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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
3 September 2009 (03.09.2009)

PCT

(10) International Publication Number
WO 2009/108682 A1

(51) International Patent Classification:
H04L 12/56 (2006.01) *H04L 12/28* (2006.01)

(21) International Application Number:
PCT/US2009/035112

(22) International Filing Date:
25 February 2009 (25.02.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/037,525 26 February 2008 (26.02.2008) US

(71) Applicant *(for all designated States except US)*: **CISCO TECHNOLOGY, INC.** [US/US]; 170 West Tasman Drive, San Jose, California 95134 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **RAJENDRAN, Saravanakumar** [US/US]; 3308 Provence Court, San Jose, California 95135 (US). **SMITH, Michael** [CA/US]; 401-11 Camille Circle, San Jose, California 95134 (US). **DEVIREDDY, Dileep, Kumar** [IN/US]; 4241 Voltaire Street, San Jose, California 95135 (US). **KATHAIL, Pradeep K.** [US/US]; 1201 Carmel Terrace, Los Altos, California 94024 (US). **APPANNA, Chandrashekhar** [US/US]; 10225 Colby Avenue, Cupertino, California 95014 (US). **WANG, Jeffrey, Ym** [US/US]; 14230 Lutheria Way, Saratoga, California 95070 (US). **GANDHI, Prashant, P.** [US/US]; 3481 Valley Vista Drive, San Jose, California 95148 (US).

(74) Agent: **READ, Randol, W.**; Patterson & Sheridan, L.L.P., 3040 Post Oak Blvd., Suite 1500, Houston, Texas 77056-6582 (US).

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

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR),

[Continued on next page]

(54) Title: BLADE SWITCH

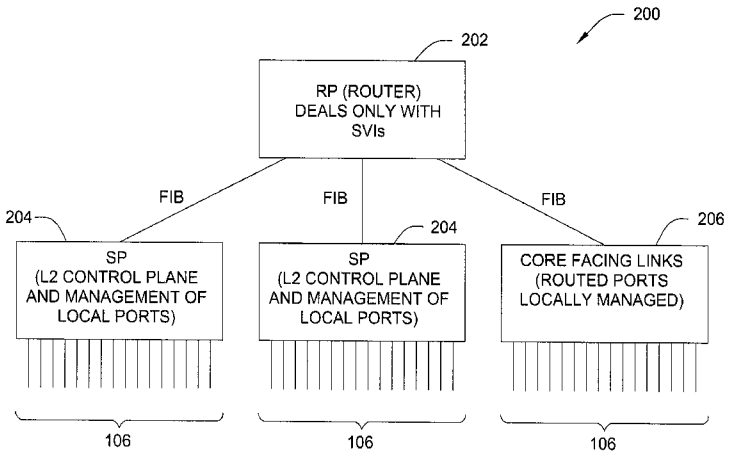


FIG. 2

(57) Abstract: A blade switch for increased interface scalability is provided. The blade switch may address interface scalability by having each of the switch linecards manage its interfaces locally and may use the concept of virtual and local interfaces for intelligent forwarding. The blade switch may appear as a single network switch having a single bridge ID from the network perspective during operation and from the customer perspective during configuration.

WO 2009/108682 A1

OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

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

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

BLADE SWITCH

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] Embodiments of the present invention generally relate to the field of network communication, and more particularly, to interface scalability when increasing the number of switch linecards within a single chassis.

Description of the Related Art

[0002] As the number of clients or network nodes having destination IP addresses increases for an Internet Service Provider (ISP) or a growing large company, network complexity and, thus, the demand for more router and switch interfaces also scale accordingly. Switching systems containing several slots for connecting switch linecards have been designed in an effort to meet the demand for network switches with a large number of interfaces, both physical and logical. However, scalability in terms of the number of interfaces and associated operations supported by a single network switch remains a problem.

[0003] Accordingly, techniques for successfully scaling the number of interfaces within a network switching system are needed.

OVERVIEW

[0004] Embodiments of the present invention generally relate to network switching systems for addressing interface scalability.

[0005] One embodiment of the present invention provides a network switch system. The system generally includes a chassis; a plurality of switch linecards disposed in the chassis and functioning as independent Layer 2 (L2) switches, wherein each of the switch linecards maintains control, forwarding, and interface management functions locally; and a controller, wherein the controller is adapted to configure the plurality of switch linecards as a single network device having a single bridge identifier (bridge ID).

[0006] Another embodiment of the present invention provides a method. The method generally includes receiving a packet in a first switch linecard of a plurality of switch linecards disposed in a chassis, wherein the plurality of switch linecards is configured as a single network switch having a single bridge ID, and forwarding the packet from the first switch linecard to a second switch linecard via a virtual interface.

[0007] Yet another embodiment of the present invention provides a method. The method generally includes configuring a forwarding information base (FIB) of a first switch linecard of a plurality of switch linecards disposed in a chassis to forward a packet to a local interface of the first switch linecard if the packet is to be forwarded within the first switch linecard and configuring the FIB of the first switch linecard to forward the packet from the first switch linecard to a virtual interface for bridging between two or more of the plurality of switch linecards if the packet is to be forwarded to a second switch linecard in the plurality.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] So that the manner in which the above recited features of the present invention can be understood in detail, a more particular description of the invention, briefly summarized above, may be had by reference to embodiments, some of which are illustrated in the appended drawings. It is to be noted, however, that the appended drawings illustrate only typical embodiments of this invention and are therefore not to be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

[0009] FIG. 1 is a block diagram of a switching network-in-a-box in accordance with an embodiment of the present invention;

[0010] FIG. 2 models the control plane for a blade switch in accordance with an embodiment of the present invention; and

[0011] FIG. 3 illustrates blade-to-blade unicast forwarding within the blade switch in accordance with an embodiment of the present invention.

