ExtremeSwitching™ Virtual Services Platform (VSP)

Episode 2: Switch Clustering incl. Link Health and Loop Prevention

Bin Han / Principal Systems Engineer

2023/01

Objectives

Explain the implementation of Switch Clusters in the context of the Fabric Connect solution:

- Virtual Inter-Switch Trunk (vIST) – non DVR-Leaf
- Virtual Inter-Switch Trunk (vIST) – DVR Leaf
- Simplified Virtual IST

Explain the operation of the Link Health and Loop Prevention protocols such as VLACP, SLPP, and SLPP Guard

ミスコミュニケーションを防ぐため、本資料はVOSS User Guide の原文(英文)をそのまま引用しておりますので、ご了承ください！

参照先

MultiLink Trunking and Split MultiLink Trunking

Virtual Inter-Switch Trunk (vIST)

Simplified Virtual-IST

MultiLink Trunking

Split MultiLink Trunking

Routed Split MultiLink Trunking

Virtual Router Redundancy Protocol

Virtual Link Aggregation Control Protocol

VLAN loop prevention - SLPP

SLPP Guard

Switch Clustering

Split Multi Link Trunking (SMLT)
Virtual Inter-Switch Trunk (vIST)
Simplified virtual IST

Extreme Networks Switch Clustering

Introduction

- ✓ Switch Clustering is a mechanism to provide Link and Nodal resilience at the edge of a Network
- ✓ Ensures no single point of failure at the network edge, thus preventing edge devices from being isolated from the network.
- ✓ Devices which form the switch cluster work in an Active-Active model with sub second stateful failover.
- ✓ All links which form the switch cluster pass traffic.
- ✓ Does not require spanning tree on switch to switch links
- ✓ Switch clustering with Extreme Networks Devices is based on the implementation of Inter-Switch-Trunk protocol, that allows the two nodes in a switch-Cluster to behave like a single switch (except the configuration and all control plane is still independently on both nodes)
  - As a result, Multi-Chassis Link aggregation can be used to connect in an active-active fashion switches or other network equipment – such connections are named Split Multi-Link Trunks (SMLT); such SMLT LAG can be built statically or dynamically (using LACP)
  - the Switch-Cluster also allows for L3 redundancy and load-balancing by providing pairs of router instances, that behave like twins – each of them is answering ARP and routing packets for the other instance, when that is “down”. This function is named “Routed SMLT” (RSMLT) and applies to both IPv4 and IPv6

Dual-Homing Support at the Edge

Dual-Homing SMLT Clustering into the Fabric

Enhancing 802.1aq by providing dual-homed active-active connectivity to the Campus Fabric for switches, appliances, servers, etc.

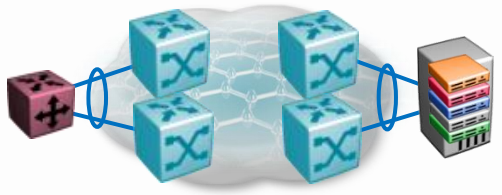

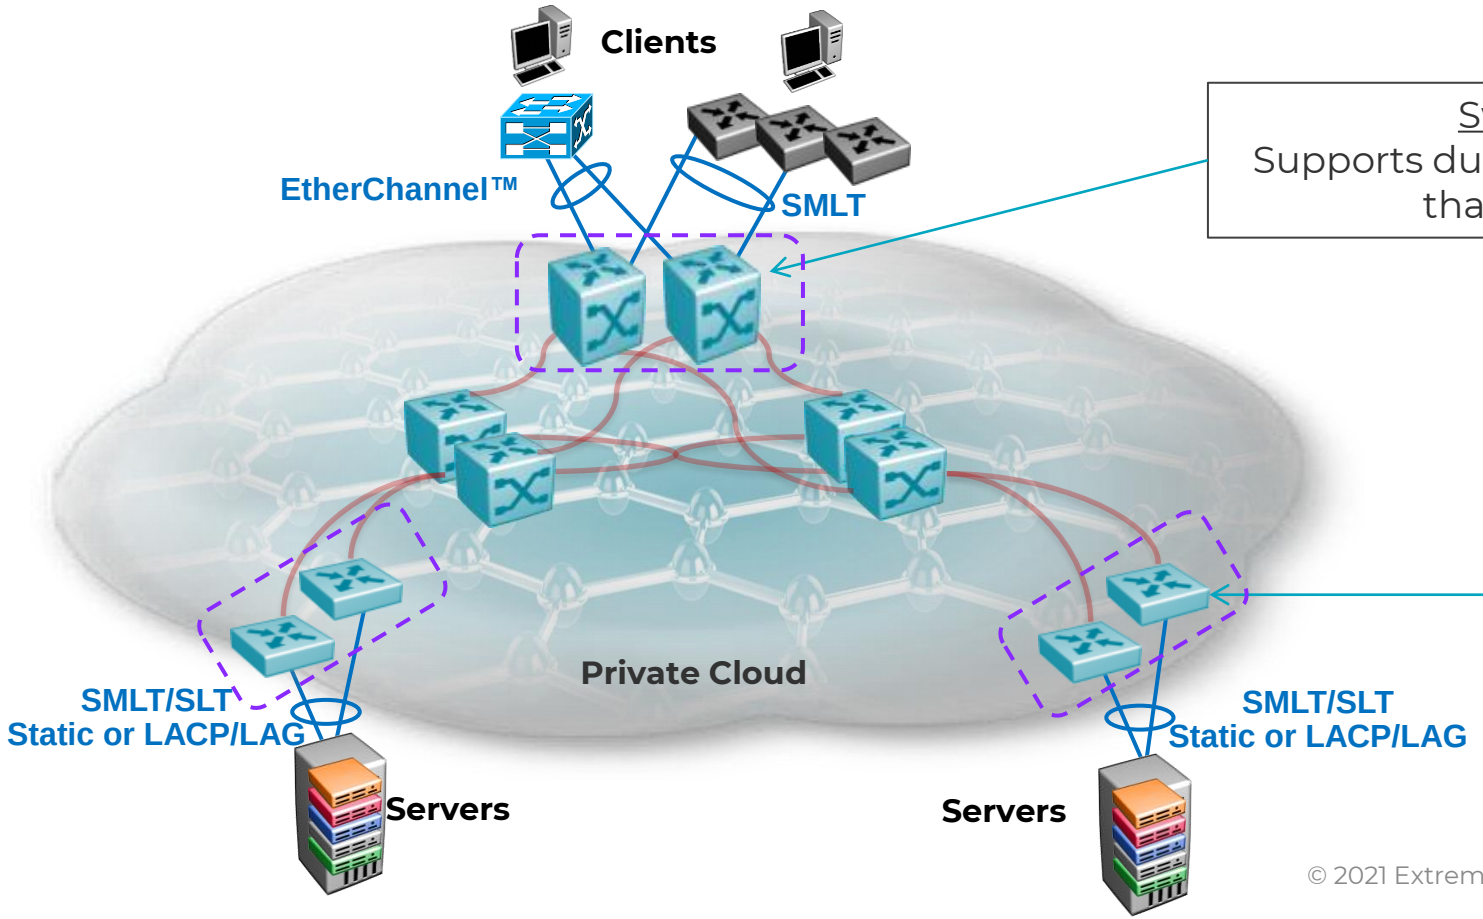

Extreme Networks Switch Clustering

How to build a Switch Cluster ?

Physical Design

- Create a Switch Cluster core with two devices
- Create the virtual Inter-Switch Trunk (IST) between the switches (virtual, because usually based on fabric L2 VSN; direct connection between nodes NOT required)
- Connect edge devices

Logical Design

- Default gateway Redundancy – VRRP or RSMLT
- Loop prevention – SLPP (covered later)
- Virtual LACP (end to end link failure detection)

Extreme Networks Switch Clustering

Two types of IST implementation possible for Switch Cluster

Cluster based on Simplified Virtual IST

- Works only in **non-SPBm** mode of operation (non-fabric network)
- works only for a single pair of switches

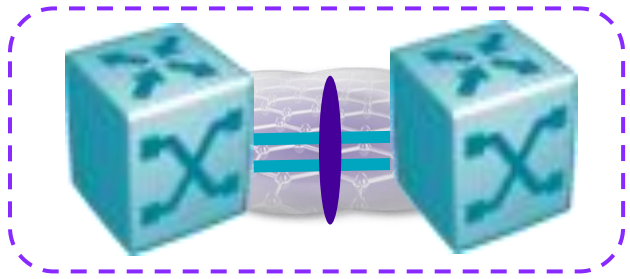

Cluster based on Virtual IST

- required in **Fabric networks** (more than 2 nodes in a Fabric exist or planned)
- Also works in conjunction with Distributed Virtual Routing (DvR)

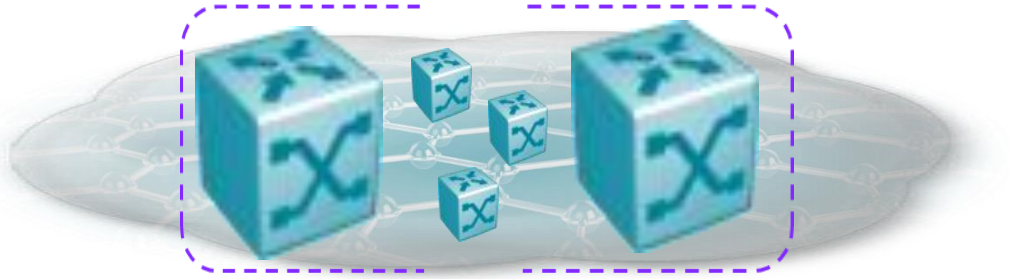

Virtual IST

Requried in Fabric Networks!!

Virtual Inter-Switch Trunk (vIST)

Connections between switch clusters can now be:

- Point-to-point MLT/DMLT OR
- A network connection over SPB

Requires SPBm connectivity between the switch cluster

- The SPBm cloud can consist of as few as two nodes

Operates between any two VOSS Devices that support vIST:

- Devices do not have to be of the same type

IST protocol traffic **uses a L2 VSN between the two nodes that form the switch cluster**

Virtual IST is used to:

- Confirm that the nodes are alive (IST hello messages)
- Exchange MAC address forwarding tables
- in case of RSMLT: to exchange IP/MAC addresses of router interfaces

Configuration:

- Configure IS-IS and SPB
- Set up vIST protocol between a cluster pair
- In case of DVR Leaf: simplified, single virtual-ist command
- in case of Simplified vIST
  - approach to configure IST w/o the need for manual configuration of IS-IS/SPB between a cluster pair

