

Multi-Chassis Endpoint for VPLS Active/ Standby Pseudowire

In This Chapter

This section provides information about multi-chassis endpoint for VPLS active/standby pseudowire.

Topics in this section include:

- [Applicability on page 404](#)
- [Overview on page 405](#)
- [Configuration on page 409](#)
- [Conclusion on page 433](#)

Applicability

The examples covered in this section are applicable to all 7x50 and 7710 SR series and were tested on Release 7.0.R6. The 7750 SR-c4 is supported from 8.0R4 and higher.

Multi-chassis endpoint peers must be the same chassis type.

Overview

When implementing a large VPLS, one of the limiting factors is the number of T-LDP sessions required for the full mesh of SDPs. Mesh SDPs are required between all PEs participating in the VPLS with a full mesh of T-LDP sessions.

This solution is not scalable, as the number of sessions grows more rapidly than the number of participating PEs. Several options exist to reduce the number of T-LDP sessions required in a large VPLS.

The first option is hierarchical VPLS (H-VPLS) with spoke SDPs. By using spoke SDPs between two clouds of fully meshed PEs, any-to-any T-LDP sessions for all participating PEs are not required.

However, if spoke SDP redundancy is required, STP must be used to avoid a loop in the VPLS. Management VPLS can be used to reduce the number of STP instances and separate customer and STP traffic ([Figure 60](#)).

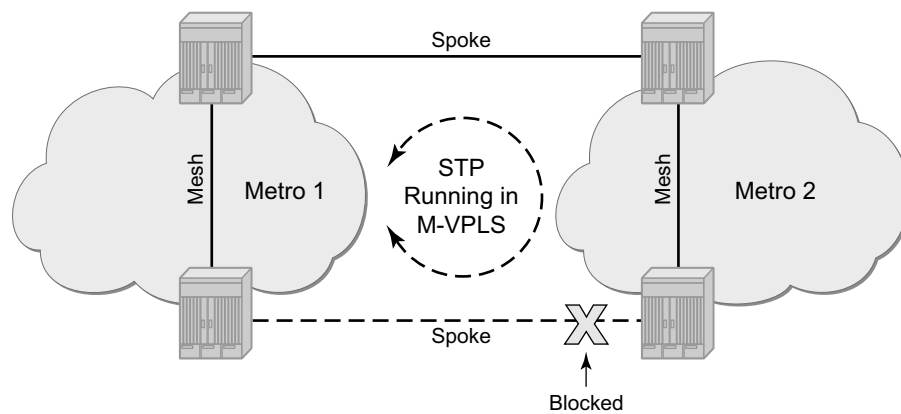


Figure 60: H-VPLS with STP

VPLS pseudowire redundancy provides H-VPLS redundant spoke connectivity. The active spoke is in forwarding state, while the standby spoke is in blocking state. Hence, STP is not needed anymore to break the loop, as illustrated in [Figure 61](#).

However, the PE implementing the active and standby spokes represents a single point of failure in the network.

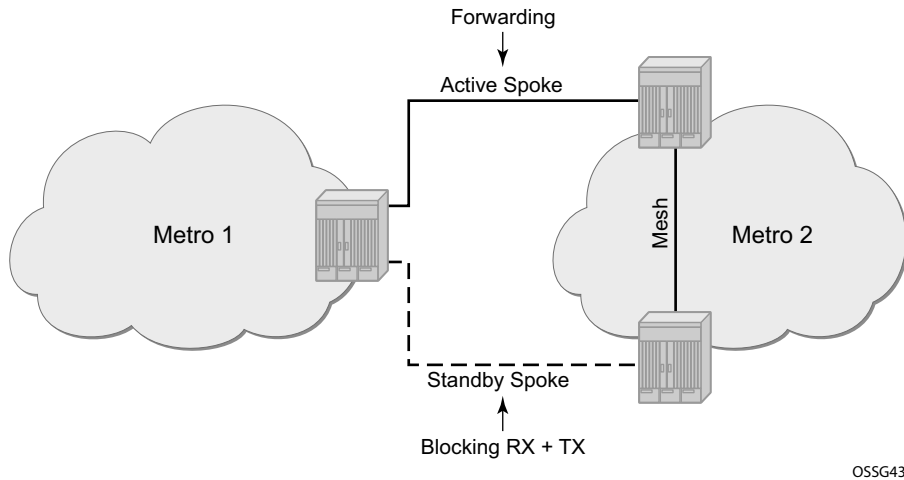


Figure 61: VPLS Pseudowire Redundancy

Multi-chassis endpoint for VPLS active/standby pseudowire expands on the VPLS pseudowire redundancy and allows the removal of the single point of failure.

