Ethernet Routing Switch 4800, 5900, 8800

Virtual Services Platform 4000, 4900, 7000, 7200, 7400, 8000, 9000

5520 Series, 5420 Series

Engineering

> Shortest Path Bridging (802.1aq) Technical Configuration Guide

Extreme Networks

Document Date: October 2021

Part Number: 9036665-00

Revision: AB

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Extreme Networks, Inc. assumes no liability for any errors. Extreme Networks, Inc. reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

“Documentation” means information published in varying mediums which may include product information, operating instructions and performance specifications that are generally made available to users of products. Documentation does not include marketing materials. Extreme Networks shall not be responsible for any modifications, additions, or deletions to the original published version of Documentation unless such modifications, additions, or deletions were performed by or on the express behalf of Extreme Networks. End User agrees to indemnify and hold harmless Extreme Networks, Extreme Networks’ agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Extreme Networks is not responsible for the contents or reliability of any linked websites referenced within this site or Documentation provided by Extreme Networks. Extreme Networks is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Extreme Networks does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Extreme Networks provides a limited warranty on Extreme Networks hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Extreme Networks’ standard warranty language, as well as information regarding support for this product while under warranty is available to Extreme Networks customers and other parties through the Extreme Networks Support website: http://www.extremenetworks.com/support under the link ““Policies” or such successor site as designated by Extreme Networks. Please note that if You acquired the product(s) from an authorized Extreme Networks Channel Partner outside of the United States and Canada, the warranty is provided to You by said Extreme Networks Channel Partner and not by Extreme Networks.

“Hosted Service” means an Extreme Networks hosted service subscription that You acquire from either Extreme Networks or an authorized Extreme Networks Channel Partner (as applicable) and which is described further in Hosted SAS or other service description documentation regarding the applicable hosted service. If You purchase a Hosted Service subscription, the foregoing limited warranty may not apply but You may be entitled to support services in connection with the Hosted Service as described further in your service description documents for the applicable Hosted Service. Contact Extreme Networks or Extreme Networks Channel Partner (as applicable) for more information.

Hosted Service

THE FOLLOWING APPLIES ONLY IF YOU PURCHASE AN EXTREME NETWORKS HOSTED SERVICE SUBSCRIPTION FROM EXTREME NETWORKS OR AN EXTREME NETWORKS CHANNEL

PARTNER (AS APPLICABLE), THE TERMS OF USE FOR HOSTED SERVICES ARE AVAILABLE ON THE EXTREME NETWORKS WEBSITE, https://extremeportal.force.com OR SUCH SUCCESSOR SITE AS DESIGNATED BY EXTREME NETWORKS, AND ARE APPLICABLE TO ANYONE WHO ACCESSES OR USES THE HOSTED SERVICE. BY ACCESSING OR USING THE HOSTED SERVICE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE DOING SO (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THE TERMS OF USE. IF YOU ARE ACCEPTING THE TERMS OF USE ON BEHALF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE AUTHORITY TO BIND SUCH ENTITY TO THESE TERMS OF USE. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO ACCEPT THESE TERMS OF USE, YOU MUST NOT ACCESS OR USE THE HOSTED SERVICE OR AUTHORIZE ANYONE TO ACCESS OR USE THE HOSTED SERVICE.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE EXTREME NETWORKS WEBSITE, https://extremeportal.force.com OR SUCH SUCCESSOR SITE AS DESIGNATED BY EXTREME NETWORKS, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS EXTREME NETWORKS SOFTWARE, PURCHASED FROM EXTREME NETWORKS, INC., ANY EXTREME NETWORKS AFFILIATE, OR AN EXTREME NETWORKS CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH EXTREME NETWORKS OR AN EXTREME NETWORKS CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY EXTREME NETWORKS IN WRITING, EXTREME NETWORKS DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN EXTREME NETWORKS, AN EXTREME NETWORKS AFFILIATE OR AN EXTREME NETWORKS CHANNEL PARTNER; EXTREME NETWORKS RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND EXTREME NETWORKS, INC. OR THE APPLICABLE EXTREME NETWORKS AFFILIATE ("EXTREME NETWORKS").

Extreme Networks grants You a license within the scope of the license types described below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License as set forth below in the Designated System(s) License (DS) section as applicable. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. "Software" means computer programs in object code, provided by Extreme Networks or an Extreme Networks Channel Partner, whether as stand-alone products, pre- installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. "Designated Processor" means a single stand-alone computing device. "Server" means a set of Designated Processors that hosts (physically or virtually) a software application to be accessed by multiple users. "Instance" means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine ("VM") or similar deployment.

License type(s)

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only: 1) on a number of Designated Processors up to the number indicated in the order; or 2) up to the number of Instances of the Software as indicated in the order, Documentation, or as authorized by Extreme Networks in writing. Extreme Networks may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Extreme Networks through electronic means established by Extreme Networks specifically for this purpose.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Extreme Networks. All content on this site, the documentation, Hosted Service, and the product provided by Extreme Networks including the selection, arrangement and design of the content is owned either by Extreme Networks or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Extreme Networks. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Extreme Networks can be a criminal, as well as a civil offense under the applicable law.

Virtualization

The following applies if the product is deployed on a virtual machine. Each product has its own ordering code and license types. Note, unless otherwise stated, that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Extreme Networks Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

Third Party Components

“Third Party Components” mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements (“Third Party Components”), which contain terms regarding the rights to use certain portions of the Software (“Third Party Terms”). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Extreme Networks’ website at:http://www.extremenetworks.com/support/policies/softwarelicensing or such successor site as designated by Extreme Networks. The open source software license terms provided as Third Party Terms are consistent with the license rights granted in these Software License Terms, and may contain additional rights benefiting You, such as modification and distribution of the open source software. The Third Party Terms shall take precedence over these Software License Terms, solely with respect to the applicable Third Party Components to the extent that these Software License Terms impose greater restrictions on You than the applicable Third Party Terms.

The following applies only if the H.264 (AVC) codec is distributed with the product. THIS PRODUCT IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD (“AVC VIDEO”) AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE HTTP://WWW.MPEGLA.COM.

Service Provider

THE FOLLOWING APPLIES TO EXTREME NETWORKS CHANNEL PARTNER’S HOSTING OF EXTREME NETWORKS PRODUCTS OR SERVICES. THE PRODUCT OR HOSTED SERVICE MAY USE THIRD PARTY COMPONENTS SUBJECT TO THIRD PARTY TERMS AND REQUIRE A SERVICE PROVIDER TO BE INDEPENDENTLY LICENSED DIRECTLY FROM THE THIRD PARTY SUPPLIER. AN EXTREME NETWORKS CHANNEL PARTNER’S HOSTING OF EXTREME NETWORKS PRODUCTS MUST BE AUTHORIZED IN WRITING BY EXTREME NETWORKS AND IF THOSE HOSTED PRODUCTS USE OR EMBED CERTAIN THIRD PARTY SOFTWARE, INCLUDING BUT NOT LIMITED TO MICROSOFT SOFTWARE OR CODECS, THE EXTREME NETWORKS CHANNEL PARTNER IS REQUIRED TO INDEPENDENTLY OBTAIN ANY APPLICABLE LICENSE AGREEMENTS,

AT THE EXTREME NETWORKS CHANNEL PARTNER'S EXPENSE, DIRECTLY FROM THE APPLICABLE THIRD PARTY SUPPLIER. WITH RESPECT TO CODECS, IF THE EXTREME NETWORKS CHANNEL PARTNER IS HOSTING ANY PRODUCTS THAT USE OR EMBED THE G.729 CODEC, H.264 CODEC, OR H.265 CODEC, THE EXTREME NETWORKS CHANNEL PARTNER ACKNOWLEDGES AND AGREES THE EXTREME NETWORKS CHANNEL PARTNER IS RESPONSIBLE FOR ANY AND ALL RELATED FEES AND/OR ROYALTIES. THE G.729 CODEC IS LICENSED BY SIPRO LAB TELECOM INC. SEE WWW.SIPRO.COM/CONTACT.HTML. THE H.264 (AVC) CODEC IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO:
(I) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (II) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION FOR H.264 (AVC) AND H.265 (HEVC) CODECS MAY BE OBTAINED FROM MPEG LA,
L.L.C. SEE HTTP:// WWW.MPEGLA.COM.

Compliance with Laws

You acknowledge and agree that it is Your responsibility for complying with any applicable laws and regulations, including, but not limited to laws and regulations related to call recording, data privacy, intellectual property, trade secret, fraud, and music performance rights, in the country or territory where the Extreme Networks product is used.

Preventing Toll Fraud

"Toll Fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Security Vulnerabilities

Information about Extreme Networks' security support policies can be found in the Global Technical Assistance Center Knowledgebase at https://gtacknowledge.extremenetworks.com/.

Downloading Documentation

For the most current versions of Documentation, see the Extreme Networks Support website: http://documentation.extremenetworks.com, or such successor site as designated by Extreme Networks.

Contact Extreme Networks Support

See the Extreme Networks Support website:http:// www.extremenetworks.com/support for product or Hosted Service notices and articles, or to report a problem with your Extreme Networks product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Extreme Networks Support website:http://www.extremenetworks.com/support/contact/ (or such successor site as designated by Extreme Networks), scroll to the bottom of the page, and select Contact Extreme Networks Support.

Contact Avaya Support

See the Avaya Support website: https://support.avaya.com for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: https://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation, Hosted Service(s), and product(s) provided by Extreme Networks are the registered or unregistered Marks of Extreme Networks, Inc., its affiliates, its licensors, its suppliers, or other third parties. Users are not

Abstract

This Technical Configuration Guide provides an overview and examples on configuring various items related to Shortest Path Bridging (SPB) support on the VSP 4000, VSP 7000, VSP 7200, VSP 9000, ERS 4800, ERS 5900, and ERS 8800.

Acronym Key

Throughout this guide the following acronyms will be used:

- AS : Autonomous System
- B-MAC : Backbone MAC
- B-VID : Backbone VLAN identifier
- BCB : Backbone Core Bridge
- BEB : Backbone Edge Bridge
- C-MAC : Customer MAC
- CFM : Connectivity Fault Management
- FA : Fabric Attach
- FC : Fabric Connect (SPBM)
- FE : Fabric Extend (SPBM over IP)
- GRT : Global Route Table
- ISID : Backbone Service Instance Identifier; IEEE 802.1ah
- IPVPN : IP Virtual Private Network
- IS-IS : Intermediate System to Intermediate System
- IST : Inter Switch Trunk (Extreme SMLT Clustering)
- L2 VSN : Layer 2 Virtual Services Network
- L3 VSN : Layer 3 Virtual Services Network
- LLDP : Link Layer Discovery Protocol; IEEE 802.1AB
- LSDB : Link State Data Base
- MAC : Media Access Control
- MLT : Multi Link Trunk
- BCB : Backbone Core Bridge

- SDN Fx : Software Defined Networking with FA, FC, FE
- SMLT : Split MLT (Extreme Clustering)
- SPB : Shortest Path Bridging
- SPBM : Shortest Path Bridging MAC
- TLV : Type Length Value
- VID : VLAN identifier
- VLACP : Virtual LACP
- VLAN : Virtual LAN
- VPN : Virtual Private Network

Revision Control

No	Date	Version	Revised By	Remarks
1	12/21/2010	1.0	PRMGT	Modifications to Software Baseline section
2	2/28/2011	1.1	PRMGT	Remove reference to BEB to BCB.
3	3/15/2011	1.2	PRMGT	Remove reference to InterISID Routing and Native IP shortcuts. Changed to InterVSN routing and GRT Shortcuts
4	4/14/2011	1.3	PRMGT	Added SPBM IP enabled for configuration example SPB L3 VSN in reference to ERS-1
5	8/25/2011	1.4	PRMGT	Changes to SPBM NNI SMLT diagrams
6	11/21/2011	1.5	John Vant Erve	Changed name from GRT shortcuts to IP Shortcuts. Added changes to CFM provisioning made in release 7.1.1.0. Note regarding SPB sys-name.
7	6/27/2012	1.6	John Vant Erve	Added SPB multicast related information and configuration examples. Added addition information pertaining to the VSP 9000.
8	4/8/2013	2.0	John Vant Erve	Added VSP 4000 and VSP 7000. Updated the configuration examples
9	8/8/2013	2.1	John Vant Erve	Adding configuration changes regarding Spanning Tree on SPB NNI ports in configuration examples
10	4/7/2014	2.2	John Vant Erve	Added ERS 4800 and VSP 8000.
11	12/5/2014	3.0	John Vant Erve	Add vIST, IPv6 Shortcuts, and ISIS accept policy.
12	4/22/2015	3.1	John Vant Erve	Added note regarding multiple port NNI MLT

				support.
13	5/21/2015	3.2	John Vant Erve	Changes to section 12.2. Updated section 5 and 6 in reference to VoSS 4.2. Added section 7.
14	7/9/2015	3.3	John Vant Erve Ludo Stevens Goeran Friedl	Added ETREE, CFM table, and added ER5900 and VSP7200 switches
15	11/19/2015	3.4	John Vant Erve	Added FA and FE.

Table of Contents

Figures 14
Tables 15
1. Overview 17
1.1 Evolution of Ethernet Bridging 17
1.2 SPB Benefits 19
2. SPB Terminology 22
2.1 SPB 22
2.2 SPBM 22
2.3 IS-IS 22
2.4 B-VLAN 22
2.5 B-MAC (System ID) 23
2.6 System ID Value 24
2.7 Backbone Edge Bridges (BEB) and Backbone Core Bridges (BCB) 25
2.8 Connectivity Fault Management (CFM) 25
3. SPB Support Topologies 28
3.1 SPB L2 VSN 28
3.2 SPB L3 VSN 29
3.3 Inter VSN Routing 30
3.4 SPB IP Shortcuts 31
4. UNI Types 32
4.1 L2VSN – C-VLAN UNI 32
4.2 L2VSN – Switched UNI 33
4.3 L2VSN – Transparent UNI 34
4.4 Flex UNI - Switched 35
4.5 Private VLAN – ETREE 36
4.6 UNI Type – Example 37
5. Summary of SPB Features and Product Release Matrix 38
6. SPB Feature and License Matrix 40
7. Scaling 41
8. Migration & Upgrades 43
8.1 Common Upgrade instructions 43
8.2 Upgrade from Pre-5.1 releases for the ERS 8800 43
8.3 Upgrade and SMLT Cluster 43
8.4 VSP 7000 44

8.5 Activating SPB 46
8.6 Migrating traffic to SPB 48
8.7 Multicast 49
8.8 Migrating a VLAN to an L2 VSN 51
8.8.1 Migrating to Inter VSN Routing 51
8.9 VSP 9000 Notes 52
9. Field Introduction & Support Specifications 53
9.1 Hardware and Deployment Specifications 53
9.2 Installation and Commissioning Specifications 54
9.3 Interoperability and Backwards / Forward Compatibility Specifications 54
10. VSP 7000 – Fabric Interconnect 55
11. ISIS Metrics - Optional 57
12. ISIS Accept Policy 58
13. ISIS External Metric 59
14. SPB over L2/L3 networks 60
14.1 Supported Networks 60
14.2 Fabric Extend Solutions 62
15. Fabric Attach 72
15.1 Fabric Attach Solution Overview 72
16. SPB SMLT BEB Design Best Practices 78
16.1 SMLT BEB – C-VLAN Guidelines for L2VSN 78
16.2 SMLT BEB – Virtual Inter-Switch Trunk (vIST) 79
16.3 SMLT BEB – ISIS Hello Timer Guidelines for ERS 8800 80
16.4 SMLT BEB – RSMLT 81
16.5 SMLT BEB – VLACP Guidelines 82
16.6 SMLT BEB – VSP 7000 Guidelines 83
16.7 SLPP Guard 83
17. SPB NNI SMLT – migrating existing SMLT network to SPB 84
18. IS-IS TLV 88
19. SPB Best Practices 89
20. SPB Configuration 91
20.1 SPB Configuration 92
20.1.1 ERS 8800 – Converting from CLI to ACLI 92
20.1.2 SPB and IS-IS Core Configuration 92
20.1.3 SPB NNI Interface Configuration 95
20.1.4 CFM Configuration 97

20.1.5 VSP 7000 – Fabric Interconnect Mesh 98
20.1.6 SMLT – Normal IST 99
20.1.7 SMLT - Virtual IST (vIST) 101
20.1.8 L2VSN Configuration 103
20.1.9 SwitchedUNI Configuration 104
20.1.10 Flex UNI Switched Configuration 105
20.1.11 Transparent UNI Configuration 106
20.1.12 Private VLAN (ETREE) Configuration 108
20.1.13 L3VSN Configuration 115
20.1.14 L3VSN – leaking routes between VRF's 117
20.1.15 IP Shortcuts 118
20.1.16 IP Shortcut– Suppress IST Network 119
20.1.17 IP Shortcuts – leaking routes between GRT and VRF 120
20.1.18 IP Shortcuts – redistribution of ISIS and OSPF 121
20.1.19 Inter-VSN Routing 126
20.1.20 IPv6 Shortcuts 128
20.1.21 SPB Multicast Configuration 130
20.1.22 Multicast 239.255.255/24 – UPnP Filtering 133
20.1.23 Connectivity Fault Management (CFM) Configuration 135
20.1.24 CFM Configuration Example – 7.1.1.x or higher 137
20.1.25 Fabric Extend Configuration 138
20.1.26 ONA: Assigning a Static IP address to the Open Network Adapter 139
20.1.27 Fabric Extend over Routed Infrastructure using VRF to interconnect to routed network 142
20.1.28 Fabric Extend over Routed Infrastructure using GRT to interconnect to routed network 147
20.1.29 Fabric Extend over E-LAN/VPLS (L2) network using Layer 3 over Layer 2 tunneling using VSP 4000 153
20.1.30 Fabric Extend over E-LAN/VPLS (L2) network using Layer 3 over Layer 2 tunneling with VSP8000 or VSP7200 158
20.1.31 Fabric Extend over E-LAN/VPLS (L2) network using VLAN Tunnels 166
20.1.32 Fabric Attach Configuration 169
20.1.33 Identity Engines – Attribute Details 172
20.1.34 Fabric Attach Base Configuration – Adding a FA Proxy and FA Server 173
20.1.35 Fabric Attach Proxy Standalone 201
20.2 Using EDM 214
20.2.1 IS-IS and SPB Configuration 214
20.2.2 VSN Configuration 217
20.2.3 Connectivity Fault Management (CFM) Configuration – release 7.0 or 7.1.1 220
21. VLAN and ISID Restrictions using TACACS+ via Identity Engines 223
21.1 TACACS+ Switch Configuration 224
22. Configuration Examples 227
22.1 SPB – Core Setup 227

22.1.1 Configuration 229
22.1.2 Configuration using EDM – Using 8005 as an example 258
22.1.3 Verify Operations 266
22.2 SMLT Configuration 292
22.2.1 Verify Operations 295
22.3 SPB L2 VSN Configuration 297
22.3.1 VLAN configuration 298
22.3.2 Layer 2 VSN configuration 299
22.4 VSP 7000 & ERS 4800 – In-band Management via L2VSN 306
22.5 Multicast over L2VSN 309
22.5.1 Enable SPB Multicast – Global 310
22.5.2 Enable IGMP 310
22.5.3 Verify Operations 312
22.6 Inter VSN Routing 319
22.7 Inter-ISID Configuration 320
22.7.1 VRF configuration 320
22.7.2 Verification 321
22.8 SPB L3 VSN – SMLT 326
22.8.1 SPB IP Enable 327
22.8.2 VLAN Configuration 328
22.8.3 IPVPN Configuration 329
22.8.4 Enable L3VSN Configuration 331
22.8.5 Enable direct interface redistribution 332
22.8.6 Verify Operations 333
22.9 Extending L3VSN to the VSP 7000 Cluster via L2VSN 344
22.9.1 L2VSN Configuration 345
22.9.2 VRF Configuration 345
22.9.3 Verify Operations 346
22.10 Multicast over L3VSN 349
22.10.1 Enable SPB Multicast – Global 350
22.10.2 Enable Multicast VPN 350
22.10.3 Enable L3 SPB Multicast 350
22.10.4 Enable IGMP 350
22.10.5 Edge Switch 351
22.10.6 Verify Operations 352
22.11.1 IS-IS Layer 3 configuration 361
22.11.2 ECMP 365
22.11.3 Local VLAN configuration 365
22.11.4 Verify Operations 367

22.12 Multicast over IP Shortcuts 371
22.12.1 IP Shortcuts Multicast configuration 372
22.12.2 Enable IP Multicast at VLAN level 372
22.13 Verify Operations 373
22.13.1 Global Settings 373
22.13.2 Verify IGMP cache/group and senders 374
22.13.3 Verify SPB Multicast Routes 376
22.13.4 Verify multicast TLV's 377
22.13.5 Trace Multicast Routes 379
23. Restrictions and Limitations 381
23.1 STP/RSTP/MSTP 381
23.2 SPB IS-IS 381
24. Reference Documentation 382

Figures

Figure 1: SPBM Service Type Encapsulations 20
Figure 2: SPB L2 VSN 28
Figure 3: SPB L3 VSN 29
Figure 4: Inter VSN Routing 30
Figure 5: SPB IP Shortcuts 31
Figure 7 – FI Rear Port Details 55
Figure 8: NNI - Triangle 84
Figure 9: NNI - SMLT Triangle A 84
Figure 10: NNI – SMLT Triangle B 85
Figure 11: NNI –Square A 85
Figure 12: NNI – Square B 85
Figure 13: NNI – SLT Square 85
Figure 14: NNI – SMLT Square 86
Figure 15: NNI – Full Mesh A 86
Figure 16: NNI – Full Mesh B 86
Figure 17: NNI – SMLT Full Mesh A 86
Figure 18: NNI – SMLT Full Mesh B 87

Tables

Table 1: IEEE Standards culminating with SPBM 18
Table 2: CFM Support 25
Table 3: UNI Type 37
Table 4: SPB Features and Product Release Matrix 39
Table 5: SPB Feature and Licence Matrix 40
Table 6: Scaling 41
Table 7: ISIDs per BEB for VSP 4000/7200/8000 42
Table 8: VSP 7000 Rear Port Mode 45
Table 9: ISIS Metric Option 57
Table 10: Devices that support Fabric Attach 75
Table 11: ISIS TLV's 88

Conventions

This section describes the text, image, and command conventions used in this document.