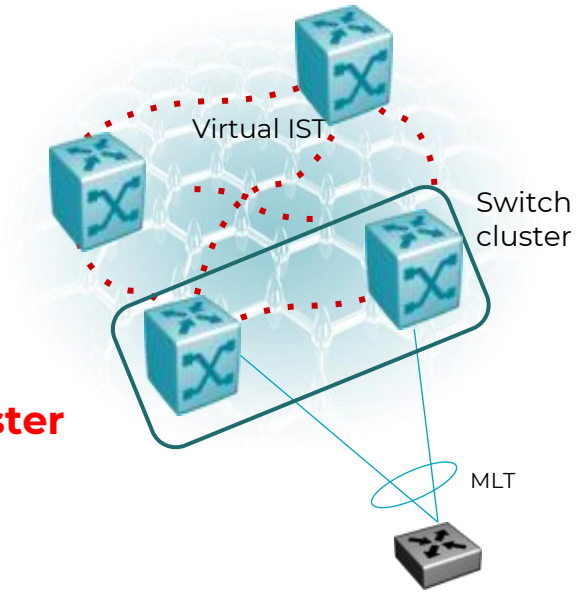

Virtual Inter-Switch Trunk

Redundancy examples

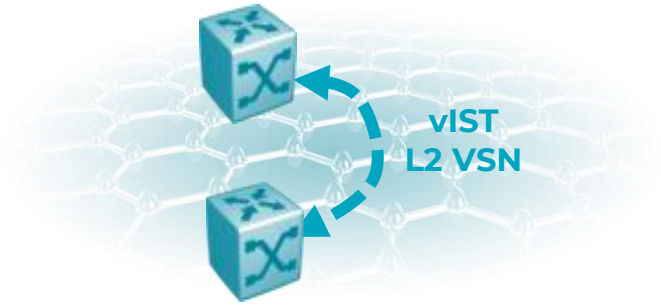

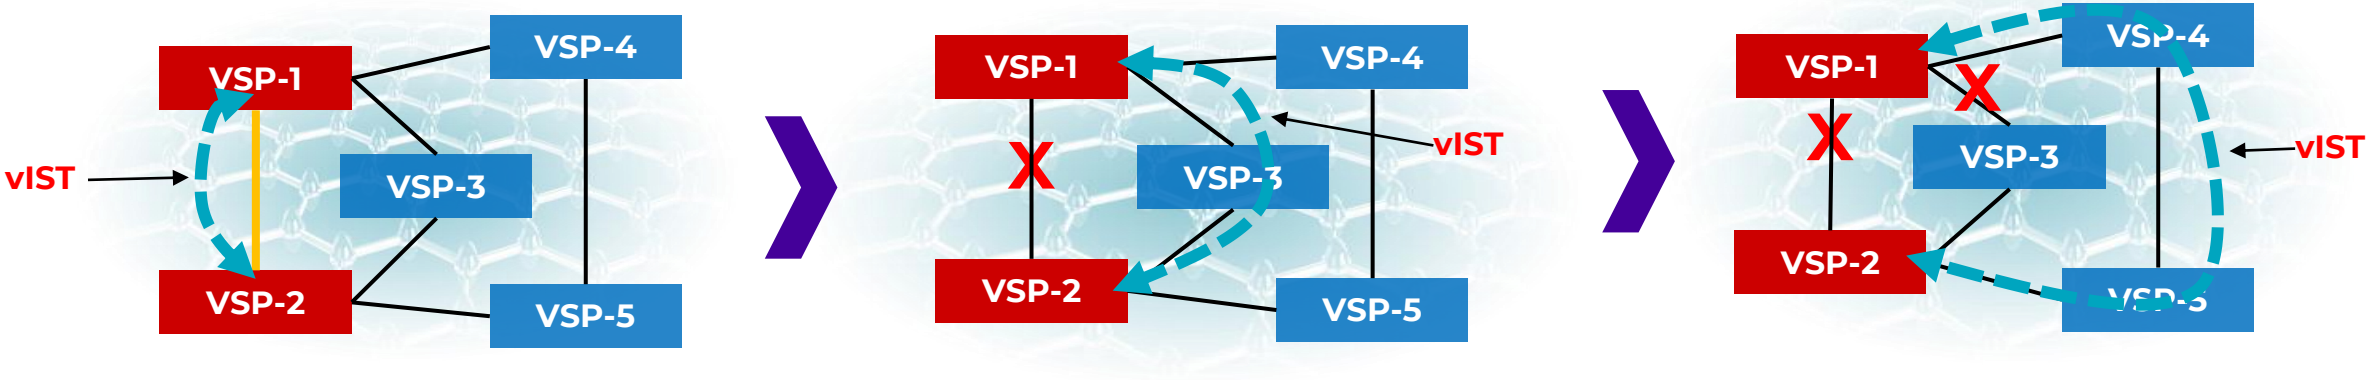

If there is a path, the Virtual IST is UP

Virtual IST

Summary

Active-Active solution protecting against link and switch failures

Sub-second failover for most Link and Switch DOWN/UP events:

- Faster failover times than MSTP/RSTP offer for Layer-2 traffic

Increased resiliency

Protection for L3 routed traffic using 3 options

- RSMLT (only 2 nodes – single Switch-Cluster)
- VRRP (any number of nodes)
- DVR (any number of nodes; optimizes traffic flows and simplifies configuration)

Virtual-IST allows mixing of node types

Virtual IST

No (direct) link between Cluster nodes is required, but possible

The IST tunnel is up if there is SPBM connectivity between the IST peers

When any link breaks, vIST used new (shortest) path over the SPBm cloud

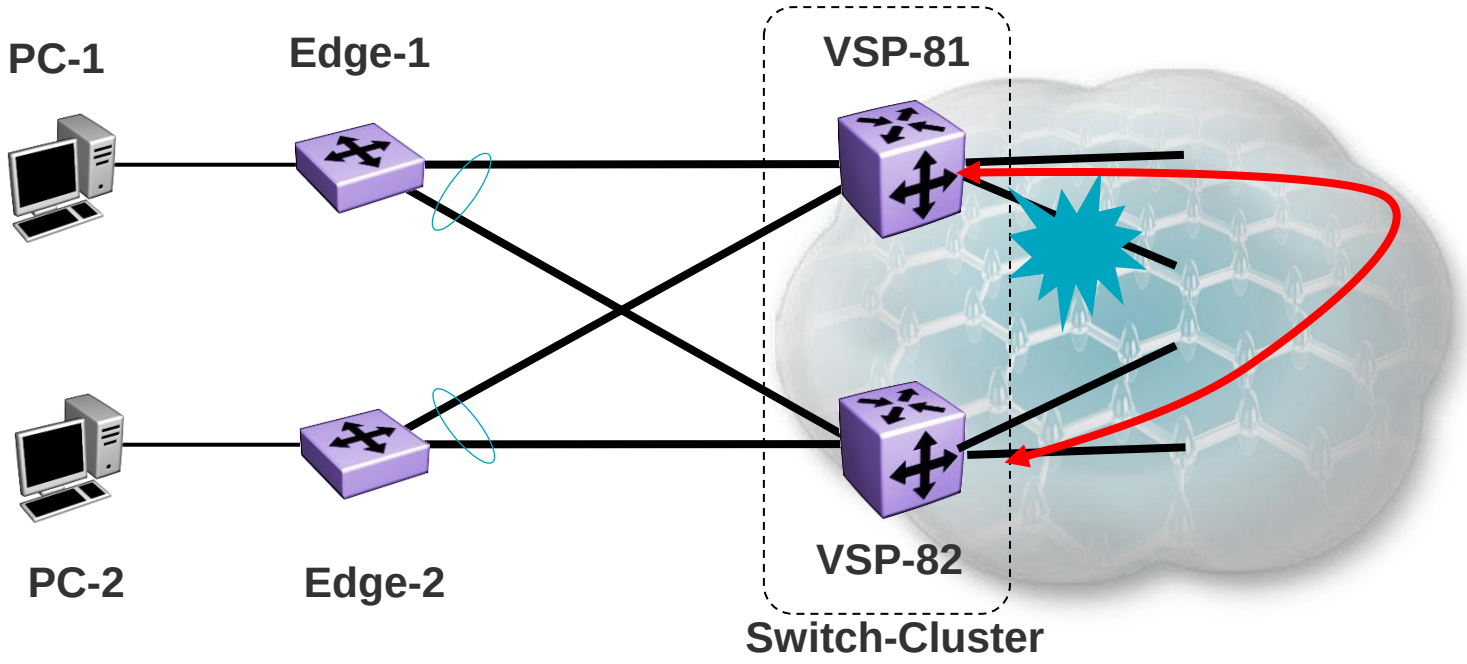

Virtual IST - Basics

Configuration Parameters

- Peer nodes “know” ISIS System-ID of each other and share Virtual BMAC
- Dedicated vIST-VLAN (port-based) with IP addresses & common I-SID
- Virtual-IST is configured with Peer IP of VIST VLAN

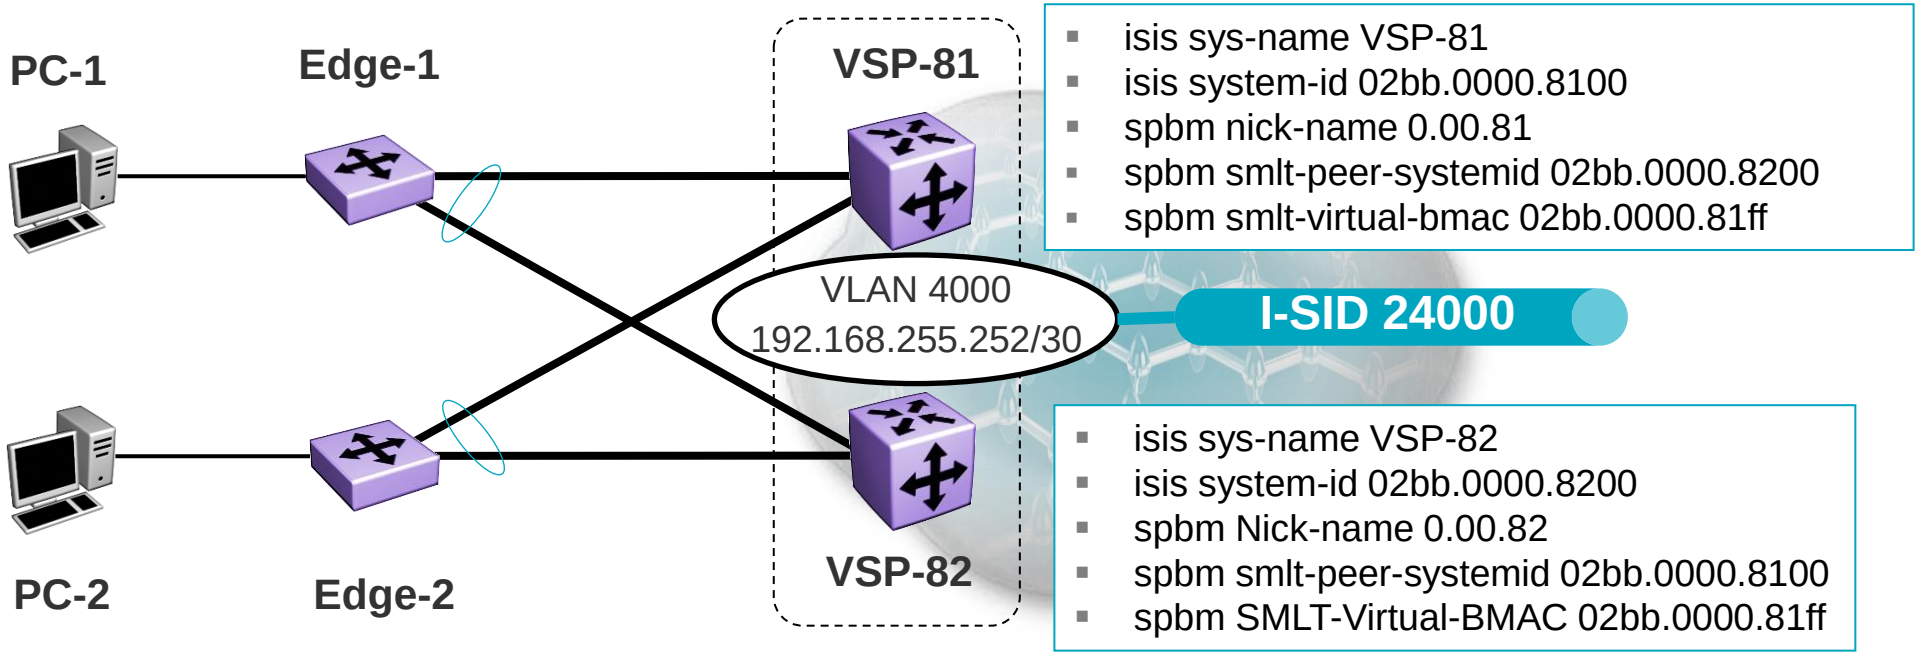

- isis sys-name VSP-81
- isis system-id 02bb.0000.8100
- spbm nick-name 0.00.81
- spbm smlt-peer-systemid 02bb.0000.8200
- spbm smlt-virtual-bmac 02bb.0000.81ff

- isis sys-name VSP-82
- isis system-id 02bb.0000.8200
- spbm Nick-name 0.00.82
- spbm smlt-peer-systemid 02bb.0000.8100
- spbm SMLT-Virtual-BMAC 02bb.0000.81ff

Virtual IST (not DVR Leaf)

Configuration

Steps	commands	remarks
Create L2 VSN for vIST	vlan create <c-vid> name vIST type port 0 vlan i-sid <c-vid> <isid>	Recommended to assign all VLANs to CIST („0“ at command line end)
Assign IP zu vIST VLAN	interface vlan <c-vid> ip address <A.B.C.D>/30 exit	This address would be peer-ip for vIST config on peer node of the cluster
Create vIST	virtual-ist peer-ip <peer> vlan <c-vid> no/default virtual-ist peer-ip	<peer> is IP on same L2VSN on peer node; <c-vid> is locally used VID from step 1
Validation	show virtual-ist show virtual-ist stat	Every second ‚hellos‘ should be seen (sent and received)
clear vIST Stats	clear virtual-ist stats	

Virtual IST (DVR Leaf)

Configuration

Special case, as DVR leaf nodes do not support basic VLAN configurations. Layer 2 connections on these nodes are always Switched UNI L2VSN, so VLANs do only exist behind interfaces, but not on the node itself (this is the concept of Switched UNI)

New command for Virtual-IST config on DVR leaf: „dvr leaf virtual-ist“.

