

Technology Overview

Frequently Asked Questions: MPLS Connectivity

Release
12.1



Published: 2012-02-28

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Technology Overview Frequently Asked Questions: MPLS Connectivity

Release 12.1

Copyright © 2012, Juniper Networks, Inc.

All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

MPLS Connectivity FAQ Overview	1
Virtual Private LAN Service on MX Series Routers FAQs	1
MPLS Layer 3 VPN on MX Series, M Series, and T Series Routers FAQs	13
Layer 2 Circuits and Layer 2 VPNs on MX Series, M Series, and T Series Routers FAQs	17

MPLS Connectivity FAQ Overview

MPLS technology is evolving, with more services being offered using MPLS connectivity. Junos OS features are also evolving to implement these services on Juniper Networks routers.

The most common MPLS connectivity services include Virtual Private LAN Service (VPLS), Layer 3 virtual private networks (VPNs), Layer 2 circuits, and Layer 2 VPNs.

- VPLS provides a multipoint Ethernet service that emulates an Ethernet LAN. From the customer edge (CE) perspective, the service provider VPLS network operates like a private Ethernet broadcast domain.

A VPLS domain consists of a set of provider edge (PE) routers that acts as a single virtual Ethernet bridge for sites connected to those PE routers on the customer side. Pseudowire tunnels are created between those PE routers to aggregate traffic from one PE router to another. The PE routers exchange the MPLS labels used for the VPLS pseudowire, using either Label Distribution Protocol (LDP) forwarding equivalence classes (FECs), as described in RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*, or BGP, as described in RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*.

- Layer 3 VPN service is a point-to-point Layer 3 VPN network built over the service provider's MPLS transport network. This type of service is based on RFC 2547bis, *BGP/MPLS IP VPNs*.
- Layer 2 circuits and Layer 2 VPNs are point-to-point Layer 2 services built over the service provider's MPLS transport. This type of service establishes pseudowires using either LDP or BGP on the service provider core network.

This document presents the most frequently asked questions about these technologies and the features used to implement these services on Juniper Networks routers using Junos OS.

Related Documentation

- [Layer 2 Circuits and Layer 2 VPNs on MX Series, M Series, and T Series Routers FAQs on page 17](#)
- [MPLS Layer 3 VPN on MX Series, M Series, and T Series Routers FAQs on page 13](#)
- [Virtual Private LAN Service on MX Series Routers FAQs on page 1](#)

Virtual Private LAN Service on MX Series Routers FAQs

This section presents frequently asked questions and answers related to VPLS configurations on Juniper Networks MX Series routers.

What are the Juniper Networks solutions for Metro Ethernet Forum (MEF) 6.1 Ethernet services definitions, including Ethernet private line, Ethernet virtual private line, Ethernet LAN, Ethernet lines, and Ethernet trees?

- Ethernet private line (EPL) and Ethernet virtual private line (EVPL) are emulated by Juniper Networks Layer 2 circuit and Layer 2 VPN configurations.
- Ethernet LAN (E-LAN) is emulated by any Juniper Networks VPLS configuration solution.
- Ethernet lines (E-Line) service is emulated by using Juniper Networks Layer 2 circuit service configuration. The encapsulation type for the E-Line can be full Ethernet or only VLAN.
- Ethernet trees (E-Tree) service can be built using point-to-multipoint, or a more sophisticated implementation can be built using hub-and-spoke communities with BGP VPLS.

How can I prevent the receipt of hub routes on a directly-connected customer premises equipment spoke interface in a hub provider edge router? Can a core-facing statement be used for this purpose on MX Series routers?

To prevent spokes from talking to each other, spoke provider edge (PE) routers export different VPN routing and forwarding (VRF) targets, and import only the hub VRF targets. Using this method, spokes only exchange routes with the hub, however, the hub PE router imports all of the spoke VRF route targets. For example, assume that an interface residing on the hub PE router needs to connect with a spoke customer provider edge (CPE) router. When this spoke interface is added to the hub routing instance, it will receive hub routes, which is not expected in a hub-and-spoke topology.

A solution is to configure a core-facing statement under the spoke interface on the hub PE router. This will prevent route advertisements from the hub PE router from going to the directly-connected spoke CPE. Use the following configuration command to have broadcasts go only to the hub CE from the spoke CE attached to the hub PE router.


```
user@host# set interfaces ge-0/0/0.0 family vpls core-facing
```

How does qualified learning for virtual private LAN service operate? On which Juniper Networks platforms is it supported?


