

**UNITED STATES DISTRICT COURT
WESTERN DISTRICT OF TEXAS
WACO DIVISION**


SMART PATH CONNECTIONS, LLC

Plaintiff,

V.

JUNIPER NETWORKS, INC.

Defendant.



Civil. Action No. 6:21-CV-00328-ADA

JURY TRIAL DEMANDED

FIRST AMENDED COMPLAINT FOR PATENT INFRINGEMENT

Smart Path Connections, LLC (“Smart Path” or “Plaintiff”), by and through its attorneys, for its Complaint for patent infringement against Juniper Networks, Inc. (“Juniper” or “Defendant”), and demanding trial by jury, hereby alleges, on information and belief with regard to the actions of Defendant and on knowledge with regard to its own actions, as follows:

I. NATURE OF THE ACTION

1. This is an action for patent infringement arising under the patent laws of the United States, 35 U.S.C. §§ 271, *et seq.*, to enjoin and obtain damages resulting from Defendant's unauthorized use, sale, and offer to sell in the United States, of products, methods, processes, services and/or systems that infringe Plaintiff's United States patents, as described herein.

2. Defendant manufactures, provides, uses, sells, offers for sale, imports, and/or distributes infringing products and services, and encourages others to use its products and services in an infringing manner, as set forth herein.

3. Plaintiff seeks past and future damages and prejudgment and post-judgment interest for Defendant's infringement of the Asserted Patents, as defined below.

II. PARTIES

4. Plaintiff Smart Path is a limited liability company organized and existing under the law of the State of Delaware, with its principal place of business located at 601 Quail Valley Drive, Georgetown, TX 78626.

5. Smart Path is the owner of the entire right, title, and interest of the Asserted Patents, as defined below.

6. Juniper Networks, Inc. ("Juniper"), is a Delaware corporation with its principal place of business at 1133 Innovation Way, Sunnyvale, California 94089. Juniper may be served through its registered agent CT Corporation System, 1999 Bryan Street, Suite 900, Dallas, Texas 75201. On information and belief, Juniper is registered to do business in the State of Texas and has been since at least April 27, 2017.

III. JURISDICTION AND VENUE

7. This is an action for patent infringement which arises under the patent laws of the United States, in particular, 35 U.S.C. §§ 271, 281, 283, 284, and 285.

8. This Court has exclusive jurisdiction over the subject matter of this action under 28 U.S.C. §§ 1331 and 1338(a).

9. This Court has personal jurisdiction over Juniper in this action because Juniper has committed acts within the Western District of Texas giving rise to this

action and has established minimum contacts with this forum such that the exercise of jurisdiction over Juniper would not offend traditional notions of fair play and substantial justice. Defendant Juniper, directly and/or through subsidiaries or intermediaries (including distributors, retailers, and others), has committed and continues to commit acts of infringement in this District by, among other things, offering to sell and selling products and/or services that infringe the Asserted Patents. Moreover, Juniper is registered to do business in the State of Texas, has offices and facilities in the State of Texas, and actively directs its activities to customers located in the State of Texas.

10. Venue is proper in this district under 28 U.S.C. §§ 1391(b)–(d) and 1400(b). Defendant Juniper is registered to do business in the State of Texas, has offices in the State of Texas, and upon information and belief, has transacted business in the Western District of Texas and has committed acts of direct and indirect infringement in the Western District of Texas. Juniper maintains a regular and established place of business in the Western District of Texas, including an office located at 1120 South Capital of Texas Highway, Suite 120, First Floor, Building 2, Austin, Texas 78746.

IV. COUNTS OF PATENT INFRINGEMENT

11. Plaintiff alleges that Defendant has infringed and continue to infringe the following United States patents (collectively the “Asserted Patents”):

United States Patent No. 7,386,010 (the “010 Patent”) (Exhibit A)
United States Patent No. 7,463,580 (the “580 Patent”) (Exhibit B)
United States Patent No. 7,551,599 (the “599 Patent”) (Exhibit C)
United States Patent No. 7,697,525 (the “525 Patent”) (Exhibit D)

United States Patent No. 7,961,755 (the “755 Patent”) (Exhibit E)

COUNT ONE
INFRINGEMENT OF U.S. PATENT 7,386,010

12. Plaintiff incorporates by reference the allegations in all preceding paragraphs as if fully set forth herein.

13. The '010 Patent, entitled “Multiprotocol media conversion,” was filed on June 13, 2003 and issued on June 10, 2008.

14. Plaintiff is the assignee and owner of all rights, title and interest to the '010 Patent, including the right to recover for past infringements, and has the legal right to enforce the patent, sue for infringement, and seek equitable relief and damages.

Technical Description

15. The '010 Patent addresses problems in the prior art of “providing different types of Layer 2 network service over a common packet network infrastructure.” 1:12-14.

16. The '010 discloses a solution to this problem in which “interworking of Layer 2 services enables endpoints using disparate protocols to communicate with one another over the same VPN.” 1:62-64.

Direct Infringement

17. Defendant, without authorization or license from Plaintiff, has been and is directly infringing the '010 Patent, either literally or equivalently, as infringement is defined by 35 U.S.C. § 271, including through making, using (including for testing purposes), importing, selling and offering for sale telecommunications equipment

that infringes one or more claims of the '010 Patent. Defendant develops, designs, manufactures, and distributes telecommunications equipment that infringes one or more claims of the '010 Patent. Defendant further provides services that practice methods that infringe one or more claims of the '010 Patent. Defendant is thus liable for direct infringement pursuant to 35 U.S.C. § 271. Exemplary infringing instrumentalities include Juniper Networks MX5 Universal Routing Platform and all other substantially similar products (collectively the "010 Accused Products").

18. Smart Path names this exemplary infringing instrumentality to serve as notice of Defendant's infringing acts, but Smart Path reserves the right to name additional infringing products, known to or learned by Smart Path or revealed during discovery, and include them in the definition of '010 Accused Products.

19. Defendant is liable for direct infringement pursuant to 35 U.S.C. § 271 for the manufacture, sale, offer for sale, importation, or distribution of Defendant's MX5 Universal Routing Platform.

20. Defendant's MX5 Universal Routing Platform is a non-limiting example of an apparatus that meets all limitations of claim 14 of the '010 Patent, either literally or equivalently.

21. The MX5 Universal Routing Platform includes a method for data communications.

Overview

The compact, agile, and full featured MX5 Universal Routing Platform is ideally suited for enterprise and service provider applications, and is optimized for space and power constrained facilities. It is equipped with a Gigabit Ethernet Modular Interface Card (MIC) that permits flexible network connectivity, and a MS-MIC that provides dedicated support for comprehensive flow monitoring.

The MX5 is also software upgradeable to deliver higher performance, port density, and additional services. This investment protecting approach enables customers to increase bandwidth, subscriber, and services scale while minimizing upfront costs.

<https://www.juniper.net/us/en/products-services/routing/mx-series/mx5/>

MX SERIES UNIVERSAL ROUTING PLATFORMS

Product Description

The continuous expansion of mobile, video, and cloud-based services is disrupting traditional networks and impacting the businesses that rely on them. While annual double-digit traffic growth requires massive resource investments to prevent congestion and accommodate unpredictable traffic spikes, capturing return on that investment can be elusive. Emerging trends such as 5G mobility, Internet of Things (IoT) communications, and the continued growth of cloud networking promise even greater network challenges in the near future. The Juniper Networks® MX Series Universal Routing Platform delivers the industry's first end-to-end infrastructure security solution for enterprises as they look to move business-critical applications to public clouds. Delivering features, functionality, and secure services at scale in the 5G era with no compromises, the MX Series is a critical part of the network evolution happening now.

Utilizing state-of-the-art software and hardware innovations, MX Series Universal Routing Platforms are helping network operators worldwide successfully transform their networks and services. Powered by the Juniper Networks Junos® operating system and the programmable Trio chipset, MX Series platforms support a broad set of automation tools and telemetry capabilities that enable a rich set of business- and consumer-oriented services with predictable low latency and wire-rate forwarding at scale, while providing the reliability needed to meet strict service-level agreements (SLAs).

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000597-en.pdf> (Page 1 of PDF)

ATM-to-Ethernet interworking

Supported on M120, M320, and T Series routers; and supported on MX Series routers with aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces. This feature is available on all Enhanced Queuing (EQ) DPCs and Enhanced DPCS for MX Series routers.

This feature is supported on the following products/applications:

Product/Application	Introduced Release
MX5	Junos OS 11.2R4 [†]
MX10	Junos OS 11.2R4 [†]
MX40	Junos OS 11.2R4 [†]
MX80	Junos OS 10.2R1 [†]
MX104	Junos OS 13.2R2 [†]
MX240	Junos OS 10.0R1 [†]
MX480	Junos OS 10.0R1 [†]
MX960	Junos OS 10.0R1 [†]
MX2008	Junos OS 15.1F7
MX2010	Junos OS 12.3R2
MX2020	Junos OS 12.3R1
T640	Junos OS 10.0R1 [†]
T1600	Junos OS 10.0R1 [†]
T4000	Junos OS 12.1R1 [†]
TX Matrix	Junos OS 10.0R1 [†]
TX Matrix Plus	Junos OS 10.0R1 [†]

<https://apps.juniper.net/feature-explorer/feature-info.html?fKey=1544&fn=ATM-to-Ethernet%20interworking>

ATM-to-Ethernet Interworking

The ATM-to-Ethernet interworking feature is useful where ATM2 interfaces are used to terminate ATM DSLAM traffic. The ATM traffic can be forwarded with encapsulation type **ccc** (circuit cross-connect) to a local or remote Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E interface or label-switched path (LSP). The ATM VPI and VCI are converted to stacked VLAN inner and outer VLAN tags.

These ATM-to-Ethernet interworking circuits can be mapped to individual logical interfaces configured on an ATM2 IQ interface and Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E physical interface.

The ATM-to-Ethernet interworking cross-connect essentially provides Layer 2 switching, and statistics are reported at the logical interface level.

During conversion from ATM to Ethernet, the least significant 12 bits of the ATM cell VCI are copied to the Ethernet frame inner VLAN tag. Cells received on an ATM logical interface configured with encapsulation type **vlan-vci-ccc** and falling within the configured VCI range are reassembled into packets and forwarded to a designated Ethernet logical interface that is configured with encapsulation type **vlan-vci-ccc**.

During conversion from Ethernet to ATM, the Ethernet frame inner VLAN tags that fall within the configured range, are copied to the least significant 12 bits of the ATM cell VCI. The ATM logical interface uses its configured VPI when segmenting the Ethernet packets into cells.

ATM-to-Ethernet interworking is supported on M120, M320, and T Series routers.

ATM-to-Ethernet interworking is supported on MX Series routers with aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces. This feature is available on all Enhanced Queuing (EQ) DPCs and Enhanced DPCS for MX Series routers.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Page 174 of PDF)

22. The MX5 Universal Routing Platform comprises linking a plurality of edge devices to communicate with a hub via a network in accordance with a packet-oriented Layer 2 communication protocol.

Example: Configuring ATM-to-Ethernet Interworking

The following example shows the configuration of the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross connect. In the example ATM DSLAM traffic is terminated on an ATM2 interface. The ATM traffic is forwarded using encapsulation type `vlan-vci-ccc` to a local Ethernet IQ2 and IQ2-E interface. See the topology in [Figure 5 on page 159](#).

Figure 5: ATM-to-Ethernet Interworking



In this example, the ATM traffic comes from the DSLAM to the router on ATM interface `at-4/0/0` and is forwarded out on Ethernet interface `ge-2/2/1`.

ATM-To-Ethernet Interworking on ATM MICs

ATM-to-Ethernet interworking supports transmission of ATM packets over Ethernet. It specifically provides support for exchange of Layer 2 and Layer 3 Protocol Data Units (PDUs) between ATM and Ethernet domains. On MX Series 5G Universal Routing Platforms with ATM MICs, you can exchange Ethernet frames between ATM and Ethernet domains over a MPLS pseudowire or a Layer 2 cross-connect by using translational cross connect (TCC). For more information about TCC, see *Circuit and Translational Cross-Connects Overview*.

Consider the following basic ATM-to-Ethernet Interworking topology where the provider edge router PE1 is connected to an ATM domain and the Provider Edge router PE2 is connected to an Ethernet domain

(see Figure 1). The customer edge routers CE1 and CE2 are customer-managed devices. The PE routers are connected by means of an MPLS pseudowire. The ATM traffic on the PE1-CE1 link can comprise untagged Ethernet frames over ATM format. The Ethernet traffic on PE2-CE2 link can comprise untagged, single-VLAN or double-VLAN tagged Ethernet frames depending on the configuration of the PE2 router.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Pages 178 - 180 of PDF)

Example: Configuring ATM-to-Ethernet Interworking on ATM MIC

IN THIS SECTION

- Requirements | 162
- Overview | 162
- Configuration | 163

This example shows how to configure the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross-connect.

Requirements

This example uses the following hardware and software components:

- One MX Series router with ATM MIC
- One MX Series router with Ethernet MIC
- Junos OS Release 16.1R1 or later release

Overview

Configuring ATM-to-Ethernet Interworking enables exchange of Ethernet frames between an ATM domain and an Ethernet domain on MX Series routers with ATM MIC. The ATM domain can be connected to the Ethernet domain over an MPLS pseudowire.

Topology

Consider a sample topology in which provider edge (PE) router (ATMRouter) is an MX Series router with an ATM MIC and PE router (EthernetRouter) is an MX Series router with an Ethernet MIC. CE1 and CE2 are the customer edge routers or customer-managed devices. ATMRouter and EthernetRouter are connected by means of an MPLS pseudowire. The ATM traffic between ATMRouter and CE1 comprises untagged Ethernet over ATM cells. The Ethernet traffic between EthernetRouter and CE2 comprises double-VLAN-tagged Ethernet frames.

When a packet is sent from CE1 to CE2 (ATM-to-Ethernet), ATMRouter accepts ATM cells from CE1 with virtual circuit identifier (VCI) in the range 10/50 to 10/100 and reassembles ATM cells into AAL5 frames. ATMRouter extracts the Ethernet frame from the AAL5 frame payload. ATMRouter adds two VLAN tags with VLAN IDs corresponding to the virtual path identifier (VPI) and VCI of the received ATM cell. The dual-tagged-Ethernet frame is then encapsulated into a MPLS packet and sent over the pseudowire to EthernetRouter.

EthernetRouter strips the MPLS encapsulation and the dual-VLAN-tagged Ethernet frame is sent to CE2. The outer VLAN ID is rewritten to 20 and the inner VLAN ID remains the same. The packet arrives at CE2.

The reverse happens when a packet is sent from CE2 to CE1.


https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (**Pages 181 and 182 of PDF**)

23. The MX5 Universal Routing Platform comprises, at each of the plurality of edge devices, receiving incoming data frames from client nodes in accordance with respective native Layer 2 protocols, at least one of which is different from the packet-oriented Layer 2 communication protocol.

Example: Configuring ATM-to-Ethernet Interworking

The following example shows the configuration of the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross connect. In the example ATM DSLAM traffic is terminated on an ATM2 interface. The ATM traffic is forwarded using encapsulation type `vlan-vci-ccc` to a local Ethernet IQ2 and IQ2-E interface. See the topology in [Figure 5 on page 159](#).

Figure 5: ATM-to-Ethernet Interworking



In this example, the ATM traffic comes from the DSLAM to the router on ATM interface `at-4/0/0` and is forwarded out on Ethernet interface `ge-2/2/1`.

ATM-To-Ethernet Interworking on ATM MICs

ATM-to-Ethernet interworking supports transmission of ATM packets over Ethernet. It specifically provides support for exchange of Layer 2 and Layer 3 Protocol Data Units (PDUs) between ATM and Ethernet domains. On MX Series 5G Universal Routing Platforms with ATM MICs, you can exchange Ethernet frames between ATM and Ethernet domains over a MPLS pseudowire or a Layer 2 cross-connect by using translational cross connect (TCC). For more information about TCC, see *Circuit and Translational Cross-Connects Overview*.

Consider the following basic ATM-to-Ethernet Interworking topology where the provider edge router PE1 is connected to an ATM domain and the Provider Edge router PE2 is connected to an Ethernet domain

(see Figure 1). The customer edge routers CE1 and CE2 are customer-managed devices. The PE routers are connected by means of an MPLS pseudowire. The ATM traffic on the PE1-CE1 link can comprise untagged Ethernet frames over ATM format. The Ethernet traffic on PE2-CE2 link can comprise untagged, single-VLAN or double-VLAN tagged Ethernet frames depending on the configuration of the PE2 router.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Pages 178 - 180 of PDF)

Example: Configuring ATM-to-Ethernet Interworking on ATM MIC

IN THIS SECTION

- Requirements | 162
- Overview | 162
- Configuration | 163

This example shows how to configure the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross-connect.

Requirements

This example uses the following hardware and software components:

- One MX Series router with ATM MIC
- One MX Series router with Ethernet MIC
- Junos OS Release 16.1R1 or later release

Overview

Configuring ATM-to-Ethernet Interworking enables exchange of Ethernet frames between an ATM domain and an Ethernet domain on MX Series routers with ATM MIC. The ATM domain can be connected to the Ethernet domain over an MPLS pseudowire.

Topology

Consider a sample topology in which provider edge (PE) router (ATMRouter) is an MX Series router with an ATM MIC and PE router (EthernetRouter) is an MX Series router with an Ethernet MIC. CE1 and CE2 are the customer edge routers or customer-managed devices. ATMRouter and EthernetRouter are connected by means of an MPLS pseudowire. The ATM traffic between ATMRouter and CE1 comprises untagged Ethernet over ATM cells. The Ethernet traffic between EthernetRouter and CE2 comprises double-VLAN-tagged Ethernet frames.

When a packet is sent from CE1 to CE2 (ATM-to-Ethernet), ATMRouter accepts ATM cells from CE1 with virtual circuit identifier (VCI) in the range 10/50 to 10/100 and reassembles ATM cells into AAL5 frames. ATMRouter extracts the Ethernet frame from the AAL5 frame payload. ATMRouter adds two VLAN tags with VLAN IDs corresponding to the virtual path identifier (VPI) and VCI of the received ATM cell. The dual-tagged-Ethernet frame is then encapsulated into a MPLS packet and sent over the pseudowire to EthernetRouter.

EthernetRouter strips the MPLS encapsulation and the dual-VLAN-tagged Ethernet frame is sent to CE2. The outer VLAN ID is rewritten to 20 and the inner VLAN ID remains the same. The packet arrives at CE2.

