



(11) **EP 3 582 449 B9**

(12) **CORRECTED EUROPEAN PATENT SPECIFICATION**

(15) Correction information:
Corrected version no 1 (W1 B1)
Corrections, see
Claims EN 1

(48) Corrigendum issued on:
02.03.2022 Bulletin 2022/09

(45) Date of publication and mention
of the grant of the patent:
05.01.2022 Bulletin 2022/01

(21) Application number: **17900752.1**

(22) Date of filing: **17.11.2017**

(51) International Patent Classification (IPC):
H04L 12/46^(2006.01)

(52) Cooperative Patent Classification (CPC):
H04L 12/4633; H04L 12/4641

(86) International application number:
PCT/CN2017/111608

(87) International publication number:
WO 2018/166233 (20.09.2018 Gazette 2018/38)

(54) **ROUTE PROCESSING IN A VXLAN**

ROUTENVERARBEITUNG IN EINEM VXLAN

TRAITEMENT DES ROUTES DANS UN VXLAN

(84) Designated Contracting States:
**AL AT BE BG CH CY CZ DE DK EE ES FI FR GB
GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO
PL PT RO RS SE SI SK SM TR**

(30) Priority: **14.03.2017 CN 201710151473**

(43) Date of publication of application:
18.12.2019 Bulletin 2019/51

(60) Divisional application:
21209223.3

(73) Proprietor: **Huawei Technologies Co., Ltd.**
Longgang District
Shenzhen, Guangdong 518129 (CN)

(72) Inventors:
• **WANG, Haibo**
Shenzhen
Guangdong 518129 (CN)
• **CHEN, Dapeng**
Shenzhen
Guangdong 518129 (CN)
• **ZHUANG, Shunwan**
Shenzhen
Guangdong 518129 (CN)

(74) Representative: **Epping - Hermann - Fischer**
Patentanwaltsgesellschaft mbH
Schloßschmidstraße 5
80639 München (DE)

(56) References cited:
CN-A- 104 702 476 CN-A- 105 591 868
CN-A- 105 991 432 US-A1- 2015 003 458
US-A1- 2016 134 513 US-A1- 2016 248 601
US-A1- 2017 019 331

• **SAJASSI (EDITOR) CISCO J DRAKE (EDITOR)**
JUNIPER N BITAR NOKIA R SHEKHAR JUNIPER
J UTTARO AT&T W HENDERICKX NOKIA A: "A
Network Virtualization Overlay Solution using
EVPN; draft-ietf-bess-evpn-overlay-07.txt", A
NETWORK VIRTUALIZATION OVERLAY
SOLUTION USING EVPN;
DRAFT-IETF-BESS-EVPN-OVERLAY-07.TXT,
INTERNET ENGINEERING TASK FORCE, IETF;
STANDARDWORKINGDRAFT, INTERNET
SOCIETY (ISOC) 4, RUE DES FALAISES CH- 1205
GENEVA, SWITZERLAND, 30 November 2016
(2016-11-30), pages 1-28, XP015116866,
[retrieved on 2016-11-30]

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 3 582 449 B9

Description

TECHNICAL FIELD

[0001] This application relates to the field of communications technologies, and in particular, to a route processing method, a device, and a system. More specifically, this application relates to a Virtual Extensible Local Area Network (VXLAN) technology that has an Ethernet virtual private network (EVPN) control plane.

BACKGROUND

[0002] A VXLAN is a technology in which a layer 2 packet is encapsulated by using a layer 3 protocol. The VXLAN technology relates to a packet in a MAC-in-UDP format. Specifically, an Ethernet frame based on a Media Access Control (MAC) protocol is encapsulated in a User Datagram Protocol (UDP) packet. Further, the UDP packet is encapsulated in an Internet Protocol (IP) packet, and the IP packet may be transmitted in a layer 3 network. Therefore, the Ethernet frame is transmitted in the layer 3 network. In the VXLAN technology, a VXLAN network identifier (VNI) is used to identify a VXLAN segment. Different VXLAN segments are corresponding to different VNIs. Different VXLAN segments are isolated from each other. Two virtual machines (VM) in a same VNI do not need to use a VXLAN layer 3 gateway (VXLAN L3 Gateway) during communication with each other. Two VMs separately in different VNIs need to communicate with each other by using the VXLAN layer 3 gateway. A VNI field includes 24 bits. One management domain may include a maximum of 2^{16} VXLAN segments. A virtual extensible local area network tunnel endpoint (VXLAN Tunnel Endpoint, VTEP) may be integrated into a network virtualization edge (NVE) device, and used as an edge device in a VXLAN. The NVE device transmits traffic of the VXLAN through a VXLAN tunnel. The VXLAN tunnel is a logical point-to-point tunnel between two NVE devices.

[0003] An EVPN is a layer 2 virtual private network (VPN) technology. The EVPN connects customer sites in different regions by using an IP/Multiprotocol Label Switching (MPLS) bearer network. This is, these customer sites are in a same local area network (LAN).

[0004] In an application scenario, the VXLAN does not have a control plane. The VXLAN performs VTEP discovery and remote host-address learning by using a data plane through multicast-based flooding. As a result, a larger amount of flooding traffic exists in a data center network. The EVPN based on the Border Gateway Protocol (BGP) and MPLS may be used to implement the control plane of the VXLAN. In this specification, the VXLAN scenario with an EVPN control plane is named as an EVPN VXLAN.

[0005] In an EVPN VXLAN application scenario, two NVE devices are directly connected to each other by using a physical link. For example, two NVE devices are

connected to each other by using a peer link. However, an EVPN application scenario is limited by the manner in which two NVE devices are directly connected to each other by using a physical link, and there is a problem of low reliability and complex deployment.

[0006] US2016/134513A1 relates generally to multi-destination data forwarding, specifically in a joint TRILL Fabric and VXLAN/IP Fabric data centers.

[0007] The document "A Network Virtualization Overlay Solution using EVPN (draft-ietf-bess-evpn-overlay-07)" describes how Ethernet VPN (EVPN) [RFC7432] can be used as an Network Virtualization Overlay (NVO) solution and explores the various tunnel encapsulation options over IP and their impact on the EVPN control-plane and procedures. In particular, the following encapsulation options are analyzed: VXLAN, NVGRE, and MPLS over GRE.

SUMMARY

[0008] The present invention is defined in the appended independent claims to which reference should be made. Advantageous features are set out in the appended dependent claims.

BRIEF DESCRIPTION OF DRAWINGS

[0009]

FIG. 1 is a schematic structural diagram of an EVPN VXLAN according to an embodiment of this application;

FIG. 2A and FIG. 2B are a flowchart of a route processing method according to an embodiment of this application;

FIG. 3 is a schematic structural diagram of another EVPN VXLAN according to an embodiment of this application;

FIG. 4 is a schematic structural diagram of still another EVPN VXLAN according to an embodiment of this application;

FIG. 5 is a schematic diagram of a format of a VXLAN packet according to an embodiment of this application;

FIG. 6 is a schematic diagram of a format of a VXLAN attribute according to an embodiment of this application;

FIG. 7 is a schematic diagram of a format of a VTEP VNI attribute according to an embodiment of this application;

FIG. 8 is a schematic diagram of a format of a MAC-VLAN attribute according to an embodiment of this application;

FIG. 9 is a schematic structural diagram of still another EVPN VXLAN according to an embodiment of this application;

FIG. 10 is a schematic structural diagram of still another EVPN VXLAN according to an embodiment of

this application;

FIG. 11 is a schematic structural diagram of a first NVE device according to an embodiment of this application;

FIG. 12 is a schematic diagram of a hardware structure of a first NVE device according to an embodiment of this application;

FIG. 13 is a schematic diagram of a hardware structure of another first NVE device according to an embodiment of this application; and

FIG. 14 is a schematic diagram of a hardware structure of still another first NVE device according to an embodiment of this application.

DESCRIPTION OF EMBODIMENTS

[0010] Embodiments of this application provide a route processing method, a device, and a system, to be applied to an EVPN VXLAN application scenario. An EVPN application scenario is not limited by a restriction condition that a physical direct link needs to be used as a link between NVE devices. This helps extend the EVPN application scenario and also helps improve reliability and reduce deployment complexity.

[0011] The following provides detailed descriptions by using specific embodiments separately.

[0012] For an EVPN technology in this application, refer to descriptions of the Internet Engineering Task Force (IETF) Request for Comments (RFC) 7432.

[0013] In the embodiments of this application, the expression "a device includes a VTEP address" or the expression "a VTEP address in a device" means that a VTEP address is stored in a device. The expression "a device includes a common VTEP address" or the expression "a common VTEP address in a device" means that a common VTEP address is stored in a device. The device may be an NVE device, a provider edge (PE) device, or a VTEP.

[0014] FIG. 1 is a schematic structural diagram of an EVPN VXLAN according to an embodiment of this application. The EVPN VXLAN is a VXLAN having an EVPN control plane. The EVPN VXLAN includes an NVE device set, and the NVE device set includes at least two NVE devices. As shown in FIG. 1, the NVE device set includes a first NVE device, a second NVE device, and a third NVE device. It should be understood that in this application, a quantity of NVE devices included in the NVE device set is not limited, and there may be two, three or more NVE devices. Any two NVE devices in the NVE device set are a pair of BGP peers. In an EVPN scenario, the BGP peer may also be referred to as an EVPN peer. FIG. 1 is used as an example. The first NVE device and the second NVE device are a pair of BGP peers, the first NVE device and the third NVE device are a pair of BGP peers, and the second NVE device and the third NVE device are a pair of BGP peers. "A pair of BGP peers" may be understood as that one device is a BGP peer of the other device. For example, that the first NVE device

and the second NVE device are a pair of BGP peers may be understood as that the first NVE device is a BGP peer of the second NVE device, or understood as that the second NVE device is a BGP peer of the first NVE device.

5 The BGP peer may also be referred to as a BGP neighbor. Correspondingly, the EVPN peer may also be referred to as an EVPN neighbor. In this application, for ease of description, the BGP peer is used in all subsequent embodiments. The BGP peer is established by using an open message specified in the BGP, and the established BGP peer is maintained by using a keepalive message. For implementation of the OPEN message and the KEEPALIVE message, refer to related descriptions of the IETF RFC 2858 and the IETF RFC 1771. In addition, 10 a route reflector (RR) may be deployed on each of two end devices that establish the BGP peer, so that establishment of the BGP peer is completed by using the RRs. An NVE device included in an NVE device set may also be referred to as a PE device. Unless otherwise specified in this application, the NVE device is equivalent to the PE device. In descriptions of this application, a device included in an NVE device set is named as an NVE device. Actually, the NVE device included in the NVE device set may also be referred to as a PE device. Correspondingly, the NVE device set may also be referred to as a PE device set.

[0015] The EVPN VXLAN further includes a PE device. As shown in FIG. 1, the PE device is outside the NVE device set. The PE device and each NVE device in the NVE device set may form a pair of BGP peers. FIG. 1 is used as an example. The PE device is a BGP peer of each of the first NVE device, the second NVE device, and the third NVE device. The PE device may be a physical device, or may be a logical device including a plurality of physical devices. In addition, the PE device may be connected to another network device. For example, the PE device is connected to one or more customer edge (CE) devices.

[0016] The EVPN VXLAN further includes at least one CE device. FIG. 1 shows three CE devices as an example: a first CE device, a second CE device, and a third CE device. The first CE device is multi-homed to the first NVE device, the second NVE device, and the third NVE device, that is, the first CE device communicates with the first NVE device, the second NVE device, and the third NVE device separately. The second CE device is dual-homed to the first NVE device and the second NVE device, that is, the second CE device communicates with the first NVE device and the second NVE device separately. The third CE device is single-homed to the second NVE device, that is, the third CE device communicates with the second NVE device. In addition, it should be understood that multi-homing may be dual-homing. As shown in the foregoing specification, the two expressions coexist to distinguish between quantities of NVE devices connected to the CE devices.

[0017] In the EVPN VXLAN, the CE device and the NVE device are connected to each other by using an

Ethernet link. In addition, all Ethernet links connected to a same CE device form an Ethernet segment (ES). FIG. 1 is used as an example. The first CE device is multi-homed to the first NVE device, the second NVE device, and the third NVE device by using three Ethernet links. The three Ethernet links form an ES, and the ES is represented by an ES 1 in FIG. 1. An Ethernet segment identifier (ESI) is used to identify a corresponding ES. FIG. 1 is used as an example. ESI values of ESs 1 of the three Ethernet links that connect the first CE device to the first NVE device, the second NVE device, and the third NVE device are a same value. ESI values of ESs 2 of two Ethernet links that connect the second CE device to the first NVE device and the second NVE device are a same value. The ESI value of the ES 1 and the ESI value of the ES 2 are unequal, and both are non-zero values. An ESI value of an ES 3 is zero. The ESI includes a type field and an ESI value field, and the type field is used to indicate a generation manner of the ESI. Two commonly used generation manners are a type 0 and a type 1. The type 0 indicates that the ESI is generated through manual configuration, and the type 1 indicates that the ESI is generated by using the Link Aggregation Control Protocol (LACP) running between an NVE device and a CE device. A value of the ESI value field ranges from 0 to 0xFF, and "0x" represents hexadecimal. For generation and setting of the ES and the ESI, refer to descriptions of Chapter 5 of the RFC 7432.

[0018] In the EVPN VXLAN shown in FIG. 1, the NVE device and the PE device each may be a router or a layer 3 switch. The CE device may be a router, a switch, or a host. The NVE device, the PE device, and the CE device in this embodiment of this application are corresponding devices defined in the RFC 7432. When the CE device is a router or a switch, the CE device may be connected to one or more hosts. The host may be a physical device or a VM. The EVPN VXLAN shown in FIG. 1 may be applied to a plurality of scenarios. For example, the EVPN VXLAN is applied to a mobile bearer network, and a typical mobile bearer network is an Internet Protocol radio access network (IP RAN). In the mobile bearer network, the first CE device, the second CE device, and the third CE device may be base transceiver stations (BTS); the PE device may be connected to a base station controller (BSC) or a radio network controller (RNC); the first NVE device, the second NVE device, and the third NVE device may be cell site gateways (CSG); and the PE device may be a radio network controller site gateway (RSG). For another example, the EVPN VXLAN is applied to a fixed network. In the fixed network, the first CE device, the second CE device, and the third CE device may be user-side stations; the first NVE device, the second NVE device, and the third NVE device may be digital subscriber line access multiplexers (DSLAM); and the PE device may be a broadband access server (BAS).

[0019] In the EVPN VXLAN shown in FIG. 1, a VXLAN tunnel is established between the PE device and the NVE device set. For the PE device, the NVE device set is like

a single NVE device, that is, all the NVE devices in the NVE device set are virtualized as a single NVE device. In a possible implementation scenario, actual traffic forwarded from the PE device to the first NVE device, the second NVE device, and the third NVE device is carried on a plurality of physical links (shown as dashed lines in FIG. 1 that separately connect the PE device to the first NVE device, the second NVE device, and the third NVE device) from the PE device to the first NVE device, the second NVE device, and the third NVE device. In another possible implementation scenario, a router is included between the PE device and the NVE device set. The router performs IP forwarding to implement a traffic off-loading function. The PE device is connected to the router by using a physical link. The router is connected to the first NVE device, the second NVE device, and the third NVE device by using a plurality of physical links. In the foregoing two possible implementation scenarios, physical links from the PE device to a plurality of NVE devices in the NVE device set are virtualized as a single VXLAN tunnel. Therefore, when sending a packet to the NVE device set, the PE device considers the NVE device set as a single NVE device. To achieve the foregoing objective, the NVE device set includes a common VTEP. The common VTEP includes a common VTEP address. The common VTEP address is used to identify the common VTEP. The common VTEP address may be an IP address. In addition, the common VTEP is deployed on each NVE device in the NVE device set. In this way, each NVE device in the NVE device set includes the common VTEP. The NVE devices in the NVE device set include a same VTEP, and the NVE devices in the NVE device set include a same VTEP address. Therefore, in this application, a set including NVE devices that include a same VTEP is referred to as an NVE device set. After receiving an inclusive multicast Ethernet tag (IMET) route sent by each NVE device in the NVE device set, the PE device determines that only one VTEP address (the common VTEP address) exists at a transmit end of the IMET route. Therefore, for the PE device, the NVE device set is like an NVE device. The PE device establishes a VXLAN tunnel by using the common VTEP as an endpoint of the VXLAN tunnel. FIG. 1 is used as an example. A value of the common VTEP address of the NVE device set is 9.9.9.9. The common VTEP including the common VTEP address is deployed on each of the first NVE device, the second NVE device, and the third NVE device. The PE device may establish a VXLAN tunnel from the PE device to the NVE device set based on the received IMET route from the NVE device. For a specific implementation of establishing the VXLAN tunnel, refer to descriptions of a subsequent embodiment of this application. In addition, from a perspective of the PE device, the NVE device set is a virtual NVE device. Use of "set" in the technical term "NVE device set" constitutes no limitation to this application. For example, "group" may be used to replace "set", that is, "NVE device group".

[0020] In the NVE device set, a VXLAN tunnel is es-

established between NVE devices. That is, a VXLAN tunnel is established between each pair of BGP peers in the NVE device set. FIG. 1 is used as an example. A first VXLAN tunnel is established between the first NVE device and the second NVE device. A second VXLAN tunnel is established between the second NVE device and the third NVE device. A third VXLAN tunnel is established between the first NVE device and the third NVE device. To achieve the foregoing objective, each NVE device in the NVE device set further includes a VTEP, and the VTEP includes a VTEP address. In addition, the VTEP address is unique in the EVPN VXLAN. The VTEP address is used to identify a VTEP in a corresponding NVE device. The VTEP address may be an IP address. Further, the common VTEP address of the NVE device set is different from an endpoint address of a VXLAN tunnel between any two NVE devices in the NVE device set. FIG. 1 is used as an example. The first NVE device includes the common VTEP mentioned above, and a value of the common VTEP address included in the common VTEP is 9.9.9.9. The first NVE device further includes a first VTEP, and a value of a first VTEP address included in the first VTEP is 1.1.1.1. The second NVE device includes the common VTEP mentioned above, and the value of the common VTEP address is 9.9.9.9. The second NVE device further includes a second VTEP, and a value of a second VTEP address included in the second VTEP is 2.2.2.2. The third NVE device includes the common VTEP mentioned above, and the value of the common VTEP address is 9.9.9.9. The third NVE device further includes a third VTEP, and a value of a third VTEP address included in the third VTEP is 3.3.3.3. An NVE device in the NVE device set may establish a VXLAN tunnel from the NVE device to a BGP peer of the NVE device based on a VTEP address and the common VTEP address by using an IMET route. For a specific implementation of establishing the VXLAN tunnel, refer to descriptions of a subsequent embodiment of this application.

[0021] In the foregoing implementation, in an EVPN VXLAN application scenario, an NVE device set including at least two NVE devices includes a common VTEP address. The common VTEP address is used to identify a common VTEP, and the common VTEP is deployed on each NVE device in the NVE device set. In addition, each NVE device in the NVE device set includes a VTEP address, and the VTEP address is used to identify a VTEP included in a corresponding NVE device. An NVE device in the NVE device set establishes a VXLAN tunnel from the NVE device to a BGP peer of the NVE device based on a VTEP address and the common VTEP address by using an IMET route. Therefore, in the EVPN VXLAN application scenario, NVE devices in the NVE device set are interconnected through a VXLAN tunnel, so that a link between NVE devices in the NVE device set is not limited to a physical direct link. The VXLAN tunnel is used to implement interconnection between the NVE devices, so that even if a switching node exists between the NVE devices in the NVE device set, implementation of the

EVPN VXLAN is not affected. When the NVE device in the NVE device set is deployed, the NVE device is not affected by a deployment region and distance, thereby improving deployment flexibility. Further, a link between the NVE devices in the NVE device set is not limited to a physical direct link, thereby helping improve deployment flexibility.

[0022] FIG. 2A and FIG. 2B are a flowchart of a route processing method according to an embodiment of this application. The method shown in FIG. 2A and FIG. 2B may be applied to the EVPN VXLAN shown in FIG. 1. That is, the method shown in FIG. 2A and FIG. 2B may be applied to the VXLAN having an EVPN control plane shown in FIG. 1. The EVPN VXLAN includes a first NVE device and a second NVE device. The first NVE device and the second NVE device are a pair of BGP peers. For an implementation of establishing a BGP peer, refer to descriptions of the foregoing embodiment. Details are not described herein again. FIG. 2A shows an NVE device set, and the first NVE device and the second NVE device are included in the NVE device set. It should be understood that in an actual application scenario, from a perspective of a PE device, the NVE device set is a single virtual NVE device. The first NVE device and the second NVE device each include a common VTEP. For ease of description, the expression of the NVE device set is used in this embodiment of this application. FIG. 2A shows two NVE devices. It should be understood that in an actual application scenario, a quantity of NVE devices in the NVE device set may be greater than 2. When the quantity of NVE devices in the NVE device set can be greater than 2, any two NVE devices in the NVE device set can implement the route processing method shown in FIG. 2A and FIG. 2B. The method shown in FIG. 2A and FIG. 2B include S101 to S104, and may further include S105.

[0023] S101. The second NVE device sends a second IMET route to the first NVE device.

[0024] S102. The first NVE device receives the second IMET route from the second NVE device, where an originating router's IP address field in the second IMET route includes a common VTEP address, the second IMET route further includes a second VTEP address, and the common VTEP address included in the originating router's IP address field in the second IMET route is different from the second VTEP address.

[0025] Based on the descriptions of the scenario shown in FIG. 1 in the foregoing embodiment, the NVE device set includes a common VTEP, the common VTEP includes a common VTEP address, and the VTEP address may be an IP address. The common VTEP is deployed on each NVE device in the NVE device set. With reference to FIG. 2A and FIG. 2B, the common VTEP is deployed on each of the first NVE device and the second NVE device. Optionally, the common VTEP address may be deployed in the common VTEP on each NVE device in the NVE device set in a static configuration manner. In this way, based on the common VTEP address, from the perspective of the PE device, the NVE device set

may be virtualized as a single NVE device.

[0026] An NVE device in the NVE device set further includes a VTEP, the VTEP includes a VTEP address, and the VTEP address may be an IP address. With reference to FIG. 2A and FIG. 2B, the first NVE device further includes a first VTEP, and the first VTEP includes a first VTEP address. The second NVE device further includes a second VTEP, and the second VTEP includes a second VTEP address. The first VTEP address and the second VTEP address are unique in the EVPN VXLAN, and neither a value of the first VTEP address nor a value of the second VTEP address is equal to a value of the common VTEP address. Based on the foregoing, the NVE device in the NVE device set includes both the common VTEP address and a VTEP address of the NVE device. For example, the common VTEP address and the VTEP address are stored in a memory of the NVE device, so that a processor of the NVE device can access the common VTEP address and the VTEP address.

[0027] The second NVE device sends the second IMET route to the first NVE device. An IMET route mentioned in the RFC 7432 is used as the IMET route used in this embodiment of this application, and the IMET route mentioned in the RFC 7432 is extended in this embodiment of this application. For a format and a function of the IMET route, refer to descriptions of the RFC 7432. Specifically, the second NVE device may send the second IMET route to the first NVE device in the following manner. The second NVE device may send an update message to the first NVE device. The update message carries a multiprotocol reachable network layer reachability information (MP_REACH_NLRI) attribute, and a full name of the NLRI is network layer reachability information. The MP_REACH_NLRI attribute includes an EVPN NLRI field and a next hop field. The EVPN NLRI field is used to carry the second IMET route. For definitions of the update message and the MP_REACH_NLRI attribute, refer to descriptions of the RFC 4760. For a definition of the EVPN NLRI field, refer to descriptions of Chapter 7 of the RFC 7432.

[0028] The second IMET route includes the originating router's IP address field. In this embodiment of this application, both a value of the next hop field and a value of an originating router's IP address in the second IMET route are set to a common VTEP address 9.9.9.9. After the second IMET route is advertised to another NVE device in the NVE device set and the PE device (shown in FIG. 1) outside the NVE device set, it is ensured that the second NVE device can receive a packet transmitted by the PE device and the another NVE device.

[0029] In this embodiment of this application, the IMET route mentioned in the RFC 7432 is further extended, so that the IMET route can carry a VTEP address of an NVE device that sends the IMET route. For example, a management IP address (or referred to as a device IP address) of the NVE device may be used as the VTEP address. In this setting, the second IMET route carries both the common VTEP address 9.9.9.9 and a second VTEP

address 1.1.1.1.

[0030] In a solution, a new attribute is extended in this embodiment of this application to carry the VTEP address. As shown in FIG. 6, a VXLAN attribute is extended in this embodiment of this application to carry the VTEP address. A type field and a sub type field are used to indicate a type of the VXLAN attribute, and specific values of the type field and the sub type field each may be set according to a requirement of a standard organization. A flag field may be used to indicate whether the VTEP address in the VXLAN attribute is allowed to be read. For example, if the flag field is 0, it indicates that the VTEP address in the VXLAN attribute is not allowed to be read; or if the flag field is 1, it indicates that the VTEP address in the VXLAN attribute is allowed to be read. It should be understood that the foregoing functions of the flag field are used as an example, and a function of the flag field may correspondingly change based on a use scenario. A VTEP address field is used to carry a value of the VTEP address. A reserved field is used to implement further extension. The VXLAN attribute may be carried in the IMET route as an attribute, so as to implement the solution in this embodiment of this application. The second IMET route is used as an example. The second IMET route carries the VXLAN attribute, and the VXLAN attribute includes the second VTEP address.

[0031] The second IMET route may further carry a VNI, and the VNI is used to indicate a BD to which the second NVE device belongs. More specifically, the VNI is used to indicate a BD to which a port used by the second NVE device to send the second IMET route belongs. In this application, the BD may be referred to as a broadcast domain (BD) or a bridge domain (BD). The VNI may be carried in a provider multicast service interface (PMSI) tunnel attribute of the second IMET route. The PMSI tunnel attribute includes an MPLS label field, and the MPLS label field is used to carry the VNI.