[Figure 1 on page 3](#) illustrates topology examples used to answer this question.

Figure 1: VPLS Learning Examples

VPLS Qualified Learning on MX Series Routers



VPLS Learning on M Series and T Series Routers



Traffic type and direction: unknown unicast from CE1 VLAN700 to CE2 VLAN700


In this example, first assume that Router PE2 is a Juniper Networks M Series Multiservice Edge Router. M Series routers do not process VLAN information, so the packets from Router CE1 are flooded out from Router PE2 to all ports and VLANs, not only to those connected to VLAN700. You can use the **monitor interface** and **show interface queue** commands to verify this behavior.

If, however, Router PE2 is a Juniper Networks MX Series 3D Universal Edge Router, the packets from Router CE1 for VLAN700 are flooded from Router PE2 only on the ports connected to VLAN700. This behavior is VPLS qualified learning. MX Series routers perform qualified lookups and qualified learning.

When one of the sites goes down in a single VPLS instance, why is there traffic loss on another site? How is a VPLS label block allocated when multiple sites are configured in a single VPLS instance?

Figure 2 on page 4 displays the topology used to illustrate this scenario. In it, a PE router has several site IDs configured within the same VPLS instance. The PE router advertises a separate label block for each of these site IDs to the remote PE routers. The label block associated with the lowest site ID number is selected to build pseudowires between the local PE router and the remote PE routers. You can use the **show vpls connections** command to verify that the pseudowire that has "Up" status between two PE routers is associated with the lowest site IDs on both PE routers.

Figure 2: VPLS Label Block Allocation



Router PE1 derives its labels from the label block associated with Site 2, which comes through Router PE2.

If Site 2 goes down and Site 3 remains operational, Router PE1 changes its label allocation to use the label block associated with Site 3. Consequently, a failure of Site 2 causes a loss of traffic between Site 1 and Site 3.

Additionally, traffic disruption can occur if Site 3 is provisioned initially, and Site 2 is provisioned at a later time. When Site 2 comes up, this causes a change in the labels used by Router PE1 for both Sites 2 and 3, causing traffic disruption for Site 3.

Solution: Router PE1 only programs its forwarding table to push labels that are advertised by a Site 2 label block. The forwarding table for Router PE2 expects traffic to come with the label advertised in a Site 2 label block. Neither Router PE1 nor Router PE2 has programmed their forwarding tables to use labels from a Site 3 label block, consequently, there will be traffic loss if Site 2 goes down. Traffic will resume after Router PE1 and Router PE2 program their forwarding tables to use labels from a Site 3 label block.

This is the intended behavior and is working as designed. For more information about VPLS label blocks, see *Technology Overview Understanding VPLS Label Blocks Operation*.

How can I control the behavior of VPLS class-based forwarding with multiple label-switched paths (LSPs) when one LSP goes down?

The following configuration is an example in which class-based forwarding is applied to a VPLS. It works as expected when both RSVP LSPs are up. If one of the LSPs goes down, all traffic uses the remaining LSP. For example, if the **to-brg2-private** LSP or tunnel goes down, all of the **PRIVATE** class traffic then uses the **REALTIME** tunnel.

```
next-hop-map NHOP-LSP-MAP {
  forwarding-class REALTIME {
    lsp-next-hop to-brg2-realtime;
  }
  forwarding-class PRIVATE {
    lsp-next-hop to-brg2-private;
  }
}
```

To change this behavior, apply a filter that only allows traffic of the specific forwarding class on each LSP:

```
set protocols mpls label-switched-path name policing filter name
```

What is the maximum number of VPLS instances supported on MX Series routers?

The maximum number of VPLS instances supported on MX Series routers is 8,000.

This number is derived as follows: the logical interface limit on an MX Series router is 64,000. In BGP VPLS, the default label block size is eight, which provides eight pseudowires per VPLS. Divide the number of logical interfaces by the number of pseudowires to get the maximum number of VPLS instances supported: $64,000/8 = 8,000$.

In Junos OS 10.0 and later, use the **label-block-size size** statement to configure VPLS label block size as 2, 4, 8, or 16. If the block size is increased, the number of VPLS instances can be increased. With LDP VPLS, there is no concept of label block, so the number of logical interfaces used is based on the number of sites attached to the VPLS.

Theoretically, if there are fewer than eight sites per LDP VPLS, that network could scale to 8,000 logical interfaces and greater. However, the qualified tested maximum number is 8,000.

What is the maximum number of sites supported per VPLS instance?