DESCRIPTION OF EXAMPLE EMBODIMENTS

[0012] Embodiments of the present invention provide a blade switch to address interface scalability and permit an increased number of switch linecards supported within a single chassis.

An Example Blade Switch

[0013] FIG. 1 is a block diagram of an example blade switch 100 according to an embodiment of the invention. The blade switch 100 may comprise a chassis 102 having a backplane and a plurality of slots. Switch linecards 104 (also known as switch blades and labeled S1-S4 in FIG. 1) may be disposed in the slots of the chassis 102 and electrically coupled to the backplane, typically through plug-in interfaces. The number of slots may dictate the number of switch linecards 104 the chassis 102 may hold. Although only four switch linecards 104 are shown in FIG. 1, this is representative of any number of linecards that may be disposed within a blade switch chassis (*e.g.*, 44 or 128 linecards). Moreover, an operable blade switch may have empty slots where the chassis 102 has more available slots than switch linecards 104 used for a given application.

[0014] Each switch linecard 104 may have a plurality of external interfaces 106 for coupling the switch linecard 104 to an external network, such as a router, an external bridge, a server, or a client. These external interfaces 106 may be for wireless transmission or physical electrical connection to any suitable cable, such as Cat-5 cable, for high-speed data transmission. The physical external interfaces 106 and any logical interfaces for a given switch linecard 104 may be considered as the local interfaces for that switch linecard 104. Each switch linecard 104 may also be coupled to an internal local area network (LAN) segment 108 for routing between the linecards 104. Physically, the internal LAN segment 108 may be located within the backplane (*e.g.*, as a trace within the printed circuit board (PCB) of the backplane), and the individual switch linecards 104 may be coupled to the internal LAN segment 108 by being connected to the backplane.

[0015] Furthermore, a controller 110 may be coupled to the switch linecards 104 in an effort to manage and configure the routing information for the individual switch linecards 104. To manage the switch linecards 104, the controller 110 may run an operating system. The controller 110 may be disposed within a separate area of the chassis 102 from the switch linecards 104 or have a designated slot. For some embodiments, the controller 110 may be coupled to the linecards 104 via the internal LAN segment 108 as shown.

[0016] The chassis 102 may be composed of any suitable material, such as metal, for providing a rigid structure in which to contain and protect elements disposed within (*e.g.*, the switch linecards 104) and for acting as a Faraday shield against electrical noise. For some embodiments, the chassis 102 may also contain one or more fans (not shown) in an effort to provide air circulation and, thus, forced convective cooling of the switch linecards 104. Furthermore, a power supply (not shown) may be disposed within the chassis 102 and may serve to convert AC power, such as 120 V_{rms} from a wall outlet, to one or more DC voltages for use by the switch linecards 104, the controller 110, and the fan(s). The power supply may be coupled to the backplane for power distribution to the switch linecards 104 and the controller 110.

[0017] In the blade switch 100, each switch linecard 104 may own the interfaces locally. In other words, each switch linecard 104 may store and maintain its own forwarding information base (FIB), also known as a forwarding table, for making switching/routing decisions. Hence, the controller 110 need not control the data packets during operation of the blade switch 100. This is in contrast to conventional systems where the controller typically distributes the FIB to the blades and controls the data switching. Used to find the proper interface to which an input interface should send a packet to be transmitted by the switch linecard, the FIB may be directed toward fast lookup of destination addresses. From a data plane perspective, the blade switch 100 may behave substantially similar to a single network switch in a network, even though there may tens to hundreds of switch linecards 104 contained within the chassis 102, thereby leading to the concept of a

"network-in-a-box" so to speak. By conceptually having all of the switch linecards in the chassis 102 function as a single network switch, the system software may be scaled to support a large number of interfaces.

[0018] To configure the blade switch 100, the customer may interface with the controller 110, which handles the global aspects of the blade switch 100, via software, such as the controller's operating system. From the customer, control plane, and data plane perspectives, the blade switch 100 may appear and behave as a single network switch having a single bridge identifier (bridge ID). However, once the customer has completed configuration of the blade switch 100, the FIB for each switch linecard 104 may be updated and stored locally. For some embodiments, the controller 110 may act as a router and download the FIB contents into the switches. The blade switch 100 may configure only a single virtual interface, such as the internal LAN segment 108, for the network fabric within the chassis 102.

[0019] With the virtual interface configured by the controller 110, the blade switch 100 may employ intelligent forwarding through the fabric by dividing the routes into two groups: local interfaces and virtual interfaces. Local interfaces, as used herein, may be generally defined as interfaces (both physical and logical) that are local to a given switch linecard 104. A virtual interface, as used herein, may be generally defined as an interface for the blade switch model that identifies the shared segment connecting the linecards. With these two groups, the intelligent forwarding may be known as two-stage, or split, forwarding. Based on the destination, the FIB for each switch linecard 104 may indicate routing incoming data packets to a local interface, local to that particular linecard, or to the virtual interface for routing to another switch linecard 104 within the chassis 102 of the blade switch 100.

[0020] In other words, the virtual interface may be a global internal virtual LAN (VLAN) used for bridging across the linecards 104. These local and virtual interfaces may operate on Layer 2 (L2, the data link layer) of the Open Systems Interconnection Reference Model (OSI model). Conventionally on a shared LAN segment, the data packets are received by every node on the segment. With intelligent forwarding, however, packet forwarding is performed as if the linecards

are connected with point-to-point links, which will be described in further detail below. Furthermore, mechanisms may be implemented in an effort to ensure that the internal routes of the virtual interface are preferred over any possible routes existing outside of the blade switch's chassis 102.