Symbols

Tip – Highlights a configuration or technical tip.

Note – Highlights important information to the reader.

Warning – Highlights important information about an action that may result in equipment damage, configuration or data loss.

Text

Bold text indicates emphasis.

Italic text in a Courier New font indicates text the user must enter or select in a menu item, button or command:

```
ERS5520-48T# show running-config
```

Output examples from Extreme devices are displayed in a Lucida Console font:

```
ERS5520-48T# show sys-info

Operation Mode:          Switch
MAC Address:             00-12-83-93-B0-00
PoE Module FW:           6370.4
Reset Count:             83
Last Reset Type:         Management Factory Reset
Power Status:            Primary Power
Autotopology:            Enabled
Pluggable Port 45:       None
Pluggable Port 46:       None
Pluggable Port 47:       None
Pluggable Port 48:       None
Base Unit Selection:     Non-base unit using rear-panel switch
sysDescr:                Ethernet Routing Switch 5520-48T-PWR
                         HW:02       FW:6.0.0.10  SW:v6.2.0.009
                         Mfg Date:12042004    HW Dev:H/W rev.02
```

1. Overview

1.1 Evolution of Ethernet Bridging

The evolution of Ethernet technologies continues with the IEEE 802.1aq standard of Shortest Path Bridging. This next generation virtualization technology will revolutionize the design, deployment and operations of the Enterprise Campus core networks along with the Enterprise Data Centre. The benefits of the technology will be clearly evident in its ability to provide massive scalability while at the same time reducing the complexity of the network. This will make network virtualization a much easier paradigm to deploy within the Enterprise environment.

Shortest Path Bridging brings the features and benefits required by Carrier grade deployments to the Enterprise market without the complexity of alternative technologies traditionally used in Carrier deployments (typically MPLS).

The IEEE has been working on Layer 2 virtualization techniques over the last decade. It had standardized a set of solutions that built on each other and continuously addressed the predecessor's disadvantages.

In 1998, IEEE 802.1Q provided a simple way to virtualize Layer 2 broadcast domains by using VLAN tagging to form Virtual LANs. The 12 bits that are available in the 802.1Q defined header provided the ability to separately transport 4096 individual virtual LANs.

The loop free topology had been provided through IEEE 802.1D spanning tree and later rapid spanning tree (RSTP) and multiple spanning tree (MSTP) extensions. However, spanning tree is not the technology of choice for large network deployments.

Carrier deployments wanted to leverage the cost points of Ethernet and wanted to use the virtual LAN technology. In order to improve scalability, the IEEE introduced the QinQ approach, where the header had been extended to provide a carrier tag attached to a customer tag (QinQ). This allowed the carrier to transport customer tagged traffic over its Ethernet based 802.1ad backbone. However in large deployments this technology did not scale well, because the carrier's backbone still "saw", and thus learned, all the end-customer MAC addresses (C-MAC).

In order to overcome this scaling limitation, the IEEE standardized 802.1ah (also known as Provider Backbone Bridging – BCB) in 2008 which introduced a new header encapsulation to hide the customer MAC addresses inside an additional backbone MAC header (MACinMAC encapsulation).

In addition to this, the new header also includes a service instance identifier (ISID) with a length of 24 bits. This ISID can be used to identify any virtualized traffic across an 802.1ah encapsulated frame. In 802.1ah, these ISIDs are used to virtualize VLANs across a BCB network. The "hiding/encapsulating" of customer MAC addresses in backbone MAC addresses greatly improves network scalability (no end-user C-MAC learning required in the core) and also significantly improves network robustness (loops have no effect on the backbone infrastructure).

So BCB addressed the scaling issues of virtualizing and transporting VLANs across a provider backbone. Yet, within that backbone, even with BCB, the loop free topology still had to be provided by 802.1D Spanning Tree (or RSTP or MSTP).

With the latest 802.1aq Shortest Path Bridging MacInMac (SPBM) standard this final limitation is being lifted via the development of a new link-stated based technology.

Standard	Year	Name	Loop free Topology by:	Service IDs	Provisioning	Virtualization of
IEEE 802.1Q	1998	Virtual LANs (VLAN Tagging)	Spanning Tree SMLT	4096	Edge and Core	Layer 2
IEEE 802.1ad	2005	Provider Bridging (QinQ)	Spanning Tree SMLT	4096x4096	Edge and Core	Layer 2
IEEE 802.1ah	2008	Provider Backbone Bridging (MacInMac)	Spanning Tree SMLT	16 Million	Edge and Core	Layer 2
IEEE 802.1aq	2012	Shortest Path Bridging (SPBM)	Link-State-Protocol (IS-IS)	16 Million	Only Service Access Points	IEEE: Layer 2 IETF draft: Layer 3 Unicast & Multicast

Table 1: IEEE Standards culminating with SPBM

SPBM is based on the 802.1ah encapsulation schema but does not depend on spanning tree to provide a loop free Layer 2 domain, instead it uses the nodal based IS-IS topology protocol. The IEEE is reworking the spanning tree specification 802.1D to include the new SPB solution. The intention is that once the standard is implemented in network products, the network operator will be able to choose a shortest path bridging topology protocol or the legacy root tree based option.

In addition to the Layer 2 virtualization support that SPBM provides, the model is being extended to also support Layer 3 virtualization via the IETF Draft IP/SPB-Unbehagen. Where L2 virtualization associates an ISID to an edge VLAN in such a way as to extend that VLAN across the backbone, with the L3 extension a VRF can also be associated to an ISID in such a way as to extend a virtualized L3 routing instance across the backbone.

Extreme also enhanced the SPBM capability by adding multicast support which greatly simplify the multicast deployment and provide resiliency to multicast at the same time.

In summary, SPBM brings to the Enterprise network the features, functionalities and scalability demanded by carriers via the use of a single simple and dynamic link state routing protocol which is IS-IS.

1.2 SPB Benefits

The benefits that SPB brings to the Enterprise network can be listed as follows.

- Backbone provisioning simplicity

Provisioning an SPB core is as simple as enabling SPB and IS-IS globally on all the nodes and on the core facing links. The IS-IS protocol operates at layer 2, it does not need IP addresses configured on the links to form IS-IS adjacencies with neighboring switches (like OSPF does). Hence there is no need to configure any IP addresses on any of the core links.

- Natively provides virtualized Layer 2 services

Layer 2 virtualization is handled by the Backbone Edge Bridges (BEBs) where the end-user VLAN is mapped into a Backbone Service Instance Identifier (ISID) by local provisioning. Any BEB that has the same ISID configured can participate in the same L2 virtual services network (VSN). IS-IS within the SPB backbone is used as the Layer 2 routing protocol to forward traffic between the BEB and Provider Backbone Core Bridges (BCBs). Only the BEB has knowledge of the L2 VSN and corresponding MAC addresses. The BCB only has knowledge of each Backbone MAC address (B-MAC) used to send traffic across an SPB network.

- Natively provides virtualized routing services

Layer 3 virtualized routing is handled by the Backbone Edge Bridges (BEBs) where the end-user IPv4 enabled VLAN or VLANs are mapped to a Virtualized Routing and Forwarding (VRF) instance. The VRF in turn is mapped into a Backbone Service Instance Identifier (ISID) by local provisioning. Any BEB that has the same ISID configured can participate in the same L3 virtual service network (VSN). IS-IS within the SPB backbone is used as the Layer 2 routing protocol to forward traffic between the BEB and Backbone Core Bridges (BCB). Only the BEB has knowledge of the L3 VSN and corresponding IP/ARP/MAC addresses. The BCB only has knowledge of each Backbone MAC address (B-MAC) used send traffic across an SPB network.

- Adapts to any physical layer / fibre plant

IS-IS is a link-state protocol which will compute the shortest open path just like OSPF does. It can therefore be deployed on any regular (e.g. square or fully meshed core-to-distribution topologies) or irregular (e.g. ring topologies) fibre plants.

Whereas OSPF computes the shortest path to destination subnets and then populates the IP routing table with the results, IS-IS (as used with SPB) computes the shortest path to backbone node MAC addresses (B-MACs) and then populates the backbone MAC tables.

- Robust/Scalable link-state routing applied to MAC tables

With SPB, the MAC table is now only populated by the IS-IS control plane. The conventional Ethernet bridging behavior which consisted of (a) “learning” the MAC tables with the source MAC address of packets seen arriving on local ports and (b) flooding unknown and broadcast traffic to all ports no longer apply in an SPB backbone.

Furthermore, with SPB, IS-IS is leveraged to build source based forwarding trees for the delivery of multicast and broadcast traffic across the SPB backbone in such a way that the replication of broadcast/multicast traffic within the core is optimized to follow the shortest path from source to leaf nodes.

- Separation between Services and Backbone

Since SPB leverages the MACinMAC encapsulation of 802.1ah (BCB) only the nodes at the edge of the SPB backbone (the Backbone Edge Bridges - BEBs) need to learn the MAC addresses (C-MACs) used within the transported Customer VLANs (L2VSNs). These same nodes, when forwarding traffic into the SPB core will always re-encapsulate the service traffic in a Backbone MAC header with a destination B-

MAC corresponding to the destination SPB node across the backbone where the service traffic will get de-capsulated. The encapsulation used is shown in Figure 1. As such, the nodes within the SPB backbone will have no knowledge of the addresses used within the service VSNs (C-MACs or IP addresses) transported across and only need to provide reachability to the B-MAC addresses within the backbone.

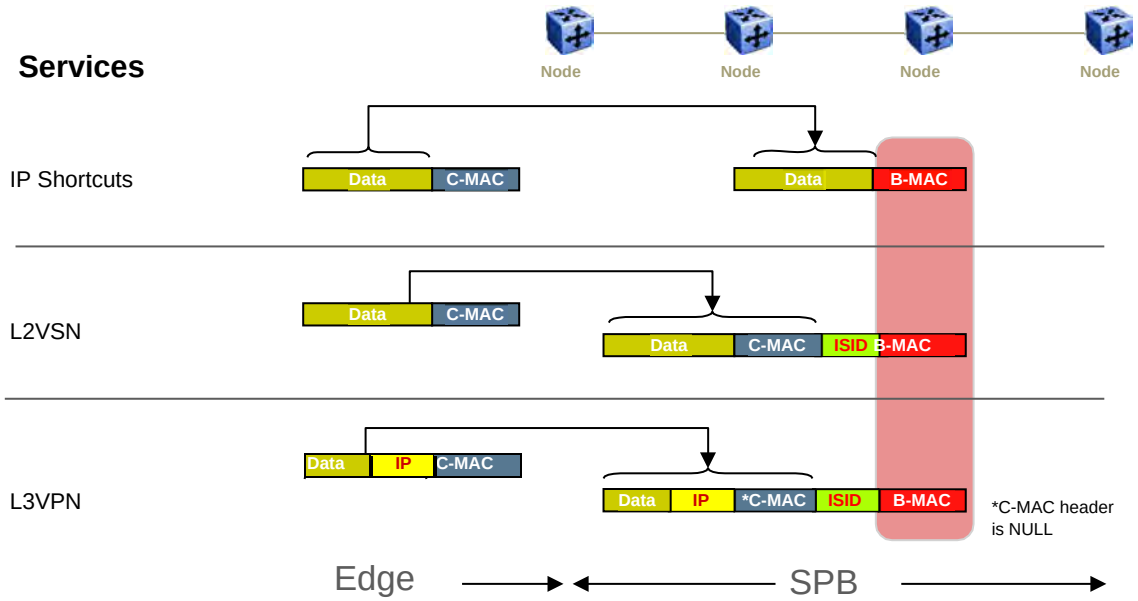

Figure 1: SPBM Service Type Encapsulations

- Connectivity Fault Management

Connectivity Fault Management (CFM) offers loopbacks and link trace for troubleshooting, and continuity checks for fast fault detection. Presently only the loopback and link trace features of CFM are supported. These commands allow operators, service providers and customers to verify the connectivity that they provide or utilize and to debug systems. This is accomplished through:

- Loopback messaging to an intermediate or endpoint within a domain for the purpose of fault verification. (LBM)
- Linktrace messaging to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. (LTM)
- End-point provisioning

The boundary between the MACinMAC SPB domain and 802.1Q domain is handled by the Backbone Edge Bridges (BEBs). At the BEBs, VLANs or VRFs are mapped into ISIDs based on the local service provisioning.

Services (whether L2 or L3 VSNs) only need to be configured at the edge of the SPB backbone (on the BEBs). There is no provisioning needed on the core SPB nodes. This provides a robust carrier grade architecture where configuration on the core nodes never needs to be touched when adding new services.

- Service provisioning simplicity

The same simplicity extends to provisioning the services to run above the SPB backbone. Creating an L2VSN is as simple as associating an ISID number with an edge VLAN; creating an L3VSN is as simple as associating an ISID number with a VRF and configuring the desired IS-IS IP route redistribution within the newly created L3VSN.

- Multicast

Multicast over SPB L2 VSNs, L3 VSNs, or IP Shortcuts is supported on the ERS 8800 beginning in the 7.2 release, on the VSP 4000 in the 3.1 release, the VSP 7200 in the 4.2.1 release, the VSP 9000 in the 3.4 release and the VSP 8000 in the 4.1 release. The ERS 4800 in release 5.9 and the ERS 5900 in release 7.0 support multicast over L2 VSNs only; please note multicast is constrained to within the VSN only. Multicast is supported over SPB by globally enabling the feature and just enabling IGMP at the SPB edge. There is no need for any multicast routing protocols such as PIM, hence, multicast over SPB greatly simplifies multicast deployment. A multicast stream can be forwarded anywhere in a SPB network where IS-IS is used to advertise the stream to the rest of the fabric. Note that the stream is not forwarded until a receiver requests to join a specific multicast group and it is only forwarded to those receivers who requested it.

2. SPB Terminology

2.1 SPB

Shortest Path Bridging (SPB) is being standardized by the IEEE as the next evolution step. It provides shortest path forwarding using layer 2 to provide shortest path forwarding. SPB uses the IS-IS protocol operating at layer 2 allowing for large networks with fast convergence, equal cost paths, and easy provisioning without having to add complex addition protocols in the core to support virtualization of VLAN's or VRF's. In summary, all that is needed is to enable SPB and IS-IS in the core and all the virtualization is done on the edge.

2.2 SPBM

The 802.1aq standard supports two modes, SPB VID (SPBV) and SPB MAC (SPBM). Only SPBM supports true virtualization via the use of the 802.1ah MAC-in-MAC encapsulation. SPBV offers shortest path forwarding but with reduced functionality using 802.1ad Q-in-Q tagging for devices which cannot support the 802.1ah MAC-in-MAC encapsulation. All Extreme SPB capable switches support exclusively SPBM. SPBM virtualized services are delineated by ISIDs where the ISID is simply assigned at the BEB to either a VLAN for virtualized layer 2 services, a multicast group for virtualized multicast services, or to a VRF for virtualized layer 3 services.

In a SPBM network, each bridge advertises its own unique MAC address using IS-IS which is known as the system-id. The system-id can also be manually provisioned to ease in trouble shooting when looking at the layer 2 forwarding table.

2.3 IS-IS

Provisioning an SPB core is as simple as enabling SPB and IS-IS globally on all the nodes and configuring SPB IS-IS interfaces on the core facing links (NNI links). The IS-IS protocol operates at layer 2, it does not need IP addresses configured on the links to form IS-IS adjacencies with neighboring switches (like OSPF does). Hence, there is no need to configure any IP addresses on any of the core links.

IS-IS is a link-state protocol which will compute the shortest open path just like OSPF does. It can therefore be deployed on any regular (e.g. square or fully meshed core-to-distribution topologies) or irregular (e.g. ring topologies) fibre plants.

2.4 B-VLAN

Each SPBM network instance is associated with at least one backbone VLAN (B-VLAN) in the core SPBM network.

This VLAN is used for both control plane traffic and dataplane traffic.

Extreme recommends to always configuring two B-VLANs in the core to allow load distribution over both B-VLANs.

SPBM alters the behavior of the VLAN. When a B-VLAN is associated with an SPBM network the following VLAN attributes and behaviors are modified for the B-VLAN:

- Flooding is disabled
- Broadcasting is disabled

- Source address learning is disabled
- Unknown mac discard is disabled

Essentially the VLAN becomes a header indicating the SPBM network to use.

Modification of the VLAN behavior is necessary to ensure proper control over the SPBM traffic.

Although it is recommended to use BVIDs that are in the upper range, using a BVID less than 4000 may have to be used if tunneling SPB across an MPLS or IP network via a router GRE tunnel. For example, the Ayava Secure Router supports VLAN tunneling via GRE with a restriction of allowing only VLAN ID's of less than 4000.

2.5 B-MAC (System ID)

Whereas OSPF computes the shortest path to destination subnets and then populates the IP routing table with the results, IS-IS (as used with SPB) computes the shortest path to backbone node MAC addresses (B-MACs) and then populates the backbone MAC tables. The B-MAC addresses are advertised in IS-IS via one or more backbone VLAN IDs (B-VIDs). In summary, frames are forwarded using the System-Id as the Backbone Source Access (B-SA) to a specific node using the Backbone Destination Address (B-DA). Note that the backbone nodes will know how to reach all the B-MACs (IS-IS will have programmed the B-VID MAC tables according) while the Customer MACs (C-MACs) will only be learned on the appropriate BEB nodes which terminate the virtual services.

The SPB forwarding database (FDB) will contain a combination of unicast and multicast MAC addresses.

SPB uses source specific multicast trees. There has to be a unique multicast tree for every BEB across all B-VIDs provisioned and for every Service Instance (ISID) which requires delivery of multicast/broadcast (L2VSNs and only L3VSNs if enabled for IP Multicast). In terms of IS-IS computation there will be as many multicast SPT trees as there are SPB nodes across each B-VID. These trees will then be further pruned into Service (ISID) specific multicast SPTs based on which BEBs are configured with the corresponding ISID. In the data plane every individual Service Specific multicast SPT will have a unique Multicast MAC address defined which is obtained by combining the ingress BEB Nick-name (referred to as the SP SourceID; 20 bits) with the ISID service identifier (24 bits).

SPB Unicast FDB

```
CLI
ERS-8800:5# show isis spbm unicast-fib
ERS-8800:5# show isis spbm unicast-fib vlan <vlan-id>
ACLI
ERS-8800:5#show isis spbm unicast-fib
ERS-8800:5#show isis spbm unicast-fib vlan <vlan-id>
================================================================================
                          SPBM UNICAST FIB ENTRY INFO
================================================================================
DESTINATION        BVLAN  SYSID           HOST-NAME         OUTGOING     COST

ADDRESS                                                     INTERFACE
--------------------------------------------------------------------------------
```

```
02:be:b0:00:00:02  40     00be.b000.0002   ERS-2          2/2          10
02:be:b0:00:00:02  41     00be.b000.0002   ERS-2          2/2          10
02:be:b0:00:00:30  40     00be.b000.0030   ERS-3          2/2          20
02:be:b1:00:03:04  40     00be.b000.0030   ERS-3          2/2          20
02:be:b0:00:00:30  41     00be.b000.0030   ERS-3          2/2          20
02:be:b0:00:00:40  40     00be.b000.0040   ERS-4          2/2          20
02:be:b0:00:00:40  41     00be.b000.0040   ERS-4          2/2          20
02:be:b1:00:03:04  41     00be.b000.0040   ERS-4          2/2          20
```

SPB Multicast FDB

```
CLI
ERS-8800:5# show isis spbm multicast-fib
ERS-8800:5# show isis spbm multicast-fib vlan <vlan-id>
ACLI
ERS-8800:5#show isis spbm multicast-fib
ERS-8800:5#show isis spbm multicast-fib vlan <vlan-id>
================================================================================
                        SPBM MULTICAST FIB ENTRY INFO
================================================================================
MCAST DA           ISID   BVLAN  SYSID               HOST-NAME         OUTGOING
-INTERFACES
--------------------------------------------------------------------------------------------------------
03:00:01:00:03:e8  1000   40     0001.8128.87df      ERS-1             2/2,3/11
03:00:01:00:03:e9  1001   40     0001.8128.87df      ERS-1             2/2,3/12,3/13
03:00:04:00:03:e8  1000   40     0001.8129.1fdf      ERS-4             3/11
03:00:04:00:03:e9  1001   40     0001.8129.1fdf      ERS-4             3/12,3/13
03:00:03:00:03:e8  1000   40     0080.2dbe.23df      ERS-3             3/11
03:00:03:00:03:e9  1001   40     0080.2dbe.23df      ERS-3             3/12,3/13
```

2.6 System ID Value

The default switch behavior regarding System-Id is to use a MAC address within the MAC address range reserved for the switch. This ensures that there will be no de-stabilizing System-Id conflicts in the network. Extreme recommends the use of default System-Id values for this reason. To allow greater flexibility to customers, the use of configured System-Id values is also supported. When using configured System-Id values it is very critical to ensure that each SPB enabled switch in the network uses a unique ISIS System-Id value.

If you do decide to change the System ID, it is recommended to set the locally administered bit. The second least significant bit of the most significant byte of the MAC address should be set to 1 to indicate the MAC address as locally administered. For more details, please go to http://en.wikipedia.org/wiki/MAC_address.

2.7 Backbone Edge Bridges (BEB) and Backbone Core Bridges (BCB)

The BEB provides the boundary between the MACinMAC SPBM domain and virtualized service domain. At the BEBs, VLANs or VRFs are mapped into ISIDs based on the local service provisioning. As such, all nodes within the SPBM backbone will have no knowledge of the addresses within the virtualized services VSNs (C-MAC or IP addresses). Only the BEB nodes will contain a C-MAC table (or FDB), and if configured, a VRF IP forwarding table. All backbone nodes will have no knowledge of the virtualized service VSNs, C-MAC and VRF addresses.