The reverse happens when a packet is sent from CE2 to CE1.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Pages 181 and 182 of PDF)

24. The MX5 Universal Routing Platform comprises converting the received incoming data frames at each of the edge devices from at least a first format specified by the native Layer 2 protocols to a second format specified by the packet-oriented Layer 2 communication protocol.

ATM-to-Ethernet Interworking

The ATM-to-Ethernet interworking feature is useful where ATM2 interfaces are used to terminate ATM DSLAM traffic. The ATM traffic can be forwarded with encapsulation type **ccc** (circuit cross-connect) to a local or remote Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E interface or label-switched path (LSP). The ATM VPI and VCI are converted to stacked VLAN inner and outer VLAN tags.

These ATM-to-Ethernet interworking circuits can be mapped to individual logical interfaces configured on an ATM2 IQ interface and Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E physical interface.

The ATM-to-Ethernet interworking cross-connect essentially provides Layer 2 switching, and statistics are reported at the logical interface level.

During conversion from ATM to Ethernet, the least significant 12 bits of the ATM cell VCI are copied to the Ethernet frame inner VLAN tag. Cells received on an ATM logical interface configured with encapsulation type **vlan-vci-ccc** and falling within the configured VCI range are reassembled into packets and forwarded to a designated Ethernet logical interface that is configured with encapsulation type **vlan-vci-ccc**.

During conversion from Ethernet to ATM, the Ethernet frame inner VLAN tags that fall within the configured range, are copied to the least significant 12 bits of the ATM cell VCI. The ATM logical interface uses its configured VPI when segmenting the Ethernet packets into cells.

ATM-to-Ethernet interworking is supported on M120, M320, and T Series routers.


ATM-to-Ethernet interworking is supported on MX Series routers with aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces. This feature is available on all Enhanced Queuing (EQ) DPCs and Enhanced DPCS for MX Series routers.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Page 174 of PDF)

Example: Configuring ATM-to-Ethernet Interworking

The following example shows the configuration of the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross connect. In the example ATM DSLAM traffic is terminated on an ATM2 interface. The ATM traffic is forwarded using encapsulation type **vlan-vci-ccc** to a local Ethernet IQ2 and IQ2-E interface. See the topology in [Figure 5 on page 159](#).

Figure 5: ATM-to-Ethernet Interworking



In this example, the ATM traffic comes from the DSLAM to the router on ATM interface **at-4/0/0** and is forwarded out on Ethernet interface **ge-2/2/1**.


```

[edit interfaces]
ge-2/2/1 {
  vlan-vci-tagging;
  encapsulation vlan-vci-ccc;
  unit 0 {
    encapsulation vlan-vci-ccc;
    vlan-id 100;
    inner-vlan-id-range start 100 end 500;
  }
}
at-4/0/0 {
  atm-options {
    vpi 100;
  }
  unit 0 {
    encapsulation vlan-vci-ccc;
    family ccc;
    vpi 100;
    vci-range start 100 end 500;
  }
}
  
```

Example: Configuring ATM-to-Ethernet Interworking

The following example shows the configuration of the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross connect. In the example ATM DSLAM traffic is terminated on an ATM2 interface. The ATM traffic is forwarded using encapsulation type **vlan-vci-ccc** to a local Ethernet IQ2 and IQ2-E interface. See the topology in [Figure 5 on page 159](#).

Figure 5: ATM-to-Ethernet Interworking



In this example, the ATM traffic comes from the DSLAM to the router on ATM interface **at-4/0/0** and is forwarded out on Ethernet interface **ge-2/2/1**.

```

[edit interfaces]
ge-2/2/1 {
  vlan-vci-tagging;
  encapsulation vlan-vci-ccc;
  unit 0 {
    encapsulation vlan-vci-ccc;
    vlan-id 100;
    inner-vlan-id-range start 100 end 500;
  }
}
at-4/0/0 {
  atm-options {
    vpi 100;
  }
  unit 0 {
    encapsulation vlan-vci-ccc;
    family ccc;
    vpi 100;
    vci-range start 100 end 500;
  }
}

```

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (**Pages 178 - 180 of PDF**)

Example: Configuring ATM-to-Ethernet Interworking on ATM MIC

IN THIS SECTION

- Requirements | 162
- Overview | 162
- Configuration | 163

This example shows how to configure the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross-connect.

Requirements

This example uses the following hardware and software components:

- One MX Series router with ATM MIC
- One MX Series router with Ethernet MIC
- Junos OS Release 16.1R1 or later release

Overview

Configuring ATM-to-Ethernet Interworking enables exchange of Ethernet frames between an ATM domain and an Ethernet domain on MX Series routers with ATM MIC. The ATM domain can be connected to the Ethernet domain over an MPLS pseudowire.

Topology

Consider a sample topology in which provider edge (PE) router (ATMRouter) is an MX Series router with an ATM MIC and PE router (EthernetRouter) is an MX Series router with an Ethernet MIC. CE1 and CE2 are the customer edge routers or customer-managed devices. ATMRouter and EthernetRouter are connected by means of an MPLS pseudowire. The ATM traffic between ATMRouter and CE1 comprises untagged Ethernet over ATM cells. The Ethernet traffic between EthernetRouter and CE2 comprises double-VLAN-tagged Ethernet frames.

When a packet is sent from CE1 to CE2 (ATM-to-Ethernet), ATMRouter accepts ATM cells from CE1 with virtual circuit identifier (VCI) in the range 10/50 to 10/100 and reassembles ATM cells into AAL5 frames. ATMRouter extracts the Ethernet frame from the AAL5 frame payload. ATMRouter adds two VLAN tags with VLAN IDs corresponding to the virtual path identifier (VPI) and VCI of the received ATM cell. The dual-tagged-Ethernet frame is then encapsulated into a MPLS packet and sent over the pseudowire to EthernetRouter.

EthernetRouter strips the MPLS encapsulation and the dual-VLAN-tagged Ethernet frame is sent to CE2. The outer VLAN ID is rewritten to 20 and the inner VLAN ID remains the same. The packet arrives at CE2.

The reverse happens when a packet is sent from CE2 to CE1.


https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Pages 181 and 182 of PDF)

25. The MX5 Universal Routing Platform comprises transmitting the incoming data frames in the second format via the network to the hub.

Example: Configuring ATM-to-Ethernet Interworking

The following example shows the configuration of the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross connect. In the example ATM DSLAM traffic is terminated on an ATM2 interface. The ATM traffic is forwarded using encapsulation type `vlan-vci-ccc` to a local Ethernet IQ2 and IQ2-E interface. See the topology in [Figure 5 on page 159](#).

Figure 5: ATM-to-Ethernet Interworking



In this example, the ATM traffic comes from the DSLAM to the router on ATM interface **at-4/0/0** and is forwarded out on Ethernet interface **ge-2/2/1**.

ATM-To-Ethernet Interworking on ATM MICs

ATM-to-Ethernet interworking supports transmission of ATM packets over Ethernet. It specifically provides support for exchange of Layer 2 and Layer 3 Protocol Data Units (PDUs) between ATM and Ethernet domains. On MX Series 5G Universal Routing Platforms with ATM MICs, you can exchange Ethernet frames between ATM and Ethernet domains over a MPLS pseudowire or a Layer 2 cross-connect by using translational cross connect (TCC). For more information about TCC, see *Circuit and Translational Cross-Connects Overview*.

Consider the following basic ATM-to-Ethernet Interworking topology where the provider edge router PE1 is connected to an ATM domain and the Provider Edge router PE2 is connected to an Ethernet domain

(see Figure 1). The customer edge routers CE1 and CE2 are customer-managed devices. The PE routers are connected by means of an MPLS pseudowire. The ATM traffic on the PE1-CE1 link can comprise untagged Ethernet frames over ATM format. The Ethernet traffic on PE2-CE2 link can comprise untagged, single-VLAN or double-VLAN tagged Ethernet frames depending on the configuration of the PE2 router.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Pages 178 - 180 of PDF)

Example: Configuring ATM-to-Ethernet Interworking on ATM MIC

IN THIS SECTION

- Requirements | 162
- Overview | 162
- Configuration | 163

This example shows how to configure the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross-connect.

Requirements

This example uses the following hardware and software components:

- One MX Series router with ATM MIC
- One MX Series router with Ethernet MIC
- Junos OS Release 16.1R1 or later release

Overview

Configuring ATM-to-Ethernet Interworking enables exchange of Ethernet frames between an ATM domain and an Ethernet domain on MX Series routers with ATM MIC. The ATM domain can be connected to the Ethernet domain over an MPLS pseudowire.

Topology

Consider a sample topology in which provider edge (PE) router (ATMRouter) is an MX Series router with an ATM MIC and PE router (EthernetRouter) is an MX Series router with an Ethernet MIC. CE1 and CE2 are the customer edge routers or customer-managed devices. ATMRouter and EthernetRouter are connected by means of an MPLS pseudowire. The ATM traffic between ATMRouter and CE1 comprises untagged Ethernet over ATM cells. The Ethernet traffic between EthernetRouter and CE2 comprises double-VLAN-tagged Ethernet frames.

When a packet is sent from CE1 to CE2 (ATM-to-Ethernet), ATMRouter accepts ATM cells from CE1 with virtual circuit identifier (VCI) in the range 10/50 to 10/100 and reassembles ATM cells into AAL5 frames. ATMRouter extracts the Ethernet frame from the AAL5 frame payload. ATMRouter adds two VLAN tags with VLAN IDs corresponding to the virtual path identifier (VPI) and VCI of the received ATM cell. The dual-tagged-Ethernet frame is then encapsulated into a MPLS packet and sent over the pseudowire to EthernetRouter.

EthernetRouter strips the MPLS encapsulation and the dual-VLAN-tagged Ethernet frame is sent to CE2. The outer VLAN ID is rewritten to 20 and the inner VLAN ID remains the same. The packet arrives at CE2.

The reverse happens when a packet is sent from CE2 to CE1.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Pages 181 and 182 of PDF)

26. The MX5 Universal Routing Platform comprises transmitting the incoming data frames received from two or more of the client nodes to one of the ports of the hub, and such that converting the received incoming data frames comprises associating the two or more of the client nodes with different, respective Virtual Local Area Networks (VLANs) on the network, such that receiving the incoming data frames comprises receiving the incoming frames through at least one of a time domain multiplexed (TDM) interface and a serial interface, and such that transmitting the incoming data frames comprises transmitting the incoming data frames through an Ethernet port.

ATM-to-Ethernet Interworking

The ATM-to-Ethernet interworking feature is useful where ATM2 interfaces are used to terminate ATM DSLAM traffic. The ATM traffic can be forwarded with encapsulation type **ccc** (circuit cross-connect) to a local or remote Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E interface or label-switched path (LSP). The ATM VPI and VCI are converted to stacked VLAN inner and outer VLAN tags.

These ATM-to-Ethernet interworking circuits can be mapped to individual logical interfaces configured on an ATM2 IQ interface and Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E physical interface.

The ATM-to-Ethernet interworking cross-connect essentially provides Layer 2 switching, and statistics are reported at the logical interface level.

During conversion from ATM to Ethernet, the least significant 12 bits of the ATM cell VCI are copied to the Ethernet frame inner VLAN tag. Cells received on an ATM logical interface configured with encapsulation type **vlan-vci-ccc** and falling within the configured VCI range are reassembled into packets and forwarded to a designated Ethernet logical interface that is configured with encapsulation type **vlan-vci-ccc**.

During conversion from Ethernet to ATM, the Ethernet frame inner VLAN tags that fall within the configured range, are copied to the least significant 12 bits of the ATM cell VCI. The ATM logical interface uses its configured VPI when segmenting the Ethernet packets into cells.

ATM-to-Ethernet interworking is supported on M120, M320, and T Series routers.

ATM-to-Ethernet interworking is supported on MX Series routers with aggregated Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces. This feature is available on all Enhanced Queuing (EQ) DPCs and Enhanced DPCS for MX Series routers.

NOTE: This feature is *not* supported on MX Series routers with ATM interfaces.

For more information on MX Series ATM-to-Ethernet interworking, see the *MX Series Solutions Guide*.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Page 174 of PDF)

Configuring the Inner VLAN Identifier Range

Configure the Ethernet logical interface inner VLAN ID range by including the **inner-vlan-id-range** statement and specifying the starting VLAN ID and ending VLAN ID at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
inner-vlan-id-range start start-id end end-id;
```

VLAN IDs 0 and 4095 are reserved by IEEE 801.1q and must not be used for the inner or outer VLAN ID.

VCIs 0 through 31 are reserved for ATM management purposes by convention. Therefore inner VLAN IDs 1 through 31 should not be used.

VLAN ID 1 might be used by Ethernet switches for certain bridge management services, so using VLAN ID 1 for the inner or outer VLAN ID is discouraged.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Page 176 of PDF)

Example: Configuring ATM-to-Ethernet Interworking

The following example shows the configuration of the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross connect. In the example ATM DSLAM traffic is terminated on an ATM2 interface. The ATM traffic is forwarded using encapsulation type **vlan-vci-ccc** to a local Ethernet IQ2 and IQ2-E interface. See the topology in [Figure 5 on page 159](#).

Figure 5: ATM-to-Ethernet Interworking



In this example, the ATM traffic comes from the DSLAM to the router on ATM interface **at-4/0/0** and is forwarded out on Ethernet interface **ge-2/2/1**.