Currently, the flooding in VPLS is limited to 4,000 sites per VPLS mesh group. MX Series routers support 14 user-defined mesh groups. Therefore, the maximum number of sites possible is 56,000 ($4000 \times 14 = 56,000$).

How many Layer 2 circuits can be terminated into a single VRF instance?

On an MX Series router, 48 Layer 2 circuits can be terminated into a single VRF, per each tunnel PIC.

The limiting factor is the way MAC addresses are assigned to logical tunnel (**lt**) interface units. On an MX Series router there are 1984 MAC addresses on the EEPROM. This is statically divided by the number of slots (12) and PICs (4) to determine there are 48 per PIC. That number is the same on each platform although MAC space might be smaller.

Additional limitations are presented by the availability of bandwidth and the lack of resiliency because these solutions are bound to a tunnel PIC. If the dense port concentrator (DPC) associated with the tunnel goes down, all Layer 2 circuits are gone.

What are the Junos OS solutions for terminating Martini Layer 2 circuits into VPLS?

In Junos OS, Martini Layer 2 circuits can be terminated into VPLS by using logical tunnel (**lt**) interfaces. In Junos OS 9.2 and later, this can also be done without a tunnel PIC by using mesh groups.

What are the label-switched interface (LSI) label value assignments for different services?

VPLS uses a dynamic virtual tunnel logical interface on a tunnel PIC to model traffic from a remote PE router site in a VPLS domain. All traffic coming from the remote site is treated as if it is coming over the virtual port that represents this remote site, for the purposes of Ethernet flooding, forwarding, and learning.

In this approach, an MPLS lookup based on the inner VPN label is done on a PE router in the CF chip or R-chip. The label is stripped and the Layer 2 Ethernet frame it contained is forwarded to a tunnel PIC. The tunnel PIC loops the packet back, then a lookup is performed based on Ethernet MAC addresses.

Drawbacks to this approach include creating a bottleneck at the tunnel PIC and requiring the PE router to perform two lookups.

By default, VPLS uses a tunnel PIC. If the **no-tunnel-services** statement is configured under a VPLS instance, it uses an LSI and does not require a tunnel PIC.

The current label numbering assignment is listed in [Table 1 on page 6](#).

Table 1: Label Numbering Assignments

Services	Label Values
Reserved	0 - 15
LSI VPN	16 - 2047
LSP VPLS	2048 - 4095
Unassigned	4096 - 10,000
Static LSP	10,000-99,999
Global	100000-799999
Block Allocation	800000-899999
Per Intf	900000-999999

I have a non-MX Series router for VPLS. Is there an alternative to the MX-specific **no-local-switching** statement that I can use in this configuration?

Yes, as an alternative to the **no-local-switching** statement, you can configure the CE interfaces in single mesh groups on M Series or T Series routers.

The **no-local-switching** statement completely blocks the CE-to-CE communication in a multihomed VPLS instance. If the PE router has multiple CE interfaces, the CE routers are not able to switch traffic among themselves. This saves bandwidth to the CE link by not broadcasting unwanted flooded traffic.

Sample Configuration

user@host# show routing-instances

```
vpls-vlan100 {
  instance-type vpls;
  interface ge-2/0/1.0;
  interface ge-2/1/0.0;
  route-distinguisher 10.255.168.185:100;
  vrf-target target:100:1;
  protocols {
    vpls {
      site-range 10;
      no-tunnel-services;
      site vlan100-volt-PE {
        site-identifier 4;
      }
      mesh-group CE-MG {
        interface ge-2/0/1.0;
        interface ge-2/1/0.0;
      }
    }
  }
}
```

user@host# show vpls connections

Layer-2 VPN connections:

Legend for connection status (St)

EI -- encapsulation invalid NC -- interface encapsulation not CCC/TCC/VPLS

EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	- -- only outbound connection is up
CN -- circuit not provisioned	<- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision
LN -- local site not designated	LM -- local site ID not minimum designated
RN -- remote site not designated	RM -- remote site ID not minimum designated
XX -- unknown connection status	IL -- no incoming label
MM -- MTU mismatch	MI -- Mesh-Group ID not available
BK -- Backup connection	ST -- Standby connection
PF -- Profile parse failure	PB -- Profile busy