Switch creates appropriate VLAN 4002 plus IP implicitly. It uses reserved ISID, starting with 16677216 (for Cluster ID = 1)

Steps	commands	remarks
Create	dvr leaf virtual-ist <own-IP> <mask> peer-ip <peer-IP> cluster-id <id>	Node will create VLAN 4002 with IP and assign I-SID 16677215+<cluster-id>
Validation	show virtual-ist show virtual-ist stat	Every second ‚hellos‘ should be seen (sent and received)
clear vIST stats	clear virtual-ist stats	

CVLAN Considerations

All CVLAN's must be mapped to an I-SID, even for single attached scenario <u>on both nodes</u>

- In case of single attached node (see Edge-2 below) VLAN 200 on VSP-81 does not have port-member, but I-SID assigned

Note: CVLAN's must not (and cannot) be mapped to the vIST I-SID

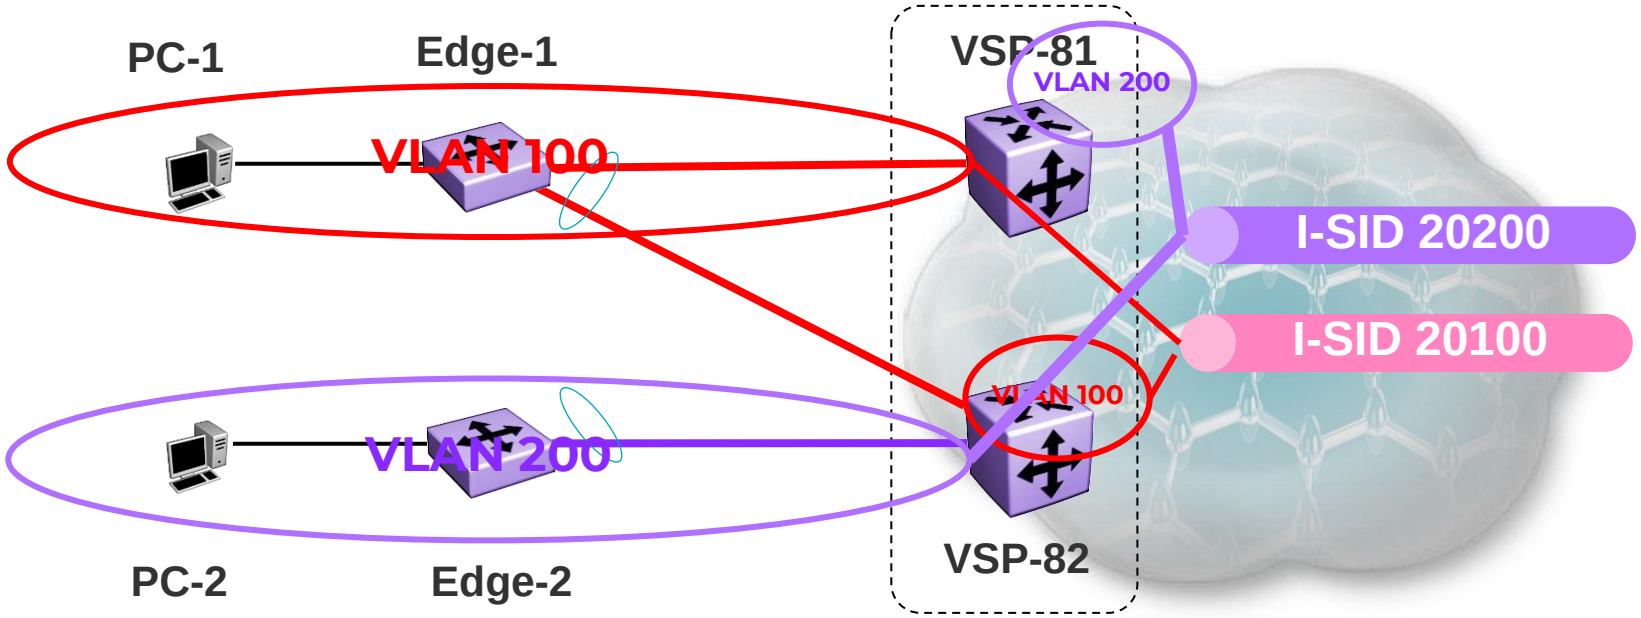

Virtual IST

Guidelines

Do not enable loop-detect on NNI ports:

- The system does not allow enabling loop-detect on an existing vIST port
- The system does allow adding a port with loop-detect enabled to a vIST
  - Such port can potentially cause system errors

When creating a vIST:

- All C-VLANs associated with it must have an I-SID
- Do not use the I-SID used for the vIST anywhere else in the network

Virtual IST

Best Practices

- Use a private address space with 30 bit mask for vIST VLAN IP’s
- Do NOT use the vIST IP addresses as the next hop address for any static routes
- Disable Spanning Tree – Uplinks and vIST Ports and NNI interfaces
- Disable Spanning tree – Uplink ports on Edge
- Enable FastStart (or MSTP Edge Ports) on all other ports on the Edge Switch
- Use Loop prevention features SLPP and SLPP Guard
- Use Link Health monitoring feature like VLACP

Example: SMLT – vIST Configuration

Hints

With the introduction of Zero Touch Fabric (ZTF; VOSS 8.3 and above) Switches could already show-up with SPBM and ISIS enabled, so no need to create BVLAN, assign system-ids etc.

Nick-names could be assigned dynamically from one (or more) VSP nodes, that have nick-name server functionality enabled.

ISIS area will be learned from neighbors, that run Fabric connect (on at least 8.3 release) dynamically.

Still it‘s good practise to modify ISIS system-id on nodes, that should work as switch cluster, as for RMA of switches the use of chassis MAC as system-id creates trouble.

- Note: when you do change the ISIS system-id you must also change the nick-name (if statically assigned)!

Example: SMLT – vIST Configuration : VSP-81

```
### BASE CONFIG for SWITCH CLUSTER
spbm
router isis
  system-id 02bb.0000.8100
  manual-area 49.0000
  spbm 1
  spbm 1 nick-name 0.00.81
  spbm 1 b-vid 4051,4052 primary 4051
  spbm 1 smlt-virtual-bmac 02:bb:00:00:81:ff
  spbm 1 smlt-peer-system-id 02bb.0000.8200
exit
vlan create 4051 name "B-VLAN-1" type spbm-
bvlan
vlan create 4052 name "B-VLAN-2" type spbm-
bvlan
interface GigabitEthernet 2/40
  isis
  isis spbm 1
  isis enable
  no spanning-tree mstp force-port-state enable
exit
```

```
### VIRTUAL IST configuration

vlan create 4000 name VIST type port-mstprstp 0
vlan i-sid 4000 24000
interface vlan 4000
  ip address 192.168.255.253/30
exit
virtual-ist peer-ip 192.168.255.254 vlan 4000

router isis enable
```

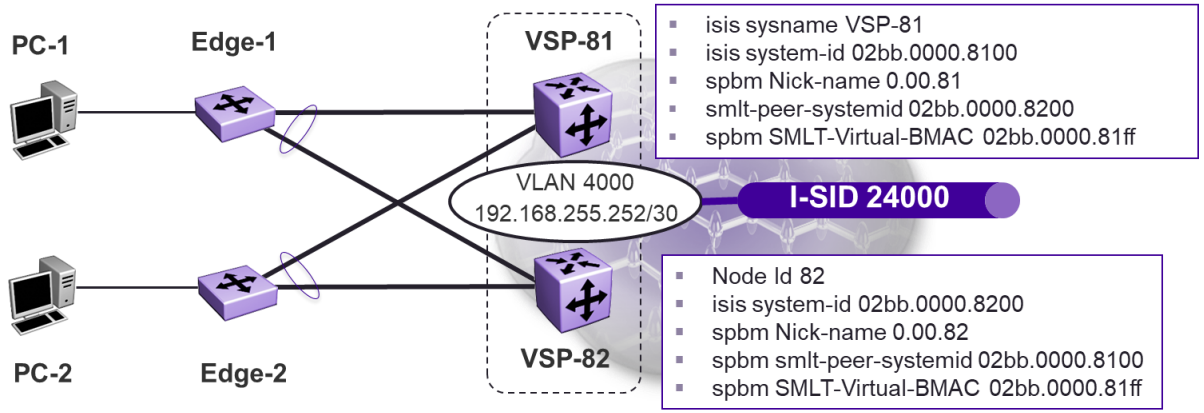

Example: SMLT – vIST Configuration : VSP-82

```
### BASE CONFIG for SWITCH CLUSTER
spbm
router isis
  system-id 02bb.0000.8200
  manual-area 49.0000
  spbm 1
  spbm 1 nick-name 0.00.82
  spbm 1 b-vid 4051,4052 primary 4051
  spbm 1 smlt-virtual-bmac 02:bb:00:00:81:ff
  spbm 1 smlt-peer-system-id 02bb.0000.8100
exit
vlan create 4051 name "B-VLAN-1" type spbm-
bvlan
vlan create 4052 name "B-VLAN-2" type spbm-
bvlan
interface GigabitEthernet 2/40
  isis
  isis spbm 1
  isis enable
  no spanning-tree mstp force-port-state enable
exit
```

```
### VIRTUAL IST configuration

vlan create 4000 name VIST type port-mstprstp 0
vlan i-sid 4000 24000
interface vlan 4000
  ip address 192.168.255.254/30
exit
virtual-ist peer-ip 192.168.255.253 vlan 4000

router isis enable
```

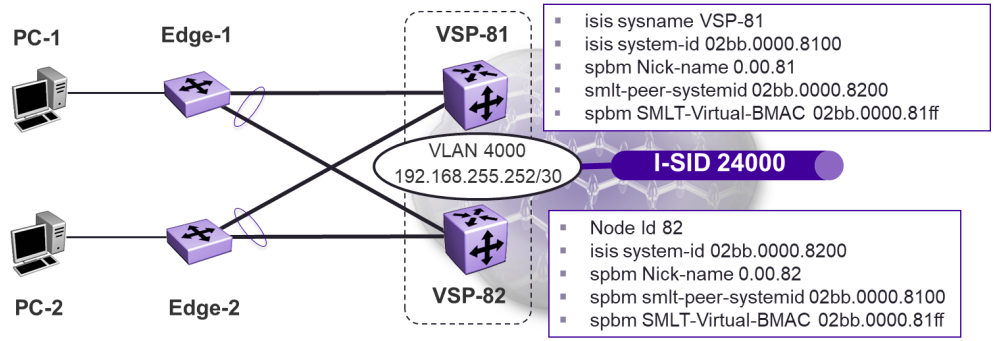

Example: C-VLAN Configuration on Switch Cluster

You have to create edge VLAN (User-1) on the switch itself

You cannot assign edge VLANs to IST link as there is no physical ports

You have to assign an I-SID to edge VLAN.