There is only one spoke in forwarding state, all standby spokes are in blocking state. Mesh and square resiliency are supported.

A mesh resiliency can protect against simultaneous node failure in the core and in the MC-EP (double failure), but requires more SDPs (and thus more T-LDP sessions). Mesh resiliency is illustrated in [Figure 62](#).




Figure 62: Multi-Chassis Endpoint with Mesh Resiliency

Multi-Chassis Endpoint for VPLS Active/Standby PW

A square resiliency provides single failure node protection, and requires less SDPs (and thus less T-LDP sessions). Square resiliency is illustrated in [Figure 63](#).

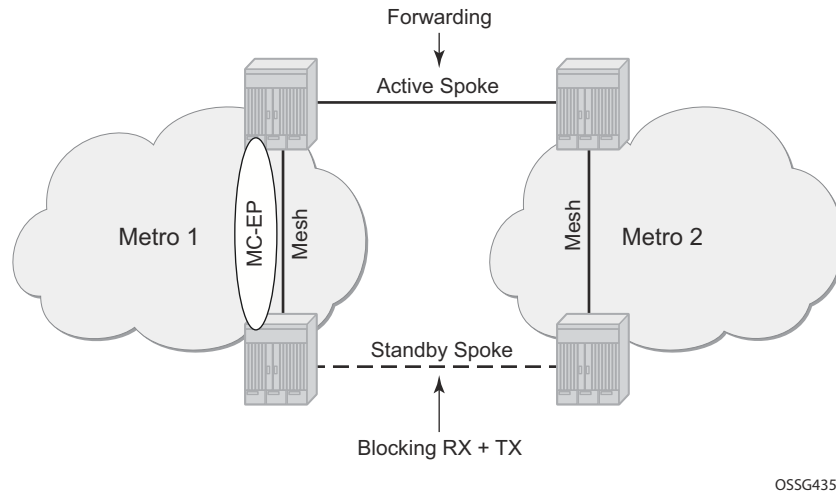


Figure 63: Multi-Chassis Endpoint with Square Resiliency

Network Topology

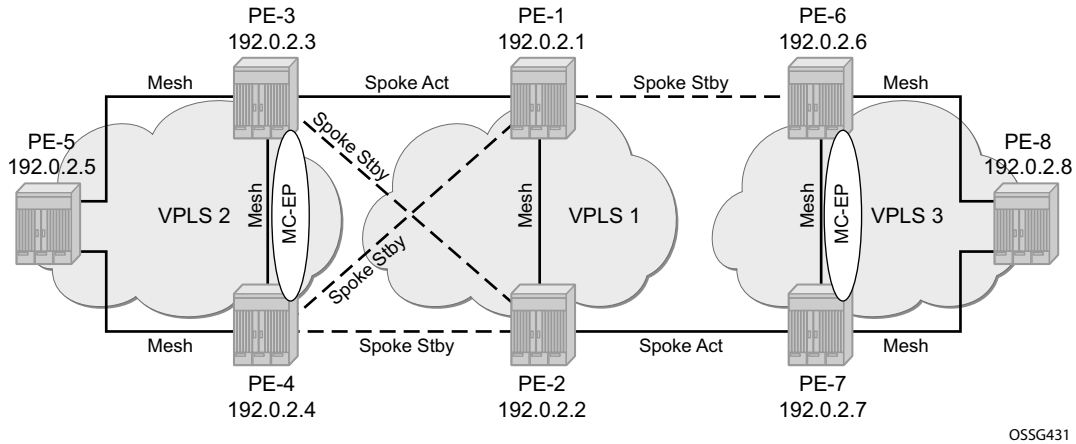


Figure 64: Network Topology

The network topology is displayed in [Figure 64](#).

The setup consists of:

- Two core nodes (PE-1 and PE-2), and three nodes for each metro area (PE-3, PE-4, PE-5 and PE-6, PE-7, PE-8, respectively).
- VPLS 1 is the core VPLS, used to interconnect the two metro areas represented by VPLS 2 and VPLS 3.
- VPLS 2 will be connected to the core VPLS in mesh resiliency.
- VPLS 3 will be connected to the core VPLS in square resiliency.

Note that three separate VPLS identifiers are used for clarity, however, the same identifier could be used for each. For interoperation, only the same VC-ID is required to be used on both ends of the spoke SDPs.

The following configuration tasks should be done first:

- ISIS or OSPF throughout the network.
- RSVP or LDP-signalled LSPs over the paths used for mesh/spoke SDPs.

Configuration

SDP Configuration

On each PE, SDPs are created to match the topology described in [Figure 64](#).

The convention for the SDP naming is: XY where X is the originating node and Y the target node.