Legend for interface status

Up -- operational

Dn -- down

Instance: vpls-vlan100

Local site: vlan100-volt-PE (4)

connection-site	Type	St	Time last up	# Up trans
3	rmt	Up	Jun 17 08:11:59 2009	1

Remote PE: 10.255.171.30, Negotiated control-word: No

Incoming label: 262147, Outgoing label: 800003

Local interface: lsi.1048576, Status: Up, Encapsulation: VPLS

Description: Intf - vpls vpls-vlan100 local site 4 remote site 3

Use the **show vpls-flood extensive** command to verify that traffic from CE interfaces will only travel on pseudowires:

```
user@host# show vpls flood extensive
Name: vpls-vlan100
CEs: 2
VEs: 1
  Flood route prefix: 0x4a/32
  Flood route type: IFF_FLOOD
  Flood route owner: lsi.1048576
  Flood group name: __ves__
  Flood group index: 0
  Nexthop type: comp
  Nexthop index: 568
  Flooding to:
    Name      Type      NhType      Index
    CE-MG     Group     comp        565
    Composition: flood-to-all
    Flooding to:
      Name      Type      NhType      Index
      ge-2/0/1.0 CE       ucst        515
      ge-2/1/0.0 CE       ucst        571

  Flood route prefix: 0x46/32
  Flood route type: IFF_FLOOD
  Flood route owner: ge-2/0/1.0
  Flood group name: CE-MG
  Flood group index: 2
  Nexthop type: comp
  Nexthop index: 556
  Flooding to:
    Name      Type      NhType      Index
    __ves__   Group     comp        549
    Composition: flood-to-all
    Flooding to:
      Name      Type      NhType      Index
      lsi.1048576 VE       indr        1048574

  Flood route prefix: 0x4b/32
  Flood route type: IFF_FLOOD
  Flood route owner: ge-2/1/0.0
  Flood group name: CE-MG
  Flood group index: 2
  Nexthop type: comp
  Nexthop index: 556
  Flooding to:
    Name      Type      NhType      Index
    __ves__   Group     comp        549
    Composition: flood-to-all
    Flooding to:
      Name      Type      NhType      Index
      lsi.1048576 VE       indr        1048574
```

Is there an example configuration for terminating Layer 2 circuits into VPLS using mesh groups?

The following example terminates Layer 2 circuits into VPLS using mesh groups. In the example, the **mesh-group mx-pw** statement is the mesh group configuration:

```
user@host# show configuration routing-instances vpls1
```

```

instance-type vpls;
vlan-id 200;
interface ge-0/0/3.1;
route-distinguisher 1.1.1.121:1;
vrf-target target:64577:1;
protocols {
vpls {
no-tunnel-services;
site site3 {
site-identifier 3;
interface ge-0/0/3.1;
}
site site-mx-pw {
site-identifier 11;
mesh-group mx-pw;
}
vpls-id 123;
no-vlan-id-validate;
mesh-group mx-pw {
neighbor 69.158.196.218 {
no-vlan-id-validate;
ignore-encapsulation-mismatch;
}
}
}
}
}

```

Which components are used in calculating the load balancing (hashing) algorithm for a link aggregation group?

For Layer 3 traffic, the default components used to calculate the load balancing (hashing) algorithm for a link aggregation group (LAG) are the source address, destination address, and interface indexes. If you include the **hash-key** statement under the **[edit forwarding-options hash-key]** hierarchy level, the information in that level is used to compute the hash key for load balancing.

What is the maximum supported number of link aggregation group bundles and members per bundle on MX Series routers?

Junos OS supports a maximum of 128 LAG bundles with 16 members each on MX Series routers.

What is the interworking scalability for mesh groups in VPLS LDP-BGP?

You can have up to 14 mesh groups per VPLS on MX Series routers and up to 126 on M Series routers. For example, for a given VPLS instance, if you need LDP-BGP interworking, there must be one mesh group for every LDP-VPLS cloud that is connected to the BGP-VPLS cloud, at the interworking router.

Is the label block size configurable per VPLS and per router or logical system?

In Junos OS 10.0 and later, use the **label-block-size size** statement to configure VPLS label block size. Use this to increase the VPLS scaling. If the label block is configured as

two, this will increase the VPLS scaling four times. You can allocate the label block size in increments of 2, 4, 8, or 16.

Is indirect next hop supported for VPLS in MX Series routers with I-chip-based DPCs?

Prior to Junos OS Release 10.3, indirect next hop is only supported for Layer 3 VPNs. However, in Junos OS Release 10.3 and later, a constraint of not allowing VPLS configuration when using indirect next hop for other services is removed. With this constraint removed, you can configure VPLS with indirect next hops and the routing protocol process (RPD) will ignore the VPLS configuration.