[0032] In this embodiment of this application, the VNI may be configured in two forms: a common VTEP VNI and a VTEP VNI. The common VTEP VNI is used to bind to the common VTEP address, and the VTEP VNI is used to bind to the VTEP address. With reference to FIG. 1, the common VTEP address is used as an endpoint address of a VXLAN tunnel established between the PE device and the NVE device set. Therefore, the VXLAN tunnel established between the PE device and the NVE device set is corresponding to the common VTEP VNI. The VTEP address is used as an endpoint address of a VXLAN tunnel established between the NVE devices in the NVE device set. Therefore, the VXLAN tunnel established between the NVE devices is corresponding to the VTEP VNI. In this embodiment of this application, the common VTEP VNI and the VTEP VNI may be set to be the same, or may be set to be different.

[0033] If the common VTEP VNI and the VTEP VNI are set to be the same, the VXLAN tunnel established between the PE device and the NVE device set and the VXLAN tunnel established between the NVE devices in

the NVE device set are corresponding to a same BD. When the common VTEP VNI and the VTEP VNI are set to be the same, VNI information may be carried according to the foregoing implementation. To be specific, the VNI is carried in the MPLS label field included in the PMSI tunnel attribute, and the VNI is used to indicate both the common VTEP VNI and the VTEP VNI. The VNI carried in the MPLS label field included in the PMSI tunnel attribute is used in both a process of establishing the VXLAN tunnel between the PE device and the NVE device set and a process of establishing the VXLAN tunnel between the NVE devices in the NVE device set.

[0034] If the common VTEP VNI and the VTEP VNI are set to be different, the VXLAN tunnel established between the PE device and the NVE device set and the VXLAN tunnel established between the NVE devices in the NVE device set are corresponding to different BDs. For example, the VXLAN tunnel established between the PE device and the NVE device set is corresponding to a BD 1, and the VXLAN tunnel established between the NVE devices in the NVE device set is corresponding to a BD 2. That "the common VTEP VNI and the VTEP VNI are set to be different" described above means that a same VTEP VNI is used for all VXLAN tunnels established between the NVE devices in the NVE device set, and the VTEP VNI is different from the common VTEP VNI. When the common VTEP VNI and the VTEP VNI are set to be different, VNI information may be carried according to the following implementation. The common VTEP VNI is carried in the MPLS label field included in the PMSI tunnel attribute. In addition, a new attribute is extended in this embodiment of this application to carry the VTEP VNI. As shown in FIG. 7, a VTEP VNI attribute is extended in this embodiment of this application to carry the VTEP VNI. A type field and a sub type field are used to indicate a type of the VTEP VNI attribute, and specific values of the type field and the sub type field each may be set according to a requirement of a standard organization. A flag field may be used to indicate whether the VTEP VNI in the VTEP VNI attribute is allowed to be read. For example, if the flag field is 0, it indicates that the VTEP VNI in the VTEP VNI attribute is not allowed to be read; or if the flag field is 1, it indicates that the VTEP VNI in the VTEP VNI attribute is allowed to be read. It should be understood that the foregoing functions of the flag field are used as an example, and a function of the flag field may correspondingly change based on a use scenario. A VTEP VNI field is used to carry a value of the VTEP VNI. A reserved field is used to implement further extension. The VTEP VNI attribute may be carried in an IMET route as an attribute. When receiving an IMET route advertised by an NVE device in the NVE device set, the PE device establishes a VXLAN tunnel based on the common VTEP VNI, and does not parse the VTEP VNI attribute, so as to ignore the VTEP VNI. Correspondingly, after receiving an IMET route advertised by another NVE device, the NVE device in the NVE device set preferably establishes a VXLAN tunnel based on the VTEP

VNI in the VTEP VNI attribute. If the NVE device in the NVE device set determines that the VTEP VNI attribute does not exist in the received IMET route, the NVE device establishes a corresponding VXLAN tunnel by using the common VTEP VNI carried in the MPLS label field included in the PMSI tunnel attribute.

[0035] Further, the VXLAN tunnels established between the NVE devices in the NVE device set may use different VTEP VNIs. For example, in FIG. 1, the first VXLAN tunnel uses a VTEP VNI 1 and is corresponding to a BD 21, the second VXLAN tunnel uses a VTEP VNI 2 and is corresponding to a BD 22, and the third VXLAN tunnel uses a VTEP VNI 3 and is corresponding to a BD 23. In this way, a plurality of VNIs are allowed to be configured for the NVE device, so that the NVE device can establish VXLAN tunnels corresponding to different BDs. In this embodiment of this application, to make the technical solution easier to understand, the common VTEP VNI and the VTEP VNI are set to be the same, and the VNI may be used to represent the common VTEP VNI and the VTEP VNI in subsequent descriptions. It should be understood that in this embodiment of this application, in addition to setting the common VTEP VNI and the VTEP VNI to be the same, the common VTEP VNI and the VTEP VNI may be set to be different. Further, different VTEP VNIs may be set, and this embodiment of this application is implemented with reference to the foregoing descriptions.

[0036] A route distinguisher (RD) field in the second IMET route is used to indicate an EVPN instance (EVI) corresponding to the second NVE device. An Ethernet tag field in the second IMET route may be set to a virtual local area network identifier (VLAN ID or VID) corresponding to a MAC address included in a CE device connected to the second NVE device or set to a BD corresponding to the second NVE device. An IP address length field in the second IMET route is used to indicate a length of the originating router's IP address.

[0037] The PMSI tunnel attribute of the second IMET route further includes a tunnel type field, and the tunnel type field is used to indicate a type of an established tunnel. In this embodiment of this application, a value of the tunnel type field is set to 6, and this is used to identify that a type of an established tunnel is ingress replication. The second IMET route further includes a VXLAN tunnel attribute, and a format of the VXLAN tunnel attribute is similar to the attribute format shown in FIG. 6 or FIG. 7. A difference lies in that the VTEP address field in FIG. 6 or the VTEP VNI field in FIG. 7 is replaced with a VXLAN tunnel type field. The VXLAN tunnel type field is used to indicate that an established tunnel is a VXLAN tunnel. A value of the VXLAN tunnel type field is specifically determined by a standard organization. For example, the value of the VXLAN tunnel type field is 8. For an implementation in which the first NVE device establishes a VXLAN tunnel by using the second IMET route, refer to subsequent descriptions of this embodiment. The PMSI tunnel attribute of the second IMET route further includes a tunnel iden-

tifier field. In this embodiment of this application, the common VTEP address may be filled in the tunnel identifier field.

[0038] S103. The first NVE device determines whether the common VTEP address included in the originating router's IP address field in the second IMET route is the same as a common VTEP address stored in the first NVE device.

[0039] The first NVE device can send, to a BGP peer of the first NVE device, an IMET route whose originating router's IP address field carries the common VTEP address stored in the first NVE device. The common VTEP address included in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device.

[0040] The second NVE device can advertise the second IMET route to the outside, and the first NVE device and the second NVE device are a pair of BGP peers. Therefore, the first NVE device can receive the second IMET route.

[0041] After receiving the second IMET route, the first NVE device parses the second IMET route to obtain the common VTEP address included in the originating router's IP address field in the second IMET route. The first NVE device determines whether the common VTEP address included in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device. In this embodiment of this application, the first NVE device and the second NVE device each include the common VTEP. Therefore, the common VTEP address included in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device. When advertising an IMET route to the outside, the first NVE device adds the common VTEP address stored in the first NVE device to an originating router's IP address field in the IMET route. To be specific, the common VTEP address stored in the first NVE device is the common VTEP address in the common VTEP included in the first NVE device. If the first NVE device determines that the common VTEP address included in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device, the first NVE device may determine that the first NVE device and the second NVE device belong to a same NVE device set.

[0042] S104. When the first NVE device determines that the common VTEP address included in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device, the first NVE device establishes a VXLAN tunnel from the first NVE device to the second NVE device based on the second VTEP address in the second IMET route.

[0043] Based on the foregoing result of determining, the first NVE device can establish the VXLAN tunnel from the first NVE device to the second NVE device based on the second VTEP address. Specifically, the first NVE de-

vice continues to parse the second IMET route to obtain the second VTEP address in the second IMET route, and establishes the VXLAN tunnel from the first NVE device to the second NVE device by using the second VTEP address as a destination endpoint address of the VXLAN tunnel.

[0044] In a possible implementation, with reference to the descriptions of S101 and S102, the second VTEP address is carried in a VXLAN attribute. Based on the foregoing result of determining, that is, the first NVE device determines that the common VTEP address included in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device, the first NVE device checks whether the second IMET route carries the VXLAN attribute. If the first NVE device determines that the second IMET route carries the VXLAN attribute, the first NVE device obtains the second VTEP address in the VXLAN attribute, and establishes the VXLAN tunnel from the first NVE device to the second NVE device by using the second VTEP address as the destination endpoint address of the VXLAN tunnel.

[0045] For example, based on that the value of the tunnel type field in the PMSI tunnel attribute of the second IMET route is 6, the first NVE device determines that the type of the established tunnel is ingress replication. The first NVE device parses the VXLAN tunnel attribute of the second IMET route to obtain a value of the VXLAN tunnel type field in the VXLAN tunnel attribute, so as to determine that the established tunnel is a VXLAN tunnel. After determining that the established tunnel is a VXLAN tunnel, the first NVE device generates a VXLAN tunnel table and a VXLAN tunnel ingress replication table. The VXLAN tunnel table includes a tunnel identifier, a source endpoint address, and a destination endpoint address. The tunnel identifier is used to identify a VXLAN tunnel. The source endpoint address is the first VTEP address included in the first NVE device, for example, 1.1.1.1. The destination endpoint address is the second VTEP address obtained from the second IMET route, for example, 2.2.2.2. The VXLAN tunnel ingress replication table includes a VNI and a BGP peer address. Based on the foregoing descriptions, the VNI may be obtained from the MPLS label field included in the PMSI tunnel attribute or from the VTEP VNI attribute. The BGP peer address is the second VTEP address obtained from the second IMET route, for example, 2.2.2.2. Therefore, the first NVE device establishes the VXLAN tunnel from the first NVE device to the second NVE device.

[0046] If the first NVE device determines that a VTEP address included in an originating router's IP address field in a received IMET route is different from the common VTEP address stored in the first NVE device, the first NVE device may determine that the first NVE device and a device that sends the IMET route do not belong to a same NVE device set. The first NVE device establishes, based on the VTEP address included in the originating router's IP address field in the received IMET route, a

VXLAN tunnel from the first NVE device to the device that sends the IMET route, instead of establishing the VXLAN tunnel based on a VTEP address in a VXLAN attribute. Similarly, for details, refer to a process of establishing a VXLAN tunnel between a PE device and an NVE device set described in a subsequent embodiment.

[0047] If the first NVE device determines that the common VTEP address included in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device, but cannot further find the second VTEP address or the VXLAN attribute, the first NVE device discards the second IMET route, and returns an error message to the second NVE device. Further, optionally, the first NVE device may feed back the error message to a network management device of the EVPN VXLAN. This setting helps improve reliability of establishing a VXLAN tunnel.

[0048] Optionally, S105. The first NVE device sends a first IMET route to the second NVE device, where an originating router's IP address field in the first IMET route carries the common VTEP address stored in the first NVE device, the first IMET route further includes a first VTEP address, the first VTEP address is used by the second NVE device to establish the VXLAN tunnel from the second NVE device to the first NVE device, and the common VTEP address stored in the first NVE device is different from the first VTEP address.

[0049] For an implementation process in which the first NVE device sends the first IMET route to the second NVE device, and an implementation of the first IMET route, refer to the descriptions of S101 and S102 in this embodiment. Details are not described herein again.

[0050] After receiving the first IMET route, the second NVE device may establish the VXLAN tunnel from the second NVE device to the first NVE device based on the first IMET route. For an implementation process in which the second NVE device establishes the VXLAN tunnel from the second NVE device to the first NVE device based on the first IMET route, refer to the descriptions of S103 and S104 in this embodiment. Details are not described herein again.

[0051] It should be understood that the descriptions of the foregoing embodiment constitute no limitation on a sequence of an execution process in S105 and execution processes in S101 to S104. For example, S105 may be performed before S101, S105 and S101 may be simultaneously triggered, or S105 may be performed after S101.

[0052] Through the foregoing process, the VXLAN tunnel can be established between the first NVE device and the second NVE device, and endpoint addresses of two ends of the VXLAN tunnel are the first VTEP address and the second VTEP address.

[0053] In the foregoing implementation, in an EVPN VXLAN application scenario, an NVE device set including at least two NVE devices includes a common VTEP address. The common VTEP address is used to identify a common VTEP, and the common VTEP is deployed on

each NVE device in the NVE device set. In addition, each NVE device in the NVE device set includes a VTEP address, and the VTEP address is used to identify a VTEP included in a corresponding NVE device. An NVE device in the NVE device set establishes a VXLAN tunnel between the NVE devices in the NVE device set based on a VTEP address and the common VTEP address by using an IMET route. Therefore, in the EVPN VXLAN application scenario, NVE devices in the NVE device set are interconnected through a VXLAN tunnel, so that a link between NVE devices in the NVE device set is not limited to a physical direct link. The VXLAN tunnel is used to implement interconnection, so that even if a switching node exists between the NVE devices in the NVE device set, implementation of the EVPN VXLAN is not affected. In addition, when the NVE device in the NVE device set is deployed, the NVE device is not affected by a deployment region and distance, thereby reducing deployment complexity. Further, an interconnection link between the NVE devices in the NVE device set is no longer affected by a physical direct link, thereby helping improve deployment flexibility.

[0054] Optionally, based on the descriptions of FIG. 1 in the foregoing embodiment, the EVPN VXLAN further includes a PE device. The PE device and the first NVE device are a pair of BGP peers, the PE device and the second NVE device are a pair of BGP peers, and the PE device and the third NVE device are a pair of BGP peers. A process of establishing a VXLAN tunnel between the PE device and the NVE device set is described below by using FIG. 1 as an example. The route processing method further includes the following steps.

[0055] S201. The first NVE device sends the first IMET route to the PE device.

[0056] S202. The second NVE device sends the second IMET route to the PE device.

[0057] S203. The third NVE device sends the third IMET route to the PE device.

[0058] Based on the descriptions of FIG. 2A and FIG. 2B in the foregoing embodiment, as NVE devices in the NVE device set, the first NVE device, the second NVE device, and the third NVE device advertise IMET routes to respective BGP peers of the devices. To be specific, the first NVE device advertises the first IMET route to a BGP peer of the first NVE device, the second NVE device advertises the second IMET route to a BGP peer of the second NVE device, and the third NVE device advertises the third IMET route to a BGP peer of the third NVE device. For implementations of the first IMET route, the second IMET route, and the third IMET route, refer to descriptions of the foregoing embodiment. Details are not described herein again.

[0059] S204. The PE device receives the first IMET route, the second IMET route, and the third IMET route.

[0060] S205. The PE device establishes a VXLAN tunnel from the PE device to the NVE device set based on the first IMET route, the second IMET route, and the third IMET route.

[0061] Because the PE device is a BGP peer of each NVE device in the NVE device set, the PE device can receive the first IMET route, the second IMET route, and the third IMET route. S204 and S205 may be implemented in two possible manners.

[0062] In one possible implementation, after receiving the IMET routes sent by all the BGP peers, the PE device establishes a VXLAN tunnel. After receiving the three IMET routes, the PE device parses the three IMET routes separately to obtain common VTEP addresses included in originating router's IP address fields in the three IMET routes. If the PE device determines that the common VTEP addresses included in the originating router's IP address fields in the three IMET routes are the same, the PE device may determine that the NVE devices that send the three IMET routes belong to a same NVE device set. The PE device randomly selects an IMET route from the three IMET routes, and the PE device establishes, by using the common VTEP address included in the selected IMET route as a destination endpoint address of a VXLAN tunnel, the VXLAN tunnel from the PE device to a common VTEP including the common VTEP address.

[0063] In the other possible implementation, after receiving any IMET route in the three IMET routes, the PE device establishes a VXLAN tunnel. For example, after receiving the first IMET route sent by the first NVE device, the first PE device parses the first IMET route to obtain a common VTEP address included in an originating router's IP address field in the first IMET route. The PE device establishes, by using the common VTEP address included in the first IMET route as a destination endpoint address of a VXLAN tunnel, the VXLAN tunnel from the PE device to a common VTEP including the common VTEP address. After the VXLAN tunnel is established, when the PE device receives the second IMET route sent by the second NVE device or the third IMET route sent by the third NVE device, and determines that a common VTEP address included in an originating router's IP address field in the second IMET route or the third IMET route is the same as the common VTEP address used to establish the VXLAN tunnel, the PE device no longer establishes the VXLAN tunnel from the PE device to the common VTEP including the common VTEP address.

[0064] In the foregoing two implementations, for a specific implementation of establishing a VXLAN tunnel, refer to the descriptions of the foregoing embodiments. Details are not described herein again. In the foregoing two implementations, the common VTEP addresses included in the originating router's IP address fields in the three IMET routes are the same. After the PE device receives IMET routes having a same common VTEP address, the PE device considers that only one VTEP corresponding to the common VTEP address exists at a peer end of the PE device. Therefore, the PE device establishes a VXLAN tunnel for the common VTEP address only once, instead of establishing a VXLAN tunnel for the common VTEP address stored in each NVE device in an NVE device set. In this way, a VXLAN tunnel that is from the

PE device to the common VTEP including the common VTEP address and that is established by the PE device by using the common VTEP address as the destination endpoint address of the VXLAN tunnel may be referred to as a VXLAN tunnel from the PE device to the NVE device set.

[0065] Optionally, the three IMET routes each include a VNI. For establishment of a VXLAN tunnel, the VNI is the common VTEP VNI described above. After determining that the three IMET routes include the same common VTEP address, the PE device further determines that the three IMET routes include a same common VTEP VNI, so as to establish the VXLAN tunnel from the PE device to the NVE device set. For explanations of the common VTEP VNI, refer to the corresponding descriptions of the foregoing embodiments. Details are not described herein again.

[0066] S206. The PE device sends an IMET route to each of the first NVE device, the second NVE device, and the third NVE device, where a next hop field and an originating router's IP address field that is in the IMET route each include a VTEP address of the PE device, and the VTEP address of the PE device is used to identify a VTEP included in the PE device.

[0067] The PE device may advertise the IMET route to a BGP peer of the PE device. For a format and basic usage of the IMET route, refer to descriptions of the RFC 7432. For example, in this embodiment of this application, a value of the VTEP address included in the VTEP of the PE device is set to 4.4.4.4. An example in which the first NVE device receives the IMET route sent by the PE device is used as an example below for description. It should be understood that the second NVE device and the third NVE device may perform a same processing process.

[0068] S207. The first NVE device receives the IMET route from the PE device.

[0069] S208. The first NVE device establishes a VXLAN tunnel from the first NVE device to the PE device based on the VTEP address from the PE device, where the VTEP address of the PE device is not the same as a common VTEP address stored in the first NVE device.

[0070] After receiving the IMET route sent by the PE device, the first NVE device parses the IMET route to obtain the VTEP address of the PE device that is carried in the originating router's IP address field in the IMET route, and establishes the VXLAN tunnel from the first NVE device to the PE device by using the VTEP address as a destination endpoint address of the VXLAN tunnel. Optionally, the IMET route further includes a common VTEP VNI. The first NVE device obtains the common VTEP VNI, and establishes a VXLAN tunnel that is corresponding to the common VTEP VNI and that is from the first NVE device to the PE device.

[0071] Optionally, after obtaining the VTEP address of the PE device, the first NVE device may perform a similar determining process as S103 and S104. To be specific, if the first NVE device determines that the VTEP address

of the PE device is different from the common VTEP address stored in the first NVE device, the first NVE device may determine that the first NVE device and the PE device do not belong to a same NVE device set. The first NVE device establishes the VXLAN tunnel from the first NVE device to the PE device by using the common VTEP address as a source endpoint address, instead of establishing the VXLAN tunnel from the first NVE device to the PE device by using a first VTEP address as the source endpoint address.

[0072] Based on S207 and S208, the second NVE device establishes a VXLAN tunnel from the second NVE device to the PE device, and the third NVE device establishes a VXLAN tunnel from the third NVE device to the PE device. With reference to the foregoing, the three VXLAN tunnels established to the PE device by the first NVE device, the second NVE device, and the third NVE device are the same, that is, source endpoint addresses of the three VXLAN tunnels each are the common VTEP address, destination endpoint addresses of the three VXLAN tunnels each are the VTEP address of the PE device, and the three VXLAN tunnels each are corresponding to the common VTEP VNI. Correspondingly, the PE device establishes a VXLAN tunnel from the PE device to the NVE device set, a source endpoint address is the VTEP address of the PE device, and a destination endpoint address is the common VTEP address. In this way, the PE device considers that only one NVE device exists at the peer end. Actually, the NVE device is the NVE device set. In addition, the first NVE device, the second NVE device, and the third NVE device can still send traffic through the VXLAN tunnels respectively established to the PE device.

[0073] Based on the foregoing embodiment, this embodiment of this application may be applied to a network scenario in which more than two NVE devices form an NVE device set. Using FIG. 1 as an example, the NVE device set includes three NVE devices. According to the method shown in FIG. 2A and FIG. 2B, a first VXLAN tunnel is established between the first NVE device and the second NVE device, a second VXLAN tunnel is established between the second NVE device and the third NVE device, and a third VXLAN tunnel is established between the first NVE device and the third NVE device. In this way, according to the technical solution in this application, a network scenario in which one CE device is connected to three or more NVE devices may be implemented, that is, the CE device is multi-homed to three or more NVE devices. Using FIG. 1 as an example, the first CE device is multi-homed to the first NVE device, the second NVE device, and the third NVE device. The foregoing implementation is an all-active redundancy mode of the EVPN VXLAN. To be specific, a plurality of NVE devices are connected to at least one CE device by using ESs, and the plurality of NVE devices form a device set in the active-active redundancy mode and are all used to forward traffic. In comparison, in an existing all-active redundancy mode of the EVPN VXLAN, a dual-active

redundancy mode is used, and a CE device can be dual-homed to a maximum of two NVEs. In addition, the two NVE devices are directly connected to each other by using a physical link. For example, the two NVE devices are connected to each other by using a peer link. Therefore, a further beneficial effect of this embodiment of this application lies in that the active-active redundancy mode supports more NVE devices, thereby helping improve reliability of the EVPN VXLAN and flexibility of network deployment.

[0074] FIG. 3 is a schematic structural diagram of another EVPN VXLAN according to an embodiment of this application. FIG. 3 shows the following: After a VXLAN tunnel between NVE devices in an NVE device set is established by using the technical solution provided in the foregoing embodiments, a designated forwarder (DF) is elected, and broadcast, unknown unicast, and multi-cast (BUM) traffic are forwarded. FIG. 3 shows three CE devices as an example: a first CE device, a second CE device, and a third CE device. The first CE device is multi-homed to the first NVE device, the second NVE device, and the third NVE device, that is, the first CE device communicates with the first NVE device, the second NVE device, and the third NVE device separately. The second CE device is dual-homed to the first NVE device and the second NVE device, that is, the second CE device communicates with the first NVE device and the second NVE device separately. The third CE device is single-homed to the second NVE device, that is, the third CE device communicates with the second NVE device. An example in which BUM traffic is forwarded from a PE device to the first CE device, the second CE device, and the third CE device is used to describe FIG. 3.

[0075] The first CE device is multi-homed to the first NVE device, the second NVE device, and the third NVE device by using a plurality of Ethernet links, and therefore the plurality of Ethernet links form an ES, as shown by an ES 1 in FIG. 3. ESI values of ESs 1 of three Ethernet links that connect the first CE device to the first NVE device, the second NVE device, and the third NVE device are a same value. The first NVE device, the second NVE device, and the third NVE device separately send Ethernet segment (ES) routes to respective BGP peers of the devices, and the ES routes are used for DF election. The ES route includes an RD field, an ESI field, an IP address length field, and an originating router's IP address field. Definitions of the fields are similar to that of the foregoing IMET route. Details are not described herein again. In addition, for a format and basic usage of the ES route, refer to descriptions of Chapter 7.4 and Chapter 8.5 of the RFC 7432. Using the first CE device in FIG. 3 as an example, the first NVE device receives ES routes from the second NVE device and the third NVE device. The first NVE device determines, based on the ES routes received from the second NVE device and the third NVE device, the first NVE device as a DF in the ES that includes the plurality of Ethernet links. Likewise, the second NVE device is determined as a non-designated for-

warder (NDF) in the ES that includes the plurality of Ethernet links, and the third NVE device is determined as an NDF in the ES that includes the plurality of Ethernet links. As shown in FIG. 3, the ES 1 between the first CE device and the first NVE device is a DF, the ES 1 between the first CE device and the second NVE device is an NDF, and the ES 1 between the first CE device and the third NVE device is an NDF. The second CE device is dual-homed to the first NVE device and the second NVE device by using two Ethernet links, and the two Ethernet links form an ES, as shown by an ES 2 in FIG. 3. ESI values of ESs 2 of the two Ethernet links that connect the second CE device to the first NVE device and the second NVE device are a same value. Likewise, the second NVE device is determined as a DF in the ES that includes the two Ethernet links, and the first NVE device is determined as an NDF in the ES that includes the two Ethernet links. The third CE device is single-homed to the second NVE device, and therefore has no DF election process.

[0076] The NVE device may send the ES route to a BGP peer of the NVE device in the following manner. The NVE device may send an update message to the BGP peer of the NVE device, and the update message carries an MP_REACH_NLRI attribute. The MP_REACH_NLRI attribute includes an EVPN NLRI field and a next hop field. The EVPN NLRI field is used to carry the ES route.

[0077] The PE device sends BUM traffic to an NVE device set based on a VXLAN tunnel. Specifically, the PE device sends a VXLAN packet to the NVE device set according to a hash rule, and the VXLAN packet carries the BUM traffic. Optionally, the VXLAN packet further carries a BUM mark, and the BUM mark is used to identify that traffic carried in the VXLAN packet is the BUM traffic. In a possible implementation, the BUM mark may be carried in a VXLAN header in the VXLAN packet.

[0078] An example in which the VXLAN packet is hashed to the second NVE device is used for description. In this embodiment and a subsequent embodiment, to distinguish between the VXLAN packet transmitted between the PE device and the NVE device set and a VXLAN packet transmitted between NVE devices in the NVE device set, the VXLAN packet transmitted between the PE device and the NVE device set is named as a PE VXLAN packet. The second NVE device receives the PE VXLAN packet from the PE device, parses the PE VXLAN packet, and may determine, based on the BUM mark carried in the PE VXLAN packet, that traffic carried in the PE VXLAN packet is BUM traffic. The second NVE device forwards the BUM traffic to a CE device that has a connection relationship with the second NVE device, and the second NVE device further forwards the BUM traffic to another NVE device in the NVE device set.