An example of the SDP configuration in PE-3 (using LDP):

```
configure service
  sdp 31 mpls create
    far-end 192.0.2.1
    ldp
    keep-alive
    shutdown
  exit
  no shutdown
exit
sdp 32 mpls create
  far-end 192.0.2.2
  ldp
  keep-alive
  shutdown
  exit
  no shutdown
exit
sdp 34 mpls create
  far-end 192.0.2.4
  ldp
  keep-alive
  shutdown
  exit
  no shutdown
exit
sdp 35 mpls create
  far-end 192.0.2.5
  ldp
  keep-alive
  shutdown
  exit
  no shutdown
exit
exit
```

SDP Configuration

Verification of the SDPs on PE-3:

```
A:PE-3# show service sdp
```

```
=====
```

```
Services: Service Destination Points
```

```
=====
```

SdpId	Adm MTU	Opr MTU	IP address	Adm	Opr	Deliver	Signal
31	0	1556	192.0.2.1	Up	Up	LDP	TLDP
32	0	1556	192.0.2.2	Up	Up	LDP	TLDP
34	0	1556	192.0.2.4	Up	Up	LDP	TLDP
35	0	1556	192.0.2.5	Up	Up	LDP	TLDP

```
-----
```

```
Number of SDPs : 4
```

```
-----
```

```
=====
```

```
A:PE-3#
```


Full Mesh VPLS Configuration

Next, three fully meshed VPLS services are configured.

- VPLS 1 is the core VPLS, on PE-1 and PE-2
- VPLS 2 is the metro 1 VPLS, on PE-3, PE-4 and PE-5
- VPLS 3 is the metro 2 VPLS, on PE-6, PE-7 and PE-8

On PE-1 (similar configuration on PE-2):

```
configure service
  vpls 1 customer 1 create
    description "core VPLS"
    stp
      shutdown
    exit
  mesh-sdp 12:1 create
    exit
  no shutdown
exit
exit
```

On PE-3 (similar configuration on PE-4 and PE-5):

```
configure service
  description "Metro 1 VPLS"
  vpls 2 customer 1 create
    stp
      shutdown
    exit
  mesh-sdp 34:2 create
    exit
  mesh-sdp 35:2 create
    exit
  no shutdown
exit
exit
```

On PE-6 (similar configuration on PE-7 and PE-8):

```
configure service
  description "Metro 2 VPLS"
  vpls 3 customer 1 create
    stp
      shutdown
    exit
  mesh-sdp 67:3 create
    exit
  mesh-sdp 68:3 create
```

Full Mesh VPLS Configuration

```
        exit
    no shutdown
    exit
exit
```

Verification of the VPLS

- The service must be operationally up.
- All mesh SDPs must be up in the VPLS service.

On PE-6 (similar on other nodes):

```
A:PE-6# show service id 3 base
=====
Service Basic Information
=====
Service Id       : 3                Vpn Id           : 0
Service Type     : VPLS
Description      : (Not Specified)
Customer Id     : 1
Last Status Change: 12/03/2009 14:05:31
Last Mgmt Change : 12/03/2009 14:05:17
Admin State      : Up                Oper State       : Up
MTU              : 1514              Def. Mesh VC Id : 3
SAP Count       : 0                  SDP Bind Count  : 3
Snd Flush on Fail : Disabled         Host Conn Verify : Disabled
Propagate MacFlush: Disabled
Def. Gateway IP  : None
Def. Gateway MAC : None
=====
Service Access & Destination Points
=====
Identifier                               Type           AdmMTU  OprMTU  Adm  Opr
-----
sdp:67:3 M(192.0.2.7)                    n/a           0       1556   Up   Up
sdp:68:3 M(192.0.2.8)                    n/a           0       1556   Up   Up
=====
A:PE-6#
```

Multi-Chassis Configuration

Multi-chassis will be configured on the MC peers PE-3, PE-4 and PE-6, PE-7. The peer system address is configured, and mc-endpoint will be enabled.

On PE-3 (similar configuration on PE-4, PE-6 and PE-7):

```
configure redundancy
  multi-chassis
    peer 192.0.2.4 create
      mc-endpoint
      no shutdown
    exit
  no shutdown
exit
exit
exit
```

Verification of the multi-chassis synchronization:

If the multi-chassis synchronization fails, both nodes will fall back to single-chassis mode. In that case, two spoke SDPs could become active at the same time. It is important to verify the multi-chassis synchronization before enabling the redundant spoke SDPs.

```
A:PE-3# show redundancy multi-chassis mc-endpoint peer 192.0.2.4
```

```
=====
Multi-Chassis MC-Endpoint
=====
```

Peer Addr	: 192.0.2.4	Peer Name	:
Admin State	: up	Oper State	: up
Last State chg	: 12/03/2009 14:08:45	Source Addr	: 0.0.0.0
System Id	: 26:54:ff:00:00:00	Sys Priority	: 0
Keep Alive Intvl	: 10	Hold on Nbr Fail	: 3
Passive Mode	: disabled	Psv Mode Oper	: No
Boot Timer	: 300	BFD	: disabled
Last update	: 12/08/2009 09:17:41	MC-EP Count	: 1

```
=====
A:PE-3#
```

Mesh Resiliency Configuration

PE-3 and PE-4 will be connected to the core VPLS in mesh resiliency.