Note: this type of L2VSN is needed, if IP interface should be used on VLAN!

VSP8200-1 & VSP8200-2

```
vlan create 100 name "User-1" type port 1
vlan i-sid 100 20100
vlan member remove 1 1/1,2/1
vlan member add 100 1/1,2/1
exit
interface GigabitEthernet 1/1,2/1
  no shutdown
exit
```

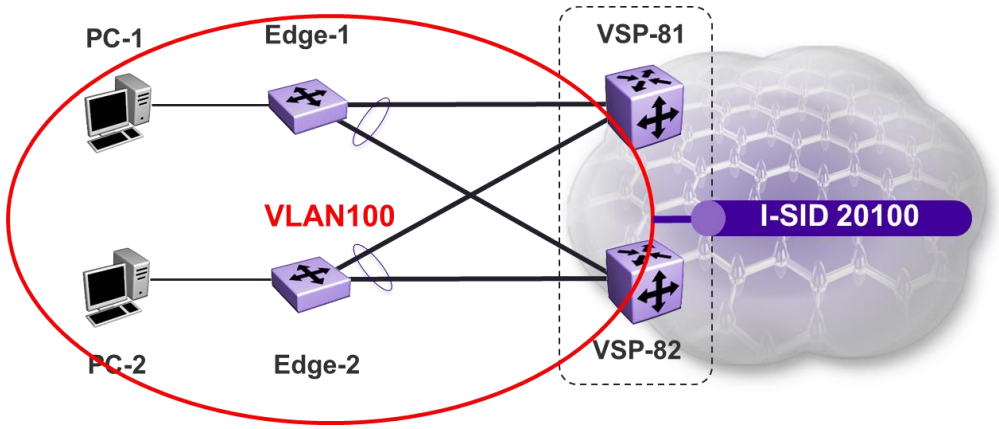

Example: Switched UNI L2VSN Configuration on Switch Cluster

For Switched UNI you don't need a VLAN on the switch itself

We will just assign VLAN-ID and Port to I-ISID(s).

Ports need to be set to flex-uni" mode to support this.

Note: Both types of L2VSN can share the same I-SID!

VSP8200-1 & VSP8200-2

```
vlan member remove 1 1/1,2/1
interface GigabitEthernet 1/1,2/1
  flex-uni enable
  no shutdown
Exit
i-sid 20100
  c-vid 100 port 1/1
  c-vid 100 port 2/1
  exit
```

