use-case where a FIB lookup in a specific tenant table at the egress PE is not required.

End.DX4: Endpoint with Decapsulation and IPv4 Cross-Connect

**[0181]** One embodiment uses the “Endpoint with decapsulation and cross-connect to an array of IPv4 adjacencies” function (End.DX4 for short).

**[0182]** When N receives a packet destined to S and S is a local End.DX4 SID, N does:

---

```

IF NH=SRH and SL > 0
  drop the packet                ;; Ref1
ELSE IF ENH = 4                 ;; Ref2
  pop the (outer)IPv6 header and its extension headers
  forward to layer-3 adjacency bound to the SID S ;; Ref3
ELSE
  drop the packet

```

---

**[0183]** Ref1: An End. DX4 SID must always be the last SID, or it can be the Destination Address of an IPv6 packet with no SRH header.

**[0184]** Ref2: 4 refers to IPv4 encapsulation as defined by IANA allocation for Internet Protocol Numbers.

**[0185]** Ref3: Selected based on a hash of the packet’s header (at least SA, DA, Flow Label in one embodiment).

**[0186]** One of the applications of the End.DX4 function of one embodiment is the L3VPN use-case where a FIB lookup in a specific tenant table at the egress PE is not required.

End.DT6: Endpoint with Decapsulation and Specific IPv6 Table Lookup

**[0187]** One embodiment uses the “Endpoint with decapsulation and specific IPv6 table lookup” function (End.DT6 for short).

**[0188]** When N receives a packet destined to S and S is a local End.DT6 SID, N does:

---

```

IF NH=SRH and SL > 0
  drop the packet                ;; Ref1
ELSE IF ENH = 41                ;; Ref2
  pop the (outer)IPv6 header and its extension headers
  lookup the exposed inner IPv6 DA in IPv6 table T
  forward via the matched table entry
ELSE
  drop the packet

```

---

**[0189]** Ref1: An End. DT6 SID must always be the last SID, or it can be the Destination Address of an IPv6 packet with no SRH header.

**[0190]** Ref2: 41 refers to IPv6 encapsulation as defined by IANA allocation for Internet Protocol Numbers.

**[0191]** One of the applications of the End.DT6 function of one embodiment is the L3VPN use-case where a FIB lookup in a specific tenant table at the egress PE is required. Note that an End.DT6 may be defined for the main IPv6 table in which case and End.DT6 supports the equivalent of an IPv6inIPv6 decaps (without VPN/tenant implication).

End.DT4: Endpoint with Decapsulation and Specific IPv4 Table Lookup

**[0192]** One embodiment uses the “Endpoint with decapsulation and specific IPv4 table lookup” function (End.DT4 for short).

**[0193]** When N receives a packet destined to S and S is a local End.DT4 SID, N does:

---

```

IF NH=SRH and SL > 0
  drop the packet                ;; Ref1
ELSE IF ENH = 4                 ;; Ref2
  pop the (outer)IPv6 header and its extension headers
  lookup the exposed inner IPv4 DA in IPv4 table T
  forward via the matched table entry
ELSE
  drop the packet

```

---

**[0194]** Ref1: An End. DT4 SID must always be the last SID, or it can be the Destination Address of an IPv6 packet with no SRH header.

**[0195]** Ref2: 4 refers to IPv4 encapsulation as defined by IANA allocation for Internet Protocol Numbers.

**[0196]** One of the applications of the End.DT4 function of one embodiment is the L3VPN use-case where a FIB lookup in a specific tenant table at the egress PE is required. Note that an End.DT4 may be defined for the main IPv4 table in

which case and End.DT4 supports the equivalent of an IPv4inIPv6 decaps (without VPN/tenant implication).

T.Encaps.L2: Transit with Encapsulation of L2 Frames

**[0197]** In one embodiment, T.Encaps.L2 encapsulates the received L2 frame (i.e. the received Ethernet header and its optional VLAN header is in the payload of the outer packet). If the outer header is pushed without SRH then the DA must be a SID of type End.DX2, End.DXV2, End.DT2U or End.DT2M and the next-header must be 59 (IPv6 NoNextHeader). The received Ethernet frame follows the IPv6 header and its extension headers. Else, if the outer header is pushed with a SRH, then the last SID of the SRH must be of type End.DX2, End.DXV2, End.DT2U or End.DT2M and the next-header of the SRH must be 59 (IPv6 NoNextHeader). The received Ethernet frame follows the IPv6 header and its extension headers.

**[0198]** In one embodiment, T.Encaps.L2 encapsulates the received L2 frame as follows.

---

IF outer header is pushed without SRH  
 IPv6 DA must be a SID of type Layer-2 ;; Ref2  
 IPv6 Next-header equals 59 (IPv6 NoNextHeader)  
 received L2 frame follows IPv6 header  
 ELSE (i.e. outer header is pushed with IPv6+SRH)  
 The 0th SID of the SRH must be a SID of type Layer-2 ;; Ref2  
 SRH Next-Header equals 59 (IPv6 NoNextHeader)  
 received L2 frame follows IPv6 header and its extension headers

---

**[0199]** Ref1: received L2 frame=received Ethernet header and its optional VLAN header.

**[0200]** Ref2: 4 Layer-2=End.DX2, End.DX2V, End.DT2U or End.DT2M.

#### Error Handling of BGP SRv6 SID Updates

**[0201]** In one embodiment, when a BGP Speaker receives a BGP Update message containing a malformed SRv6-VPN SID TLV, it ignores the received BGP attributes and does not pass it to other BGP peers. This is equivalent to the—attribute discard—action specified in RFC 7606. In one embodiment, when discarding an attribute, a BGP speaker logs or does not log an error for further analysis.