- First an endpoint is configured.
- The “no suppress-standby-signaling” is needed to block the standby spoke SDP.
- The multi-chassis endpoint peer is configured. The mc-endpoint ID must match between the two peers.

On PE-3 (similar on PE-4):

```
configure service
  vpls 2 customer 1 create
    endpoint "core" create
    no suppress-standby-signaling
    mc-endpoint 1
      mc-ep-peer 192.0.2.4
    exit
  exit
exit
```

Two spoke SDPs are configured on each peer of the multi-chassis to the two nodes of the core VPLS (mesh resiliency). Each spoke SDP refers to the endpoint “core”.

The precedence is defined on the spoke SDPs as follows:

- Spoke SDP 31 on PE-3 will be active. It is configured as primary (= precedence 0).
- Spoke SDP 32 on PE-3 will be the first backup. It is configured with precedence 1.
- Spoke SDP 41 on PE-4 will be the second backup. It is configured with precedence 2.
- Spoke SDP 42 on PE-4 will be the third backup. It is configured with precedence 3.

On PE-3 (similar on PE-4):

```
configure service
  vpls 2 customer 1 create
    spoke-sdp 31:1 endpoint "core" create
      stp
        shutdown
      exit
      precedence primary
    exit
    spoke-sdp 32:1 endpoint "core" create
      stp
        shutdown
      exit
      precedence 1
    exit
  exit
exit
```

Verification of the spoke SDPs:

On PE-3 and PE-4, the spoke SDPs must be up.

```
A:PE-3# show service id 2 sdp
=====
Services: Service Destination Points
=====
SdpId          Type IP address   Adm   Opr    I.Lbl    E.Lbl
-----
31:1           Spok 192.0.2.1    Up    Up     131059   131062
32:1           Spok 192.0.2.2    Up    Up     131064   131067
34:2           Mesh 192.0.2.4    Up    Up     131068   131067
35:2           Mesh 192.0.2.5    Up    Up     131063   131063
-----
Number of SDPs : 4
-----
=====
A:PE-3#
```

The endpoints on PE-3 and PE-4 can be verified. One spoke SDP is in Tx-Active mode (31 on PE-1 because it is configured as primary).

```
A:PE-3# show service id 2 endpoint core | match "Tx Active"
Tx Active           : 31:1
Tx Active Up Time   : 0d 00:03:04
Tx Active Change Count : 15
Last Tx Active Change : 12/08/2009 15:24:54
```

There is no active spoke SDP on PE-4.

```
A:PE-4# show service id 2 endpoint | match "Tx Active"
Tx Active           : none
Tx Active Up Time   : 0d 00:00:00
Tx Active Change Count : 18
```

Full Mesh VPLS Configuration

On PE-1 and PE-2, the spoke SDPs are operationally up.

```
A:PE-1# show service id 1 sdp
=====
Services: Service Destination Points
=====
SdpId          Type IP address   Adm   Opr     I.Lbl   E.Lbl
-----
12:1           Mesh 192.0.2.2     Up    Up      131060  131060
13:1           Spok 192.0.2.3      Up    Up      131062  131059
14:1           Spok 192.0.2.4      Up    Up      131063  131060
-----
Number of SDPs : 3
-----
A:PE-1#
```

However, because pseudowire signaling has been enabled, only one spoke SDP will be active, the others are set in standby.

On PE-1, spoke SDP 13 is active (no pseudowire bit signalled from PE-3).

And the spoke SDP 14 is signalled in standby by PE-4.

```
A:PE-1# show service id 1 sdp 13:1 detail | match "Peer Pw Bits"
Peer Pw Bits      : None
A:PE-1# show service id 1 sdp 14:1 detail | match "Peer Pw Bits"
Peer Pw Bits      : pwFwdingStandby
```

On PE-2, both spoke SDPs are signalled in standby.

```
A:PE-2# show service id 1 sdp 23:1 detail | match "Peer Pw Bits"
Peer Pw Bits      : pwFwdingStandby
A:PE-2# show service id 1 sdp 24:1 detail | match "Peer Pw Bits"
Peer Pw Bits      : pwFwdingStandby
```

There is one active and three standby spoke SDPs.