[0079] That the second NVE device forwards the BUM traffic to a CE device that has a connection relationship with the second NVE device includes the following: The second NVE device determines that the ES 2 between

the second NVE device and the second CE device is in a DF state, and the second NVE device transmits the BUM traffic to an outbound interface that connects the second NVE device and the ES 2, and sends the BUM traffic to the second CE device by using the ES 2. The second NVE device determines that the ES 1 between the second NVE device and the first CE device is in an NDF state, and the second NVE device does not transmit the BUM traffic to an outbound interface that connects the second NVE device and the ES 1, so that the second NVE device avoids sending the BUM traffic to the first CE device by using the ES 1. In addition, the second NVE device determines that an ES 3 between the second NVE device and the third CE device does not have a DF or an NDF state, and directly sends the BUM traffic to the third CE device by using the ES 3.

[0080] That the second NVE device forwards the BUM traffic to another NVE device in the NVE device set includes the following: The second NVE device parses the PE VXLAN packet from the PE device to obtain the BUM traffic. The second NVE device determines that the BUM traffic is from the PE device, encapsulates the BUM traffic as a VXLAN packet, sends the VXLAN packet to the first NVE device through a first VXLAN tunnel, and sends the VXLAN packet to the second NVE device through a second VXLAN tunnel.

[0081] After receiving the VXLAN packet through the first VXLAN tunnel, the first NVE device can determine that the BUM traffic in the VXLAN packet is received through the first VXLAN tunnel, and no longer forwards the BUM traffic to a third VXLAN tunnel connected to the first NVE device. As shown in FIG. 3, after receiving the VXLAN packet through the first VXLAN tunnel, the first NVE device parses the VXLAN packet to obtain the BUM traffic, and no longer forwards the BUM traffic to the third NVE device through a third VXLAN tunnel. To be specific, the first NVE device avoids forwarding the BUM traffic to the third NVE device through the VXLAN tunnel from the first NVE device to the third NVE device, based on that the BUM traffic is received through the first VXLAN tunnel. A beneficial effect based on the foregoing lies in that the BUM traffic is prevented from being dual-transmitted.

[0082] After receiving the VXLAN packet through the first VXLAN tunnel, the first NVE device parses the VXLAN packet to obtain the BUM traffic, and further sends the BUM traffic to a CE device connected to the first NVE device. As shown in FIG. 3, the first NVE device determines that the ES 1 between the first NVE device and the first CE device is in a DF state, and the first NVE device transmits the BUM traffic to an outbound interface that connects the first NVE device and the ES 1, and sends the BUM traffic to the first CE device by using the ES 1. The first NVE device determines that the ES 2 between the first NVE device and the second CE device is in an NDF state, and the first NVE device does not send the BUM traffic to an outbound interface that connects the first NVE device and the ES 2, so that the first NVE device avoids sending the BUM traffic to the second

CE device by using the ES 2.

[0083] After receiving the VXLAN packet through the second VXLAN tunnel, the third NVE device parses the VXLAN packet to obtain the BUM traffic. In one aspect, the third NVE device avoids forwarding the BUM traffic to the first NVE device through the third VXLAN tunnel. In another aspect, the third NVE device determines that the ES 1 between the third NVE device and the first CE device is in an NDF state, and the third NVE device does not transmit the BUM traffic to an outbound interface that connects the third NVE device and the ES 1, so that the third NVE device avoids sending the BUM traffic to the first CE device by using the ES 1.

[0084] In the foregoing implementation, not only BUM traffic can be normally forwarded through a VXLAN tunnel between NVE devices, but also forwarding of the BUM traffic to a VXLAN tunnel between other NVE devices can be avoided, so as to avoid dual transmission of the BUM traffic.

[0085] Likewise, a processing of forwarding BUM traffic from a CE device to a PE device and another CE device is an inverse process of the foregoing implementation. Therefore, forwarding of the BUM traffic may be completed by referring to the foregoing manner. For example, BUM traffic is forwarded from the second CE device to the PE device, the first CE device, and the third CE device. The BUM traffic arrives at the second NVE device by using the ES 2 between the second NVE device and the second CE device. The second NVE device encapsulates the BUM traffic as a PE VXLAN packet, and forwards the PE VXLAN packet to the PE device through a VXLAN tunnel between the second NVE device and the PE device. The second NVE device further encapsulates the BUM traffic as a VXLAN packet, and forwards the VXLAN packet to the first NVE device and the third NVE device through the first VXLAN tunnel and the second VXLAN tunnel respectively. The second NVE device further directly forwards the BUM traffic to the third CE device by using the ES 3, but avoids forwarding the BUM traffic to the first CE device by using the ES 1 between the second NVE device and the first CE device. After receiving the BUM traffic from the second NVE device, the first NVE device avoids forwarding the BUM traffic to the third NVE device through the third VXLAN tunnel, and avoids forwarding the BUM traffic to the PE device through the VXLAN tunnel. The first NVE device forwards the BUM traffic to the first CE device by using the ES 1, but avoids forwarding the BUM traffic to the second CE device by using the ES 2. After receiving the BUM traffic from the second NVE device, the third NVE device avoids forwarding the BUM traffic to the first NVE device through the third VXLAN tunnel, avoids forwarding the BUM traffic to the first CE device by using the ES 1, and further avoids forwarding the BUM traffic to the PE device through the VXLAN tunnel.

[0086] In the foregoing enumerated process in which the BUM traffic is forwarded from the second CE device to the PE device, the first CE device, and the third CE

device, the BUM traffic sent by the second CE device arrives at the second NVE device by using the ES 2 between the second NVE device and the second CE device. However, in an actual scenario, the BUM traffic sent by the second CE device may arrive at the first NVE device by using the ES 2 between the first NVE device and the second CE device, because the second CE device does not care about a DF state or an NDF state of a communication link. As shown in FIG. 4, when the second CE device forwards the BUM traffic to the first NVE device by using the ES 2, the first NVE device encapsulates the BUM traffic as a VXLAN packet, and forwards the VXLAN packet to the second NVE device through the first VXLAN tunnel. Because the ES 2 between the second NVE device and the second CE device is in a DF state, the second NVE device forwards the BUM traffic to the second CE device by using the ES 2, as shown by a dashed line arrow between the second NVE device and the second CE device in FIG. 4. In this way, the second CE device, the first NVE device, and the second NVE device form a BUM traffic loop. It should be understood that to more clearly present a loop problem, a BUM traffic forwarding line shown in FIG. 4 is only a part at which a loop occurs, and FIG. 4 does not show an entire forwarding line of the BUM traffic in the EVPN VXLAN.

[0087] A method for resolving the problem of the BUM traffic loop including the second CE device, the first NVE device, and the second NVE device is specifically described below with reference to FIG. 4. The second CE device is dual-homed to the first NVE device and the second NVE device by using two Ethernet links, and the two Ethernet links form an ES, as shown by an ES 2 in FIG 4. ESI values of ESs 2 of the two Ethernet links that connect the second CE device to the first NVE device and the second NVE device are a same value. The second NVE device determines, based on an ES route received from the first NVE device, the second NVE device as a DF in the ES that includes the two Ethernet links.

[0088] S301. The second NVE device sends an Ethernet Auto-Discovery per Ethernet segment (Ethernet A-D per ES) route to the first NVE device, where the Ethernet A-D per ES route carries an ESI label allocated by the second NVE device to an Ethernet link between the second CE device and the second NVE device.

[0089] S302. The first NVE device receives the Ethernet A-D per ES route.

[0090] Because the first NVE device and the second NVE device are a pair of BGP peers, the second NVE device can send the Ethernet A-D per ES route to the first NVE device. The Ethernet A-D per ES route is an Ethernet Auto-Discovery route. For a format and basic usage of the Ethernet Auto-Discovery route, refer to descriptions of Chapter 7.1 and Chapter 8.2 of the RFC 7432. The Ethernet A-D per ES route may carry an ESI label extended community attribute, and the ESI label extended community attribute may carry the ESI label. After receiving the Ethernet A-D per ES route, the first NVE device may determine, based on an ESI included

in the Ethernet A-D per ES route, that the second CE device is not only connected to the first NVE device, but also connected to the second NVE device. In addition, the first NVE device may determine, based on the ESI label carried in the Ethernet A-D per ES route, a value of the ESI label allocated by the second NVE device to the Ethernet link between the second CE device and the second NVE device.

[0091] The NVE device may send the Ethernet A-D per ES route to a BGP peer of the NVE device in the following manner. The NVE device may send an update message to the BGP peer of the NVE device, and the update message carries an MP_REACH_NLRI attribute. The MP_REACH_NLRI attribute includes an EVPN NLRI field and a next hop field. The EVPN NLRI field is used to carry the Ethernet A-D per ES route.

[0092] S303. The first NVE device sends a VXLAN packet to the second NVE device through the first VXLAN tunnel, where an original Ethernet payload in the VXLAN packet includes BUM traffic, the VXLAN packet carries the ESI label, and the BUM traffic is from the CE device.

[0093] S304. The second NVE device receives, through the first VXLAN tunnel, the VXLAN packet sent by the first NVE device.

[0094] As shown in FIG. 4, when the second CE device forwards the BUM traffic to the first NVE device by using the ES 2, after receiving the BUM traffic, the first NVE device determines an ESI of an ES 1 between the second CE device and the first NVE device, and may further determine that the second CE device is not only connected to the first NVE device, but also connected to the second NVE device. Therefore, the first NVE device encapsulates the BUM traffic as a VXLAN packet. The original Ethernet payload in the VXLAN packet includes the BUM traffic, and the VXLAN packet carries the ESI label. In other words, the VXLAN packet sent by the first NVE device to the second NVE device through the first VXLAN tunnel not only includes the BUM traffic, but also carries the ESI label allocated by the second NVE device to the Ethernet link between the second CE device and the second NVE device. For a definition of the original Ethernet payload, refer to descriptions of the RFC 7348.

[0095] Optionally, the VXLAN packet may carry the ESI label based on a format shown in FIG. 5. As shown in FIG. 5, the VXLAN packet includes a VXLAN tunnel part and the original Ethernet payload, the VXLAN tunnel part includes an outer Eth header (Ethernet header), an outer IP header, an outer UDP header, and a VXLAN header, and the original Ethernet payload carries the BUM traffic. The VXLAN packet further includes the ESI label, and the ESI label is encapsulated between the VXLAN header in the VXLAN packet and the original Ethernet payload. In addition, the VXLAN header carries indication information, and the indication information is used to indicate that the VXLAN packet carries the ESI label.

[0096] S305. The second NVE device avoids forwarding the BUM traffic to the CE device by using an ES between the second CE device and the second NVE device,

based on that the VXLAN packet carries the ESI label.

[0097] The second NVE device parses the VXLAN packet to obtain the ESI label carried in the VXLAN packet, and may determine, based on the value of the ESI label, that the ESI label is the ESI label allocated by the second NVE device to the Ethernet link between the second CE device and the second NVE device. The second NVE device no longer transmits the BUM traffic to an outbound interface that connects the second NVE device and the ES 2, so as to avoid forwarding the BUM traffic to the CE device by using the ES between the second CE device and the second NVE device.

[0098] The foregoing method is explained and described by using the loop shown in FIG. 4 as an example. It should be understood that the foregoing method is not limited to resolving only the problem of the BUM traffic loop including the second CE device, the first NVE device, and the second NVE device, and can be used to resolve all loop problems generated in an EVPN VXLAN application scenario according to the foregoing method.

[0099] In the foregoing implementation, a format of the VXLAN packet is extended, so that the VXLAN packet can carry the ESI label, thereby resolving a loop problem generated in a process of forwarding the BUM traffic.

[0100] The EVPN VXLAN in this embodiment of this application may be further used to process a MAC route, so that after a VXLAN tunnel is established according to the method shown in FIG. 2A and FIG. 2B, the EVPN VXLAN may be used to guide forwarding of unicast traffic. A method for processing a MAC route by using the EVPN VXLAN is specifically described below with reference to FIG. 1. The second CE device is dual-homed to the first NVE device and the second NVE device by using two Ethernet links, and the two Ethernet links form an ES. ESI values of ESs 2 of the two Ethernet links that connect the second CE device to the first NVE device and the second NVE device are a same value.

[0101] S401. The first NVE device receives a Media Access Control/Internet Protocol advertisement (MAC/IP advertisement) route from the second NVE device, where the MAC/IP advertisement route includes a Media Access Control-virtual local area network identifier (MAC-VLAN ID), the MAC-VLAN ID is used to indicate a VLAN to which a MAC address carried in the MAC/IP advertisement route belongs, and the MAC address carried in the MAC/IP advertisement route is a MAC address of the CE device, or the MAC address carried in the MAC/IP advertisement route is a MAC address of a host in an Ethernet virtual private network (EVPN) site administered by the CE device.

[0102] Because the first NVE device and the second NVE device are a pair of BGP peers, the second NVE device can send the MAC/IP advertisement route to the first NVE device. Correspondingly, the first NVE device receives the MAC/IP advertisement route from the second NVE device. For a format and basic usage of the MAC/IP advertisement route, refer to descriptions of Chapter 7.2, Chapter 9, and Chapter 14 of the RFC 7432.

In this embodiment of this application, a MAC/IP advertisement route described in the RFC 7432 is extended, so that the MAC/IP advertisement route can carry the MAC-VLAN ID. The MAC-VLAN ID is used to indicate the VLAN to which the MAC address carried in the MAC/IP advertisement route belongs. The MAC address is used to identify the second CE device, or the MAC address is used to identify a host connected to the second CE device. The host is a physical device or a VM

[0103] Optionally, a new attribute is extended in this embodiment of this application to carry the MAC-VLAN ID. As shown in FIG. 8, a MAC-VLAN attribute is extended in this embodiment of this application to carry the MAC-VLAN ID. A type field and a sub type field are used to indicate a type of the MAC-VLAN attribute, and specific values of the type field and the sub type field each may be set according to a requirement of a standard organization. A flag field may be used to indicate whether the MAC-VLAN ID in the MAC-VLAN attribute is allowed to be read. For example, if the flag field is 0, it indicates that the MAC-VLAN ID in the MAC-VLAN attribute is not allowed to be read; or if the flag field is 1, it indicates that the MAC-VLAN ID in the MAC-VLAN attribute is allowed to be read. It should be understood that the foregoing functions of the flag field are an example, and a function of the flag field may correspondingly change based on a use scenario. The MAC-VLAN ID field is used to carry a value of the MAC-VLAN ID. A reserved field is used to implement further extension. The MAC-VLAN attribute may be carried in the MAC/IP advertisement route as an attribute, so as to implement the solution in this embodiment of this application.

[0104] S402. The first NVE device determines, based on an ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route, a local interface that is of the first NVE device and that connects to an Ethernet link between the second CE device and the first NVE device, where the ESI carried in the MAC/IP advertisement route is used to indicate an ES between the second CE device and the second NVE device.

[0105] The MAC/IP advertisement route sent by the second NVE device to the first NVE device includes an ESI field, and the ESI field is used to carry the ESI that indicates the Ethernet link between the second CE device and the second NVE device. An ESI value of an ES 2 between the second CE device and the first NVE device is equal to an ESI value of an ES 2 between the second CE device and the second NVE device. Therefore, the first NVE device can determine, based on the ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route, the local interface that is of the first NVE device and that connects the Ethernet link between the second CE device and the first NVE device. In other words, the first NVE device may determine a local interface on the first NVE device based on the ESI and the MAC-VLAN ID, and the local interface is used to communicate with a device identified by the MAC address carried in the MAC/IP advertisement route.

[0106] S403. A control plane of the first NVE device sends a first MAC entry to a forwarding plane of the first NVE device, where an outbound interface in the first MAC entry is the local interface of the first NVE device, and the first MAC entry includes the MAC address and the MAC-VLAN ID that are carried in the MAC/IP advertisement route.

[0107] The first NVE device delivers, from the control plane of the first NVE device to the forwarding plane of the first NVE device, the first MAC entry that includes the MAC address and the MAC-VLAN ID, so that the first NVE device can guide, based on the first MAC entry, forwarding of unicast traffic. The outbound interface in the first MAC entry is a local interface that is of the first NVE device and that connects to an Ethernet link between the second CE device and the first NVE device.

[0108] For example, as shown in FIG. 1, the PE device sends unicast traffic to the second CE device. It is assumed that the PE device hashes the unicast traffic to the second NVE device. The second NVE device receives a VXLAN packet from the PE device, and the VXLAN packet carries the unicast traffic. The second NVE device obtains the unicast traffic in the VXLAN packet, and forwards the unicast traffic to the second CE device based on a MAC entry stored in the second NVE device.

[0109] As shown in FIG. 9, if the ES 2 between the second NVE device and the second CE device is faulty, the unicast traffic cannot arrive at the second CE device by using the ES 2 between the second NVE device and the second CE device, as shown by a dashed line arrow in FIG. 9. The first NVE device and the second NVE device have performed S401 to S403. Therefore, the first NVE device has delivered, from the control plane of the first NVE device to the forwarding plane of the first NVE device, the first MAC entry that includes the MAC address and the MAC-VLAN ID. The second NVE device sends, to the first NVE device through the first VXLAN tunnel, the VXLAN packet that carries the unicast traffic. After receiving the VXLAN packet, the first NVE device obtains the unicast traffic, and forwards the unicast traffic to the second CE device based on the first MAC entry by using the ES 2 between the first NVE device and the second CE device.

[0110] In the foregoing implementation, a forwarding path may be redirected by using the MAC/IP advertisement route and a VXLAN tunnel between NVE devices, so that when a link fault occurs on an Ethernet link, normal forwarding of unicast traffic can still be ensured. The foregoing method is described by using an example in which a CE device is dual-homed to two NVE devices. It should be understood that the foregoing method may also be applied to a scenario in which a CE device is multi-homed to more than two NVE devices.

[0111] Optionally, the foregoing MAC route processing method further includes the following step:

S404. The MAC/IP advertisement route is carried in MP_REACH_NLRI, a next hop field in the MP_REACH_NLRI includes a common VTEP address,

the MAC/IP advertisement route further includes the second VTEP address, and the common VTEP address included in the next hop field in the MP_REACH_NLRI is the same as the common VTEP address included in the originating router's IP address field in the second IMET route. The method further includes the following: The control plane of the first NVE device sends a second MAC entry to the forwarding plane of the first NVE device based on the common VTEP address included in the next hop field in the MP_REACH_NLRI and the second VTEP address included in the MAC/IP advertisement route, where an outbound interface in the second MAC entry is a local interface that is of the first NVE device and that connects to the VXLAN tunnel from the first NVE device to the second NVE device, and the second MAC entry includes the MAC address carried in the MAC/IP advertisement route.

[0112] The second NVE device may send the MAC/IP advertisement route to the first NVE device in the following manner: The second NVE device may send an update message to the first NVE device, and the update message carries an MP_REACH_NLRI attribute. The MP_REACH_NLRI attribute includes an EVPN NLRI field and the next hop field. The EVPN NLRI field is used to carry the MAC/IP advertisement route. The next hop field includes the common VTEP address, and the MAC/IP advertisement route further includes the second VTEP address. The second NVE device may determine, based on the common VTEP address, that the first NVE device and the second NVE device belong to a same NVE device set. The second NVE device may determine, based on the second VTEP address, that the MAC/IP advertisement route is advertised by the second NVE device. The first NVE device sends the second MAC entry to the forwarding plane of the first NVE device. The outbound interface in the second MAC entry is the local interface that is of the first NVE device and that connects to the VXLAN tunnel from the first NVE device to the second NVE device, and the second MAC entry includes the MAC address carried in the MAC/IP advertisement route.

[0113] For example, as shown in FIG. 1, the PE device sends unicast traffic to the second CE device. It is assumed that the PE device hashes the unicast traffic to the first NVE device. The first NVE device receives a PE VXLAN packet from the PE device, and the PE VXLAN packet carries the unicast traffic. The first NVE device obtains the unicast traffic in the PE VXLAN packet, and forwards the unicast traffic to the second CE device based on a MAC entry stored in the first NVE device and by using the ES 2 between the first NVE device and the second CE device. In addition, the first NVE device has sent the second MAC entry to the forwarding plane of the first NVE device. The outbound interface in the second MAC entry is the local interface that is of the first NVE device and that connects to the VXLAN tunnel from the first NVE device to the second NVE device, and the second MAC entry includes the MAC address. Therefore,

the first NVE device may forward the unicast traffic to the second NVE device through the first VXLAN tunnel. The second NVE device forwards the unicast traffic to the second CE device based on the MAC entry stored in the second NVE device and by using the ES 2 between the second NVE device and the second CE device.

[0114] In the foregoing implementation, load sharing of the unicast traffic may be further implemented, thereby helping increase transmission bandwidth of the unicast traffic.

[0115] The foregoing MAC route processing is applicable to a scenario in which a CE device is connected to two NVE devices through dual homing, or a scenario in which a CE device is multi-homed to more than two NVE devices. A method for processing a MAC route by using the EVPN VXLAN used in a scenario in which a CE device is single-homed to one NVE device is specifically described below with reference to FIG. 1. As shown in FIG. 1, the third CE device is homed to the second NVE device.

[0116] S501. The first NVE device receives a MAC/IP advertisement route from the second NVE device, where the MAC/IP advertisement route is carried in MP_REACH_NLRI, a next hop field in the MP_REACH_NLRI includes a common VTEP address, the MAC/IP advertisement route includes a MAC-VLAN ID and the second VTEP address, the MAC-VLAN ID is used to indicate a VLAN to which a MAC address carried in the MAC/IP advertisement route belongs, and the MAC address carried in the MAC/IP advertisement route is a MAC address of the CE device, or the MAC address carried in the MAC/IP advertisement route is a MAC address of a host in an EVPN site administered by the CE device.

[0117] For an execution process in S501, refer to the corresponding explanations of S401 in the foregoing embodiment. Details are not described herein again.

[0118] S502. The first NVE device determines, based on an ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route, that the first NVE device has no local interface that connects to an Ethernet link between the CE device and the first NVE device, where the ESI carried in the MAC/IP advertisement route is used to indicate an ES between the CE device and the second NVE device.

[0119] With reference to the corresponding explanations of S402 in the foregoing embodiment, because the third CE device is connected to the second NVE device through single homing, an ESI value of an ES 3 between the second NVE device and the third CE device is 0 or an invalid value. After receiving the MAC/IP advertisement route from the second NVE device, the first NVE device obtains the ESI carried in the MAC/IP advertisement route. If the first NVE device finds that the ESI value is 0 or an invalid value, the first NVE device considers that the ESI does not exist. Therefore, the first NVE device determines, based on the ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route, that the first NVE device has no local interface that connects to the Ethernet link between the CE device and the

first NVE device.

[0120] S503. A control plane of the first NVE device sends a MAC entry to a forwarding plane of the first NVE device based on the common VTEP address included in the next hop field in the MP REACH NLRI and the second VTEP address included in the MAC/IP advertisement route. An outbound interface in the MAC entry is a local interface that is of the first NVE device and that connects to the VXLAN tunnel from the first NVE device to the second NVE device. The MAC entry includes the MAC address carried in the MAC/IP advertisement route.

[0121] The first NVE device performs S503 after determining that the first NVE device has no local interface that connects to the Ethernet link between the CE device and the first NVE device does not exist. For an execution process in S503, refer to the corresponding explanations of S404 in the foregoing embodiment. Details are not described herein again.

[0122] For example, as shown in FIG. 10, the PE device sends unicast traffic to the third CE device. It is assumed that the PE device hashes the unicast traffic to the first NVE device. The first NVE device receives a PE VXLAN packet from the PE device, and the PE VXLAN packet carries the unicast traffic. The first NVE device obtains the unicast traffic in the PE VXLAN packet, and the first NVE device locally does not find a MAC entry from the first NVE device to the third CE device. The first NVE device has performed S501 to S503, and therefore the first NVE device may forward the unicast traffic to the second NVE device through the first VXLAN tunnel. The second NVE device forwards the unicast traffic to the third CE device based on a MAC entry stored in the second NVE device and by using an ES 3 between the second NVE device and the third CE device.

[0123] In the foregoing implementation, it can be ensured that the CE device connected to the NVE device through single homing can normally receive the unicast traffic.

[0124] The EVPN VXLAN in this embodiment of this application may be further applied to a layer 3 forwarding scenario. Specifically, as shown in FIG. 1, virtual routing and forwarding (VRF) instances are separately deployed on the PE device, the first NVE device, the second NVE device, and the third NVE device. In this embodiment, for ease of description, the VRF instances deployed on the PE device, the first NVE device, the second NVE device, and the third NVE device are set to a same VRF instance. It should be understood that a plurality of VRF instances may be configured on each of the devices, for example, different VRF instances are corresponding to different services. The first NVE device, the second NVE device, and the third NVE device advertise Internet Protocol prefix (IP Prefix) routes to respective BGP peers of the devices. The IP prefix route includes an RD field, an ESI field, an Ethernet tag ID field, an IP prefix length field, an IP prefix field, a gateway Internet Protocol address (GW IP Address) field, and an MPLS label field. Definitions of the foregoing fields are similar to the definitions

in the foregoing embodiments of this application. The IP prefix field is used to indicate an IP address of a CE device or of a host connected to a CE device. The MPLS label field is used to indicate a layer 3 VNI, and the layer 3 VNI is corresponding to a VRF instance. The GW IP address field is set to 0, that is, an invalid value.

[0125] The NVE device may send the IP prefix route to a BGP peer of the NVE device in the following manner. The NVE device may send an update message to the BGP peer of the NVE device, and the update message carries an MP_REACH_NLRI attribute. The MP_REACH_NLRI attribute includes an EVPN NLRI field and a next hop field. The EVPN NLRI field is used to carry the IP prefix route.