I want to define multiple port mirroring interfaces on a single chassis and select which traffic is mirrored to each interface. How can I accomplish this on MX Series, M Series, and T Series routers?

Junos OS supports multiple port mirror destinations in MX Series routers, and on M Series M120 and M320 routers. For more information about configuring port-mirroring on MX Series routers, refer to Layer 2 Port Mirroring in the current version of the Junos OS[®] Software MX Series Ethernet Services Routers Layer 2 Configuration Guide, http://www.juniper.net/techpubs/en_US/junos10.1/information-products/pathway-pages/layer-2/layer-2-port-mirroring.html

In T Series routers, there is only one chassis-wide destination. All platforms support multiple source ports (interfaces that need to be mirrored).

Can traffic be port-mirrored from a Layer 2 VPN on T Series routers for PE routers?

No, this is not supported in T Series routers at this time. It is only supported in MX Series routers, and in M Series routers M120 and M320 with E3 FPCs (I-chip-based).

Can the outgoing and incoming traffic of a given T Series or MX Series interface be port-mirrored?

Yes, this is supported on all Juniper Networks platforms.

Can integrated routing and bridging traffic be port-mirrored?

In Junos OS 9.6R2 and later, if the following conditions are met, an integrated routing and bridging (IRB) packet can be mirrored as a Layer 2 packet:

1. The IRB is associated with the **bridge-domain** or **vpls routing-instance**.
2. The **bridge-domain** has a forwarding table filter configured with an action of **then port-mirror** or **then port-mirror-instance instance**.

This Layer 2 IRB port mirroring can be disabled using the **no-irb-layer-2-copy** statement at the **bridge-domain** or the **vpls routing-instance** hierarchy level.

What are the differences in packet processing between a VT interface and an LSI interface (vrf-table-label or no-tunnel-services)?

The LSI interface provides better packet processing performance than the VT interface, unless there are core-facing interface restrictions or loss of ingress forwarding functionality, because the frame is sent only once through the route lookup.

In a VT interface, the packet loops back, with a first pass as **mpls-vrf** and a second pass for frame processing. The VT interface is limited by an overall tunnel bandwidth of 1/10 Gbps. The LSI interface is limited by line rate.

What is an example configuration for using an entire site as primary and backup for VPLS multihoming?

The following example shows the **VPLS_CUST_101** routing instance configured as the primary site and the **VPLS_CUST_102** routing instance configured as the backup site. Note that in the portion of the example that shows the other PE router, the **VPLS_CUST_101** routing instance is the backup site and the **VPLS_CUST_102** routing instance is the primary site.

```
user@host# show routing-instances
```

```
VPLS_CUST_101 {
  instance-type vpls;
  vlan-id 101;
  interface ge-11/0/0.101;
  route-distinguisher 3.3.3.3:101;
  vrf-target target:100:101;
  protocols {
    vpls {
      site-range 10;
      site 3 {
        site-identifier 3;
        multi-homing;
        site-preference primary;
      }
    }
  }
}
VPLS_CUST_102 {
  instance-type vpls;
  vlan-id 102;
  interface ge-11/0/0.102;
  route-distinguisher 3.3.3.3:102;
  vrf-target target:100:102;
  protocols {
    vpls {
      site-range 10;
      site 3 {
        site-identifier 3;
        multi-homing;
        site-preference backup;
      }
    }
  }
}
```

```
user@host# show vpls connections
```

```
...
```

```
Instance: VPLS_CUST_101
Local site: 3 (3)
```

```

connection-site      Type St      Time last up      # Up trans
1                    rmt  Up      Feb 18 09:16:06 2009      1
  Remote PE: 1.1.1.1, Negotiated control-word: No
  Incoming label: 800256, Outgoing label: 800274
  Local interface: vt-11/3/10.1051392, Status: Up, Encapsulation: VPLS
  Description: Intf - vpls VPLS_CUST_101 local site 3 remote site 1
3                    rmt  SC      site-collision

Instance: VPLS_CUST_102
Local site: 3 (3)
connection-site      Type St      Time last up      # Up trans
1                    rmt  LN
3                    rmt  SC      site collision

user@host# show vpls connections
...

Instance: VPLS_CUST_101
Local site: 4 (3)
connection-site      Type St      Time last up      # Up trans
1                    rmt  LN
3                    rmt  SC

Instance: VPLS_CUST_102
Local site: 4 (3)
connection-site      Type St      Time last up      # Up trans
1                    rmt  Up      Feb 18 19:00:28 2009      1
  Local interface: vt-11/3/10.1048579, Status: Up, Encapsulation: VPLS
  Description: Intf - vpls VPLS_CUST_102 local site 3 remote site 1
  Remote PE: 1.1.1.1, Negotiated control-word: No
  Incoming label: 800016, Outgoing label: 800266
3                    rmt  SC

```

Which classification methods are supported for ingress queuing on an Enhanced Queuing Dense Port Concentrator on MX Series routers?