[0021] With the concept of intelligent forwarding, local switching within a certain switch linecard 104 in the blade switch 100 may most likely have the same performance as in a conventional switching system with the same switch linecard. However, because of the two lookups in intelligent forwarding when packets are routed between linecards in the blade switch 100, the performance may be slightly reduced when compared to a conventional switching system with the same switch linecards. However, the scalability advantage may outweigh the slight performance reduction for certain applications.

[0022] By localizing the switching decisions for a linecard in the blade switch 100, each of the switch linecards 104 may run a separate instance of an operating system, such as Cisco's Internetwork Operating System (IOS). Moreover, in the blade switch 100, each of the switch linecards 104 within the chassis 102 may run a different version of the operating system. This may allow for scaling the number of switch linecards 104 in the blade switch 100 in conjunction with operating system feature/code development. Furthermore, because forwarding is localized on each linecard, each of the switch linecards 104 may support a different version of the forwarding engine, such as Cisco's Enhanced Address Recognition Logic (EARL), within the routing application specific integrated circuit (ASIC). In other words, mixed mode EARLs may be supported.

[0023] Allowing for different operating systems and different forwarding engines may permit a "pay-as-you-go" model, where customers may upscale their blade switch by adding new switch linecards potentially with upgraded features without discarding legacy linecards. In this manner, scalability is not hindered by the costs of having to replace legacy linecards when upgrading a network switching system, such as the blade switch 100.

[0024] Referring now to the control plane model 200 of FIG. 2, ports (including the external interfaces 106) that are aggregated may be L2 ports. Each port may be assigned a private VLAN, and multiple ports may be grouped into a routable VLAN with a switch virtual interface (SVI) defined. An SVI represents a VLAN of switch ports as a single interface to a routing or bridging function of a network switch, such as the blade switch 100.

[0025] Routing protocols need not be run on the aggregated ports. Rather, routing protocols may be run in a single place, such as in the route processor (RP) 202, or router. The route processor 202 may be part of the controller 110 of the blade switch 100, or the controller 110 may function as the route processor 202. Since there is one single router control plane and one single router view to the external world, EtherChannels may be supported in the blade switch 100. Each switch linecard 104 may contain a switch processor (SP) 204 for L2 control plane functionality, as well as management of ports local to that particular linecard. In other words, each switch linecard 104 may act like an L2 switch. Each blade may expose a certain number of SVIs to the route processor 202.

[0026] The blade switch 100 may also be intended to interface with a number of core-facing uplink ports, and the number of uplink ports (*i.e.*, real routed ports 206) may most likely be much less than the total number of aggregated ports. The routed ports 206, which function on Layer 3 (L3, the network layer) of the OSI model, may most likely reside on a separate switch linecard 104 than the L2 switches. The routed ports 206 may also be locally managed. The physical external interfaces 106 may be scaled by distributing forwarding information to the switch linecards 104, as illustrated by the FIBs (*e.g.*, as part of Distributed Cisco Express Forwarding, or dCEF) being distributed to the switch processors 204 and routed ports 206 in FIG. 2.

[0027] From the management plane perspective, the blade switch 100 may be centrally managed at the controller 110, and more specifically, at the route processor 202. The controller 110 may have a single internet protocol (IP) address with session support, and commands and/or features for blade switch management may be executed on the controller 110. Furthermore, some embodiments may

include an option to telnet directly to a certain linecard 104 for debugging and network management. In other words, the blade switch 100 may allow a user to connect to the controller 110 and then select an individual switch linecard 104 for management and/or debugging.

[0028] For some embodiments, the blade switch 100 may allow a user to connect to the controller 110 and manage either all or a subset of the plurality of switch linecards 104 as a single entity. In the management plane, the blade switch 100 may allow the switch linecards 104 to be partitioned based on the operating system version the linecards are running for some embodiments. By allowing linecards running the same version and image of the operating system to be managed as a single entity, a command translation table need not be maintained. Also for some embodiments, the blade switch 100 may have provisions for specifying a subset of the switch linecards 104 as the destination for specified commands.

[0029] An Interface Descriptor Block (IDB) is a special control structure internal to the Cisco IOS software that contains information such as the IP address, interface state, and packet statistics. Cisco IOS software may maintain one IDB for each interface present on a platform. While the L2 IDBs may be local to the switch linecards 104, the SVI IDBs may reside in the route processor 202. L3 commands may be processed locally at the route processor 202, but L2 commands may be sent to one or more switch linecard(s) 104.

[0030] In view of the above description, the blade switch 100 may have each of the switch linecards 104 maintain its control/forwarding/management functions locally while appearing as a single switch with a single bridge ID from the customer and network perspectives.

Unicast Forwarding Example

[0031] The blade switch 100 may support several different types of linecard-to-linecard (blade-to-blade) forwarding from the data plane perspective. FIG. 3 illustrates blade-to-blade unicast forwarding within the blade switch according to some embodiments of the present invention. In computer networks, unicast

forwarding is the sending of data packets to a single destination, and duplicate data streams may be transmitted if more than one destination is to receive a particular data packet.