[0126] In this embodiment, the next hop field in the MP_REACH_NLRI attribute that carries the IP prefix route and that is advertised by each of the first NVE device, the second NVE device, and the third NVE device includes the common VTEP address in the foregoing embodiment. As an extension of the IP prefix route, the IP prefix route further carries a VTEP address of an NVE device that sends the IP prefix route. Further, optionally, the IP prefix route may carry the VXLAN attribute described in the foregoing embodiment, and the VXLAN attribute carries a VTEP address. In a process of processing the IP prefix route, an IP route can be learned of and an IP routing entry can be delivered to a forwarding plane, so as to guide layer 3 forwarding of data traffic. The IP prefix route processing process and the layer 3 traffic forwarding process in this embodiment each are similar to the MAC route processing method in the foregoing embodiment. Details are not described herein again.

[0127] FIG. 11 is a schematic structural diagram of a first NVE device 1000 according to an embodiment of this application. The first NVE device 1000 shown in FIG. 11 may perform corresponding steps performed by the first NVE device in the method in the foregoing embodiment. As shown in FIG. 11, the first NVE device 1000 includes a receiving unit 1002 and a processing unit 1004.

[0128] The receiving unit 1002 is configured to receive a second IMET route from the second NVE device. An originating router's IP address field in the second IMET route includes a common VTEP address, the second IMET route further includes a second VTEP address, and the common VTEP address included in the originating router's IP address field in the second IMET route is different from the second VTEP address.

[0129] The processing unit 1004 is configured to determine whether the common VTEP address included in the originating router's IP address field in the second IMET route is the same as a common VTEP address stored in the first NVE device. The first NVE device can send, to a BGP peer of the first NVE device, an IMET route whose originating router's IP address field carries the common VTEP address stored in the first NVE device. The common VTEP address included in the originating router's IP address field in the second IMET route is the

same as the common VTEP address stored in the first NVE device.

[0130] When the processing unit 1004 determines that the common VTEP address included in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device, the processing unit 1004 is further configured to establish a VXLAN tunnel from the first NVE device to the second NVE device based on the second VTEP address in the second IMET route. Further, optionally, a VXLAN attribute in the second IMET route carries the second VTEP address.

[0131] Optionally, the first NVE device further includes a sending unit 1006. The sending unit 1006 is configured to send a first IMET route to the second NVE device. An originating router's IP address field in the first IMET route carries the common VTEP address stored in the first NVE device. The first IMET route further includes a first VTEP address, and the first VTEP address is used by the second NVE device to establish the VXLAN tunnel from the second NVE device to the first NVE device. The common VTEP address stored in the first NVE device is different from the first VTEP address. Optionally, a VXLAN attribute in the first IMET route carries the first VTEP address.

[0132] Optionally, the sending unit 1006 is further configured to send the first IMET route to a PE device. The receiving unit 1002 is further configured to receive an IMET route from the PE device. An originating router's IP address field in the IMET route from the PE device carries a VTEP address of the PE device. The processing unit 1004 is further configured to establish a VXLAN tunnel from the first NVE device to the PE device based on the VTEP address of the PE device that is in the IMET route from the PE device. The VTEP address of the PE device is not the same as the common VTEP address stored in the first NVE device.

[0133] Optionally, the receiving unit 1002 is further configured to receive a third IMET route from the third NVE device. An originating router's IP address field in the third IMET route includes a common VTEP address. The third IMET route further includes a third VTEP address. The common VTEP address included in the originating router's IP address field in the third IMET route is different from the third VTEP address. The processing unit 1004 is further configured to determine whether the common VTEP address included in the originating router's IP address field in the third IMET route is the same as the common VTEP address stored in the first NVE device. The common VTEP address included in the originating router's IP address field in the third IMET route is the same as the common VTEP address stored in the first NVE device. When the processing unit 1004 determines that the common VTEP address included in the originating router's IP address field in the third IMET route is the same as the common VTEP address stored in the first NVE device, the processing unit 1004 is further configured to establish a VXLAN tunnel from the first NVE

device to the third NVE device based on the third VTEP address in the third IMET route. Optionally, a VXLAN attribute in the third IMET route carries the third VTEP address.

[0134] Optionally, a CE device is multi-homed to the first NVE device, the second NVE device, and the third NVE device by using a plurality of Ethernet links, and the plurality of Ethernet links form an ES. The processing unit 1004 is further configured to determine the first NVE device as a DF in the ES based on ES routes from the second NVE device and the third NVE device.

[0135] Optionally, the receiving unit 1002 is further configured to receive, through the VXLAN tunnel from the second NVE device to the first NVE device, BUM traffic sent by the second NVE device. The processing unit 1004 is further configured to avoid forwarding the BUM traffic to the third NVE device through the VXLAN tunnel from the first NVE device to the third NVE device, based on that the BUM traffic is received by the first NVE device through the VXLAN tunnel from the second NVE device to the first NVE device.

[0136] Optionally, a CE device is dual-homed to the first NVE device and the second NVE device by using a plurality of Ethernet links, and the plurality of Ethernet links form an ES. The processing unit 1004 is further configured to determine the first NVE device as a DF in the ES based on an ES route from the second NVE device. The sending unit 1006 is further configured to send an Ethernet A-D per ES route to the second NVE device. The Ethernet A-D per ES route carries an ESI label allocated by the first NVE device to an Ethernet link between the CE device and the first NVE device. The receiving unit 1002 is further configured to receive, through the VXLAN tunnel from the second NVE device to the first NVE device, a VXLAN packet sent by the second NVE device. An original Ethernet payload in the VXLAN packet includes BUM traffic, the VXLAN packet carries the ESI label, and the BUM traffic is from the CE device. The processing unit 1004 is further configured to avoid, based on that the VXLAN packet carries the ESI label, forwarding the BUM traffic to the CE device by using the ES between the CE device and the first NVE device. Further, optionally, the ESI label is encapsulated between a VXLAN header in the VXLAN packet and the original Ethernet payload, the VXLAN header carries indication information, and the indication information is used to indicate that the VXLAN packet carries the ESI label.

[0137] Optionally, a CE device is dual-homed to the first NVE device and the second NVE device by using a plurality of Ethernet links, and the plurality of Ethernet links form an ES. The receiving unit 1002 is further configured to receive a MAC/IP advertisement route from the second NVE device. The MAC/IP advertisement route includes a MAC-VLAN ID, and the MAC-VLAN ID is used to indicate a VLAN to which a MAC address carried in the MAC/IP advertisement route belongs. The MAC address carried in the MAC/IP advertisement route is a MAC address of the CE device, or the MAC address

carried in the MAC/IP advertisement route is a MAC address of a host in an Ethernet virtual private network EVPN site administered by the CE device. The processing unit 1004 is further configured to determine, based on an ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route, a local interface that is of the first NVE device and that connects to an Ethernet link between the CE device and the first NVE device. The ESI carried in the MAC/IP advertisement route is used to indicate the ES between the CE device and the second NVE device. The processing unit 1004 is further configured to trigger sending of a first MAC entry to a forwarding plane of the first NVE device. An outbound interface in the first MAC entry is the local interface of the first NVE device, and the first MAC entry includes a MAC address and the MAC-VLAN ID that are carried in the MAC/IP advertisement route.

[0138] Optionally, the MAC/IP advertisement route is carried in MP_REACH_NLRI. A next hop field in the MP_REACH_NLRI includes a common VTEP address. The MAC/IP advertisement route further includes the second VTEP address, and the common VTEP address included in the next hop field in the MP_REACH_NLRI is the same as the common VTEP address included in the originating router's IP address field in the second IMET route. The processing unit 1004 is further configured to trigger sending of a second MAC entry to the forwarding plane of the first NVE device based on the common VTEP address included in the next hop field in the MP_REACH_NLRI and the second VTEP address included in the MAC/IP advertisement route. An outbound interface in the second MAC entry is a local interface that is of the first NVE device and that connects to the VXLAN tunnel from the first NVE device to the second NVE device, and the second MAC entry includes the MAC address carried in the MAC/IP advertisement route.

[0139] Optionally, a CE device is connected to the second NVE device. The receiving unit 1002 is further configured to receive a MAC/IP advertisement route from the second NVE device. The MAC/IP advertisement route is carried in MP_REACH_NLRI, a next hop field in the MP_REACH_NLRI includes a common VTEP address, the MAC/IP advertisement route includes a MAC-VLAN ID and the second VTEP address, the MAC-VLAN ID is used to indicate a VLAN to which a MAC address carried in the MAC/IP advertisement route belongs, and the MAC address carried in the MAC/IP advertisement route is a MAC address of the CE device, or the MAC address carried in the MAC/IP advertisement route is a MAC address of a host in an EVPN site administered by the CE device. The processing unit 1004 is further configured to determine, based on an ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route, that the first NVE device has no local interface that connects to an Ethernet link between the CE device and the first NVE device. The ESI carried in the MAC/IP advertisement route is used to indicate an ES between the CE device and the second NVE device. The processing

unit 1004 is further configured to trigger sending of a MAC entry to a forwarding plane of the first NVE device based on the common VTEP address included in the next hop field in the MP_REACH_NLRI and the second VTEP address included in the MAC/IP advertisement route. An outbound interface in the MAC entry is a local interface that is of the first NVE device and that connects to the VXLAN tunnel from the first NVE device to the second NVE device. The MAC entry includes the MAC address carried in the MAC/IP advertisement route.

[0140] The first NVE device shown in FIG. 11 may perform corresponding steps performed by the first NVE device in the method in the foregoing embodiment. This ensures that an EVPN application scenario is not limited by a restriction condition that a physical direct link needs to be used as a link between NVE devices, and helps extend the EVPN application scenario and also helps improve reliability and reduce deployment complexity. It should be understood that the structure in FIG. 11 is also applicable to the second NVE device and the third NVE device in FIG. 1.

[0141] FIG. 12 is a schematic diagram of a hardware structure of a first NVE device 1100 according to an embodiment of this application. The first NVE device 1100 shown in FIG. 12 may perform corresponding steps performed by the first NVE device in the method in the foregoing embodiment.

[0142] As shown in FIG. 12, the first NVE device 1100 includes a processor 1101, a memory 1102, an interface 1103, and a bus 1104. The interface 1103 may be implemented in a wireless or wired manner. Specifically, the interface 1103 may be a network adapter. The processor 1101, the memory 1102, and the interface 1103 are connected by using the bus 1104.

[0143] The interface 1103 may specifically include a transmitter and a receiver, configured for information transmission or receiving between the first NVE device and each of the second NVE device and the third NVE device in the foregoing embodiment; or information transmission or receiving between the first NVE device and a PE device; or information transmission or receiving between the first NVE device and a CE device connected to the first NVE device. For example, the interface 1103 is configured to support processes S102 and S105 in FIG. 2A and FIG. 2B. The processor 1101 is configured to perform processing performed by the first NVE device in the foregoing embodiment. For example, the processor 1101 is configured to: determine whether a common VTEP address included in an originating router's IP address field in a received IMET route is the same as a common VTEP address stored in the memory 1102; establish a VXLAN tunnel based on the VTEP address in the received IMET route and based on that the common VTEP address in the originating router's IP address field in the received IMET route is the same as the common VTEP address stored in the memory 1102; process a received ES route, MAC/IP advertisement route, or Ethernet A-D per ES route; and process BUM traffic or uni-

cast traffic, and/or another process of the technology described in this specification. For example, the processor 1101 is configured to support processes S103 and S104 in FIG. 2A and FIG. 2B. The memory 1102 includes an operating system 11021 and an application program 11022, and is configured to store a program, code, or an instruction. When executing the program, the code, or the instruction, the processor or a hardware device may complete a processing process related to the first NVE device in the method embodiment. Optionally, the memory 1102 may include a read-only memory (ROM) and a random access memory (RAM). The ROM includes a basic input/output system (BIOS) or an embedded system. The RAM includes an application program and an operating system. When the first NVE device 1100 needs to be run, the BIOS or a bootloader in the embedded system that is built into the ROM is used to lead a system to start, and lead the first NVE device 1100 to enter a normal running state. After entering the normal running state, the first NVE device 1100 runs the application program and the operating system in the RAM, so as to complete a processing process related to the first NVE device in the method embodiment.

[0144] It may be understood that FIG. 12 merely shows a simplified design of the first NVE device. In actual application, the first NVE device may include any quantity of interfaces, processors, or memories. In addition, only the first NVE device is used as an example for a description of this embodiment. It should be understood that the second NVE device, the third NVE device, or more NVE devices have same functions as the first NVE device. Details are not described herein again.

[0145] FIG. 13 is a schematic diagram of a hardware structure of another first NVE device 1200 according to an embodiment of this application. The first NVE device 1200 shown in FIG. 13 may perform corresponding steps performed by the first NVE device in the method in the foregoing embodiment.

[0146] As shown in FIG. 13, the first NVE device 1200 includes a main control board 1210, an interface board 1230, a switching board 1220, and an interface board 1240. The main control board 1210 is configured to complete functions such as system management, device maintenance, and protocol processing. The switching board 1220 is configured to complete data exchange between interface boards (the interface board is also referred to as a line card or a service board). The interface boards 1230 and 1240 are configured to: provide various service interfaces (for example, a POS interface, a GE interface, and an ATM interface), and forward a data packet. The main control board 1210, the interface boards 1230 and 1240, and the switching board 1220 are connected to a system backplane by using a system bus to communicate with each other. A central processing unit 1231 on the interface board 1230 is configured to: control and manage the interface board, and communicate with a central processing unit on the main control board.

[0147] A physical interface card 1233 on the interface board 1230 receives a second IMET route from the second NVE device, and sends the second IMET route to a central processing unit 1211 on the main control board 1210 by using the central processing unit 1231 on the interface board 1230.

[0148] The central processing unit 1211 on the main control board 1210 is configured to obtain the second IMET route. The central processing unit 1211 is further configured to determine whether a common VTEP address included in an originating router's IP address field in the second IMET route is the same as a common VTEP address stored in the first NVE device. When the central processing unit 1211 determines that the common VTEP address included in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device, the central processing unit 1211 establishes a virtual extensible local area network VXLAN tunnel from the first NVE device to the second NVE device based on the second VTEP address in the second IMET route. The common VTEP address included in the originating router's IP address field in the second IMET route is different from the second VTEP address.

[0149] The central processing unit 1211 on the main control board 1210 is further configured to generate a first IMET route. An originating router's IP address field in the first IMET route carries the common VTEP address stored in the first NVE device. The first IMET route further includes a first VTEP address. The first VTEP address is used by the second NVE device to establish the VXLAN tunnel from the second NVE device to the first NVE device. The common VTEP address stored in the first NVE device is different from the first VTEP address. The central processing unit 1211 on the main control board 1210 sends the generated first IMET route to the physical interface card 1233 by using the central processing unit 1231 on the interface board 1230. The physical interface card 1233 on the interface board 1230 sends the first IMET route to a BGP peer of the first NVE device.

[0150] The central processing unit 1211 on the main control board 1210 is further configured to obtain, from the physical interface card 1233 on the interface board 1230, ES routes from the second NVE device and the third NVE device. The central processing unit 1211 is further configured to determine the first NVE device as a DF or an NDF in an ES based on the ES routes.

[0151] The central processing unit 1211 on the main control board 1210 is further configured to control the interface board 1230 to forward and process BUM traffic or unicast traffic.

[0152] The central processing unit 1211 on the main control board 1210 is further configured to: generate an Ethernet A-D per ES route, and send the Ethernet A-D per ES route to the BGP peer of the first NVE device by using the physical interface card 1233 on the interface board 1230.

[0153] The central processing unit 1211 on the main

control board 1210 is further configured to: obtain, from the physical interface card 1233 on the interface board 1230, a MAC/IP advertisement route from the second NVE device, and generate a MAC entry based on an ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route. The central processing unit 1211 on the main control board 1210 transmits the MAC entry to a forwarding entry memory 1234 on the interface board 1230 by using the central processing unit 1231 on the interface board 1230.

[0154] The forwarding entry memory 1234 on the interface board 1230 is configured to store the MAC entry. The central processing unit 1231 on the interface board 1230 is configured to control a network memory 1232 to obtain the MAC entry in the forwarding entry memory 1234. In addition, the central processing unit 1231 is configured to control the network memory 1232 to receive and send traffic by using the physical interface card 1233.

[0155] It should be understood that an operation on the interface board 1240 is consistent with an operation on the interface board 1230 in this embodiment of this application. For brevity, details are not described again. It should be understood that the first NVE device 1200 in this embodiment may correspondingly implement functions of the first NVE device and/or steps performed by the first NVE device in the foregoing method embodiment. For brevity, details are not described herein again. In addition, only the first NVE device is used as an example for a description of this embodiment. It should be understood that the second NVE device, the third NVE device, or more NVE devices have same functions as the first NVE device. Details are not described herein again.

[0156] In addition, it should be noted that there may be one or more main control boards. When there are a plurality of main control boards, a primary main control board and a secondary main control board may be included. There may be one or more interface boards, and a stronger data processing capability of the first NVE device indicates a larger quantity of provided interface boards. There may be one or more physical interface cards on the interface board. There may be no switching board, or there may be one or more switching boards. When there are a plurality of switching boards, load sharing and redundancy backup may be jointly implemented. In a centralized forwarding architecture, the first NVE device may not need a switching board, and the interface board bears a service data processing function in an entire system. In a distributed forwarding architecture, the first NVE device may have at least one switching board, and data is exchanged between a plurality of interface boards by using the switching board, so as to provide large-capacity data exchanging and processing capability. Therefore, the first NVE device has a better data access and processing capability in the distributed architecture than in the centralized architecture. A specific architecture to be used depends on a specific networking deployment scenario. This is not limited herein.

[0157] FIG. 14 is a schematic diagram of a hardware structure of still another first NVE device 1300 according to an embodiment of this application. The first NVE device 1300 shown in FIG. 14 may perform corresponding steps performed by the first NVE device in the method in the foregoing embodiment.

[0158] This product form of the first NVE device 1300 is applicable to a network architecture (for example, software-defined networking (SDN)) in which control and forwarding are separated. In the SDN, the main control board 1210 of the first NVE device 1200 shown in FIG. 13 is separated from the device, and forms a new independent physical device (that is, a controller 1210A shown in FIG. 14), and remaining components form another independent physical device (that is, a first NVE forwarding device 1200A shown in FIG. 14). The controller 1210A interacts with the first NVE forwarding device 1200A by using a control channel protocol. The control channel protocol may be the OpenFlow protocol, the Path Computation Element Communication Protocol (PCEP), the BGP, the Interface to the Routing System (I2RS), or the like. In other words, compared with the embodiment corresponding to FIG. 13, the first NVE device 1300 in this embodiment includes the controller 1210A separated from the device and the first NVE forwarding device 1200A. That is, in this embodiment, the first NVE device 1300 may also be considered as a system.

[0159] The controller 1210A may be implemented based on a general purpose physical server or a dedicated hardware structure. In a design example, the controller includes a receiver, a processor, a transmitter, a RAM, a ROM, and a bus (not shown in the figure). The processor is separately coupled to the receiver, the transmitter, the RAM, and the ROM by using the bus. When the controller needs to be run, a BIOS or a bootloader in an embedded system that is built into the ROM is used to lead a system to start, and lead the controller to enter a normal running state. After entering the normal running state, the controller runs an application program and an operating system in the RAM, so that the processor performs all functions and steps of the main control board 1210 in FIG. 13.

[0160] The first NVE forwarding device 1200A may be implemented based on a dedicated hardware structure. Functions and structures of the first NVE forwarding device 1200A remain the same as functions and structures of the interface board 1230, the interface board 1240, and the switching board 1220 in FIG. 13, so as to perform a corresponding function and step. Alternatively, the first NVE forwarding device 1200A may be a virtual first NVE forwarding device implemented based on a general purpose physical server and a network functions virtualization (NFV) technology, and the virtual first NVE forwarding device is a virtual router. In a scenario of the virtual first NVE forwarding device, the interface board, the switching board, and the processor included in the first NVE forwarding device mentioned in the embodiment of the physical first NVE forwarding device may be consid-

ered, in a virtual environment, as an interface resource, a network resource, and a processing resource that are allocated by the first NVE forwarding device to the virtual first NVE forwarding device based on the general purpose physical server. For a specific implementation of implementing functions or steps of the first forwarding NVE device by using the general physical server, or implementing functions or steps of the first forwarding NVE device by using the general physical server and the NFV technology, refer to the embodiment in FIG. 12.

[0161] It should be understood that in this embodiment, the controller 1210A and the first NVE forwarding device 1200A in the first NVE device 1300 may implement various functions and steps implemented by the first NVE device in the method embodiment. For brevity, details are not described herein again. In addition, only the first NVE device is used as an example for a description of this embodiment. It should be understood that the second NVE device, the third NVE device, or more NVE devices have same functions as the first NVE device. Details are not described herein again.

[0162] In addition, an embodiment of this application provides a computer storage medium, configured to store a computer software instruction used by the first NVE device, and the computer software instruction includes a program designed for performing the foregoing method embodiment.

[0163] As shown in FIG. 1, an embodiment of this application further includes a network system for processing a route. The network system includes at least two network virtualization edge NVE devices, and each of the at least two NVE devices is the first NVE device in FIG. 11, FIG. 12, FIG. 13, or FIG. 14.

[0164] Method or algorithm steps described in combination with the content disclosed in this application may be implemented by hardware, or may be implemented by a processor by executing a software instruction. The software instruction may include a corresponding software module. The software module may be located in a RAM memory, a flash memory, a ROM memory, an EPROM memory, an EEPROM memory, a register, a hard disk, a removable hard disk, a CD-ROM, or a storage medium of any other form known in the art. For example, a storage medium is coupled to a processor, so that the processor can read information from the storage medium or write information into the storage medium. Certainly, the storage medium may be a component of the processor. The processor and the storage medium may be located in the ASIC. In addition, the ASIC may be located in user equipment. Certainly, the processor and the storage medium may exist in the user equipment as discrete components.

[0165] A person skilled in the art should be aware that in the foregoing one or more examples, functions described in this application may be implemented by hardware, software, firmware, or any combination thereof. When the present invention is implemented by software, the foregoing functions may be stored in a computer-

readable medium or transmitted as one or more instructions or code in the computer-readable medium. The computer-readable medium includes a computer storage medium and a communications medium, where the communications medium includes any medium that enables a computer program to be transmitted from one place to another. The storage medium may be any available medium accessible to a general-purpose or dedicated computer.

Claims

1. A route processing method, wherein the method is applied to a Virtual Extensible Local Area Network, VXLAN, having an Ethernet Virtual Private Network, EVPN, control plane, the VXLAN comprises a first network virtualization edge, NVE, device and a second NVE device, and the method comprises:

receiving (S102), by the first NVE device, a second inclusive multicast Ethernet tag, IMET, route from the second NVE device, wherein an originating router's IP address field in the second IMET route comprises a common virtual extensible local area network tunnel endpoint, VTEP, address, the second IMET route further comprises a second VTEP address, and the common VTEP address comprised in the originating router's IP address field in the second IMET route is different from the second VTEP address, the common VTEP address comprised in the originating router's IP address field in the second IMET route indicates a common VTEP included in the second NVE device, the second VTEP address indicates a second VTEP included in the second NVE device;

determining (S103), by the first NVE device, whether the common VTEP address comprised in the originating router's IP address field in the second IMET route is the same as a common VTEP address stored in the first NVE device, wherein the first NVE device and the second NVE device are a pair of Border Gateway Protocol peers, BGP peers, for exchanging IMET routes, the common VTEP address stored in the first NVE device indicates a common VTEP included in the first NVE device; and

when the first NVE device determines that the common VTEP address comprised in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device, establishing (S104), by the first NVE device, a VXLAN tunnel from the first NVE device to the second NVE device based on the second VTEP address in the second IMET route.

2. The method according to claim 1, wherein the method further comprises:

sending (S105), by the first NVE device, a first IMET route to the second NVE device, wherein an originating router's IP address field in the first IMET route carries the common VTEP address stored in the first NVE device, the first IMET route further comprises a first VTEP address, the first VTEP address is used by the second NVE device to establish the VXLAN tunnel from the second NVE device to the first NVE device, the first VTEP address indicates a first VTEP included in the first NVE device, and the common VTEP address stored in the first NVE device is different from the first VTEP address.

3. The method according to claim 2, wherein the VXLAN further comprises a provider edge, PE, device, and the method further comprises:

sending, by the first NVE device, the first IMET route to the PE device;
receiving, by the first NVE device, an IMET route from the PE device, wherein an originating router's IP address field in the IMET route from the PE device carries a VTEP address of the PE device; and
establishing, by the first NVE device, a VXLAN tunnel from the first NVE device to the PE device based on the VTEP address of the PE device that is in the IMET route from the PE device, wherein the VTEP address of the PE device is not the same as the common VTEP address stored in the first NVE device.

4. The method according to claim 2 or 3, wherein

a VXLAN attribute in the first IMET route carries the first VTEP address; and
a VXLAN attribute in the second IMET route carries the second VTEP address.

5. The method according to any one of claims 1 to 4, wherein the VXLAN further comprises a third NVE device, and the method further comprises:

receiving, by the first NVE device, a third IMET route from the third NVE device, wherein an originating router's IP address field in the third IMET route comprises a common VTEP address, the third IMET route further comprises a third VTEP address, and the common VTEP address comprised in the originating router's IP address field in the third IMET route is different from the third VTEP address;
determining, by the first NVE device, whether the common VTEP address comprised in the originating router's IP address field in the third IMET route is the same as the common VTEP

address stored in the first NVE device, wherein the common VTEP address comprised in the originating router's IP address field in the third IMET route is the same as the common VTEP address stored in the first NVE device; and
when the first NVE device determines that the common VTEP address comprised in the originating router's IP address field in the third IMET route is the same as the common VTEP address stored in the first NVE device, establishing, by the first NVE device, a VXLAN tunnel from the first NVE device to the third NVE device based on the third VTEP address in the third IMET route.

6. The method according to claim 5, wherein a customer edge, CE, device is multi-homed to the first NVE device, the second NVE device, and the third NVE device by using a plurality of Ethernet links, and the plurality of Ethernet links form an Ethernet segment, ES; and
the method further comprises:
determining, by the first NVE device, the first NVE device as a designated forwarder, DF, in the ES based on ES routes from the second NVE device and the third NVE device.

7. The method according to claim 5 or 6, wherein the method further comprises:

receiving, by the first NVE device through the VXLAN tunnel from the second NVE device to the first NVE device, broadcast, unknown unicast, and multicast, BUM, traffic sent by the second NVE device; and
avoiding, by the first NVE device, forwarding the BUM traffic to the third NVE device through the VXLAN tunnel from the first NVE device to the third NVE device, based on that the BUM traffic is received by the first NVE device through the VXLAN tunnel from the second NVE device to the first NVE device.

