(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) **EP 3 108 625 B1**

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention of the grant of the patent:
01.08.2018 Bulletin 2018/31

(21) Application number: **15760834.0**

(22) Date of filing: **06.03.2015**

(51) Int Cl.:
H04L 12/46 *(2006.01)* ***H04L 12/64*** *(2006.01)*
H04L 12/715 *(2013.01)* ***H04L 12/717*** *(2013.01)*
H04L 12/751 *(2013.01)*

(86) International application number:
PCT/CN2015/073791

(87) International publication number:
WO 2015/135444 (17.09.2015 Gazette 2015/37)

(54) **VIRTUAL PRIVATE NETWORK MIGRATION AND MANAGEMENT IN CENTRALLY CONTROLLED NETWORKS**

MIGRATION UND VERWALTUNG EINES VIRTUELLEN PRIVATEN NETZWERKES IN ZENTRAL GESTEUERTEN NETZWERKEN

MIGRATION ET GESTION DE RÉSEAU PRIVÉ VIRTUEL, ET GESTION DANS DES RÉSEAUX À COMMANDE CENTRALE

(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: **11.03.2014 US 201414204827**

(43) Date of publication of application:
28.12.2016 Bulletin 2016/52

(73) Proprietor: **HUAWEI TECHNOLOGIES CO., LTD.**
Shenzhen Guangdong 518129 (CN)

(72) Inventors:
- **CHEN, Huaimo**
  Bolton, MA 01740 (US)
- **LI, Renwei**
  Fremont, CA 94539 (US)
- **DONG, Xuesong**
  Shenzhen
  Guangdong 518129 (CN)

(74) Representative: **Körber, Martin Hans**
Mitscherlich PartmbB
Patent- und Rechtsanwälte
Sonnenstrasse 33
80331 München (DE)

(56) References cited:
CN-A- 1 812 363
CN-A- 101 471 879
CN-A- 102 449 964
US-A1- 2008 002 697
US-A1- 2011 142 053

- **LUYUAN FANG DAVID WARD REX FERNANDO CISCO MARIA NAPIERALA AT&T NABIL BITAR VERIZON: "BGP/MPLS IP VPN Virtual PE; draft-fang-l3vpn-virtual-pe-04.txt", BGP/MPLS IP VPN VIRTUAL PE; DRAFT-FANG-L3VPN-VIRTUAL-PE-04.TXT, INTERNET ENGINEERING TASK FORCE, IETF; STANDARDWORKINGDRAFT, INTERNET SOCIETY (ISOC) 4, RUE DES FALAISES CH- 1205 GENEVA, SWITZERLAND, 21 October 2013 (2013-10-21), pages 1-24, XP015095580, [retrieved on 2013-10-21]**
- **LUYUAN FANG JOHN EVANS DAVID WARD REX FERNANDO JOHN MULLOOLY CISCO NING SO TATA COMMUNICATIONS NABIL BITAR VERIZON: "BGP/MPLS IP VPN Virtual CE; draft-fang-l3vpn-virtual-ce-02.txt", BGP/MPLS IP VPN VIRTUAL CE; DRAFT-FANG-L3VPN-VIRTUAL-CE-02.TXT, INTERNET ENGINEERING TASK FORCE, IETF; STANDARDWORKINGDRAFT, INTERNET SOCIETY (ISOC) 4, RUE DES FALAISES CH- 1205 GENEVA, SWITZERLAND, 21 October 2013 (2013-10-21), pages 1-18, XP015095579, [retrieved on 2013-10-21]**

Printed by Jouve, 75001 PARIS (FR)

(Cont. next page)

- **MARIA NAPIERALA AT&T LUYUAN FANG CISCO: "Requirements for Extending BGP/MPLS VPNs to End-Systems; draft-fang-l3vpn-end-system-requirements-02.txt", REQUIREMENTS FOR EXTENDING BGP/MPLS VPNS TO END-SYSTEMS; DRAFT-FANG-L3VPN-END-SYSTEM-REQUIREMENTS-0 2.TXT, INTERNET ENGINEERING TASK FORCE, IETF; STANDARDWORKINGDRAFT, INTERNET SOCIETY (ISOC) 4, RUE DES FALAISES CH- 1205 GENEVA, SWITZERLAND, 22 October 2013 (2013-10-22), pages 1-15, XP015095850, [retrieved on 2013-10-22]**

Description

TECHNICAL FIELD

[0001] The present invention relates generally to virtual private network (VPN), and in particular embodiments, to techniques and mechanisms for VPN migration and management in centrally controlled networks.

BACKGROUND

[0002] A virtual private network (VPN) may refer to the extension of a private network across a service provider's network. A VPN may allow a computer (or other network node) to communicate with a private network across shared networks as if the computer was directly connected to the private network, while benefiting from the functionality, security and/or management policies of the private network.

[0003] A typical VPN utilizes distributed control of the network. However, with increasing interest in software-defined networking (SDN), centralized control of networks is becoming more of interest because centralized control of networks is a feature of SDN. However, it is not currently known how to migrate a distributed or decentralized VPN to a centrally controlled VPN or how to manage a centrally controlled VPN. Thus, there is a need for migration of control of a VPN from distributed control to centralized control, especially without interrupting service.

[0004] US 2011/0142053 A1 discloses a method communicatively coupling virtual private networks to virtual machines within distributive computing networks. A disclosed example method includes receiving a request to provision a virtual machine from a virtual private network, determining a host for the virtual machine within a distributive computing network, creating the virtual machine within the host, communicatively coupling the virtual machine to a virtual local area network switch within the distributive computing network, configuring a portion of a router to be communicatively coupled to the virtual machine via the virtual local area network switch by specifying an address space within the router associated with at least one of the virtual machine or the virtual private network communicatively coupled to the router, and communicatively coupling the portion of the router to the virtual private network.

[0005] LUYUAN FANG DAVID WARD REX FERNANDO CISCO MARIA NAPIERALA AT&T NABIL BITAR VERIZON, "BGP/MPLS IP VPN Virtual PE; draft-fang-13vpn-virtual-pe-04.txt", BGP/MPLS IP VPN VIRTUAL PE; DRAFT-FANG-L3VPN-VIRTUAL-PE-04.TXT, INTERNET ENGINEERING TASK FORCE, IETF; STANDARDWORKINGDRAFT, INTERNET SOCIETY (ISOC) 4, RUE DES FALAISES CH- 1205 GENEVA, SWITZERLAND, (20131021), pages 1 - 24, XP015095580, describes the architecture solutions for BGP/MPLS IP Virtual Private Networks (VPNs) with virtual Provider Edge (vPE) routers.

SUMMARY

[0006] Disclosed herein are example embodiments for migrating a decentralized VPN to a VPN controlled by a central controller. The present invention is defined in the independent claims. Preferred embodiments are defined in the dependent claims.

[0007] The migration may be performed without service interruption. Prior to migration, a native border gateway protocol (BGP) stack exists on each edge node and may be used to distribute VPN routes and labels in the decentralized network. In example embodiments, a central controller may manage an installation of two software agents on each edge node in a backbone network across which at least one VPN has been established. A first software agent on a first edge node may take over a native BGP peer session with a second edge node. After taking over of the BGP peer session, there may be no need for any BGP software. The first edge node may be coupled to a first site, and the second edge node may be coupled to a second site. The second software agent may form an adjacency with the first site and communicates label and route information between the first site, the central controller, and the second edge node. After software agents on each edge node have been installed and information from each site distributed to the central controller, there may be no need for a BGP session any longer as the duties of the BGP session are accomplished by the central controller. Thus, the agents that are responsible for the BGP sessions are removed from the edge nodes and software agents that provide for communication with the central controller remain on the edge nodes to manage the at least one VPN that has been migrated.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] For a more complete understanding of this disclosure, reference is now made to the following brief description, taken in connection with the accompanying drawings and detailed description, wherein like reference numerals represent like parts.

FIG. 1 is an example embodiment of a network.
FIG. 2 is an example embodiment of a network.
FIG. 3 is a flowchart of an example embodiment of a method.
FIG. 4 is an example embodiment of a network.
FIG. 5 is an example embodiment of an edge node.
FIG. 6 is an example embodiment of a central controller.

DETAILED DESCRIPTION

[0009] It should be understood at the outset that, although an illustrative implementation of one or more ex-

ample embodiments are provided below, the disclosed systems and/or methods may be implemented using any number of techniques, whether currently known or in existence. The disclosure should in no way be limited to the illustrative implementations, drawings, and techniques illustrated below, including the exemplary designs and implementations illustrated and described herein, but may be modified within the scope of the appended claims along with their full scope of equivalents.