```
[edit interfaces]
ge-2/2/1 {
  vlan-vci-tagging;
  encapsulation vlan-vci-ccc;
  unit 0 {
    encapsulation vlan-vci-ccc;
    vlan-id 100;
    inner-vlan-id-range start 100 end 500;
  }
}
at-4/0/0 {
  atm-options {
    vpi 100;
  }
  unit 0 {
    encapsulation vlan-vci-ccc;
    family ccc;
    vpi 100;
    vci-range start 100 end 500;
  }
}
```

RELATED DOCUMENTATION

| [Configuring ATM-to-Ethernet Interworking](#) | 154

ATM-To-Ethernet Interworking on ATM MICs

ATM-to-Ethernet interworking supports transmission of ATM packets over Ethernet. It specifically provides support for exchange of Layer 2 and Layer 3 Protocol Data Units (PDUs) between ATM and Ethernet domains. On MX Series 5G Universal Routing Platforms with ATM MICs, you can exchange Ethernet frames between ATM and Ethernet domains over a MPLS pseudowire or a Layer 2 cross-connect by using translational cross connect (TCC). For more information about TCC, see *Circuit and Translational Cross-Connects Overview*.

Consider the following basic ATM-to-Ethernet Interworking topology where the provider edge router PE1 is connected to an ATM domain and the Provider Edge router PE2 is connected to an Ethernet domain

(see Figure 1). The customer edge routers CE1 and CE2 are customer-managed devices. The PE routers are connected by means of an MPLS pseudowire. The ATM traffic on the PE1–CE1 link can comprise untagged Ethernet frames over ATM format. The Ethernet traffic on PE2–CE2 link can comprise untagged, single-VLAN or double-VLAN tagged Ethernet frames depending on the configuration of the PE2 router.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (**Pages 178 - 180 of PDF**)

Example: Configuring ATM-to-Ethernet Interworking on ATM MIC

IN THIS SECTION

- [Requirements | 162](#)
- [Overview | 162](#)
- [Configuration | 163](#)

This example shows how to configure the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross-connect.

Requirements

This example uses the following hardware and software components:

- One MX Series router with ATM MIC
- One MX Series router with Ethernet MIC
- Junos OS Release 16.1R1 or later release

Overview

Configuring ATM-to-Ethernet Interworking enables exchange of Ethernet frames between an ATM domain and an Ethernet domain on MX Series routers with ATM MIC. The ATM domain can be connected to the Ethernet domain over an MPLS pseudowire.

Topology

Consider a sample topology in which provider edge (PE) router (ATMRouter) is an MX Series router with an ATM MIC and PE router (EthernetRouter) is an MX Series router with an Ethernet MIC. CE1 and CE2 are the customer edge routers or customer-managed devices. ATMRouter and EthernetRouter are connected by means of an MPLS pseudowire. The ATM traffic between ATMRouter and CE1 comprises untagged Ethernet over ATM cells. The Ethernet traffic between EthernetRouter and CE2 comprises double-VLAN-tagged Ethernet frames.

When a packet is sent from CE1 to CE2 (ATM-to-Ethernet), ATMRouter accepts ATM cells from CE1 with virtual circuit identifier (VCI) in the range 10/50 to 10/100 and reassembles ATM cells into AAL5 frames. ATMRouter extracts the Ethernet frame from the AAL5 frame payload. ATMRouter adds two VLAN tags with VLAN IDs corresponding to the virtual path identifier (VPI) and VCI of the received ATM cell. The dual-tagged-Ethernet frame is then encapsulated into a MPLS packet and sent over the pseudowire to EthernetRouter.

EthernetRouter strips the MPLS encapsulation and the dual-VLAN-tagged Ethernet frame is sent to CE2. The outer VLAN ID is rewritten to 20 and the inner VLAN ID remains the same. The packet arrives at CE2.

The reverse happens when a packet is sent from CE2 to CE1.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Pages 181 and 182 of PDF)

Understanding Asynchronous Transfer Mode

Asynchronous Transfer Mode (ATM) is a high-speed networking technology that handles data in fixed-size units called cells. It enables high-speed communication between edge routers and core routers in an ATM network.

ATM is designed to facilitate the simultaneous handling of various types of traffic streams (voice, data, and video) at very high speeds over a dedicated connection. ATM uses asynchronous time-division multiplexing (TDM) and it encodes data into 53-byte cells, thereby simplifying the design of hardware and enabling it to quickly determine the destination address of each cell. ATM operates over either fiber optic cables or twisted-pair cables. Each ATM PIC is assigned an ATM switch ID that displays the switch's IP address and the local interface names of the adjacent Fore ATM switches. For information about ATM PICs, see the platform-specific *Hardware Guide*.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf (Page 24 of PDF)

Overview

The compact, agile, and full featured MX5 Universal Routing Platform is ideally suited for enterprise and service provider applications, and is optimized for space and power constrained facilities. It is equipped with a Gigabit Ethernet Modular Interface Card (MIC) that permits flexible network connectivity, and a MS-MIC that provides dedicated support for comprehensive flow monitoring.

The MX5 is also software upgradeable to deliver higher performance, port density, and additional services. This investment protecting approach enables customers to increase bandwidth, subscriber, and services scale while minimizing upfront costs.



20 Gbps
Capacity

20
Gigabit Ethernet
Ports

MS-MIC
with Junos Traffic
Vision

Upgradeable
to
MX10, MX40,
MX80

<https://www.juniper.net/us/en/products-services/routing/mx-series/mx5/>

Willful Infringement

27. On May 9, 2017 Orckit IP LLC sent a letter to Defendant (“Notice Letter”), which informed Defendant of the Asserted Patents.

28. Defendant has had actual knowledge of the ’010 Patent and its infringement thereof at least as of receipt of the Notice Letter.

29. Defendant has had actual knowledge of the ’010 Patent and its infringement thereof at least as of service or other receipt of Plaintiff’s Complaint.

30. Defendant’s infringement of the Asserted Patents was either known or was so obvious that it should have been known to Defendant.

31. Notwithstanding this knowledge, Defendant has knowingly or with reckless disregard infringed the ’010 Patent. Defendant continued to commit acts of infringement despite being on notice of an objectively high likelihood that its actions constituted infringement of Plaintiff’s valid patent rights, either literally or equivalently.

32. Defendant is therefore liable for willful infringement. Accordingly, Plaintiff seeks enhanced damages pursuant to 35 U.S.C. §§ 284 and 285.

Indirect Infringement

33. Defendant has induced and is knowingly inducing its distributors, testers, trainers, customers and/or end users to directly infringe the ’010 Patent, with the specific intent to induce acts constituting infringement, and knowing that the induced acts constitute patent infringement, either literally or equivalently.

34. Defendant has knowingly contributed to direct infringement by its customers and end users by having imported, sold, and/or offered for sale, and knowingly importing, selling, and/or offering to sell within the United States the accused products which are not suitable for substantial non-infringing use and which are especially made or especially adapted for use by its customers in an infringement of the asserted patent.

35. Defendant's indirect infringement includes, for example, providing data sheets, technical guides, demonstrations, software and hardware specifications, installation guides, and other forms of support that induce its customers and/or end users to directly infringe the '010 Patent, including:

- <https://www.juniper.net/us/en/products-services/routing/mx-series/mx5/>
- <https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000597-en.pdf>
- <https://apps.juniper.net/feature-explorer/feature-info.html?fKey=1544&fn=ATM-to-Ethernet%20interworking>
- https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-network-interfaces/atm-interfaces.pdf

36. Defendant's indirect infringement additionally includes marketing its products for import by its customers into the United States. Defendant's indirect infringement further includes providing application notes instructing its customers on infringing uses of the '010 Accused Products. The '010 Accused Products are designed in such a way that when they are used for their intended purpose, the user infringes the '010 Patent, either literally or equivalently. Defendant knows and intends that customers who purchase the '010 Accused Products will use those

products for their intended purpose. For example, Defendant's United States website, <https://www.juniper.net>, instructs customers to use the '010 Accused Products in numerous infringing applications. Defendant's customers directly infringe the '010 patent when they follow Defendant's provided instructions on websites, videos, and elsewhere. Defendant's customers who follow Defendant's provided instructions directly infringe claims of the '010 Patent.

37. In addition, Defendant specifically intends that its customers, such as United States distributors, retailers and consumer product companies, will import, use, and sell infringing products in the United States to serve and develop the United States market for Defendant's infringing products. Defendant knows following its instructions directly infringes claims of the '010 Patent, including for example Claim 1.

38. As a result of Defendant's infringement, Plaintiff has suffered monetary damages, and is entitled to an award of damages adequate to compensate it for such infringement which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT TWO
INFRINGEMENT OF U.S. PATENT 7,463,580

39. Plaintiff incorporates by reference the allegations in preceding paragraphs 1-11 as if fully set forth herein.

40. The '580 Patent, entitled "Resource sharing among network tunnels" was filed on December 15, 2005 and issued on December 9, 2008.

41. Plaintiff is the assignee and owner of all rights, title and interest to the '580 Patent, including the right to recover for past infringements, and has the legal right to enforce the patent, sue for infringement, and seek equitable relief and damages.

Technical Description

42. The '580 Patent addresses problems in the art of Multiprotocol label switching (MPLS), specifically that “tunnel-oriented resource reservation protocols such as RSVP-TE and CR-LDP cited above are typically unable to share resources among communication paths, such as protected paths (except for resource sharing between different instances of the same path, which are not considered to be separate communication paths in this context).” 2:22-27.

43. The '580 Patent discloses that, “[t]he methods and systems described hereinbelow enable resource allocations in network segments and network elements to be shared between two or more communication paths, thus overcoming these shortcomings of the prior art.” 2:27-31.

Direct Infringement

44. Defendant, without authorization or license from Plaintiff, has been and is directly infringing the '580 Patent, either literally or equivalently, as infringement is defined by 35 U.S.C. § 271, including through making, using (including for testing purposes), importing, selling and offering for sale telecommunications equipment that infringes one or more claims of the '580 Patent. Defendant develops, designs, manufactures, and distributes telecommunications equipment that infringes one or more claims of the '580 Patent. Defendant further provides services that practice

methods that infringe one or more claims of the '580 Patent. Defendant is thus liable for direct infringement pursuant to 35 U.S.C. § 271. Exemplary infringing instrumentalities include Juniper Networks MX5 Universal Routing Platform, and all other substantially similar products (collectively the "580 Accused Products").

45. Smart Path names this exemplary infringing instrumentality to serve as notice of Defendant's infringing acts, but Smart Path reserves the right to name additional infringing products, known to or learned by Smart Path or revealed during discovery, and include them in the definition of '580 Accused Products.

46. Defendant is liable for direct infringement pursuant to 35 U.S.C. § 271 for the manufacture, sale, offer for sale, importation, or distribution of Defendant's Juniper Networks MX5 Universal Routing Platform.

47. Defendant's MX5 Universal Routing Platform is a non-limiting example of an apparatus that meets all limitations of claim 8 of the '580 Patent, either literally or equivalently.

48. Defendant's MX5 Universal Routing Platform comprises a network element.



<https://www.juniper.net/us/en/products-services/routing/mx-series/datasheets/1000597.page>

49. Defendant's MX5 Universal Routing Platform comprises a network interface for communicating with other elements in a communication network

The latest generation of line card hardware for the MX960, MX480, and MX240 platforms delivers multi-terabit crypto capabilities with 256-bit encryption complying with AES-GCM encapsulation per RFC4303; AES-GCM encapsulation per RFC4106; AES-GMAC encapsulation per RFC4543; and AES-GMAC (IPv4/v6) encapsulation per RFCs 4302 and 4543. Along with multi-terabit routing, the latest MPC also delivers integrated Layer 2 MACsec features at flexible interface rates of 10GbE, 40GbE, and 100GbE.

<https://www.juniper.net/us/en/products-services/routing/mx-series/datasheets/1000597.page>

50. Defendant's MX5 Universal Routing Platform comprises a processor, which is arranged to accept, via the network interface, a notification distributed over the communication network of an affiliation with a resource-sharing group of at least first and second tunnels, which have respective origin network elements and termination network elements and which traverse different routes through the network.

System Virtualization	Enhanced SLA and queuing	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Junos Fusion Edge (AD)	✓	✓	✓	✓	No	✓	✓	✓	✓	✓	✓
	Logical systems	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Virtual router/switch	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Path Computation Element Protocol (PCEP)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

<https://www.juniper.net/us/en/products-services/routing/mx-series/datasheets/1000597.page>

Path Computation Element Protocol (PCEP) for RSVP-TE

The MPLS RSVP-TE functionality is extended to provide a partial client-side implementation of the stateful Path Computation Element (PCE) architecture (draft-ietf-pce-stateful-pce). The PCE computes path for the traffic engineered LSPs (TE LSPs) of ingress routers that have been configured for external control. The ingress router that connects to a PCE is called a Path Computation Client (PCC). The PCC is configured with the Path Computation Client Protocol (PCEP) (defined in RFC 5440, but limited to the functionality supported on a stateful PCE only) to facilitate external path computing by a PCE. In this new functionality, the active stateful PCE sets parameters for the PCC's TE LSPs, such as bandwidth, path (ERO), and priority. The TE LSP parameters configured from the PCC's CLI are overridden by the PCE-provided parameters. The PCC re-signals the TE LSPs based on the path specified by the PCE. Since the PCE has a global view of the bandwidth demand in the network and performs external path computations after looking up the traffic engineering database, this feature provides a mechanism for offline control of TE LSPs in the MPLS RSVP TE enabled network. To enable external path computing by a PCE, include the `lsp-external-controller` statement on the PCC at the `[edit mpls]` and `[edit mpls lsp lsp-name]` hierarchy levels. To enable PCE to PCC communication, configure `pcep` on the PCC at the `[edit protocols]` hierarchy level.

[https://apps.juniper.net/feature-explorer/feature-info.html?fKey=5872&fn=Path%20Computation%20Element%20Protocol%20\(PCEP\)%20for%20RSVP-TE](https://apps.juniper.net/feature-explorer/feature-info.html?fKey=5872&fn=Path%20Computation%20Element%20Protocol%20(PCEP)%20for%20RSVP-TE)

51. Defendant's MX5 Universal Routing Platform comprises tunnels meet at least one condition selected from a group of conditions consisting of:

- the respective origin network elements of the first and second tunnels are different; and
- the respective termination network elements of the first and second tunnels are different.

PCEP Overview

A Path Computation Element (PCE) is an entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints. A Path Computation Client (PCC) is any client application requesting a path computation to be performed by a PCE. The Path Computation Element Protocol (PCEP) enables communications between a PCC and a PCE, or between two PCEs (defined in RFC 5440).

PCEP is a TCP-based protocol defined by the IETF PCE Working Group, and defines a set of messages and objects used to manage PCEP sessions and to request and send paths for multidomain traffic engineered LSPs (TE LSPs). It provides a mechanism for a PCE to perform path computation for a PCC's external LSPs. The PCEP interactions include LSP status reports sent by the PCC to the PCE, and PCE updates for the external LSPs.

Thus, the PCEP functions include:

- LSP tunnel state synchronization between a PCC and a stateful PCE—When an active stateful PCE connection is detected, a PCC tries to delegate all LSPs to this PCE in a procedure called LSP state synchronization. PCEP enables synchronization of the PCC LSP state to the PCE.
- Delegation of control over LSP tunnels to a stateful PCE—An active stateful PCE controls one or more LSP attributes for computing paths, such as bandwidth, path (ERO), and priority (setup and hold). PCEP enables such delegation of LSPs for path computation.

<https://www.juniper.net/documentation/us/en/software/junos/mpls/mpls.pdf> (Pages 1789 and 1790 of PDF)

Path Computation Element

A Path Computation Element (PCE) can be any entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints. However, a PCE can compute the path for only those TE LSPs of a PCC that have been configured for external control.

A PCE can either be stateful or stateless.

- **Stateful PCE**—A stateful PCE maintains strict synchronization between the PCE and network states (in terms of topology and resource information), along with the set of computed paths and reserved resources in use in the network. In other words, a stateful PCE utilizes information from the traffic engineering database as well as information about existing paths (for example, TE LSPs) in the network when processing new requests from the PCC.

A stateful PCE is of two types:

- **Passive stateful PCE**—Maintains synchronization with the PCC and learns the PCC LSP states to better optimize path calculations, but does not have control over them.
- **Active stateful PCE**—Actively modifies the PCC LSPs, in addition to learning about the PCC LSP states.

A stateful PCE provides the following functions:

- Offers offline LSP path computation.
- Triggers LSP re-route when there is a need to re-optimize the network.
- Changes LSP bandwidth when there is an increase in bandwidth demand from an application.
- Modifies other LSP attributes on the router, such as ERO, setup priority, and hold priority.

A PCE has a global view of the bandwidth demand in the network and maintains a traffic-engineered database to perform path computations. It performs statistics collection from all the routers in the MPLS domain using SNMP and NETCONF. This provides a mechanism for offline control of the PCC's TE LSPs. Although an offline LSP path computation system can be embedded in a network controller, the PCE acts like a full-fledged network controller that provides control over the PCC's TE LSPs, in addition to computing paths.


Although a stateful PCE allows for optimal path computation and increased path computation success, it requires reliable state synchronization mechanisms, with potentially significant control plane overhead and the maintenance of a large amount of data in terms of states, as in the case of a full mesh of TE LSPs.

<https://www.juniper.net/documentation/us/en/software/junos/mpls/mpls.pdf> (Pages 1794 and 1795 of PDF)

Considering the topology used in Figure 132 on page 1748, Figure 134 on page 1754 illustrates the partial client-side PCE implementation in the MPLS RSVP-TE enabled network. The ingress routers A and G are the PCCs that are configured to connect to the external stateful PCE through a TCP connection.

The PCE has a global view of the bandwidth demand in the network and performs external path computations after looking up the traffic engineering database. The active stateful PCE then modifies one or more LSP attributes and sends an update to the PCC. The PCC uses the parameters it receives from the PCE to re-signal the LSP.

Figure 134: Example PCE for MPLS RSVP-TE



<https://www.juniper.net/documentation/us/en/software/junos/mpls/mpls.pdf> (Pages 1798 and 1799 of PDF)

52. Defendant's MX5 Universal Routing Platform comprises the processor comprising a call admission control (CAC) module, which is arranged, when the network element is traversed by at least some of the tunnels in the resource-sharing group, to allocate a resource associated with the network element so as to share an

allocation of the resource among the at least some of the tunnels responsively to the notification.

Description

Configure the bandwidth pools for the class-of-service call admission control (CoS-CAC). Configuring a bandwidth pool provides sufficient bandwidth for bearers to be created or modified. The call admission control (CAC) uses the bandwidth pools to negotiate and reserve bandwidth.

https://www.juniper.net/documentation/en_US/junos-mobility11.2/topics/reference/configuration-statement/bandwidth-pools-edit-unified-edge-cos-cac.html

Path Computation Element

A Path Computation Element (PCE) can be any entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints. However, a PCE can compute the path for only those TE LSPs of a PCC that have been configured for external control.

A PCE can either be stateful or stateless.

- **Stateful PCE**—A stateful PCE maintains strict synchronization between the PCE and network states (in terms of topology and resource information), along with the set of computed paths and reserved resources in use in the network. In other words, a stateful PCE utilizes information from the traffic engineering database as well as information about existing paths (for example, TE LSPs) in the network when processing new requests from the PCC.

A stateful PCE is of two types:

- **Passive stateful PCE**—Maintains synchronization with the PCC and learns the PCC LSP states to better optimize path calculations, but does not have control over them.
- **Active stateful PCE**—Actively modifies the PCC LSPs, in addition to learning about the PCC LSP states.

A stateful PCE provides the following functions:

- Offers offline LSP path computation.
- Triggers LSP re-route when there is a need to re-optimize the network.
- Changes LSP bandwidth when there is an increase in bandwidth demand from an application.
- Modifies other LSP attributes on the router, such as ERO, setup priority, and hold priority.

A PCE has a global view of the bandwidth demand in the network and maintains a traffic-engineered database to perform path computations. It performs statistics collection from all the routers in the MPLS domain using SNMP and NETCONF. This provides a mechanism for offline control of the PCC's TE LSPs. Although an offline LSP path computation system can be embedded in a network controller, the PCE acts like a full-fledged network controller that provides control over the PCC's TE LSPs, in addition to computing paths.

Although a stateful PCE allows for optimal path computation and increased path computation success, it requires reliable state synchronization mechanisms, with potentially significant control plane overhead and the maintenance of a large amount of data in terms of states, as in the case of a full mesh of TE LSPs.

<https://www.juniper.net/documentation/us/en/software/junos/mpls/mpls.pdf> (Pages 1794 and 1795 of PDF)


2. LSP delegation—After the LSP state information is synchronized, the PCC then delegates the external LSPs to one PCE, which is the main active stateful PCE. Only the main PCE can set parameters for the external LSP. The parameters that the main PCE modifies include bandwidth, path (ERO), and priority (setup and hold). The parameters specified in the local configuration are overridden by the parameters that are set by the main PCE.

<https://www.juniper.net/documentation/us/en/software/junos/mpls/mpls.pdf> (Page 1797 of PDF)

Considering the topology used in Figure 132 on page 1748, Figure 134 on page 1754 illustrates the partial client-side PCE implementation in the MPLS RSVP-TE enabled network. The ingress routers A and G are the PCCs that are configured to connect to the external stateful PCE through a TCP connection.

The PCE has a global view of the bandwidth demand in the network and performs external path computations after looking up the traffic engineering database. The active stateful PCE then modifies one or more LSP attributes and sends an update to the PCC. The PCC uses the parameters it receives from the PCE to re-signal the LSP.

Figure 134: Example PCE for MPLS RSVP-TE



<https://www.juniper.net/documentation/us/en/software/junos/mpls/mpls.pdf> (Pages 1798 and 1799 of PDF)

Willful Infringement

53. Defendant has had actual knowledge of the '580 Patent and its infringement thereof at least as of receipt of the Notice Letter.

54. Defendant has had actual knowledge of the '580 Patent and its infringement thereof at least as of service or other receipt of Plaintiff's Complaint.

55. Defendant's infringement of the Asserted Patents was either known or was so obvious that it should have been known to Defendant.

56. Notwithstanding this knowledge, Defendant has knowingly or with reckless disregard infringed the '580 Patent. Defendant continued to commit acts of infringement despite being on notice of an objectively high likelihood that its actions constituted infringement of Plaintiff's valid patent rights, either literally or equivalently.

57. Defendant is therefore liable for willful infringement. Accordingly, Plaintiff seeks enhanced damages pursuant to 35 U.S.C. §§ 284 and 285.

Indirect Infringement

58. Defendant has induced and is knowingly inducing its distributors, testers, trainers, customers and/or end users to directly infringe the '580 Patent, with the specific intent to induce acts constituting infringement, and knowing that the induced acts constitute patent infringement, either literally or equivalently.

59. Defendant has knowingly contributed to direct infringement by its customers and end users by having imported, sold, and/or offered for sale, and knowingly importing, selling, and/or offering to sell within the United States the accused products which are not suitable for substantial non-infringing use and which are especially made or especially adapted for use by its customers in an infringement of the asserted patent.

60. Defendant's indirect infringement includes, for example, providing data sheets, technical guides, demonstrations, software and hardware specifications, installation guides, and other forms of support that induce its customers and/or end users to directly infringe '580 Patent.

61. Defendant's indirect infringement additionally includes marketing its products for import by its customers into the United States. Defendant's indirect infringement further includes providing application notes instructing its customers on infringing uses of the '580 Accused Products. The '580 Accused Products are designed in such a way that when they are used for their intended purpose, the user infringes the '580 Patent, either literally or equivalently. Defendant knows and intends that customers who purchase the '580 Accused Products will use those products for their intended purpose. For example, Defendant's United States website, <https://www.juniper.net>, instructs customers to use the '580 Accused Products in numerous infringing applications. Defendant's customers directly infringe the '580 patent when they follow Defendant's provided instructions on website, videos, and elsewhere. Defendant's customers who follow Defendant's provided instructions directly infringe claims of the '580 Patent.

62. In addition, Defendant specifically intends that its customers, such as United States distributors, retailers and consumer product companies, will import, use, and sell infringing products in the United States to serve and develop the United States market for Defendant's infringing products. Defendant knows following its

instructions directly infringes claims of the '580 Patent, including for example Claims 1 - 7.

63. As a result of Defendant's infringement, Plaintiff has suffered monetary damages, and is entitled to an award of damages adequate to compensate it for such infringement which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT THREE
INFRINGEMENT OF U.S. PATENT 7,551,599

64. Plaintiff incorporates by reference the allegations in preceding paragraphs 1-11 as if fully set forth herein.

65. The '599 Patent, entitled "Layer-3 network routing with RPR layer-2 visibility" was filed on March 29, 2004 and issued on June 23, 2009.

66. Plaintiff is the assignee and owner of all rights, title and interest to the '599 Patent, including the right to recover for past infringements, and has the legal right to enforce the patent, sue for infringement, and seek equitable relief and damages.

Technical Description

67. The '599 Patent addresses technical problems in the prior art, including that "[c]urrently, layer-3 routing protocols, such as RIP and OSPF, are unaware of the topology of layer-2 RPR networks with which they must interact. A routing table allows the router to forward packets from source to destination via the most suitable path, i.e., lowest cost, minimum number of hops. The routing table is updated via the routing protocol, which dynamically discovers currently available paths. The routing table may also be updated via static routes, or can be built using a local interface

configuration, which is updated by the network administrator. However, the RPR ingress and egress nodes chosen in the operation of automatic routing protocols do not take into account the internal links within the RPR ring, and may therefore cause load imbalances within the RPR subnet, which generally results in suboptimum performance of the larger network.” 3:65-4:12.

68. To address these issues, the ’599 Patent discloses “methods and systems are provided for the manipulation of layer-3 network nodes, external routers, routing tables and elements of layer-2 ring networks, such as RPR networks, enabling the layer-3 elements to view the topology of a layer-2 ring subnet. This feature permits routers to choose optimal entry points to the layer-2 subnet for different routes that pass into or through the layer-2 subnet. This enables virtual tunnels or routing paths to utilize all existing entry links to the subnet and to minimize cost factors, such as the number of spans required to traverse the subnet from the entry point to a destination node of the subnet.” 4:17-27.

Direct Infringement

69. Defendant, without authorization or license from Plaintiff, has been and is directly infringing the ’599 Patent, either literally or equivalently, as infringement is defined by 35 U.S.C. § 271, including through making, using (including for testing purposes), importing, selling and offering for sale methods, devices, and networks infringing one or more claims of the ’599 Patent. Defendant develops, designs, manufactures, and distributes telecommunications equipment that infringes one or more claims of the ’599 Patent. Defendant further provides services that practice

methods that infringe one or more claims of the '465 Patent. Defendant is thus liable for direct infringement pursuant to 35 U.S.C. § 271. Exemplary infringing instrumentalities include Juniper MX5 Universal Routing Platform, and all other substantially similar products (collectively the "599 Accused Products").

70. Smart Path names these exemplary infringing instrumentalities to serve as notice of Defendant's infringing acts, however Smart Path reserves the right to name additional infringing products, known to or learned by Smart Path or revealed during discovery, and include them in the definition of '599 Accused Products.

71. Defendant is liable for direct infringement pursuant to 35 U.S.C. § 271 for the use, manufacture, sale, offer of sale, importation, or distribution of Defendant's MX5 Universal Routing Platform.

72. Defendant's MX5 Universal Routing Platform is a non-limiting example of an ethernet switch that meets all limitations of claim 47 of the '599 Patent, either literally or equivalently.

73. Defendant's MX5 Universal Routing Platform comprises a method for obtaining egress from a layer-2 ring network to an external layer-3 network.

Overview

The compact, agile, and full featured MX5 Universal Routing Platform is ideally suited for enterprise and service provider applications, and is optimized for space and power constrained facilities. It is equipped with a Gigabit Ethernet Modular Interface Card (MIC) that permits flexible network connectivity, and a MS-MIC that provides dedicated support for comprehensive flow monitoring.

The MX5 is also software upgradeable to deliver higher performance, port density, and additional services. This investment protecting approach enables customers to increase bandwidth, subscriber, and services scale while minimizing upfront costs.

<https://www.juniper.net/us/en/products-services/routing/mx-series/mx5/>

MX SERIES 5G UNIVERSAL ROUTING PLATFORMS

Product Description

The continuous expansion of mobile, video, and cloud-based services is disrupting traditional networks and impacting the businesses that rely on them. While annual double-digit traffic growth requires massive resource investments to prevent congestion and accommodate unpredictable traffic spikes, capturing return on that investment can be elusive. Emerging trends such as 5G mobility, Internet of Things (IoT) communications, and the continued growth of cloud networking promise even greater network challenges in the near future. The Juniper Networks® MX Series 5G Universal Routing Platform delivers the industry's first end-to-end infrastructure security solution for enterprises as they look to move business-critical applications to public clouds. Delivering features, functionality, and secure services at scale in the 5G era with no compromises, the MX Series is a critical part of the network evolution happening now.

Utilizing state-of-the-art software and hardware innovations, MX Series 5G Universal Routing Platforms are helping network operators worldwide successfully transform their networks and services. Powered by the Juniper Networks Junos® operating system and the programmable Trio chipset, MX Series platforms support a broad set of automation tools and telemetry capabilities that enable a rich set of business- and consumer-oriented services with predictable low latency and wire-rate forwarding at scale, while providing the reliability needed to meet strict service-level agreements (SLAs).

- **Mobile Backhaul:** In addition to switching, routing, and security features, MX Series platforms support highly scalable and reliable hardware-based timing that meets the strictest LTE requirements, including Synchronous Ethernet for frequency and the Precision Time Protocol (PTP) for frequency and phase synchronization. In addition, the MX104 is ETSI 300-compliant, enabling deployment in next-generation mobile applications such as 5G.

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000597-en.pdf> (Page 1 of PDF)

Pseudowire redundancy in mobile backhaul scenarios

Operators are migrating to an Ethernet-based higher capacity infrastructure in the backhaul portion of 3G/LTE topologies. Modern base stations provide Ethernet backhaul connectivity, allowing pseudowire technologies to transport end-user content to the desired destination. As part of this transition, better resiliency mechanisms are needed to cover the gap with those features provided by legacy technologies.

This Junos OS Release 13.2 and later feature enables you to configure pseudowire redundancy where Layer 2 and Layer 3 segments are interconnected in a mobile backhaul scenario. This is useful for deploying packet-based backhaul networks that offer increased capacity at lower cost, while providing the necessary service reliability and quality of experience that users expect.

More information about this feature: [Technical Documentation](#) 

This feature is supported on the following products/applications:

Product/Application	Introduced Release
MX5	Junos OS 13.2R1†
MX10	Junos OS 13.2R1†
MX40	Junos OS 13.2R1†
MX80	Junos OS 13.2R1†
MX150	Junos OS 17.3R1
MX204	Junos OS 17.4R1
MX240	Junos OS 13.2R1†
MX480	Junos OS 13.2R1†
MX960	Junos OS 13.2R1†
MX2008	Junos OS 15.1F7
MX2010	Junos OS 13.2R1†
MX2020	Junos OS 13.2R1†
MX10003	Junos OS 17.3R1
MX10008	Junos OS 18.2R1
vMX	Junos OS 14.1R5†
T640	Junos OS 13.2R1†
T1600	Junos OS 13.2R1†
T4000	Junos OS 13.2R1†
TX Matrix	Junos OS 13.2R1†
TX Matrix Plus	Junos OS 13.2R1†

<https://apps.juniper.net/feature-explorer/feature-info.html?fKey=6113&fn=Pseudowire%20redundancy%20in%20mobile%20backhaul%20scenarios>

Understanding Pseudowire Redundancy Mobile Backhaul Scenarios

IN THIS SECTION

- Sample Topology | 353
- Benefits of Pseudowire Redundancy Mobile Backhaul | 353
- Layer 2 Virtual Circuit Status TLV Extension | 354
- How It Works | 355


With the rising demand for mobile broadband services, telecommunication providers are seeing a sharp increase in bandwidth requirements. To keep pace with demand, operators are deploying packet-based backhaul networks that offer increased capacity at a lower cost, while providing the necessary service reliability and quality of experience that users expect.

Most of the legacy backhaul infrastructure has been traditionally built over PDH microwave, TDM T1/E1, or ATM-over-DSL links. Service providers have traditionally added subsequent TDM links to their base stations when needed to deal with bandwidth constraint scenarios. This expansion model has proven to be inefficient for the unprecedented traffic demands required by 3G and Long Term Evolution (LTE) services. As a direct consequence, operators are gradually migrating to an Ethernet-based higher capacity infrastructure in the backhaul portion of 3G and LTE topologies. Modern base stations now provide Ethernet backhaul connectivity, allowing pseudowire technologies to transport end-user content to the desired destination. As part of this Ethernet transition, service providers are increasingly demanding better resiliency mechanisms to cover the existence gap with those features provided by previous legacy technologies. With that goal in mind, Junos OS provides efficient pseudowire redundancy capabilities to those topologies where Layer 2 and Layer 3 segments are interconnected.

Sample Topology

Figure 29 on page 353 shows a sample topology.


Figure 29: Pseudowire Redundancy Mobile Backhaul Sample Topology



https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-vpns/config-guide-vpns-layer-2.pdf (Page 386 and 387 of PDF)

74. Defendant's MX5 Universal Routing Platform comprises, in nodes of said ring network creating entries in a host table, each of said entries comprising an address of a respective one of said nodes of said ring network and a metric determined responsively to a topology of the ring network.

Figure 29: Pseudowire Redundancy Mobile Backhaul Sample Topology



https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-vpns/config-guide-vpns-layer-2.pdf (Page 387 of PDF)

An attachment circuit (AC) is a physical or virtual circuit (VC) that attaches a CE device to a PE device. Local preference is used to provide better information than the multiple exit discriminator (MED) value provides for a packet's path selection. You can configure the local preference attribute so that it has a higher value for prefixes received from a router that provides a desired path than prefixes received from a router that provides a less desirable path. The higher the value, the more preferred the route. The local preference attribute is the metric most often used in practice to express preferences for one set of paths over another.

If the Layer 2 circuit is primary, the corresponding PE device advertises the AC's subnet with the higher local preference. All aggregation PE devices initially advertise the AC's subnet with the same local preference. You can configure a routing policy to allow a higher local preference value to be advertised if the Layer 2 VC is active.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-vpns/config-guide-vpns-layer-2.pdf (Page 390 of PDF)

Device A1