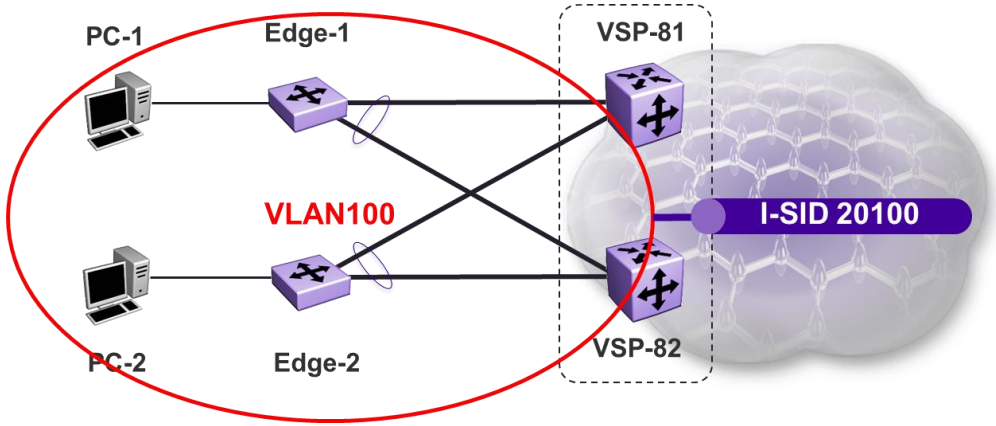

XMC view – Example with 2 Switch clusters & attached Access

V1 & V2 is a Cluster and V3 &V4 as well …

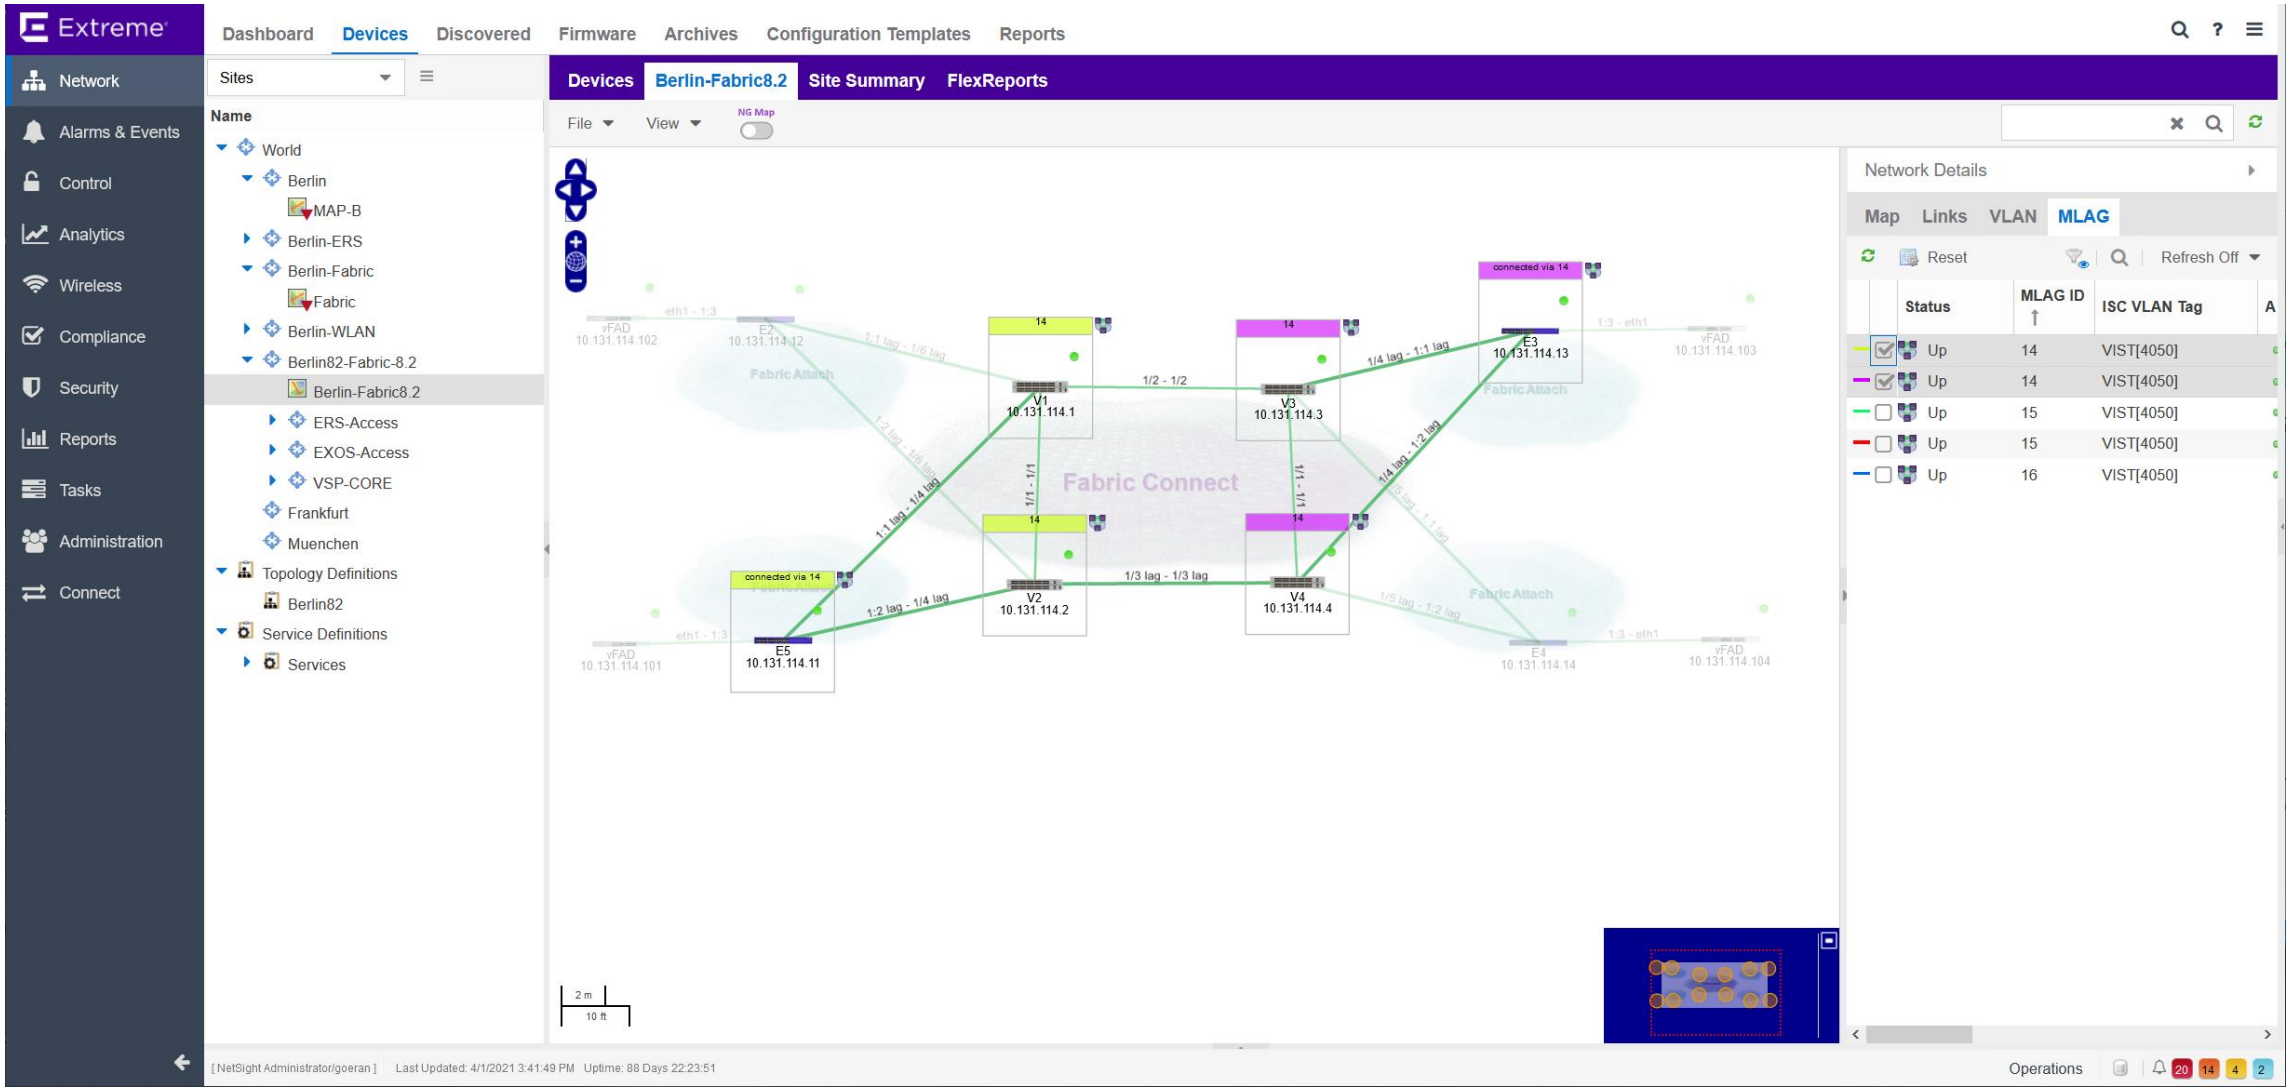

Provisioning – Virtual IST

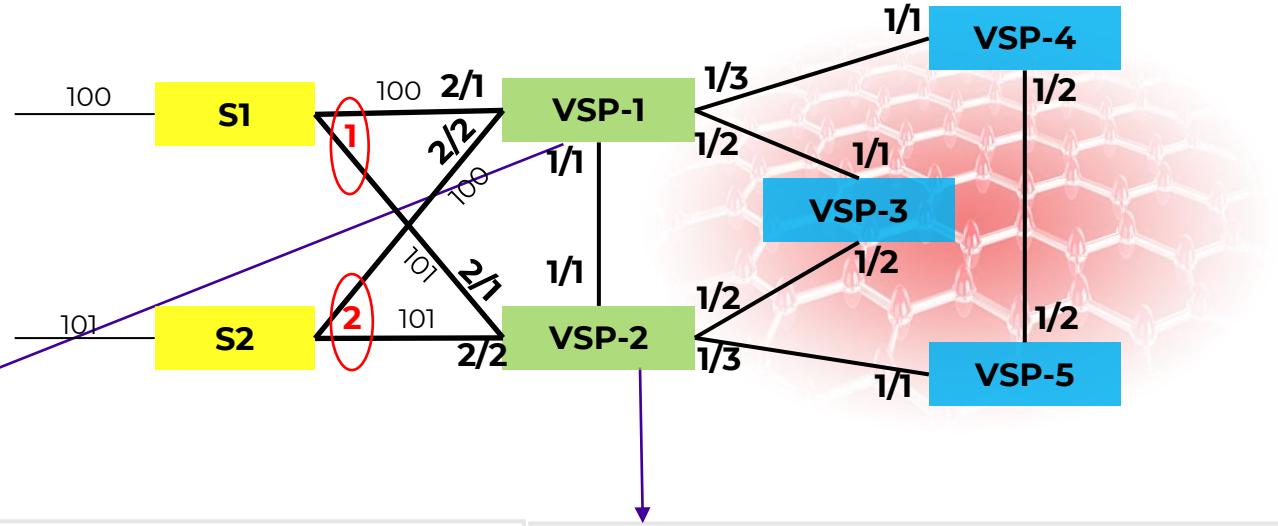


```
VSP-1:1# config terminal
prompt VSP-1
spbm
router isis
system-id 00be.b100.0000
spbm 1
spbm 1 nick-name 0.00.01
spbm 1 b-vid 4051,4052 primary 4051
spbm 1 smlt-peer-system-id 00be.b200.0000
spbm 1 smlt-virtual-bmac 00be.b012.0000
manual-area 33.0001
vlan create 4051 type spbm-bvlan
vlan create 4052 type spbm-bvlan
vlan create 2 type port-mstprstp 0
vlan i-sid 2 2
interface vlan 2
ip address 2.2.2.1/30
virtual-ist peer-ip 2.2.2.2 vlan 2
router isis enable
```

```
VSP-2:1# config terminal
prompt VSP-2
spbm
router isis
system-id 00be.b200.0000
spbm 1
spbm 1 nick-name 0.00.02
spbm 1 b-vid 4051,4052 primary 4051
spbm 1 smlt-peer-system-id 00be.b100.0000
spbm 1 smlt-virtual-bmac 00be.b012.0000
manual-area 33.0001
vlan create 4051 type spbm-bvlan
vlan create 4052 type spbm-bvlan
vlan create 2 type port-mstprstp 0
vlan i-sid 2 2
interface vlan 2
ip address 2.2.2.2/30
virtual-ist peer-ip 2.2.2.1 vlan 2
router isis enable
```

Provisioning – ISIS Interfaces

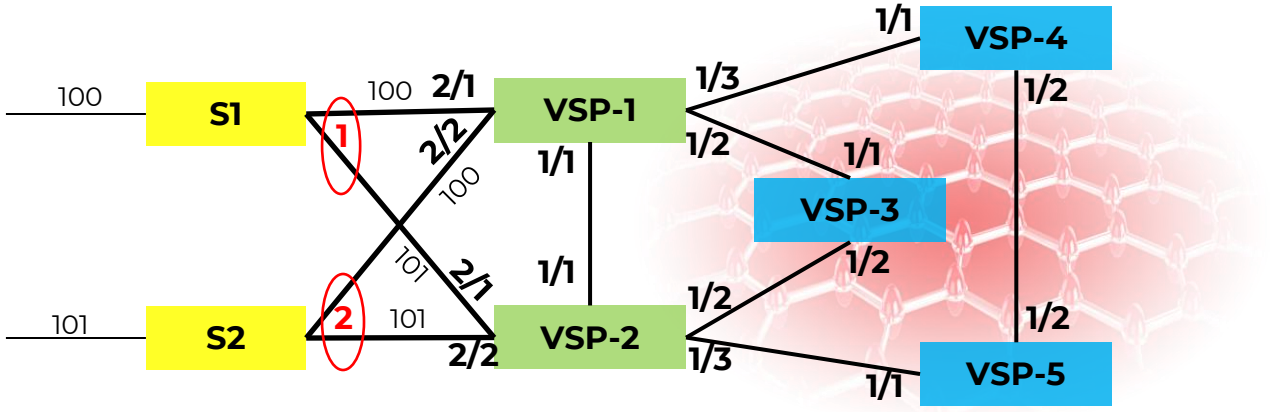

VSP-1 and VSP-2

```
interface gigabitethernet 1/1-1/3
encapsulation dot1q
no shutdown
isis
isis spbm 1
isis enable
```

Provisioning – SMLT

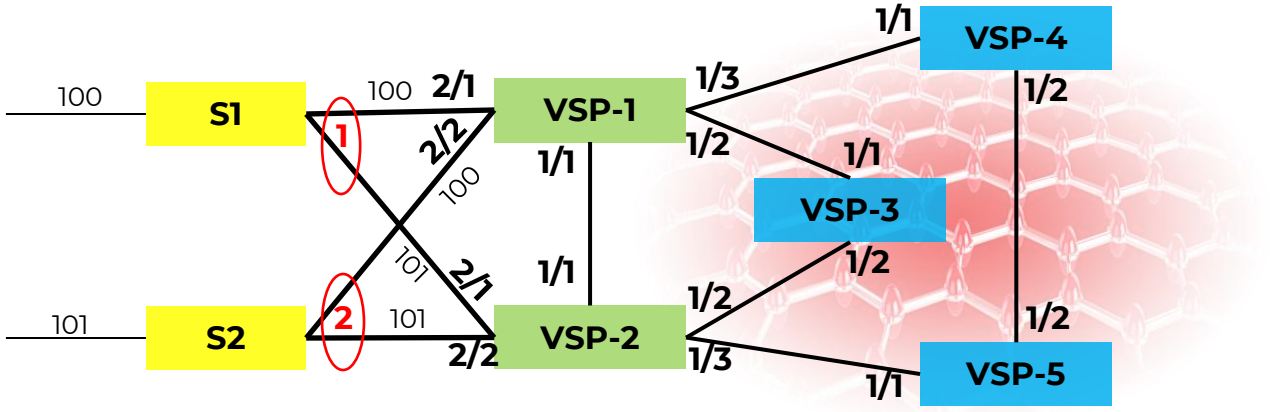

VSP-1 and VSP-2

```
mlt 1
mlt 1 member 2/1
mlt 1 encapsulation dot1q
interface mlt 1
smlt

mlt 2
mlt 2 member 2/2
mlt 2 encapsulation dot1q
interface mlt 2
smlt
```

Provisioning – VLAN, IP and RSMLT

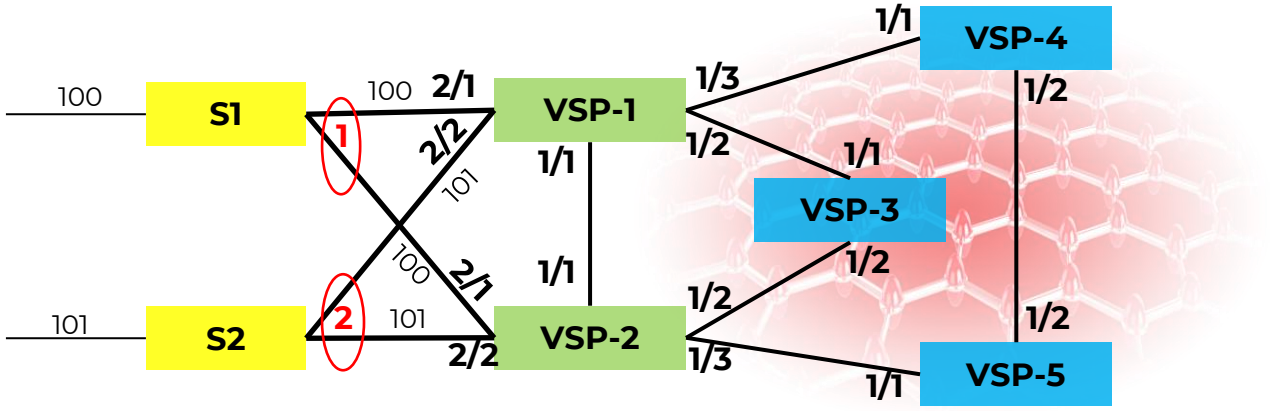

VSP-1

```
vlan create 100 type port-mstprstp 0
vlan members add 100 2/1
vlan i-sid 100 100
interface vlan 100
ip address 100.1.1.1/24
ip rsmlt
ip rsmlt holdup-timer 9999

vlan create 101 type port-mstprstp 0
vlan members add 100 2/2
vlan i-sid 101 101
interface vlan 101
ip address 101.1.1.1/24
ip rsmlt
ip rsmlt holdup-timer 9999

ip rsmlt edge-support
```

VSP-2

```
vlan create 100 type port-mstprstp 0
vlan members add 100 2/1
vlan i-sid 100 100
interface vlan 100
ip address 100.1.1.2/24
ip rsmlt
ip rsmlt holdup-timer 9999

vlan create 101 type port-mstprstp 0
vlan members add 100 2/2
vlan i-sid 101 101
interface vlan 101
ip address 101.1.1.2/24
ip rsmlt
ip rsmlt holdup-timer 9999

ip rsmlt edge-support
```

Default Gateway reduncancy - RSMLT

Procedure

1. Enter VLAN Interface Configuration mode:
   `enable`

   `configure terminal`

   `interface vlan <1-4059>`

2. Configure the holddown timer:
   `ip rsmlt holddown-timer <0-3600>`

3. Configure the holdup timer:
   `ip rsmlt holdup-timer <0-9999>`

4. Enable RSMLT on the VLAN:
   `ip rsmlt`

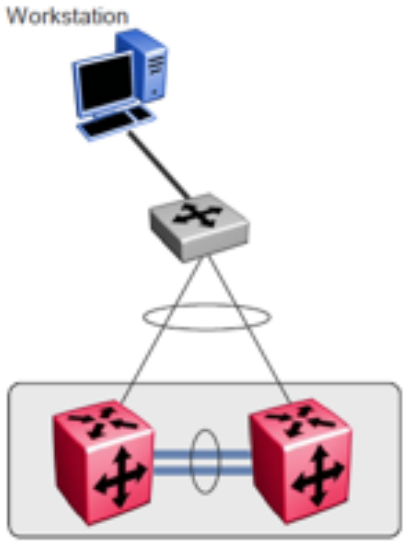


```
VSP4000-A:1(config)#show ip rsmlt edge-support

RSMLT Peer Info:
        rsmlt-peer-forwarding : enable

                        Peer Mac : b4:a9:5a:2d:5c:82
                              IP : 10.0.13.2
                            Vlan : 13

                        Peer Mac : b4:a9:5a:2d:5c:83
                              IP : 10.0.14.2
                            Vlan : 14
```