In the SPBM core, the bridges are referred to as Backbone Core Bridges (BCBs). BCBs forward encapsulated traffic based on the BMAC-DA.

2.8 Connectivity Fault Management (CFM)

Platform (current Rel.) Feature	Modular OS		VSP OS (VOSS)		Stackable OS (BOSS)		
	ERS 8800 (7.2)	VSP 9000 (4.1)	VSP 8000 VSP 7200 (5.0)	VSP 4000 (5.0)	VSP7000 (10.4)	ERS4800 (5.9)	ERS5900 (7.0)
CFM for SPB BVLAN	●	●	●	●	●	●	●
CFM for CVLAN	● (7.1.1)	● (3.4)	tbd	●	-	-	-
Simplified CFM configuration	●	● (3.4)	●	● (3.1)	●	●	●
Full EDM support	● (7.1.1)	●	●	●	●	●	●

Table 2: CFM Support

Connectivity Fault Management (CFM) offers loopbacks and link trace for troubleshooting and continuity checks for fast fault detection - loopback and link trace features of CFM are supported. These commands allow operators, service providers and customers to verify the connectivity that they provide or utilize and to debug systems. This is accomplished through:

- Loopback messaging to an intermediate or endpoint within a domain for the purpose of fault verification. (LBM)
- Linktrace messaging to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for the purpose of fault isolation. (LTM)

IEEE 802.1ag – Connectivity Fault Management (Per Service/VLAN OAM)

- Maintenance Domain – MD
  - MD is management domain on a network, typically owned and operated by a single entity MD are configured with Names and Levels, where the eight levels range from 0 to 7.

  - Hierarchal relationship exists between domains based on levels.
  - Recommended values of levels are as follows
    - Customers – Largest (e.g., 7)
    - Providers – In between (e.g., 3)
    - Operators – Smallest (e.g., 1)

- Maintenance Association
  - Maintenance Association (MA) is a set of MEPs, all of which are configured with the same MAID (Maintenance Association Identifier) and MD Level, each of which is configured with a MEPID unique within that MAID and MD Level, and all of which are configured with the complete list of MEPIDs”
- Maintenance End Point
  - Maintenance End Point (MEP), are Points at the edge of the domain, define the boundary for the domain A MEP sends and receives CFM frames through the relay function, drops all CFM frames of its level or lower that come from the wire side
- Maintenance Intermediate Point
  - Maintenance Intermediate Point (MIP), are Points internal to a domain, not at the boundary. MIPs are Passive points, respond only when triggered by CFM trace route and loop-back messages
- There are 5 message types for CFM
  - Continuity Check Message (CCM) – Not implemented.
  - Loopback Message (LBM)
  - Loopback Reply (LBR)
  - Linktrace Message (LTM)
  - Linktrace Reply (LTR)
- Raw loopback and linktrace messages can be generated using the following ACLI commands:
  - loopback <mdName> <maName> <mepId> <rmepMac>
  - linktrace <mdName> <maName> <mepId> <rmepMac>
- However, a set of L2 OAM commands which leverage the underlying CFM loopback and linktrace messages but provide a more user friendly interface and a simplified summary of the information which is carried in the CFM messages. These commands can also be executed against a target system name or IP address instead of a MAC address of MD.MA.MEP-id.
  - l2 ping <vlan> mac <SystemIdMac>
  - l2 ping <vlan> routernodename <RouterNodeName >
  - l2 traceroute vlan <vlan> mac <SystemIdMac>
  - l2 traceroute vlan <vlan> routernodename <RouterNodeName >
  - l2 traceroute ip-address < ipaddress> ?

    priority Priority <0-7>
    source-mode Source mode<nodal|noVlanMac|smltVirtual>
    ttl-value Ttl value <1-255>
    vrf Vrf
    <cr>

- Starting in software release 7.1.1 for the ERS 8800, release 3.4 for the VSP 9000, release 3.0 for the VSP 4000, 5.7 for the ERS 4800, and 10.2 for the VSP 7000, CFM commands will now automatically create a MEP and a MIP at a specific level for every SPB B-VLAN provisioned on the switch. Hence, you no longer have to configure explicit MEPs and MIPs and associated VLANs with MEPs and MIPs. Previously configured MIPs and MEPs will still continue to work if you upgrade from either 7.0 or 7.1 to release 7.1.1.x. In summary:
  - Auto-generated CFM commands create a MEP and a MIP at a specified level for every SPBM B-VLAN on the chassis
  - No more having to configure explicit MEPs and MIPs and associate multiple VLANs with MEPs and MIPs
    - Previously configured MEPs and MIPS will continue to work
  - Auto-generated MEPs and MIPs respond to l2ping, l2traceroute, and l2tracetree in the same manner as in 7.1
  - CFM extended to support C-VLANs in addition to existing support for B-VLANs.
    - This enables you to isolate a connectivity fault in either the SPBM cloud or in a customer domain.

3. SPB Support Topologies

3.1 SPB L2 VSN

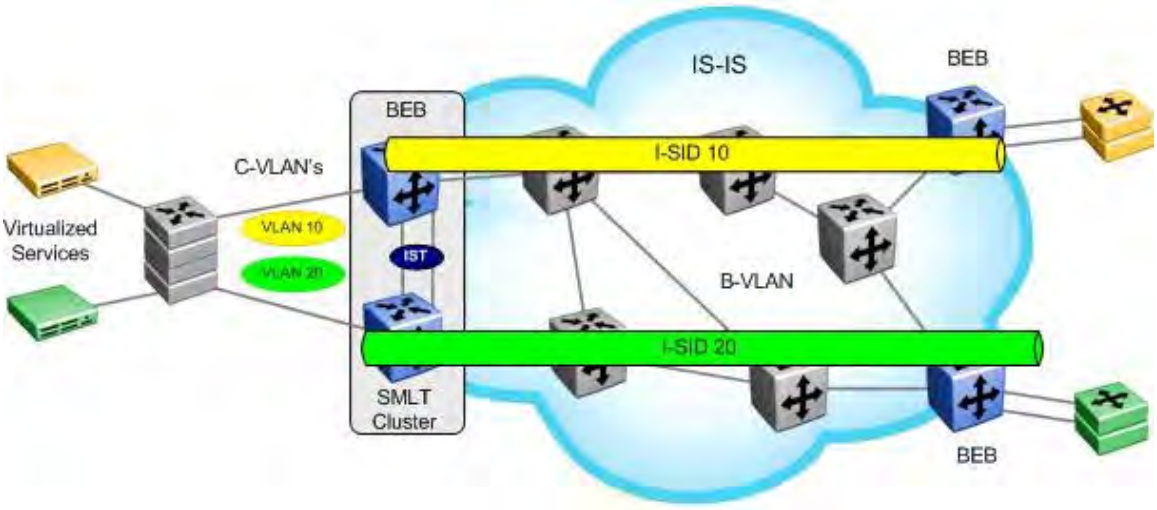

Figure 2: SPB L2 VSN

A SPB L2 VSN topology is simply made up of a number of Backbone Edge Bridges (BEB) used to terminate Layer 2 VSNs. The control plane uses IS-IS for forwarding at a Layer 2 level. Only the BEB bridges are aware of any VSN and associated MAC addresses while the backbone bridges simply forward traffic at the Backbone MAC (B-MAC) level. The backbone switches will know how to reach every B-MACs using the shortest path determined by IS-IS. Note that the backbone System ID or B-MAC can be manually provisioned to help ease trouble-shooting when looking at the B-MAC unicast forwarding table. In summary, all switches in the backbone will only learn B-MAC addresses to make forwarding decisions while the BEB will learn both the B-MACs and Customer MACs (C-MAC) for each VSN. A Backbone Service Instance Identifier (ISID) will be assigned on the BEB to each VLAN. All VLANs in the network that share the same ISID will be able to participate in the same VSN. If SMLT clusters are used, two backbone VLANs (B-VLAN) are required with a primary B-VLAN and a secondary B-VLAN. . In general two backbone VLANs should always be used (even if no SMLT cluster is in use) since the use of 2 backbone VLANs allows IS-IS to compute equal cost trees where if 2 shortest equal cost paths exist, SPB will load balance VSN traffic across both paths.

In summary:

- At minimum, one B-VLAN must be assigned to each SPB switch
    - For SMLT, two B-VLANs are required
- TLVs and sub-TLVs are used to identify SPB instance, link metric's, B-VLAN, B-MAC, and number of ISID's

3.2 SPB L3 VSN

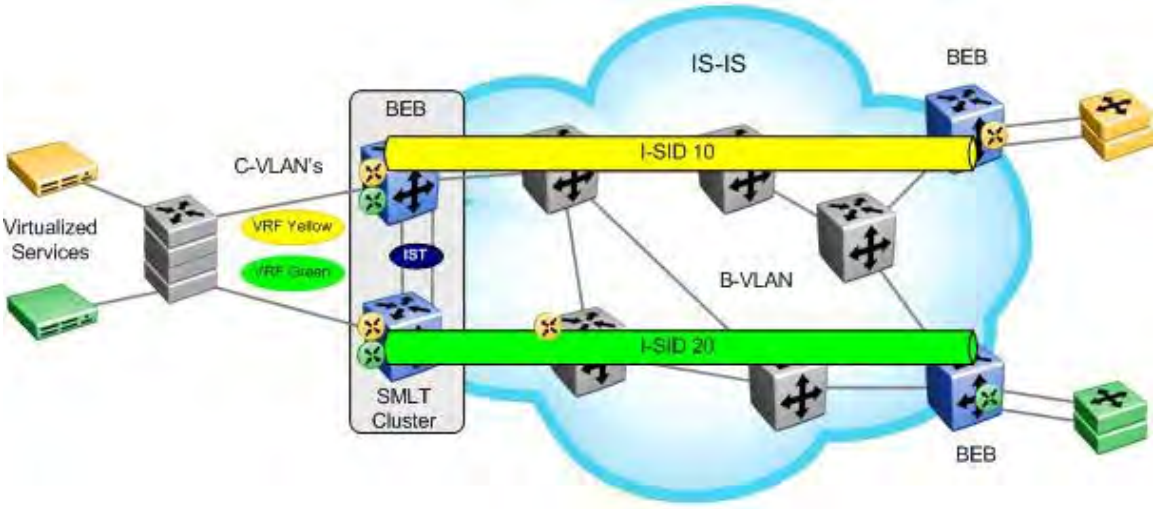

Figure 3: SPB L3 VSN

A SPB L3 VSN topology is very similar to a SPB L2 VSN topology with the exception that a Backbone Service Instance Identifier (ISID) will be assigned at a Virtual Router (VRF) level instead of at a VLAN level. All VRFs in the network that share the same ISID will be able to participate in the same VSN.

In summary:

- One or more VRFs are created on the BEB switches with an assigned ISID
  - All VRFs that share the same ISID can participate in the same VSN
- Route distribution of direct interfaces on VRF instances must be enabled to distribute VRF networks into IS-IS between BEB switches
- IS-IS IP routing must be enabled

3.3 Inter VSN Routing

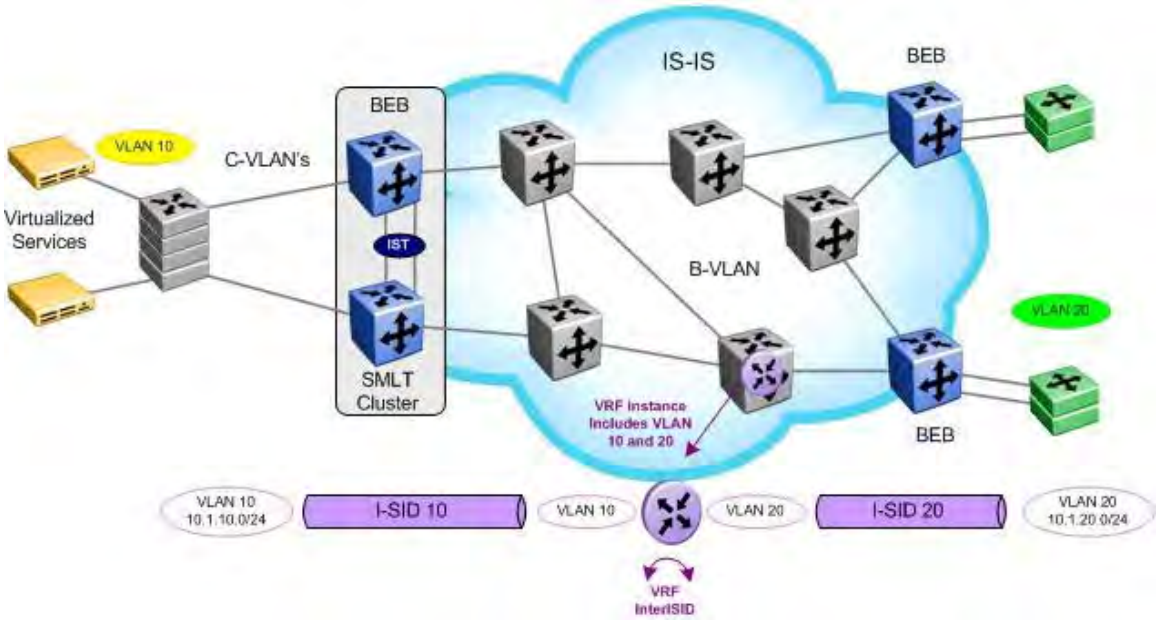

Figure 4: Inter VSN Routing

Inter VSN allows routing between IP networks on Layer 2 VLANs with different ISIDs. As illustrated in the diagram above, routing between VLANs 10 and 20 occurs on one of the SPB core switches shown in the middle of the diagram. End users from the BEB switches as shown on the right and left of the diagram are able to forward traffic between the yellow and green VLANs (VLANs 10 & 20) via the VRF instance configured on the switch shown. Although the diagram illustrates a VRF configured on a BCB switch, Inter VSN can also be performed via GRT. Also, for redundancy, Inter VSN can also be configured on another switch with VRRP to eliminate a single point of failure.

Please note Inter VSN routing is only typically used when you have to extend a VLAN as L2VSNs for applications such as vMotion. Normally, is it recommended to route when you can by using either IP shortcuts or L3VSNs. As one of the requirements for vMotion is a shared network for the ESX hosts, we have no choice but to bridge traffic between the ESX hosts. In order to forward the server traffic to the clients and vise-versa, it is necessary to IP route the traffic either via IP shortcuts or via a VRF L3VSN.

3.4 SPB IP Shortcuts

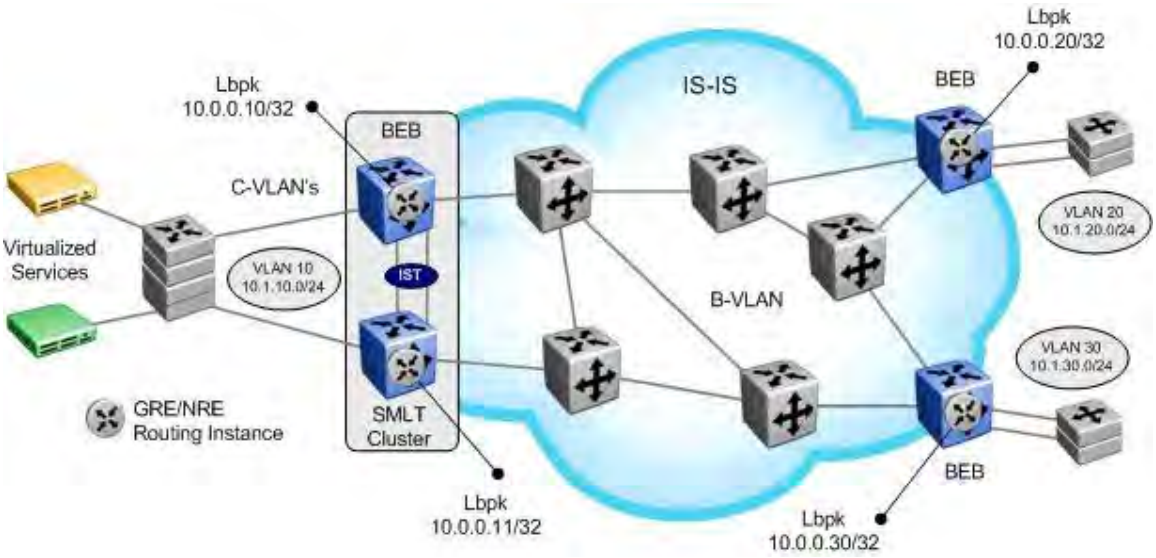

Figure 5: SPB IP Shortcuts

IP shortcuts allow routing between VLANs in the global routing table/network routing engine (GRT/NRE/VRF-0). No ISID configuration is used. IP is enabled on the B-VLAN IS-IS instance on the BEB switches. This provides normal IP forwarding between BEB sites over an IS-IS backbone.

In summary:

- IP must be enabled on IS-IS where the IS-IS source address, which must be configured, is a circuitless/loopback IP address
    - The IS-IS source address is automatically injected into IS-IS
- IS-IS redistribution of direct (or OSPF, RIP, Static, BGP…) IP routes may be enabled as a simple mechanism to forward those networks between BEB neighbors
    - This will inject all direct (or OSPF, RIP, Static, BGP…) IP routes into IS-IS
    - In a SMLT cluster, in the case of direct IP route redistribution, a route policy (CLI) or route-map (ACLI) must be configured to match the IST IP subnet to prevent it from being advertised
- The Extended IP Reachability TLV 135 is used to distribute IP reachability between IS-IS peers

4. UNI Types

4.1 L2VSN – C-VLAN UNI

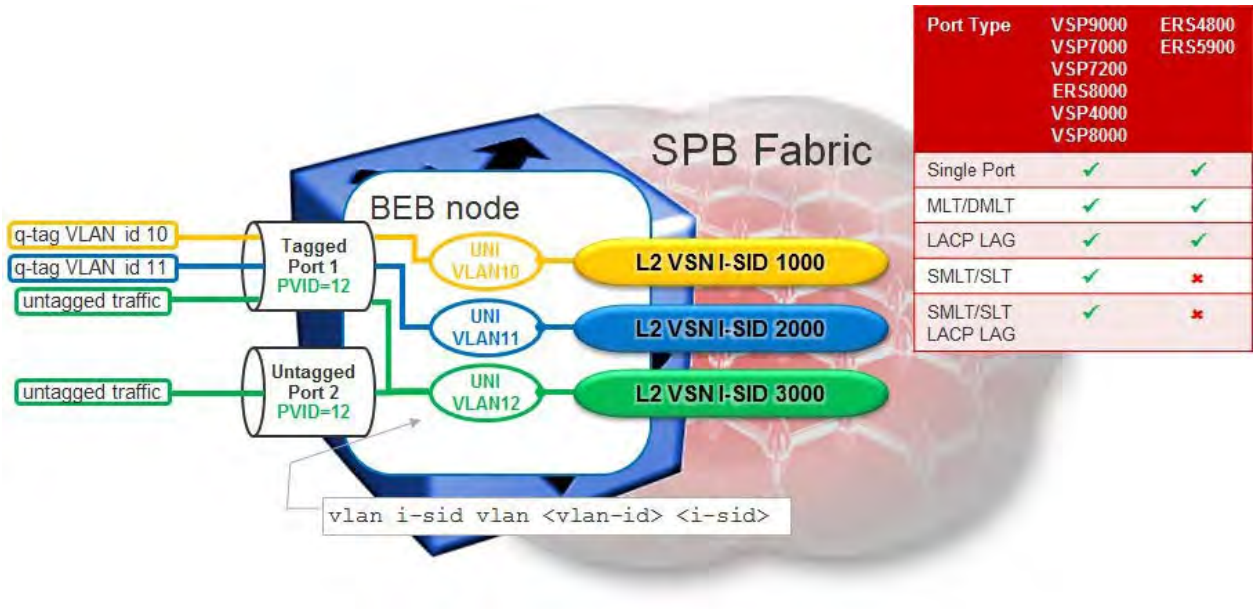

- UNI is a VLAN (Customer VLAN = C-VLAN)
- VLAN has global significance on the BEB
- VLAN performs L2 switching on local VLAN port members & transports over L2VSN for remote end-points
- Untagged traffic is assigned to VLAN corresponding to PVID configured on port
    - On tagged port, use UntagPVIDOnly mode to force PVID traffic to also go out untagged
- Supported on all SPBM capable platforms
- Switched UNIs and CVLAN UNIs can be assigned to the same ISID
- ETREE UNI and CVLAN UNI can be assigned to the same ISID
- You cannot mix Transparent UNI with C-VLAN UNI