Square Resiliency Configuration

PE-6 and PE-7 will be connected to the core VPLS in square resiliency.

- First an endpoint is configured.
- The “no suppress-standby-signaling” is needed to block the standby spoke SDP.
- The multi-chassis endpoint peer is configured. The mc-endpoint ID must match between the two peers.

On PE-7 (similar on PE-6):

```
configure service
  vpls 3 customer 1 create
    endpoint "core" create
      no suppress-standby-signaling
      mc-endpoint 1
      mc-ep-peer 192.0.2.6
    exit
  exit
exit
```

One spoke SDP is configured on each peer of the multi-chassis to one node of the core VPLS (square resiliency). Each spoke SDP refers to the endpoint core.

The precedence will be defined on the spoke SDPs as follows:

- Spoke SDP 72 on PE-7 will be active. It is configured as primary (= precedence 0)
- Spoke SDP 61 on PE-6 will be the first backup. It is configured with precedence 1.

On PE-7 (similar on PE-6):

```
configure service
  vpls 3 customer 1 create
    spoke SDP 72:1 endpoint "core" create
      stp
      shutdown
    exit
    precedence primary
  exit
exit
```

Verification of the spoke SDPs.

On PE-6 and PE-7, the spoke SDPs must be up.

```
A:PE-7# show service id 3 sdp
=====
Services: Service Destination Points
=====
SdpId          Type IP address   Adm   Opr    I.Lbl   E.Lbl
-----
72:1           Spok 192.0.2.2     Up    Up     131063  131062
76:3           Mesh 192.0.2.6     Up    Up     131062  131064
78:3           Mesh 192.0.2.8     Up    Up     131068  131066
-----
Number of SDPs : 3
-----
=====
A:PE-7#
```

The endpoints on PE-7 and PE-6 can be verified. One spoke SDP is in Tx-Active mode (72 on PE-7 because it is configured as primary).

```
A:PE-7# show service id 3 endpoint | match "Tx Active"
Tx Active           : 72:1
Tx Active Up Time   : 0d 04:19:01
Tx Active Change Count : 3
Last Tx Active Change : 12/08/2009 11:02:10
```

There are no active spoke SDP on PE-6.

```
A:PE-6# show service id 3 endpoint | match "Tx Active"
Tx Active           : none
Tx Active Up Time   : 0d 00:00:00
Tx Active Change Count : 2
Last Tx Active Change : 12/08/2009 11:17:31
```

The output shows that on PE-1, spoke SDP 16 is signalled with peer in standby mode.

```
A:PE-1# show service id 1 sdp 16:1 detail | match "Peer Pw Bits"
Peer Pw Bits       : pwFwdingStandby
```


On PE-2, the spoke SDP 27 is signalled with peer active (no pseudowire bits).

```
A:PE-2# show service id 1 sdp 27:1 detail | match "Peer Pw Bits"
Peer Pw Bits          : None
```

There is one active and one standby spoke SDPs.

Additional Parameters

Multi-Chassis

Peer Failure Detection

The default mechanism is based on the keepalives exchanged between the peers.

```
A:PE-3# configure redundancy multi-chassis peer 192.0.2.4 mc-endpoint
[no] hold-on-neighb* - Configure hold time applied on neighbor failure
[no] keep-alive-int* - Configure keep alive interval for this MC-Endpoint
```

Keep-alive-interval is the interval at which keepalives are sent to the MC peer. It is set in tenths of a second from 5 to 500), with a default value of 5.

Hold-on-neighbor-failure is the number of keep-alive-intervals that the node will wait for a packet from the peer before assuming it has failed. After this interval, the node will revert to single chassis behavior. It can be set from 2 to 25 with a default value of 3.

BFD Session

BFD is another peer failure mechanism. It can be used to speed up the convergence in case of peer loss.

```
A:PE-3# configure redundancy multi-chassis peer 192.0.2.4 mc-endpoint
[no] bfd-enable - Configure BFD
```

It is using the centralized BFD session. BFD must be enabled on the system interface.

```
configure router
  interface "system"
    address 192.0.2.3/32
    bfd 100 receive 100 multiplier 3
  exit
exit
```

Verification of the BFD session

```
A:PE-3# show router bfd session
=====
BFD Session
=====
Interface          State          Tx Intvl  Rx Intvl  Mult
  Remote Address   Protocol
-----
system             Up (3)        100       100       3
  192.0.2.4       mcep         65        65
-----
No. of BFD sessions: 1
=====
A:PE-3#
```

Boot Timer

```
A:PE-3# configure redundancy multi-chassis peer 192.0.2.4 mc-endpoint
[no] boot-timer      - Configure boot timer interval
```

The **boot-timer** command specifies the time after a reboot that the node will try to establish a connection with the MC peer before assuming a peer failure. In case of failure, the node will revert to single chassis behavior.