```

set interfaces ge-1/3/0 unit 0 family inet address 10.20.0.100/24
set interfaces ge-1/3/0 unit 0 family iso
set interfaces ge-1/3/0 unit 0 family mpls
set interfaces ge-1/3/1 unit 0 family inet address 10.10.0.100/24
set interfaces ge-1/3/1 unit 0 family iso
set interfaces ge-1/3/1 unit 0 family mpls
set interfaces ge-1/3/2 vlan-tagging
set interfaces ge-1/3/2 encapsulation vlan-ccc
set interfaces ge-1/3/2 unit 600 encapsulation vlan-ccc
set interfaces ge-1/3/2 unit 600 vlan-id 600
set interfaces ge-1/3/2 unit 600 family ccc
set interfaces lo0 unit 0 family inet address 192.168.0.100/32 primary
set interfaces lo0 unit 0 family iso address 49.0002.0192.0168.0100.00
set routing-options router-id 192.168.0.100
set routing-options autonomous-system 64510
set routing-options forwarding-table export pplb
set protocols rsvp interface ge-1/3/0.0
set protocols rsvp interface ge-1/3/1.0
set protocols rsvp interface lo0.0
set protocols mpls interface ge-1/3/0.0
set protocols mpls interface ge-1/3/1.0
set protocols isis interface ge-1/3/0.0
set protocols isis interface ge-1/3/1.0
set protocols isis interface lo0.0
set protocols ldp interface ge-1/3/0.0
set protocols ldp interface ge-1/3/1.0
set protocols ldp interface lo0.0
set protocols l2circuit neighbor 192.168.0.101 interface ge-1/3/2.600 virtual-circuit-id 1
set protocols l2circuit neighbor 192.168.0.101 interface ge-1/3/2.600 pseudowire-status-tlv
set protocols l2circuit neighbor 192.168.0.101 interface ge-1/3/2.600 revert-time 10 maximum 60
set protocols l2circuit neighbor 192.168.0.101 interface ge-1/3/2.600 backup-neighbor 192.168.0.102
    virtual-circuit-id 2
set protocols l2circuit neighbor 192.168.0.101 interface ge-1/3/2.600 backup-neighbor 192.168.0.102
    hot-standby
set policy-options policy-statement pplb then load-balance per-packet

```

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-vpns/config-guide-vpns-layer-2.pdf (Pages 393 and 394 of PDF)

17. Configure the routing instance.

This routing instance is in the Layer 2 domain where Device PE1 and Device PE2 are interconnected to the metro ring over multiaccess media (Ethernet). You must include the **vrf-table-label** statement on Device PE1 and Device PE2 to enable advertisement of the direct subnet prefix corresponding to the logical tunnel (lt-) interface toward the Layer 3 domain.

Device PE1 and Device PE2 use OSPF for Layer 3 VPN communication with Device CE1.

```
[edit routing-instances l3vpn]
user@PE1# set instance-type vrf
```

```
user@PE1# set interface lt-1/2/0.601
user@PE1# set interface lo0.1
user@PE1# set route-distinguisher 192.168.1.101:64511
user@PE1# set vrf-import l3vpn_import
user@PE1# set vrf-export l3vpn_export
user@PE1# set vrf-table-label
user@PE1# set protocols ospf export ospf_export
user@PE1# set protocols ospf area 0.0.0.0 interface lt-1/2/0.601
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.1
```

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-vpns/config-guide-vpns-layer-2.pdf (Pages 407 and 408 of PDF)

OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions, making route calculations using the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm). Each router running OSPF floods link-state advertisements throughout the AS or area that contain information about that router's attached interfaces and routing metrics. Each router uses the information in these link-state advertisements to calculate the least cost path to each network and create a routing table for the protocol.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-routing/config-guide-ospf.pdf (Page 24 of PDF)

Understanding OSPF External Metrics

When OSPF exports route information from external autonomous systems (ASs), it includes a cost, or *external metric*, in the route. OSPF supports two types of external metrics: Type 1 and Type 2. The difference between the two metrics is how OSPF calculates the cost of the route.

- Type 1 external metrics are equivalent to the link-state metric, where the cost is equal to the sum of the internal costs plus the external cost. This means that Type 1 external metrics include the external cost to the destination as well as the cost (metric) to reach the AS boundary router.
- Type 2 external metrics are greater than the cost of any path internal to the AS. Type 2 external metrics use only the external cost to the destination and ignore the cost (metric) to reach the AS boundary router.

By default, OSPF uses the Type 2 external metric.


Both Type 1 and Type 2 external metrics can be present in the AS at the same time. In that event, Type 1 external metrics always takes the precedence.

Type 1 external paths are always preferred over Type 2 external paths. When all paths are Type 2 external paths, the paths with the smallest advertised Type 2 metric are always preferred.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-routing/config-guide-ospf.pdf (Page 31 of PDF)

75. Defendant's MX5 Universal Routing Platform comprises defining paths from said nodes through egress nodes of said ring network to external elements in said external layer-3 network.


Figure 29: Pseudowire Redundancy Mobile Backhaul Sample Topology



https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-vpns/config-guide-vpns-layer-2.pdf (Page 387 of PDF)

76. Defendant's MX5 Universal Routing Platform comprises selecting one of said paths responsively to said metric.

Figure 29: Pseudowire Redundancy Mobile Backhaul Sample Topology



https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-vpns/config-guide-vpns-layer-2.pdf (Page 387 of PDF)

An attachment circuit (AC) is a physical or virtual circuit (VC) that attaches a CE device to a PE device. Local preference is used to provide better information than the multiple exit discriminator (MED) value provides for a packet's path selection. You can configure the local preference attribute so that it has a higher value for prefixes received from a router that provides a desired path than prefixes received from a router that provides a less desirable path. The higher the value, the more preferred the route. The local preference attribute is the metric most often used in practice to express preferences for one set of paths over another.

If the Layer 2 circuit is primary, the corresponding PE device advertises the AC's subnet with the higher local preference. All aggregation PE devices initially advertise the AC's subnet with the same local preference. You can configure a routing policy to allow a higher local preference value to be advertised if the Layer 2 VC is active.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-vpns/config-guide-vpns-layer-2.pdf (Page 390 of PDF)

17. Configure the routing instance.

This routing instance is in the Layer 2 domain where Device PE1 and Device PE2 are interconnected to the metro ring over multiaccess media (Ethernet). You must include the **vrf-table-label** statement on Device PE1 and Device PE2 to enable advertisement of the direct subnet prefix corresponding to the logical tunnel (lt-) interface toward the Layer 3 domain.

Device PE1 and Device PE2 use OSPF for Layer 3 VPN communication with Device CE1.