4.2 L2VSN – Switched UNI

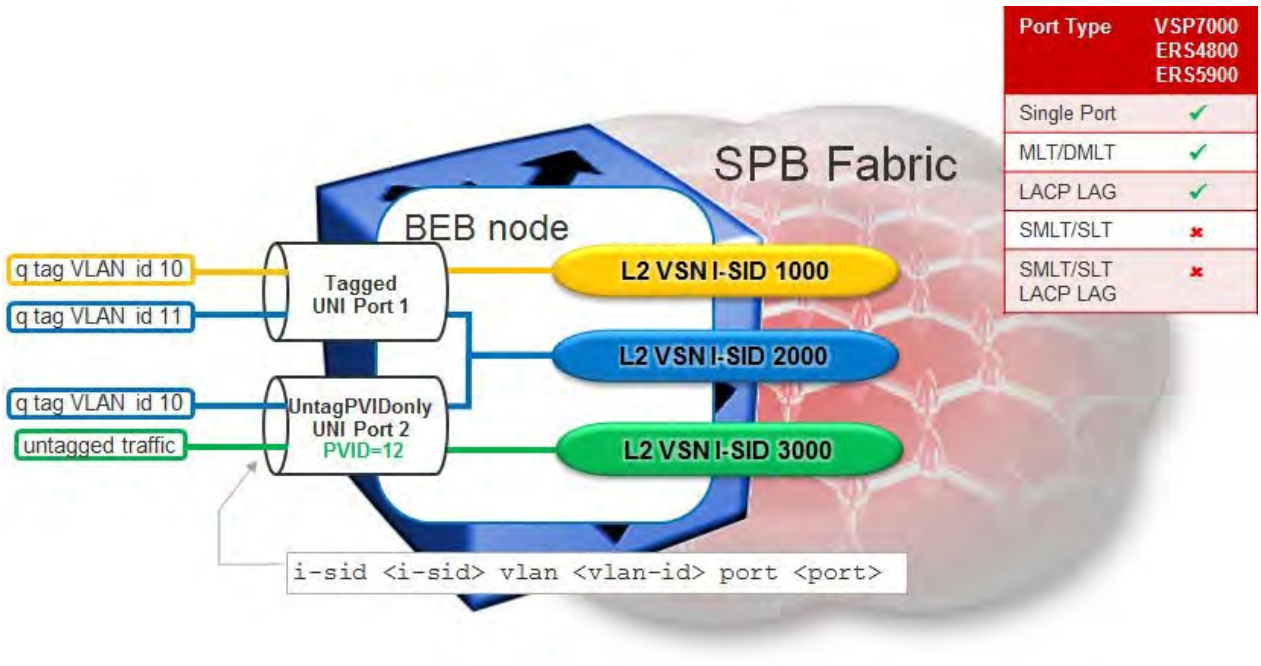

- UNI is a VLAN-id on an Ethernet port / MLT
- VLAN id has local significance on the Ethernet port / MLT
- Same VLAN-id can be re-used on different ports and belong to a different ISID
- Different VLAN-id on same or different ports can be assigned to same ISID
  - can do VLAN Mapping on local switch
- Untagged traffic can be picked up by setting the port to UntagPVIDonly and setting the PVID on the port
- Switched UNIs and CVLAN UNIs can be assigned to the same ISID
- Supported in VSP 7000 release 10.2, ERS 4800 release 5.7, and ERS 5900 in release 7.0

4.3 L2VSN – Transparent UNI

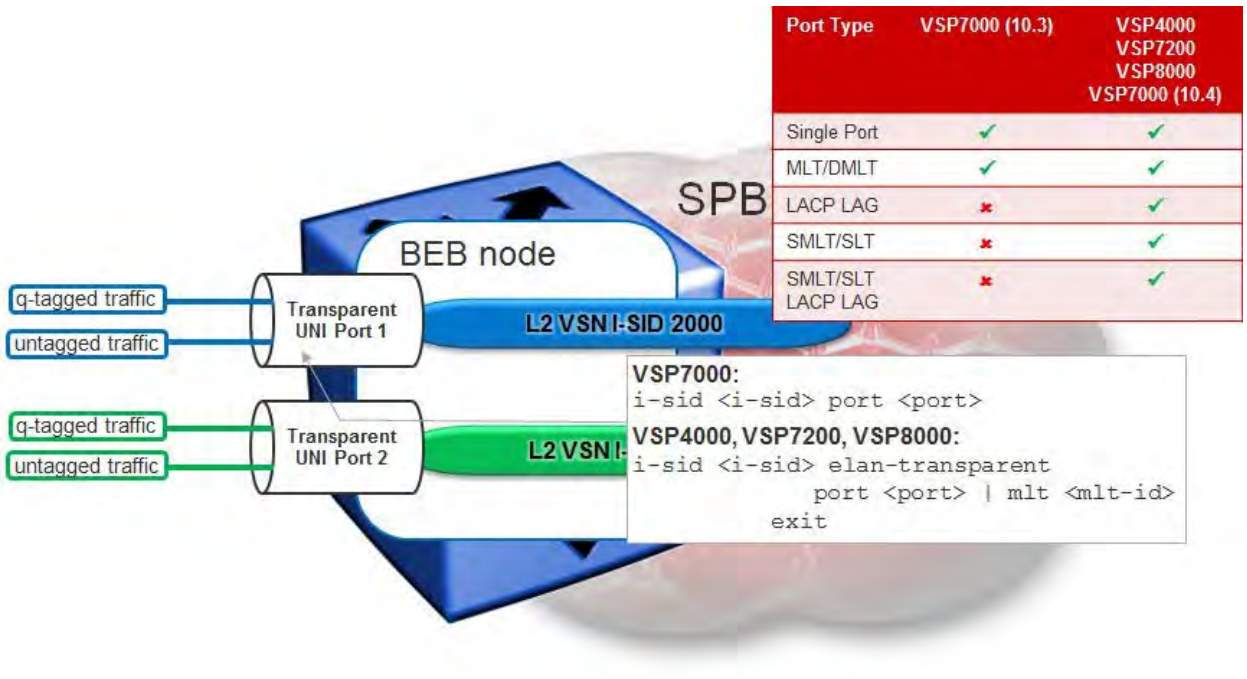

- UNI is an Ethernet port / MLT
- Ethernet UNI port / MLT is not VLAN tag aware
- Packets with or without a VLAN q-tag are transported into the L2VSN
- Untagged control traffic (STP, VLACP, LACP, LLDP, etc) is transparently forwarded
  - VLACP/LACP PDUs are forwarded (VSP 4000/7200/8000: unless configured on UNI port / MLT)
  - Flow Control Pause frames remain link local and are not transported
- Reverse MAC learning is still used, so can be used with 3 or more end-points
- Supported in VSP 7000 release 10.3 with support for SMLT in release 10.4, VSP 4000 release 3.1, VSP 8000 in release 4.2, and VSP 7200 in release 4.2.1
- MLT Transparent UNI ports are supported (on VSP 4000/7200/8000 even with LACP)
- Transparent UNIs should not be assigned to the same ISID as Switched UNI or CVLAN UNIs
- Beginning in the 4.1 release for the VSP 4000, Transparent UNI over SMLT is supported
- Transparent UNI cannot be mixed with any other type of UNI
  - You cannot use an ISID assigned to a Transparent UNI with either an C-VLAN UNI or an Switch UNI
  - Only Transparent UNIs can be assigned to the same ISID

4.4 Flex UNI - Switched

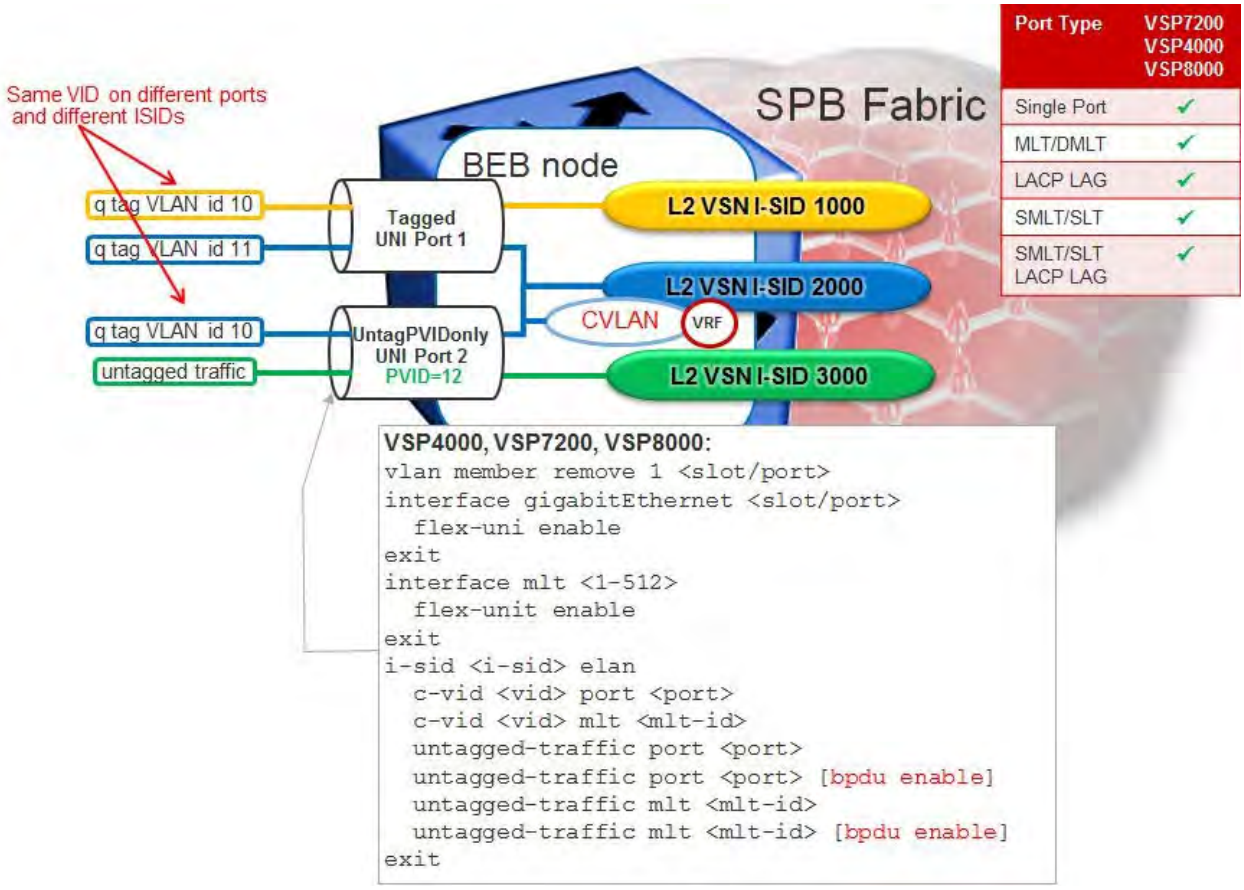

- A VLAN ID (c-vid) and a given port or MLT is mapped to a L2VSN ISID
- The c-vid is not a platform VLAN, it is simply just a VLAN ID on a Flex UNI port or MLT
  - A platform VLAN is a VLAN created using the *vlan create <2-4059> type port-mstprstp <instance>* ACLI command
  - VLAN ID only have local significance on the Ethernet port or MLT
- The same VLAN ID can re-used on different ports and belong to different ISID's
- Switched UNIs and one CVLAN can be assigned to the same ISID
  - This allows you to add an IP address to the CVLAN to enable routing of the Flex UNI
  - To receive untagged traffic, the *untagged-traffic* option must be enabled plus Spanning Tree BPDU's are either dropped or flooded depending on if the *bpdu enable* option is enabled
    - Enabling the untagged-traffic option will forward control traffic such as LACP and VLACP

To enable the BPDU option, one must enable the untagged-traffic setting without the BPDU enable option as shown in the configuration above.

4.5 Private VLAN – ETREE

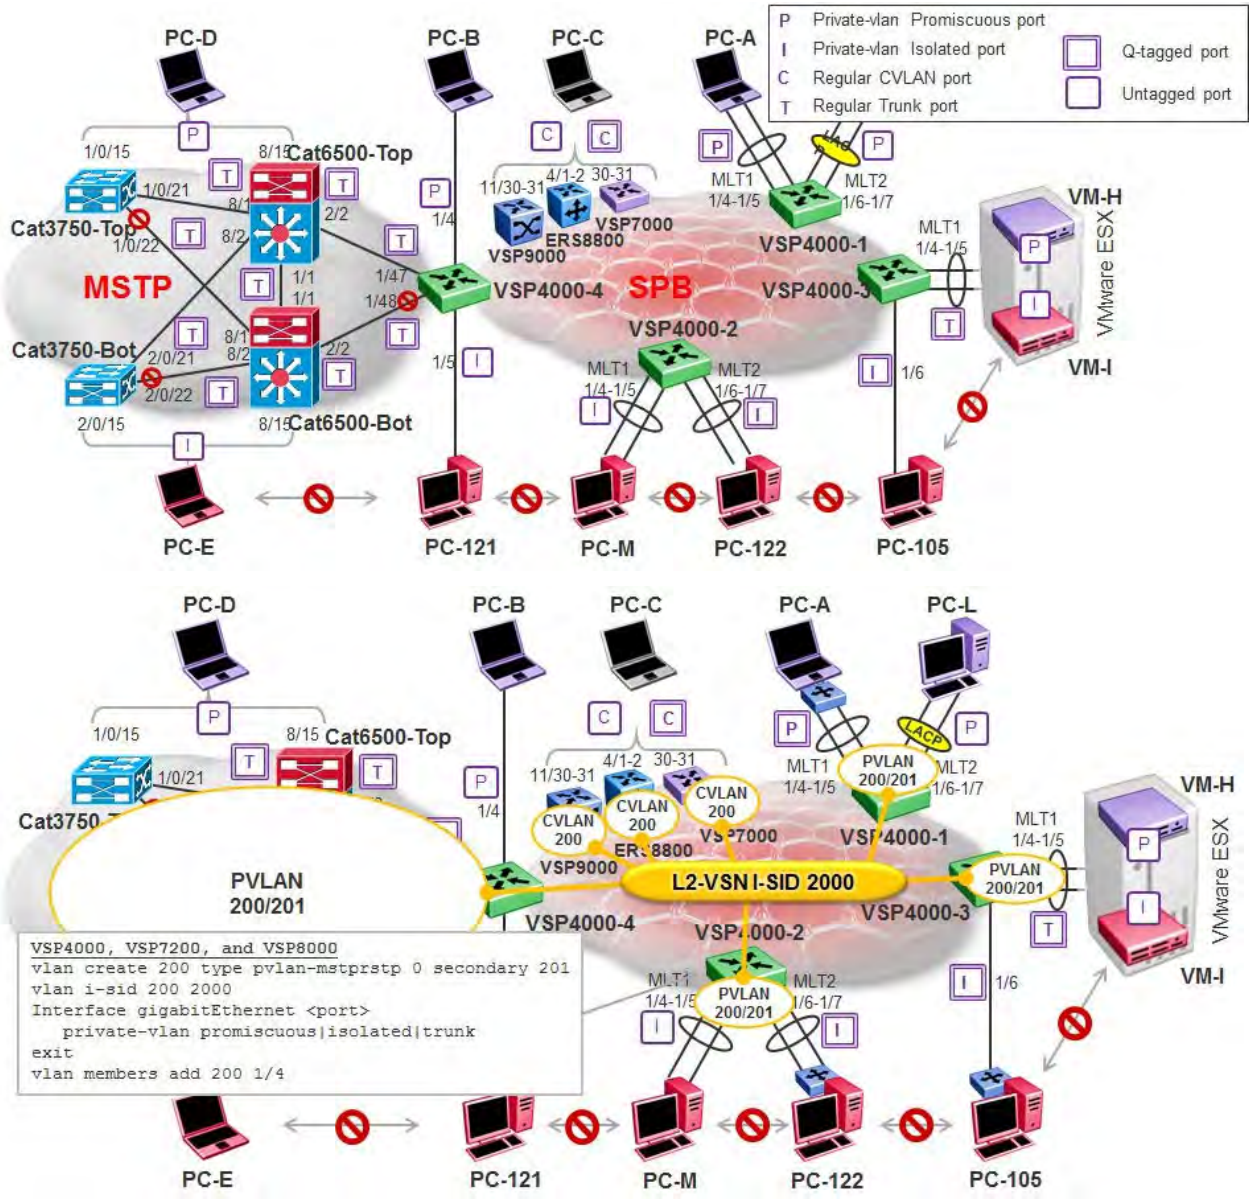

- Private VLAN type of requirement for Layer 2 services only
- Devices connected to Isolated ports (red PCs/VMs) cannot talk to each other
- UNI is a Private VLAN (PVLAN) which allocates 2 VLAN-ids
  - Compatible with Private VLAN as supported by Cisco, VmWare ESX, and many other vendors
- VSP 4000, VSP 8000, and VSP 7200 can associate a PVLAN to an ISID (ETREE L2VSN / like CVLAN UNI type)
- CVLAN devices assigned to same ISID have Promiscuous connectivity within the segment

4.6 UNI Type – Example

UNI Type	Port	VLAN	ISID	
Switched	1	10	1000	Each endpoint is uniquely identified by (Port, VLAN). The same port can send traffic to different ISIDs from different VLANs. The same VLAN can map to one ISID on one port and to another ISID on another.
	1	12	2000	
	2	11	1000	
	2	12	3000	
C-VLAN	All	14	4000	Map entire VLAN to an ISID. All member ports can send and receive traffic to / from ISID.
	All	15	1000	
Transparent	5	All	5000	All traffic from the port that creates the transparent UNI goes to a single ISID, regardless of VLAN.
	10	All	5000	
ETREE	P	PVLAN	6000	P – Promiscuous Ports will see all traffic on PVLAN
	I	PVLAN	6000	I – Isolated ports will only see traffic from Promiscuous port on PVLAN

Table 3: UNI Type

- ISID 1000 will receive traffic from: Port 1 on VLAN 10, port 2 on VLAN 11 and from all port members of VLAN 15.
- ISID 2000 will receive traffic from: Port 1 on VLAN 12
- ISID 3000 will receive traffic from: Port 2 on VLAN 12
- ISID 4000 will receive traffic from: All member ports for VLAN 14.
- ISID 5000 will receive traffic from: All traffic from ports port 5 & Port 10
- ISID 6000 will receive traffic from all ports, but Isolated ports in private VLAN will not communicate with each other; same ISID could be attached to regular VLAN as well and all port within that regular VLAN would behave like promiscuous ports on PVLAN

5. Summary of SPB Features and Product Release Matrix

Capability Feature Matrix	Modular OS		VSP OS (VOSS)		Stackable OS (BOSS)		
	ERS 8800 (7.2)	VSP 9000 (4.1)	VSP 8000 VSP 7200 (5.0)	VSP 4000 (5.0)	VSP7000 (10.4)	ERS4800 (5.9)	ERS5900 (7.0)
L2 VSN	Y	Y	Y	Y	Y	Y	Y
L2 VSN with Multicast (IGMP)	Y	Y	Y	Y	N	Y	Y
L3 VSN	Y	Y	Y	Y	N	N	N
L3 VSN with Multicast (IGMP)	Y	Y	Y	Y	N	N	N
IP Shortcut Routing	Y	Y	Y	Y	N	N	N
IPv6 Shortcut Routing	N	N	Y	Y	N	N	N
IP Shortcut Routing with Multicast	Y	Y	Y	Y	N	N	N
Inter-VSN Routing	Y	Y	Y	Y	N	N	N
IPv6 Inter-VSN Routing	N	N	Y	Y	N	N	N
IPVPN-Lite over SPB	Y	N	N	N	N	N	N
Enterprise Fabric & Switch Cluster Interoperability	Y	Y	Y	Y	Y	N	N
Enterprise Fabric & Stackable Chassis Interoperability	N/A	N/A	N/A	N	Y	Y	Y

Capability Feature Matrix	Modular OS		VSP OS (VOSS)		Stackable OS (BOSS)		
	ERS 8800 (7.2)	VSP 9000 (4.1)	VSP 8000 VSP 7200 (5.0)	VSP 4000 (5.0)	VSP7000 (10.4)	ERS4800 (5.9)	ERS5900 (7.0)
Enterprise Fabric Connectivity Management (802.1ag)	Y	Y	Y	Y	Y	Y	Y
CFM, L2 Ping, Traceroute, and Tracetree	Y	Y	Y	Y	Y	Y	Y
BCB Mode (NNI-NNI)	Y	Y	Y	Y	Y	N	N
L2 Ping for Access VLAN (CVLAN)	Y	Y	N	Y	N	N	N
Switched UNI	N	tbd	Y	Y	Y	Y	Y
Transparent UNI	N	N	Y	Y	Y	N	N
ETREE	N	N	Y	Y	N	N	N
vIST	N	N	Y	Y	N	N	N
Fabric Attach Server	N	N	Y[1]	Y[1]	N	N	N
Fabric Attach Proxy Standalone	N	N	N	N	N	Y	Y
Fabric Attach Proxy	N	N	N	N	N	Y	Y
Fabric Extend	N	N	Y	Y[2]	N	N	N

Table 4: SPB Features and Product Release Matrix

[1]Fabric Attach Server with standalone and SMLT support

[2]Reguires Open Network Adapter (ONA)

6. SPB Feature and License Matrix

Capability Feature Matrix	Modular OS		VSP OS (VOSS)		Stackable OS (BOSS)		
	ERS 8800 (7.2)	VSP 9000 (4.1)	VSP 8000 VSP 7200 (5.0)	VSP 4000 (5.0	VSP7000 (10.4)	ERS4800 (5.9)	ERS5900 (7.0)
L2 VSN	Premier	Base	Base	Base	Base	Base	Base
L3 VSN	Premier	Premier	Premier	Premier	N/A	N/A	N/A
IPv4 Shortcuts	Premier	Base	Base	Base	N/A	N/A	N/A
Multicast L2 VSN	Premier	Premier	Base	Base	Base	Base (5.9)	Base
Multicast L3 VSN	Premier	Premier	Premier	Premier	TBD	TBD	TBD
Multicast IP Shortcuts	Premier	Base	Base	Base	TBD	TBD	TBD
BCB Mode (NNI-NNI)	Premier	Base	Base	Base	Base	N/A	N/A
Inter-ISID Routing	Premier	Premier	Base	Base	N/A	N/A	N/A
ETREE	TBD	TBD	Base	Base	TBD	TBD	TBD
Switched UNI	N/A	TBD	Base	Base	Base	Base	Base
Transparent UNI	TBD	TBD	N/A	Base	Base	TBD	TBD
VRF support	Premier	Base	Base	Base	N/A	N/A	N/A
Dual homing into a Fabric (SMLT Edge)	Premier	Base	Base	Base	Base	TBD	TBD
IPv6 Shortcuts	TBD	TBD	Base	Base	N/A	N/A	N/A
IPv6 Inter-ISID Routing	TBD	TBD	Base	Base	N/A	N/A	N/A
Fabric Extend	N/A	N/A	Premier	Premier	N/A	N/A	N/A
Fabric Attach	N/A	N/A	Base	Base	N/A	Base	Base

Table 5: SPB Feature and License Matrix

7. Scaling

SPB Capabilities	Modular OS		VSP OS (VOSS)			Stackable OS (BOSS)		
	ERS 8000 (7.2)	VSP 9000 (4.1)	VSP8000 (5.0)	VSP7200 (5.0)	VSP4000 (5.0)	VSP7000 (10.4)	ERS4800 (5.9)	ERS5900 (7.0)
SPBM Operational Mode	Chassis	Chassis	Unit	Unit	Unit	Stack	Stack	Unit
[1]SPBM Nodes per region	500	1000	2000	2000	2000	1000	450	450
ISIS adjacencies	64	128	250	250	250/24[2]	24	4	4
Logical ISIS adjacencies	-	-	250	250	250/24[5]	-	-	-
SPBM L2VSNs	2000 / 1700[3]	4000	4059	4059	1000	500 / 800[4]	500	500
SPBM L3VSNs	255	511	24	24	24	N/A	N/A	N/A
SPBM Transparent-UNI	N/A	N/A	N/A	N/A	48	500[5]	N/A	N/A
E-Tree (private VLANs)	N/A	N/A	4059	4059	1000	N/A	N/A	N/A

Table 6: Scaling

[1] The total number of nodes per region is reduced by 1 for each SMLT cluster; an additional virtual B-MAC is created for each time a SMLT cluster is created. For the VSP 4000, the maximum scaling to remote BEB nodes for ISIDs is 2,000 while for the VSP 8000 it is 512 nodes

[2] Up to 250 logical ISIS adjacencies for the VSP 4850 and 24 for the VSP 4850

[3] 2000 L2VSNs without SMLT, 1700 with SMLT

[4] 500 for stacks with or without SMLT and standalone with SMLT. 800 for standalone without SMLT.