System Priority

```
A:PE-3# configure redundancy multi-chassis peer 192.0.2.4 mc-endpoint
[no] system-priority - Configure system priority
```

The system priority influences the selection of the MC master. The lowest priority node will become the master.

In case of equal priorities, the lowest system-id (=chassis MAC address) will become the master.

VPLS Endpoint and Spoke SDP

Ignore Standby Pseudowire Bits

```
A:PE-1# configure service vpls 1 spoke-sdp 14:1
    ignore-standby-signaling
```

The peer pseudowire status bits are ignored and traffic is forwarded over the spoke SDP.

It can speed up convergence for multicast traffic in case of spoke SDP failure.

Traffic sent over the standby spoke SDP will be discarded by the peer.

In this topology, if the **ignore-standby-signaling** command is enabled on PE-1, it sends MC traffic to PE-3 and PE-4 (and to PE-6). If PE-3 fails, PE-4 can start forwarding traffic in the VPLS as soon as it detects PE-3 being down. There is no signaling needed between PE-1 and PE-4.

Block-on-Mesh-Failure

```
A:PE-3# configure service vpls 2 endpoint core
    block-on-mesh-failure
```

In case a PE loses all the mesh SDPs of a VPLS, it should block the spoke SDPs to the core VPLS, and inform the MC-EP peer who can activate one of its spoke SDPs.

If block-on-mesh-failure is enabled, the PE will signal all the pseudowires of the endpoint in standby.

In this topology, if PE3 does not have any valid mesh SDP to the VPLS 2 mesh, it will set the spoke SDPs under endpoint "core" in standby.

When block-on-mesh-failure is activated under an endpoint, it is automatically set under the spoke SDPs belonging to this endpoint.

```
A:PE-3>config>service>vpls# info
-----
    endpoint "core" create
    exit
    spoke-sdp 31:1 endpoint "core" create
        precedence primary
    exit
    spoke-sdp 32:1 endpoint "core" create
        precedence 1
    exit
    mesh-sdp 34:2 create
    exit
    mesh-sdp 35:2 create
```

Multi-Chassis Endpoint for VPLS Active/Standby PW

```
exit
no shutdown
-----
A:PE-3>config>service>vpls# endpoint core block-on-mesh-failure
A:PE-3>config>service>vpls# info
-----
endpoint "core" create
    block-on-mesh-failure
exit
spoke-sdp 31:1 endpoint "core" create
    block-on-mesh-failure
    precedence primary
exit
spoke-sdp 32:1 endpoint "core" create
    block-on-mesh-failure
    precedence 1
exit
mesh-sdp 34:2 create
exit
mesh-sdp 35:2 create
exit
no shutdown
-----
A:PE-3>config>service>vpls#
```

Additional Parameters

Precedence

```
A:PE-3# configure service vpls 2 spoke-sdp 31:1
[no] precedence      - Configure the spoke-sdp precedence
```

The precedence is used to indicate in which order the spoke SDPs should be used. The value is from 0 to 4 (0 being primary), the lowest having higher priority. The default value is 4.

Revert-Time

```
A:PE-3# configure service vpls 2 endpoint core
[no] revert-time     - Configure the time to wait before reverting to primary
                    spoke-sdp
```

If the precedence is equal between the spoke SDPs, there is no revertive behavior. Changing the precedence of a spoke SDP will not trigger a revert. The default is **no revert**.

MAC-Flush Parameters

When a spoke SDP goes from standby to active (due to the active spoke SDP failure), the node will send a **flush-all-but-mine** message.

After a restoration of the spoke SDP, a new **flush-all-but-mine** message will be sent.

A node configured with **propagate MAC flush** will forward the flush messages received on the spoke-SDP to its other mesh/spoke SDPs.

```
A:PE-1# configure service vpls 1
propagate-mac-flush
```

A node configured with **send flush on failure** will send a **flush-all-from-me** message when one of its SDPs goes down.

```
A:PE-1# configure service vpls 1
send-flush-on-failure
```

Failure Scenarios

For the subsequent failure scenarios, the configuration of the nodes is as described in the [Configuration on page 409](#).

Core Node Failure

When the Core Node PE-1 fails, the spoke SDPs from PE-3 and PE-4 go down.

Because the spoke SDP 31 between PE-3 and PE-4 was active, the MC master (PE-3 in this case) will select the next best spoke SDP, which will be 32 between PE-3 and PE-2 (precedence 1). See [Figure 65](#).