```
[edit routing-instances l3vpn]
user@PE1# set instance-type vrf
```

```
user@PE1# set interface lt-1/2/0.601
user@PE1# set interface lo0.1
user@PE1# set route-distinguisher 192.168.1.101:64511
user@PE1# set vrf-import l3vpn_import
user@PE1# set vrf-export l3vpn_export
user@PE1# set vrf-table-label
user@PE1# set protocols ospf export ospf_export
user@PE1# set protocols ospf area 0.0.0.0 interface lt-1/2/0.601
user@PE1# set protocols ospf area 0.0.0.0 interface lo0.1
```

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-vpns/config-guide-vpns-layer-2.pdf (Pages 407 and 408 of PDF)

OSPF is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). OSPF uses link-state information to make routing decisions, making route calculations using the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm). Each router running OSPF floods link-state advertisements throughout the AS or area that contain information about that router's attached interfaces and routing metrics. Each router uses the information in these link-state advertisements to calculate the least cost path to each network and create a routing table for the protocol.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-routing/config-guide-ospf.pdf (Page 24 of PDF)

OSPF Routing Algorithm

OSPF uses the shortest-path-first (SPF) algorithm, also referred to as the Dijkstra algorithm, to determine the route to each destination. All routing devices in an area run this algorithm in parallel, storing the results in their individual topological databases. Routing devices with interfaces to multiple areas run multiple copies of the algorithm. This section provides a brief summary of how the SPF algorithm works.

When a routing device starts, it initializes OSPF and waits for indications from lower-level protocols that the router interfaces are functional. The routing device then uses the OSPF hello protocol to acquire neighbors, by sending hello packets to its neighbors and receiving their hello packets.

On broadcast or nonbroadcast multiaccess networks (physical networks that support the attachment of more than two routing devices), the OSPF hello protocol elects a designated router for the network. This routing device is responsible for sending *link-state advertisements* (LSAs) that describe the network, which reduces the amount of network traffic and the size of the routing devices' topological databases.

The routing device then attempts to form *adjacencies* with some of its newly acquired neighbors. (On multiaccess networks, only the designated router and backup designated router form adjacencies with other routing devices.) Adjacencies determine the distribution of routing protocol packets. Routing protocol packets are sent and received only on adjacencies, and topological database updates are sent only along adjacencies. When adjacencies have been established, pairs of adjacent routers synchronize their topological databases.

A routing device sends LSA packets to advertise its state periodically and when its state changes. These packets include information about the routing device's adjacencies, which allows detection of nonoperational routing devices.

Using a reliable algorithm, the routing device floods LSAs throughout the area, which ensures that all routing devices in an area have exactly the same topological database. Each routing device uses the information in its topological database to calculate a shortest-path tree, with itself as the root. The routing device then uses this tree to route network traffic.


The description of the SPF algorithm up to this point has explained how the algorithm works within a single area (*intra-area routing*). For internal routers to be able to route to destinations outside the area (*interarea routing*), the area border routers must inject additional routing information into the area. Because the area border routers are connected to the backbone, they have access to complete topological data about the backbone. The area border routers use this information to calculate paths to all destinations outside its area and then advertise these paths to the area's internal routers.

Autonomous system (AS) boundary routers flood information about external autonomous systems throughout the AS, except to stub areas. Area border routers are responsible for advertising the paths to all AS boundary routers.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-routing/config-guide-ospf.pdf (Pages 25 and 26 of PDF)

77. Defendant's MX5 Universal Routing Platform comprises transmitting data from at least one of said nodes via said selected one of said paths to network elements that are external to said ring network.

Figure 29: Pseudowire Redundancy Mobile Backhaul Sample Topology



https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-vpns/config-guide-vpns-layer-2.pdf (Page 387 of PDF)

Willful Infringement

78. Defendant has had actual knowledge of the '599 Patent and its infringement thereof at least as of receipt of the Notice Letter.

79. Defendant has had actual knowledge of the '599 Patent and its infringement thereof at least as of service or other receipt of Plaintiff's Complaint.

80. Defendant's risk of infringement of the Asserted Patents was either known or was so obvious that it should have been known to Defendant.

81. Notwithstanding this knowledge, Defendant has knowingly or with reckless disregard willfully infringed the '599 Patent. Defendant has thus had actual notice

of the infringement of the '599 Patent and acted despite an objectively high likelihood that its actions constituted infringement of Plaintiff's valid patent rights, either literally or equivalently.

82. This objective risk was either known or so obvious that it should have been known to Defendant. Accordingly, Plaintiff seeks enhanced damages pursuant to 35 U.S.C. §§ 284 and 285.

Indirect Infringement

83. Defendant has induced and is knowingly inducing its customers and/or end users to directly infringe the '599 Patent, with the specific intent to encourage such infringement, and knowing that the induced acts constitute patent infringement, either literally or equivalently.

84. Defendant has knowingly contributed to direct infringement by its customers by having imported, sold, and/or offered for sale, and knowingly importing, selling, and/or offering to sell within the United States the '599 Accused Products which are not suitable for substantial non-infringing use and which are especially made or especially adapted for use by its customers in an infringement of the asserted patent.

85. Defendant's indirect infringement includes, for example, providing data sheets, technical guides, demonstrations, software and hardware specifications, installation guides, and other forms of support that induce its customers and/or end users to directly infringe '599 Patent.

86. Defendant's indirect infringement additionally includes marketing its products for import by its customers into the United States. Defendant's indirect

infringement further includes providing application notes instructing its customers on infringing uses of the accused products. The '599 Accused Products are designed in such a way that when they are used for their intended purpose, the user infringes the '599 Patent, either literally or equivalently. Defendant knows and intends that customers who purchase the '599 Accused Products will use those products for their intended purpose. For example, Defendant's United States website <https://www.juniper.net>, instructs customers to use the '599 Accused Products in numerous infringing applications. Defendant's customers directly infringe the '599 Patent when they follow Defendant's provided instructions on its website, videos, and elsewhere. Defendant's customers who follow Defendant's provided instructions directly infringe claims of the '599 Patent.

87. In addition, Defendant specifically intends that its customers, such as United States distributors, retailers and consumer product companies, will import, use, and sell infringing products in the United States to serve and develop the United States market for Defendant's infringing products.

88. As a result of Defendant's infringement, Plaintiff has suffered monetary damages, and is entitled to an award of damages adequate to compensate it for such infringement which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT FOUR
INFRINGEMENT OF U.S. PATENT 7,697,525

89. Plaintiff incorporates by reference the allegations in preceding paragraphs 1-11 as if fully set forth herein.

90. The '525 Patent, entitled "Forwarding multicast traffic over link aggregation ports" was filed on December 21, 2006 and issued on April 13, 2010.

91. Plaintiff is the assignee and owner of all rights, title and interest to the '525 Patent, including the right to recover for past infringements, and has the legal right to enforce the patent, sue for infringement, and seek equitable relief and damages.

Technical Description

92. The '525 Patent provides a solution to the problems in the prior art as follows, "[u]nlike some known methods and systems in which all multicast packets are sent to the same LAG group port, the methods and systems described herein distribute multicast packets approximately evenly among the different output ports of the LAG group. Thus, the traffic load within the group is balanced, and distribution of additional unicast traffic across the group is simplified." 3:54-60.

Direct Infringement

93. Defendant, without authorization or license from Plaintiff, has been and is directly infringing the '525 Patent, either literally or equivalently, as infringement is defined by 35 U.S.C. § 271, including through making, using (including for testing purposes), importing, selling and offering for sale methods, devices, and networks infringing one or more claims of the '525 Patent. Defendant develops, designs, manufactures, and distributes telecommunications equipment that infringes one or more claims of the '525 Patent. Defendant further provides services that practice methods that infringe one or more claims of the '525 Patent. Defendant is thus liable for direct infringement pursuant to 35 U.S.C. § 271. Exemplary infringing

instrumentalities include Juniper Networks EX4600 Ethernet Switches, and all other substantially similar products (collectively the “’525 Accused Products”).

94. Smart Path names this exemplary infringing instrumentality to serve as notice of Defendant’s infringing acts, but Smart Path reserves the right to name additional infringing products, known to or learned by Smart Path or revealed during discovery, and include them in the definition of ’525 Accused Products.

95. Defendant is liable for direct infringement pursuant to 35 U.S.C. § 271 for the use, manufacture, sale, offer of sale, importation, or distribution of Defendant’s EX4600 Ethernet Switches.

96. Defendant’s EX4600 Ethernet Switch is a non-limiting example of an ethernet switch that meets all limitations of claim 12 of the ’525 Patent, either literally or equivalently.

97. Defendant’s EX4600 Ethernet Switch comprises a method for communication.

EX4600

The EX4600 Ethernet Switch offers a compact, highly scalable, high-performance 10GbE solution for enterprise campus distribution deployments as well as low-density [data center](#) top-of-rack environments. The EX4600 switch delivers rich telemetry data for [Wired Assurance](#), bringing [AI-powered automation](#) and service levels to access switching.

The EX4600 Ethernet Switch offers a compact, highly scalable, high-performance 10GbE solution for enterprise campus distribution deployments as well as low-density [data center](#) top-of-rack environments. The EX4600 is cloud-ready and ZTP-enabled, so it can be onboarded, provisioned and managed with [Wired Assurance](#), to deliver better experiences for connected devices.

Key Features

Port Density	24 x 1/10GbE and 4 QSFP+ 40GbE plus expansion slots
Form Factor	1 RU
MACsec	400 Gbps of near-line encryption
Switch capacity	720 Gbps
Fabric	Virtual Chassis, MC-LAG

<https://www.juniper.net/us/en/products-services/switching/ex-series/ex4600/>

EX4600 ETHERNET SWITCH

Product Description

Featuring up to 72 wire-speed 10GbE small form-factor pluggable and pluggable plus transceiver (SFP/SFP+) ports, and up to 12 wire-speed 40GbE quad SFP+ transceiver (QSFP+) ports in a compact one rack unit (1 U) platform, the Juniper Networks® EX4600 Ethernet Switch delivers 1.44 Tbps of Layer 2 and Layer 3 connectivity to networked devices such as secure routers, servers, and other switches. The EX4600 base switch provides 24 fixed 1GbE SFP/10GbE SFP+ ports¹ and 4 fixed 40GbE QSFP+ ports, providing the flexibility to support mixed 1GbE, 10GbE and 40GbE environments. A total of four models are available: two featuring AC power supplies and front-to-back or back-to-front airflow; and two featuring DC power supplies and front-to-back or back-to-front airflow. Each model includes dual power supplies.

All versions feature two expansion slots that can accommodate optional expansion modules, providing tremendous configuration and deployment flexibility for enterprise distribution networks. Two expansion modules are available:

- 8xGBASE/10GBASE SFP/SFP+ fiber expansion module²
- 4x40GbE QSFP+ expansion module³

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 1 of PDF)

- MAC learning disable
- Link Aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)
- MAC notification
- MAC address aging configuration
- MAC address filtering
- Persistent MAC (sticky MAC)

Link Aggregation

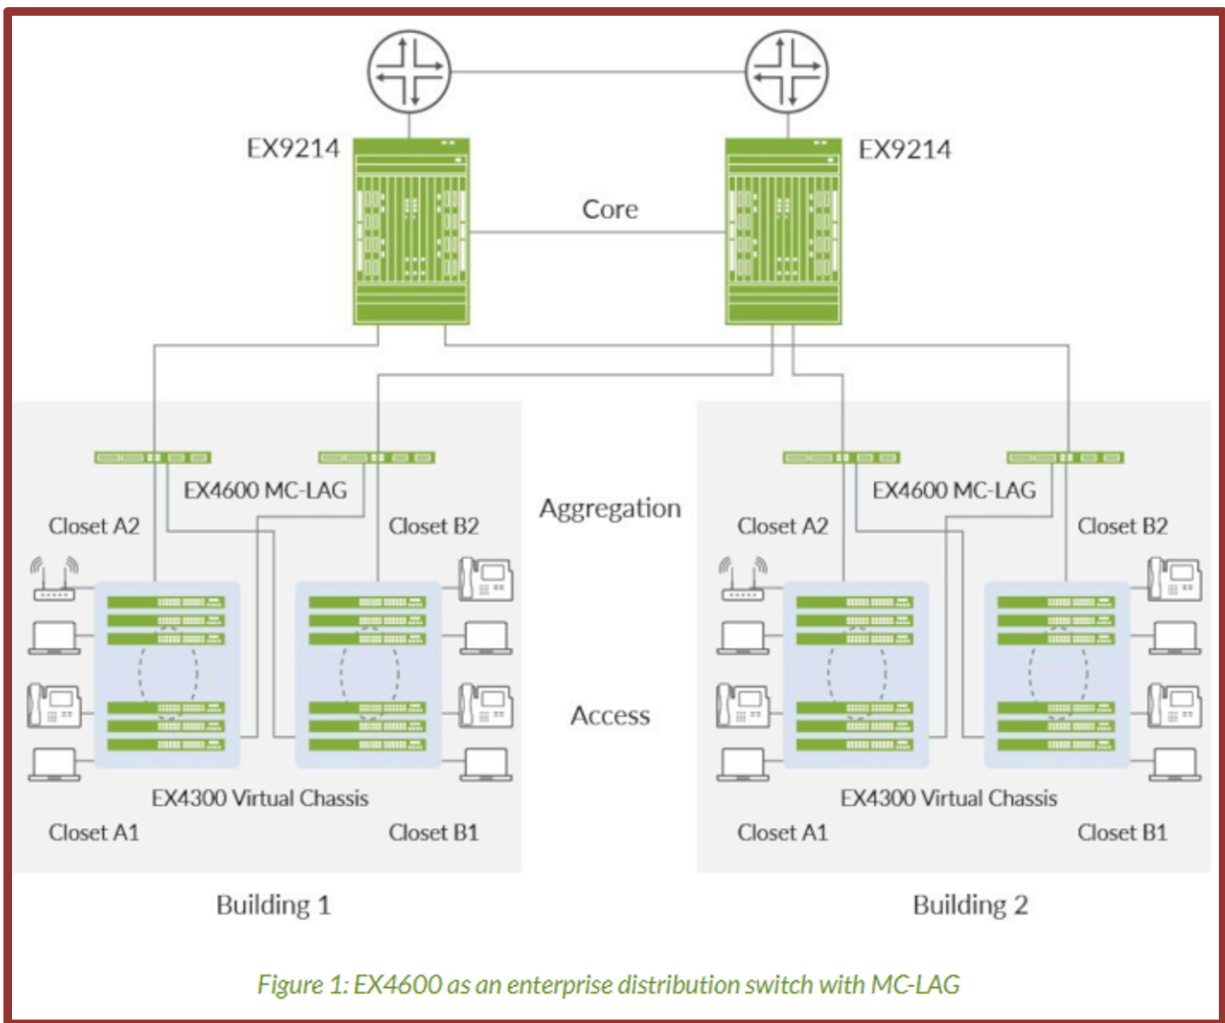
- Multichassis link aggregation (MC-LAG) - Layer 2, Layer 3, VRRP, STP
- Redundant trunk group (RTG)
- LAG load sharing algorithm—bridged or routed (unicast or multicast) traffic:
 - IP: SIP, Dynamic Internet Protocol (DIP), TCP/UDP source port, TCP/UDP destination port
 - Layer 2 and non-IP: MAC SA, MAC DA, Ethertype, VLAN ID, source port
 - FCoE packet: Source ID (SID), destination ID (DID), originator exchange ID (OXID), source port

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 8 of PDF)

Performance Scale (Unidimensional)

- MAC addresses per system: 288,000*
- VLAN IDs: 4,091
- Number of ports per LAG: 32
- FCoE scale:
 - Number of FCoE VLANs/FC virtual fabrics: 4,095
- Firewall filters: 4,000
- IPv4 unicast routes: 128,000 prefixes; 208,000 host routes; 64 ECMP paths (roadmap)
- IPv4 multicast routes: 104,000
- IPv6 multicast routes: 52,000
- IPv6 unicast routes: 64,000 prefixes
- Address Resolution Protocol (ARP) entries: 48,000
- Jumbo frame: 9,216 bytes

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 7 of PDF)



<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 2 of PDF)

Table 52: Maximum Interfaces per LAG and Maximum LAGs per Switch (continued)

Switch	Maximum Interfaces per LAG	Maximum LAGs
EX3400	16	128
EX4200 and EX4200 Virtual Chassis	8	111
EX4300 and EX4300 Virtual Chassis	16	128
EX4500, EX4500 Virtual Chassis, EX4550, and EX4550 Virtual Chassis	8	111
EX4600	32	128
EX6200	8	111

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Page 264 of PDF)

Multichassis link aggregation groups (MC-LAGs) enable a client device to form a logical LAG interface between two MC-LAG peers. An MC-LAG provides redundancy and load balancing between the two MC-LAG peers, multihoming support, and a loop-free Layer 2 network without running STP.


On one end of an MC-LAG, there is an MC-LAG client device, such as a server, that has one or more physical links in a link aggregation group (LAG). This client device uses the link as a LAG. On the other side of the MC-LAG, there can be a maximum of two MC-LAG peers. Each of the MC-LAG peers has one or more physical links connected to a single client device.

The MC-LAG peers use the Inter-Chassis Control Protocol (ICCP) to exchange control information and coordinate with each other to ensure that data traffic is forwarded properly.

The Link Aggregation Control Protocol (LACP) is a subcomponent of the IEEE 802.3ad standard. LACP is used to discover multiple links from a client device connected to an MC-LAG peer. LACP must be configured on both MC-LAG peers for an MC-LAG to work correctly.

NOTE: You must specify a service identifier (service-id) at the global level; otherwise, multichassis link aggregation will not work.

Figure 1: Basic MC-LAG Topology



https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/mc-lag/multichassis-link-aggregation-groups.pdf (Pages 20 and 21 of PDF)

98. Defendant's EX4600 Ethernet Switch comprises receiving data packets having respective destination addresses that specify forwarding the packets to groups of

multiple recipients through at least one of a plurality of ports, at least a subset of which is grouped in a link aggregation group (LAG).

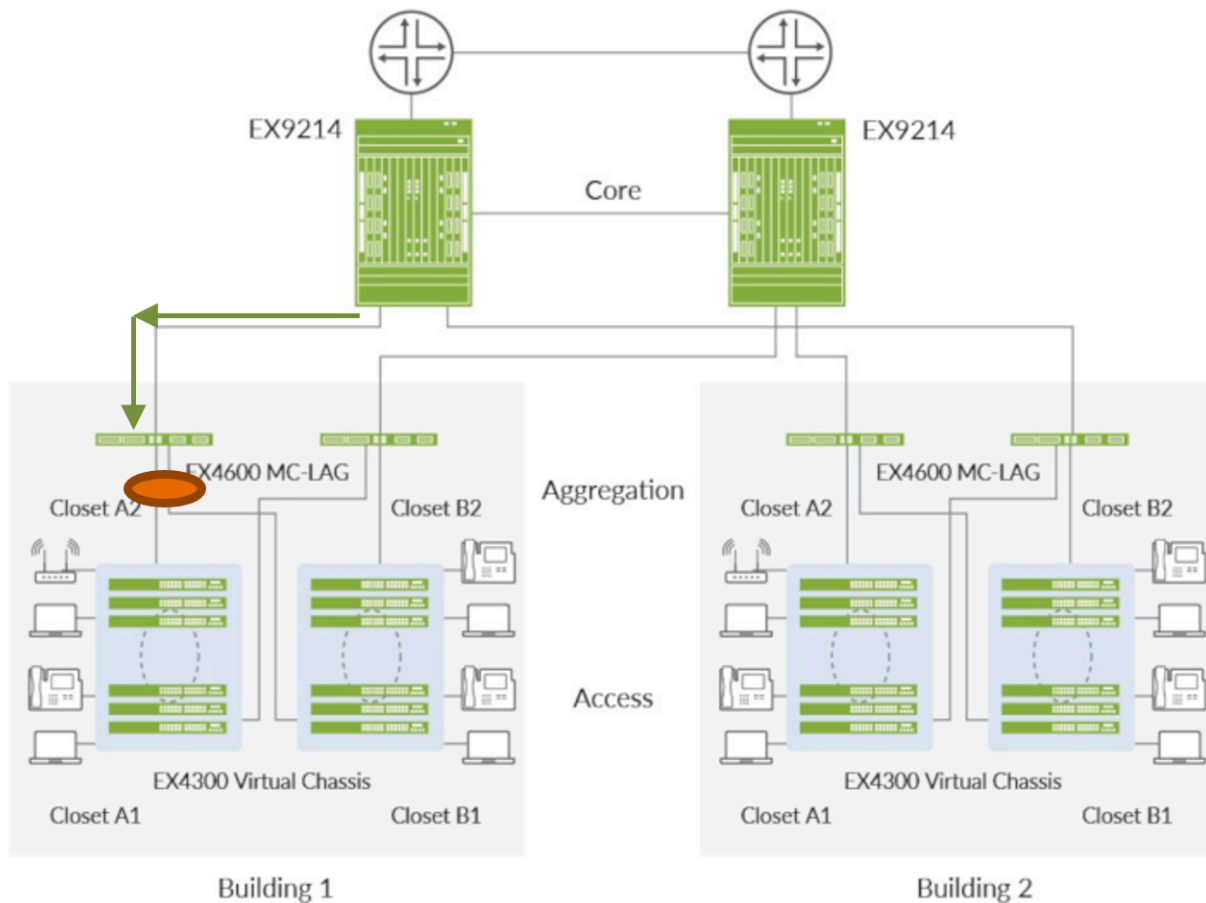


Figure 1: EX4600 as an enterprise distribution switch with MC-LAG

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 2 of PDF)

Enterprise Deployments

The EX4600 offers an economical, power-efficient, and compact solution for aggregating 10GbE expansions from access devices in building and enterprise deployments. The switch's dual-speed interfaces also support environments transitioning from 1GbE to 10GbE. The EX4600 can be deployed in the distribution layer with multichassis link aggregation (MC-LAG) (see Figure 1) to deliver higher resiliency with a distributed control plane, NSB, NSR, and unified ISSU. Multichassis LAG enables two EX4600 switches to act as separate devices with their own control planes, while eliminating STP by allowing link aggregation on the connected devices. In addition, unified ISSU allows each of the EX4600 switches to be upgraded individually without service interruption.

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 3 of PDF)

- MAC learning disable
- Link Aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)
- MAC notification
- MAC address aging configuration
- MAC address filtering
- Persistent MAC (sticky MAC)

Link Aggregation

- Multichassis link aggregation (MC-LAG) - Layer 2, Layer 3, VRRP, STP
- Redundant trunk group (RTG)
- LAG load sharing algorithm—bridged or routed (unicast or multicast) traffic:
- IP: SIP, Dynamic Internet Protocol (DIP), TCP/UDP source port, TCP/UDP destination port
- Layer 2 and non-IP: MAC SA, MAC DA, Ethertype, VLAN ID, source port
- FCoE packet: Source ID (SID), destination ID (DID), originator exchange ID (OXID), source port

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 8 of PDF)

Performance Scale (Unidimensional)

- MAC addresses per system: 288,000*
- VLAN IDs: 4,091
- Number of ports per LAG: 32
- FCoE scale:
 - Number of FCoE VLANs/FC virtual fabrics: 4,095
- Firewall filters: 4,000
- IPv4 unicast routes: 128,000 prefixes; 208,000 host routes; 64 ECMP paths (roadmap)
- IPv4 multicast routes: 104,000
- IPv6 multicast routes: 52,000
- IPv6 unicast routes: 64,000 prefixes
- Address Resolution Protocol (ARP) entries: 48,000
- Jumbo frame: 9,216 bytes

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 7 of PDF)

99. Defendant's EX4600 Ethernet Switch comprises processing the data packets so as to forward only a single copy of each of the data packets via the output ports in the subset while distributing forwarded copies of the data packets among the output ports in the subset so as to balance a traffic load within the LAG group.

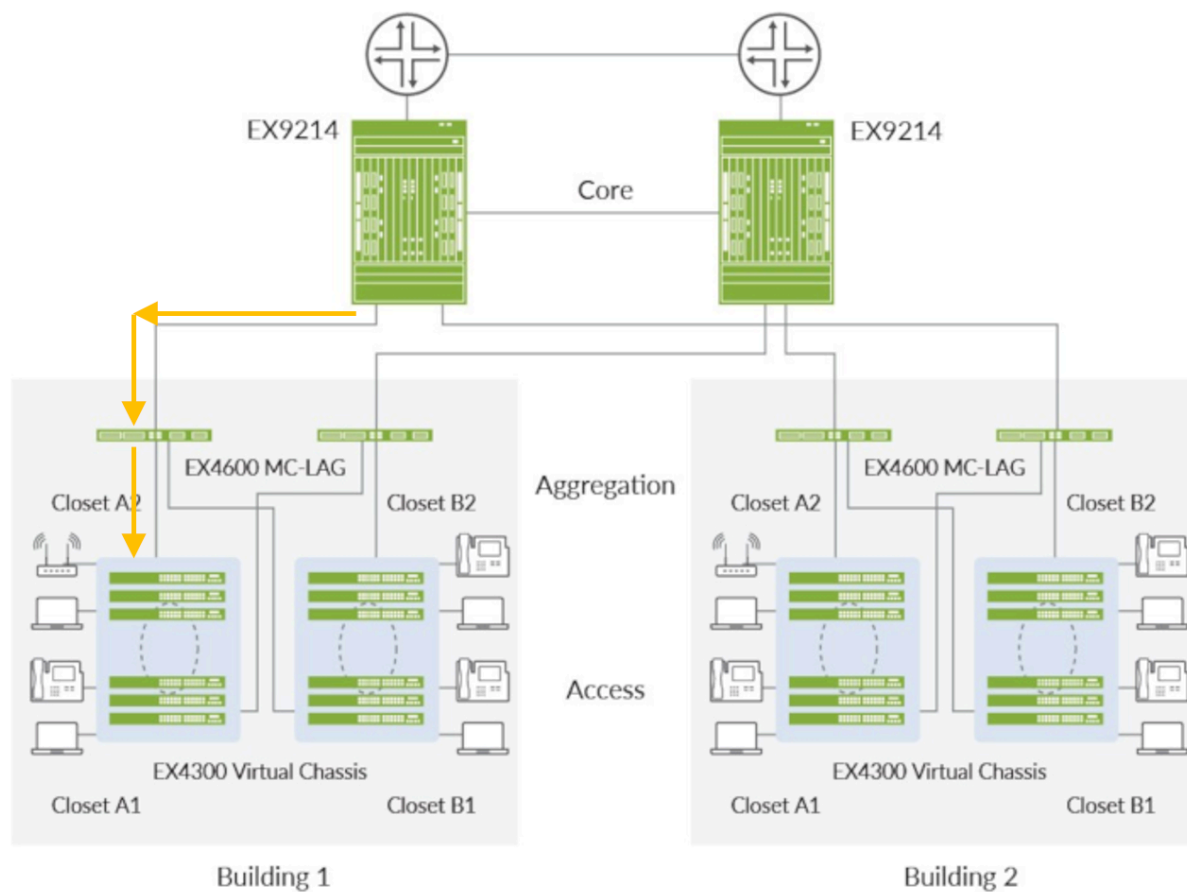


Figure 1: EX4600 as an enterprise distribution switch with MC-LAG

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 2 of PDF)

- MAC learning disable
- Link Aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)
- MAC notification
- MAC address aging configuration
- MAC address filtering
- Persistent MAC (sticky MAC)

Link Aggregation

- Multichassis link aggregation (MC-LAG) - Layer 2, Layer 3, VRRP, STP
- Redundant trunk group (RTG)
- LAG load sharing algorithm—bridged or routed (unicast or multicast) traffic:
- IP: SIP, Dynamic Internet Protocol (DIP), TCP/UDP source port, TCP/UDP destination port
- Layer 2 and non-IP: MAC SA, MAC DA, Ethertype, VLAN ID, source port
- FCoE packet: Source ID (SID), destination ID (DID), originator exchange ID (OXID), source port

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 8 of PDF)

When a packet is received on the ingress interface of a device, the packet forwarding engine (PFE) performs a look up to identify the forwarding next hop. If there are multiple equal-cost paths (ECMPs) to the same next-hop destination, the ingress PFE can be configured to distribute the flow between the next hops. Likewise, distribution of traffic may be required between the member links of an aggregated interface such as aggregated Ethernet. The selection of the actual forwarding next-hop is based on the hash computation result over select packet header fields and several internal fields such as **interface index**. You can configure some of the fields that are used by the hashing algorithm.

Junos supports different types of load balancing.

- *Per-prefix load balancing* –Each prefix is mapped to only one forwarding next-hop.
- *Per-packet load balancing*–All next-hop addresses for a destination in the active route are installed in the forwarding table (the term *per-packet* load balancing in Junos is equivalent to what other vendors may call *per-flow* load balancing). See "[Configuring Per-Packet Load Balancing](#)" on page 92 for more information.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-sfm/config-guide-sampling-forwarding-monitoring.pdf (Page 78 of PDF)

If you include both the **layer 3** and **layer 4** statements, the device uses the following Layer 3 and Layer 4 information to load-balance:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Incoming interface index
- IP type of service

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-sfm/config-guide-sampling-forwarding-monitoring.pdf (Page 93 of PDF)