[5] SMLT/SLT with Transparent UNI is supported in release 10.4

On VSP 4000, 7200 and 8000 the following table should be used to determine the number of ISIDs supported per BEB. I-SIDs are used for Layer 2 VSN, Layer 3 VSN, Transparent-UNI, E-Tree, and Multicast.

Number of ISIS interfaces (NNIs)	VSP 4000 (5.0)		VSP 7200 (5.0)		VSP 8000 (5.0)	
	vIST used	vIST not used	vIST used	vIST not used	vIST used	vIST not used
4	1000	1000	4000	4000	4000	4000
6	1000	1000	3500	4000	4000	4000
10	650	1000	2900	4000	4000	4000
20	350	700	2000	4000	2450	4000
48	150	300	1000	2000	1100	2200
72			750	1500	800	1600
100					600	1200
128					475	950

Table 7: ISIDs per BEB for VSP 4000/7200/8000

8. Migration & Upgrades

This section describes the procedures and restrictions that apply when upgrading the software load from a prior ERS/VSP software release not supporting SPB. Also described are the procedures to follow when services are being migrated to a configuration that exercises the SPB features. These should be interpreted as additional and NOT as a replacement for procedures and restrictions that may be imposed by prior releases.

8.1 Common Upgrade instructions

- Verify that the hardware requirements are met.
- If the switch is an ERS 8800 and uses 2 CPU cards – both CPU cards need to be rebooted. Using 2 CPU cards with each CPU running a different release of the software is not supported.

8.2 Upgrade from Pre-5.1 releases for the ERS 8800

If the switch being upgraded is not an IST switch - SPB does not impose any additional upgrade procedures. If the switch being upgraded is an IST switch – then both IST peers need to up-graded simultaneously. Standard SMLT resiliency for services is not available until both the IST switches are up and running with the new version of software.

8.3 Upgrade and SMLT Cluster

If the switch being upgraded is not an IST switch - SPB does not impose any additional upgrade procedures. If the switch being upgraded is an IST switch – then it is possible to upgrade one IST peer at a time while providing SMLT based resiliency to services configured on the IST peer switches. While SMLT resiliency is provided during the upgrade – it is recommended that the both the IST peers should be upgraded in a single maintenance window.

8.4 VSP 7000

The core of the Extreme VSP 7000 is a fifth generation Layer 3 Switching ASIC rated at 1,280Gbps. This provides the Extreme VSP 7000 with incredible capacity to support wire speed I/O and Extreme FI (Fabric Interconnect) Stacking concurrently.

The VSP 7000 delivers a new take on the traditional Top-of-Rack Switch requirement. For modest scenarios, switches can be horizontally interconnected, creating a single logical system spanning eight units/racks, or hundreds of VSP 7000s can be flexibly meshed for massive scale-out that uniquely delivers multi-hop and low-latency. Forming a single-tier, Extreme's Distributed ToR is a connectivity solution for the Data Center's primary requirement: high-performance, low- latency, Layer 2 east-west traffic. Utilizing the high-speed virtual backplane capacity, and invoking Ethernet's plug & play advantage, the VSP 7000 empowers simplified, one-touch, edge-only provisioning.

Fabric Interconnect can be used in two mutually exclusive modes:

- **Fabric Interconnect Stacking** where the rear ports are set as Fabric Interconnect Stack-mode. Up to 8 units create a vToR (virtual Top of Rack) or 16 units in a dToR (distributed Top of Rack) delivering up to 10Tbps using two SMLT clusters of 8 switches. For Fabric Interconnect Stack, the stack operates in the same manner as other Extreme stackable products and features many of the associated benefits of stacking (single IP address for FI stack, hot swap unit replacement, and distributed uplinks with distributed MLT and LAGs). By default, Fabric Interconnect (FI) ports on the rear of the VSP 7000 are configured for Stack-mode.
- **Fabric Interconnect Mesh** where the rear ports are configured as "rear-port" in either Standard (Raw) or SPB modes in which the Fabric Interconnect ports operate as multiple high-speed interconnects, allowing the creation of a fully flexible and scalable network mesh. Depending on the software release, SPB and/or SMLT is supported on the rear-ports as highlighted in the chart below. Standard is a Raw-mode that can support various port configurations and protocols, such as Inter-Switch Trunking (IST) for Switch Clustering and SMLT. Please note that rear-port Standard Mode does not support SPB on the rear ports.

Desired Deployment Model	Needed Rear-port Mode	SMLT (IST) Needed	SPB enabled	Virtual Servers (e.g. ESX) NIC teaming	Server NIC teaming (LACP)	Minimum Required Software
vToR FI Stacking	**Disabled** (= Stacking enabled)	**No**	**Can be**	**Yes** – Vport hashing on non-SLT ports	**Yes** – On DMLT ports (with or without LACP)	10.1.0 (10.3.0 if need to run SPB on uplinks)
dToR FI Stacking with SMLT	**Disabled** (= Stacking enabled)	**Yes**	**Can be**	**Yes** – Vport hashing on non-SLT ports - OR - IP hashing on SLT ports	**Yes** – On SLT ports (with or without LACP)	10.2.0 (10.3.0 if need to run SPB on uplinks & IST)
FI Mesh with SMLT	**Enabled in raw mode**	**Yes**	**No**	**Yes** – Vport hashing on non-SLT ports - OR - IP hashing on SLT ports	**Yes** – On SLT ports (with or without LACP)	10.2.0
SPB Mesh	**Enabled in SPBM mode**	**No**	**Yes**	**Yes** – Vport hashing on non-SLT ports	**No** – Use Active Standby NICs	10.2.0
SPB Mesh with SMLT	**Enabled in SPBM mode**	**Yes**	**Yes**	**Yes** – Vport hashing on non-SLT ports - OR - IP hashing on SLT ports	**Yes** – On SLT ports (with or without LACP)	10.3.0

Table 8: VSP 7000 Rear Port Mode

In Fabric Interconnecting Stacking, with SPB enabled, 10.2.1 supports a maximum stack of 2; in the 10.3 release, a stack of 8 is supported. SMLT or IST over rear port Raw-mode is supported starting in release 10.2.1, however, SPB is not supported in rear port Raw-mode. LACP must be disabled on the rear ports prior to enabling an IST on them. In rear port SPB mode, IP routing cannot be enabled and the total number of rear ports is reduced by one - the switch uses port 40 as loopback when rear port SPB is enabled; please see section 9 for rear port numbering.

8.5 Activating SPB

Once the network is upgraded the following minimum steps must be followed before any services can be provisioned. SPB leverages the usage of default parameters and link metrics, system-id values etc., to minimize the number of configurations steps. Customers that desire to use non default parameters should do so in accordance with the configuration and engineering guidelines.

- ERS 8800
    - Activating SPB infrastructure reserves 600 multicast group ID (MGIDs) for SPB operation on the ERS 8800. This is in addition to any MGIDs that may be used for VLANs and IP multicast services. The “show sys mgid-usage” command should be used to check if the MGIDs required for SPB are available.
    - If the ERS 8800 is running in STG mode, verify that STG-63 is not current in use. SPB will use this STG.
    - If running in MSTP mode – verify that MSTI-62 is not currently in use. SPB will use this MSTI.
- VSP 9000
    - Activating SPB infrastructure reserves 100 multicast group ID (MGIDs) for SPB operation on the VSP 9000.
- All models
    - SPB is not supported if the switch is in RSTP mode. Note that there is no reason to use the RSTP mode since it provides a sub-set of the functionality of MSTP mode and MSTP mode is able to operate in RSTP mode if it sees adjacent switches sending RSTP BPDUs.
- Identify two VLAN-ids to be used as B-VLANs by SPB, primary and secondary B-VLAN
    - Note, the same primary and secondary VLAN IDs must be provisioned on all SPB enabled switches so that all SPB bridges will load balance traffic accordingly
    - The IS-IS adjacencies will not come up if there is a discrepancy in the B-VLAN ids configuration between 2 nodes.
    - The Primary VLAN IDs is also used on all IS-IS messages
- Enable SPB globally.
- Assign a unique nickname to each switch.
    - An alarm will be logged if a duplicate nickname is provisioned in the network
- Assign a common Area ID
    - Note, the same Area ID must be provisioned on all SPB enabled switches in the same domain
- Assign a unique IS-IS system-name to each switch. While this is not strictly required – it will greatly aid in validating connectivity and when troubleshooting.
    - If the IS-IS sys-name is not provisioned, by default, the global system name is used as the IS-IS sys-name. If you do wish to set the IS-IS sys-name, it must be set to a value different than global system name.
- If configuring an IST switch, configure the system-id of the IST peer.
- Identify all the intended NNIs and configure and enable IS-IS on these ports (or MLTs).

  - Please note that only one adjacency is supported between a pair of SPB bridges (one physical port or one MLT instance)

- Enable IS-IS globally.
- Configure IEEE 802.1ag (a.k.a CFM) in order to enable network connectivity troubleshooting tools.
- Verify basic SPB connectivity by checking the SBPM unicast-fib and the FDB entries for the B-VLANs.
- Verify basic SPB unicast connectivity using the l2ping and l2traceroute commands between all the switches in the network for both the B-VLANs.
- Verify that the path reported by the l2traceroute command is the same as the one calculated by IS-IS (use the *show isis spbm unicast-fib* command).
- SMLT Operation

ERS 8800 & VSP 7000	VSP 9000
Does not require you to configure C-VLANs on the IST MLT.	Requires the inclusion of IST MLT in the C-VLAN.
Traffic can pass between single-homed VLANs attached to IST peers if the IST is down.	Traffic cannot pass between single-homed VLANs attached to IST peers if the IST is down.
Decapsulates MAC-in-MAC traffic at the primary BEB or secondary BEB irrespective of whether the traffic is from the primary BVLAN or secondary B-VLAN.	Decapsulates MAC-in-MAC traffic at the primary BEB from the primary B-VLAN and traffic at the secondary BEB from the secondary B-VLAN. Requires the IST to be up to pass traffic between both IST switches for single-homed VLANs.

- vIST Operation
  - The vIST uses an SPBM tunnel to virtually connect two nodes that can be anywhere in the SPBM network
  - A L2VSN is required for the vIST VLAN
    - An IP address is added to the VLAN as you would do with a normal IST VLAN and you need to peer with the IP address to the neighboring IST switch
  - All C-VLANs must have an ISID
    - Before you add the SMLT MLT instance or SMLT port members to a C-VLAN, an ISID must be assigned to the C-VLAN
      - If you try to add a C-VLAN to an SMLT MLT instance or port members that become to an SMLT instance prior to adding an ISID to the VLAN, you will be prompted with the following error message: *Error: SMLT VLAN must have ISID association.*

8.6 Migrating traffic to SPB

Pre migration checks for configuration migration to SPB should include an audit to determine if the desired configuration and traffic is something that is supported by SPB. The following kinds of traffic are supported by SPB.

- Layer-2 bridged traffic.
- IPv4 unicast routed traffic on the Global Router.
- IPv4 unicast routed traffic using a VRF.
- IPv4 Unicast routed traffic using an IPVPN (ERS 8800 only).
- IPv4 multicast routed traffic. ERS8000 software release 7.2, VSP 4000 software release 3.1, VSP 9000 release 3.4 and VSP 72000/8000 release 4.1 add support for L2VSN, L3VSN and IP Shortcut.
    - If a PIM router is connected to an SPB bridge, either use IGMP or static IGMP entries

The following traffic is supported on selected platforms – please refer to chapter 6.

- IPv6 unicast routed traffic

Traffic which is not yet supported by SPB can continue to exist in parallel to SPB. Please note that only the ERS 8800 and VSP 9000 can support SPBM in addition to a classical SMLT deployment.

8.7 Multicast

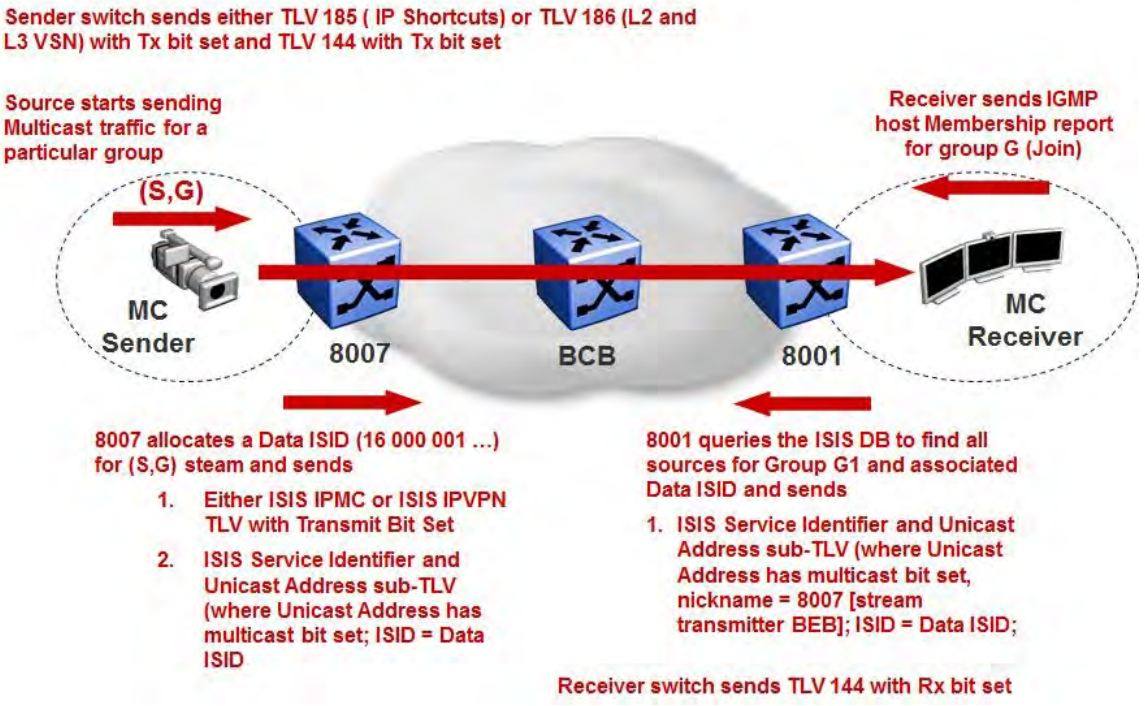

Multicast over SPB for L2 VSN, IP Shortcuts, and L3 VSN is supported in the 7.2 release for the ERS 8800, 3.1 for the VSP 4000, 4.1 for the VSP8000, 4.2.1 for the VSP 7200, and the 3.4 release for the VSP 9000. Multicast over SPB for L2 VSN is supported in release 5.9 for the ERS4800 and release 7.0 for the ERS 5900. If the VSP 9000 is used in the network, ensure that it is operating at release 3.3.x or higher.

- SPB multicast supported over L2VSN
  - Simple provision by enabling SPB multicast globally and IGMP snooping at L2VSN VLAN level
  - Traffic does not cross L2VSN service boundary
  - By default, on a BEB UNI port, an IGMP querier address of 0.0.0.0 will be used. If the L2 edge switch does not support a 0.0.0.0 query address, any IP address can be provisioned the L2VSN VLAN as the query address
  - Any device in a L2VSN boundary can start a multicast stream
    - Note that you can still use any of the various IGMP features on the L2VSN VLAN such as allow or deny certain IGMP group addresses
  - Single-Homed BEB hashes between the two BVLAN's based on the ISID; odd ISID transmitted on Primary BVLAN, even ISID transmitted on Secondary BVLAN
  - SMLT BEB transmit on a single BVLAN; Primary SMLT BEB on primary BVLAN and Secondary SMLT BEB on Secondary BVLAN
- SPB multicast supported over L3VSN
  - All or a subset of VLANs within a L3VSN can exchange IP multicast traffic between themselves

- Simply provision by enabling SPB multicast globally, enable MVPN on the VRF, and enable IP SPB Multicast on some or all of the VLANs within the L3VSN
  - Only those VLANs that have IP SPB Multicast enabled can pass multicast traffic
- It is not a requirement to enable IP Shortcuts in order to support IP Multicast in the L3VSN
- IPVPN creation and ISID assignment for the IPVPN is required – but the IPVPN does not need to be enabled
- Any device in the L3VSN can start a multicast stream
  - Note that you can still use any of the various IGMP features on the VRF VLAN such as allow or deny certain IGMP group addresses
- SMLT operation – does not apply to the ERS 4800 or ERS 5900
  - On the Primary IST BEB
    - Primary BVLAN traffic is forwarded to the SMLT and Single Homed UNIs
    - Secondary BVLAN traffic is forwarded only to Single Homed UNIs
    - With SMLT Down on the Primary IST BEB, the primary IST BEB does not forward any primary BVLAN traffic to the SMLT
    - With SMLT Down on the Primary IST BEB, the Secondary IST BEB forwards both primary and secondary BVLAN traffic to the SMLT
  - On the Secondary IST BEB
    - Primary BVLAN traffic is forwarded only to Single Homed UNIs
    - Secondary BVLAN traffic is forwarded to the SMLT and Single Homed UNIs

- IP Shortcuts (GRT) with IP Multicast
  - All or a subset of VLANs within GRT can exchange IP multicast traffic between themselves
  - Simple provision by enabling SPB multicast globally and enabling IP SPB Multicast on some or all of the GRT VLANs
    - Only those VLANs that have IP SPB Multicast enabled can pass multicast traffic
  - It is not a requirement to enable IP Shortcuts in order to support IP multicast in the GRT using SPB
  - Any device in the GRT can start a multicast stream
    - Note that you can still use any of the various IGMP features on the GRT VLAN such as allow or deny certain IGMP group addresses

8.8 Migrating a VLAN to an L2 VSN

The following procedure can be used to provide L2 connectivity for a VLAN across the SPB core.

- Follow the pre-migration procedures checks described in the section “Common Procedures and Exclusions on Migration”
- Identify the UNI and NNI ports that are currently port members of the VLAN on all the switches in the network.
- On all the switches in the network which are currently connected by the VLAN – remove the NNI ports from the membership list of the VLAN. This step will cause service interruption.
- Make the VLAN an L2VSN using the “*vlan <vlan id> ISID <isid value>*” ACLI command. Use the same value of ISID on all the switches. This step should restore service.
- SMLT deployment
    - For the ERS 8800 & VSP 7000, the L2VSN VLAN cannot be a member of the IST. An error message will be recorded and logged if you try to add VLAN to the IST MLT instance
    - For the VSP 9000, the L2VSN VLAN must be a member of the IST. A warning message to this effect will always be displayed when an ISID is assigned to a VLAN.
    - For a vIST, for the L2VSN C-VLAN, only the local SMLT ports are added. There is no physical IST as the vIST is virtualized.

8.8.1 Migrating to Inter VSN Routing

Inter VSN provides the ability to route traffic between extended VLANs where the VLANs have different ISIDs. All of the traditional IPv4 unicast routing and gateway redundancy protocols (OSPF, RIP, BGP, VRRP, RSMLT etc) are supported on top of any VLAN that is mapped to an ISID. Please note that RSMLT will only work if the switches acting as redundant gateways are IST connected together.

Currently the only protocols which will not work on an IP interface assigned to a L2VSN VLAN are the following:

- VSP 4000, VSP 7200, VSP 8000
    - IPv6 multicast routing (MLD)
    - IPv4 multicast routing (IGMP, PIM-SM, PIM-SSM)
- VSP 9000, ERS 8800
    - IPv6 unicast & multicast routing (OSPFv3, MLD)
    - IPv4 multicast routing (IGMP, PIM-SM, PIM-SSM)

Please note that RSMLT will only work if the switches acting as redundant gateways are IST connected together.

The high level procedure to migrate a configuration to use Inter-ISID routing is described below.

- Follow the pre-migration procedures checks described in the section “Common Procedures and Exclusions on Migration”
- For each VLAN in the SPB core
    - On all the switches where the VLAN is configured - remove all NNI ports
        - This will cause service interruption.

- On all the switches where the VLAN is configured – map the VLAN to an ISID. This will restore L2 connectivity (the l2tracetree command can be used to validate L2 connectivity within the VLAN at this point). L3 will be restored once the routing protocols configured on top of the VLAN converge.
- Once all the VLANs identified for migration have been assigned an ISID – the configuration part of the migration is completed. At this point all the traffic flows should be back to normal.