Procedure

1. Enter Global Configuration mode:
   `enable`

   `configure terminal`

2. Enable RSMLT Edge support:
   `ip rsmlt edge-support`

```
VSP4000-A:1(config)#show ip rsmlt
================================================================================
                    Ip Rsmlt Local Info - GlobalRouter
================================================================================
VID   IP              MAC               ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------
13    10.0.13.1       6c:a8:49:71:8b:82  Enable  Up    60     infinity
14    10.0.14.1       6c:a8:49:71:8b:83  Enable  Up    60     infinity

VID   SMLT ID
--------------------------------------------------------------------------------
13      1
14

VID   IPv6            MAC               ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------

VID   SMLT ID
--------------------------------------------------------------------------------


================================================================================
                    Ip Rsmlt Peer Info - GlobalRouter
================================================================================
VID   IP              MAC               ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------
13    10.0.13.2       b4:a9:5a:2d:5c:82  Enable  Up    60     infinity
14    10.0.14.2       b4:a9:5a:2d:5c:83  Enable  Up    60     infinity

VID   HDT REMAIN  HUT REMAIN  SMLT ID
--------------------------------------------------------------------------------
13    60          infinity      1
14    60          infinity
```

```
VSP4000-B:1(config)#show ip rsmlt
================================================================================
                    Ip Rsmlt Local Info - GlobalRouter
================================================================================
VID   IP              MAC               ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------
13    10.0.13.2       b4:a9:5a:2d:5c:82  Enable  Up    60     infinity
14    10.0.14.2       b4:a9:5a:2d:5c:83  Enable  Up    60     infinity

VID   SMLT ID
--------------------------------------------------------------------------------
13      1
14

VID   IPv6            MAC               ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------

VID   SMLT ID
--------------------------------------------------------------------------------


================================================================================
                    Ip Rsmlt Peer Info - GlobalRouter
================================================================================
VID   IP              MAC               ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------
13    10.0.13.1       6c:a8:49:71:8b:82  Enable  Up    60     infinity
14    10.0.14.1       6c:a8:49:71:8b:83  Enable  Up    60     infinity

VID   HDT REMAIN  HUT REMAIN  SMLT ID
--------------------------------------------------------------------------------
13    60          infinity      1
14    60          infinity
```

Provisioning – VLAN, IP and VRRP

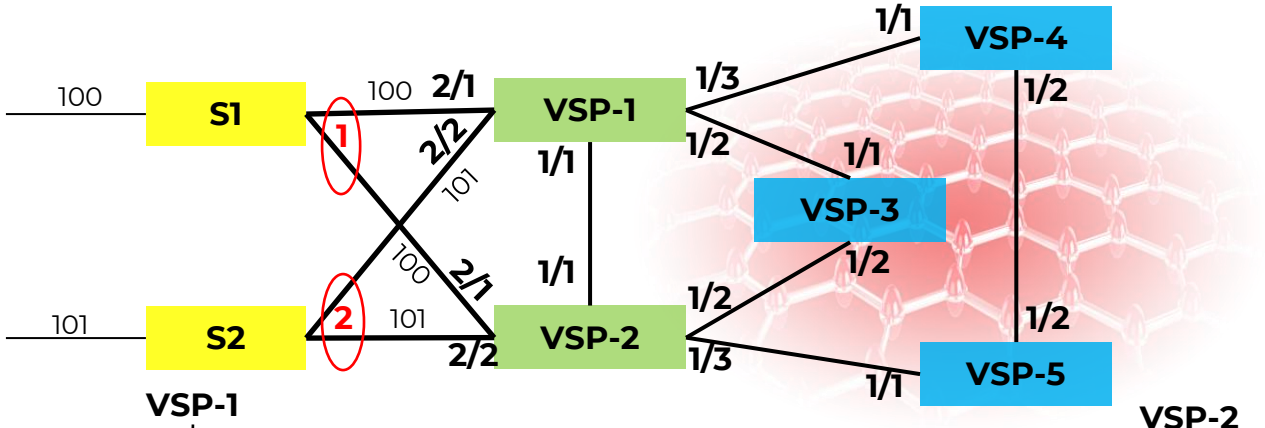

VSP-1

```
vlan create 100 type port-mstprstp 0
vlan members add 100 2/1
vlan i-sid 100 100
interface vlan 100
ip address 100.1.1.1/24
Ip vrrp version 3
ip vrrp address 100 100.1.1.254
ip vrrp 100 backup-master enable
ip vrrp 100 priority 101
ip vrrp 100 enable

vlan create 101 type port-mstprstp 0
vlan members add 100 2/2
vlan i-sid 101 101
interface vlan 101
ip address 101.1.1.1/24
ip vrrp version 3
ip vrrp address 101 101.1.1.254
ip vrrp 101 backup-master enable
ip vrrp 101 priority 101
ip vrrp 101 enable
```

VSP-2

```
vlan create 100 type port-mstprstp 0
vlan members add 100 2/1
vlan i-sid 100 100
interface vlan 100
ip address 100.1.1.2/24
ip vrrp version 3
ip vrrp address 100 100.1.1.254
ip vrrp 100 backup-master enable
ip vrrp 100 priority 100
ip vrrp 100 enable

vlan create 101 type port-mstprstp 0
vlan members add 100 2/2
vlan i-sid 101 101
interface vlan 101
ip address 101.1.1.2/24
ip vrrp version 3
ip vrrp address 101 101.1.1.254
ip vrrp 101 backup-master enable
ip vrrp 101 priority 100
ip vrrp 101 enable
```

Default Gateway reduncancy - VRRP

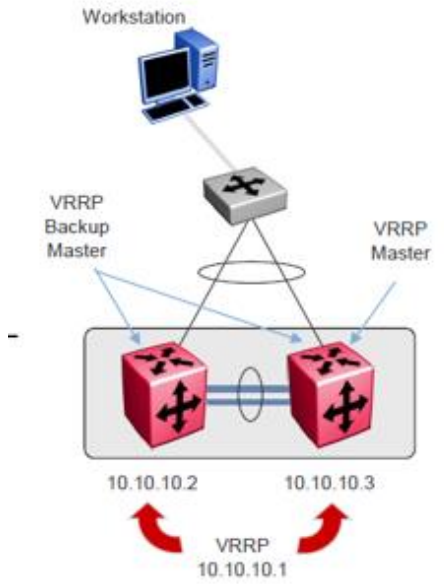

Procedure

1. Enter GigabitEthernet Interface Configuration mode:
   ```
   enable

   configure terminal

   interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
   ```

   NOTE

   If the platform supports channelization and the port is channelized, you must also specify the su
   slot/port/sub-port.

2. Configure a backup VRRP address:
   ```
   ip vrrp address <1-255> <A.B.C.D>
   ```

3. Configure VRRP on a port:
   ```
   ip vrrp <1-255> enable
   ```

4. Show the global VRRP configuration:
   ```
   show ip vrrp
   ```

```
VSP-4450GSX-PWR+:1(config-if)#show ip vrrp address

================================================================================
                          VRRP Info - GlobalRouter
================================================================================

VRRP ID  P/V      IP              MAC               STATE    CONTROL  PRIO  ADV VERSION
--------------------------------------------------------------------------------
101      101      101.1.1.254     00:00:5e:00:01:65 Master   Enabled  101   1   3

1 out of 1 Total Num of VRRP Address Entries displayed.


VRRP ID  P/V      MASTER          UP TIME              HLD DWN  CRITICAL IP(ENABLED)  VERSION
--------------------------------------------------------------------------------
101      101      101.1.1.1       0 day(s), 00:00:30   0        0.0.0.0         (No)  3

1 out of 1 Total Num of VRRP Address Entries displayed.


VRRP ID  P/V      BACKUP MASTER   BACKUP MASTER STATE   FAST ADV (ENABLED)     VERSION
--------------------------------------------------------------------------------
101      101      enable          down                  200       (NO)         3

1 out of 1 Total Num of VRRP Address Entries displayed.
```

```
VSP-4850GTS-PWR+:1(config-if)#show ip vrrp address

================================================================================
                          VRRP Info - GlobalRouter
================================================================================

VRRP ID  P/V      IP              MAC               STATE    CONTROL  PRIO  ADV VERSION
--------------------------------------------------------------------------------
101      101      101.1.1.254     00:00:5e:00:01:65 Backup   Enabled  100   1   3

1 out of 1 Total Num of VRRP Address Entries displayed.


VRRP ID  P/V      MASTER          UP TIME              HLD DWN  CRITICAL IP(ENABLED)  VERSION
--------------------------------------------------------------------------------
101      101      101.1.1.1       0 day(s), 00:00:17   0        0.0.0.0         (No)  3

1 out of 1 Total Num of VRRP Address Entries displayed.


VRRP ID  P/V      BACKUP MASTER   BACKUP MASTER STATE   FAST ADV (ENABLED)     VERSION
--------------------------------------------------------------------------------
101      101      enable          up                    200       (NO)         3

1 out of 1 Total Num of VRRP Address Entries displayed.
```

vIST – Verification

Verify vIST Operation and SMLT Peer Info

show virtual-ist → State should be up

show isis spbm → Verifty SMLT Peer info

show isis adjacencies → two adjacencies up

show isis interface → two interfaces up

show isis lsdb → more entries (neighbors)

show isis spbm ? [For Additional Commands]

On the VSP clustering- Verify MLT & SMLT State

- show mlt → Note MLT Admin vs. MLT Current
- show smlt mlt → Note Admin vs. Current

**The current state will show 'norm' if the other side of the link is down, hasn't yet been configured or is misconfigured.*

Simplified Virtual IST

Only works in NON-SPBM Mode!!

Simplified Virtual IST - Introduction

All current VSP series support a **simplified Virtual-IST** configuration in order to enable seamless migration of legacy IST based SMLT to Virtual-IST based SMLT

This feature can be enabled by first **disabling SPBM** and then enabling simplified virtual-IST at the Interface MLT level

Disabling SPBm is done through a boot flag “spbm-config-mode”

The “spbm-config-mode” makes SPBm and PIM functionality mutually exclusive as shown in the next slide

```
VSP-8284XSQ:1(config)# boot config flags spbm-config-mode
```

```
VSP-8284XSQ:1(config)# no boot config flags spbm-config-mode
```

Boot Flag: spbm-config-mode

Differentiates SPBm/Non-SPBm deployments

Feature availability dependent on spbm-config-mode boot flag

Feature availability	spbm-config-mode Enabled	spbm-config-mode Disabled
SPBm Provisioning	✔	✘
CFM Provisioning (SPBm-BVLAN)	✔	✘
IGMP v1/v2/v3	✔	✔
Multicast over SPBm	✔	✘
PIM-SM, PIM-SSM	✘	✔
Simplified vIST configuration	✘	✔
All other features	✔	✔

Simplified Virtual IST Configuration

Simplified Virtual IST does require SPB functionality

However all SPB parameters are auto-configured,

- MLT requires creation.
- No explicit SPB configurations are required
- SPBm and IS-IS show commands work in this mode to be used for debugging

A new command as shown below is introduced under the interface-MLT

Enabled by disabling SPBm and then enabling simplified vIST on the MLT

```
(config)#no boot config flags spbm-config-mode
(config)#interface mlt <mlt id>
(config-if)#virtual-ist enable
```

Reminder:
Simplified vIST is available ONLY for legacy multicast deployments when the boot flag (spbm-config-mode) is disabled.