[0010] FIG. 1 is an example embodiment of a network 100 in which at least one VPN may be established. The network 100 comprises a local network 140, sometimes referred to as Site A, a local network 150, sometimes referred to as Site B, and a backbone network 130. The backbone network 130 may be any sort of backbone network, such as a mobile backhaul network. The backbone network 130 comprises edge nodes, denoted as node A 110 and node B 120, and internal nodes 105. The nodes in backbone network 130 may be referred to herein as backbone nodes. Three internal nodes 105 are shown for illustrative purposes, but there may be any number of internal nodes. As used herein, "node" may be synonymous with "router" or "switch". The local networks 140 and 150 each may comprise one or more nodes or end devices (not shown) in addition to the edge nodes 115 and 125. If the backbone network 130 is a mobile backhaul network, node A 110 may be a cell site gateway, internal nodes 105 may be aggregation site gateways, and node B 120 may be a radio network controller (RNC) side gateway.

[0011] A VPN may be established between local networks 140 and 150 via backbone network 130. The VPN may be a layer 3 (L3) VPN established using multiprocotol label switching (MPLS) and border gateway protocol (BGP) as understood by a person of ordinary skill in the art. For example, each route within a VPN may be assigned an MPLS label, and when BGP distributes a VPN route, BGP also distributes an MPLS label for that route. The VPN may connect local area networks 140 and 150 via the path illustrated connecting nodes A 110 and B 120 via internal nodes 105. In such a VPN node A 110 and node B 120 may be referred to as VPN provider edge nodes. Control of the VPN is performed in a distributed manner.

[0012] Disclosed herein are systems, methods, apparatuses, and computer program products for migrating and managing control of one or more VPNs from distributed control to centralized control without interrupting service. After migration, all VPNs in a service area may be controlled by a central controller. Native protocols, such as BGP, on backbone nodes may be replaced by communication software installed on backbone nodes by a central controller to provide for migration of one or more VPNs. After migration is complete, part of the communication software may be removed and VPNs that backbone nodes are a part of may be managed by the central controller.

[0013] Software-defined networking (SDN) is a relatively new technology, which, among other things, introduces the notion of a central controller. Conventional packet networks utilize distributed control, but in SDN part of the control of a network may be placed under the control of a central controller. The portion of a network controlled by a central controller may be referred to as an SDN domain, and the portion of a network not under the control of a central controller may be referred to as a non-SDN domain.

[0014] FIG. 2 is an example embodiment of a network 200 that may be the same as network 100, except for the introduction of a central controller 210. In the interest of conciseness, the description of elements of network 200 that are the same as the elements of network 100 are not described. In addition to the elements described earlier, the network 200 comprises a central controller 210 as shown in FIG. 2. The central controller 210 may be capable of communicating with each of nodes A 110, node B 120, and internal nodes 105 and performing a migration of a VPN running over backbone network 130 and connecting local networks 140 and 150. There are a variety of ways known to a person of ordinary skill in the art in which the central controller 210 may communicate with node A 110, node B 120, and internal nodes 105, and one of those ways is through the use of the OpenFlow protocol.

[0015] FIG. 3 is a flowchart of an example embodiment of a method 300 for migrating a VPN. The blocks of FIG. 3 are discussed with reference to FIG. 2 as an example network architecture for migrating a VPN in which the steps of the method 300 may be performed. Step 310 is performed using coordinated effort of a central controller and a first edge node, such as central controller 210 and node A 110, respectively. In step 310 a central controller initiates and manages installation of a first software agent and a second software agent on a first edge node. The installation of the first and second software agents may be performed as an in-service software upgrade (ISSU). The central controller may send the first software agent and the second software agent to the first edge node for installation on the first edge node. The first edge node may be responsible for helping to provide a VPN across a backbone network, such as network 130. The first software agent may be referred to as a backward compatible agent for a VPN (BCAV), and the second software agent may be referred to as a backward compatible agent for a site (BCAS). A software agent may also be referred to as a client.

[0016] The first software agent performs or causes the first edge node, such as node A 110, to take over a BGP peer session with a remote VPN edge node, such as node B 120. The remote VPN edge node may also be referred to as a second edge node. The BGP peer session may be used to exchange information between the first edge node and the remote VPN edge node. The first software agent may be a replacement of or substitute for native BGP software residing on the first edge node, and the BGP peer session may take the place of a BGP peer

session (which may be referred to as a native BGP peer session) running using the native BGP software. In one embodiment, the first software agent performs or causes the first edge node to (1) obtain a VPN label for a VPN from the central controller and send this label to the remote VPN edge node (e.g., via the BGP peer session); (2) receive a VPN label from the remote VPN edge node (e.g., via the BGP peer session) and send the label to the central controller; and (3) receive remote VPN routes from the remote VPN edge node (e.g., via the BGP peer session) and send them to the central controller. The remote VPN routes may be VPN routes for a site to which the remote VPN edge node is connected. In another embodiment, the first software agent performs or causes the first edge node to (1) obtain a VPN label for a VPN route from the central controller and send this label with the VPN route to the remote VPN edge node; and (2) receive a VPN label with a VPN route from the remote VPN edge node and send the label with the VPN route to the central controller.

[0017] The second software agent performs or causes the first edge node to perform the following tasks: (1) take over a protocol adjacency with a site to which the first edge node is connected, such as site A 140; (2) obtain VPN routes from the site through the protocol adjacency and send the routes to the central controller; (3) receive VPN routes from the remote VPN edge node via the central controller and send them to the site.

[0018] In step 310, the central controller may control the first software agent (in the first edge node) for taking over a native BGP peer session with the second edge node (i.e., migrating the native BGP peer session to a BGP peer session controlled by the first software agent smoothly) and control the second software agent for taking over a native protocol adjacency with the site (i.e., migrating the native adjacency to an adjacency controlled by the second software agent smoothly). The central controller may install the first software agent and the second software agent on the first edge node via an ISSU. In an embodiment, the central controller may allocate a VPN label for a VPN and send it to the first software agent, as well as send local VPN routes to the first software agent. The central controller may also: (1) receive the VPN label for the second edge node from the first software agent; (2) add a forwarding entry to the VPN routing and forwarding (VRF) table in the first edge node when receiving a new remote VPN route; and (3) delete a forwarding entry from the VRF in the first edge node when receiving a remote VPN route withdrawal. In another embodiment, the central controller may allocate a VPN label for a VPN route and send the VPN route with the VPN label to the first software agent. The central controller may also: (1) receive a VPN route with a VPN label for the second edge node from the first software agent; (2) add a forwarding entry to the VRF table in the first edge node if the VPN route is a new route; and (3) delete a forwarding entry from the VRF in the first edge node when receiving a remote VPN route withdrawal. One implication to note is that remote VPN routes from the second edge router are communicated via the first software agent in the first edge router to the central controller. Then the central controller sends the remote VPN routes to the second software agent, and the second software agent sends them to the first site. A new remote VPN route may be added subsequent to completion of step 340. A VRF table may be stored in the first edge node and used by the first edge node for routing and/or forwarding VPN packets received by the first edge node.

[0019] In step 320, an ISSU may be performed on nodes internal to the backbone network, e.g., nodes 105 as shown in FIG. 1, to turn over control of the VPN on the internal nodes to the central controller. During the migration of an internal node, an MPLS LSP tunnel between the first edge node and the second edge node is maintained and is not affected. In the meantime, the first and second software agents on the first edge node continue to operate, e.g., the BGP session between the first and second edge nodes is maintained during upgrade of the internal nodes.

[0020] Step 330 is performed using coordinated effort of the central controller and the second edge node, e.g., node B 120. In step 330 the central controller initiates and manages installation of a third software agent and a fourth software agent on the second edge node. The installation of the third and fourth software agents may be performed as an in-service software upgrade (ISSU). The second edge node may be responsible for helping to provide a VPN across the backbone network, such as network 130. The third software agent may be a BCAV, and the fourth software agent may be a BCAS. The third software agent may perform the same functions on the second edge node as the first software agent performs on the first edge node. Also, the fourth software agent may perform the same functions on the second edge node as the second software agent performs on the first edge node.

[0021] The third software agent performs or causes the second edge node to take over a BGP peer session with a remote VPN edge node (i.e., the first edge node), such as node A 110. Note that in step 310 the first edge node takes over a BGP peer session with the second edge node. The peer session in step 310 was a peer session with the second edge node using the native BGP capability of the second edge node. The third software agent maintains the BGP peer session with the first edge node while allowing for the removal of the native BGP capability from the second edge node. After step 340 is complete, the third software agent is removed or deleted from the second edge node after the central controller can (1) receive a VPN route from the second software agent on the first edge node and send the VPN route to the fourth software agent on the second edge node, and (2) receive a VPN route from the fourth software agent on the second edge node and send the VPN route to the second software agent on the first edge node because it is no longer needed. In one embodiment, the third soft-

ware agent also performs or causes the second edge node to: (1) obtain a VPN label for a VPN from the central controller and send this label to the remote VPN edge node; (2) receive a VPN label from the remote VPN edge node and send the label to the central controller; and (3) receive remote VPN routes from the remote VPN edge node and send them to the central controller. In another embodiment, the third software agent performs or causes the second edge node to (1) obtain a VPN label for a VPN route from the central controller and send the VPN route with the label to the remote VPN edge node; and (2) receive a VPN route with a VPN label from the remote VPN edge node and send the VPN route with the VPN label to the central controller.