8. The method according to any one of claims 2 to 4, wherein a CE device is dual-homed to the first NVE device and the second NVE device by using a plurality of Ethernet links, and the plurality of Ethernet links form an ES; and
the method further comprises:

determining, by the first NVE device, the first NVE device as a DF in the ES based on an ES route from the second NVE device;
sending, by the first NVE device, an Ethernet Auto-Discovery per Ethernet segment, Ethernet A-D per ES, route to the second NVE device, wherein the Ethernet A-D per ES route carries an Ethernet segment identifier, ESI, label allo-

cated by the first NVE device to an Ethernet link between the CE device and the first NVE device; receiving, by the first NVE device through the VXLAN tunnel from the second NVE device to the first NVE device, a VXLAN packet sent by the second NVE device, wherein an original Ethernet payload in the VXLAN packet comprises BUM traffic, the VXLAN packet carries the ESI label, and the BUM traffic is from the CE device; and
 avoiding, by the first NVE device based on that the VXLAN packet carries the ESI label, forwarding the BUM traffic to the CE device by using the ES between the CE device and the first NVE device.

9. The method according to claim 8, wherein the ESI label is encapsulated between a VXLAN header in the VXLAN packet and the original Ethernet payload, the VXLAN header carries indication information, and the indication information is used to indicate that the VXLAN packet carries the ESI label.
10. The method according to any one of claims 1 to 4, wherein a CE device is dual-homed to the first NVE device and the second NVE device by using a plurality of Ethernet links, and the plurality of Ethernet links form an ES; and
 the method further comprises:

receiving, by the first NVE device, a Media Access Control/Internet Protocol advertisement, MAC/IP advertisement, route from the second NVE device, wherein the MAC/IP advertisement route comprises a Media Access Control-virtual local area network identifier, MAC-VLAN ID, the MAC-VLAN ID is used to indicate a VLAN to which a MAC address carried in the MAC/IP advertisement route belongs, and the MAC address carried in the MAC/IP advertisement route is a MAC address of the CE device, or the MAC address carried in the MAC/IP advertisement route is a MAC address of a host in an Ethernet virtual private network, EVPN, site administered by the CE device;
 determining, by the first NVE device based on an ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route, a local interface that is of the first NVE device and that connects to an Ethernet link between the CE device and the first NVE device, wherein the ESI carried in the MAC/IP advertisement route is used to indicate the ES between the CE device and the second NVE device; and
 sending, by a control plane of the first NVE device, a first MAC entry to a forwarding plane of the first NVE device, wherein an outbound interface in the first MAC entry is the local interface

of the first NVE device, and the first MAC entry comprises the MAC address and the MAC-VLAN ID that are carried in the MAC/IP advertisement route.

11. The method according to claim 10, wherein the MAC/IP advertisement route is carried in multiprotocol reachable network layer reachability information, MP_REACH_NLRI, a next hop field in the MP_REACH_NLRI comprises a common VTEP address, the MAC/IP advertisement route further comprises the second VTEP address, and the common VTEP address comprised in the next hop field in the MP_REACH_NLRI is the same as the common VTEP address comprised in the originating router's IP address field in the second IMET route; and
 the method further comprises:
 sending, by the control plane of the first NVE device, a second MAC entry to the forwarding plane of the first NVE device based on the common VTEP address comprised in the next hop field in the MP_REACH_NLRI and the second VTEP address comprised in the MAC/IP advertisement route, wherein an outbound interface in the second MAC entry is a local interface that is of the first NVE device and that connects to the VXLAN tunnel from the first NVE device to the second NVE device, and the second MAC entry comprises the MAC address carried in the MAC/IP advertisement route.
12. The method according to any one of claims 1 to 4, wherein a CE device is connected to the second NVE device; and
 the method further comprises:

receiving, by the first NVE device, a MAC/IP advertisement route from the second NVE device, wherein the MAC/IP advertisement route is carried in MP_REACH_NLRI, a next hop field in the MP_REACH_NLRI comprises a common VTEP address, the MAC/IP advertisement route comprises a MAC-VLAN ID and the second VTEP address, the MAC-VLAN ID is used to indicate a VLAN to which a MAC address carried in the MAC/IP advertisement route belongs, and the MAC address carried in the MAC/IP advertisement route is a MAC address of the CE device, or the MAC address carried in the MAC/IP advertisement route is a MAC address of a host in an EVPN site administered by the CE device;
 determining, by the first NVE device based on an ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route, that the first NVE device has no local interface that connects to an Ethernet link between the CE device and the first NVE device, wherein the ESI carried in the MAC/IP advertisement route is used to indicate an ES between the CE device and the sec-

ond NVE device; and
 sending, by a control plane of the first NVE device, a MAC entry to a forwarding plane of the first NVE device based on the common VTEP address comprised in the next hop field in the MP_REACH_NLRI and the second VTEP address comprised in the MAC/IP advertisement route, wherein an outbound interface in the MAC entry is a local interface that is of the first NVE device and that connects to the VXLAN tunnel from the first NVE device to the second NVE device, and the MAC entry comprises the MAC address carried in the MAC/IP advertisement route.

13. A first network virtualization edge, NVE, device, wherein the first NVE is applied to a Virtual Extensible Local Area Network, VXLAN, having an Ethernet Virtual Private Network, EVPN, control plane, the VXLAN further comprises a second NVE device, and the first NVE device comprises:

a receiving unit (1002), configured to receive a second inclusive multicast Ethernet tag, IMET, route from the
 second NVE device, wherein an originating router's IP address field in the second IMET route comprises a common virtual extensible local area network tunnel endpoint, VTEP, address, the second IMET route further comprises a second VTEP address, and the common VTEP address comprised in the originating router's IP address field in the second IMET route is different from the second VTEP address, the common VTEP address comprised in the originating router's IP address field in the second IMET route indicates a common VTEP included in the second NVE device, the second VTEP address indicates a second VTEP included in the second NVE device; and
 a processing unit (1004), configured to determine whether the common VTEP address comprised in the originating router's IP address field in the second IMET route is the same as a common VTEP address stored in the first NVE device, wherein the first NVE device and the second NVE device are a pair of Border Gateway Protocol peers, BGP peers, for exchanging IMET routes, the common VTEP address stored in the first NVE device indicates a common VTEP included in the first NVE device; wherein when the processing unit (1004) determines that the common VTEP address comprised in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device, the processing unit (1004) is further configured to establish a VXLAN tunnel from the first NVE de-

vice to the second NVE device based on the second VTEP address in the second IMET route.

14. The first NVE device according to claim 13, wherein the first NVE device further comprises:
 a sending unit (1006), configured to send a first IMET route to the second NVE device, wherein an originating router's IP address field in the first IMET route carries the common VTEP address stored in the first NVE device, the first IMET route further comprises a first VTEP address, the first VTEP address is used by the second NVE device to establish the VXLAN tunnel from the second NVE device to the first NVE device, the first VTEP address indicates a first VTEP included in the first NVE device, and the common VTEP address stored in the first NVE device is different from the first VTEP address.

15. The first NVE device according to claim 14, wherein
 the sending unit is further configured to send the first IMET route to a provider edge, PE, device; the receiving unit (1002) is further configured to receive an IMET route from the PE device, wherein an originating router's IP address field in the IMET route from the PE device carries a VTEP address of the PE device; and
 the processing unit (1004) is further configured to establish a VXLAN tunnel from the first NVE device to the PE device based on the VTEP address of the PE device that is in the IMET route from the PE device, wherein the VTEP address of the PE device is not the same as the common VTEP address stored in the first NVE device.

16. The first NVE device according to any one of claims 13 to 15, wherein

the receiving unit (1002) is further configured to receive a third IMET route from a third NVE device, wherein an originating router's IP address field in the third IMET route comprises a common VTEP address, the third IMET route further comprises a third VTEP address, and the common VTEP address comprised in the originating router's IP address field in the third IMET route is different from the third VTEP address;
 the processing unit (1004) is further configured to determine whether the common VTEP address comprised in the originating router's IP address field in the third IMET route is the same as the common VTEP address stored in the first NVE device, wherein the common VTEP address comprised in the originating router's IP address field in the third IMET route is the same as the common VTEP address stored in the first NVE device; and

when the processing unit (1004) determines that the common VTEP address comprised in the originating router's IP address field in the third IMET route is the same as the common VTEP address stored in the first NVE device, the processing unit (1004) is further configured to establish a VXLAN tunnel from the first NVE device to the third NVE device based on the third VTEP address in the third IMET route.

17. The first NVE device according to any one of claims 14 to 15, wherein a CE device is dual-homed to the first NVE device and the second NVE device by using a plurality of Ethernet links, and the plurality of Ethernet links form an ES;

the processing unit (1004) is further configured to determine the first NVE device as a DF in the ES based on an ES route from the second NVE device;

the sending unit (1006) is further configured to send an Ethernet Auto-Discovery per Ethernet segment, Ethernet A-D per ES, route to the second NVE device, wherein the Ethernet A-D per ES route carries an Ethernet segment identifier, ESI, label allocated by the first NVE device to an Ethernet link between the CE device and the first NVE device;

the receiving unit (1002) is further configured to receive, through the VXLAN tunnel from the second NVE device to the first NVE device, a VXLAN packet sent by the second NVE device, wherein an original Ethernet payload in the VXLAN packet comprises BUM traffic, the VXLAN packet carries the ESI label, and the BUM traffic is from the CE device; and

the processing unit (1004) is further configured to avoid, based on that the VXLAN packet carries the ESI label, forwarding the BUM traffic to the CE device by using the ES between the CE device and the first NVE device.

18. The first NVE device according to any one of claims 13 to 15, wherein a CE device is dual-homed to the first NVE device and the second NVE device by using a plurality of Ethernet links, and the plurality of Ethernet links form an ES;

the receiving unit (1002) is further configured to receive a Media Access Control/Internet Protocol advertisement, MAC/IP advertisement, route from the second NVE device, wherein the MAC/IP advertisement route comprises a Media Access Control-virtual local area network identifier, MAC-VLAN ID, the MAC-VLAN ID is used to indicate a VLAN to which a MAC address carried in the MAC/IP advertisement route belongs, and the MAC address carried in the MAC/IP ad-

vertisement route is a MAC address of the CE device, or the MAC address carried in the MAC/IP advertisement route is a MAC address of a host in an Ethernet virtual private network, EVPN, site administered by the CE device; the processing unit (1004) is further configured to determine, based on an ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route, a local interface that is of the first NVE device and that connects to an Ethernet link between the CE device and the first NVE device, wherein the ESI carried in the MAC/IP advertisement route is used to indicate the ES between the CE device and the second NVE device; and

the processing unit (1004) is further configured to trigger sending of a first MAC entry to a forwarding plane of the first NVE device, wherein an outbound interface in the first MAC entry is the local interface of the first NVE device, and the first MAC entry comprises the MAC address and the MAC-VLAN ID that are carried in the MAC/IP advertisement route.

19. The first NVE device according to claim 18, wherein the MAC/IP advertisement route is carried in multi-protocol reachable network layer reachability information, MP_REACH_NLRI, a next hop field in the MP_REACH_NLRI comprises a common VTEP address, the MAC/IP advertisement route further comprises the second VTEP address, and the common VTEP address comprised in the next hop field in the MP_REACH_NLRI is the same as the common VTEP address comprised in the originating router's IP address field in the second IMET route; and the processing unit (1004) is further configured to trigger sending of a second MAC entry to the forwarding plane of the first NVE device based on the common VTEP address comprised in the next hop field in the MP_REACH_NLRI and the second VTEP address comprised in the MAC/IP advertisement route, wherein an outbound interface in the second MAC entry is a local interface that is of the first NVE device and that connects to the VXLAN tunnel from the first NVE device to the second NVE device, and the second MAC entry comprises the MAC address carried in the MAC/IP advertisement route.

20. The first NVE device according to any one of claims 13 to 15, wherein a CE device is connected to the second NVE device;

the receiving unit (1002) is further configured to receive a MAC/IP advertisement route from the second NVE device, wherein the MAC/IP advertisement route is carried in MP REACH NLRI, a next hop field in the MP REACH NLRI comprises a common VTEP address, the MAC/IP adver-

tisement route comprises a MAC-VLAN ID and the second VTEP address, the MAC-VLAN ID is used to indicate a VLAN to which a MAC address carried in the MAC/IP advertisement route belongs, and the MAC address carried in the MAC/IP advertisement route is a MAC address of the CE device, or the MAC address carried in the MAC/IP advertisement route is a MAC address of a host in an EVPN site administered by the CE device;

the processing unit (1004) is further configured to determine, based on an ESI and the MAC-VLAN ID that are carried in the MAC/IP advertisement route, that the first NVE device has no local interface that connects to an Ethernet link between the CE device and the first NVE device, wherein the ESI carried in the MAC/IP advertisement route is used to indicate an ES between the CE device and the second NVE device; and the processing unit (1004) is further configured to trigger sending of a MAC entry to a forwarding plane of the first NVE device based on the common VTEP address comprised in the next hop field in the MP_REACH_NLRI and the second VTEP address comprised in the MAC/IP advertisement route, wherein an outbound interface in the MAC entry is a local interface that is of the first NVE device and that connects to the VXLAN tunnel from the first NVE device to the second NVE device, and the MAC entry comprises the MAC address carried in the MAC/IP advertisement route.

Patentansprüche

1. Routenverarbeitungsverfahren, wobei das Verfahren auf ein "Virtual Extensible Local Area Network" VXLAN angewandt wird, das eine "Ethernet Virtual Private Network"- bzw. EVPN-Steuerebene aufweist, das VXLAN eine erste Netzwerk-Virtualisierungsebene- bzw. NVE-Vorrichtung und eine zweite NVE-Vorrichtung umfasst und das Verfahren Folgendes umfasst:

Empfangen (S102) einer zweiten "Inclusive Multicast Ethernet Tag"- bzw. IMET-Route durch die erste NVE-Vorrichtung von der zweiten NVE-Vorrichtung, wobei ein IP-Adressenfeld des Ursprungsrouteurs in der zweiten IMET-Route eine gemeinsame "Virtual Extensible Local Area Network Tunnel Endpoint"- bzw. VTEP-Adresse umfasst, die zweite IMET-Route ferner eine zweite VTEP-Adresse umfasst und die in dem IP-Adressenfeld des Ursprungsrouteurs in der zweiten IMET-Route enthaltene gemeinsame VTEP-Adresse von der zweiten VTEP-Adresse verschieden ist, die in dem IP-Adressenfeld des

Ursprungsrouteurs in der zweiten IMET-Route enthaltene gemeinsame VTEP-Adresse einen in der zweiten NVE-Vorrichtung enthaltenen gemeinsamen VTEP angibt und die zweite VTEP-Adresse einen in der zweiten NVE-Vorrichtung enthaltenen zweiten VTEP angibt;

Bestimmen (S103) durch die erste NVE-Vorrichtung, ob die in dem IP-Adressenfeld des Ursprungsrouteurs in der zweiten IMET-Route enthaltene gemeinsame VTEP-Adresse dieselbe wie eine in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist, wobei die erste NVE-Vorrichtung und die zweite NVE-Vorrichtung ein Paar von "Border Gateway Protocol"-Peers bzw. BGP-Peers zum Austauschen von IMET-Routen sind, die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse einen in der ersten NVE-Vorrichtung enthaltenen gemeinsamen VTEP angibt; und

wenn die erste NVE-Vorrichtung bestimmt, dass die in dem IP-Adressenfeld des Ursprungsrouteurs in der zweiten IMET-Route enthaltene gemeinsame VTEP-Adresse dieselbe wie die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist, Herstellen (S104) eines VXLAN-Tunnels von der ersten NVE-Vorrichtung zu der zweiten NVE-Vorrichtung durch die erste NVE-Vorrichtung auf der Basis der zweiten VTEP-Adresse in der zweiten IMET-Route.

2. Verfahren nach Anspruch 1, wobei das Verfahren ferner Folgendes umfasst:

Senden (S105) einer ersten IMET-Route zu der zweiten NVE-Vorrichtung durch die erste NVE-Vorrichtung, wobei ein IP-Adressenfeld des Ursprungsrouteurs in der ersten IMET-Route die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse führt, die erste IMET-Route ferner eine erste VTEP-Adresse umfasst, die erste VTEP-Adresse durch die zweite NVE-Vorrichtung zum Herstellen des VXLAN-Tunnels von der zweiten NVE-Vorrichtung zu der ersten NVE-Vorrichtung verwendet wird, die erste VTEP-Adresse einen in der ersten NVE-Vorrichtung enthaltenen ersten VTEP angibt und die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse von der ersten VTEP-Adresse verschieden ist.

3. Verfahren nach Anspruch 2, wobei das VXLAN ferner eine "Provider Edge"- bzw. PE-Vorrichtung umfasst und das Verfahren ferner Folgendes umfasst:

Senden der ersten IMET-Route zu der PE-Vorrichtung durch die erste NVE-Vorrichtung; Empfangen einer IMET-Route von der PE-Vor-

- richtung durch die erste NVE-Vorrichtung, wobei ein IP-Adressenfeld des Ursprungsrouters in der IMET-Route von der PE-Vorrichtung eine VTEP-Adresse der PE-Vorrichtung führt; und
Herstellen eines VXLAN-Tunnels von der ersten NVE-Vorrichtung zu der PE-Vorrichtung durch die erste NVE-Vorrichtung auf der Basis der VTEP-Adresse der PE-Vorrichtung, die sich in der IMET-Route von der PE-Vorrichtung befindet, wobei die VTEP-Adresse der PE-Vorrichtung nicht dieselbe wie die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist.
4. Verfahren nach Anspruch 2 oder 3, wobei
ein VXLAN-Attribut in der ersten IMET-Route die erste VTEP-Adresse führt; und
ein VXLAN-Attribut in der zweiten IMET-Route die zweite VTEP-Adresse führt.
5. Verfahren nach einem der Ansprüche 1 bis 4, wobei das VXLAN ferner eine dritte NVE-Vorrichtung umfasst und das Verfahren ferner Folgendes umfasst:
Empfangen einer dritten IMET-Route von der dritten NVE-Vorrichtung durch die erste NVE-Vorrichtung, wobei ein IP-Adressenfeld des Ursprungsrouters in der dritten IMET-Route eine gemeinsame VTEP-Adresse umfasst, die dritte IMET-Route ferner eine dritte VTEP-Adresse umfasst und die in dem IP-Adressenfeld des Ursprungsrouters in der dritten IMET-Route enthaltene gemeinsame VTEP-Adresse von der dritten VTEP-Adresse verschieden ist;
Bestimmen durch die erste NVE-Vorrichtung, ob die in dem IP-Adressenfeld des Ursprungsrouters in der dritten IMET-Route enthaltene gemeinsame VTEP-Adresse dieselbe wie die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist, wobei die in dem IP-Adressenfeld des Ursprungsrouters in der dritten IMET-Route enthaltene gemeinsame VTEP-Adresse dieselbe wie die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist; und
wenn die erste NVE-Vorrichtung bestimmt, dass die in dem IP-Adressenfeld des Ursprungsrouters in der dritten IMET-Route enthaltene gemeinsame VTEP-Adresse dieselbe wie die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist, Herstellen eines VXLAN-Tunnels von der ersten NVE-Vorrichtung zu der dritten NVE-Vorrichtung durch die erste NVE-Vorrichtung auf der Basis der dritten VTEP-Adresse in der dritten IMET-Route.
6. Verfahren nach Anspruch 5, wobei eine "Customer Edge"- bzw. CE-Vorrichtung bezüglich der ersten NVE-Vorrichtung, der zweiten NVE-Vorrichtung und der dritten NVE-Vorrichtung durch Verwendung mehrerer Ethernet-Verbindungen mehrere Heimaten hat und die mehreren Ethernet-Verbindungen ein Ethernet-Segment ES bilden; und das Verfahren ferner Folgendes umfasst:
Bestimmen der ersten NVE-Vorrichtung als einen designierten Weiterleiter DF in dem ES durch die erste NVE-Vorrichtung auf der Basis von ES-Router von der zweiten NVE-Vorrichtung und der dritten NVE-Vorrichtung.
7. Verfahren nach Anspruch 5 oder 6, wobei das Verfahren ferner Folgendes umfasst:
Empfangen von durch die zweite NVE-Vorrichtung gesendetem "Broadcast"-, "Unknown Unicast"- und "Multicast"- bzw. BUM-Verkehr durch die erste NVE-Vorrichtung mittels des VXLAN-Tunnels von der zweiten NVE-Vorrichtung zu der ersten NVE-Vorrichtung; und
Vermeiden von Weiterleitung des BUM-Verkehrs zu der dritten NVE-Vorrichtung mittels des VXLAN-Tunnels von der ersten NVE-Vorrichtung zu der dritten NVE-Vorrichtung durch die erste NVE-Vorrichtung auf der Basis, dass der BUM-Verkehr mittels des VXLAN-Tunnels von der zweiten NVE-Vorrichtung zu der ersten NVE-Vorrichtung durch die erste NVE-Vorrichtung empfangen wird.
8. Verfahren nach einem der Ansprüche 2 bis 4, wobei eine CE-Vorrichtung bezüglich der ersten NVE-Vorrichtung und der zweiten NVE-Vorrichtung durch Verwendung mehrerer Ethernet-Verbindungen zwei Heimaten hat und die mehreren Ethernet-Verbindungen ein ES bilden; und das Verfahren ferner Folgendes umfasst:
Bestimmen der ersten NVE-Vorrichtung als einen DF in dem ES durch die erste NVE-Vorrichtung auf der Basis einer ES-Route von der zweiten NVE-Vorrichtung;
Senden einer "Ethernet Auto-Discovery per Ethernet Segment"- bzw. Ethernet-A-D-per-ES-Route zu der zweiten NVE-Vorrichtung durch die erste NVE-Vorrichtung, wobei die Ethernet-A-D-per-ES-Route ein Ethernet-Segmentkennungs- bzw. ESI-Label trägt, das durch die erste NVE-Vorrichtung an eine Ethernet-Verbindung zwischen der CE-Vorrichtung und der ersten NVE-Vorrichtung vergeben wird;
Empfangen eines durch die zweite NVE-Vorrichtung gesendeten VXLAN-Pakets durch die erste NVE-Vorrichtung mittels des VXLAN-Tunnels von der zweiten NVE-Vorrichtung zu der ersten NVE-Vorrichtung, wobei ursprüngliche

Ethernet-Nutzinformationen in dem VXLAN-Paket BUM-Verkehr umfassen, das VXLAN-Paket das ESI-Label trägt und der BUM-Verkehr von der CE-Vorrichtung kommt; und

Vermeiden von Weiterleitung des BUM-Verkehrs zu der CE-Vorrichtung durch die erste NVE-Vorrichtung auf der Basis, dass das VXLAN-Paket das ESI-Label trägt, durch Verwendung des ES zwischen der CE-Vorrichtung und der ersten NVE-Vorrichtung.

9. Verfahren nach Anspruch 8, wobei das ESI-Label zwischen einem VXLAN-Header in dem VXLAN-Paket und den ursprünglichen Ethernet-Nutzinformationen eingekapselt ist, der VXLAN-Header Angabeinformationen trägt und die Angabeinformationen verwendet werden, um anzugeben, dass das VXLAN-Paket das ESI-Label trägt.

10. Verfahren nach einem der Ansprüche 1 bis 4, wobei eine CE-Vorrichtung bezüglich der ersten NVE-Vorrichtung und der zweiten NVE-Vorrichtung durch Verwendung mehrerer Ethernet-Verbindungen zwei Heimaten hat und die mehreren Ethernet-Verbindungen ein ES bilden; und das Verfahren ferner Folgendes umfasst:

Empfangen einer "Media Access Control/Internet Protocol"-Ankündigungs- bzw. MAC/IP-Ankündigungsroute von der zweiten NVE-Vorrichtung durch die erste NVE-Vorrichtung, wobei die MAC/IP-Ankündigungsroute eine "Media Access Control-Virtual Local Area Network Identifier" bzw. MAC-VLAN-ID umfasst, die MAC-VLAN-ID verwendet wird, um ein VLAN anzugeben, zu dem eine in der MAC/IP-Ankündigungsroute geführte MAC-Adresse gehört, und die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse eine MAC-Adresse der CE-Vorrichtung ist oder die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse eine MAC-Adresse eines Hosts in einem durch die CE-Vorrichtung administrierten "Ethernet Virtual Private Network"- bzw. EVPN-Standort ist;

Bestimmen durch die erste NVE-Vorrichtung auf der Basis einer ESI und der MAC-VLAN-ID, die in der MAC/IP-Ankündigungsroute geführt wird, einer lokalen Schnittstelle, die von der ersten NVE-Vorrichtung ist und die mit einer Ethernet-Verbindung zwischen der CE-Vorrichtung und der ersten NVE-Vorrichtung verbindet, wobei die in der MAC/IP-Ankündigungsroute geführte ESI verwendet wird, um das ES zwischen der CE-Vorrichtung und der zweiten NVE-Vorrichtung anzugeben; und

Senden eines ersten MAC-Eintrags zu einer Weiterleitungsebene der ersten NVE-Vorrichtung durch eine Steuerebene der ersten NVE-

Vorrichtung, wobei eine abgehende Schnittstelle in dem ersten MAC-Eintrag die lokale Schnittstelle der ersten NVE-Vorrichtung ist und der erste MAC-Eintrag die MAC-Adresse und die MAC-VLAN-ID umfasst, die in der MAC/IP-Ankündigungsroute geführt werden.