Support for ingress queuing on EQ-DPC includes the following methods, as defined in Institute of Electrical and Electronics Engineers (IEEE) 802.1p:

- IP Differentiated Services code point (DSCP) precedence for IPv4 interfaces
- MPLS EXP for MPLS interfaces
- Tagged VPLS, bridge, and circuit cross-connect (CCC) interfaces

There is no support for ingress queuing for untagged VPLS and CCC interfaces.

DSCP is not supported as a classification option for ingress queuing on VPLS interfaces.

Is BGP autodiscovery supported for LDP-based VPLS as defined in draft-ietf-l2vpn-signaling-xx?

Yes. Junos OS Release 11.2 and later supports BGP autodiscovery for LDP VPLS using forwarding equivalence class (FEC) 129.

As a workaround for older releases, you can manually provision each LDP PE router in an LDP mesh group by including the **neighbor** statement.

Is IGMP snooping supported on NG-VPLS when point-to-multipoint is in the core?

No, this is not supported.

**Related
Documentation**

- [Layer 2 Circuits and Layer 2 VPNs on MX Series, M Series, and T Series Routers FAQs on page 17](#)
- [MPLS Connectivity FAQ Overview on page 1](#)
- [MPLS Layer 3 VPN on MX Series, M Series, and T Series Routers FAQs on page 13](#)

MPLS Layer 3 VPN on MX Series, M Series, and T Series Routers FAQs

This section presents frequently asked questions and answers related to MPLS Layer 3 VPNs on Juniper MX Series, M Series, and T Series routers.

Can I manually select the upstream multicast hop in NG MPVN instead of having it default to the highest IP address?

When a multicast source connects to two PE routers, the receiver PE router selects the upstream multicast PE router. By default, it selects the router whose upstream address is numerically highest.

In Junos OS 9.6 and later, it is possible to change the upstream PE router selection behavior, for example, if the primary sender PE router has a lower address than the backup sender PE router.

You can select the upstream multicast hop (UMH) based on the unicast route preference, using the **unicast-umh-election** statement as shown:

```
set routing-instances routing-instance name protocols mvpn unicast-umh-election
```

Cautions on the use of the **unicast-umh-election** statement include:

- The unicast preference-based single forwarder election only works with the default deterministic path-selection algorithm in BGP. When a PE router is configured with the **cisco-non-deterministic** path selection algorithm in BGP, the unicast preference-based single forwarder election might fail. The **unicast-umh-election** statement can only be configured when **cisco-non-deterministic** path selection algorithm is not configured.
- All PE routers must prefer the same route to the upstream multicast hop. The unicast preferences for routes to the sources must not depend on any BGP path selection criteria (such as lowest IGP metric) that will cause one PE router to choose one UMH while another PE router chooses a different UMH.

Is four-label push/pop for MPLS supported on Juniper Networks devices?

Although there is no inherent limit to the total number of labels that a given packet can support, the number of labels that a router can push/pop is limited.


Four-label push/pop is not supported for MPLS for Juniper Networks devices. Support for four-label push/pop would only be needed for VPN PE routers that could be the ingress for an RSVP LSP as well as for an LDP LSP that tunnels over the RSVP, and then egresses at a router beyond the endpoint of the RSVP LSP.

For Juniper Networks devices, the VPN PE routers only need to push two labels (VPN and LDP) and the core routers only need to swap and then push up to two labels (swap LDP, push RSVP, and bypass). This provides LDP tunneling in RSVP LSPs, along with link protection.

When NG-MVPN with Protocol Independent Multicast (PIM) any-source multicast or generic routing encapsulation tunneling is used, traffic will not switch over from shared tree to shortest-path tree. Although a register packet is sent to the customer rendezvous point initially, multicast traffic is sent to the receiver through the shortest-path tree from the root. Is this the correct behavior?

Yes, this is the expected behavior of an NG-MVPN implementation, see [Figure 3 on page 14](#). This is an optimization designed to minimize state in the provider core. A drawback to this optimization is that it requires the customer rendezvous point (C-RP) to either be placed in the VRF of the PE router, or it requires a Multicast Source Discovery Protocol (MSDP) or PIM anycast-PIM connection between the C-RP and the PE router. This is because at least one PE router must know about all active sources in the VPN. This behavior is detailed in *draft-ietf-l3vpn-2547bis-mcast-bgp-07: Multicast in MPLS/BGP IP VPNs*.