Load Balancing

Load balancing of network traffic between MC-LAG peers is 100 percent local bias. Load balancing of network traffic between multiple LAG members in a local MC-LAG node is achieved through a standard LAG hashing algorithm.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/mc-lag/multichassis-link-aggregation-groups.pdf (Page 24 of PDF)

Load Balancing and Ethernet Link Aggregation Overview

You can create a link aggregation group (LAG) for a group of Ethernet ports. Layer 2 bridging traffic is load balanced across the member links of this group, making the configuration attractive for congestion concerns as well as for redundancy. You can configure up to 128 LAG bundles on M Series, and T Series routers, and 480 LAG bundles on MX Series routers and EX9200 switches. Each LAG bundle contains up to 16 links. (Platform support depends on the Junos OS release in your installation.)

By default, the hash key mechanism to load-balance frames across LAG interfaces is based on Layer 2 fields (such as frame source and destination address) as well as the input logical interface (unit). The default LAG algorithm is optimized for Layer 2 switching. Starting with Junos OS Release 10.1, you can also configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers using the **payload** statement. However, note that the load-balancing behavior is platform-specific and based on appropriate hash-key configurations.

For more information, see "Configuring Load Balancing on a LAG Link" on page 355. In a Layer 2 switch, one link is overutilized and other links are underutilized.

SEE ALSO

| [payload](#)

Configuring Load Balancing Based on MAC Addresses

The hash key mechanism for load-balancing uses Layer 2 media access control (MAC) information such as frame source and destination address. To load-balance traffic based on Layer 2 MAC information, include

https://www.juniper.net/documentation/en_US/junos/information-on-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Page 353 of PDF)

Configuring Load Balancing on a LAG Link

You can configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers inside the frame payload for load-balancing purposes using the **payload** statement. You can configure the statement to look at **layer-3** (and **source-ip-only** or **destination-ip-only** packet header fields) or **layer-4** fields. You configure this statement at the **[edit forwarding-options hash-key family multiservice]** hierarchy level.

You can configure Layer 3 or Layer 4 options, or both. The **source-ip-only** or **destination-ip-only** options are mutually exclusive. The **layer-3-only** statement is not available on MX Series routers.

By default, Junos implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet.

For more information about link aggregation group (LAG) configuration, see the *Junos OS Network Interfaces Library for Routing Devices*.

https://www.juniper.net/documentation/en_US/junos/information-on-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Page 355 of PDF)

How Does Multicast Load Balancing Work?

Juniper Networks Junos operating system (Junos OS) supports the Link Aggregation Control Protocol (LACP), which is a subcomponent of IEEE 802.3ad. LACP provides additional functionality for LAGs and is supported only on Layer 3 interfaces. When traffic can use multiple member links, traffic that is part of the same stream must always be on the same link.

Multicast load balancing uses one of seven available hashing algorithms and a technique called queue shuffling (alternating between two queues) to distribute and balance the data, directing streams over all available aggregated links. You can select one of the seven algorithms when you configure multicast load balancing, or you can use the default algorithm, `crc-sgip`, which uses a cyclic redundancy check (CRC) algorithm on the multicast packets' group IP address. We recommend that you start with the `crc-sgip` default and try other options if this algorithm does not evenly distribute the Layer 3 routed multicast traffic. Six of the algorithms are based on the hashed value of IP addresses (IPv4 or IPv6) and will produce the same result each time they are used. Only the balanced mode option produces results that vary depending on the order in which streams are added. See [Table 59 on page 363](#) for more information.

Table 59: Hashing Algorithms Used by Multicast Load Balancing

Hashing Algorithms	Based On	Best Use
<code>crc-sgip</code>	Cyclic redundancy check of multicast packets' source and group IP address	Default—high-performance management of IP traffic on 10-Gigabit Ethernet network. Predictable assignment to the same link each time. This mode is complex but yields a good distributed hash.
<code>crc-gip</code>	Cyclic redundancy check of multicast packets' group IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
<code>crc-sip</code>	Cyclic redundancy check of multicast packets' source IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the stream sources vary.
<code>simple-sgip</code>	XOR calculation of multicast packets' source and group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-sgip</code> yields. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic.

Table 59: Hashing Algorithms Used by Multicast Load Balancing (*continued*)

Hashing Algorithms	Based On	Best Use
<code>simple-gip</code>	XOR calculation of multicast packets' group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-gip</code> yields. Try this when <code>crc-gip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
<code>simple-sip</code>	XOR calculation of multicast packets' source IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-sip</code> yields. Try this mode when <code>crc-sip</code> does not evenly distribute the Layer 3 routed multicast traffic and stream sources vary.
<code>balanced</code>	Round-robin calculation method used to identify multicast links with the least amount of traffic	Best balance is achieved, but you cannot predict which link will be consistently used because that depends on the order in which streams come online. Use when consistent assignment is not needed after every reboot.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Pages 363 and 364 of PDF)

100. Defendant's EX4600 Ethernet Switch comprises allocating to each of the received data packets a fabric multicast identification (FMID) value selected from a range of possible FMID values, the FMID value being associated with one of the ports in the subset, and forwarding the single copy to the port associated with the allocated FMID value.

When a packet is received on the ingress interface of a device, the packet forwarding engine (PFE) performs a look up to identify the forwarding next hop. If there are multiple equal-cost paths (ECMPs) to the same next-hop destination, the ingress PFE can be configured to distribute the flow between the next hops. Likewise, distribution of traffic may be required between the member links of an aggregated interface such as aggregated Ethernet. The selection of the actual forwarding next-hop is based on the hash computation result over select packet header fields and several internal fields such as **interface index**. You can configure some of the fields that are used by the hashing algorithm.

Junos supports different types of load balancing.

- *Per-prefix load balancing* –Each prefix is mapped to only one forwarding next-hop.
- *Per-packet load balancing*–All next-hop addresses for a destination in the active route are installed in the forwarding table (the term *per-packet* load balancing in Junos is equivalent to what other vendors may call *per-flow* load balancing). See "[Configuring Per-Packet Load Balancing](#)" on page 92 for more information.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-sfm/config-guide-sampling-forwarding-monitoring.pdf (Page 78 of PDF)

If you include both the **layer 3** and **layer 4** statements, the device uses the following Layer 3 and Layer 4 information to load-balance:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Incoming interface index
- IP type of service

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-sfm/config-guide-sampling-forwarding-monitoring.pdf (Page 93 of PDF)

Load Balancing

Load balancing of network traffic between MC-LAG peers is 100 percent local bias. Load balancing of network traffic between multiple LAG members in a local MC-LAG node is achieved through a standard LAG hashing algorithm.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/mc-lag/multichassis-link-aggregation-groups.pdf (Page 24 of PDF)

Load Balancing and Ethernet Link Aggregation Overview

You can create a link aggregation group (LAG) for a group of Ethernet ports. Layer 2 bridging traffic is load balanced across the member links of this group, making the configuration attractive for congestion concerns as well as for redundancy. You can configure up to 128 LAG bundles on M Series, and T Series routers, and 480 LAG bundles on MX Series routers and EX9200 switches. Each LAG bundle contains up to 16 links. (Platform support depends on the Junos OS release in your installation.)

By default, the hash key mechanism to load-balance frames across LAG interfaces is based on Layer 2 fields (such as frame source and destination address) as well as the input logical interface (unit). The default LAG algorithm is optimized for Layer 2 switching. Starting with Junos OS Release 10.1, you can also configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers using the **payload** statement. However, note that the load-balancing behavior is platform-specific and based on appropriate hash-key configurations.

For more information, see "[Configuring Load Balancing on a LAG Link](#)" on page 355. In a Layer 2 switch, one link is overutilized and other links are underutilized.

SEE ALSO

| [payload](#)

Configuring Load Balancing Based on MAC Addresses

The hash key mechanism for load-balancing uses Layer 2 media access control (MAC) information such as frame source and destination address. To load-balance traffic based on Layer 2 MAC information, include

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Page 353 of PDF)

Configuring Load Balancing on a LAG Link

You can configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers inside the frame payload for load-balancing purposes using the **payload** statement. You can configure the statement to look at **layer-3** (and **source-ip-only** or **destination-ip-only** packet header fields) or **layer-4** fields. You configure this statement at the **[edit forwarding-options hash-key family multiservice]** hierarchy level.

You can configure Layer 3 or Layer 4 options, or both. The **source-ip-only** or **destination-ip-only** options are mutually exclusive. The **layer-3-only** statement is not available on MX Series routers.

By default, Junos implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet.

For more information about link aggregation group (LAG) configuration, see the *Junos OS Network Interfaces Library for Routing Devices*.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Page 355 of PDF)

How Does Multicast Load Balancing Work?

Juniper Networks Junos operating system (Junos OS) supports the Link Aggregation Control Protocol (LACP), which is a subcomponent of IEEE 802.3ad. LACP provides additional functionality for LAGs and is supported only on Layer 3 interfaces. When traffic can use multiple member links, traffic that is part of the same stream must always be on the same link.

Multicast load balancing uses one of seven available hashing algorithms and a technique called queue shuffling (alternating between two queues) to distribute and balance the data, directing streams over all available aggregated links. You can select one of the seven algorithms when you configure multicast load balancing, or you can use the default algorithm, `crc-sgip`, which uses a cyclic redundancy check (CRC) algorithm on the multicast packets' group IP address. We recommend that you start with the `crc-sgip` default and try other options if this algorithm does not evenly distribute the Layer 3 routed multicast traffic. Six of the algorithms are based on the hashed value of IP addresses (IPv4 or IPv6) and will produce the same result each time they are used. Only the balanced mode option produces results that vary depending on the order in which streams are added. See [Table 59 on page 363](#) for more information.

Table 59: Hashing Algorithms Used by Multicast Load Balancing

Hashing Algorithms	Based On	Best Use
<code>crc-sgip</code>	Cyclic redundancy check of multicast packets' source and group IP address	Default—high-performance management of IP traffic on 10-Gigabit Ethernet network. Predictable assignment to the same link each time. This mode is complex but yields a good distributed hash.
<code>crc-gip</code>	Cyclic redundancy check of multicast packets' group IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
<code>crc-sip</code>	Cyclic redundancy check of multicast packets' source IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the stream sources vary.
<code>simple-sgip</code>	XOR calculation of multicast packets' source and group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-sgip</code> yields. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic.

Table 59: Hashing Algorithms Used by Multicast Load Balancing (*continued*)

Hashing Algorithms	Based On	Best Use
<code>simple-gip</code>	XOR calculation of multicast packets' group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-gip</code> yields. Try this when <code>crc-gip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
<code>simple-sip</code>	XOR calculation of multicast packets' source IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-sip</code> yields. Try this mode when <code>crc-sip</code> does not evenly distribute the Layer 3 routed multicast traffic and stream sources vary.
<code>balanced</code>	Round-robin calculation method used to identify multicast links with the least amount of traffic	Best balance is achieved, but you cannot predict which link will be consistently used because that depends on the order in which streams come online. Use when consistent assignment is not needed after every reboot.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Pages 363 and 364 of PDF)

101. Defendant's EX4600 Ethernet Switch comprises for each of the received data packets, processing the data packets comprises allocating the FMID value to the data packet by a first line card connected to a first port via which the data packet is received, and configuring a second line card connected to a second port to which the data packet is to be sent and a switching fabric interconnecting the first and second line cards to forward the data packet responsively to the FMID value.

Key Features

Port Density	24 x 1/10GbE and 4 QSFP+ 40GbE plus expansion slots
Form Factor	1 RU
MACsec	400 Gbps of near-line encryption
Switch capacity	720 Gbps
Fabric	Virtual Chassis, MC-LAG

<https://www.juniper.net/us/en/products-services/switching/ex-series/ex4600/>

Interface Options

- 1GbE SFP: 24(40) (with 10GbE expansion modules)
- 10GbE SFP+: 24(40/72) (with 10GbE expansion modules/with fixed 40GbE ports using breakout cables)
- 40GbE QSFP+: 4(12) (with expansion modules)
 - Each fixed QSFP+ port can be configured as a 4x10GbE interface
 - Each QSFP+ port can be configured as a 40 Gbps port
 - USB port
 - Console port
 - 2 management ports: 1 RJ-45 and 1 SFP
 - Supported transceiver and direct attach cable
 - SFP+ 10GbE optical modules
 - SFP+ DAC cables: 1/3/5 m direct-attached copper and 1/3/5/7/10 m active direct-attached copper
 - SFP GbE optical and copper module
 - QSFP+ to SFP+ 10GbE direct attach break-out copper (1/3 m direct-attached copper cable)

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 7 of PDF)

EX4600 Hardware

The EX4600 switch is a compact 1U model that provides wire-speed packet performance, very low latency, and a rich set of Layer 2 and Layer 3 features. In addition to a high-throughput Packet Forwarding Engine, the performance of the control plane running on the EX4600 model is enhanced by the 1.5 -GHz dual-core Intel CPU with 8 GB of memory and 32 GB of solid-state drive (SSD) storage.

The port panel of the EX4600 features 24 fixed small form-factor pluggable (SFP) or SFP+ access ports and 4 fixed quad SFP+ (QSFP+) high-speed uplinks.


Figure 1: EX4600 Port Panel with Expansion Bays



In addition, the switch has two module bays where you can install optional expansion modules. The EX4600 switch supports two expansion modules to increase port density:


- QFX-EM-4Q—Adds four additional QSFP+ ports to the chassis. When fully populated with QFX-EM-4Q expansion modules, the EX4600 is equivalent to one with 72 interfaces (24 + 16 + 16 + 16). See [Figure 2 on page 20](#).

Figure 2: QFX-EM-4Q Expansion Module



- EX4600-EM-8F—Adds a total of eight additional SFP+ ports to the chassis. When fully populated with EX4600-EM-8F expansion modules, the EX4600 is equivalent to one with 56 interfaces (24 + 16 + 8 + 8). See [Figure 3 on page 21](#).

Figure 3: EX4600-EM-8F Expansion Module



https://www.juniper.net/documentation/en_US/release-independent/junos/information-products/pathway-pages/ex-series/ex4600/ex4600.pdf (Pages 20 and 21 of PDF)

When a packet is received on the ingress interface of a device, the packet forwarding engine (PFE) performs a look up to identify the forwarding next hop. If there are multiple equal-cost paths (ECMPs) to the same next-hop destination, the ingress PFE can be configured to distribute the flow between the next hops. Likewise, distribution of traffic may be required between the member links of an aggregated interface such as aggregated Ethernet. The selection of the actual forwarding next-hop is based on the hash computation result over select packet header fields and several internal fields such as **interface index**. You can configure some of the fields that are used by the hashing algorithm.

Junos supports different types of load balancing.

- *Per-prefix load balancing* –Each prefix is mapped to only one forwarding next-hop.
- *Per-packet load balancing*–All next-hop addresses for a destination in the active route are installed in the forwarding table (the term *per-packet* load balancing in Junos is equivalent to what other vendors may call *per-flow* load balancing). See "[Configuring Per-Packet Load Balancing](#)" on page 92 for more information.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-sfm/config-guide-sampling-forwarding-monitoring.pdf (Page 78 of PDF)

If you include both the **layer 3** and **layer 4** statements, the device uses the following Layer 3 and Layer 4 information to load-balance:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Incoming interface index
- IP type of service

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-sfm/config-guide-sampling-forwarding-monitoring.pdf (**Page 93 of PDF**)

Load Balancing

Load balancing of network traffic between MC-LAG peers is 100 percent local bias. Load balancing of network traffic between multiple LAG members in a local MC-LAG node is achieved through a standard LAG hashing algorithm.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/mc-lag/multichassis-link-aggregation-groups.pdf (**Page 24 of PDF**)

Load Balancing and Ethernet Link Aggregation Overview

You can create a link aggregation group (LAG) for a group of Ethernet ports. Layer 2 bridging traffic is load balanced across the member links of this group, making the configuration attractive for congestion concerns as well as for redundancy. You can configure up to 128 LAG bundles on M Series, and T Series routers, and 480 LAG bundles on MX Series routers and EX9200 switches. Each LAG bundle contains up to 16 links. (Platform support depends on the Junos OS release in your installation.)

By default, the hash key mechanism to load-balance frames across LAG interfaces is based on Layer 2 fields (such as frame source and destination address) as well as the input logical interface (unit). The default LAG algorithm is optimized for Layer 2 switching. Starting with Junos OS Release 10.1, you can also configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers using the **payload** statement. However, note that the load-balancing behavior is platform-specific and based on appropriate hash-key configurations.

For more information, see "[Configuring Load Balancing on a LAG Link](#)" on [page 355](#). In a Layer 2 switch, one link is overutilized and other links are underutilized.

SEE ALSO

| *payload*

Configuring Load Balancing Based on MAC Addresses

The hash key mechanism for load-balancing uses Layer 2 media access control (MAC) information such as frame source and destination address. To load-balance traffic based on Layer 2 MAC information, include

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (**Page 353 of PDF**)

Configuring Load Balancing on a LAG Link

You can configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers inside the frame payload for load-balancing purposes using the **payload** statement. You can configure the statement to look at **layer-3** (and **source-ip-only** or **destination-ip-only** packet header fields) or **layer-4** fields. You configure this statement at the **[edit forwarding-options hash-key family multiservice]** hierarchy level.

You can configure Layer 3 or Layer 4 options, or both. The **source-ip-only** or **destination-ip-only** options are mutually exclusive. The **layer-3-only** statement is not available on MX Series routers.

By default, Junos implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet.

For more information about link aggregation group (LAG) configuration, see the *Junos OS Network Interfaces Library for Routing Devices*.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Page 355 of PDF)