11. Verfahren nach Anspruch 10, wobei die MAC/IP-Ankündigungsroute in Mehrprotokoll-erreichbare-Netzwerkschicht-Erreichbarkeitsinformationen MP_REACH_NLRI geführt wird, ein Nächster-Sprung-Feld in den MP REACH NLRI eine gemeinsame VTEP-Adresse umfasst, die MAC/IP-Ankündigungsroute ferner die zweite VTEP-Adresse umfasst und die in dem Nächster-Sprung-Feld in dem MP_REACH_NLRI enthaltene gemeinsame VTEP-Adresse dieselbe wie die in dem IP-Adressenfeld des Ursprungsroute in der zweiten IMET-Route enthaltene gemeinsame VTEP-Adresse ist; und das Verfahren ferner Folgendes umfasst: Senden eines zweiten MAC-Eintrags zu der Weiterleitungsebene der ersten NVE-Vorrichtung durch die Steuerebene der ersten NVE-Vorrichtung auf der Basis der in dem Nächster-Sprung-Feld in den MP_REACH_NLRI enthaltenen gemeinsamen VTEP-Adresse und der in der MAC/IP-Ankündigungsroute enthaltenen zweiten VTEP-Adresse, wobei eine abgehende Schnittstelle in dem zweiten MAC-Eintrag eine lokale Schnittstelle ist, die von der ersten NVE-Vorrichtung ist und die mit dem VXLAN-Tunnel von der ersten NVE-Vorrichtung zu der zweiten NVE-Vorrichtung verbindet, und der zweite MAC-Eintrag die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse umfasst.

12. Verfahren nach einem der Ansprüche 1 bis 4, wobei eine CE-Vorrichtung mit der zweiten NVE-Vorrichtung verbunden ist; und das Verfahren ferner Folgendes umfasst:

Empfangen einer MAC/IP-Ankündigungsroute von der zweiten NVE-Vorrichtung durch die erste NVE-Vorrichtung, wobei die MAC/IP-Ankündigungsroute in MP_REACH_NLRI geführt wird, ein Nächster-Sprung-Feld in den MP REACH NLRI eine gemeinsame VTEP-Adresse umfasst, die MAC/IP-Ankündigungsroute eine MAC/VLAN-ID und die zweite VTEP-Adresse umfasst, die MAC-VLAN-ID verwendet wird, um ein VLAN anzugeben, zu dem eine in der MAC/IP-Ankündigungsroute geführte MAC-Adresse gehört, und die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse eine MAC-Adresse der CE-Vorrichtung ist oder die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse eine MAC-Adresse eines Hosts in einem durch die CE-Vorrichtung administrierten EVPN-Standort ist;

Bestimmen durch die erste NVE-Vorrichtung auf der Basis einer ESI und der MAC-VLAN-ID, die in der MAC/IP-Ankündigungsroute geführt werden, dass die erste NVE-Vorrichtung keine lokale Schnittstelle aufweist, die mit einer Ethernet-Verbindung zwischen der CE-Vorrichtung und der ersten NVE-Vorrichtung verbindet, wobei die in der MAC/IP-Ankündigungsroute geführte ESI verwendet wird, um ein ES zwischen der CE-Vorrichtung und der zweiten NVE-Vorrichtung anzugeben; und
 Senden eines MAC-Eintrags zu einer Weiterleitungsebene der ersten NVE-Vorrichtung durch eine Steuerebene der ersten NVE-Vorrichtung auf der Basis der in dem Nächster-Sprung-Feld in den MP_REACH_NLRI enthaltenen gemeinsamen VTEP-Adresse und der in der MAC/IP-Ankündigungsroute enthaltenen zweiten VTEP-Adresse, wobei eine abgehende Schnittstelle in dem MAC-Eintrag eine lokale Schnittstelle ist, die von der ersten NVE-Vorrichtung ist und die mit dem VXLAN-Tunnel von der ersten NVE-Vorrichtung zu der zweiten NVE-Vorrichtung verbindet, und der MAC-Eintrag die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse umfasst.

13. Erste "Network Virtualization Edge"- bzw. NVE-Vorrichtung, wobei die erste NVE auf ein "Virtual Extensible Local Area Network" VXLAN angewandt wird, das eine "Ethernet Virtual Private Network"- bzw. EVPN-Steuerebene aufweist, das VXLAN ferner eine zweite NVE-Vorrichtung umfasst und die erste NVE-Vorrichtung Folgendes umfasst:

eine Empfangseinheit (1002), ausgelegt zum Empfangen einer zweiten "Inclusive Multicast Ethernet Tag"- bzw. IMET-Route von der zweiten NVE-Vorrichtung, wobei ein IP-Adressenfeld des Ursprungsrouteurs in der zweiten IMET-Route eine gemeinsame "Virtual Extensible Local Area Network Tunnel Endpoint"- bzw. VTEP-Adresse umfasst, die zweite IMET-Route ferner eine zweite VTEP-Adresse umfasst und die in dem IP-Adressenfeld des Ursprungsrouteurs in der zweiten IMET-Route enthaltene gemeinsame VTEP-Adresse von der zweiten VTEP-Adresse verschieden ist, die in dem IP-Adressenfeld des Ursprungsrouteurs in der zweiten IMET-Route enthaltene gemeinsame VTEP-Adresse einen in der zweiten NVE-Vorrichtung enthaltenen gemeinsamen VTEP angibt und die zweite VTEP-Adresse einen in der zweiten NVE-Vorrichtung enthaltenen zweiten VTEP angibt; und
 eine Verarbeitungseinheit (1004), ausgelegt zum Bestimmen, ob die in dem IP-Adressenfeld des Ursprungsrouteurs in der zweiten IMET-Rou-

te enthaltene gemeinsame VTEP-Adresse dieselbe wie eine in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist, wobei die erste NVE-Vorrichtung und die zweite NVE-Vorrichtung ein Paar von "Border Gateway Protocol"-Peers bzw. BGP-Peers zum Austauschen von IMET-Routen sind, die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse einen in der ersten NVE-Vorrichtung enthaltenen gemeinsamen VTEP angibt; wobei,
 wenn die Verarbeitungseinheit (1004) bestimmt, dass die in dem IP-Adressenfeld des Ursprungsrouteurs in der zweiten IMET-Route enthaltene gemeinsame VTEP-Adresse dieselbe wie die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist, die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Herstellen eines VXLAN-Tunnels von der ersten NVE-Vorrichtung zu der zweiten NVE-Vorrichtung auf der Basis der zweiten VTEP-Adresse in der zweiten IMET-Route.

14. Erste NVE-Vorrichtung nach Anspruch 13, wobei die erste NVE-Vorrichtung ferner Folgendes umfasst:

eine Sendeeinheit (1006), ausgelegt zum Senden einer ersten IMET-Route zu der zweiten NVE-Vorrichtung, wobei ein IP-Adressenfeld des Ursprungsrouteurs in der ersten IMET-Route die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse führt, die erste IMET-Route ferner eine erste VTEP-Adresse umfasst, die erste VTEP-Adresse durch die zweite NVE-Vorrichtung zum Herstellen des VXLAN-Tunnels von der zweiten NVE-Vorrichtung zu der ersten NVE-Vorrichtung verwendet wird, die erste VTEP-Adresse einen in der ersten NVE-Vorrichtung enthaltenen ersten VTEP angibt und die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse von der ersten VTEP-Adresse verschieden ist.

15. Erste NVE-Vorrichtung nach Anspruch 14, wobei

die Sendeeinheit ferner ausgelegt ist zum Senden der ersten IMET-Route zu einer "Provider Edge"- bzw. PE-Vorrichtung;
 die Empfangseinheit (1002) ferner ausgelegt ist zum Empfangen einer IMET-Route von der PE-Vorrichtung, wobei ein IP-Adressenfeld des Ursprungsrouteurs in der IMET-Route von der PE-Vorrichtung eine VTEP-Adresse der PE-Vorrichtung führt, und
 die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Herstellen eines VXLAN-Tunnels von der ersten NVE-Vorrichtung zu der PE-Vor-

richtung auf der Basis der VTEP-Adresse der PE-Vorrichtung, die sich in der IMET-Route von der PE-Vorrichtung befindet, wobei die VTEP-Adresse der PE-Vorrichtung nicht dieselbe wie die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist.

16. Erste NVE-Vorrichtung nach einem der Ansprüche 13 bis 15, wobei

die Empfangseinheit (1002) ferner ausgelegt ist zum Empfangen einer dritten IMET-Route von einer dritten NVE-Vorrichtung, wobei ein IP-Adressenfeld des Ursprungsrouteurs in der dritten IMET-Route eine gemeinsame VTEP-Adresse umfasst, die dritte IMET-Route ferner eine dritte VTEP-Adresse umfasst und die in dem IP-Adressenfeld des Ursprungsrouteurs in der dritten IMET-Route enthaltene gemeinsame VTEP-Adresse von der dritten VTEP-Adresse verschieden ist;

die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Bestimmen, ob die in dem IP-Adressenfeld des Ursprungsrouteurs in der dritten IMET-Route enthaltene gemeinsame VTEP-Adresse dieselbe wie die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist, wobei die in dem IP-Adressenfeld des Ursprungsrouteurs in der dritten IMET-Route enthaltene gemeinsame VTEP-Adresse dieselbe wie die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist; und,

wenn die Verarbeitungseinheit (1004) bestimmt, dass die in dem IP-Adressenfeld des Ursprungsrouteurs in der dritten IMET-Route enthaltene gemeinsame VTEP-Adresse dieselbe wie die in der ersten NVE-Vorrichtung gespeicherte gemeinsame VTEP-Adresse ist, die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Herstellen eines VXLAN-Tunnels von der ersten NVE-Vorrichtung zu der dritten NVE-Vorrichtung auf der Basis der dritten VTEP-Adresse in der dritten IMET-Route.

17. Erste NVE-Vorrichtung nach einem der Ansprüche 14 bis 15, wobei eine CE-Vorrichtung bezüglich der ersten NVE-Vorrichtung und der zweiten NVE-Vorrichtung durch Verwendung mehrerer Ethernet-Verbindungen zwei Heimaten hat und die mehreren Ethernet-Verbindungen ein ES bilden;

die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Bestimmen der ersten NVE-Vorrichtung als einen DF in dem ES auf der Basis einer ES-Route von der zweiten NVE-Vorrichtung; die Sendeeinheit (1006) ferner ausgelegt ist zum Senden einer "Ethernet Auto-Discovery per

Ethernet Segment"- bzw. Ethernet-A-D-per-ES-Route zu der zweiten NVE-Vorrichtung, wobei die Ethernet-A-D-per-ES-Route ein Ethernet-Segmentkennungs- bzw. ESI-Label trägt, das durch die erste NVE-Vorrichtung an eine Ethernet-Verbindung zwischen der CE-Vorrichtung und der ersten NVE-Vorrichtung vergeben wird; die Empfangseinheit (1002) ferner ausgelegt ist zum Empfangen eines durch die zweite NVE-Vorrichtung gesendeten VXLAN-Pakets mittels des VXLAN-Tunnels von der zweiten NVE-Vorrichtung zu der ersten NVE-Vorrichtung, wobei ursprüngliche Ethernet-Nutzinformationen in dem VXLAN-Paket BUM-Verkehr umfassen, das VXLAN-Paket das ESI-Label trägt und der BUM-Verkehr von der CE-Vorrichtung kommt; und

die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Vermeiden von Weiterleitung des BUM-Verkehrs zu der CE-Vorrichtung auf der Basis, dass das VXLAN-Paket das ESI-Label trägt, durch Verwendung des ES zwischen der CE-Vorrichtung und der ersten NVE-Vorrichtung.

18. Erste NVE-Vorrichtung nach einem der Ansprüche 13 bis 15, wobei eine CE-Vorrichtung bezüglich der ersten NVE-Vorrichtung und der zweiten NVE-Vorrichtung durch Verwendung mehrerer Ethernet-Verbindungen zwei Heimaten hat und die mehreren Ethernet-Verbindungen ein ES bilden;

die Empfangseinheit (1002) ferner ausgelegt ist zum Empfangen einer "Media Access Control/Internet Protocol"-Ankündigungs- bzw. MAC/IP-Ankündigungsroute von der zweiten NVE-Vorrichtung, wobei die MAC/IP-Ankündigungsroute eine "Media Access Control-Virtual Local Area Network Identifier" bzw. MAC-VLAN-ID umfasst, die MAC-VLAN-ID verwendet wird, um ein VLAN anzugeben, zu dem eine in der MAC/IP-Ankündigungsroute geführte MAC-Adresse gehört, und die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse eine MAC-Adresse der CE-Vorrichtung ist oder die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse eine MAC-Adresse eines Hosts in einem durch die CE-Vorrichtung administrierten "Ethernet Virtual Private Network"- bzw. EVPN-Standort ist;

die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Bestimmen auf der Basis einer ESI und der MAC-VLAN-ID, die in der MAC/IP-Ankündigungsroute geführt wird, einer lokalen Schnittstelle, die von der ersten NVE-Vorrichtung ist und die mit einer Ethernet-Verbindung zwischen der CE-Vorrichtung und der ersten NVE-Vorrichtung verbindet, wobei die in der

MAC/IP-Ankündigungsroute geführte ESI verwendet wird, um das ES zwischen der CE-Vorrichtung und der zweiten NVE-Vorrichtung anzugeben; und

die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Triggern des Sendens eines ersten MAC-Eintrags zu einer Weiterleitungsebene der ersten NVE-Vorrichtung, wobei eine abgehende Schnittstelle in dem ersten MAC-Eintrag die lokale Schnittstelle der ersten NVE-Vorrichtung ist und der erste MAC-Eintrag die MAC-Adresse und die MAC-VLAN-ID umfasst, die in der MAC/IP-Ankündigungsroute geführt werden.

19. Erste NVE-Vorrichtung nach Anspruch 18, wobei die MAC/IP-Ankündigungsroute in Mehrprotokoll-erreichbare-Netzwerkschicht-Erreichbarkeitsinformationen MP_REACH_NLRI geführt wird, ein Nächster-Sprung-Feld in den MP_REACH_NLRI eine gemeinsame VTEP-Adresse umfasst, die MAC/IP-Ankündigungsroute ferner die zweite VTEP-Adresse umfasst und die in dem Nächster-Sprung-Feld in dem MP_REACH_NLRI enthaltene gemeinsame VTEP-Adresse dieselbe wie die in dem IP-Adressenfeld des Ursprungsrouteurs in der zweiten IMET-Route enthaltene gemeinsame VTEP-Adresse ist; und

die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Triggern des Sendens eines zweiten MAC-Eintrags zu der Weiterleitungsebene der ersten NVE-Vorrichtung auf der Basis der in dem Nächster-Sprung-Feld in den MP_REACH_NLRI enthaltenen gemeinsamen VTEP-Adresse und der in der MAC/IP-Ankündigungsroute enthaltenen zweiten VTEP-Adresse, wobei eine abgehende Schnittstelle in dem zweiten MAC-Eintrag eine lokale Schnittstelle ist, die von der ersten NVE-Vorrichtung ist und die mit dem VXLAN-Tunnel von der ersten NVE-Vorrichtung zu der zweiten NVE-Vorrichtung verbindet, und der zweite MAC-Eintrag die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse umfasst.

20. Erste NVE-Vorrichtung nach einem der Ansprüche 13 bis 15, wobei eine CE-Vorrichtung mit der zweiten NVE-Vorrichtung verbunden ist;

die Empfangseinheit (1002) ferner ausgelegt ist zum Empfangen einer MAC/IP-Ankündigungsroute von der zweiten NVE-Vorrichtung, wobei die MAC/IP-Ankündigungsroute in MP_REACH_NLRI geführt wird, ein Nächster-Sprung-Feld in den MP_REACH_NLRI eine gemeinsame VTEP-Adresse umfasst, die MAC/IP-Ankündigungsroute eine MAC/VLAN-ID und die zweite VTEP-Adresse umfasst, die MAC-VLAN-ID verwendet wird, um ein VLAN anzugeben, zu dem eine in der MAC/IP-Ankündigungsroute geführte MAC-Adresse gehört,

und die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse eine MAC-Adresse der CE-Vorrichtung ist oder die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse eine MAC-Adresse eines Hosts in einem durch die CE-Vorrichtung administrierten EVPN-Standort ist;

die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Bestimmen auf der Basis einer ESI und der MAC-VLAN-ID, die in der MAC/IP-Ankündigungsroute geführt werden, dass die erste NVE-Vorrichtung keine lokale Schnittstelle aufweist, die mit einer Ethernet-Verbindung zwischen der CE-Vorrichtung und der ersten NVE-Vorrichtung verbindet, wobei die in der MAC/IP-Ankündigungsroute geführte ESI verwendet wird, um ein ES zwischen der CE-Vorrichtung und der zweiten NVE-Vorrichtung anzugeben; und

die Verarbeitungseinheit (1004) ferner ausgelegt ist zum Triggern des Sendens eines MAC-Eintrags zu einer Weiterleitungsebene der ersten NVE-Vorrichtung auf der Basis der in dem Nächster-Sprung-Feld in den MP_REACH_NLRI enthaltenen gemeinsamen VTEP-Adresse und der in der MAC/IP-Ankündigungsroute enthaltenen VTEP-Adresse, wobei eine abgehende Schnittstelle in dem MAC-Eintrag eine lokale Schnittstelle ist, die von der ersten NVE-Vorrichtung ist und die mit dem VXLAN-Tunnel von der ersten NVE-Vorrichtung zu der zweiten NVE-Vorrichtung verbindet, und der MAC-Eintrag die in der MAC/IP-Ankündigungsroute geführte MAC-Adresse umfasst.

Revendications

1. Procédé de traitement d'itinéraire, dans lequel le procédé est appliqué à un Réseau Local Extensible Virtuel, VXLAN, ayant un plan de contrôle de Réseau Privé Virtuel Ethernet, EVPN, le VXLAN comprenant un premier dispositif de périphérie de virtualisation de réseau, NVE, et un deuxième dispositif NVE, et le procédé comprenant les étapes suivantes :

recevoir (S102), par le premier dispositif NVE, un deuxième itinéraire de balise Ethernet de multidiffusion inclusive, IMET, en provenance du deuxième dispositif NVE, où un champ d'adresse IP de routeur d'origine dans le deuxième itinéraire IMET comprend une adresse de point d'extrémité de tunnel de réseau local extensible virtuel commun, VTEP, le deuxième itinéraire IMET comprend en outre une deuxième adresse VTEP, et l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le deuxième itinéraire IMET est

- différente de la deuxième adresse VTEP, l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le deuxième itinéraire IMET indique un VTEP commun inclus dans le deuxième dispositif NVE, la deuxième adresse VTEP indique un deuxième VTEP inclus dans le deuxième dispositif NVE ;
- déterminer (S103), par le premier dispositif NVE, si l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le deuxième itinéraire IMET est la même qu'une adresse VTEP commune stockée dans le premier dispositif NVE, où le premier dispositif NVE et le deuxième dispositif NVE sont une paire de pairs de protocole de passerelle de frontière, pairs BGP, pour échanger des itinéraires IMET, l'adresse VTEP commune stockée dans le premier dispositif NVE indiquant un VTEP commun inclus dans le premier dispositif NVE ; et
- lorsque le premier dispositif NVE détermine que l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le deuxième itinéraire IMET est la même que l'adresse VTEP commune stockée dans le premier dispositif NVE, établir (S104), par le premier dispositif NVE, un tunnel VXLAN du premier dispositif NVE au deuxième dispositif NVE sur la base de la deuxième adresse VTEP dans le deuxième itinéraire IMET.
2. Procédé selon la revendication 1, dans lequel le procédé comprend en outre l'étape suivante :
- envoyer (S105), par le premier dispositif NVE, un premier itinéraire IMET au deuxième dispositif NVE, où un champ d'adresse IP du routeur d'origine dans le premier itinéraire IMET porte l'adresse VTEP commune stockée dans le premier dispositif NVE, le premier itinéraire IMET comprend en outre une première adresse VTEP, la première adresse VTEP est utilisée par le deuxième dispositif NVE pour établir le tunnel VXLAN du deuxième dispositif NVE au premier dispositif NVE, la première adresse VTEP indique un premier VTEP inclus dans le premier dispositif NVE, et l'adresse VTEP commune stockée dans le premier dispositif NVE est différente de la première adresse VTEP.
3. Procédé selon la revendication 2, dans lequel le VXLAN comprend en outre un dispositif de périphérie fournisseur, PE, et le procédé comprend en outre les étapes suivantes :
- envoyer, par le premier dispositif NVE, le premier itinéraire IMET au dispositif PE ; recevoir,
- par le premier dispositif NVE, un itinéraire IMET en provenance du dispositif PE, où un champ d'adresse IP du routeur d'origine dans l'itinéraire IMET en provenance du dispositif PE porte une adresse VTEP du dispositif PE ; et
- établir, par le premier dispositif NVE, un tunnel VXLAN du premier dispositif NVE au dispositif PE sur la base de l'adresse VTEP du dispositif PE qui est dans l'itinéraire IMET du dispositif PE, où l'adresse VTEP du dispositif PE n'est pas la même que l'adresse VTEP commune stockée dans le premier dispositif NVE.
4. Procédé selon la revendication 2 ou la revendication 3, dans lequel :
- un attribut VXLAN dans le premier itinéraire IMET porte la première adresse VTEP ; et
- un attribut VXLAN dans le deuxième itinéraire IMET porte la deuxième adresse VTEP.
5. Procédé selon l'une quelconque des revendications 1 à 4, dans lequel le VXLAN comprend en outre un troisième dispositif NVE, et le procédé comprend en outre les étapes suivantes :
- recevoir, par le premier dispositif NVE, un troisième itinéraire IMET en provenance du troisième dispositif NVE, où un champ d'adresse IP du routeur d'origine dans le troisième itinéraire IMET comprend une adresse VTEP commune, le troisième itinéraire IMET comprend en outre une troisième adresse VTEP, et l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le troisième itinéraire IMET est différente de la troisième adresse VTEP ;
- déterminer, par le premier dispositif NVE, si l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le troisième itinéraire IMET est la même que l'adresse VTEP commune stockée dans le premier dispositif NVE, où l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le troisième itinéraire IMET est la même que l'adresse VTEP commune stockée dans le premier dispositif NVE ; et
- lorsque le premier dispositif NVE détermine que l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le troisième itinéraire IMET est la même que l'adresse VTEP commune stockée dans le premier dispositif NVE, établir, par le premier dispositif NVE, un tunnel VXLAN du premier dispositif NVE au troisième dispositif NVE sur la base de la troisième adresse VTEP dans le troisième itinéraire IMET.

6. Procédé selon la revendication 5, dans lequel un dispositif de périphérie client, CE, est rattaché de manière multiple au premier dispositif NVE, au deuxième dispositif NVE, et au troisième dispositif NVE en utilisant une pluralité de liaisons Ethernet, et la pluralité de liaisons Ethernet forme un segment Ethernet, ES ; et
le procédé comprend en outre l'étape suivante :
déterminer, par le premier dispositif NVE, le premier dispositif NVE en tant que transitaire désigné, DF, dans l'ES sur la base des itinéraires ES provenant du deuxième dispositif NVE et du troisième dispositif NVE.
7. Procédé selon la revendication 5 ou la revendication 6, dans lequel le procédé comprend en outre les étapes suivantes :
- recevoir, par le premier dispositif NVE à travers le tunnel VXLAN du deuxième dispositif NVE au premier dispositif NVE, du trafic de diffusion, de monodiffusion inconnue et de multidiffusion, BUM, envoyé par le deuxième dispositif NVE ; et éviter, par le premier dispositif NVE, de transmettre le trafic BUM au troisième dispositif NVE par le tunnel VXLAN du premier dispositif NVE au troisième dispositif NVE, sur la base du fait que le trafic BUM est reçu par le premier dispositif NVE par le tunnel VXLAN du deuxième dispositif NVE au premier dispositif NVE.
8. Procédé selon l'une quelconque des revendications 2 à 4, dans lequel un dispositif CE est doublement rattaché au premier dispositif NVE et au deuxième dispositif NVE en utilisant une pluralité de liaisons Ethernet, et la pluralité de liaisons Ethernet forme un ES ; et
le procédé comprend en outre les étapes suivantes :
- déterminer, par le premier dispositif NVE, le premier dispositif NVE en tant que DF dans l'ES sur la base d'un itinéraire ES provenant du deuxième dispositif NVE ;
envoyer, par le premier dispositif NVE, un itinéraire d'auto-découverte Ethernet par segment Ethernet, Ethernet A-D par ES, au deuxième dispositif NVE, où l'itinéraire Ethernet A-D par ES porte une étiquette d'identifiant de segment Ethernet, ESI, allouée par le premier dispositif NVE à une liaison Ethernet entre le dispositif CE et le premier dispositif NVE ;
recevoir, par le premier dispositif NVE à travers le tunnel VXLAN du deuxième dispositif NVE au premier dispositif NVE, un paquet VXLAN envoyé par le deuxième dispositif NVE, où une charge utile Ethernet d'origine dans le paquet VXLAN comprend un trafic BUM, le paquet VXLAN porte l'étiquette ESI, et le trafic BUM pro-

vient du dispositif CE ; et
éviter, par le premier dispositif NVE, sur la base du fait que le paquet VXLAN porte l'étiquette ESI, de transférer le trafic BUM au dispositif CE en utilisant l'ES entre le dispositif CE et le premier dispositif NVE.