Figure 65: Core Node Failure

```
A:PE-3# show service id 2 endpoint
=====
Service 2 endpoints
=====
Endpoint name           : core
Description             : (Not Specified)
Revert time             : 0
Act Hold Delay          : 0
Ignore Standby Signaling : false
Suppress Standby Signaling : false
Block On Mesh Fail      : false
Multi-Chassis Endpoint  : 1
MC Endpoint Peer Addr   : 192.0.2.4
Psv Mode Active         : No
Tx Active                : 32:1
Tx Active Up Time       : 0d 00:01:45
Tx Active Change Count  : 19
Last Tx Active Change   : 12/08/2009 16:09:07
```

Additional Parameters

```
-----  
Members  
-----  
Spoke-sdp: 31:1 Prec:0                               Oper Status: Down  
Spoke-sdp: 32:1 Prec:1                               Oper Status: Up  
=====
```

A:PE-3#

A:PE-4# show service id 2 endpoint

```
-----  
Service 2 endpoints  
=====
```

Endpoint name	: core
Description	: (Not Specified)
Revert time	: 0
Act Hold Delay	: 0
Ignore Standby Signaling	: false
Suppress Standby Signaling	: false
Block On Mesh Fail	: false
Multi-Chassis Endpoint	: 1
MC Endpoint Peer Addr	: 192.0.2.3
Psv Mode Active	: No
Tx Active	: none
Tx Active Up Time	: 0d 00:00:00
Tx Active Change Count	: 21
Last Tx Active Change	: 12/08/2009 16:08:12

```
-----  
Members  
-----  
Spoke-sdp: 41:1 Prec:2                               Oper Status: Down  
Spoke-sdp: 42:1 Prec:3                               Oper Status: Up  
=====
```

A:PE-4#

Multi-Chassis Node Failure

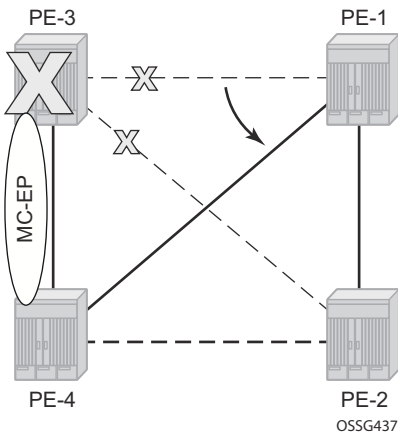


Figure 66: Multi-Chassis Node Failure

When the multi-chassis node PE-3 fails, both spoke SDPs from PE-3 go down.

PE-4 reverts to single chassis mode and selects the best spoke SDP, which will be 41 between PE-4 and PE-1 (precedence 2). See [Figure 66](#).

```
A:PE-4# show redundancy multi-chassis mc-endpoint peer 192.0.2.3
```

```
=====
Multi-Chassis MC-Endpoint
=====
```

```
Peer Addr       : 192.0.2.3           Peer Name       :
Admin State     : up                  Oper State      : down
Last State chg  : 12/08/2009 15:41:58 Source Addr     : 0.0.0.0
System Id       : 1e:28:ff:00:00:00  Sys Priority    : 0
Keep Alive Intvl: 10                  Hold on Nbr Fail : 3
Passive Mode    : disabled            Psv Mode Oper   : No
Boot Timer      : 300                 BFD             : disabled
Last update     : 12/08/2009 15:05:08 MC-EP Count     : 1
=====
```

```
A:PE-4#
```

```
A:PE-4# show service id 2 endpoint
```

```
=====
Service 2 endpoints
=====
```

```
Endpoint name      : core
Description        : (Not Specified)
Revert time        : 0
Act Hold Delay     : 0
Ignore Standby Signaling : false
Suppress Standby Signaling : false
Block On Mesh Fail : false
Multi-Chassis Endpoint : 1
MC Endpoint Peer Addr : 192.0.2.3
```

Additional Parameters

```
Psv Mode Active          : No
Tx Active                : 41:1
Tx Active Up Time       : 0d 00:00:13
Tx Active Change Count  : 19
Last Tx Active Change   : 12/08/2009 15:41:58
```

Members

```
-----  
Spoke-sdp: 41:1 Prec:2          Oper Status: Up  
Spoke-sdp: 42:1 Prec:3          Oper Status: Up  
=====
```

A:PE-4#

Multi-Chassis Communication Failure

If the multi-chassis communication is interrupted, both nodes will revert to single chassis mode.