[0032] As depicted in FIG. 3, the blade switch 100 may contain a first switch linecard 300 labeled BLADE 1 and a second switch linecard 302 labeled BLADE 2. The first switch linecard 300 may have two local ports 304 labeled I_0 and I_1, and the second switch linecard 202 may have two local ports 304 labeled I_6 and I_7. Although FIG. 3 depicts only two local ports per switch linecard, each linecard may have many more local ports for interfacing with a network external to the blade switch 100. Interface I_1 may be coupled to a network 306 having an IP subnet of 2.x.x.x, for example. Interface I_6 may be coupled to a network 308 having an IP subnet of 1.x.x.x. The first and second switch linecards 300, 302 may be coupled by the virtual interface (V) 310 for blade-to-blade forwarding.

[0033] The forwarding engine 312, such as EARL, within the first switch linecard 300 may contain a forwarding information base (FIB) 314 after configuration of the blade switch 100 that indicates routing incoming data packets to a local interface, such as I_0 and I_1, to a switch virtual interface (SVI) such as SVI1 or SVI2, or to the virtual interface 310. For example, if a data packet has a destination address in the subnet of 1.x.x.x, there may not be a local interface on BLADE 1 to route the data packet to reach the network 308. Instead the data packet should be routed over the fabric link. Thus, the FIB 314 may contain an entry to rewrite data packets with a destination address in the subnet of 1.x.x.x to the virtual interface 310 based on the Media Access Control (MAC) address of BLADE 2 or, in this case, the SVI associated with the destination network (*e.g.*, SVI2).

[0034] The MAC address may be based on the slot number of the switch linecard 104 within the chassis 102 of the blade switch 100 having the correct local port coupled to the destination network. According to the L2 MAC table 316 for the first switch linecard 300 (L2 standing for Layer 2, the data link layer of the OSI model), the MAC address for BLADE 2, the second switch linecard 302, is 2.

[0035] The FIB 314 entry for data packets with a destination address in the subnet of 1.x.x.x should be to rewrite data packets to SVI2, through the virtual interface 310, as illustrated. Thus, even though multiple switch linecards may be coupled to the virtual interface 310, the SVI addressing used in conjunction with the virtual interface routing in the first portion of intelligent forwarding allows data packets to be forwarded as if the linecards (such as BLADE 1 and BLADE 2) were connected with point-to-point links.

[0036] In a similar manner, the forwarding engine 318, such as EARL, within BLADE 2, the second switch linecard 302, may contain an FIB 320 after configuration of the blade switch 100 that indicates routing incoming data packets to a local interface, such as I_6 and I_7, to an SVI such as SVI1 or SVI2, or to the virtual interface 310. For example, if a data packet is received in the second switch linecard 302 with a destination address in the subnet of 1.x.x.x, the FIB 320 may contain routing instructions to rewrite the packet to SVI2 (using local interface I_6). However, if the data packet has a destination address in the subnet of 2.x.x.x, there may not be a local interface on BLADE 2 to route the data packet to reach the network 306. Thus, the FIB 320 may contain an entry to rewrite data packets with a destination address in the subnet of 2.x.x.x to SVI1 (using virtual interface 310). According to the L2 MAC table 322 for the second switch linecard 302, the MAC address for BLADE 1, the first switch linecard 300, is 1.

[0037] During unicast forwarding according to embodiments of the invention and as illustrated in FIG. 3 as an example, L3 forwarding may occur on the ingress blade along with various input/output features, and then L2 bridging may occur on the output blade. First, an incoming data packet may be received from network 306 on interface I_1 in SVI1. An Access Control List (ACL) and Quality of Service (QoS) lookup may be performed on the input interface (I_1 in this example). ACL, as used herein, may be generally defined as a list of permissions attached to an object specifying who or what is allowed to access the object and what operations are permitted to be performed on the object. QoS, as used herein, may be generally defined as a resource reservation control mechanism, which can provide different

priority to different users or data flows or guarantee a certain level of performance to a data flow in accordance with requests from the application program or the Internet Service Provider (ISP) policy.

[0038] After the ACL and QoS policies are established for the data packet, the ingress forwarding engine 312, such as EARL, may perform an FIB lookup into FIB 314. Since the packet has a destination address in the subnet of 1.x.x.x, the FIB adjacency may rewrite the packet to interface SVI2 through the virtual interface (V) 310 as described above. There need not be any output ACL/QoS policies on V 310, or any input ACL/QoS policies either. The egress forwarding engine 318 may perform a second FIB lookup, and FIB adjacency may rewrite the data packet to SVI2 through interface I_6 according to FIB 320. The data packet may then leave the blade switch 100 out of interface I_6 bound for network 308 with an IP address in the subnet of 1.x.x.x.

Conclusion

[0039] The blade switch may address interface scalability by having each of the linecards maintain its control/forwarding/management functions locally and uses the concept of virtual and local interfaces for intelligent forwarding. From the customer and network perspectives, the blade switch may appear as a single network switch having a single bridge ID. The blade switch may operate as a distributed L2 device with a single L3 interface, where one route processor (RP) interacts with multiple switch processors (SP).

[0040] While the foregoing is directed to embodiments of the present invention, other and further embodiments of the invention may be devised without departing from the basic scope thereof, and the scope thereof is determined by the claims that follow.

Claims:

1. A network switch system, comprising:

a chassis;

a plurality of switch linecards disposed in the chassis and functioning as independent Layer 2 (L2) switches, wherein each of the switch linecards maintains control, forwarding, and interface management functions locally; and

a controller, wherein the controller is adapted to configure the plurality of switch linecards as a single network device having a single bridge identifier (bridge ID).

2. The system of claim 1, wherein each of the switch linecards maintains a forwarding information base (FIB) for directing incoming packets to a local interface or to a virtual interface of another switch linecard in the plurality.