How Does Multicast Load Balancing Work?

Juniper Networks Junos operating system (Junos OS) supports the Link Aggregation Control Protocol (LACP), which is a subcomponent of IEEE 802.3ad. LACP provides additional functionality for LAGs and is supported only on Layer 3 interfaces. When traffic can use multiple member links, traffic that is part of the same stream must always be on the same link.

Multicast load balancing uses one of seven available hashing algorithms and a technique called queue shuffling (alternating between two queues) to distribute and balance the data, directing streams over all available aggregated links. You can select one of the seven algorithms when you configure multicast load balancing, or you can use the default algorithm, `crc-sgip`, which uses a cyclic redundancy check (CRC) algorithm on the multicast packets' group IP address. We recommend that you start with the `crc-sgip` default and try other options if this algorithm does not evenly distribute the Layer 3 routed multicast traffic. Six of the algorithms are based on the hashed value of IP addresses (IPv4 or IPv6) and will produce the same result each time they are used. Only the balanced mode option produces results that vary depending on the order in which streams are added. See [Table 59 on page 363](#) for more information.

Table 59: Hashing Algorithms Used by Multicast Load Balancing

Hashing Algorithms	Based On	Best Use
<code>crc-sgip</code>	Cyclic redundancy check of multicast packets' source and group IP address	Default—high-performance management of IP traffic on 10-Gigabit Ethernet network. Predictable assignment to the same link each time. This mode is complex but yields a good distributed hash.
<code>crc-gip</code>	Cyclic redundancy check of multicast packets' group IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
<code>crc-sip</code>	Cyclic redundancy check of multicast packets' source IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the stream sources vary.
<code>simple-sgip</code>	XOR calculation of multicast packets' source and group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-sgip</code> yields. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic.

Table 59: Hashing Algorithms Used by Multicast Load Balancing (*continued*)

Hashing Algorithms	Based On	Best Use
<code>simple-gip</code>	XOR calculation of multicast packets' group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-gip</code> yields. Try this when <code>crc-gip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
<code>simple-sip</code>	XOR calculation of multicast packets' source IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-sip</code> yields. Try this mode when <code>crc-sip</code> does not evenly distribute the Layer 3 routed multicast traffic and stream sources vary.
<code>balanced</code>	Round-robin calculation method used to identify multicast links with the least amount of traffic	Best balance is achieved, but you cannot predict which link will be consistently used because that depends on the order in which streams come online. Use when consistent assignment is not needed after every reboot.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Pages 363 and 364 of PDF)

102. Defendant's EX4600 Ethernet Switch comprises allocating the FMID value comprises assigning to the data packets line card FMID (LC-FMID values selected from a first range of possible LC-FMID values, and wherein configuring the switching fabric comprises mapping the LC-FMID values to respective central FMID (C-FMID) values selected from a second range of possible C-FMID values that is smaller than the first range and forwarding the data packets responsively to the C-FMID values.

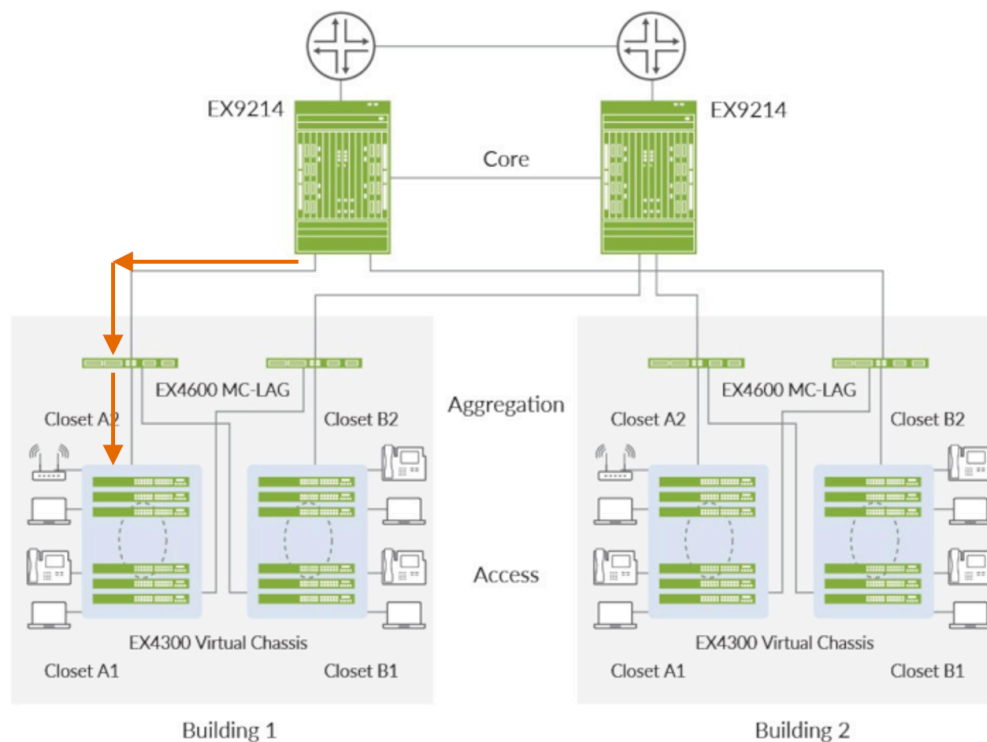


Figure 1: EX4600 as an enterprise distribution switch with MC-LAG

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf> (Page 2 of PDF)

When a packet is received on the ingress interface of a device, the packet forwarding engine (PFE) performs a look up to identify the forwarding next hop. If there are multiple equal-cost paths (ECMPs) to the same next-hop destination, the ingress PFE can be configured to distribute the flow between the next hops. Likewise, distribution of traffic may be required between the member links of an aggregated interface such as aggregated Ethernet. The selection of the actual forwarding next-hop is based on the hash computation result over select packet header fields and several internal fields such as **interface index**. You can configure some of the fields that are used by the hashing algorithm.

Junos supports different types of load balancing.

- *Per-prefix load balancing* – Each prefix is mapped to only one forwarding next-hop.
- *Per-packet load balancing* – All next-hop addresses for a destination in the active route are installed in the forwarding table (the term *per-packet* load balancing in Junos is equivalent to what other vendors may call *per-flow* load balancing). See "[Configuring Per-Packet Load Balancing](#)" on page 92 for more information.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-sfm/config-guide-sampling-forwarding-monitoring.pdf (Page 78 of PDF)

If you include both the **layer 3** and **layer 4** statements, the device uses the following Layer 3 and Layer 4 information to load-balance:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Incoming interface index
- IP type of service

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-sfm/config-guide-sampling-forwarding-monitoring.pdf (Page 93 of PDF)

Load Balancing

Load balancing of network traffic between MC-LAG peers is 100 percent local bias. Load balancing of network traffic between multiple LAG members in a local MC-LAG node is achieved through a standard LAG hashing algorithm.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/mc-lag/multichassis-link-aggregation-groups.pdf (Page 24 of PDF)

Load Balancing and Ethernet Link Aggregation Overview

You can create a link aggregation group (LAG) for a group of Ethernet ports. Layer 2 bridging traffic is load balanced across the member links of this group, making the configuration attractive for congestion concerns as well as for redundancy. You can configure up to 128 LAG bundles on M Series, and T Series routers, and 480 LAG bundles on MX Series routers and EX9200 switches. Each LAG bundle contains up to 16 links. (Platform support depends on the Junos OS release in your installation.)

By default, the hash key mechanism to load-balance frames across LAG interfaces is based on Layer 2 fields (such as frame source and destination address) as well as the input logical interface (unit). The default LAG algorithm is optimized for Layer 2 switching. Starting with Junos OS Release 10.1, you can also configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers using the **payload** statement. However, note that the load-balancing behavior is platform-specific and based on appropriate hash-key configurations.

For more information, see "[Configuring Load Balancing on a LAG Link](#)" on page 355. In a Layer 2 switch, one link is overutilized and other links are underutilized.

SEE ALSO

| [payload](#)

Configuring Load Balancing Based on MAC Addresses

The hash key mechanism for load-balancing uses Layer 2 media access control (MAC) information such as frame source and destination address. To load-balance traffic based on Layer 2 MAC information, include

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Page 353 of PDF)

Configuring Load Balancing on a LAG Link

You can configure the load balancing hash key for Layer 2 traffic to use fields in the Layer 3 and Layer 4 headers inside the frame payload for load-balancing purposes using the **payload** statement. You can configure the statement to look at **layer-3** (and **source-ip-only** or **destination-ip-only** packet header fields) or **layer-4** fields. You configure this statement at the [edit forwarding-options hash-key family multiservice] hierarchy level.

You can configure Layer 3 or Layer 4 options, or both. The **source-ip-only** or **destination-ip-only** options are mutually exclusive. The **layer-3-only** statement is not available on MX Series routers.

By default, Junos implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet.

For more information about link aggregation group (LAG) configuration, see the *Junos OS Network Interfaces Library for Routing Devices*.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Page 355 of PDF)

How Does Multicast Load Balancing Work?

Juniper Networks Junos operating system (Junos OS) supports the Link Aggregation Control Protocol (LACP), which is a subcomponent of IEEE 802.3ad. LACP provides additional functionality for LAGs and is supported only on Layer 3 interfaces. When traffic can use multiple member links, traffic that is part of the same stream must always be on the same link.

Multicast load balancing uses one of seven available hashing algorithms and a technique called queue shuffling (alternating between two queues) to distribute and balance the data, directing streams over all available aggregated links. You can select one of the seven algorithms when you configure multicast load balancing, or you can use the default algorithm, `crc-sgip`, which uses a cyclic redundancy check (CRC) algorithm on the multicast packets' group IP address. We recommend that you start with the `crc-sgip` default and try other options if this algorithm does not evenly distribute the Layer 3 routed multicast traffic. Six of the algorithms are based on the hashed value of IP addresses (IPv4 or IPv6) and will produce the same result each time they are used. Only the balanced mode option produces results that vary depending on the order in which streams are added. See [Table 59 on page 363](#) for more information.

Table 59: Hashing Algorithms Used by Multicast Load Balancing

Hashing Algorithms	Based On	Best Use
<code>crc-sgip</code>	Cyclic redundancy check of multicast packets' source and group IP address	Default—high-performance management of IP traffic on 10-Gigabit Ethernet network. Predictable assignment to the same link each time. This mode is complex but yields a good distributed hash.
<code>crc-gip</code>	Cyclic redundancy check of multicast packets' group IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
<code>crc-sip</code>	Cyclic redundancy check of multicast packets' source IP address	Predictable assignment to the same link each time. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic and the stream sources vary.
<code>simple-sgip</code>	XOR calculation of multicast packets' source and group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-sgip</code> yields. Try this mode when <code>crc-sgip</code> does not evenly distribute the Layer 3 routed multicast traffic.

Table 59: Hashing Algorithms Used by Multicast Load Balancing (*continued*)

Hashing Algorithms	Based On	Best Use
<code>simple-gip</code>	XOR calculation of multicast packets' group IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-gip</code> yields. Try this when <code>crc-gip</code> does not evenly distribute the Layer 3 routed multicast traffic and the group IP addresses vary.
<code>simple-sip</code>	XOR calculation of multicast packets' source IP address	Predictable assignment to the same link each time. This is a simple hashing method that might not yield as even a distribution as <code>crc-sip</code> yields. Try this mode when <code>crc-sip</code> does not evenly distribute the Layer 3 routed multicast traffic and stream sources vary.
<code>balanced</code>	Round-robin calculation method used to identify multicast links with the least amount of traffic	Best balance is achieved, but you cannot predict which link will be consistently used because that depends on the order in which streams come online. Use when consistent assignment is not needed after every reboot.

https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf (Pages 363 and 364 of PDF)

Willful Infringement

103. Defendant has had actual knowledge the '525 Patent and its infringement thereof at least as of receipt of the Notice Letter.

104. Defendant has had actual knowledge of the '525 Patent and its infringement thereof at least as of service or other receipt of Plaintiff's Complaint.

105. Defendant's risk of infringement of the Asserted Patents was either known or was so obvious that it should have been known to Defendant.

106. Notwithstanding this knowledge, Defendant has knowingly or with reckless disregard willfully infringed the '525 Patent. Defendant has thus had actual notice of the infringement of the '525 Patent and acted despite an objectively high likelihood that its actions constituted infringement of Plaintiff's valid patent rights, either literally or equivalently.

107. This objective risk was either known or so obvious that it should have been known to Defendant. Accordingly, Plaintiff seeks enhanced damages pursuant to 35 U.S.C. §§ 284 and 285.

Indirect Infringement

108. Defendant has induced and is knowingly inducing its customers and/or end users to directly infringe the '525 Patent, with the specific intent to encourage such infringement, and knowing that the induced acts constitute patent infringement, either literally or equivalently.

109. Defendant has knowingly contributed to direct infringement by its customers by having imported, sold, and/or offered for sale, and knowingly importing, selling, and/or offering to sell within the United States the '525 Accused Products which are not suitable for substantial non-infringing use and which are especially made or especially adapted for use by its customers in an infringement of the asserted patent.

110. Defendant's indirect infringement includes, for example, providing data sheets, technical guides, demonstrations, software and hardware specifications, installation guides, and other forms of support that induce its customers and/or end users to directly infringe the '525 Patent, for example:

- <https://www.juniper.net/us/en/products-services/switching/ex-series/ex4600/>
- <https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000511-en.pdf>
- https://www.juniper.net/documentation/en_US/release-independent/junos/information-products/pathway-pages/ex-series/ex4600/ex4600.pdf
- https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/mc-lag/multichassis-link-aggregation-groups.pdf
- https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/qfx-series/ethernet-interfaces-switches.pdf
- https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/config-guide-sfm/config-guide-sampling-forwarding-monitoring.pdf

111. Defendant's indirect infringement additionally includes marketing its products for import by its customers into the United States. Defendant's indirect infringement further includes providing application notes instructing its customers on infringing uses of the '525 Accused Products. The '525 Accused Products are

designed in such a way that when they are used for their intended purpose, the user infringes the '525 Patent, either literally or equivalently. Defendant knows and intends that customers who purchase the '525 Accused Products will use those products for their intended purpose. For example, Defendant's United States website: <https://www.juniper.net>, instructs customers to use the '525 Accused Products in numerous infringing applications. Defendant's customers directly infringe the '525 patent when they follow Defendant's provided instructions on website, videos, and elsewhere. Defendant's customers who follow Defendant's provided instructions directly infringe claims of the '525 Patent.

112. In addition, Defendant specifically intends that its customers, such as United States distributors, retailers and consumer product companies, will import, use, and sell infringing products in the United States to serve and develop the United States market for Defendant's infringing products. Defendant knows following its instructions directly infringes claims of the '525 Patent, including for example Claim 12.

113. Defendant's customers who follow Defendant's provided instructions directly infringe the method of claim 12 of the '525 Patent.

114. As a result of Defendant's infringement, Plaintiff has suffered monetary damages, and is entitled to an award of damages adequate to compensate it for such infringement which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT FIVE
INFRINGEMENT OF U.S. PATENT 7,961,755

115. Plaintiff incorporates by reference the allegations in preceding paragraphs 1-11 as if fully set forth herein.

116. The '755 Patent, entitled "Efficient transport of TDM services over packet networks" was filed on March 16, 2009 and issued on June 14, 2011.

117. Plaintiff is the assignee and owner of all rights, title and interest to the '755 Patent, including the right to recover for past infringements, and has the legal right to enforce the patent, sue for infringement, and seek equitable relief and damages.

Technical Description

118. The '755 Patent addresses problems in the prior art, such as "transport of TDM signals over a packet network, based on adaptive encapsulation of the TDM signal payloads. 4:40-42.

Direct Infringement

119. Defendant, without authorization or license from Plaintiff, has been and is directly infringing the '755 Patent, either literally or equivalently, as infringement is defined by 35 U.S.C. § 271, including through making, using (including for testing purposes), importing, selling and offering for sale telecommunications equipment that infringes one or more claims of the '755 Patent. Defendant develops, designs, manufactures, and distributes telecommunications equipment that infringes one or more claims of the '755 Patent. Defendant further provides services that practice methods that infringe one or more claims of the '755 Patent. Defendant is thus liable for direct infringement pursuant to 35 U.S.C. § 271. Exemplary infringing

instrumentalities include the M Series and MX Series Universal Routing Platform, and all other substantially similar products (collectively the “’755 Accused Products”).

120. Smart Path names this exemplary infringing instrumentality to serve as notice of Defendant’s infringing acts, but Smart Path reserves the right to name additional infringing products, known to or learned by Smart Path or revealed during discovery, and include them in the definition of ‘755 Accused Products.

121. Defendant is liable for direct infringement pursuant to 35 U.S.C. § 271 for the manufacture, sale, offer for sale, importation, use or distribution of Defendant’s M Series and MX Series Universal Routing.

122. Defendant’s M Series and MX Series Universal Routing is a non-limiting example of an apparatus that meets all limitations of claim 1 of the ’755 Patent, either literally or equivalently.

123. Defendant’s M Series and MX Series Universal Routers comprise a method for data communications.

MX SERIES UNIVERSAL ROUTING PLATFORMS

Product Description

The continuous expansion of mobile, video, and cloud-based services is disrupting traditional networks and impacting the businesses that rely on them. While annual double-digit traffic growth requires massive resource investments to prevent congestion and accommodate unpredictable traffic spikes, capturing return on that investment can be elusive. Emerging trends such as 5G mobility, Internet of Things (IoT) communications, and the continued growth of cloud networking promise even greater network challenges in the near future. The Juniper Networks® MX Series Universal Routing Platform delivers the industry's first end-to-end infrastructure security solution for enterprises as they look to move business-critical applications to public clouds. Delivering features, functionality, and secure services at scale in the 5G era with no compromises, the MX Series is a critical part of the network evolution happening now.

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000597-en.pdf>

124. Defendant's M Series and MX Series Universal Routers comprise receiving a time-division-multiplexed (TDM) input signal carrying a payload comprising data in at least one Synchronous Optical Network (SONET) frame having a plurality of virtual tributaries (VT's) including at least one unused VT.

Table 1. MX Series Universal Routing Platforms at a Glance

	MX960	MX480	MX240	MX204	MX150	MX104	MX80	MX40	MX10	MX5
Rack units	16	8	5	1	1	3.5	2	2	2	2
Systems per rack	3	6	9	48	48	13	24	24	24	24
Slots	11 MPCs	6 MPCs	2 MPCs	8 10GbE, 4 100GbE	NA	4 10GbE, 4 MIC slots	4 10GbE, 3 ⁺ MIC slots	2 10GbE, 3 ⁺ MIC slots	3 ⁺ MIC slots	3 ⁺ MIC slots
Per slot capacity	1.5 Tbps	1.5 Tbps	1.5 Tbps	NA	NA	NA	NA	NA	NA	NA
Maximum system throughput^a	12 Tbps	7.5 Tbps	3 Tbps	400 Gbps	20 Gbps	80 Gbps	80 Gbps	60 Gbps	40 Gbps	20 Gbps
PDH	Yes	Yes	Yes	NA	NA	Yes	Yes	Yes	Yes	NA
Sonet/SDH	Yes	Yes	Yes	NA	NA	Yes	Yes	Yes	Yes	NA

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000597-en.pdf>

- **Mobile Backhaul:** In addition to switching, routing, and security features, MX Series platforms support highly scalable and reliable hardware-based timing that meets the strictest LTE requirements, including Synchronous Ethernet for frequency and the Precision Time Protocol (PTP) for frequency and phase synchronization. In addition, the MX104 is ETSI 300-compliant, enabling deployment in next-generation mobile applications such as 5G.

<https://www.juniper.net/assets/us/en/local/pdf/datasheets/1000597-en.pdf>

125. Defendant's M Series and MX Series Universal Routers comprise determining whether the data comprise synchronous or non-synchronous data

1. Configure the `mpls` option to process incoming IP packets that have MPLS labels for passive monitoring on ATM and SONET/SDH interfaces and 10-Gigabit Ethernet interfaces in WAN PHY mode.