Figure 3: NG-MVPN Shortest-Path Tree and Shared Tree Flow



In Junos OS 10.0 and later, there is an NG-MVPN RPT-SPT mode that implements Section 13 of the *BGP-MVPN* draft. This mode allows shared rendezvous-point trees to be signaled across the NG-MVPN core. This allows you to place the C-RP anywhere, with the cost of slightly more state in the core.

To configure the RPT-SPT mode, include the `rpt-spt` statement at the `[edit routing-instances routing-instance-name protocols mvpn mvpn-mode]` hierarchy level for all VRFs that make up the VPN. To configure a selective provider tunnel for the shared tree, include the `wildcard-group-inet`, `wildcard-group-inet6`, and `wildcard-source` statements at the `[edit routing-instances routing-instance-name provider-tunnel selective]` hierarchy level.



.....

CAUTION: When you configure RPT-SPT mode, receivers or sources directly attached to the PE router are not supported. As a workaround, place a CE router between any receiver or source and the PE router.

.....

Can VPN.inet.2 routes be used for reverse path forwarding (RPF) checks in NG-MVPN?

Yes. PIM can be configured to use VPN.inet.2 for RPF on NG-MVPN receiver site VRFs.

Is the smart-optimize-timer statement available for point-to-multipoint LSPs? Is the optimize-timer statement available?

No, the **smart-optimize-timer** statement isn't supported for point-to-multipoint LSPs. Only the **optimize-timer** statement is available for point-to-multipoint LSPs.

Is the adaptive statement supported in point-to-multipoint LSPs?

The **adaptive** statement is not supported in point-to-multipoint LSPs, because the point-to-multipoint LSP is expected to be adaptive by default.

How is composite next hop used in Layer 3 VPNs?

Composite next hop infrastructure provides optimized data structures to handle a **label-per-prefix** case and to help with convergence. Prior to Junos OS version 9.5, handling **label-per-prefix** consumed a lot of kernel memory, restricting overall chassis scale. Now, the standard Junos OS method is to use **label-per-VPN**, which is very scalable.

With chained or composite next hop, memory usage is optimized in both the kernel and the Packet Forwarding Engine. However, on the I-chip this is available only for the Packet Forwarding Engine. The scaling numbers for I-chip based platforms are:

- Junos OS versions prior to 9.5 -- 250,000 to 300,000 labels per prefix
- Junos OS version 9.5 and later -- 600,000 labels per prefix

Data structure optimization within the Packet Forwarding Engine leads to significant savings in DRAM, with the most gains in the next-hop space and in the Layer 2 descriptors.

Can multicast ping be used with a static IGMP join in Rosen MVPN?

When static IGMP is used to join any group, multicast ping does not get any response because static IGMP is not a real IGMP host. Consequently, the ping fails. If Session Announcement Protocol (SAP) is used instead of static IGMP, multicast ping can be used if interface and bypass routing is specified. To do this, include the **sap-listen** statement at the **[edit protocols]** hierarchy level on the CE router on the main configuration (not under the **routing-instance** hierarchy):

```
set protocols sap-listen 239.10.10.14
```

Is NG-MVPN supported in logical systems?

No, NG-MVPN is not supported in logical systems at this time.

Is fast reroute supported in point-to-multipoint LSP?

Fast reroute provides a mechanism for automatically rerouting traffic on an LSP if a node or link in an LSP fails, thus reducing the loss of packets traveling over the LSP. For point-to-multipoint LSPs with fast reroute, only link protection is supported; node protection is not supported.

Supported functionality for link protection on point-to-multipoint includes:

- Link protection can be provided for all constituent sub-LSPs of a point-to-multipoint LSP in the event of a link failure.
- The protection LSP is a point-to-multipoint bypass LSP, set up using existing point-to-point bypass LSP (facility backup) mechanisms. The path of the bypass point-to-point LSP can be determined using Constrained Shortest Path First (CSPF) or it can be configured.

Not supported functionality for link protection on point-to-multipoint includes:

- Global repair is not supported: If a link fails and local repair is performed using link protection, the ingress does not attempt to perform global repair. Because there is no secondary LSP or CSPF, traffic remains on the bypass LSP until the link comes back up.
- Regular fast reroute one-to-one detours are not supported because node protection is not supported; instead, link protection is provided by using bypass LSPs.