To simulate a communication failure between the two nodes, define a static route on PE-3 that will black-hole the system address of PE-4.

```
A:PE-3# configure router
A:PE-3>config>router# static-route 192.0.2.4/32 black-hole
```

Verify that the MC synchronization is down.

```
A:PE-4# show redundancy multi-chassis mc-endpoint peer 192.0.2.3
=====
Multi-Chassis MC-Endpoint
=====
Peer Addr       : 192.0.2.3           Peer Name       :
Admin State     : up                Oper State      : down
Last State chg  : 12/11/2009 15:12:25 Source Addr     : 0.0.0.0
System Id       : 1e:28:ff:00:00:00 Sys Priority     : 0
Keep Alive Intvl: 10                Hold on Nbr Fail : 3
Passive Mode    : disabled           Psv Mode Oper   : No
Boot Timer      : 300                 BFD              : disabled
Last update     : 12/08/2009 15:05:08 MC-EP Count     : 1
=====
A:PE-4#
```

The spoke SDPs are active on PE-3 and on PE-4.

```
A:PE-3# show service id 2 endpoint | match "Tx Active"
Tx Active           : 31:1
Tx Active Up Time   : 1d 23:29:04
Tx Active Change Count : 20
Last Tx Active Change : 12/09/2009 15:50:22

A:PE-4# show service id 2 endpoint | match "Tx Active"
Tx Active           : 42:1
Tx Active Up Time   : 0d 00:06:27
Tx Active Change Count : 23
```

This can potentially cause a loop in the system. The [Passive Mode on page 430](#) explains how to avoid this loop.

Passive Mode

As in [Multi-Chassis Communication Failure on page 429](#), if there is a failure in the multi-chassis communication, both nodes will assume that the peer is down and will revert to single-chassis mode. This can create loops because two spoke SDPs can become active.

One solution is to synchronize the two core nodes, and configure them in passive mode. See [Figure 67](#).

In passive mode, both peers will stay dormant as long as one active spoke SDP is signalled from the remote end. If more than one spoke SDP becomes active, the MC-EP algorithm will select the best SDP. All other spoke SDPs are blocked locally (in Rx and Tx directions). There is no signaling sent to the remote PEs.

If one peer is configured in passive mode, the other peer will be forced to passive mode as well.

The **no suppress-standby-signaling** and **no ignore-standby-signaling** commands are required.




Figure 67: Multi-Chassis Passive Mode

The following output shows the multi-chassis configuration on PE-1 (similar on PE-2).

```
A:PE-1# configure redundancy
      multi-chassis
        peer 192.0.2.2 create
          mc-endpoint
            no shutdown
            passive-mode
          exit
        no shutdown
      exit
    exit
```

The following output shows the VPLS spoke SDPs configuration on PE-1 (similar on PE-2)

```
A:PE-1# configure service vpls 1
      endpoint "metrol" create
        no suppress-standby-signaling
        mc-endpoint 1
        mc-ep-peer 192.0.2.2
      exit
    exit
  spoke-sdp 13:1 endpoint "metrol" create
    stp
      shutdown
    exit
  exit
  spoke-sdp 14:1 endpoint "metrol" create
    stp
      shutdown
    exit
  exit
```

To simulate a communication failure between the two nodes, a static route is defined on PE-3 that will black-hole the system address of PE-4.

```
A:PE-3# configure router
A:PE-3>config>router# static-route 192.0.2.4/32 black-hole
```

The spoke SDPs are active on PE-3 and on PE-4.

```
A:PE-3# show service id 2 endpoint | match "Tx Active"
Tx Active           : 31:1
Tx Active Up Time   : 1d 23:29:04
Tx Active Change Count : 20
Last Tx Active Change : 12/09/2009 15:50:22
```

```
A:PE-4# show service id 2 endpoint | match "Tx Active"
Tx Active           : 42:1
Tx Active Up Time   : 0d 00:06:27
Tx Active Change Count : 23
```

Additional Parameters

PE-1 and PE-2 have blocked one spoke SDP which avoids a loop in the VPLS.

```
A:PE-1# show service id 1 endpoint metrol | match "Tx Active"
Tx Active                : none
Tx Active Up Time       : 0d 00:00:00
Tx Active Change Count  : 4
Last Tx Active Change   : 12/11/2009 15:44:35
```

```
A:PE-2# show service id 1 endpoint metrol | match "Tx Active"
Tx Active                : 24:1
Tx Active Up Time       : 0d 00:02:39
Tx Active Change Count  : 1
Last Tx Active Change   : 12/11/2009 15:44:35
```

The passive nodes do not set the pseudowire status bits; hence, the nodes PE-3 and PE-4 are not aware that one spoke SDP is blocked.

Conclusion

Multi-chassis endpoint for VPLS active/standby pseudowire allows the building of hierarchical VPLS without single point of failure, and without requiring STP to avoid loops.

Care must be taken to avoid loops. The multi-chassis peer communication is important and should be possible on different interfaces.

Passive mode can be a solution to avoid loops in case of multi-chassis communication failure.

Conclusion