9. Procédé selon la revendication 8, dans lequel l'étiquette ESI est encapsulée entre un en-tête VXLAN dans le paquet VXLAN et la charge utile Ethernet originale, l'en-tête VXLAN portant des informations d'indication, et les informations d'indication étant utilisées pour indiquer que le paquet VXLAN porte l'étiquette ESI.
10. Procédé selon l'une quelconque des revendications 1 à 4, dans lequel un dispositif CE est doublement rattaché au premier dispositif NVE et au deuxième dispositif NVE en utilisant une pluralité de liaisons Ethernet, et la pluralité de liaisons Ethernet forme un ES ; et
le procédé comprend en outre les étapes suivantes :
- recevoir, par le premier dispositif NVE, un itinéraire d'annonce de contrôle d'accès au support/protocole Internet, annonce MAC/IP, en provenance du deuxième dispositif NVE, où l'itinéraire d'annonce MAC/IP comprend un identifiant de réseau local virtuel de contrôle d'accès au support, ID MAC-VLAN, l'ID MAC-VLAN étant utilisé pour indiquer un VLAN auquel une adresse MAC transportée dans l'itinéraire d'annonce MAC/IP appartient, et l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP est une adresse MAC du dispositif CE, ou l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP est une adresse MAC d'un hôte dans un site de réseau privé virtuel Ethernet, EVPN, administré par le dispositif CE ;
déterminer, par le premier dispositif NVE, sur la base d'un ESI et de l'ID MAC-VLAN qui sont transportés dans l'itinéraire d'annonce MAC/IP, une interface locale qui est du premier dispositif NVE et qui se connecte à une liaison Ethernet entre le dispositif CE et le premier dispositif NVE, où l'ESI transporté dans l'itinéraire d'annonce MAC/IP est utilisé pour indiquer l'ES entre le dispositif CE et le deuxième dispositif NVE ; et
envoyer, par un plan de contrôle du premier dispositif NVE, une première entrée MAC à un plan d'acheminement du premier dispositif NVE, où une interface sortante dans la première entrée MAC est l'interface locale du premier dispositif NVE, et la première entrée MAC comprend l'adresse MAC et l'ID MAC-VLAN qui sont transportés dans l'itinéraire d'annonce MAC/IP.
11. Procédé selon la revendication 10, dans lequel l'iti-

néraire d'annonce MAC/IP est transporté dans des informations d'accessibilité de couche réseau atteignable multiprotocole, MP_REACH_NLRI, un champ de saut suivant dans le MP_REACH_NLRI comprend une adresse VTEP commune, l'itinéraire d'annonce MAC/IP comprend en outre la deuxième adresse VTEP, et l'adresse VTEP commune comprise dans le champ de saut suivant dans le MP_REACH_NLRI est la même que l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le deuxième itinéraire IMET ;

et
le procédé comprend en outre l'étape suivante :
envoyer, par le plan de contrôle du premier dispositif NVE, une deuxième entrée MAC au plan d'acheminement du premier dispositif NVE sur la base de l'adresse VTEP commune comprise dans le champ de saut suivant dans le MP_REACH_NLRI et de la deuxième adresse VTEP comprise dans l'itinéraire d'annonce MAC/IP, où une interface sortante dans la deuxième entrée MAC est une interface locale qui est du premier dispositif NVE et qui se connecte au tunnel VXLAN du premier dispositif NVE au deuxième dispositif NVE, et la deuxième entrée MAC comprend l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP.

12. Procédé selon l'une quelconque des revendications 1 à 4, dans lequel un dispositif CE est connecté au deuxième dispositif NVE ; et
le procédé comprend en outre les étapes suivantes :

recevoir, par le premier dispositif NVE, un itinéraire d'annonce MAC/IP en provenance du deuxième dispositif NVE, où l'itinéraire d'annonce MAC/IP est transporté dans le MP_REACH_NLRI, un champ de saut suivant dans le MP_REACH_NLRI comprend une adresse VTEP commune, l'itinéraire d'annonce MAC/IP comprend un ID MAC-VLAN et la deuxième adresse VTEP, l'ID MAC-VLAN est utilisé pour indiquer un VLAN auquel une adresse MAC transportée dans l'itinéraire d'annonce MAC/IP appartient, et l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP est une adresse MAC du dispositif CE, ou l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP est une adresse MAC d'un hôte dans un site EVPN administré par le dispositif CE ;
déterminer, par le premier dispositif NVE, sur la base d'un ESI et de l'ID MAC-VLAN qui sont transportés dans l'itinéraire d'annonce MAC/IP, que le premier dispositif NVE n'a pas d'interface locale qui se connecte à une liaison Ethernet entre le dispositif CE et le premier dispositif NVE, où l'ESI transporté dans l'itinéraire d'annonce MAC/IP est utilisé pour indiquer un ES entre le dispositif CE et le deuxième dispositif

NVE ; et

envoyer, par un plan de contrôle du premier dispositif NVE, une entrée MAC à un plan d'acheminement du premier dispositif NVE sur la base de l'adresse VTEP commune comprise dans le champ de saut suivant dans le MP_REACH_NLRI et la deuxième adresse VTEP comprise dans l'itinéraire d'annonce MAC/IP, où une interface sortante dans l'entrée MAC est une interface locale qui est du premier dispositif NVE et qui se connecte au tunnel VXLAN du premier dispositif NVE au deuxième dispositif NVE, et l'entrée MAC comprend l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP.

13. Premier dispositif de périphérie de virtualisation de réseau, NVE, où le premier NVE est appliqué à un réseau local extensible virtuel, VXLAN, ayant un plan de contrôle de réseau privé virtuel Ethernet, EVPN, le VXLAN comprenant en outre un deuxième dispositif NVE, et le premier dispositif NVE comprenant :

une unité de réception (1002), configurée pour recevoir un deuxième itinéraire de balise Ethernet de multidiffusion inclusive, IMET, à partir du deuxième dispositif NVE, où un champ d'adresse IP de routeur d'origine dans le deuxième itinéraire IMET comprend une adresse de point d'extrémité de tunnel de réseau local extensible virtuel commun, VTEP, le deuxième itinéraire IMET comprenant en outre une deuxième adresse VTEP, et l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le deuxième itinéraire IMET étant différente de la deuxième adresse VTEP, l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le deuxième itinéraire IMET indiquant un VTEP commun inclus dans le deuxième dispositif NVE, la deuxième adresse VTEP indiquant un deuxième VTEP inclus dans le deuxième dispositif NVE ; et

une unité de traitement (1004), configurée pour déterminer si l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le deuxième itinéraire IMET est la même qu'une adresse VTEP commune stockée dans le premier dispositif NVE, où le premier dispositif NVE et le deuxième dispositif NVE sont une paire de pairs de protocole de passerelle frontière, pairs BGP, pour échanger des itinéraires IMET, l'adresse VTEP commune stockée dans le premier dispositif NVE indiquant un VTEP commun inclus dans le premier dispositif NVE ; où

lorsque l'unité de traitement (1004) détermine que l'adresse VTEP commune comprise dans

- le champ d'adresse IP du routeur d'origine dans le deuxième itinéraire IMET est la même que l'adresse VTEP commune stockée dans le premier dispositif NVE, l'unité de traitement (1004) est en outre configurée pour établir un tunnel VXLAN du premier dispositif NVE au deuxième dispositif NVE sur la base de la deuxième adresse VTEP dans le deuxième itinéraire IMET. 5
14. Premier dispositif NVE selon la revendication 13, dans lequel le premier dispositif NVE comprend en outre : 10
- une unité d'envoi (1006), configurée pour envoyer un premier itinéraire IMET au deuxième dispositif NVE, où un champ d'adresse IP du routeur d'origine dans le premier itinéraire IMET porte l'adresse VTEP commune stockée dans le premier dispositif NVE, le premier itinéraire IMET comprend en outre une première adresse VTEP, la première adresse VTEP est utilisée par le deuxième dispositif NVE pour établir le tunnel VXLAN du deuxième dispositif NVE au premier dispositif NVE, la première adresse VTEP indique un premier VTEP inclus dans le premier dispositif NVE, et l'adresse VTEP commune stockée dans le premier dispositif NVE est différente de la première adresse VTEP. 15 20 25
15. Premier dispositif NVE selon la revendication 14, dans lequel : 30
- l'unité d'envoi est en outre configurée pour envoyer le premier itinéraire IMET à un dispositif de périphérie fournisseur, PE ;
- l'unité de réception (1002) est en outre configurée pour recevoir un itinéraire IMET du dispositif PE, où un champ d'adresse IP de routeur d'origine dans l'itinéraire IMET provenant du dispositif PE porte une adresse VTEP du dispositif PE ; et 35
- l'unité de traitement (1004) est en outre configurée pour établir un tunnel VXLAN du premier dispositif NVE au dispositif PE sur la base de l'adresse VTEP du dispositif PE qui est dans l'itinéraire IMET provenant du dispositif PE, où l'adresse VTEP du dispositif PE n'est pas la même que l'adresse VTEP commune stockée dans le premier dispositif NVE. 40 45
16. Premier dispositif NVE selon l'une quelconque des revendications 13 à 15, dans lequel : 50
- l'unité de réception (1002) est en outre configurée pour recevoir un troisième itinéraire IMET à partir d'un troisième dispositif NVE, où un champ d'adresse IP de routeur d'origine dans le troisième itinéraire IMET comprend une adresse VTEP commune, le troisième itinéraire IMET comprend en outre une troisième adresse 55
- VTEP, et l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le troisième itinéraire IMET est différente de la troisième adresse VTEP ;
- l'unité de traitement (1004) est en outre configurée pour déterminer si l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le troisième itinéraire IMET est la même que l'adresse VTEP commune stockée dans le premier dispositif NVE ; et
- lorsque l'unité de traitement (1004) détermine que l'adresse VTEP commune comprise dans le champ d'adresse IP du routeur d'origine dans le troisième itinéraire IMET est la même que l'adresse VTEP commune stockée dans le premier dispositif NVE, l'unité de traitement (1004) est en outre configurée pour établir un tunnel VXLAN du premier dispositif NVE au troisième dispositif NVE sur la base de la troisième adresse VTEP dans le troisième itinéraire IMET. 17. Premier dispositif NVE selon l'une quelconque des revendications 14 à 15, dans lequel un dispositif CE est doublement rattaché au premier dispositif NVE et au deuxième dispositif NVE en utilisant une pluralité de liaisons Ethernet, et la pluralité de liaisons Ethernet forme un ES ;
- l'unité de traitement (1004) est en outre configurée pour déterminer le premier dispositif NVE comme un DF dans l'ES sur la base d'un itinéraire ES provenant du deuxième dispositif NVE ;
- l'unité d'envoi (1006) est en outre configurée pour envoyer un itinéraire d'auto-découverte Ethernet par segment Ethernet, Ethernet A-D par ES, au deuxième dispositif NVE, où l'itinéraire Ethernet A-D par ES porte une étiquette d'identifiant de segment Ethernet, ESI, allouée par le premier dispositif NVE à une liaison Ethernet entre le dispositif CE et le premier dispositif NVE ;
- l'unité de réception (1002) est en outre configurée pour recevoir, à travers le tunnel VXLAN du deuxième dispositif NVE au premier dispositif NVE, un paquet VXLAN envoyé par le deuxième dispositif NVE, où une charge utile Ethernet d'origine dans le paquet VXLAN comprend un trafic BUM, le paquet VXLAN porte l'étiquette ESI, et le trafic BUM provient du dispositif CE ; et
- l'unité de traitement (1004) est en outre configurée pour éviter, sur la base du fait que le paquet VXLAN porte l'étiquette ESI, de transférer le trafic BUM au dispositif CE en utilisant l'ES

entre le dispositif CE et le premier dispositif NVE.

18. Premier dispositif NVE selon l'une quelconque des revendications 13 à 15, dans lequel un dispositif CE est doublement rattaché au premier dispositif NVE et au deuxième dispositif NVE en utilisant une pluralité de liaisons Ethernet, et la pluralité de liaisons Ethernet forme un ES ;

l'unité de réception (1002) est en outre configurée pour recevoir un itinéraire d'annonce de contrôle d'accès au support/protocole Internet, annonce MAC/IP, en provenance du deuxième dispositif NVE, où l'itinéraire d'annonce MAC/IP comprend un identifiant de réseau local virtuel de contrôle d'accès au support, ID MAC-VLAN, l'ID MAC-VLAN étant utilisé pour indiquer un VLAN auquel une adresse MAC transportée dans l'itinéraire d'annonce MAC/IP appartient, et l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP est une adresse MAC du dispositif CE, ou l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP est une adresse MAC d'un hôte dans un site de réseau privé virtuel Ethernet, EVPN, administré par le dispositif CE ;

l'unité de traitement (1004) est en outre configurée pour déterminer, sur la base d'un ESI et de l'ID MAC-VLAN qui sont transportés dans l'itinéraire d'annonce MAC/IP, une interface locale qui est du premier dispositif NVE et qui se connecte à une liaison Ethernet entre le dispositif CE et le premier dispositif NVE, où l'ESI transporté dans l'itinéraire d'annonce MAC/IP est utilisé pour indiquer l'ES entre le dispositif CE et le deuxième dispositif NVE ; et

l'unité de traitement (1004) est en outre configurée pour déclencher l'envoi d'une première entrée MAC à un plan d'acheminement du premier dispositif NVE, où une interface sortante dans la première entrée MAC est l'interface locale du premier dispositif NVE, et la première entrée MAC comprend l'adresse MAC et l'ID MAC-VLAN qui sont transportés dans l'itinéraire d'annonce MAC/IP.

19. Premier dispositif NVE selon la revendication 18, dans lequel l'itinéraire d'annonce MAC/IP est transporté dans des informations d'accessibilité de couche réseau atteignable multiprotocole, MP_REACH_NLRI, un champ de saut suivant dans le MP_REACH_NLRI comprend une adresse VTEP commune, l'itinéraire d'annonce MAC/IP comprend en outre la deuxième adresse VTEP, et l'adresse VTEP commune comprise dans le champ de saut suivant dans le MP_REACH_NLRI est la même que l'adresse VTEP commune comprise dans le champ

d'adresse IP du routeur d'origine dans le deuxième itinéraire IMET ; et

l'unité de traitement (1004) est en outre configurée pour déclencher l'envoi d'une deuxième entrée MAC au plan d'acheminement du premier dispositif NVE sur la base de l'adresse VTEP commune comprise dans le champ de saut suivant dans le MP_REACH_NLRI et de la deuxième adresse VTEP comprise dans l'itinéraire d'annonce MAC/IP, où une interface sortante dans la deuxième entrée MAC est une interface locale qui est du premier dispositif NVE et qui se connecte au tunnel VXLAN du premier dispositif NVE au deuxième dispositif NVE, et la deuxième entrée MAC comprend l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP.

20. Premier dispositif NVE selon l'une quelconque des revendications 13 à 15, dans lequel un dispositif CE est connecté au deuxième dispositif NVE ;

l'unité de réception (1002) est en outre configurée pour recevoir un itinéraire d'annonce MAC/IP en provenance du deuxième dispositif NVE, où l'itinéraire d'annonce MAC/IP est transporté dans le MP_REACH_NLRI, un champ de saut suivant dans MP_REACH_NLRI comprend une adresse VTEP commune, l'itinéraire d'annonce MAC/IP comprend un ID MAC-VLAN et la deuxième adresse VTEP, l'ID MAC-VLAN est utilisé pour indiquer un VLAN auquel une adresse MAC transportée dans l'itinéraire d'annonce MAC/IP appartient, et l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP est une adresse MAC du dispositif CE, ou l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP est une adresse MAC d'un hôte dans un site EVPN administré par le dispositif CE ; l'unité de traitement (1004) est en outre configurée pour déterminer, sur la base d'un ESI et de l'ID MAC-VLAN qui sont transportés dans l'itinéraire d'annonce MAC/IP, que le premier dispositif NVE n'a pas d'interface locale qui se connecte à une liaison Ethernet entre le dispositif CE et le premier dispositif NVE, où l'ESI transporté dans l'itinéraire d'annonce MAC/IP est utilisé pour indiquer un ES entre le dispositif CE et le deuxième dispositif NVE ; et

l'unité de traitement (1004) est en outre configurée pour déclencher l'envoi d'une entrée MAC à un plan d'acheminement du premier dispositif NVE sur la base de l'adresse VTEP commune comprise dans le champ de saut suivant dans le MP_REACH_NLRI et de la deuxième adresse VTEP comprise dans l'itinéraire d'annonce MAC/IP, où une interface sortante dans l'entrée MAC est une interface locale qui est du premier dispositif NVE et qui se connecte au tunnel VXLAN du premier dispositif NVE au deuxième

dispositif NVE, et l'entrée MAC comprend l'adresse MAC transportée dans l'itinéraire d'annonce MAC/IP.

5

10

15

20

25

30

35

40

45

50

55

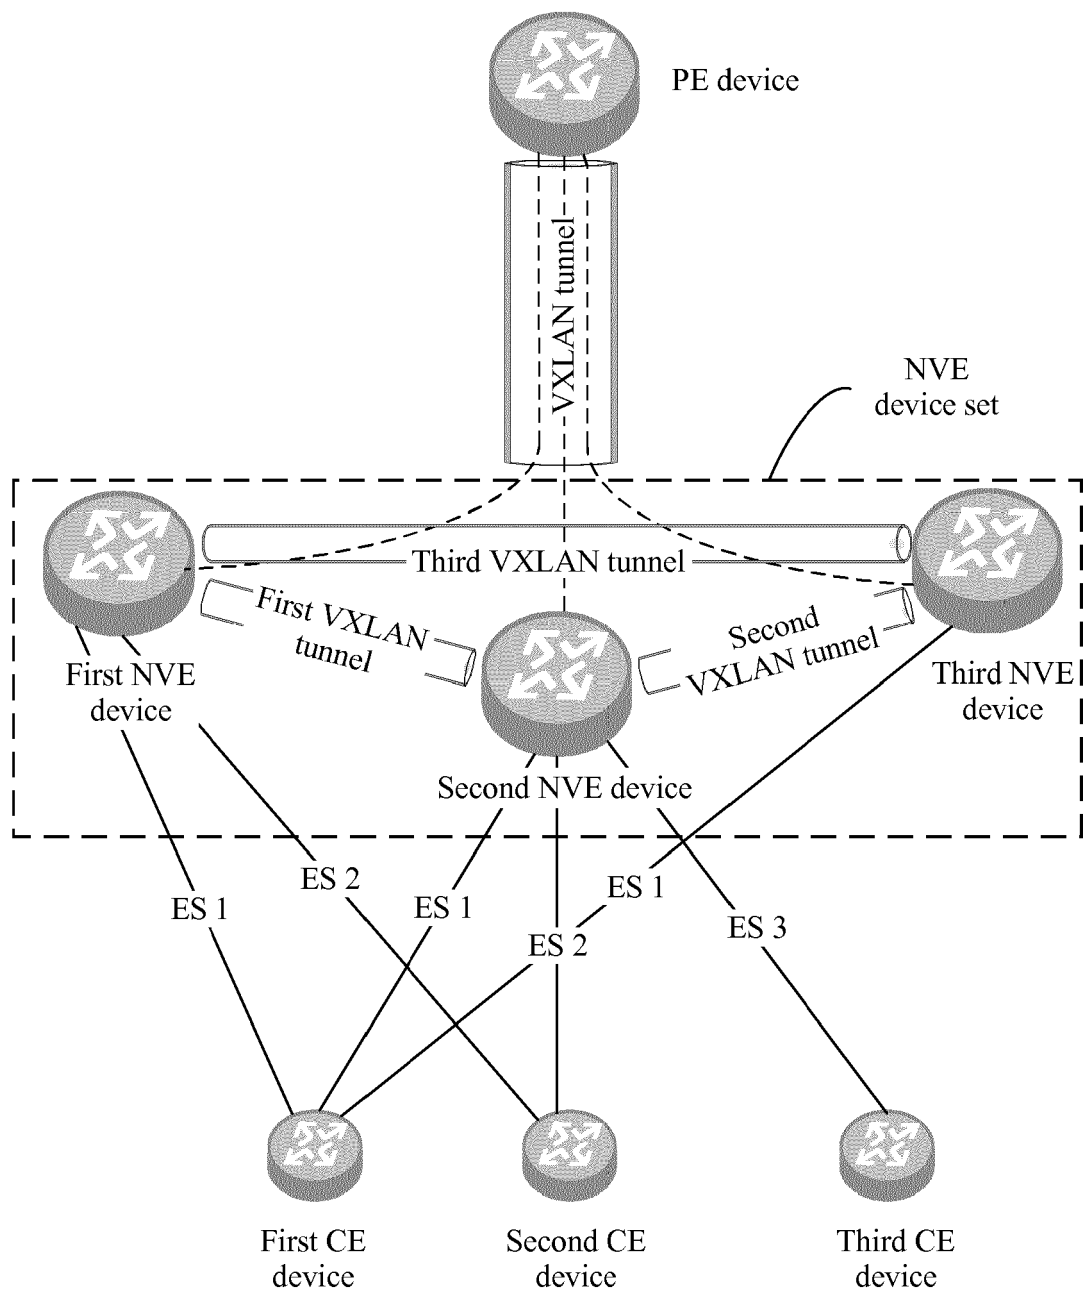


FIG. 1

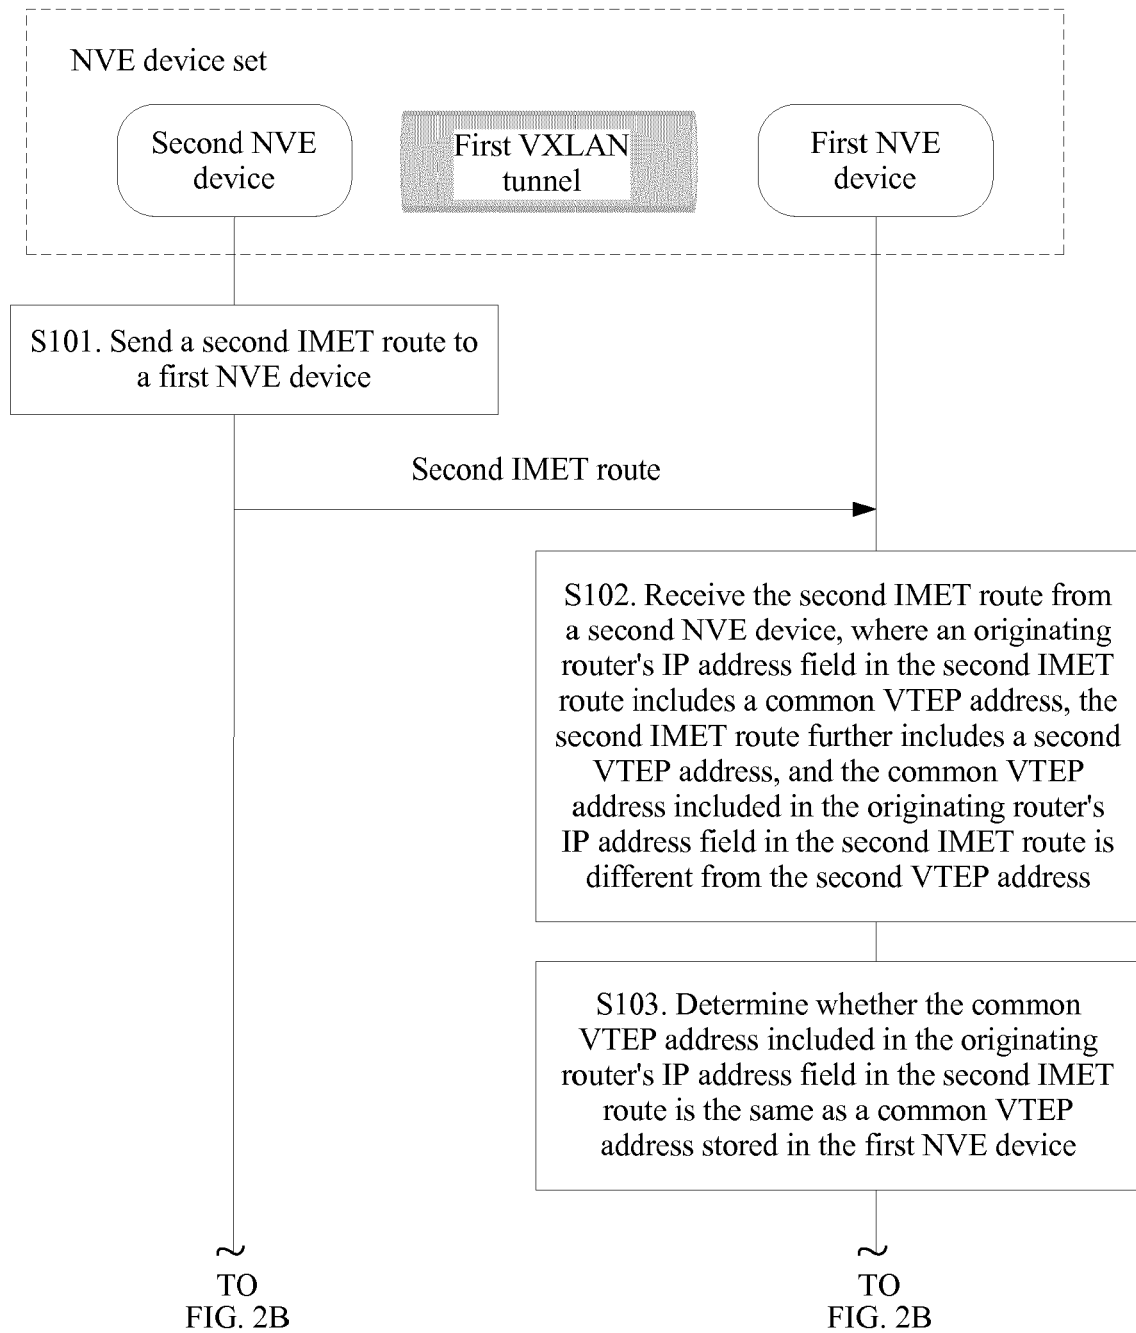


FIG. 2A

CONT.
FROM
FIG. 2A

CONT.
FROM
FIG. 2A

S104. When determining that the common VTEP address included in the originating router's IP address field in the second IMET route is the same as the common VTEP address stored in the first NVE device, the first NVE device establishes a VXLAN tunnel from the first NVE device to the second NVE device based on the second VTEP address in the second IMET route

S105. Send a first IMET route to the second NVE device, where an originating router's IP address field in the first IMET route carries the common VTEP address stored in the first NVE device, the first IMET route further includes a first VTEP address, the first VTEP address is used by the second NVE device to establish a VXLAN tunnel from the second NVE device to the first NVE device, and the common VTEP address stored in the first NVE device is different from the first VTEP address