3. The system of claim 1, wherein each of the plurality of switch linecards comprises a switch processor (SP).

4. The system of claim 1, wherein the controller comprises a route processor (RP).

5. The system of claim 1, wherein a first linecard of the plurality of switch linecards runs a different operating system than a second linecard of the plurality.

6. The system of claim 1, wherein a first linecard of the plurality of switch linecards runs a different version of a forwarding engine than a second linecard of the plurality.

7. The system of claim 1, wherein the system appears as the single network device from a control plane perspective and a management plane perspective.

8. The system of claim 1, wherein the virtual interface is an internal local area network (LAN) segment connecting all of the plurality of switch linecards.

9. A method comprising:

receiving a packet in a first switch linecard of a plurality of switch linecards disposed in a chassis, wherein the plurality of switch linecards is configured as a single network switch having a single bridge identifier (bridge ID); and

forwarding the packet from the first switch linecard to a second switch linecard via a virtual interface.

10. The method of claim 9, wherein the first switch linecard runs a different operating system than the second switch linecard.

11. The method of claim 9, further comprising forwarding the packet to a local interface of the second switch linecard after receiving the packet from the virtual interface.

12. The method of claim 11, wherein the local interface of the second switch linecard is a switch virtual interface (SVI).

13. The method of claim 9, wherein forwarding the packet comprises unicast forwarding.

14. The method of claim 9, wherein the first switch linecard maintains a first forwarding information base (FIB) and the second switch linecard maintains a second FIB.

15. The method of claim 9, wherein the virtual interface is an internal local area network (LAN) segment connecting all of the plurality of switch linecards.

16. A method comprising:

configuring a forwarding information base (FIB) of a first switch linecard of a plurality of switch linecards disposed in a chassis to forward a packet to a local interface of the first switch linecard if the packet is to be forwarded within the first switch linecard; and

configuring the FIB of the first switch linecard to forward the packet from the first switch linecard to a virtual interface for bridging between two or more of the

plurality of switch linecards if the packet is to be forwarded to a second switch linecard in the plurality.

17. The method of claim 16, further comprising configuring an FIB of the second switch linecard to forward the packet to a local interface of the second switch linecard if received from the virtual interface.

18. The method of claim 16, further comprising storing the FIB of the first switch linecard on the first switch linecard.

19. The method of claim 16, wherein configuring the FIB is performed by a controller which considers the plurality of switch linecards as a single network switch having a single bridge identifier (bridge ID) during the configuring.

20. The method of claim 16, wherein configuring the FIB of the first switch linecard comprises writing one or more entries in the FIB to forward the packet according to unicast forwarding.