```
[edit interfaces so-fpc/pic/port sonet-options]
user@host# edit mpls
```

https://www.juniper.net/documentation/en_US/junos/topics/task/configuration/sonet-physical-options-edit-interfaces.html Page 6 of 13.

126. Defendant's M Series and MX Series Universal Routers comprise selecting a first encapsulation scheme if the data comprise synchronous data, and a second encapsulation scheme if the data comprise non-synchronous data, wherein the selected encapsulation scheme includes suppression of the at least one unused VT.

Understanding Interface Encapsulation on SONET/SDH Interfaces

📅 22-Mar-20

← →

Point-to-Point Protocol (PPP) encapsulation is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. Note that if you do not configure encapsulation, PPP is used by default.

For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface. You can optionally configure an encapsulation on a logical interface, which is the encapsulation used within certain packet types.


When you configure a point-to-point encapsulation (such as PPP or Cisco HDLC) on a physical interface, the physical interface can have only one logical interface (that is, only one `unit` statement) associated with it. When you configure a multipoint encapsulation (such as Frame Relay), the physical interface can have multiple logical units, and the units can be either point to point or multipoint. Use PPP if you are running Cisco IOS Release 12.0 or later. If you need to run Cisco HDLC, the Junos OS automatically configures an ISO family MTU of 4469 in the router. This is due to an extra byte of padding used by Cisco.

https://www.juniper.net/documentation/en_US/junos/topics/concept/sonet-interface-encapsulation-understanding.html

127. Defendant's M Series and MX Series Universal Routers comprise encapsulating the data for transmission over a packet-switched network in accordance with the selected encapsulation scheme.

MPLS tunnel cross-connects between interfaces and LSPs allow you to connect two distant interface circuits of the same type by creating MPLS tunnels that use LSPs as the conduit. The topology in [Figure 127 on page 1711](#) illustrates an MPLS LSP tunnel cross-connect. In this topology, two separate networks, in this case ATM access networks, are connected through an IP backbone. CCC allows you to establish an LSP tunnel between the two domains. With LSP tunneling, you tunnel the ATM traffic from one network across a SONET backbone to the second network by using an MPLS LSP.

Figure 127: MPLS Tunnel Cross-Connect



When traffic from Router A (VC 234) reaches Router B, it is encapsulated and placed into an LSP, which is sent through the backbone to Router C. At Router C, the label is removed, and the packets are placed onto the ATM permanent virtual circuit (PVC) (VC 591) and sent to Router D. Similarly, traffic from Router D (VC 591) is sent over an LSP to Router B, then placed on VC 234 to Router A.

Junos OS MPLS Applications User Guide, P. 1711.

Willful Infringement

128. Defendant has had actual knowledge of the '755 Patent and its infringement thereof at least as of receipt of the Notice Letter.

129. Defendant has had actual knowledge of the '755 Patent and its infringement thereof at least as of service or other receipt of Plaintiff's Complaint.

130. Defendant's infringement of the Asserted Patents was either known or was so obvious that it should have been known to Defendant.

131. Notwithstanding this knowledge, Defendant has knowingly or with reckless disregard infringed the '755 Patent. Defendant continued to commit acts of

infringement despite being on notice of an objectively high likelihood that its actions constituted infringement of Plaintiff's valid patent rights, either literally or equivalently.

132. Defendant is therefore liable for willful infringement. Accordingly, Plaintiff seeks enhanced damages pursuant to 35 U.S.C. §§ 284 and 285.

Indirect Infringement

133. Defendant has induced and is knowingly inducing its distributors, testers, trainers, customers and/or end users to directly infringe the '755 Patent, with the specific intent to induce acts constituting infringement, and knowing that the induced acts constitute patent infringement, either literally or equivalently.

134. Defendant has knowingly contributed to direct infringement by its customers and end users by having imported, sold, and/or offered for sale, and knowingly importing, selling, and/or offering to sell within the United States the accused products which are not suitable for substantial non-infringing use and which are especially made or especially adapted for use by its customers in an infringement of the asserted patent.

135. Defendant's indirect infringement includes, for example, providing data sheets, technical guides, demonstrations, software and hardware specifications, installation guides, and other forms of support that induce its customers and/or end users to directly infringe '755 Patent.

136. Defendant's indirect infringement additionally includes marketing its products for import by its customers into the United States. Defendant's indirect

infringement further includes providing application notes instructing its customers on infringing uses of the '755 Accused Products. The '755 Accused Products are designed in such a way that when they are used for their intended purpose, the user infringes the '755 Patent, either literally or equivalently. Defendant knows and intends that customers who purchase the '755 Accused Products will use those products for their intended purpose. For example, Defendant's United States website, <https://www.juniper.net>, instructs customers to use the '755 Accused Products in numerous infringing applications. Defendant's customers directly infringe the '755 patent when they follow Defendant's provided instructions on website, videos, and elsewhere. Defendant's customers who follow Defendant's provided instructions directly infringe claims of the '755 Patent.

137. In addition, Defendant specifically intends that its customers, such as United States distributors, retailers and consumer product companies, will import, use, and sell infringing products in the United States to serve and develop the United States market for Defendant's infringing products. Defendant knows following its instructions directly infringes claims of the '755 Patent, including for example Claim 1.

138. As a result of Defendant's infringement, Plaintiff has suffered monetary damages, and is entitled to an award of damages adequate to compensate it for such infringement which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

V. NOTICE

139. Smart Path has complied with the notice requirement of 35 U.S.C. § 287 and does not currently distribute, sell, offer for sale, or make products embodying the Asserted Patents. This notice requirement has been complied with by all relevant persons at all relevant times.

VI. JURY DEMAND

140. Plaintiff demands a trial by jury of all matters to which it is entitled to trial by jury, pursuant to FED. R. CIV. P. 38.

VII. PRAYER FOR RELIEF

WHEREFORE, Plaintiff prays for judgment and seeks relief against Defendant as follows:

- A. That the Court determine that one or more claims of each of the Asserted Patents is infringed by Defendant, both literally and under the doctrine of equivalents;
- B. That the Court determine that one or more claims of each of the Asserted Patents is indirectly infringed by Defendant;
- C. That the Court award damages adequate to compensate Plaintiff for the patent infringement that has occurred, together with prejudgment and post-judgment interest and costs, and an ongoing royalty for continued infringement;
- D. That the Court permanently enjoin Defendant pursuant to 35 U.S.C. § 283;
- E. That the Court find this case to be exceptional pursuant to 35 U.S.C. § 285;
- F. That the Court determine that Defendant's infringements were willful;
- G. That the Court award enhanced damages against Defendant pursuant to 35 U.S.C. § 284;
- H. That the Court award reasonable attorneys' fees; and

- I. That the Court award such other relief to Plaintiff as the Court deems just and proper.

Dated: May 26, 2021

Respectfully Submitted,

/s/ Bradley D. Liddle
E. Leon Carter
lcarter@carterarnett.com
Texas Bar No. 03914300
Bradley D. Liddle
bliddle@carterarnett.com
Texas Bar No. 24074599
Scott W. Breedlove
sbreedlove@carterarnett.com
State Bar No. 00790361
Joshua J. Bennett
jbennett@carterarnett.com
Texas Bar No. 24059444
Monica Litle
mlitle@carterarnett.com
Texas Bar No. 24102101
Nathan Cox
ncox@carterarnett.com
Texas Bar No. 24105751
Seth Lindner
slindner@carterarnett.com
Texas Bar No. 24078862
Michael C. Pomeroy
mpomeroy@carterarnett.com
Texas Bar No. 24098952
CARTER ARNETT PLLC
8150 N. Central Expy, 5th Floor
Dallas, Texas 75206
Telephone No. (214) 550-8188
Facsimile No. (214) 550-8185

ATTORNEYS FOR PLAINTIFF