[0022] The fourth software agent performs or causes the second edge node to perform the following tasks: (1) take over a protocol adjacency with a second site to which the second edge node is connected, such as site B 150; (2) obtain VPN routes from the second site through the protocol adjacency and send the routes to the central controller; (3) receive routes from the first edge node via the central controller and send them to the second site.

[0023] In step 330, the central controller may control the third software agent (in the second edge node) for taking over a BGP peer session with the first edge node and control the fourth software agent for taking over a protocol adjacency with the second site. The central controller may install the third software agent and the fourth software agent on the second edge node via an ISSU. In an embodiment, the central controller may: (1) allocate a VPN label for a VPN and send it to the third software agent; (2) send local VPN routes to the third software agent; (3) receive the VPN label for the first edge node from the third software agent; (4) add a forwarding entry to the VRF table in the second edge node when receiving a new remote VPN route; and (5) delete a forwarding entry from the VRF in the second VPN edge node when receiving a remote VPN route withdrawal. In another embodiment, the central controller may: (1) allocate a VPN label for a VPN route and send the VPN route with the VPN label to the third software agent; (2) receive a VPN route with a VPN label for the first edge node from the third software agent; (3) add a forwarding entry to the VRF table in the second edge node if the VPN route is a new route; and (4) delete a forwarding entry from the VRF in the second edge node when receiving a remote VPN route withdrawal. One implication to note is that remote VPN routes from the first edge router are communicated via the third software agent in the second edge router to the central controller. Then the central controller sends the remote VPN routes to the fourth software agent, and the fourth software agent sends them to the second site.

[0024] The first and third software agents may be referred to as different instantiations of the same software agent (i.e., a BCAV). Likewise the second and fourth software agents may be referred to as different instantiations of the same software agent (i.e., a BCAS). One purpose of the BCAV is to remove protocols, such as BGP, from a network node by node and transfer information about sites to a central controller. In step 330, after installing the third software agent on the second edge node, the third software agent on the second edge node sends the information about Site A to the central controller in addition to maintaining the session between the first edge node and the second edge node. Moreover, native BGP capability on the second edge node is no longer needed and can be removed since the BGP peer session originally maintained by native BGP on the second edge node is maintained by the third software agent on the second edge node.

[0025] After performing the steps 310-330 as described above, step 340 is performed in which the peer session between the first and second edge nodes (e.g., node A 110 and node B 120) may be deleted because it is no longer needed. Also, the first software agent and the third software agent may deleted from the first and the second edge nodes, respectively, because they are no longer needed.

[0026] After steps 310-340 are complete migration of a VPN between the first and second edge nodes may be complete. The second software agent may remain on the first edge node, and the fourth software agent may remain on the second edge node. The second software agent runs on the first edge node and performs the following tasks after migration: (1) maintain a protocol adjacency with the site connected to the first edge node; (2) obtain updated routes from the site through the protocol adjacency and send them to the central controller; and (3) receive routes from the central controller (obtained from the second edge node) and send them to the site. The portion of a network controlled by a central controller may be referred to as an SDN domain. Thus, after migration, the backbone network, such as backbone network 130, may be referred to as an SDN domain, whereas during migration, the nodes that have been placed under control by a central controller may be an SDN domain, and nodes not under control by the central controller may be referred to as a non-SDN domain. As migration to a central controller progresses, the SDN domain grows and the non-SDN domain shrinks.

[0027] After migration, the central controller performs the following tasks for an existing VPN: (1) control the second software agent for maintaining protocol adjacency with the site; (2) receive routes from the second software agent and send routes to the fourth software agent on the second edge node; (3) add a forwarding entry to the VRF on the second edge node when receiving a new route; and (4) delete a forwarding entry from the VRF when receiving a route withdraw. In an embodiment, in addition to the above tasks, for a newly configured VPN, the central controller performs the following tasks: (1) allocate a VPN label for a VPN node for the VPN and store it for the other VPN nodes in the same VPN; and (2) use this VPN label when updating VRFs in the VPN edge nodes for VPN routes. For example for the network

200, for a newly configured VPN, the central controller may: (1) allocate a VPN label for the first VPN edge node and store it for the second VPN edge node in the same VPN; and (2) use this VPN label when updating VRFs in the first and second VPN edge nodes for VPN routes. In another embodiment, for a newly configured VPN, the central controller performs the following tasks: (1) allocate a VPN label for a VPN node for each VPN route in the VPN and store it for the other VPN nodes in the same VPN; and (2) use this VPN label when updating VRFs in the VPN edge nodes for the VPN route.

[0028] FIG. 4 is an example embodiment of the network 200 after step 310 has been performed. As discussed above, in step 310, first and second software agents may be installed on a first edge node. The first and second software agents may be provided to node A 110 by the central controller 210. Referring to the network 200 in FIG. 4 to illustrate step 310, the communication paths between node B 120 and the central controller 210 and between the edge node 115 and the central controller 210 are illustrated by dashed lines. Thus, information may be communicated between the edge node 115 and the central controller 210 as facilitated by the second software agent, and information may be communicated between node B 120 and the central controller as facilitated by the first software agent. In step 340, the first software agent is removed because it is no longer needed to provide for communication between node A 110 and node B 120, as the central controller 210 disseminates information between these nodes.

[0029] FIG. 5 is an example embodiment of an edge node 400. The edge node 400 comprises ports 410, a transmitter/receiver (Tx/Rx) or transceiver 412, a processor 420, and a memory 422 configured as shown in FIG. 5. The edge node 400 may be an edge node in a backbone network that can support one or more VPNs. The edge node 400 may, for example, be configured as node A 110 or node B 120 described previously. The edge node 400 may comprise one or more ports 410 coupled to transceiver 412. Although only one transceiver 412 is shown for illustrative purposes, there may be more than one transceiver, such as one transceiver per port. The edge node 400 may comprise a processor 420 coupled to the transceiver 412 and configured to process the packets or otherwise determine to which network components to send the packets. The processor 420 may be implemented using hardware or a combination of hardware and software. Although illustrated as a single processor, the processor 420 is not so limited and may comprise multiple processors. The processor 420 may be implemented as one or more central processor unit (CPU) chips, cores (e.g., a multi-core processor), field-programmable gate arrays (FPGAs), application specific integrated circuits (ASICs), and/or digital signal processors (DSPs).

[0030] The edge router 400 may further comprise a memory 422. The memory 422 may be used to store instructions for carrying out the methods described herein, e.g., for instructions for software agents and forwarding and routing tables. The memory 422 may store instructions for a first software agent 440, instructions for a second software agent 442, and a VRF 444. The first software agent 440 may be the same as the BCAV discussed previously, and the second software agent 442 may be the same as BCAS. The installation of software agents 440 and 442 may be controlled by a central controller, such as the central controller 210. Furthermore, the first software agent 440 may be installed for a period of time before being removed as described previously. The memory 422 may comprise secondary storage, random access memory (RAM), and/or read-only memory (ROM) and/or any other type of storage. The secondary storage may comprise one or more disk drives or tape drives and is used for non-volatile storage of data and as an over-flow data storage device if the RAM is not large enough to hold all working data. The secondary storage may be used to store programs that are loaded into the RAM when such programs are selected for execution. The ROM is used to store instructions and perhaps data that are read during program execution. The ROM is a non-volatile memory device that typically has a small memory capacity relative to the larger memory capacity of the secondary storage. The RAM is used to store volatile data and perhaps to store instructions. Access to both the ROM and the RAM is typically faster than to the secondary storage.

[0031] The processor 420 and the transceiver 412 may also be configured to implement or support at least part of any of the schemes and methods described above, such as the steps of the method for migrating a VPN 300 performed in an edge node. Finally, although the ports 410 are illustrated as bidirectional ports coupled to a transceiver, one of ordinary skill in the art will understand that alternatively the edge node 400 may comprise unidirectional ports coupled to a transmitter or receiver, depending on whether the port is an ingress or egress port, respectively.

[0032] FIG. 6 is an example embodiment of a central controller 500. The central controller 500 comprises ports 510, a transmitter/receiver (Tx/Rx) or transceiver 512, a processor 520, and a memory 522 configured as shown in FIG. 6. The central controller 500 may be capable of communicating with nodes in a backbone network. For example, the central controller 500 may be the same as central controller 210 described previously. The ports 510, transceiver 512, processor 520 and memory 522 may share many of the same characteristics as the ports 410, transceiver 412, processor 420 and memory 422, respectively, described previously. In the interest of conciseness, only the characteristics of these elements that are different than those described previously are described here.

[0033] The memory 522 may be used to store instructions for carrying out the methods described herein, e.g., instructions for software agents that are communicated to edge nodes, such as the first software agent 440 and

second software agent 442 described previously, represented by 542 in FIG. 6. The memory 522 may also store information on VPN routes and corresponding labels, represented by 540 in FIG. 6. The processor 520 and the transceiver 512 may also be configured to implement or support at least part of any of the schemes and methods described above, such as the steps of the method for migrating a VPN 300 performed in a central controller. The central controller 500 may be one of a plurality or a cluster of central controllers that together perform the functions of a central controller 210 in the architecture of FIG. 2.