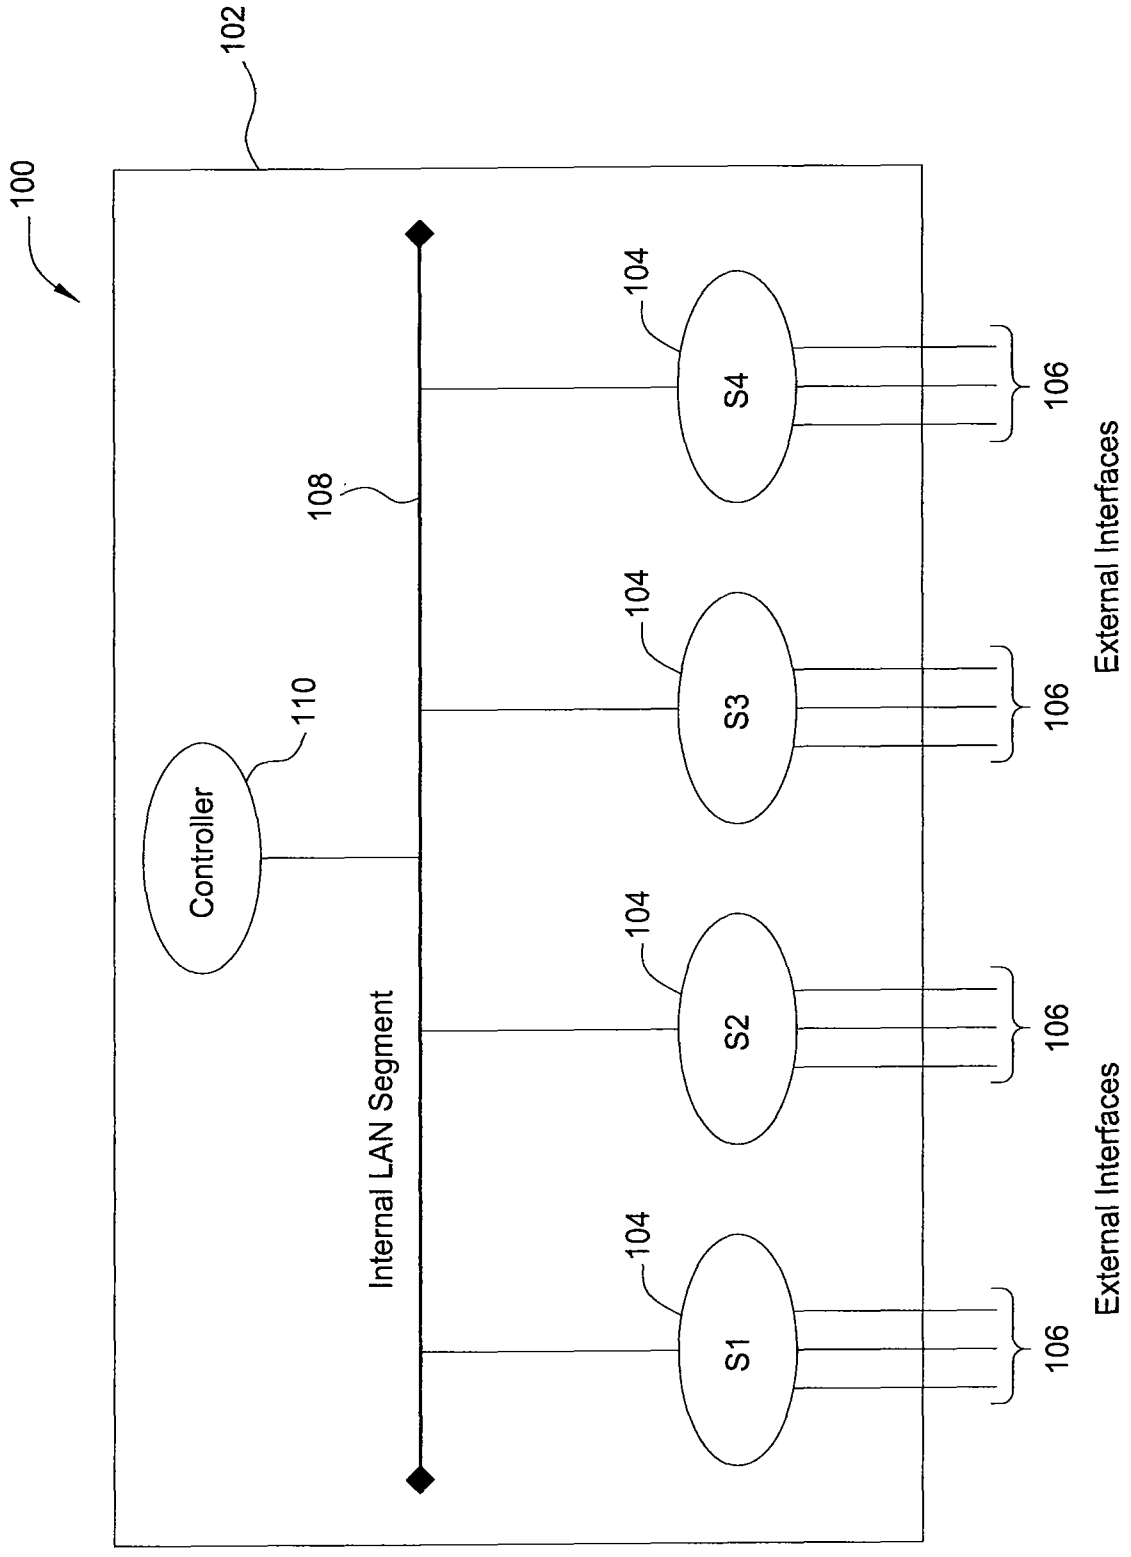

100
102
104
S4
106
108
External Interfaces
104
S3
106
110
Controller
Internal LAN Segment
104
S2
106
External Interfaces
104
S1
106