#### CONCLUDING PARAGRAPHS

**[0202]** In summary, in one embodiment, Ethernet Virtual Private Network (EVPN) is implemented using Internet Protocol Version 6 (IPv6) Segment Routing (SRv6) underlay network and SRv6-enhanced Border Gateway Protocol (BGP) signaling. A particular route associated with a particular Internet Protocol Version 6 (IPv6) Segment Routing (SRv6) Segment Identifier (SID) is advertised in a particular route advertisement message of a routing protocol (e.g., BGP). The SID includes a locator of a particular router and a function encoding representing a particular EVPN end function of the particular router, with the particular SID including a routable prefix to the particular router. The particular router receives a particular packet including the particular SID; and in response, the particular router performs the particular EVPN end function on the particular packet. In one embodiment, the particular packet includes a Segment Routing Header (SRH) including the particular SID as the currently active SID.

**[0203]** In view of the many possible embodiments to which the principles of the disclosure may be applied, it will be appreciated that the embodiments and aspects thereof

described herein with respect to the drawings/figures are only illustrative and should not be taken as limiting the scope of the disclosure. For example, and as would be apparent to one skilled in the art, many of the process block operations can be re-ordered to be performed before, after, or substantially concurrent with other operations. Also, many different forms of data structures could be used in various embodiments. The disclosure as described herein contemplates all such embodiments as may come within the scope of the following claims and equivalents thereof.

What is claimed is:

1. An apparatus comprising:

network interfaces configured to send and receive packets in a network;

one or more memory devices; and

one or more processors operatively coupled to the one or more network interfaces and the one or more memory devices;

wherein the one or more processors are configured to cause the apparatus to perform processing operations, the processing operations comprising:

maintaining in the one or more memory devices, forwarding information, the forwarding information associated with an Internet Protocol Version 6 (IPv6) Segment Routing (SRv6) Segment Identifier (SID), wherein the SID comprises an encoding of an Ethernet Virtual Private Network (EVPN) Layer 2 (L2) flooding Segment Routing end function and an encoding of an Ethernet Segment Identifier (ESI);

receiving, at a first network interface of the network interfaces, an IPv6 packet, the IPv6 packet comprising an SRv6 header, the SRv6 header comprising a destination address;

responsive to the destination address comprising the SID, performing behavior of the EVPN L2 flooding segment routing end function,

wherein the behavior of the EVPN L2 flooding segment routing end function comprises:

decapsulating an Ethernet frame from the received IPv6 packet by removing the SRv6 header; and

forwarding the Ethernet frame from each of a set of output network interfaces of the network interfaces identified in a L2 table of the forwarding information,

wherein the forwarding excludes one or more of the network interfaces associated with the ESI;

wherein the Ethernet frame is sent from at least one of the network interfaces.

2. The apparatus of claim 1 wherein the processing operations further comprise advertising the forwarding information in a route advertisement message of a routing protocol.

3. The apparatus of claim 2 wherein the routing protocol is Border Gateway Protocol (BGP).

4. The apparatus of claim 3, wherein the route advertisement message advertises a BGP route type 1, 2, 3, or 5, and associated with a particular SRv6-VPN Type Length Value (TLV) including the SID.

5. The apparatus of claim 4, wherein the route advertisement message identifies that SRv6 EVPN is supported as identified by the SRv6-VPN TLV.

6. The apparatus of claim 1, wherein the received IPv6 Packet comprises a Segment Routing Header (SRH) including the SID as the currently active SID.



7. The apparatus of claim 1, wherein the SID is a 128-bit IPv6 address.

8. A method comprising:

maintaining in one or more of memory devices, forwarding information, the forwarding information associated with an Internet Protocol Version 6 (IPv6) Segment Routing (SRv6) Segment Identifier (SID), wherein the SID comprises an encoding of an Ethernet Virtual Private Network (EVPN) Layer 2 (L2) flooding Segment Routing end function and an encoding of an Ethernet Segment Identifier (ESI);

receiving, at a first network interface of network interfaces, an IPv6 packet, the IPv6 packet comprising an SRv6 header, the SRv6 header comprising a destination address;

responsive to the destination address comprising the SID, performing behavior of the EVPN L2 flooding Segment Routing end function, wherein the behavior of the EVPN L2 flooding Segment Routing end function comprises:

decapsulating an Ethernet frame from the received IPv6 packet by removing the SRv6 header, and

forwarding the Ethernet frame from each of a set of output network interfaces of the network interfaces identified in a L2 table of the forwarding information,

wherein the forwarding excludes one or more of the network interfaces associated with the ESI;

wherein the Ethernet frame is sent from at least one of the network interfaces.

9. The method of claim 8 further comprising advertising the forwarding information in a received route advertisement message of a routing protocol.

10. The method of claim 9 wherein the routing protocol is Border Gateway Protocol (BGP).

11. The method of claim 10, wherein the route advertisement message advertises a BGP route type 1, 2, 3, or 5, and associated with a particular SRv6-VPN Type Length Value (TLV) including the SID.

12. The method of claim 11, wherein the route advertisement message identifies that SRv6 EVPN is supported as identified by the SRv6-VPN TLV.

13. The method of claim 8, wherein the received IPv6 Packet comprises a Segment Routing Header (SRH) including the SID as the currently active SID.

14. The method of claim 8, wherein the SID is a 128-bit IPv6 address.

15. The method of claim 8, further comprising:

receiving, by an ingress router, an route advertisement message;

receiving, by the ingress router, the Ethernet frame;

encapsulating the Ethernet frame into the IPv6 packet by the ingress router; and

sending the IPv6 packet into a network by the ingress router.

\* \* \* \* \*