Simplified Virtual IST Configuration Example

```
# BOOT CONFIGURATION
#
no boot config flags spbm-config-mode

# CLI CONFIGURATION
prompt "VSP-1"

mlt 2 enable name "IST-MLT"
mlt 2 member 1/10-1/11
mlt 2 encapsulation dot1q

vlan create 2 name "IST-VLAN" type port-mstprstp 1
interface vlan 2
ip address 192.168.102.253 255.255.255.252 0

# MLT CONFIGURATION
interface mlt 2
virtual-ist enable

# VIRTUAL IST CONFIGURATION
virtual-ist peer-ip 192.168.102.254 vlan 2
```

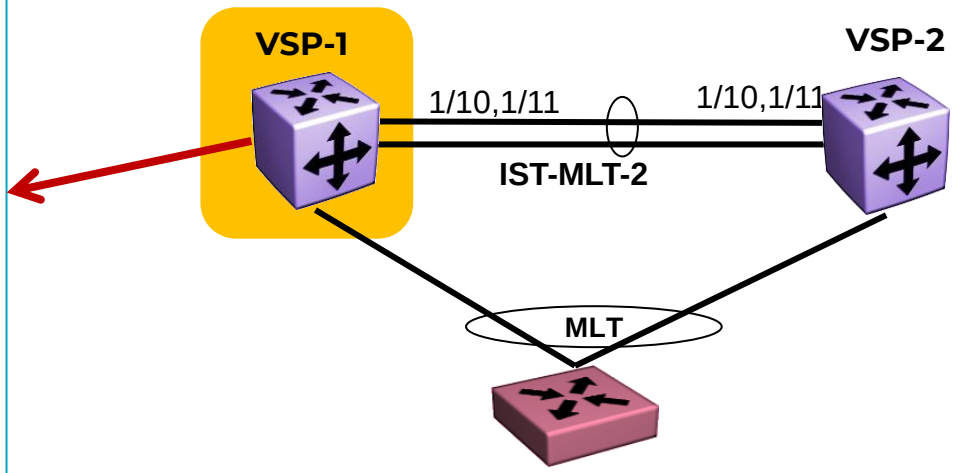

Loop Prevention and Link Health

SLPP
SLPP Guard
VLACP

Simple Loop Prevention Protocol (SLPP)

Prevents loops in a Switch Cluster network

- Loops can occur when:
  - MLT at the edge is misconfigured
  - MLT not created at the edge but links are plugged in anyway
  - MLT configuration is lost (switch set back to factory default)

SLPP uses an SLPP-PDU, which is generated by the Switch Cluster cores

- Loop detection is achieved by detecting whether the SLPP-PDU is received on the IST peer switch port or on the same switch where it originated
- If a self or SMLT peer originated SLPP PDU packet is received:
  - The port is taken down (if the packet is received on the same VLAN it originated on)
  - A log file entry is generated
  - An SNMP trap is sent
  - Once the port is down, it will stay down and needs manual intervention to be enabled

Simple Loop Prevention Protocol (SLPP)

SLPP-PDU when enabling SLPP on a VLAN

- The packet is constrained to the VLAN on which it was sent

SLPP-PDU receiving/processing only on ports where SLPP-Rx is enabled

If SLPP-PDU receiving works on a port, which is an MLT member, all port members in that MLT are taken down

The SLPP-PDU can be received by the originated switch or its IST peer

- All other switches treat the SLPP-PDU as normal multicast packet and will forward it on the VLAN

SLPP threshold based on the sum of all packets received

Port-based VLANs only

Supported on:

- VOSS (VSP), BOSS (ERS)

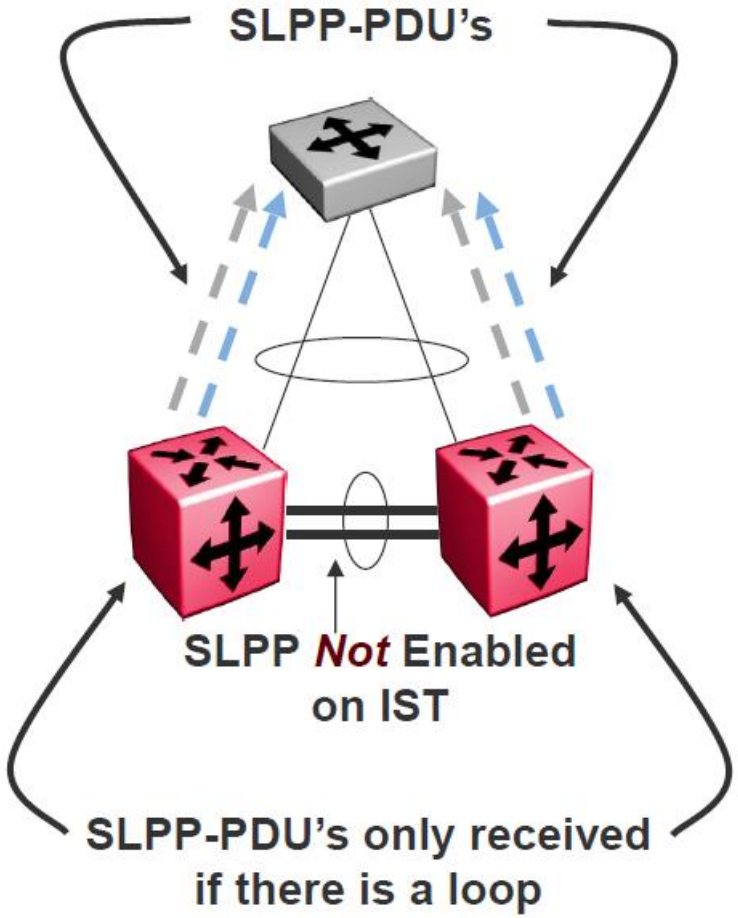

SLPP and Switch Clustering - Implementation

Enable SLPP
- Per VLAN
- Per port by setting RxThreshold

Identify one Switch peer as Primary and the other as Secondary
- Not a configurable option, strictly from a design standpoint
- Enable RxThreshold per table below on uplink ports

Do not enable auto recovery
- Once the port is down, it will stay in down state
- Port needs manual intervention to be enabled

Do not enable SLPP-Rx on IST ports
- Never want to take these ports down

Increase secondary if more VLANs are set up

Edge Switch	SLPP Holddown timer
All non-core ports (Edge only)	0

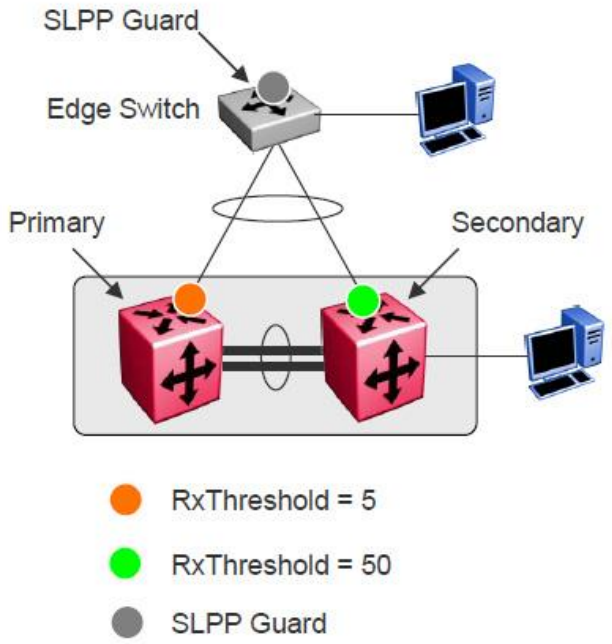

SMLT Cluster Switch	EtherType	Packet Rx Threshold	Transmission Interval
Primary	Default	5	Default (.5 seconds)
Secondary		50	

SLPP and Switch Clustering

Implementation

SLPP RxThreshold is a cumulative count.

A sequence of non-related events could lead to taking ports down.

A disable/enable of SLPP should be performed after any SLPP event to clear the counter.

- VSP switches will automatically re-arm the SLPP counters every 6 hours.

Sometimes, the secondary switch also detects the loop and reaches its SLPP RxThreshold before the primary takes the port down to stop the loop.

- Therefore, both switches take their ports down and the edge becomes isolated.

The larger the number of VLANs associated with the port, the more likely this could occur, especially for loop conditions that affect all VLANs.

The recommended step here is to increase the RxThreshold.

- Increase the RxThreshold on the primary using a multiplier of 5 for each VLAN.
- Increase this number on the secondary using a multiplier of 10.
- For example, for 5 VLANs, use SLPP RxThreshold values of 25 and 250.

SLPP Guard

In most environments there is a need for additional loop protection when used in combination with a Switch Cluster (SMLT)

SLPP Guard helps prevent loops in customers' networks by administratively disabling an edge port if they receive an SLPP packet

Loop prevention for edge ports can be provided by enabling STP on edge

- Provides protection for looping back edge ports to the same switch or stack
- STP loop prevention will not work if the attached device does not support STP

Due to moves, adds or changes, it is possible to create a loop by connecting an edge port back to a port in the switch cluster

SLPP Guard will disable a port when an SLPP packet is received on a port

It will generate a local log message, and can generate a syslog message and SNMP traps

SLPP Guard Continued

Each port has its own administrative hold-down timer

- If the port is shutdown due to reception of a SLPP packet the timer starts for that port
- When the timer reaches the configured interval, the port is re-enabled and a local log message, syslog message, and SNMP traps are generated

This timer is user configurable between 10 and 65,535 seconds

- 60 seconds is the default
- if set to 0, the port will be administratively disabled and must be manually enabled again

The default SLPP EtherType is (hex): 0x8102

- For some switches, it has used an old value of 0x8104
- You can globally configure the EtherType for SLPP guard

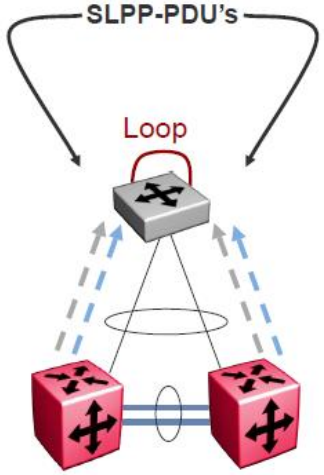

SLPP and SLPP Guard Configuration Example

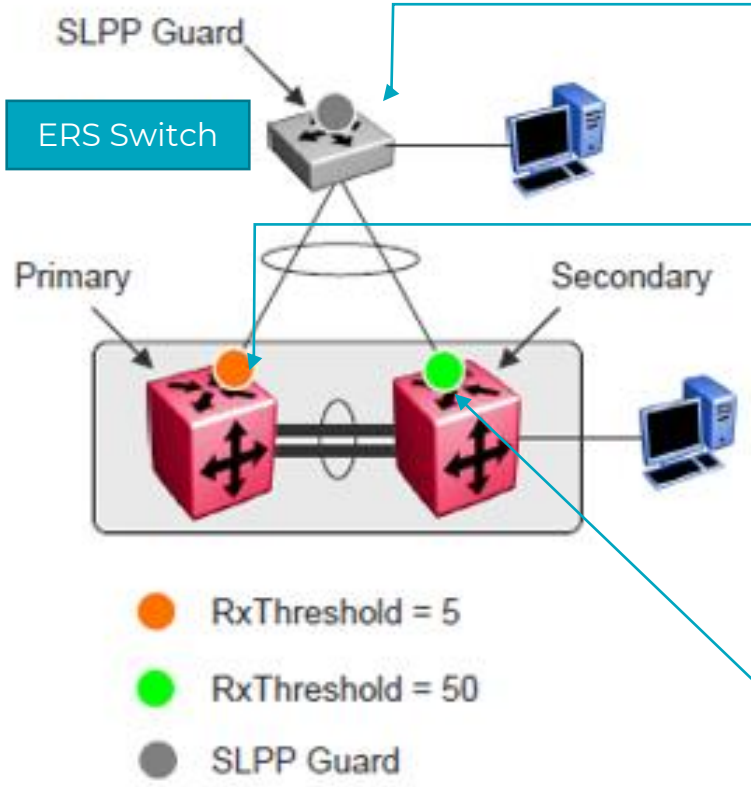


```
Interface ethernet 1-48
slpp-guard enable
slpp-guard timeout 0
```

```
Slpp enable
Slpp vid 10
Interface gigabitEthernet 1/1
slpp packet-rx
slpp packet-rx-threshold 5
```

```
Slpp enable
Slpp vid 10
Interface gigabitEthernet 1/1
slpp packet-rx
slpp packet-rx-threshold 50
```

Configuring Simple Loop Prevention Protocol (SLPP) Guard on EXOS

https://documentation.extremenetworks.com/exos_31.4/GUID-54903026-AF6E-4C79-BB2D-C85D6F34E9CC.shtml?_ga=2.50684650.2104166867.1633934435-1003111500.1632184941

Configuring Simple Loop Prevention Protocol (SLPP) Guard

For Simple Loop Prevention Protocol (SLPP) Guard to operate, SLPP must have already been configured in the core of the network.
The following information explains how to configure SLPP Guard on ExtremeXOS switches.