[0034] It is understood that by programming and/or loading executable instructions onto the edge node 400, at least one of the processor 420 and/or the memory 422 are changed, transforming the edge node 400 in part into a particular machine or apparatus (e.g., the edge nodes 110 or 120) having the functionality taught by the present disclosure. The executable instructions may be stored on the memory 422 and loaded into the processor 420 for execution. Similarly, it is understood that by programming and/or loading executable instructions onto the central controller 500, at least one of the processor 520 and/or the memory 522 are changed, transforming the central controller 500 in part into a particular machine or apparatus (e.g., the central controller 210) having the functionality taught by the present disclosure. The executable instructions may be stored on the memory 522 and loaded into the processor 520 for execution. It is fundamental to the electrical engineering and software engineering arts that functionality that can be implemented by loading executable software into a computer can be converted to a hardware implementation by well-known design rules. Decisions between implementing a concept in software versus hardware typically hinge on considerations of stability of the design and numbers of units to be produced rather than any issues involved in translating from the software domain to the hardware domain. Generally, a design that is still subject to frequent change may be preferred to be implemented in software, because re-spinning a hardware implementation is more expensive than re-spinning a software design. Generally, a design that is stable that will be produced in large volume may be preferred to be implemented in hardware, for example in an ASIC, because for large production runs the hardware implementation may be less expensive than the software implementation. Often a design may be developed and tested in a software form and later transformed, by well-known design rules, to an equivalent hardware implementation in an application specific integrated circuit that hardwires the instructions of the software. In the same manner, as a machine controlled by a new ASIC is a particular machine or apparatus, likewise a computer that has been programmed and/or loaded with executable instructions may be viewed as a particular machine or apparatus.

[0035] At least one embodiment is disclosed and variations, combinations, and/or modifications of the embodiment(s) and/or features of the embodiment(s) made by a person having ordinary skill in the art are within the scope of the disclosure. Alternative embodiments that result from combining, integrating, and/or omitting features of the embodiment(s) are also within the scope of the disclosure. Where numerical ranges or limitations are expressly stated, such express ranges or limitations may be understood to include iterative ranges or limitations of like magnitude falling within the expressly stated ranges or limitations (e.g., from about 1 to about 10 includes, 2, 3, 4, etc.; greater than 0.10 includes 0.11, 0.12, 0.13, etc.). For example, whenever a numerical range with a lower limit, Rl, and an upper limit, Ru, is disclosed, any number falling within the range is specifically disclosed. In particular, the following numbers within the range are specifically disclosed: R = Rl + k * (Ru - Rl), wherein k is a variable ranging from 1 percent to 100 percent with a 1 percent increment, i.e., k is 1 percent, 2 percent, 3 percent, 4 percent, 5 percent, ..., 50 percent, 51 percent, 52 percent, ..., 95 percent, 96 percent, 97 percent, 98 percent, 99 percent, or 100 percent. Moreover, any numerical range defined by two R numbers as defined in the above is also specifically disclosed. The use of the term "about" means +/- 10% of the subsequent number, unless otherwise stated. Use of the term "optionally" with respect to any element of a claim means that the element is required, or alternatively, the element is not required, both alternatives being within the scope of the claim. Use of broader terms such as comprises, includes, and having may be understood to provide support for narrower terms such as consisting of, consisting essentially of, and comprised substantially of. Accordingly, the scope of protection is not limited by the description set out above but is defined by the claims that follow. Each and every claim is incorporated as further disclosure into the specification and the claims are embodiment(s) of the present disclosure. The discussion of a reference in the disclosure is not an admission that it is prior art, especially any reference that has a publication date after the priority date of this application.

[0036] While several example embodiments have been provided in the present disclosure, it should be understood that the disclosed systems and methods might be embodied in many other specific forms. The present examples are to be considered as illustrative and not restrictive, and the intention is not to be limited to the details given herein. For example, the various elements or components may be combined or integrated in another system or certain features may be omitted, or not implemented.

[0037] In addition, techniques, systems, subsystems, and methods described and illustrated in the various example embodiments as discrete or separate may be combined or integrated with other systems, modules, techniques, or methods. Other items shown or discussed as coupled or directly coupled or communicating with each other may be indirectly coupled or communicating through some interface, device, or intermediate compo-

nent whether electrically, mechanically, or otherwise. Other examples of changes, substitutions, and alterations are ascertainable by one skilled in the art.

Claims

1. In a first edge node (110) coupled to a first site (140), a method of migrating control of at least one virtual private network, VPN, from a native border gateway protocol, BGP, peer session to a software-defined networking, SDN, central controller (210), wherein the native BGP peer session is ran by using a native BGP software residing on the first edge node, the method comprising:

   receiving a first software agent from the SDN central controller for installation on the first edge node;
   using the first software agent on the first edge node (110) to perform the steps of:

      taking over the native BGP peer session with a second edge node (120) coupled to a second site (150); receiving a first VPN route with a first VPN label from the SDN central controller (210) and sending the first VPN route with the first VPN label to the second edge node (120); and
      receiving a second VPN route with a second VPN label from the second edge node (120) coupled to the second site (150) and sending the second VPN route with the second VPN label to the SDN central controller (210).

2. The method of claim 1, further comprising:

   receiving a second software agent from the SDN central controller for installation on the first edge node;
   using the second software agent on the first edge node to perform the steps of:

      receiving the first VPN route from the first site and sending the first VPN route to the SDN central controller;
      receiving the second VPN route from the SDN central controller; and
      sending the second VPN route to the first site.

3. The method of claim 2, further comprising installing the first software agent and the second software agent on the first edge node, wherein the first software agent and the second software agent are received from the SDN central controller.

4. The method of claim 3, wherein after the second software agent on the first edge node and a fourth software agent on the second edge node are stable, the method further comprises:

   ending the BGP peer session with the second edge node; and
   deleting the first software agent.

5. The method of any one of claims 2 to 4, further comprising using the second software agent to perform the steps of:

   receiving an update to the first VPN route from the first site and sending the updated first VPN route to the SDN central controller;
   receiving an update to the second VPN route from the SDN central controller; and
   sending the updated second VPN route to the first site.

6. In a software-defined networking, SDN, central controller (210), a method for migrating and managing control of a virtual private network, VPN, from a native border gateway protocol, BGP, peer session to the SDN central controller (210), the VPN comprises a first edge node (110) for a first site (140) and a second edge node (120) for a second site (150), the method comprising:

   sending a first software agent to the first edge node (110) for installation on the first edge node (110); sending a second software agent to the first edge node (110) for installation on the first edge node (110); allocating a first VPN label for a first VPN route from the first site (140) via the second software agent;
   sending the first VPN route with the first VPN label to the first software agent;
   receiving a second VPN route with a second VPN label for the second edge node (120) from the first software agent; and
   sending the second VPN route for the second site (150) from the second edge node (120) to the second software agent;
   wherein the first software agent is used to perform the steps of:

      taking over the native BGP peer session with the second edge node (120) coupled to the second site (150); receiving the first VPN route with the first VPN label from the SDN central controller (210) and sending the first VPN route with the first VPN label to the second edge node (120); and
      receiving the second VPN route with the second VPN label from the second edge node (120) and sending the second VPN

route with the second VPN label to the SDN central controller (210).

7. The method of claim 6, further comprising:

adding a forwarding entry into a VPN routing and forwarding, VRF, table on the first edge node and the second edge node based on the first VPN route and the first VPN label; and
adding a forwarding entry into a VRF table on the first edge node and the second edge node based on the second VPN route and the second VPN label.

8. The method of claim 6 or 7, wherein the second software agent maintains a protocol adjacency with the first site.

9. The method of any one of claims 6 to 8, further comprising:

sending a third software agent to the second edge node for installation on the second edge node, wherein the third software agent takes over a native BGP peer session from native BGP software; and
sending a fourth software agent to the second edge node for installation on the second edge node, wherein the fourth software agent takes over a protocol adjacency with the second site.

10. The method of claim 9, wherein after the second software agent on the first edge node and the fourth software agent on the second edge node are stable, the method further comprises:

requesting that the first edge node remove the first software agent; and
requesting that the second edge node remove the third software agent.

11. The method of claim 10, wherein after the first software agent and the third software agent are removed, the method further comprises:

receiving a new VPN route for the first site via the second software agent;
allocating a VPN label for the new VPN route; and
adding a forwarding entry to a VPN routing and forwarding, VRF, table in the first edge node and the second edge node corresponding to the new VPN route and the VPN label;
or receiving a new VPN route for the second site via the fourth software agent;
allocating a VPN label for the new VPN route; and
adding a forwarding entry to a VPN routing and forwarding, VRF, table in the first edge node and the second edge node corresponding to the new VPN route and the VPN label.