FIG. 1

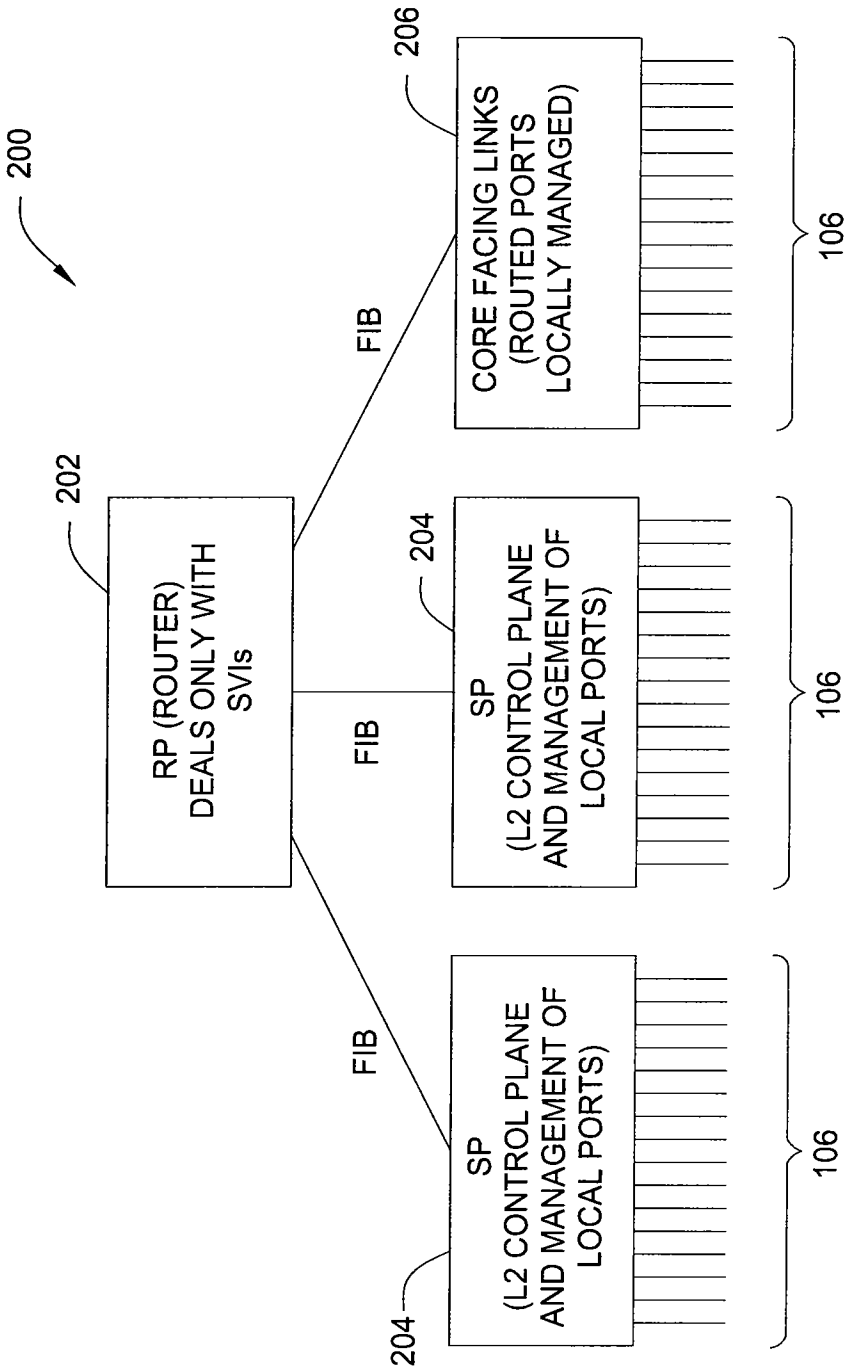

200
202
RP (ROUTER)
DEALS ONLY WITH
SVIs
FIB
FIB
FIB
204
SP
(L2 CONTROL PLANE
AND MANAGEMENT OF
LOCAL PORTS)
204
SP
(L2 CONTROL PLANE
AND MANAGEMENT OF
LOCAL PORTS)
206
CORE FACING LINKS
(ROUTED PORTS
LOCALLY MANAGED)
106
106
106