1. Enable SLPP on the desired ports using the following command:
   `enable slpp guard ports [port_list | all]`

   NOTE

   To view the SLPP status of ports, use the `show slpp guard {ports port_list} {disabled-ports}` command.

2. (Optional) Set the recovery timeout period using the following command:
   `configure slpp guard [ports [port_list | all] recovery-timeout [seconds | none]`
   After the configured timeout value set by this command expires (associated with each port), the port is automatically re-enabled.

3. (Optional) Configure the Ethertype for SLPP Guard using the following command:
   `configure slpp guard ethertype hex`
   This command configures the Ethertype that the SLPP Guard feature uses to identify SLPP PDUs.

If you need to disable SLPP Guard on a port, use the following command:

```
disable slpp guard ports [port_list | all]
```

To view SLPP Guard status for selected ports or all SLPP Guard-disabled ports, use the following command:

```
show slpp guard {ports port_list} {disabled-ports}
```

Configuring SLPP Guard on VOSS

https://documentation.extremenetworks.com/VOSS/SW/83/VOSSUserGuide/GUID-FD448A8E-046B-4848-B066-C6753339A021.shtml

Procedure

1. Enter GigabitEthernet Interface Configuration mode:
   ```
   enable

   configure terminal

   interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
   ```

   NOTE

   If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Enable SLPP Guard on the port:
   ```
   slpp-guard enable
   ```

3. **Optional:** Configure the timeout value on the port:
   ```
   slpp-guard timeout {<10-65535> | 0}
   ```

Example

Copy

```
Switch:1(config-if)#slpp-guard enable
Switch:1(config-if)#slpp-guard timeout 120
```

Virtual Link Aggregation Control Protocol (VLACP)

Detects end-to-end failure by propagating link status between ports

- Directly connected point-to-point
- Logically connected point-to-point across an intermediate network

Can detect:

- Complete link failure
- Receive or transmit link disruptions (one-way link failure)
- Transmits VLACPDU every “x” milliseconds so both ends of link maintain state

Based on LACP and is Intellectual Property of Extreme Networks

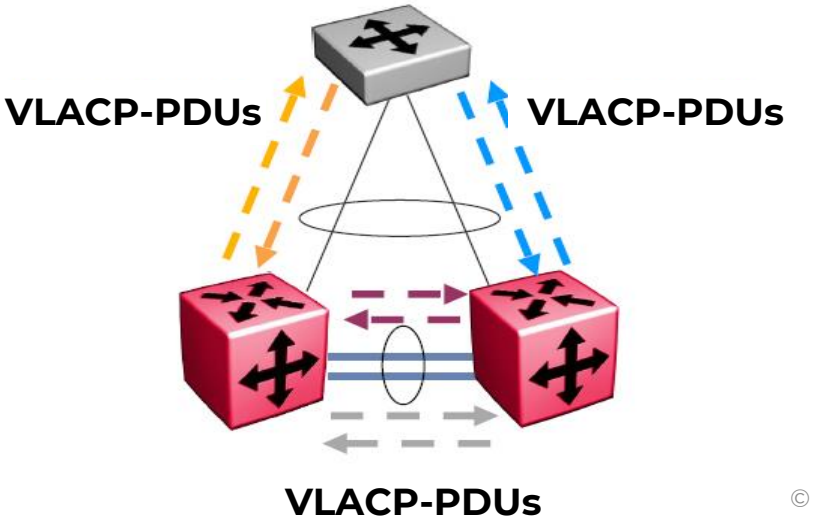

Virtual Link Aggregation Control Protocol (VLACP)

Considerations

If a VLACP-enabled port does not receive a VLACP Data Unit (VLACPDU), it enters the disabled state.

If VLACPDUs are not received on a particular link, that link is taken down after the expiry timeout occurs (timeout scale x periodic time).

There are occasions where a VLACP-enabled port does not receive a VLACPDU but remains in the forwarding state.

To avoid this, ensure that VLACP configurations at the port level are consistent:

- Either enable or disable ***both sides*** of the point-to-point connection.
- configure VLACP on each port.
- The port can be either an individual port or an MLT member.

VLACPDUs can be sent periodically on each port where VLACP is enabled to exchange VLACPDUs from an end-to-end perspective.

Virtual Link Aggregation Control Protocol (VLACP)

Best Practices

Enable VLACP globally and on each individual uplink and IST port:

- Both ends must have matching multicast MAC, EtherType, and Timers
- Do not enable VLACP and LACP on the same links
- Do not enable VLACP on IST port members on VSP 7000

For directly connected point-to-point links:

- Use reserved multicast MAC 01-80-c2-00-00-0f
- Ensures packet is not flooded across a defaulted switch

For end-to-end connections traversing intermediate networks:

- Use default MAC 01:80:c2:00:11:00

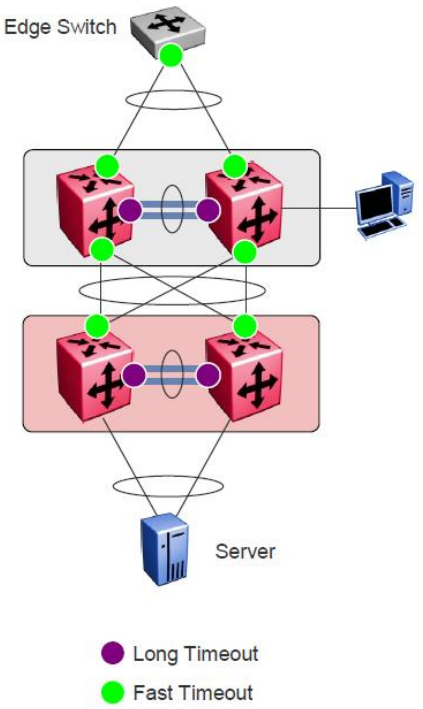

Connection type	Fast timer	Slow timer	Timeout	Timeoutscale
Uplink	500 ms	N/A	Short	5
IST	N/A	10000	Long	3
Short timeout = Timeout scale * Fast Periodic Timer Long timeout = Timeout scale * Slow Periodic Timer				

Virtual Link Aggregation Control Protocol (VLACP)

CLI commands

Option 1: use custom settings

- vlacp ethertype <1536-65535 | 0x600-0xffff>
- vlacp fast-periodic-time <100-20000>
- vlacp funcmac-addr 0x00:0x00:0x00:0x00:0x00:0x00
- vlacp slow-periodic-time <10000-30000>
- vlacp timeout long
- vlacp timeout short
- vlacp timeout-scale <2-10>

Option 2: use Default VLACP Settings

- default vlacp
- default vlacp ethertype
- default vlacp fast-periodic-time
- default vlacp funcmac-addr
- default vlacp slow-periodic-time
- default vlacp timeout
- default vlacp timeout-scale

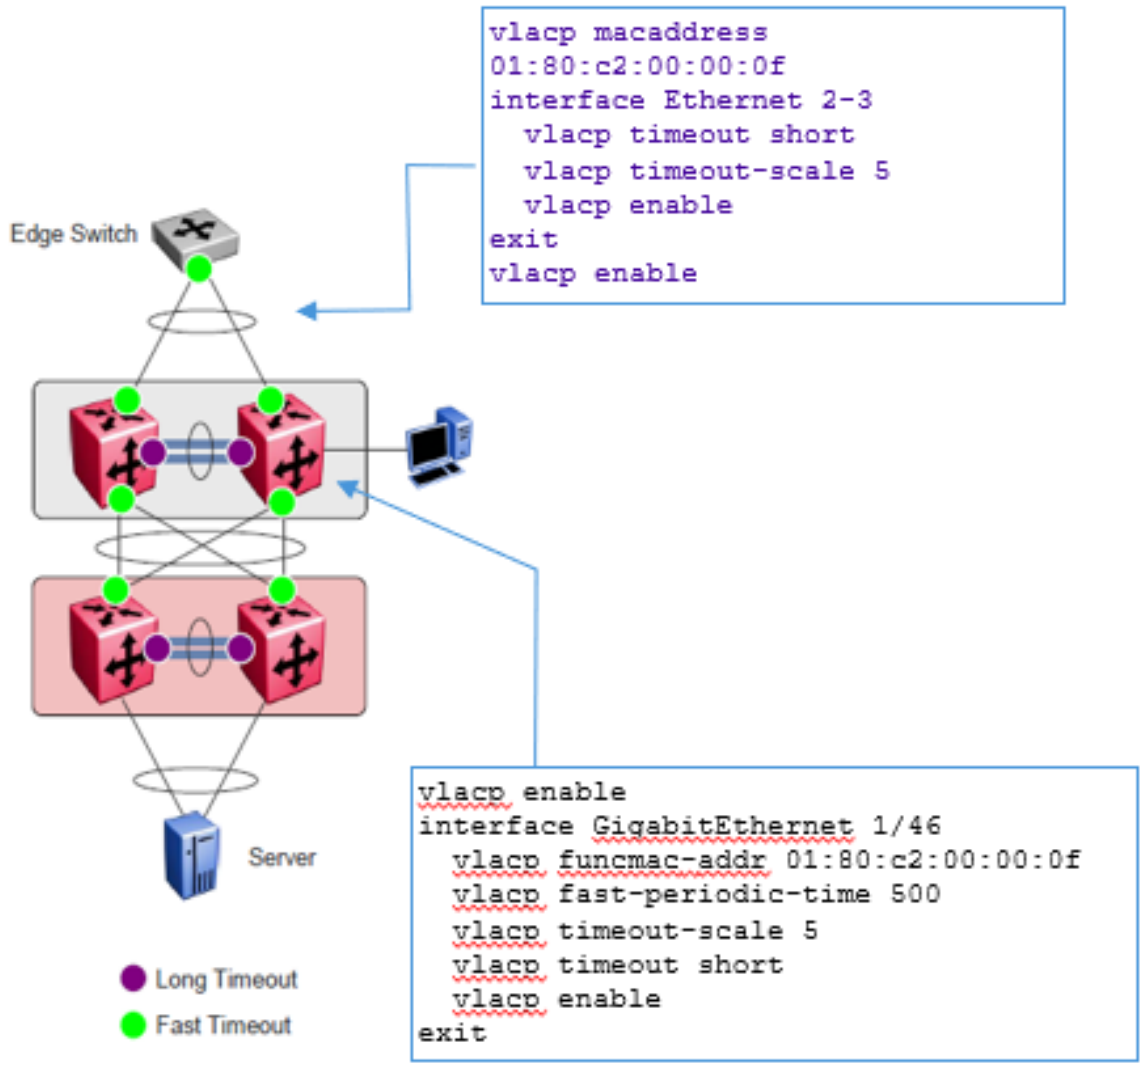

VLACP Configuration Example on VOSS

https://documentation.extremenetworks.com/VOSS/SW/83/VOSSUserGuide/GUID-00A7639F-441D-4501-BEFF-B122C09037BA.shtml

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable

configure terminal

interface GigabitEthernet {slot/port[/sub-port][-slot/port[/sub-port]][,...]}
```

NOTE

If the platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

2. Configure optional parameters for the port. If you do not configure these parameters, the system uses the default values.

   a. Configure the protocol identification for the port:

```
vlacp ethertype <1536-65535 | 0x600-0xffff> [funcmac-addr 0x00:0x00:0x00:0x00:0x00:0x00]
```

   b. Configure the fast or slow periodic times:

```
vlacp fast-periodic-time <100-20000> | slow-periodic-time <10000-30000>
```

   You can configure both parameters in the same command entry.

   c. Configure the timeout parameters:

```
vlacp timeout <long|short> timeout-scale <2-10>
```

   You can configure both parameters in the same command entry.

3. Enable VLACP on a port:

```
vlacp enable
```

Example

Configure VLACP on port 1/1:

```
Switch:1# configure terminal

Switch:1# interface GigabitEthernet 1/2

Switch:1# vlacp fast-periodic-time 400 timeout short

Switch:1# vlacp enable
```

WWW.EXTREMENETWORKS.COM