8.9 VSP 9000 Notes

- VSP 9000 supports SPB NNI Interfaces on the 9024XL, 9048XS-2, and 9012QQ-2 cards.
- For L2VSN services on an IST switch
  - If a L2VSN is configured on one IST switch, it must be configured on the peer IST switch as well (even if the IST peer has no UNI ports using the L2VSN).
  - The IST ports must be configured as member ports of the VLAN which is using the L2VSN; on the ERS 8800, the opposite is true, the IST ports must be removed from the VLAN using a L2VSN
  - You will see the messages below during configuration.
    - CAUTION: Adding ISID to a VLAN on an IST switch requires configuring this ISID-VLAN pair on both IST peers and the IST MLT must be a member of the VLAN.
    - CAUTION: All VLANs with ISIDs MUST be configured on both IST peers and IST MLT MUST be a member of all these VLANs.
- CFM simplified configuration is supported starting in release 3.4
- IPVPN-Lite over SPB is not supported in VSP 9000
- Multicast support for L2VSN, L3VSN, and IP Shortcuts is supported in the VSP 9000 3.4 release

9. Field Introduction & Support Specifications

9.1 Hardware and Deployment Specifications

Product	Specifications
ERS 8800	• Line Cards – R, RS, and 88xx modules • 8692SF with Supermezz or 8895SF
VSP 9000	• Only the 10G & 40G capable modules support SPB NNI Interfaces; any of the other line card modules can be used as UNI ports.
VSP 4850GTS(-PWR+)	• usage of conversion kit (EoS Sep 2015) with models that have HW rev 10 or higher • Not stackable
VSP 4450GSX-PWR+	• not convertable to ERS • not stackable
ERS 4800	• not limited to Rev 10 hardware (and higher) • can be used in stack configuration • loses all L3 features, when SPB is enabled • No BCB • L2 VSN w/ IGMP in Rel. 5.9 (sacrifice one SFP+ or both stacking ports)
VSP 7000	• the CVLAN cannot be configured on any NNI interface • loses all L3 features, when SPB is enabled • do not configure VLACP and „filter untagged frames“ on IST
ERS 5900	• Stackable • L2 VSN w/ IGMP (sacrifice two SFP+ or stacking ports) • loses all L3 features, when SPB is enabled • No BCB
VSP 7200	• All ports can be used as NNI or UNI
VSP 8000	• All ports can be used as NNI or UNI

9.2 Installation and Commissioning Specifications

Please check the section on upgrades and migration for information on impact on existing features when SPB features are enabled.

9.3 Interoperability and Backwards / Forward Compatibility Specifications

For the ERS 8800 only, new SPBM 802.1aq TLVs have been defined by IANA after the 7.1.0.0 release. Release 7.1.0.x and 7.1.1.x both used pre-standard (draft) TLVs. In release 7.1.3.0, both the pre-standard (draft) and new 802.1aq standard TLVs are supported. Since release 7.2 for the ERS 8800, only the new 802.1aq standard TLVs are supported.

10. VSP 7000 – Fabric Interconnect

The VSP 7000 by default operates in Fabric Interconnect stacking mode. The VSP 7000 can be provisioned in rear-port mode where the rear Fabric Interconnect ports will be treated as multiple virtual ports over the 4 physical Fabric Interconnect Ports. When in rear-port mode, the VSP 7000 operates in a standalone mode.

Two modes of operation are available in rear-port mode, standard or Shortest Path Bridging (SPB). Standard mode allows all the switch standard features minus SPB across the rear ports, i.e. Spanning Tree, OSPF, RIP, etc. In SPB mode, in the 10.2 release Shortest Path Bridging is supported while in the 10.3 both SPB and SMLT are supported. Hence, when FI Mesh is required, rear-port mode with operational state of SPB needs to be provisioned. The diagram shows the FI port speeds available depending if Standard or SPB operational state is enabled.

To provide greater plug n 'play capability over the virtual ports when rear-port mode is enabled, LACP link aggregation and VLAN tagging are automatically enabled. This ensures that multiple virtual ports which may run within a single cable or if multiple FI cables are run in parallel that all virtual ports are automatically treated as one link. This simplifies any protocol adjacency such as IS-IS or OSPF. When you issue rear-ports mode all virtual ports will have their LACP state set to true, the LACP Admin Key to 4095 and LACP hashing mode be set to advance.

Color	Physical Fabric Interconnect Port	Rear Port Mode	Throughput	Ports
Black	FI Up (right) Top	Standard	240Gbps (x3 40GbE)	34, 35, 36
		SPB	240Gbps (x3 40GbE)	
Red	FI Down (left) Top	Standard	240Gbps (x3 40GbE)	38, 39, 40
		SPB	160Gbps (x2 40GbE)	38, 39
Blue	FI Up (right) Bottom	Standard	80Gbps (x1 40GbE)	33
		SPB	80Gbps (x1 40GbE)	
Blue	FI Down (left) Bottom	Standard	80Gbps (x1 40GbE)	37
		SPB	80Gbps (x1 40GbE)	

Figure 7 – FI Rear Port Details

In FI mesh, it is recommended to connect "like" color fabrics interconnect ports together, i.e. red port to an adjacent switch red port to get maximum possible throughput. You can connect any color ports together, i.e. a red port to a blue port, however, the port throughput will drop to the lower of two ports.

Rear-port interfaces 33-40 are regular Ethernet 40 GbE interfaces. For some of the rear-ports multiple such 40 GbE interfaces are bundled together. As the rear-ports constitute a backplane

connection their throughput is shown in the table above for both transmit & receive (Full Duplex).

In rear-port SPB operational state, virtual port 40 is not available. Hence, the red port is reduced to 160Gbps.

In rear port mode, the front panel *Up* and *Down* LEDs blink in a quick pattern (125ms) to indicate rear-port mode is operational.

In the 10.2 release, SPBM is officially only supported in rear port SPB mode.

In the 10.2.1 release, SPB is supported in rear port SPB mode or in Fabric Connect Stacking mode (in a stack of two).

For more details, please refer to *Resilient Data Center Solutions Technical Configuration Guide*, publication number *NN48500-645*.

11. ISIS Metrics - Optional

You can configure the link metric to overwrite the default metric value. By configuring the metric, you can specify a preferred path. Low cost reflects high-speed media, and high cost reflects slower media. For the wide metric, the value ranges from 1 to 16,777,215.

- SPBM uses the L1-SPB metric defined in a new SPB sub-TLV
- The total cost of a path equals the sum of the cost of each link
- If a link has different metric values configured at each end of the link, SPBM will use the highest metric value
- The default value for wide metrics is 10

As an option, you can change the wide metric to the suggested values as shown in the table below to allow the switch to prefer the higher speed NNI links over the lower speed links.

Link Speed (Gbps)	Interface Type	ISIS L1-metric
1	Native Ethernet	2000
2	MLT bundle	1000
10	Native Ethernet	200
20	MLT bundle	100
40	Native Ethernet	50
80	MLT bundle / FI	25
100	Native Ethernet	20
120	MLT bundle / FI	17
160	MLT bundle	13

Table 9: ISIS Metric Option

By default, all Fabric Interconnect ports operating in rear-port mode use the same LACP key (4095) on the VSP 7000. If you modify a rear-port metric, such as the SPBM-L1–Metric, the modification applies to all ports which have the same LACP key. If different metrics are to be used on specific rear-ports, you will need to set different LACP keys on those ports.

12. ISIS Accept Policy

Beginning in software release 4.1 for the VSP 4000/8000, 4.2.1 for the VSP 7200, and release 4.0 for the VSP 9000, IS-IS accept policies for IPv4 is introduced. Prior to this release, the ISIS IPv4 routes received over the SPB cloud are installed directly into the routing table. There is no ability to filter those routes and apply incoming route policies to them. Hence, networks that are being migrated from other routing protocols to ISIS/SPB are vulnerable to routing loops. ISIS accept policy functionality provides a way to avoid such loops.

You can create an IS-IS accept policy for the Global Routing Table (GRT) or a Virtual Routing and Forwarding (VRF) instance. You can create an IS-IS accept policy on a switch that operates at a global default level or for a specific advertising BEB. You can also use the filter mechanism for ISIS accept policies to redistribute routes between different VRFs, or between a VRF and the GRT.

ISIS policies can also use either a service instance identifier (ISID) or an ISID list to filter incoming traffic. For inter-VRF route redistribution, an ISID value of 0 represents the GRT. For inter-VRF route redistribution between VRFs, the ISID is the source VRF (or remote VRF).

You can filter traffic with ISIS accept policies by:

- advertising BEB
- ISID or ISID list (ISID take precedence over an ISID list)
- route-map

You can use IS-IS accept policies to apply at a global default level for all advertising Backbone Edge Bridges (BEBs) or for a specific advertising BEB.

IS-IS accept policies also allow you to use either a service instance identifier (ISID) or an ISID list to filter routes. The switch uses ISIDs to define Virtual Services Networks (VSNs). ISIDs identify and transmit virtualized traffic in an encapsulated SPBM frame. IS-IS accept policies can use ISIDs or ISID lists to filter the incoming virtualized traffic.

IS-IS accept policies can also apply route policies to determine what incoming traffic to accept into the routing table. With route policies the device can determine which routes to accept into the routing table based on the criteria you configure, which can give precedence to advertised routes from a particular protocol, route-source, route-type, or through other criteria.

13. ISIS External Metric

Beginning in the software release 5.0 for the VSP 4000/7200/8000, ISIS external metrics is introduced. By default, internal metric is always preferred over external metrics. Hence, when advertising external networks, it is recommended to advertise these routes as external similar to what OSPF does with external type1 and type 2. Hence, when an SPBM ISIS router decides what route to install in the route table, it will always prefer ISIS routes with the shortest internal metric path to the SPBM node advertising them.

With this feature you can use IS-IS to:

- Change the external metric-type of a route when redistributing it from another protocol to IS-IS through route redistribution using a route-map.
- Change the external metric-type of a route when accepting a remote ISIS route with the help of ISIS accept policies using a route-map.
- Match the external metric-type when redistributing ISIS routes into other protocols using the match option in the route-map.
- Match the external metric-type when accepting a remote ISIS route with the help of ISIS accept policies by using a route-map
- Process the external metric-type in the route selection process.

The IS-IS metric type can also be set using the base redistribute command without using the route map.

For configuration examples, please see the section titled "*IP Shortcuts – redistribution of ISIS and OSPF*".

14. SPB over L2/L3 networks

14.1 Supported Networks

Without Fabric Extend	With Fabric Extend
Ethernet Direct Cable	[4]MPLS / Pseudo-Wire / E-Line (L2)
[1]Point-to-Point Ethernet L2 service	[4]MPLS VPLS or PBB ELAN (L3 over L2)
[2]Q-in-Q Tagged Management Services	[4]MPLS IP-VPN
CWDM/DWDM Optical Networks	[4]Campus L3 network
[1]MPLS VPLS	
[1]MPLS Pseudo-Wire	
[3]MPLS IP-VPN via a GRE router	

[1]At minimum, 1544 byte packets must be supported for the additional 22 byte packet overhead for SPB frame
[2]The customer payload must support 1544 byte packets
[3]Where an SPB NNI interface is connected to a router that can encapsulate the two B-VLANs into GRE tunnel
[4]At minimum, 1594 bytes packets must be supported for the additional SPB and VXLAN overhead if the VSP8000 is used as it does not support fragmentation. The VSP4000 with ONA is required for fragmentation for networks that do not support jumbo frames

Without Fabric Extend

The current IEEE 802.1aq SPBM implementations require that all SPBM nodes to form adjacencies over point-to-point links. Point-to-multipoint and broadcast networks cannot be used to establish SPBM adjacencies. Any network to be used for an SPBM transport must be capable of transparently supporting SPBM point-to-point adjacencies with multiple 802.1Q VLANs for each adjacency. Note that Extreme presently supports two VLANs referred to as Backbone VLANs, hence, the backbone network needs to support 2 VLANs to natively transport SPBM across the network in addition to forward an untagged default VLAN for ISIS traffic. Any point-to-point link used to transport SPBM must be capable of supporting at minimum 1544 byte packets as SPBM adds an additional 22 bytes overhead.

If SPBM is to be transported across an IP network such as a private campus network or MPLS IP-VPN service, this can be accomplished in a couple of methods. A router can be used between an SPBM node and the IP network using GRE encapsulation providing end-to-end SPBM support. The router must have the capability of GRE encapsulating the two SPBM backbone VLANs and supporting fragmentation. The limitation here is a separate router is required for each SPBM point-to-point link unless it has the capability of supporting the same two B-VLAN ID's on multiple ports. Another method could be if the router supported Pseudowire Switching, but, typically this is supported across an MPLS access network.

Another method is to redistribute ISIS into another protocol such as OSPF, BGP, Static, or RIP depending on the protocol used in the external network. The disadvantage of this method is you lose end- to-end SPBM plus you need to route policies to avoid routing loops if there are multiple redistribution points.

With Fabric Extend

Extreme's Fabric Extend solution allows extension of Fabric Connect (SPBm) over third party transport networks. It tunnels SPBM traffic over a VXLAN header allowing for end-to-end SPBM networking over

either layer 2 or layer 3 networks. Hub and spoke plus mesh topologies are supported. Fabric extend is supported natively as of release 5.0 for the VSP 8000 and VSP 7200 and with the ONA for the VSP 4000.

Please note the following with using Fabric extend:

- Premier license is required on the VSP 4000, VSP 7200, and VSP 8000 for Fabric Extend
- VSP 7200 and VSP 8000 support VXLAN Tunneling in hardware, no ONA is required
- VSP 4000 does not support VXLAN Tunneling in hardware, hence an ONA is required to do the tunneling
    - One ONA required per VSP 4000 where FE tunnels need to be terminated
- Logical Interface Count per device or port:
    - VSP 4450, VSP 7200, VSP 8000: 250
    - VSP 4850: 24
- Fragmentation and reassembly is required if IP MTU <1594 bytes
    - SPB over IP session headers and encapsulation results in an Ethernet MTU size of 1594 bytes
    - VSP 4000 + ONA combination is required for fragmentation and reassembly wherever FE tunnels need to be terminated
    - Fragmentation and reassembly is done by the ONA
    - VSP 7200 and VSP 8000 do not support fragmentation and reassembly
- Combining VSP 7200 and VSP 8000 with ONA is not supported
- Tunnels can originate and terminate on any of the hardware combinations below if fragmentation and reassembly is not a requirement:
    - (VSP 4000 + ONA), VSP 7200, and VSP 8000
- Optional Fabric Extend Manager with Extreme Fabric Orchestrator can be used to manage tunnel configuration in an FE-domain

Latency

Fabric Connect itself does not pose any stringent latency requirements on an NNI link. NNI links can stretch thousands of miles across the globe, as long as the physical Ethernet or emulated Ethernet integrity is guaranteed. IS-IS timers are typically very long (multi-seconds) thus, a link won't time-out due to extended distance between the NNI ports. If packet loss occurs on the Ethernet links, then application layers will have to retransmit packets. If excessive packet loss occurs, then links might drop due to missed IS-IS hello packets. The effect of latency on upper layers protocol and applications should be borne in mind.

14.2 Fabric Extend Solutions

Fabric Extend enables Enterprises to extend Extreme Fabric Connect technology over Layer 2 or Layer 3 core networks. With the introduction of logical IS-IS interfaces, new WAN solution deployment options will be available that allow SPB Fabric to run over IP MPLS VPNs, or to aggregate VLAN tunnels (Pseudowire-MPLS or PBB E-Lines) to be leveraged as wide area hub-and-spoke connectivity/topology models.

Components

Fabric Extend is supported on the VSP 8000, VSP 7200, and VSP 4000 platforms. For the VSP 4000, it will require an Open Network Adapter (ONA) depending on whether Layer 3 tunnels are required or not.

The number of logical interfaces supported on the VSP 8000, VSP 7200, and VSP 4450 is 250 while the VSP 4850 support 24.

Open Network Adapter (ONA-1101GT)

The ONA-1101GT provides IP tunneling Data Path capabilities for the VSP 4000. It is required for the VSP 4000 as it has not IP tunneling capabilities.

The VSP 4000 manages the ONA in the following ways:

- Controls and provisions the ONA.
- If PoE capable, the VSP 4000 supplies power to the ONA. (The ONA also supports an optional wall unit power adapter.)
- Transports traffic to and from the ONA over 1GbE ports and sets QoS appropriately to the ONA's.
- The ONA 1101GT can support basic Extreme Fabric Extend at line rate 1G traffic from the VSP 4000 at 1500 byte packet sizes.
- Oversubscription of the ONA's packet engine may result if packets are smaller than 1500 bytes or if you enable enhanced features such as fragmentation and reassembly of packets. This results in packet drop starting with lower QoS queued packets consistent with PCP and DSCP markings on packets received from the VSP 4000.

The ONA can operate in two different modes. SDN controller is Operational Mode 0 and is the default mode. Fabric Extend is Operational Mode 1 which is configured via the web configuration menu. The web configuration menu can be accessed by holding down the reset button during powering on the ONA; the ONA also provide an IP address in the 192.168.100.x/24 via DHCP to the local attached PC on the device side.

In regards to IP addressing for the ONA, this can be accomplished via DHCP or adding a static IP address. If using DHCP, the ONA requires access to a local DHCP server to automatically configure IP addresses.

Please see the ONA configuration section below for more details regarding configuration.

Fabric Extend over Layer 3 Networks

Over a Layer 3 core network, Fabric Connect can be extended by IP tunneling using VXLAN encapsulation. The VSP 4000 will require an ONA as it does not support native IP tunneling.

For IP tunneling, a VXLAN header is added to the SPBM packets. Fabric Extend in this case is supported over third party IPv4 transport networks such as MPLS IP-VPN or Campus IP backbones. Or Fabric Extend can be transported over a Layer 2 MPLS VPLS or PBB ELAN service by creating layer 3 tunnels over a L2 third party network. An Open Network Adapter (ONA) is required when using the VSP 4000 while the VSP 8000 or VSP 7200 supports Fabric Extend natively.

MTU Configuration

The VSP 8000 and VSP 7200 default MTU size is 1950 and cannot be changed. When using the VSP 4000 with the ONA, an MTU size from 750 to 1950 is supported with a default setting of 1950. For the VSP 4000 to work with a VSP 8000 or VSP 7200, the MTU size must be left at the default setting of 1950.

If the core network does not support jumbo frames, the VSP 4000 with ONA must be used on all sites. A suitable MTU size must be provisioned else an ISIS adjacency over the Fabric Extend tunnel cannot be established; please note a log message will be recorded in this case. Also note that fragmentation with the ONA is only supported over layer 3 networks and not supported over layer 2 networks.

Fabric Extend over Layer 2 Networks

Over a Layer 2 core network, Fabric Connect can be extended either by Fabric Extend by either IP tunneling or VID-tunneling.

VID-tunneling is simply encapsulating the Backbone VLANs into provider VLANs. VXLAN is not used in this scenario. As Extreme supports two B-VLANs presently, for each point-to-point connection, i.e. hub site to a spoke site, two provider VLANs will be required. Hence, for two branch sites, four provider VLANs will be required at the hub site, two each for each spoke site. For VID-tunneling, an ONA is not required for the VSP 4000.

Instead of VID-tunneling, IP tunneling can be used where only one provider VLAN is required. For the VSP 4000, an ONA will be required. Please note that if using the VSP 8000 or VSP7200, it requires a single next hop for all tunnels as explained below.

VSP8000 and VSP7200 Fabric Extend Tunnel Next Hop Limitation

The VSP8000 and VSP7200 require a single next hop (default gateway) for all tunnels. Over a L3 core network, there is no issue as there is one router next hop over which we can support multiple VXLAN tunnels to one or more remote sites. However, for example, if we wish to support L3 tunneling over a L2 core, the VSP7000 or VSP8200 without any specific configuration will only be able to support one Fabric Extend tunnel to one remote site via a given port. To get around this issue, we can create an additional VRF, VLAN and loopback interface, to get around the single next hop issue for all tunnels through a single port.

QoS over Fabric Extend

On the logical NNI interface used by Fabric Extend, QoS is maintained over both the 801.1p bit and outer IP Header DSCP value.

Failover behavior

Tunnel status & failure detection are monitored by ISIS link hello timers of 27s (hello interval 9s and hello multiplier 3). VLACP cannot be used on logical interface connections. BFD (Bidirectional Forwarding Detection) is planned for tunnel status monitoring.

Default Metrics

The default metrics for all FE logical interfaces is kept higher than port based ISIS interfaces. If there are two connections between a pair of BEB nodes, the port based ISIS interface is preferred and will be used to bring up the ISIS adjacency. The logical ISIS interface will then be used as a backup in the case the port based ISIS should go down.