12. A computer program product comprising computer executable instructions stored on a non-transitory computer readable medium such that when executed by a processor cause a central controller to perform the method according to any one of claims 6 to 11.

Patentansprüche

1. Verfahren zum Migrieren der Steuerung von mindestens einem virtuellen privaten Netzwerk (VPN) von einer nativen Border Gateway Protocol(BGP)-Peer-Sitzung zu einer Softwaredefiniertes-Netzwerk(SDN)-Zentralsteuerung (210) bei einem ersten Kantenknoten (110), der mit einer ersten Stelle (140) gekoppelt ist,
wobei die native BGP-Peer-Sitzung unter Verwendung einer nativen BGP-Software ausgeführt wird, die sich auf dem ersten Kantenknoten befindet, wobei das Verfahren umfasst:

Empfangen eines ersten Softwareagenten von der SDN-Zentralsteuerung zur Installation auf dem ersten Kantenknoten;
Verwenden des ersten Software-Agenten auf dem ersten Kantenknoten (110), um die folgenden Schritte durchzuführen:

Übernehmen der nativen BGP-Peer-Sitzung mit einem zweiten Kantenknoten (120), der mit einer zweiten Stelle (150) gekoppelt ist;
Empfangen einer ersten VPN-Route mit einem ersten VPN-Label von der SDN-Zentralsteuerung (210) und Senden der ersten VPN-Route mit dem ersten VPN-Label an den zweiten Kantenknoten (120); und
Empfangen einer zweiten VPN-Route mit einem zweiten VPN-Label von dem zweiten Kantenknoten (120), der mit der zweiten Stelle (150) gekoppelt ist, und Senden der zweiten VPN-Route mit dem zweiten VPN-Label an die SDN-Zentralsteuerung (210).

2. Verfahren nach Anspruch 1, ferner umfassend:

Empfangen eines zweiten Softwareagenten von der SDN-Zentralsteuerung zur Installation auf dem ersten Kantenknoten;
Verwenden des zweiten Softwareagenten auf dem ersten Kantenknoten, um die folgenden Schritte durchzuführen:

Empfangen der ersten VPN-Route von der ersten Stelle und Senden der ersten VPN-Route an die SDN-Zentralsteuerung;
Empfangen der zweiten VPN-Route von der SDN-Zentralsteuerung; und
Senden der zweiten VPN-Route an die erste Stelle.

3. Verfahren nach Anspruch 2, ferner umfassend das Installieren des ersten Softwareagenten und des zweiten Softwareagenten auf dem ersten Kantenknoten, wobei der erste Softwareagent und der zweite Softwareagent von der SDN-Zentralsteuerung empfangen werden.

4. Verfahren nach Anspruch 3, wobei, nachdem der zweite Softwareagent auf dem ersten Kantenknoten und ein vierter Softwareagent auf dem zweiten Kantenknoten stabil sind, das Verfahren ferner umfasst:

Beenden der BGP-Peer-Sitzung mit dem zweiten Kantenknoten; und
Löschen des ersten Softwareagenten.

5. Verfahren nach einem der Ansprüche 2 bis 4, ferner umfassend das Verwenden des zweiten Softwareagenten, um die folgenden Schritte durchzuführen:

Empfangen einer Aktualisierung der ersten VPN-Route von der ersten Stelle und Senden der aktualisierten ersten VPN-Route an die SDN-Zentralsteuerung;
Empfangen einer Aktualisierung der zweiten VPN-Route von der SDN-Zentralsteuerung; und
Senden der aktualisierten zweiten VPN-Route an die erste Stelle.

6. Verfahren zum Migrieren und Verwalten der Steuerung eines virtuellen privaten Netzwerks (VPN) von einer nativen Border Gateway Protocol (BGP)-Peer-Sitzung zu der SDN-Zentralsteuerung (210) in einer Softwaredefiniertes-Netzwerk(SDN)-Zentralsteuerung (210), wobei das VPN einen ersten Kantenknoten (110) für eine erste Stelle (140) und einen zweiten Kantenknoten (120) für eine zweite Stelle (150) umfasst, wobei das Verfahren umfasst:

Senden eines ersten Softwareagenten an den ersten Kantenknoten (110) zur Installation auf dem ersten Kantenknoten (110);
Senden eines zweiten Softwareagenten an den ersten Kantenknoten (110) zur Installation auf dem ersten Kantenknoten (110);
Zuweisen eines ersten VPN-Labels für eine erste VPN-Route von der ersten Stelle (140) über den zweiten Softwareagenten;
Senden der ersten VPN-Route mit dem ersten VPN-Label an den ersten Softwareagenten;
Empfangen einer zweiten VPN-Route mit einem zweiten VPN-Label für den zweiten Kantenknoten (120) von dem ersten Softwareagenten; und
Senden der zweiten VPN-Route für die zweite Stelle (150) von dem zweiten Kantenknoten (120) an den zweiten Softwareagenten;
wobei der erste Softwareagent verwendet wird, um die folgenden Schritte durchzuführen:

Übernehmen der nativen BGP-Peer-Sitzung mit dem zweiten Kantenknoten (120), der mit der zweiten Stelle (150) gekoppelt ist;
Empfangen der ersten VPN-Route mit dem ersten VPN-Label von der SDN-Zentralsteuerung (210) und Senden der ersten VPN-Route mit dem ersten VPN-Label an den zweiten Kantenknoten (120); und
Empfangen der zweiten VPN-Route mit dem zweiten VPN-Label von dem zweiten Kantenknoten (120) und Senden der zweiten VPN-Route mit dem zweiten VPN-Label an die SDN-Zentralsteuerung (210).

7. Verfahren nach Anspruch 6, ferner umfassend:

Hinzufügen eines Weiterleitungseintrags in eine VPN-Routing- und -Weiterleitungs(VPN Routing and Forwarding, VRF)-Tabelle auf dem ersten Kantenknoten und dem zweiten Kantenknoten basierend auf der ersten VPN-Route und dem ersten VPN-Label; und
Hinzufügen eines Weiterleitungseintrags in eine VRF-Tabelle auf dem ersten Kantenknoten und dem zweiten Kantenknoten basierend auf der zweiten VPN-Route und dem zweiten VPN-Label.

8. Verfahren nach Anspruch 6 oder 7, wobei der zweite Softwareagent eine Protokollnähe zur ersten Stelle aufrechterhält.

9. Verfahren nach einem der Ansprüche 6 bis 8, ferner umfassend:

Senden eines dritten Softwareagenten an den zweiten Kantenknoten zur Installation auf dem zweiten Kantenknoten, wobei der dritte Softwareagent eine native BGP-Peer-Sitzung von der nativen BGP-Software übernimmt; und
Senden eines vierten Softwareagenten an den zweiten Kantenknoten zur Installation auf dem zweiten Kantenknoten, wobei der vierte Softwareagent eine Protokollnähe zur zweiten Stelle übernimmt.

10. Verfahren nach Anspruch 9, wobei, nachdem der zweite Softwareagent auf dem ersten Kantenknoten

und der vierte Softwareagent auf dem zweiten Kantenknoten stabil sind, das Verfahren ferner umfasst:

Anfordern, dass der erste Kantenknoten den ersten Softwareagenten entfernt; und
Anfordern, dass der zweite Kantenknoten den dritten Softwareagenten entfernt.

11. Verfahren nach Anspruch 10, wobei, nachdem der erste Softwareagent und der dritte Softwareagent entfernt worden sind, das Verfahren ferner umfasst:

Empfangen einer neuen VPN-Route für die erste Stelle über den zweiten Softwareagenten;
Zuweisen eines VPN-Labels für die neue VPN-Route; und
Hinzufügen eines Weiterleitungseintrags zu einer VPN-Routing- und -Weiterleitungs(VPN Routing and Forwarding, VRF)-Tabelle in dem ersten Kantenknoten und dem zweiten Kantenknoten entsprechend der neuen VPN-Route und dem VPN-Label; oder Empfangen einer neuen VPN-Route für die zweite Stelle über den vierten Softwareagenten;
Zuweisen eines VPN-Labels für die neue VPN-Route; und
Hinzufügen eines Weiterleitungseintrags zu einer VPN-Routing- und -Weiterleitungs(VPN Routing and Forwarding, VRF)-Tabelle in dem ersten Kantenknoten und dem zweiten Kantenknoten entsprechend der neuen VPN-Route und dem VPN-Label.

12. Computerprogrammprodukt, umfassend computerausführbare Anweisungen, die in einem nichtflüchtigen computerlesbaren Medium gespeichert sind, sodass sie bei Ausführung durch einen Prozessor eine Zentralsteuerung veranlassen, das Verfahren nach einem der Ansprüche 6 bis 11 durchzuführen.