FIG. 2

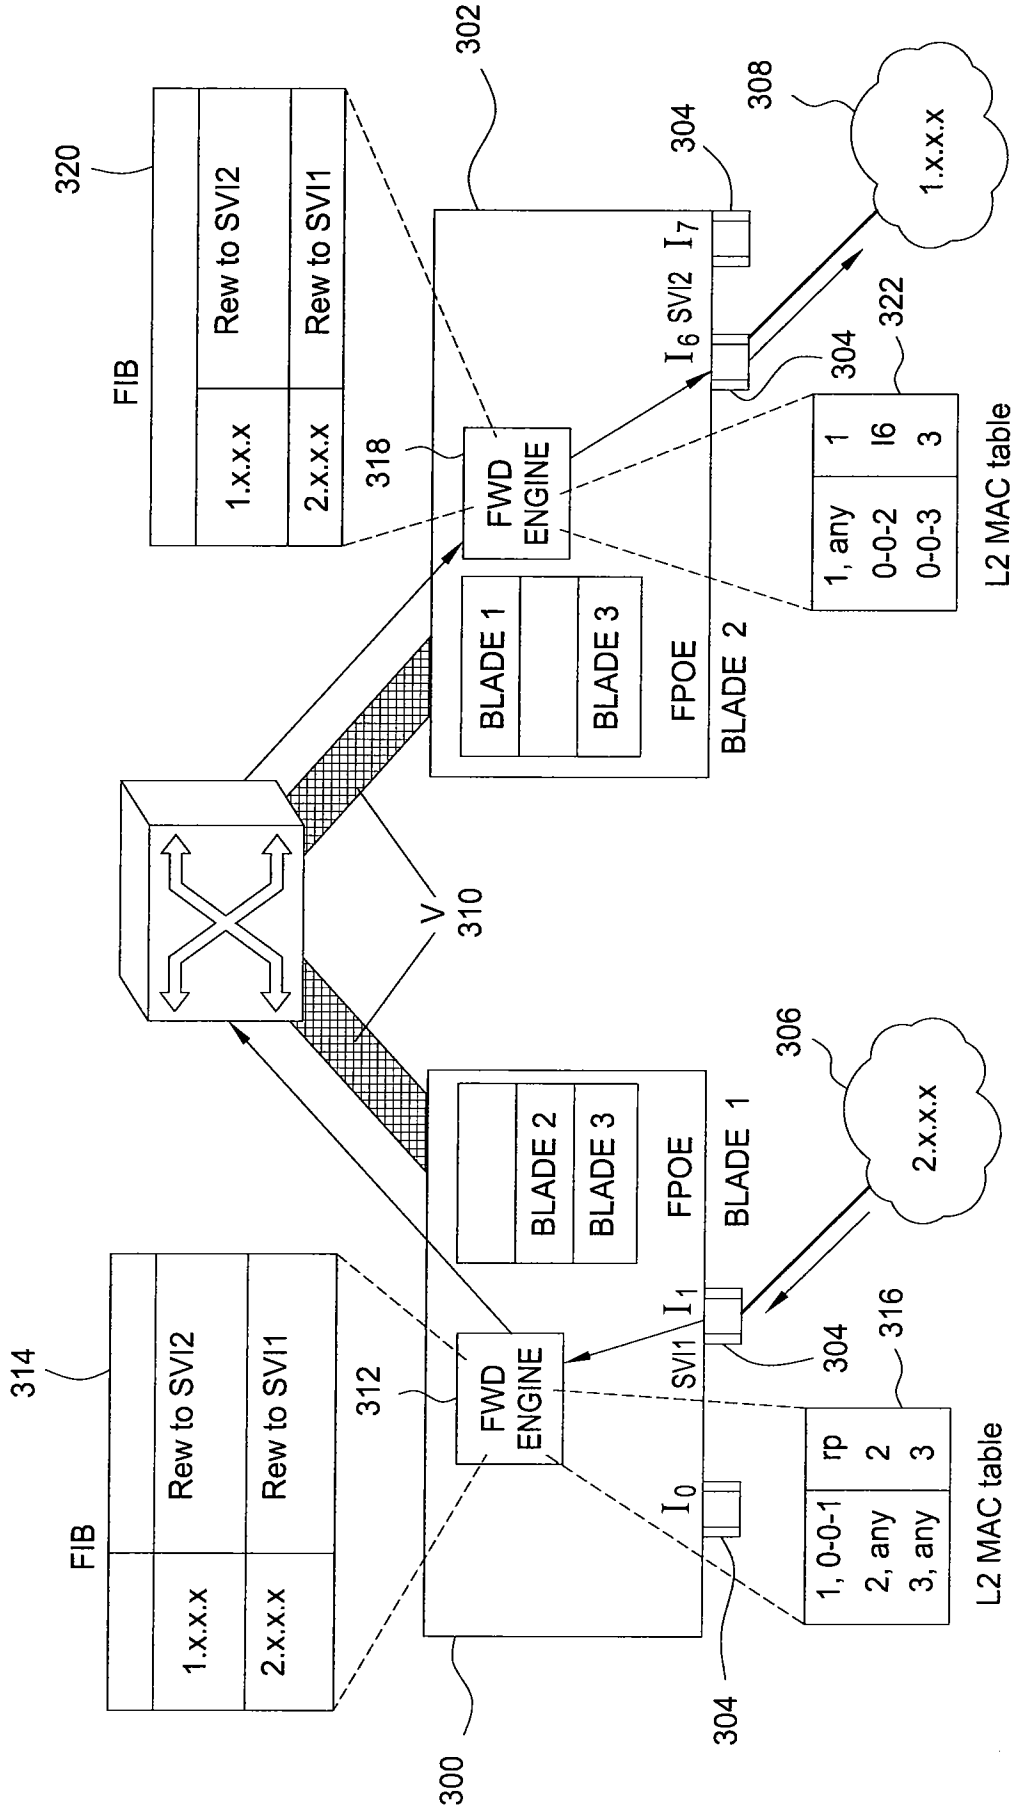

320
FIB
1.x.x.x
2.x.x.x
Rew to SVI2
Rew to SVI1
318
FWD ENGINE
302
304
308
1.x.x.x
I7
I6 SVI2
322
304
1, any 1
0-0-2 I6
0-0-3 3
L2 MAC table
BLADE 1
BLADE 3
FPOE
BLADE 2
V
310
306
2.x.x.x
BLADE 2
BLADE 3
FPOE
BLADE 1
I1
SVI1
304
316
1, 0-0-1 rp
2, any 2
3, any 3
L2 MAC table
I0
304
314
FIB
1.x.x.x
2.x.x.x
Rew to SVI2
Rew to SVI1
312
FWD ENGINE
300

FIG. 3

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2009/035112

A. CLASSIFICATION OF SUBJECT MATTER

H04L 12/56(2006.01)i, H04L 12/28(2006.01)i

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC : H04L, G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean Utility models and applications for Utility models since 1975

Japanese Utility models and applications for Utility models since 1975

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) "switch, network, linecard, interface, controller, bridge identifier"

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y A	WO 2005-029784 A2 (CISCO TECHNOLOGY, INC.) 31 March 2005 See the abstract; page 7, line 21 - page 8, line 28; figures 1, 2A and 2B; claims 1, 13, 14, 30, 33	1, 9, 16 2-8, 10-15, 17-20
Y A	US 2004-0199568 A1 (Martin Lund) 07 October 2004 See paragraph [0029] - [0030], [0032], [0042], [0047]; figures 1-2 and 4	1, 9, 16 2-8, 10-15, 17-20
Y A	US 2006-0221972 A1 (Punit Bhargava et al.) 05 October 2006 See the abstract; paragraph [0007] - [0008]; figure 2	1, 9, 16 2-8, 10-15, 17-20
A	US 2007-0291752 A1 (Nir Ben-Dvora et al.) 20 December 2007 See paragraph [0008], [0028], [0039], [0044]; figures 2A and 3	1-20

☐ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
11 AUGUST 2009 (11.08.2009)	**11 AUGUST 2009 (11.08.2009)**

Name and mailing address of the ISA/KR

Korean Intellectual Property Office
Government Complex-Daejeon, 139 Seonsa-ro, Seo-gu, Daejeon 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

KIM, Sae Young

Telephone No. 82-42-481-5685

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2009/035112

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2005-029784 A2	31.03.2005	AU 2005-239568 A1	10.11.2005
		AU 2004-305790 A1	31.03.2005
		AU 2004-305790 A1	08.09.2004
		AU 2004-305790 B2	16.07.2009
		AU 2005-239568 A1	05.04.2005
		CA 2534510 A1	31.03.2005
		CA 2562013 A1	10.11.2005
		EP 1757037 A1	28.02.2007
		EP 1665652 B1	03.12.2008
		EP 1665652 A2	07.06.2006
		US 2005-0063395 A1	24.03.2005
		US 2005-0163115 A1	28.07.2005
		WO 2005-029784 A3	31.03.2005
		WO 2005-107188 A1	10.11.2005
US 2004-0199568 A1	07.10.2004	NONE	
US 2006-0221972 A1	05.10.2006	US 2006-0222012 A1	05.10.2006
		US 07554928 B2	30.06.2009
US 2007-0291752 A1	20.12.2007	EP 2036260 A2	18.03.2009
		US 07522595 B2	21.04.2009
		WO 2007-149206 A2	27.12.2007
		WO 2007-149206 A3	27.12.2007