**Related
Documentation**

- [Layer 2 Circuits and Layer 2 VPNs on MX Series, M Series, and T Series Routers FAQs on page 17](#)
- [MPLS Connectivity FAQ Overview on page 1](#)
- [Virtual Private LAN Service on MX Series Routers FAQs on page 1](#)

Layer 2 Circuits and Layer 2 VPNs on MX Series, M Series, and T Series Routers FAQs

This section presents frequently asked questions and answers related to Layer 2 circuits and Layer 2 VPNs on Juniper MX Series, M Series, and T Series routers.

What is the solution for interprovider Option B connection between LDP-based Layer 2 circuits and BGP-based Layer 2 VPNs?

As of Junos OS Release 9.3 and later, Layer 2 circuits with interworking **iw0** interfaces can be manually stitched. If dynamic signaling is needed, use FEC 129 multisegment pseudowires if you are not using BGP-based Layer 2 VPNs. In Junos OS Release 11.2 and later, FEC 129 is supported.

What are the supported Layer 2 interworking stitching configurations?

These are the stitching options in Junos OS 9.3 and later:

- Layer 2 circuit into Layer 2 VPN.



NOTE: In Junos OS 9.4 and later, you can use an **iw** interface instead of an **lt** interface for this purpose.

- Layer 2 circuit into Layer 2 circuit.

- Layer 2 VPN into Layer 2 VPN.

Can a Layer 2 circuit be terminated into VPLS?

Yes, Junos OS 9.2 and later uses mesh groups to terminate pseudowire emulation (pseudowireE) into VPLS. You can terminate one pseudowireE per mesh group within the VPLS. A maximum of 16 mesh groups is supported per VPLS instance.

What is the behavior of a virtual tunnel interface in different services routing instances?

If you configure a virtual tunnel (**vt**) interface on the egress PE router and that router is also a transit router for the point-to-multipoint LSP, the penultimate hop router sends just one copy of each packet over the link to the egress PE router. A **vt** interface can perform two lookups on an incoming packet, one for the multicast MPLS lookup and one for the IP lookup. This applies to Layer 3 VPNs and VPLS point-to-multipoint LSPs.

This is the behavior of **vt-interface** in different services:

- **NG-MVPN**: The MPLS lookup is performed first, with MPLS packet copying. One of the copies has a null label and is sent to **vt-iftl**. The null label in this case is a label pop operation. When the packet is received from **vt-iftl**, IP lookup and copying is performed in VRF.
- **VPLS**: Initial behavior is the same as in NG-MVPN: The MPLS lookup is performed first, with MPLS packet copying. One of the copies has a null label and is sent to **vt-iftl**. The null label in this case is a label pop operation. However, in VPLS, when a packet is received from **vt-iftl**, Ethernet MAC lookup and copying is performed in the VPLS edge (VE) or the virtual switch interface (VSI).
- **CCC**: Initial behavior is the same as in NG-MVPN and VPLS: The MPLS lookup is performed first, with MPLS packet copying. One of the copies has a null label, however, this label is sent to **_egress CCC** interface. A second lookup is not needed.

What is an example configuration to map a bridge domain into a Layer 2 circuit?

The following example shows a configuration to map a bridge domain into a Layer 2 circuit.



NOTE: This configuration requires that the other end of the circuit is sending packets tagged with VLAN-ID 100.

```
lt-0/0/0 {
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 100;
    peer-unit 1;
  }
  unit 1 {
    encapsulation vlan-bridge;
    vlan-id 100;
    peer-unit 0;
  }
}
```

```
}
ge-0/1/5 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 100 {
    encapsulation vlan-bridge;
    vlan-id 100;
  }
}
ge-0/2/5 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 100 {
    encapsulation vlan-bridge;
    vlan-id 100;
  }
}
l2circuit {
  neighbor 10.1.1.1 {
    interface lt-0/0/0.0 {
      virtual-circuit-id 10;
    }
  }
}
bridge-domains {
  bridge-l2cct {
    domain-type bridge;
    interface ge-0/1/5.100;
    interface ge-0/2/5.100;
    interface lt-0/0/0.1;
  }
}
```

**Related
Documentation**

- [MPLS Connectivity FAQ Overview on page 1](#)
- [MPLS Layer 3 VPN on MX Series, M Series, and T Series Routers FAQs on page 13](#)
- [Virtual Private LAN Service on MX Series Routers FAQs on page 1](#)