Revendications

1. Dans un premier noeud périphérique (110) couplé à un premier site (140), procédé de commande de migration d'au moins un réseau privé virtuel (VPN) depuis une session entre homologues de protocole de passerelle frontière (BGP) native vers un dispositif de commande central de réseautage défini par un logiciel (SDN) (210), dans lequel la session entre homologues de protocole BGP native est mise en fonctionnement en utilisant un logiciel de protocole BGP native résidant sur le premier noeud périphérique, le procédé consistant :

à recevoir un premier agent logiciel en provenance du dispositif de commande central de réseautage SDN pour une installation sur le premier noeud périphérique ;
à utiliser le premier agent logiciel sur le premier noeud périphérique (110) pour effectuer les étapes consistant :

à prendre le pas sur la session entre homologues de protocole BGP native avec un second noeud périphérique (120) couplé à un second site (150) ;
à recevoir un premier trajet de réseau VPN ayant une première étiquette de réseau VPN en provenance du dispositif de commande central de réseautage SDN (210) et à envoyer le premier trajet de réseau VPN ayant la première étiquette de réseau VPN au second noeud périphérique (120) ; et
à recevoir un second trajet de réseau VPN ayant une seconde étiquette de réseau VPN en provenance du second noeud périphérique (120) couplé au second site (150) et à envoyer le second trajet de réseau VPN ayant la seconde étiquette de réseau VPN au dispositif de commande central de réseautage SDN (210).

2. Procédé selon la revendication 1, consistant en outre :

à recevoir un deuxième agent logiciel en provenance du dispositif de commande central de réseautage SDN pour une installation sur le premier noeud périphérique ;
à utiliser le deuxième agent logiciel sur le premier noeud périphérique pour effectuer les étapes consistant :

à recevoir le premier trajet de réseau VPN en provenance du premier site et à envoyer le premier trajet de réseau VPN au dispositif de commande central de réseautage SDN ;
à recevoir le second trajet de réseau VPN en provenance du dispositif de commande central de réseautage SDN ; et
à envoyer le second trajet de réseau VPN au premier site.

3. Procédé selon la revendication 2, consistant en outre à installer le premier agent logiciel et le deuxième agent logiciel sur le premier noeud périphérique, dans lequel le premier agent logiciel et le deuxième agent logiciel sont reçus du dispositif de commande central de réseautage SDN.

4. Procédé selon la revendication 3, dans lequel, après que le deuxième agent logiciel sur le premier noeud périphérique et un quatrième agent logiciel sur le second noeud périphérique sont stables, le procédé consiste en outre :

à terminer la session entre homologues de protocole BGP avec le second noeud périphérique ; et
à supprimer le premier agent logiciel.

5. Procédé selon l'une quelconque des revendications 2 à 4, consistant en outre à utiliser le deuxième agent logiciel pour effectuer les étapes consistant :

à recevoir une mise à jour pour le premier trajet de réseau VPN en provenance du premier site et à envoyer le premier trajet de réseau VPN mis à jour au dispositif de commande central de réseautage SDN ;
à recevoir une mise à jour pour le second trajet de réseau VPN en provenance du dispositif de commande central de réseautage SDN ; et
à envoyer le second trajet de réseau VPN mis à jour au premier site.

6. Dans un dispositif de commande central de réseautage défini par un logiciel (SDN) (210), procédé de commande de migration et de gestion d'un réseau privé virtuel (VPN) depuis une session entre homologues de protocole de passerelle frontière (BGP) native vers le dispositif de commande central de réseautage SDN (210), le réseau VPN comprend un premier noeud périphérique (110) pour un premier site (140) et un deuxième noeud périphérique (120) pour un second site (150), le procédé consistant :

à envoyer un premier agent logiciel au premier noeud périphérique (110) pour une installation sur le premier noeud périphérique (110) ;
à envoyer un deuxième agent logiciel au premier noeud périphérique (110) pour une installation sur le premier noeud périphérique (110) ;
à attribuer une première étiquette de réseau VPN pour un premier trajet de réseau VPN depuis le premier site (140) par le biais du deuxième agent logiciel ;
à envoyer le premier trajet de réseau VPN ayant la première étiquette de réseau VPN au premier agent logiciel ;
à recevoir un second trajet de réseau VPN ayant une seconde étiquette de réseau VPN pour le deuxième noeud périphérique (120) en provenance du premier agent logiciel ; et
à envoyer le second trajet de réseau VPN pour le second site (150) depuis le second noeud périphérique (120) au deuxième agent logiciel ;
dans lequel le premier agent logiciel est utilisé pour effectuer les étapes consistant :

à prendre le pas sur la session entre homologues de protocole BGP native avec un second noeud périphérique (120) couplé au second site (150) ;
à recevoir le premier trajet de réseau VPN ayant la première étiquette de réseau VPN en provenance du dispositif de commande central de réseautage SDN (210) et à envoyer le premier trajet de réseau VPN ayant la première étiquette de réseau VPN au second noeud périphérique (120) ; et
à recevoir le second trajet de réseau VPN ayant la seconde étiquette de réseau VPN en provenance du second noeud périphérique (120) et à envoyer le second trajet de réseau VPN ayant la seconde étiquette de réseau VPN au dispositif de commande central de réseautage SDN (210).

7. Procédé selon la revendication 6, consistant en outre :

à ajouter une entrée de transfert dans une table de routage et d'acheminement de réseau VPN (VRF) sur le premier noeud périphérique et le second noeud périphérique en se basant sur le premier trajet de réseau VPN et la première étiquette de réseau VPN ; et
à ajouter une entrée de transfert dans une table VRF sur le premier noeud périphérique et le second noeud périphérique en se basant sur le second trajet de réseau VPN et la seconde étiquette de réseau VPN.

8. Procédé selon la revendication 6 ou 7, dans lequel le deuxième agent logiciel maintient une proximité de protocole avec le premier site.

9. Procédé selon l'une quelconque des revendications 6 à 8, consistant en outre :

à envoyer un troisième agent logiciel au second noeud périphérique pour une installation sur le second noeud périphérique, dans lequel le troisième agent logiciel prend le pas sur une session entre homologues de protocole BGP native à partir d'un logiciel de protocole BGP native ; et
à envoyer un quatrième agent logiciel au second noeud périphérique pour une installation sur le second noeud périphérique, dans lequel le quatrième agent logiciel prend le pas sur une proximité de protocole avec le second site.

10. Procédé selon la revendication 9, dans lequel, après que le deuxième agent logiciel sur le premier noeud périphérique et le quatrième agent logiciel sur le second noeud périphérique sont stables, le procédé consiste en outre :

à demander que le premier noeud périphérique supprime le premier agent logiciel ; et
à demander que le second noeud périphérique

supprime le troisième agent logiciel.

11. Procédé selon la revendication 10, dans lequel, après que le premier agent logiciel et le troisième agent logiciel sont supprimés, le procédé consiste en outre :

à recevoir un nouveau trajet de réseau VPN pour le premier site par le biais du deuxième agent logiciel ;
à attribuer une étiquette de réseau VPN pour le nouveau trajet de réseau VPN ; et
à ajouter une entrée de transfert à une table de routage et d'acheminement de réseau VPN (VRF) dans le premier noeud périphérique et le second noeud périphérique correspondant au nouveau trajet de réseau VPN et à l'étiquette de réseau VPN ;
ou à recevoir un nouveau trajet de réseau VPN pour le second site par le biais du quatrième agent logiciel ;
à attribuer une étiquette de réseau VPN pour le nouveau trajet de réseau VPN ; et
à ajouter une entrée de transfert à une table de routage et d'acheminement de réseau VPN (VRF) dans le premier noeud périphérique et le second noeud périphérique correspondant au nouveau trajet de réseau VPN et à l'étiquette de réseau VPN.

12. Produit-programme d'ordinateur comprenant des instructions exécutables par un ordinateur stockées sur un support non transitoire lisible par un ordinateur de telle sorte que, lorsqu'elles sont exécutées par un processeur, elles contraignent un dispositif de commande central à réaliser le procédé selon l'une quelconque des revendications 6 à 11.