The default ISIS metric for FE logical interfaces are as follows:

- For IP core, the default metric is 20,000 that can be manually changed
- For Layer 2 core, the default metric is 1,000 that can be manually changed

Fabric Extend (FE) considerations

- vIST session over Fabric Extend logical interface is not supported in software release 5.0 release. However Logical interface tunnels can be terminated/originated to other BEBs from vIST peers.
- FE ISIS logical interface over IP and ISIS logical interface over VLANs will not be supported on the same node simultaneously.
- Egress shaping over Logical interface is not supported in this release.
- Do not fragment (DF) bit will not be set in the IP header of VXLAN encapsulation for data packets.
- MTU will not be auto-discovered over IP tunnel and tunnel MTU will not be automatically set to the auto-discovered value.
- Encryption
    - Switch based MACSEC encryption cannot be used with Fabric Extend IP.
- FE Logical interface Tunnel IP addresses, i.e. local and loopback interfaces used for the FE tunneling, cannot be reachable via ISIS IP shortcuts route and it has to be prevented by user by applying appropriate ISIS accept policies or using separate VRF for reachability to tunnel destination IPs.
- In FE solution using VSP 4000 with ONA, the ONA management VLAN always has to be in the GRT
- Fragmentation/re-assembly is supported on FE tunnels only between VSP 4000's with ONA's.
- Dynamic change of MTU on ONA resulting in mismatch of MTU config between VSP 4000 and VSP 8000/7200 will lead to traffic loss while adjacency remains operational
- No fragmentation/re-assembly support in L2 core or VID tunneling solution. Minimum MTU of 1544 bytes should be supported by underlying L2 core.
- In FE solution using VSP 4000 with ONA, MTU can be updated while logical interface is admin up. There will be a minimal transient loss of traffic while the MTU is updated.
- FE configuration will not be allowed if User has only configured Single B-VLAN in the SPBM. Both primary and secondary B-VLAN are mandatory for the Fabric extend feature to work.
- In Fabric extend L2 core solution ISIS logical interface vids cannot be same value as that of SPBm B-VLANs.
- In Fabric extend L2 core solution ISIS logical interfaces cannot be configured on Brouter interface.
- User created IP based filters (ACL/ACE) will not work on the physical interface with the FE VXLAN tunnels terminate

SPB Fabric over MPLS Pseudo-wire/E-Line Provider Network

Hub-and-Spoke over provider point to point VLAN Tunnels

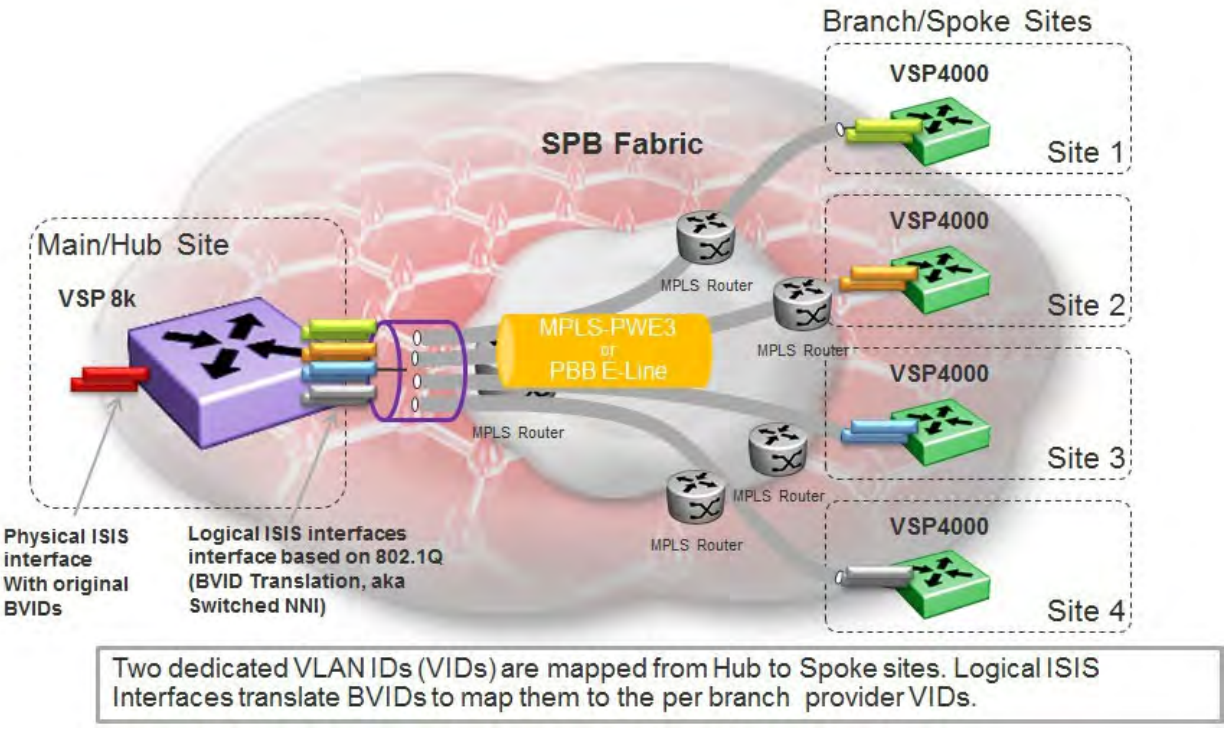

Two dedicated VLAN IDs (VIDs) are mapped from Hub to Spoke sites. Logical ISIS Interfaces translate BVIDs to map them to the per branch provider VIDs.

Hub-and-Spoke IP Tunnels over L2 segment

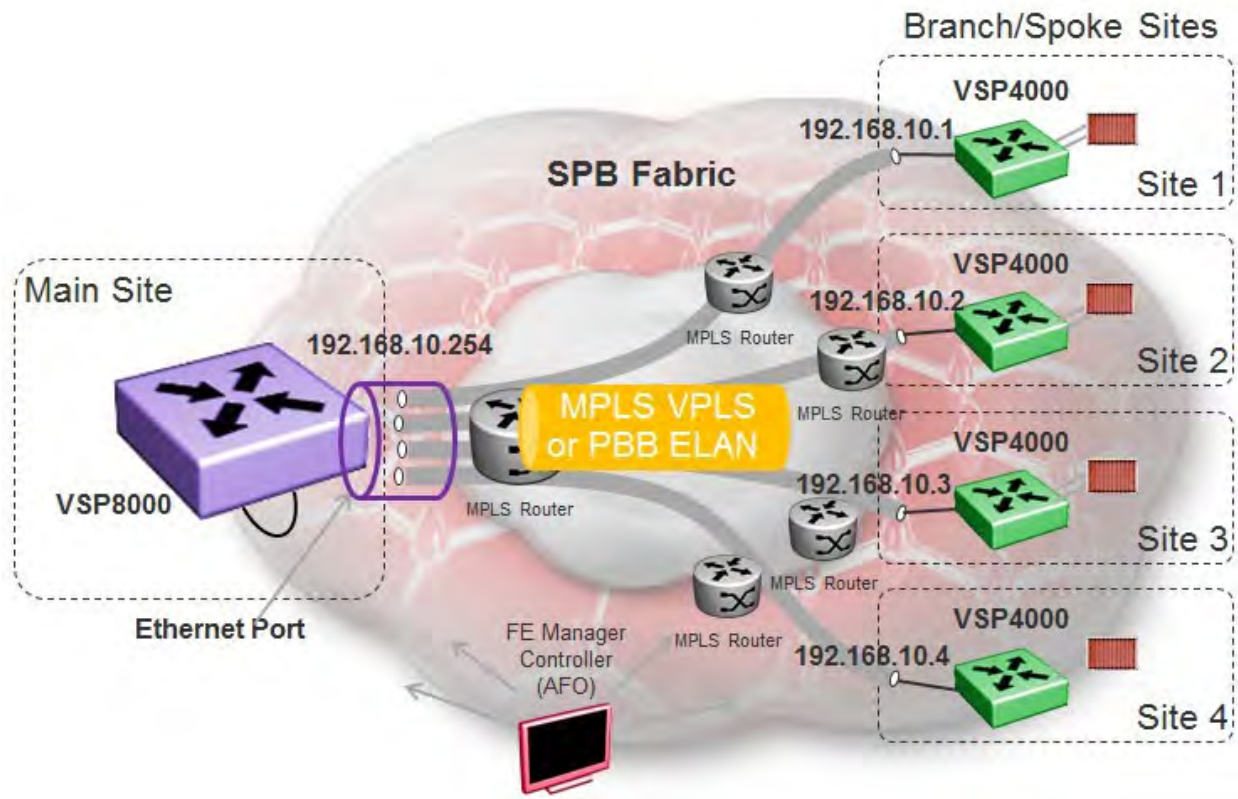

Please note the VSP 8000 and VSP 7200 only supports a single next hop IP address. If the core network is a flat layer 2 network from the main site to multiple spoke sites, a VSP 4000 with ONA can be used instead of the VSP 8000 as it does not have this restriction. However, a VSP 8000 or VSP 7200 can still be used at the main site by adding a loopback interface to a second VRF and using this interface as the tunnel source IP address. Please see the configuration example shown below under the section titled "*Fabric Extend over E-LAN/VPLS (L2) network using Layer 3 over Layer 2 tunneling with VSP8000 or VSP7200*".

SPB Fabric over MPLS IP-VPN Wide Area Provider Network (WAN)

Hub-and-Spoke over IP VPN

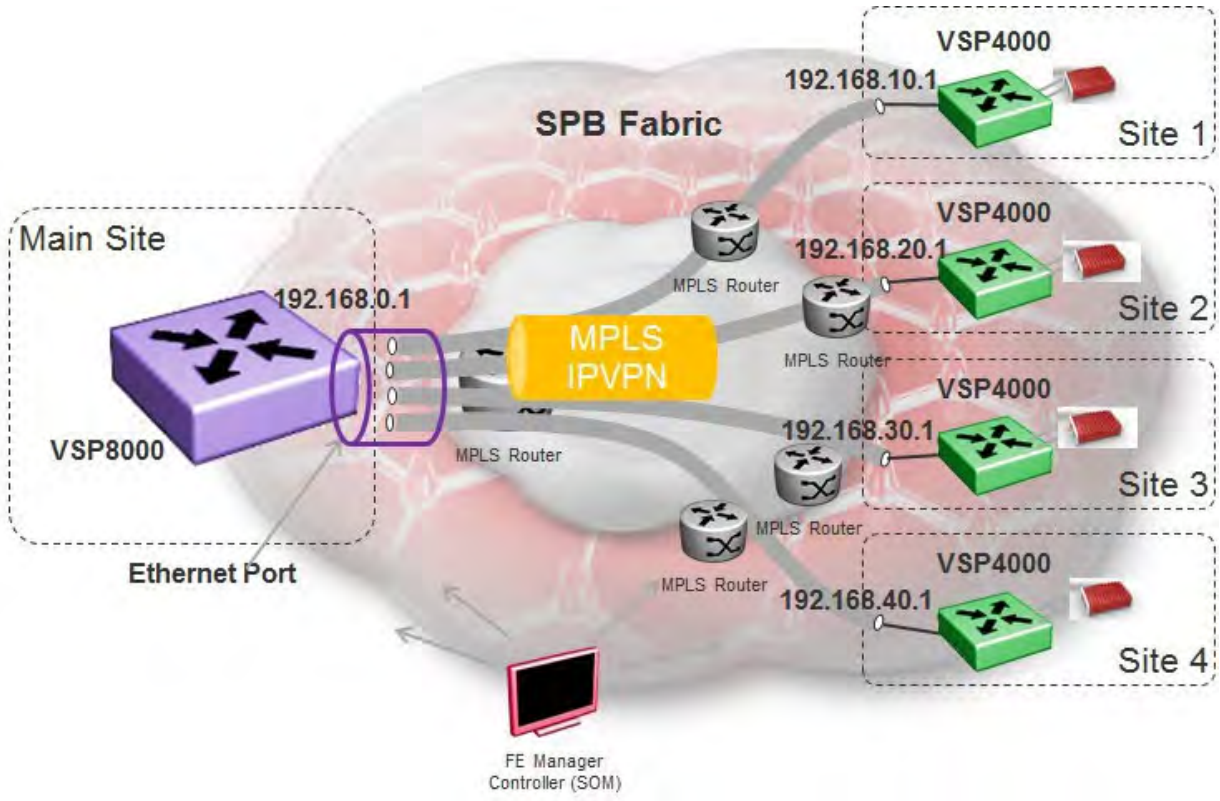

SPB Fabric over IP-Campus Network

IP Tunnel Mesh between Fabric nodes

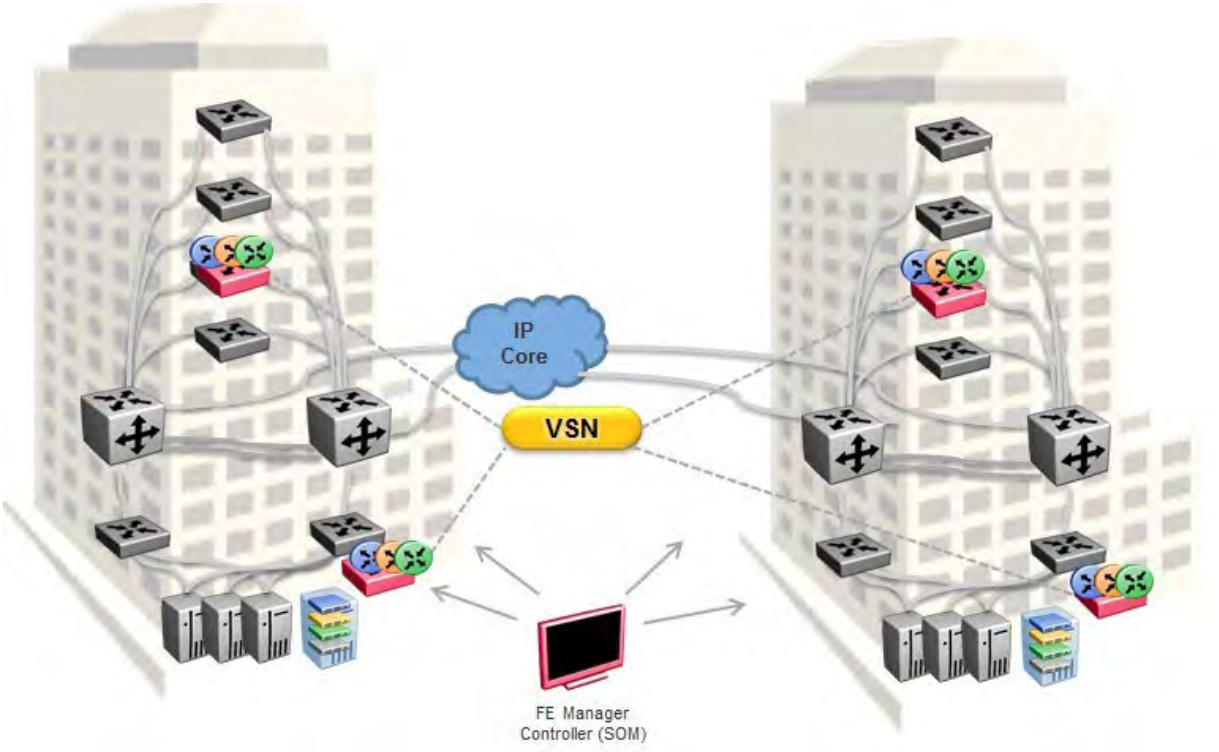

Multi-Hub Site Topology

Hub-and-Spoke over IP VPN

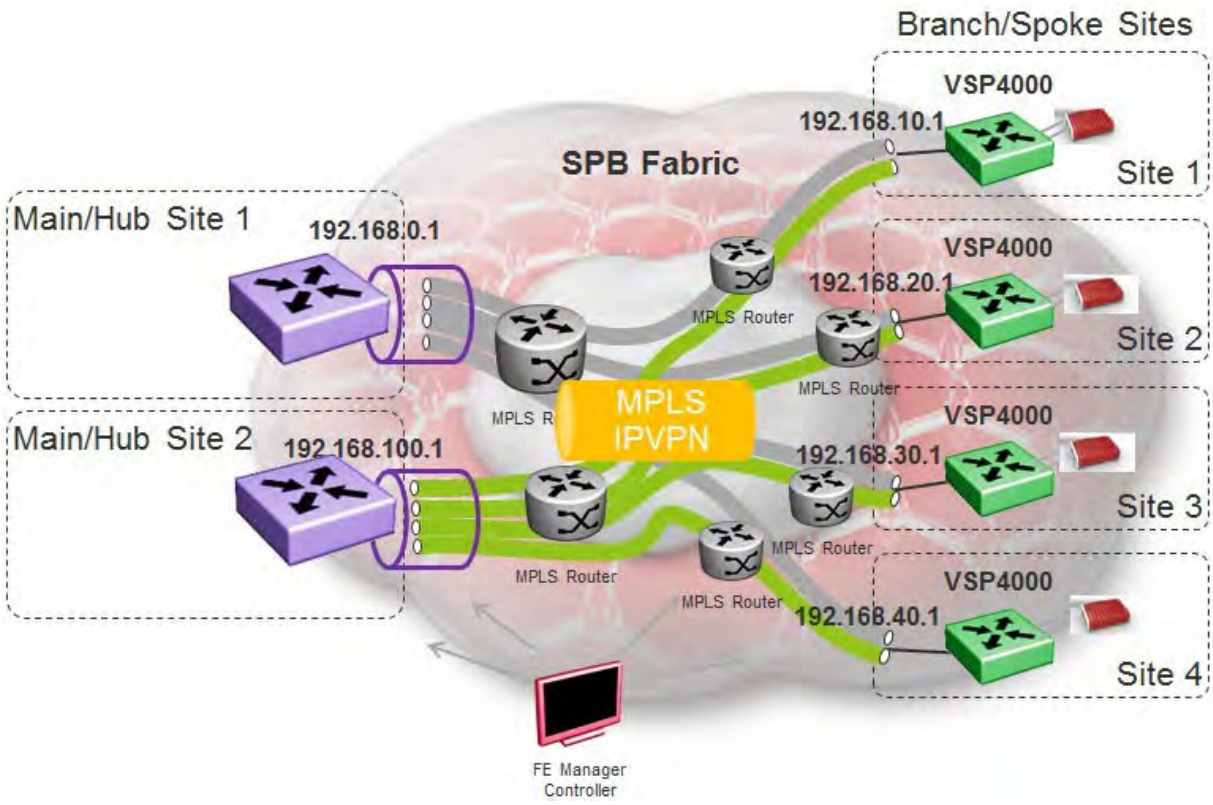

Fabric Extend – Tunnel Creation Full Mesh

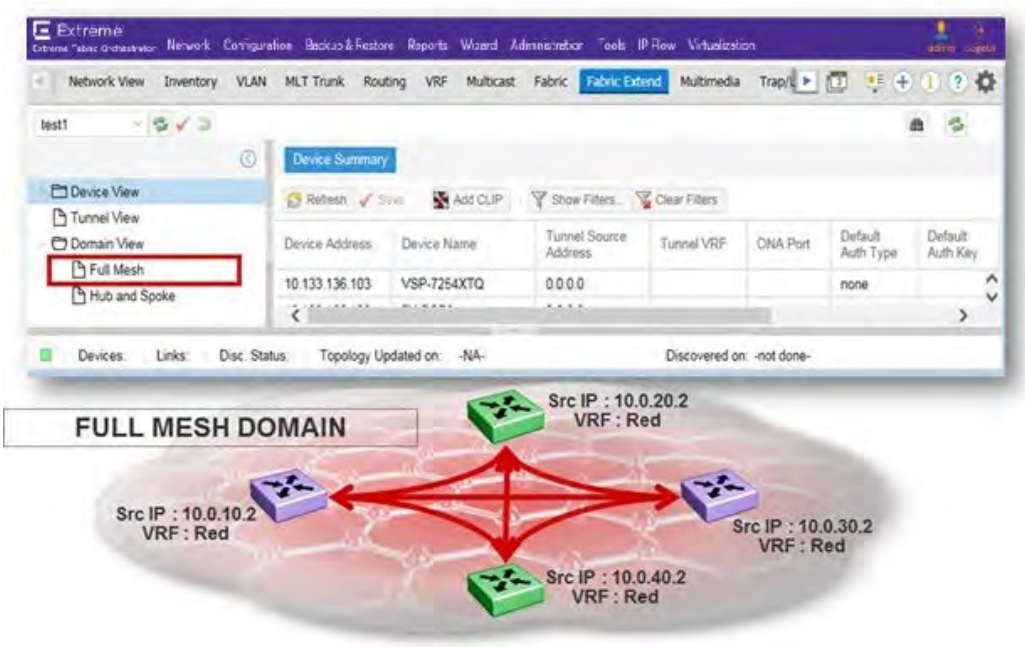

Fabric Extend – Tunnel Creation Hub and Spoke

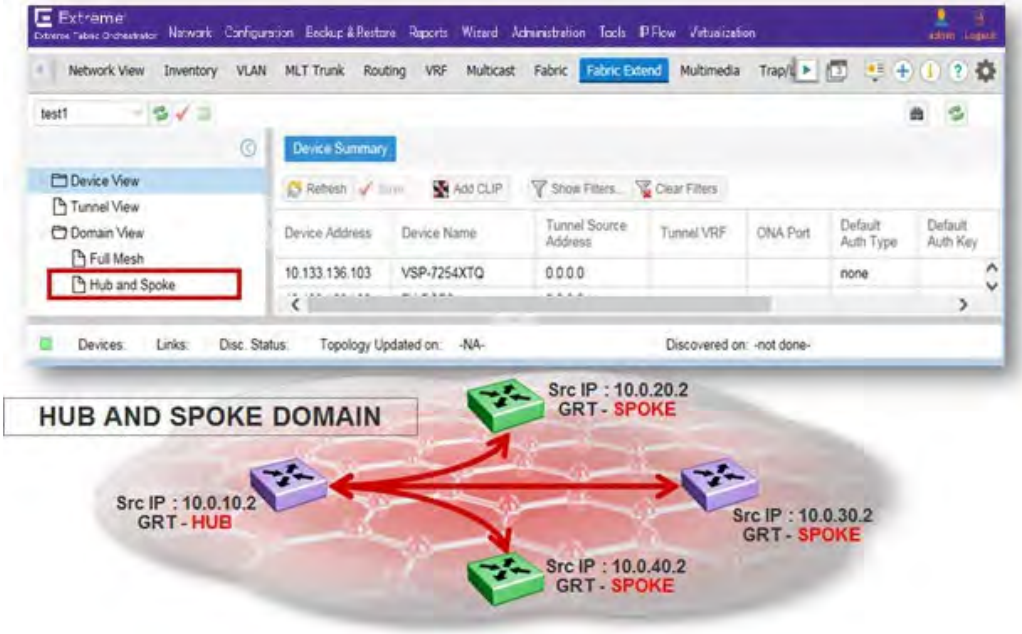

15. Fabric Attach

15.1 Fabric Attach Solution Overview

Fabric Attach (FA) is a feature that provides the service provisioning benefits of Fabric Connect to networking stub attached devices that do not support SPBM and where it makes little sense to IS-IS computing shortest paths. These devices can include WLAN Access Points, wiring closet access switches, servers, virtual machines, IP cameras, and Internet of Things (IoT). Devices supporting FA can dynamically request the network to create and provision VLAN and SPB virtual service mappings from the access layer of the network. By including Extreme Identity Engines policy engine in an FA solution, access layer provisioning can be fully automated by leveraging authentication of end users and devices, and then signaling to the access switch the VLAN and SPB service to create and assign to the user/device.

In summary, FA provides the ability to attach users and applications directly to the correct SPB virtual service (I-SID) by automatically adding the necessary VLANs on the FA components and tying those VLANs with the appropriate I-SID on the SPBM (FC) boundary nodes (BEBs enabled with FA Server).