First IMET route

FIG. 2B




FIG. 3

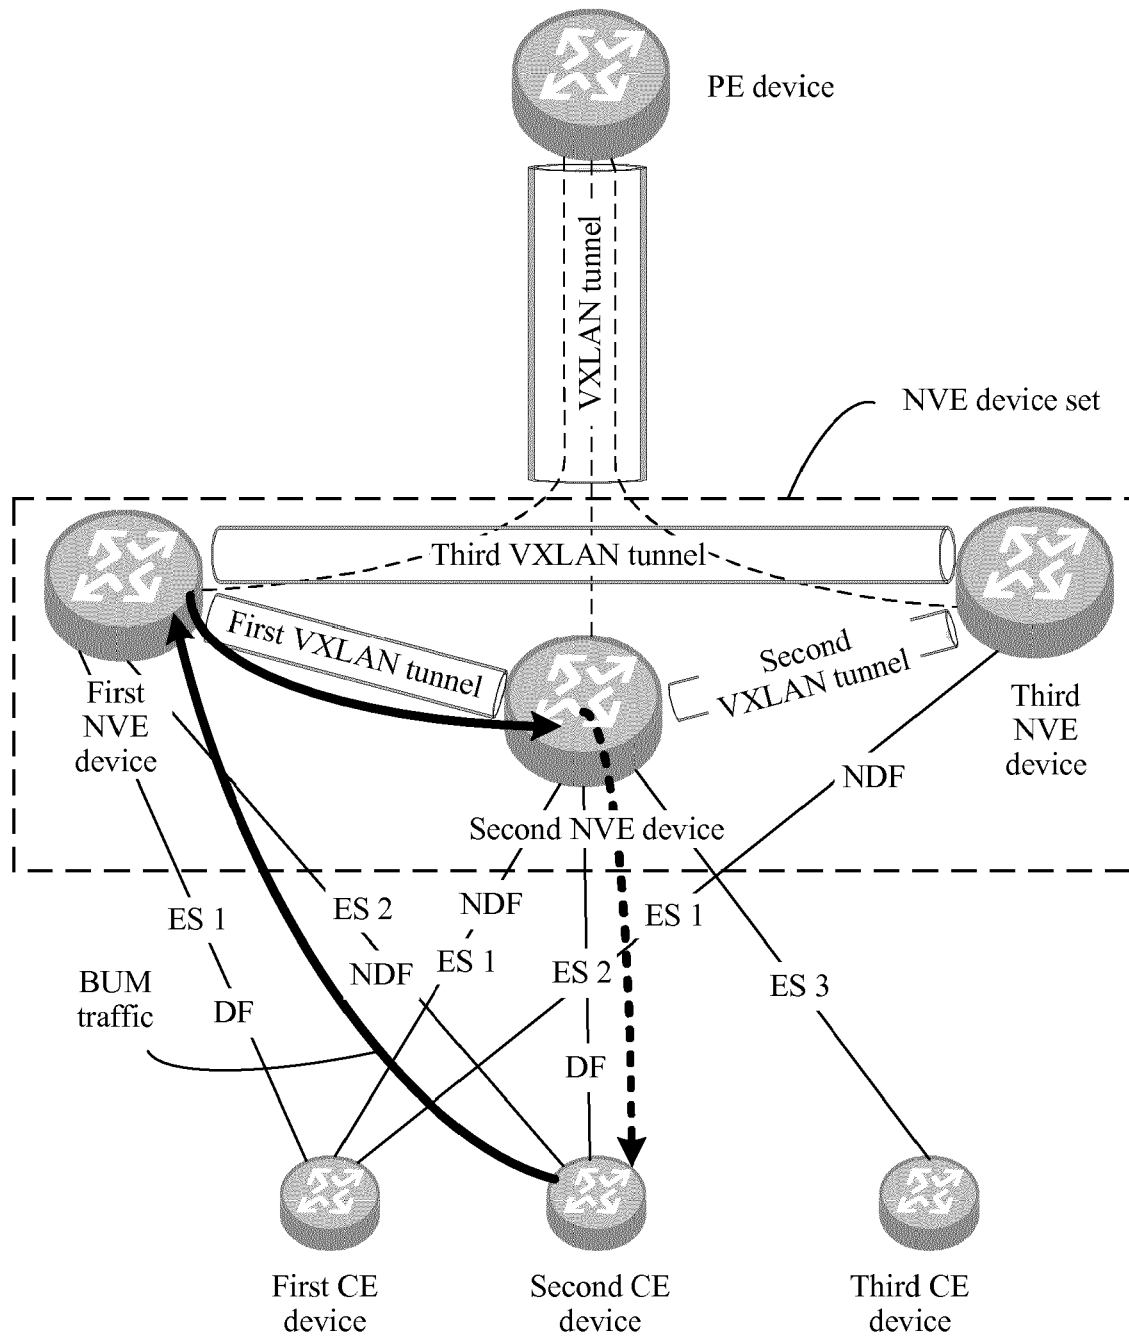


FIG. 4



FIG. 5




FIG. 6




FIG. 7




FIG. 8

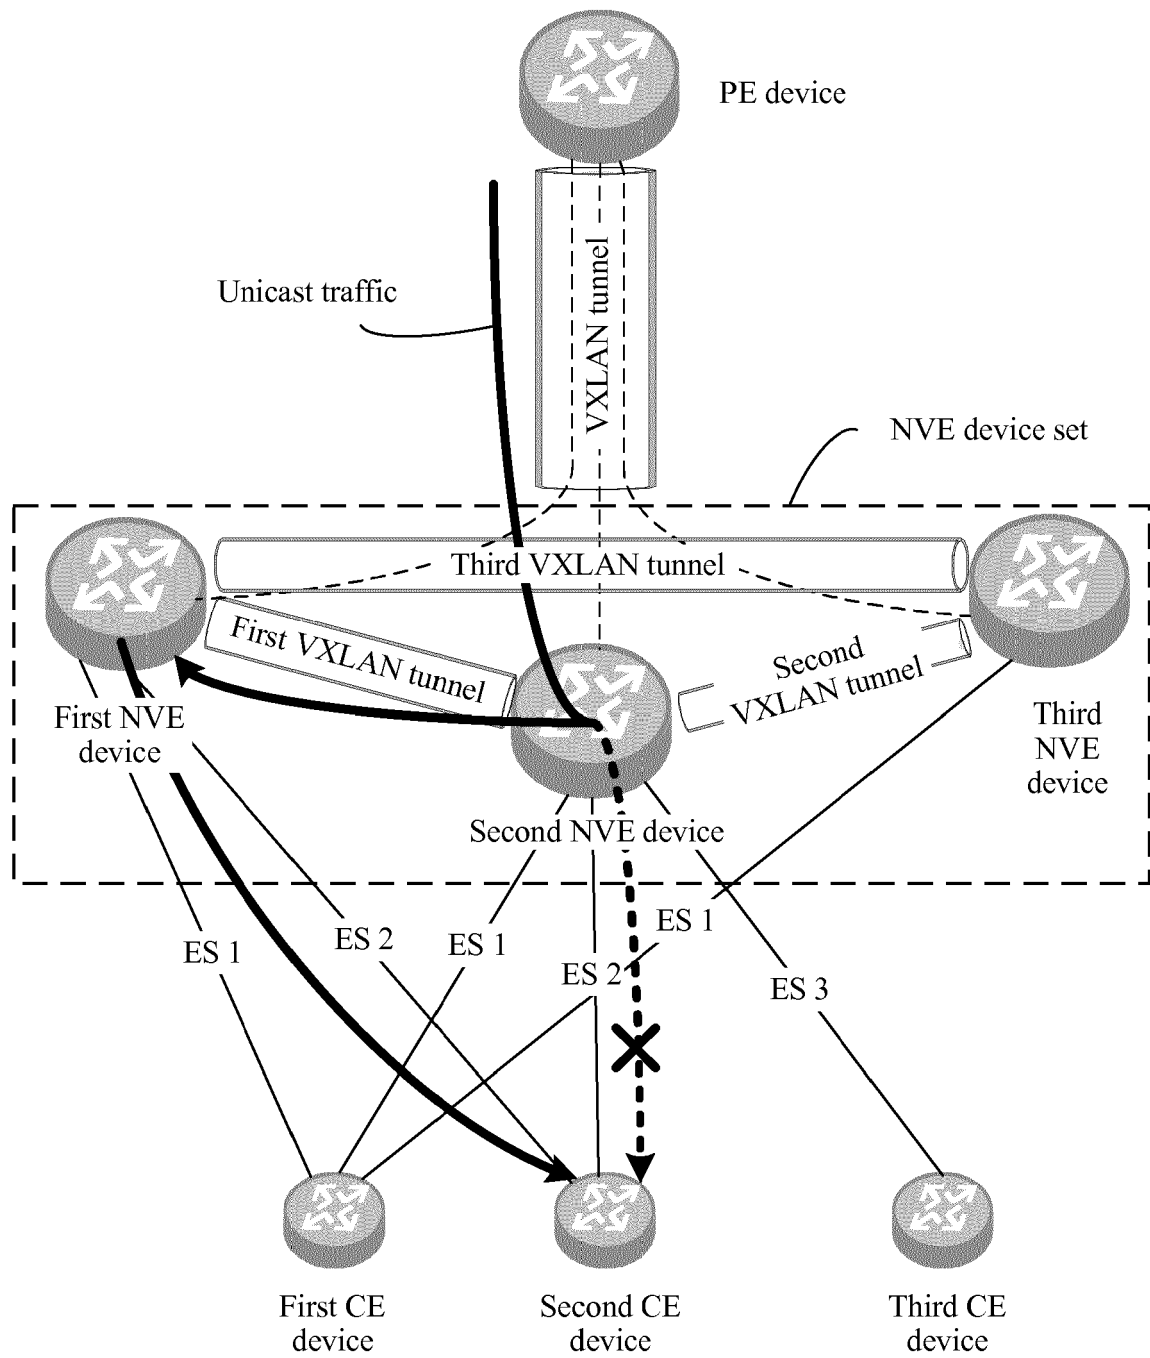


FIG. 9

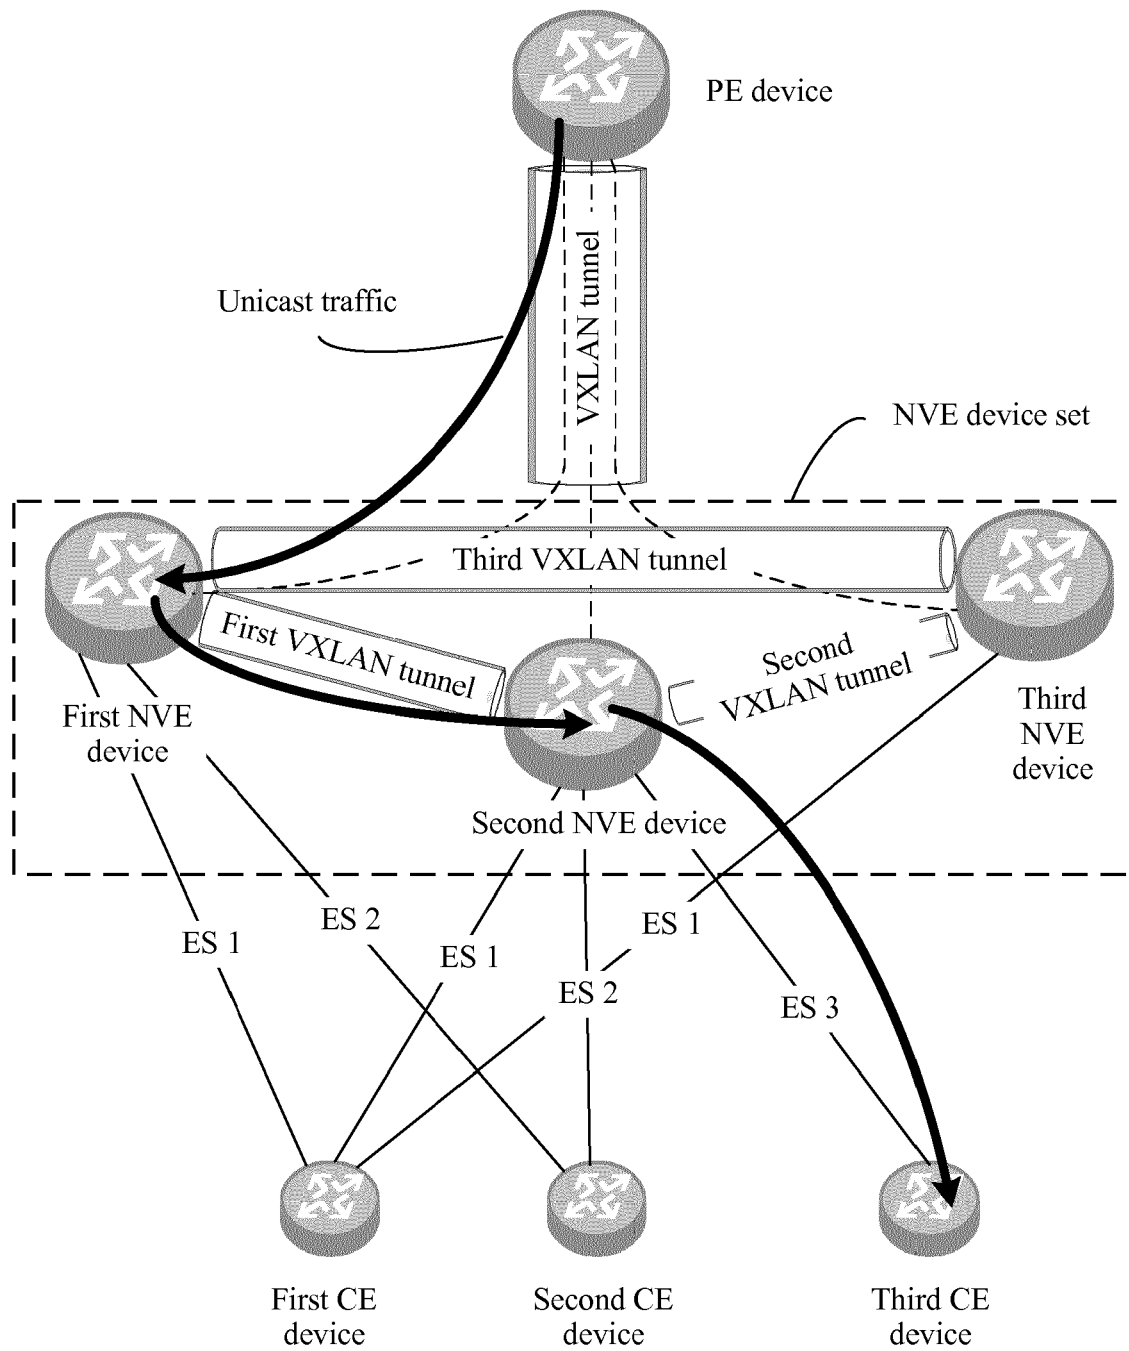


FIG. 10




FIG. 11

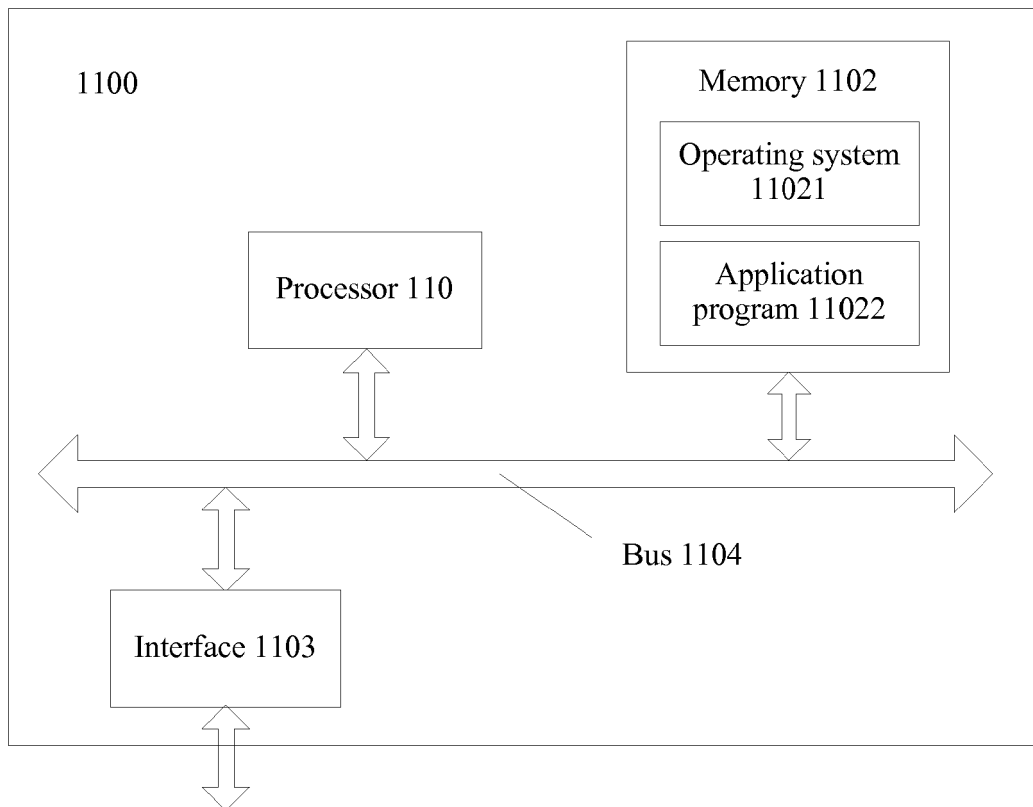


FIG. 12

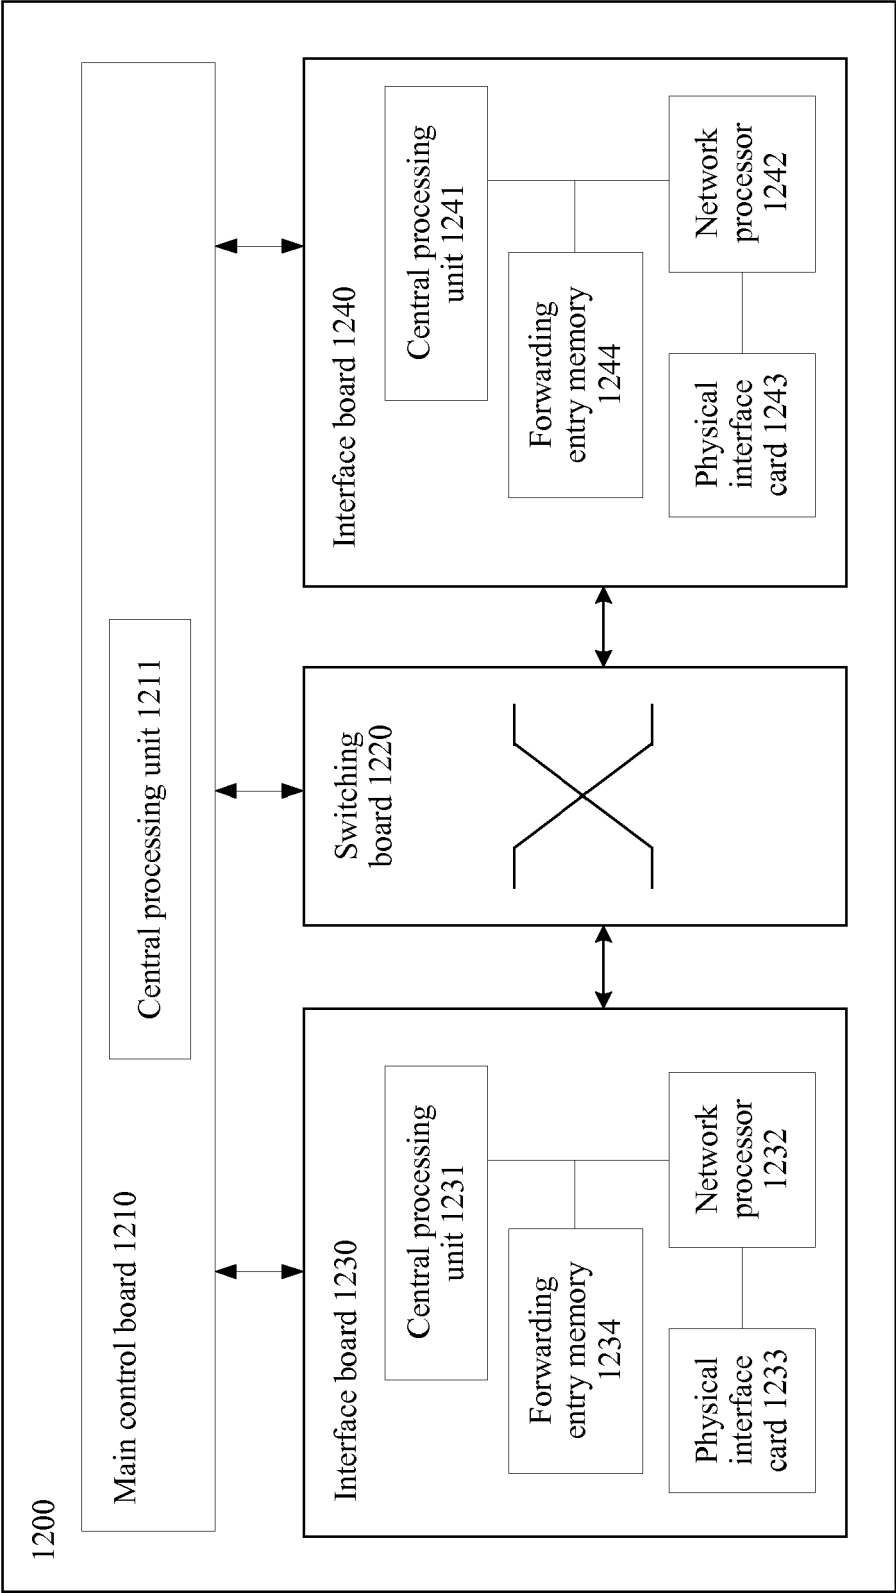


FIG. 13

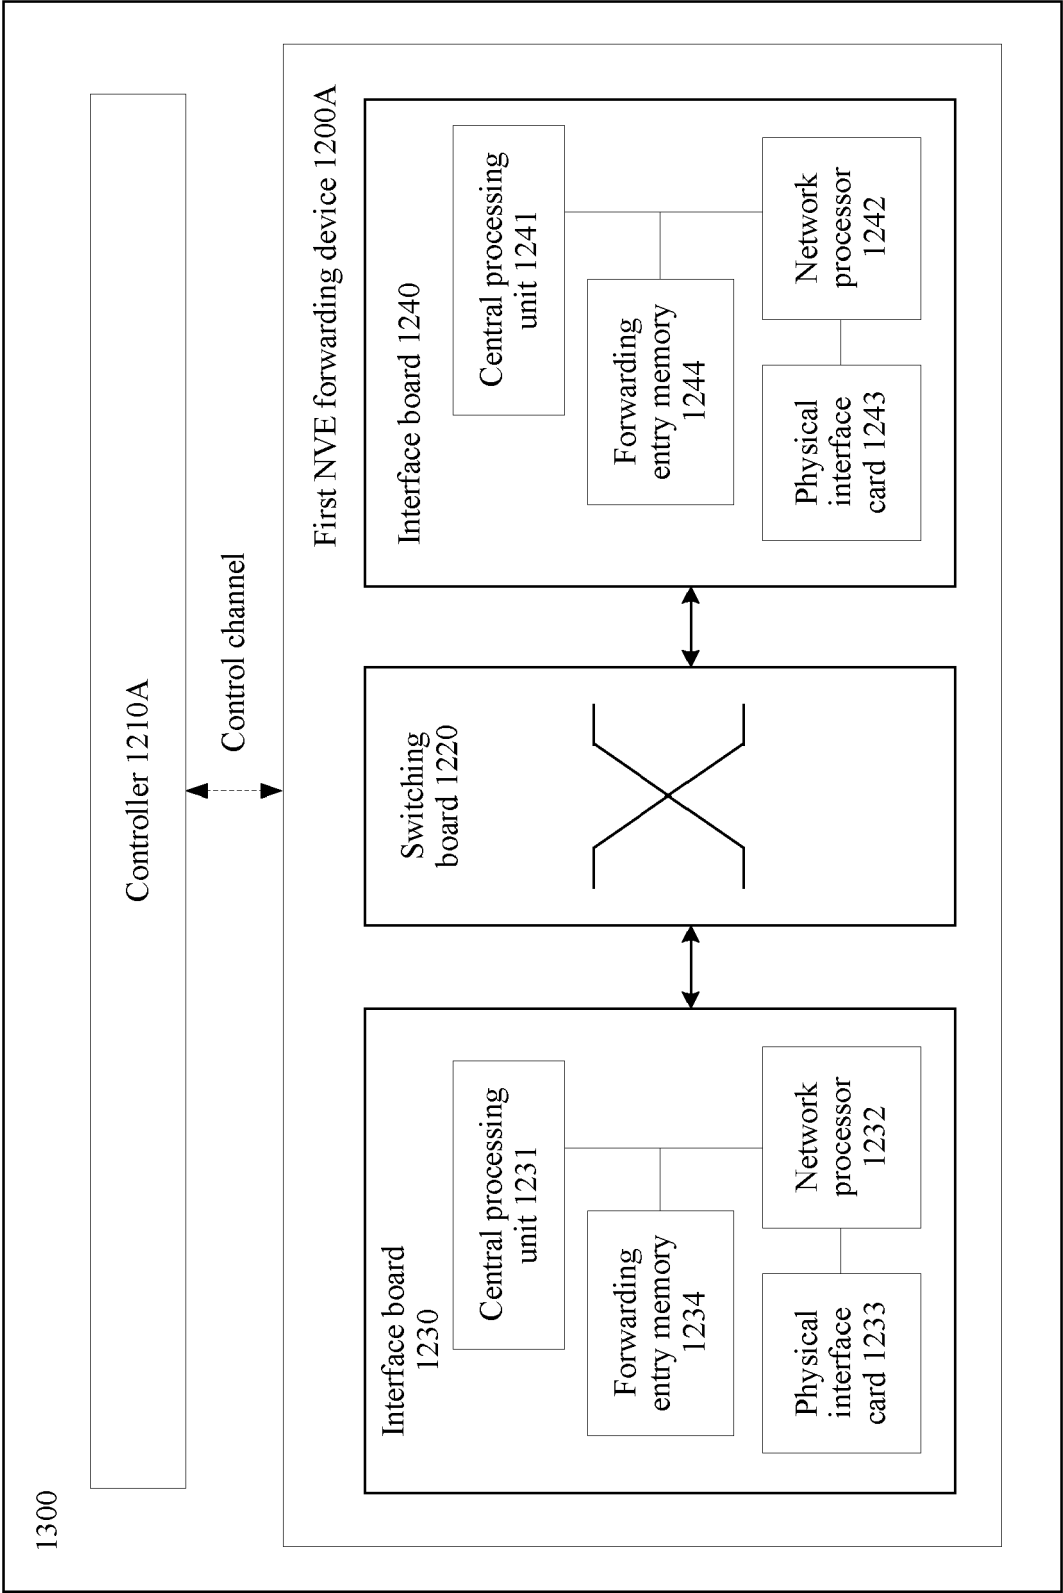


FIG. 14

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 2016134513 A1 [0006]