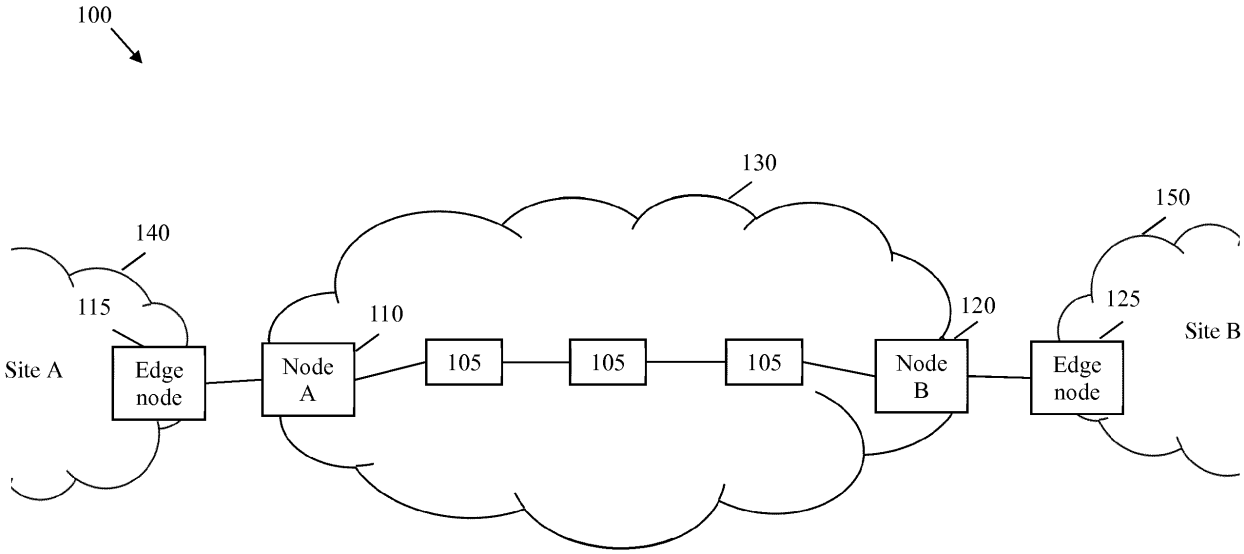
100
130
150
140
115
110
120
125
Site A
Edge node
Node A
105
105
105
Node B
Edge node
Site B

FIG. 1

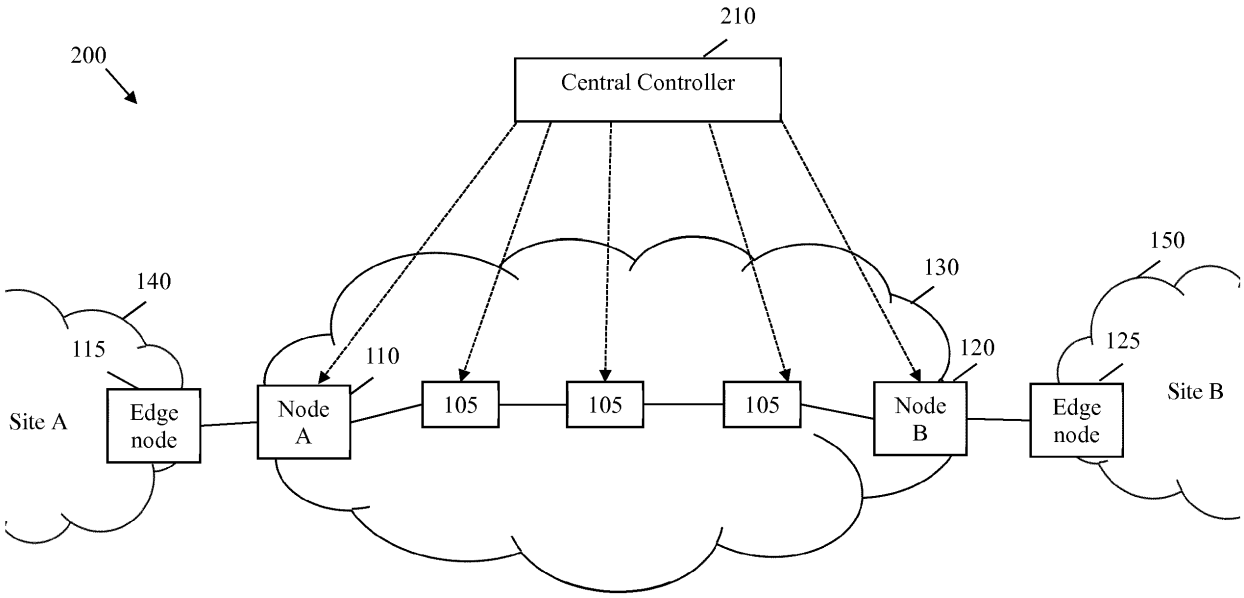
200
210
Central Controller
140
115
Site A
Edge node
Node A
110
105
105
105
130
Node B
120
150
125
Edge node
Site B

FIG. 2

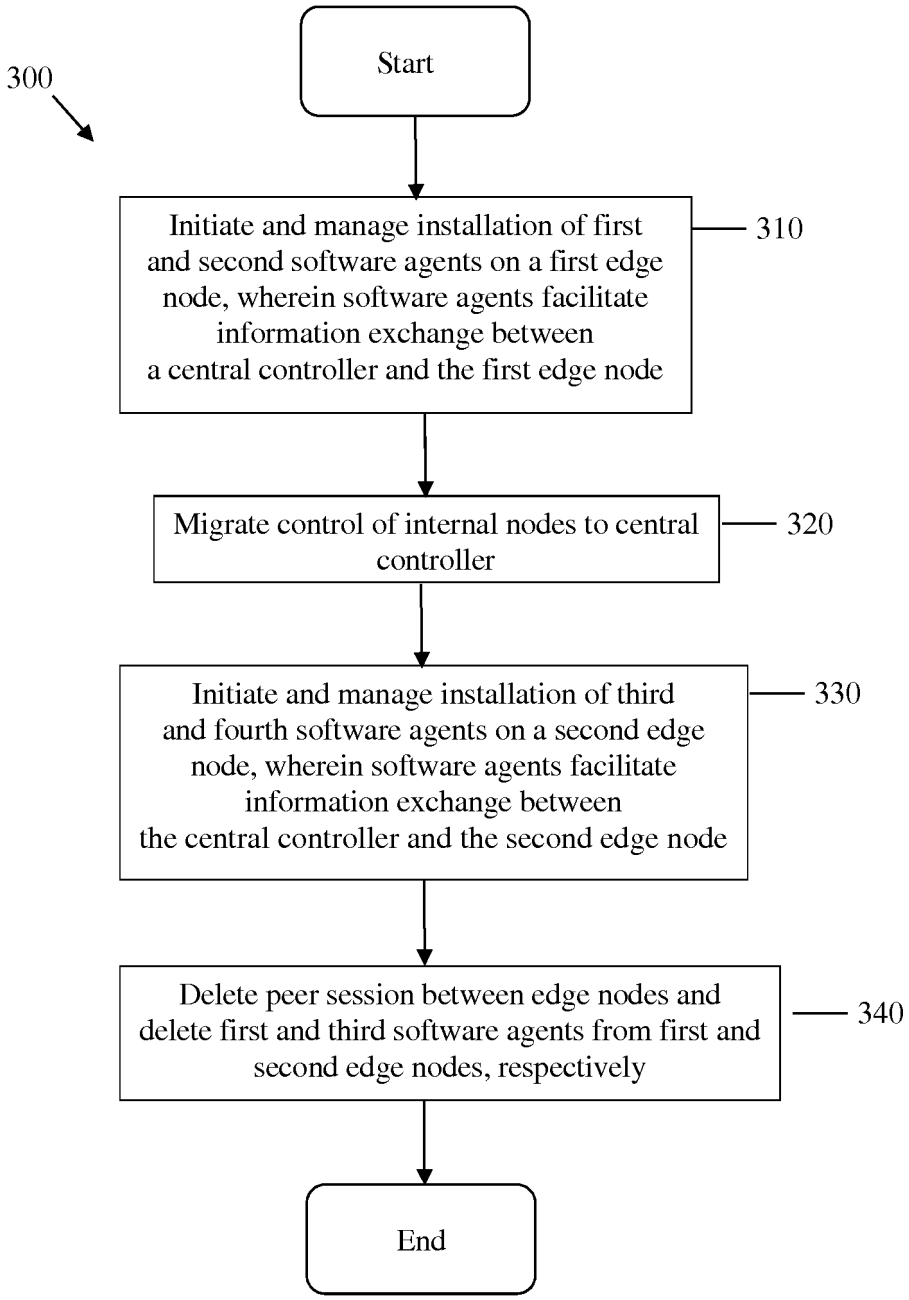
300
Start
Initiate and manage installation of first and second software agents on a first edge node, wherein software agents facilitate information exchange between a central controller and the first edge node
310
Migrate control of internal nodes to central controller
320
Initiate and manage installation of third and fourth software agents on a second edge node, wherein software agents facilitate information exchange between the central controller and the second edge node
330
Delete peer session between edge nodes and delete first and third software agents from first and second edge nodes, respectively
340
End

FIG. 3

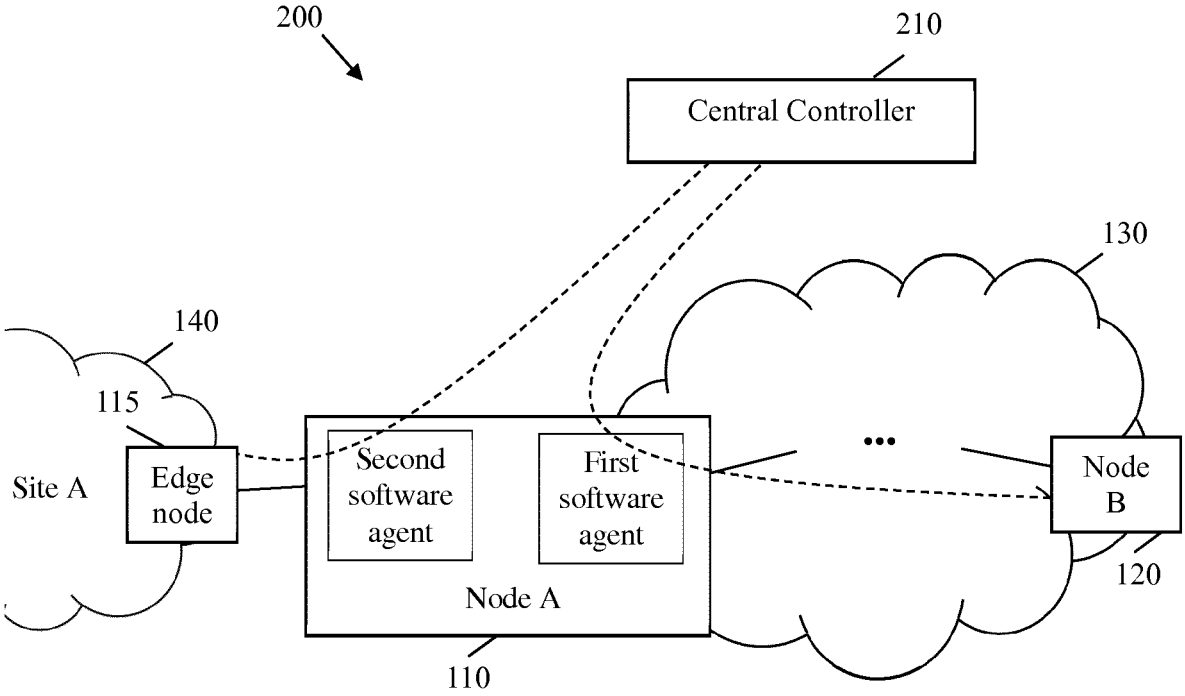
200
210
Central Controller
130
140
115
Site A
Edge node
Second software agent
First software agent
Node A
110
Node B
120

FIG. 4

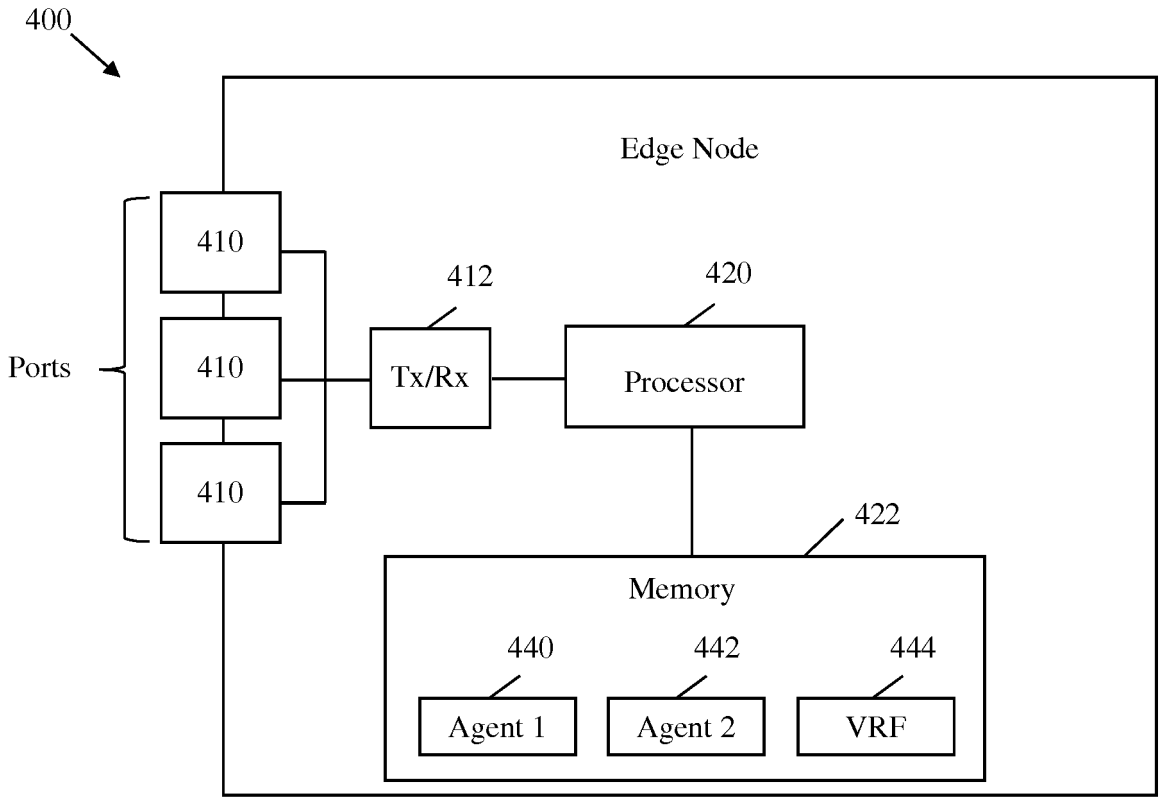
400
Edge Node
410
410
410
Ports
412
Tx/Rx
420
Processor
422
Memory
440
Agent 1
442
Agent 2
444
VRF

FIG. 5

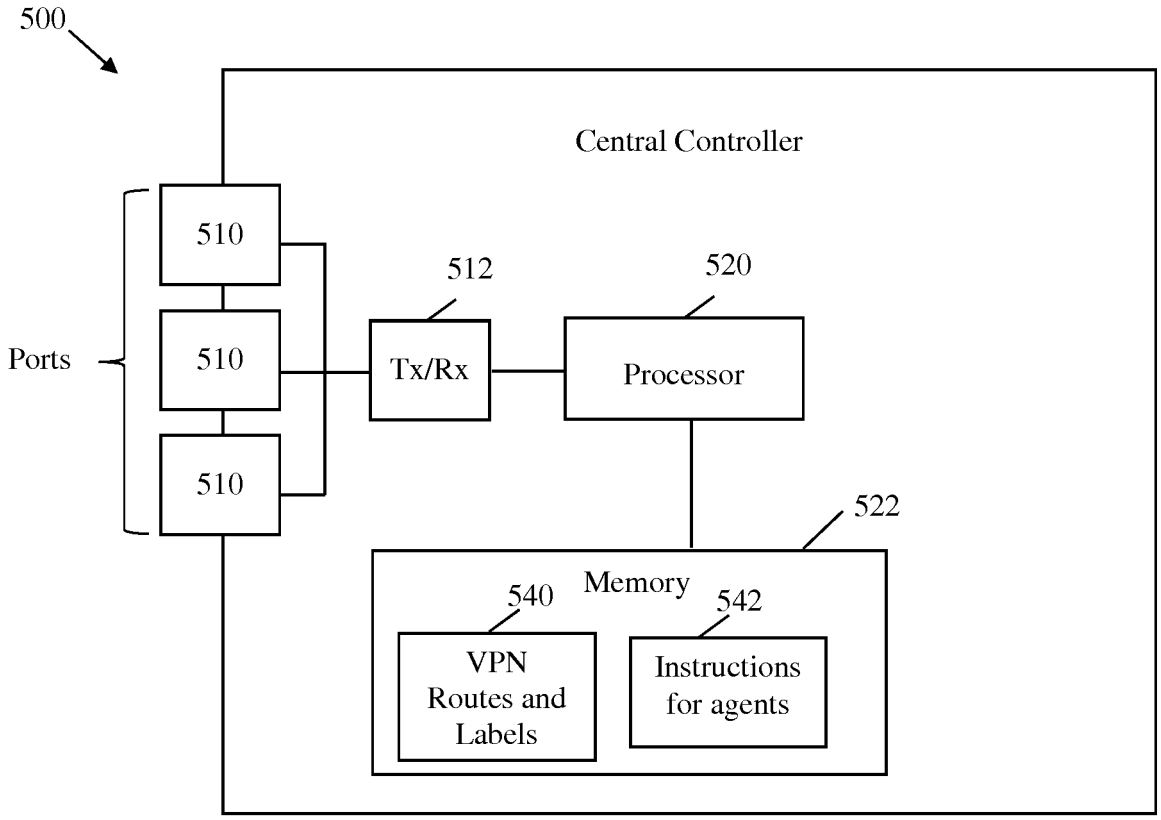
500
Central Controller
510
510
510
Ports
512
Tx/Rx
520
Processor
522
Memory
540
VPN
Routes and
Labels
542
Instructions
for agents

FIG. 6

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 20110142053 A1 **[0004]**

Non-patent literature cited in the description

- BGP/MPLS IP VPN Virtual PE; draft-fang-13vpn-virtual-pe-04.txt. **LUYUAN FANG ; DAVID WARD ; REX FERNANDO ; CISCO MARIA ; NAPIERALA AT ; T NABIL BITAR VERIZON.** BGP/MPLS IP VPN VIRTUAL PE; DRAFT-FANG-L3VPN-VIRTUAL-PE-04.TXT. INTERNET ENGINEERING TASK FORCE, IETF; STANDARDWORKINGDRAFT, INTERNET SOCIETY (ISOC), 21 October 2013, 1-24 **[0005]**