Components of Fabric Attach solution

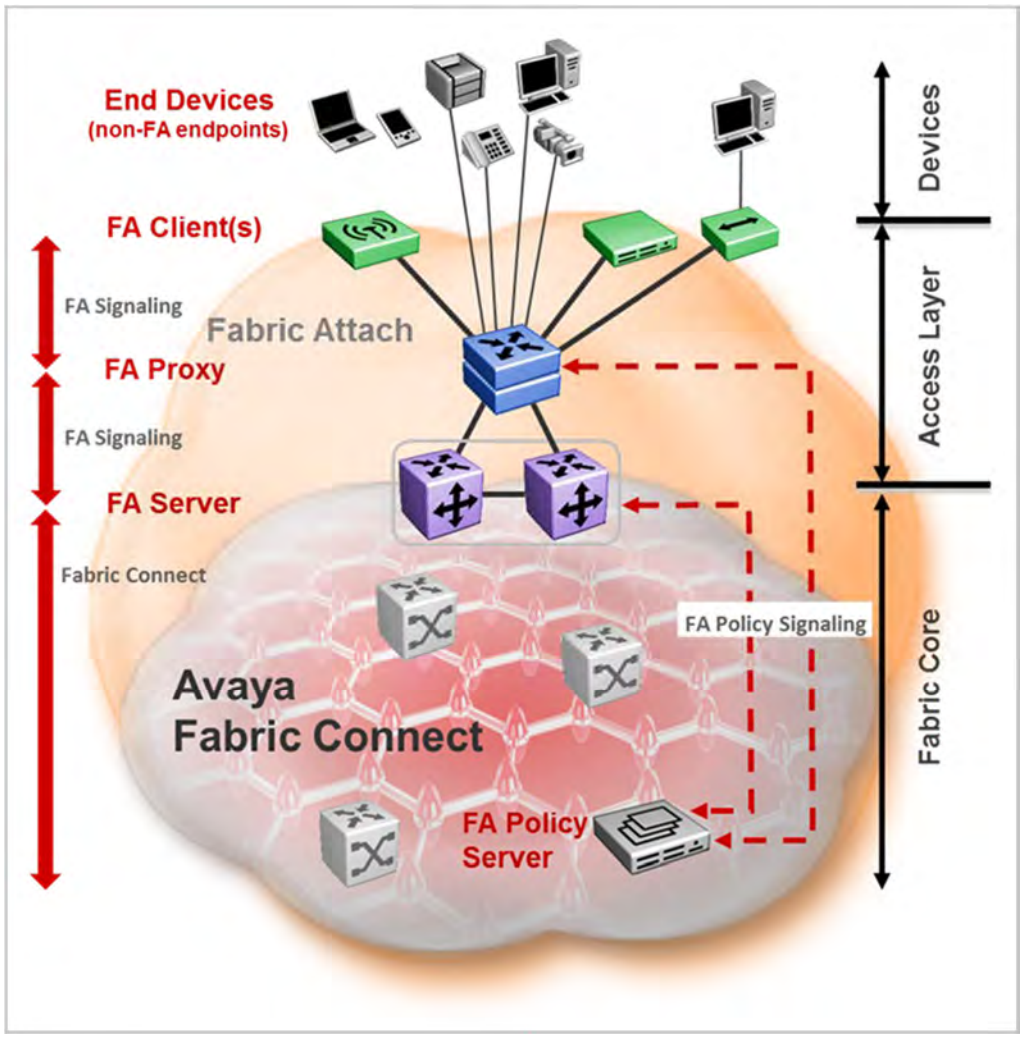

Fabric Attach element model

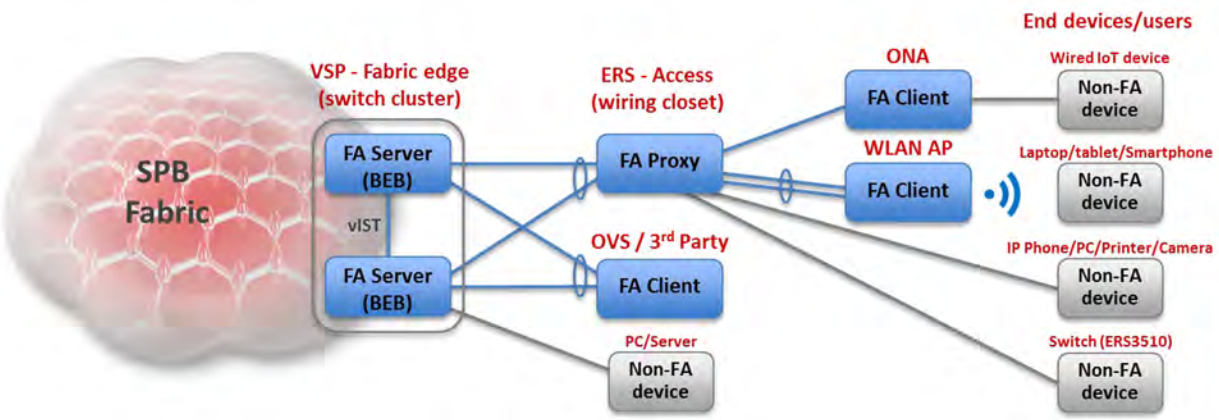

FA Server

When a switch is enabled as an FA Server, it receives IEEE 802.1AB Logical Link Discovery Protocol (LLDP) messages from FA Proxy switches and/or FA Client devices requesting the creation of Switched UNI service identifiers (I-SIDs). An FA Server can receive requests and consequently attach to multiple FA Proxy switches and/or FA Client devices. The I-SIDs thus created is required to join a Shortest Path Bridging (SPB) network.

The service created on the FA Server is an ELAN Switch UNI service with a FA I-SID to VLAN mapping. For Layer 2 or Layer 3 participation, you can create a platform VLAN with the same I-SID value as that of the ELAN I-SID value; this can be on the local FA Server node or another SPBM node in the network.

A platform VLAN is a VLAN created using the *vlan create <2-4059> type port-mstprstp <instance>* ACLI command

FA Proxy & Proxy Standalone

A FA proxy switch supports the ability to define ISID to VLAN assignments and relay this information to the FA server. This assignment can be accomplished, for example, by the local CLI or using Enterprise Device Manger. It also has the ability add I-SID to VLAN assignments after a FA client, i.e. WLAN 9100 AP, has successfully authenticated using EAP device authentication against Extreme's Identity Engines RADIUS server. In this case, after the FA Client has successfully authenticated, the policy used on Identity Engines RADIUS will contain all the various VLAN and I-SID assignments required. This information is sent via the outbound values provisioned in the Identity Engines Policy server. The I-SID assignment binding request in turn will be relayed from the FA Proxy to the FA Server where the FA Server will automatically create an ELAN (Switched UNI with c-vid and I-SID mappings).

A FA Proxy switch can be deployed in Standalone Proxy mode for scenarios where a FA Server is not available, i.e. a legacy network. In this case, in Standalone Proxy mode, the switch simply connects to a core/distribution switch or cluster via a tagged VLAN uplink port or MLT. Identity Engines is used in this case to authenticate the attached FA Clients and push down the necessary VLANs required. In this case, the policy created on Identity Engines must use a VLAN and I-SID combination where the I-SID must be a null ISID value of 0.

FA Client

A FA Client is a network attached end device supporting the Fabric Attach or (IEEE Auto Attach) agent in Client mode. FA Clients can initiate I-SID / VLAN binding requests for service creation to a FA Proxy or a FA Server. A FA Client will use FA signaling to automatically attach to fabric services that are always terminated on a FA Server. An FA Proxy switch will simply relay these requests from FA clients to the FA Server.

Please note if the FA Client is a WLAN 9100 AP, presently it is not ISID aware. It will always generate VLAN bindings with a null ISID. Hence, it can only be deployed in conjunction with FA Client NEAP(MHSA) authentication on the FA Proxy via IDE where the FA Proxy is required to relay the VLAN and ISID values, it obtained from Radius attributes, to the FA Server.

FA Policy Server

Identity Engines Policy server can be added to the solution to authenticate the FA Client WLAN 9100 AP using device authentication and push VLAN and I-SID assignments to the FA Proxy switch. Authentication of the WLAN 9100 is performed using Non-EAP Multiple Host Single Authentication.

Devices that support Fabric Attach and minimum software required

Product	Distribution Layer / TOR		Wiring Closet (SMLT TOR)		End Devices
	FA Server In SPBM mode	FA Server In VLAN mode	FA Proxy	FA Standalone Proxy	FA Client
VSP9000, ERS8000	N	N	N/A	N/A	N/A
VSP4000 / 7200 / 8000 (5.0)	Y[1]	N	N/A	N/A	N/A
VSP7000 (10.4)	N	N	Y	Y	N/A
ERS5900 (7.0.1)	Y	Y	Y	Y	N/A
ERS5600 (6.6.3)	N	N	Y	Y	N/A
ERS4800 (5.9.2)	Y	Y	Y	Y	N/A
ERS4500 (5.7.3)	N	N	Y	Y	N/A
ERS3500 (5.3)	N	N	N	Y	N/A
WLAN9100 (7.2)	N/A	N/A	N/A	N/A	Y
ONA 1.0	N/A	N/A	N/A	N/A	Y

Table 10: Devices that support Fabric Attach

[1]With or without SMLT

Fabric Attach Discovery and Signaling

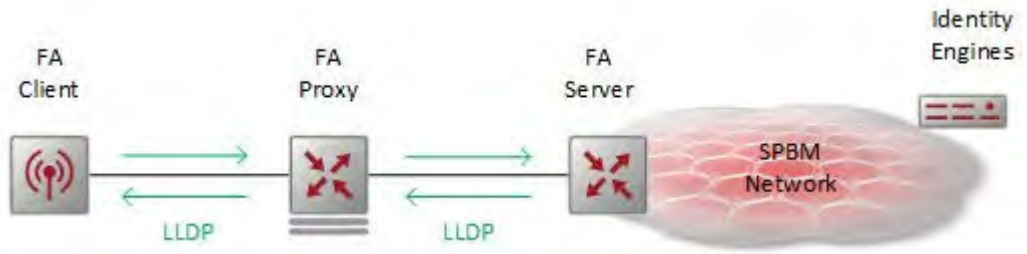

TLV Type [127]	TLV Length [50 octets]	Extreme OUI [00-04-0D]	Subtype [11]	HMAC-SHA Digest	Element Type	State	Mgmt VLAN	Rsvd	System ID
7 bits	9 bits	3 octets	1 octet	32 octets	6 bits	6 bits	12 bits	1 octet	10 octets

IEEE 802.1AB Logical Link Discovery Protocol (LLDP) is exchanged between Fabric Attach Client, Proxy, and Server components as part of an FA solution. The FA Client or Proxy will send LLDP PDUs to the FA Server switch, i.e. a Discovery Element TLV is used as the initial handshake between FA Server and FA Proxy or FA Client. LLDP is used to relay the I-SID and VLAN Mapping to the FA Server to allow the FA Server to create a Switched UNI ELAN.

The LLDP PDUs are received per port or all ports in an MLT where FA has been enabled. The LLDP Type field will indicate whether the FA element is a FA Client or FA Proxy. The Management VLAN is the value advertised by the FA Client and FA Proxy in the Discovery TLV. If the Management VLAN is configured on the FA Server, this will in return be relayed to the FA Proxy where it will automatically create the management VLAN and use DHCP to get an IP address. LLDP message authentication is enabled by default and uses a default 32 bit Extreme key. Please note that FA authentication is only available with the Secure Image if the Proxy switch is an ERS 4800 switch; the ERS 4800 supports both a secure and a non-secure image where only the secure image supports FA authentication. Please note FA message authentication can be disabled and the authentication key can also be changed if you do not wish to use the default Extreme 32 bit key.

FA Client is enabled on the WLAN 9100, FA Proxy is enabled on the ERS 4800 and ERS 5900, and FA Server is enabled on the VSP 4000/7200/8000 by default. FA is also enabled on all ports on the WLAN 9100, ERS 4800, and ERS 5900. FA needs to be enabled on a port or MLT level on the VSP 4000/7200/8000.

FA Auto Attach / Zero Touch

On an FA Server, an FA port can be configured with the management VLAN using an I-SID and VLAN ID value (c-vid). The c-vid is optional and if not specified, then the management traffic will be untagged. The I-SID is mandatory and is required for a network wide L2VSN. The FA server will announce this information in FA LLDP messages where only the management VLAN ID is announced as the ISID is not required. For untagged management, a VLAN ID of 4095 is announced.

On the FA Proxy, upon receiving the FA Message via LLDP, will create the corresponding management VLAN, set the uplink port to TagAll, add the management VLAN to the uplink port, and set QoS of trusted on the uplink port. It will also use DHCP to get an IP address if one has not already been configured – this is assuming a platform VLAN has been provisioned on the FA server with DHCP relay enabled. If you do not wish to use DHCP, this can be disabled by removing the default zero touch option *ip-addr-dhcp* on the FA Proxy switch.

Please note, in regards to the FA Proxy, FA Auto Attach does not at this time automate the following items:

- Disabling of Spanning Tree on the uplink ports to the FA Server
- MLT configuration of the uplink, i.e. if connected to an SMLT cluster of FA Servers
- VLACP configuration on the FA uplink ports if configured on the FA Server
- Removal of the default VLAN 1 from the FA uplink

16. SPB SMLT BEB Design Best Practices

16.1 SMLT BEB – C-VLAN Guidelines for L2VSN

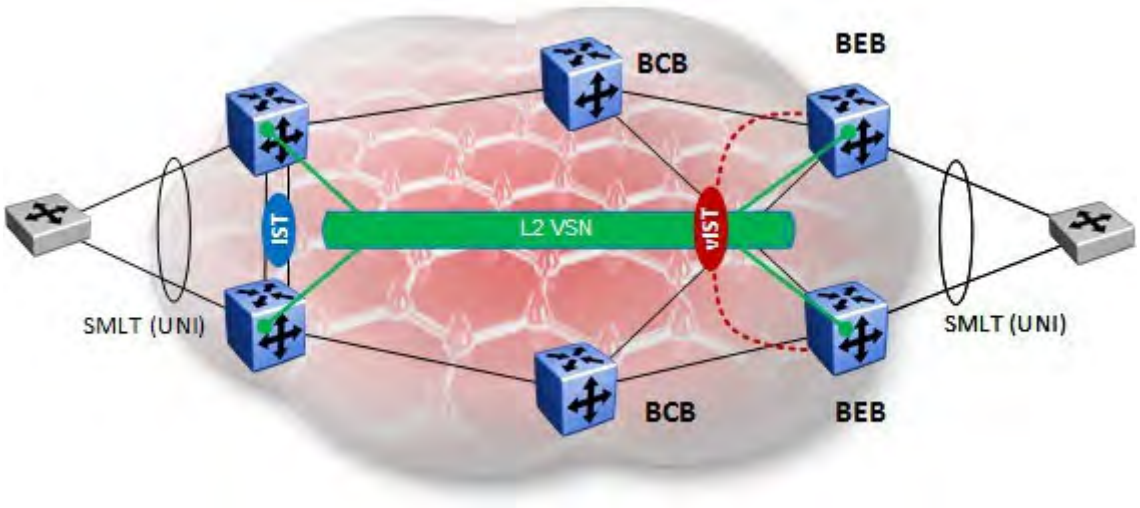

- Customer VLAN (CVLAN) has ISID assigned and is thus L2 extended with L2VSN
- On the ERS 8800 and VSP 7000 the CVLAN cannot be configured on any NNI interface (including the IST)
- On the VSP 9000 the CVLAN cannot be configured on any NNI interface (except on the IST where it MUST be configured)
- On the VSP 4000, VSP 7200, and VSP 8000 the CVLAN cannot be configured on any NNI interface
- VSP 4000, VSP 7000, VSP 7200, and VSP 8000 do not support multiple-port NNI MLTs. ISIS cannot be enabled on an interface within the MLT.

16.2SMLT BEB – Virtual Inter-Switch Trunk (vIST)

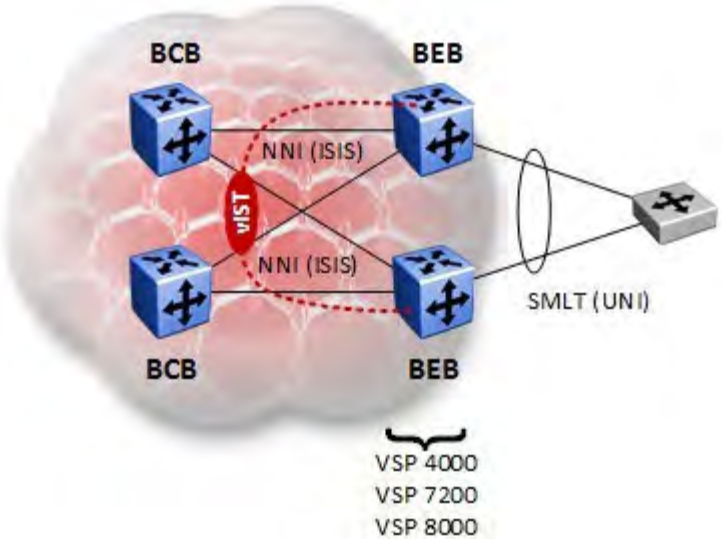

A traditional IST uses direct physical links configured as an MLT between a pair of cluster switches. Unlike a traditional IST, a vIST instead uses a virtual IST channel between a pair of cluster switches. This IST virtual channel is supported across the SPBM cloud and is not dependent on local physical ports. Hence, this eliminates the single point of failure with a dedicated MLT. The vIST is always up as long as there is SPBM connectivity between the vIST peers. Also, the vIST devices do not have to be the same type.

For the vIST, an IST VLAN must be provisioned as would be required with a normal IST with an IPv4 address. In addition, an SPBM ISID must be assigned to the vIST VLAN on both cluster switches. This ISID value must be unique and cannot be assigned to other VLANs in the network.

With vIST, any VLAN to be assigned to an SMLT interface must have an I-SID configured and thus becomes a C-VLAN. For each C-VLAN provisioned (C-VLAN = VLAN with ISID) on a vIST cluster/SMLT- pair, an SPBM ISID must be assigned on both cluster switches regardless of the additional services used on this VLAN (L2VSN, IP Shortcuts, or L3VSN). For an L2VSN service, this ISID can be used on other BEB switches in the network to form an L2VSN service. For an IP L3VSN service, a separate ISID value is used to identify the L3VSN service.

16.3SMLT BEB – ISIS Hello Timer Guidelines for ERS 8800

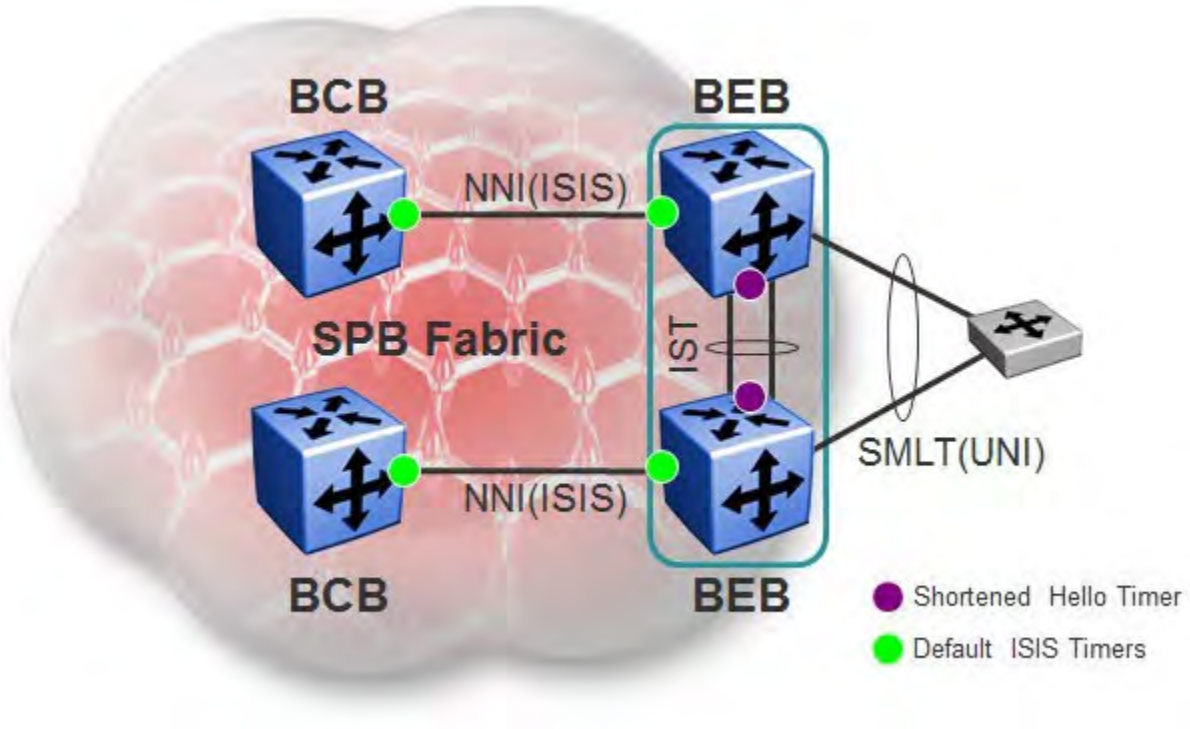

Connection Type	11-hello-interval	l1-hello-multiplier
IST – NNI ISIS	1 sec	27
NNI ISIS	9 sec (default)	3 (default)

- Please note that this recommended ISIS hello timer guideline only applies to the ERS 8800
- On the IST, ISIS is enabled on the MLT bundle
- Upon node restart, we need the ISIS adjacency over the IST MLT to come up before the IST comes up, therefore the ISIS Hello timer is reduced to 1 sec
- The hello multiplier is increased by the same factor to ensure the same time delay for an ISIS adjacency to transition in the down state
  - 1x27 = 9x3 = 27

16.4SMLT BEB – RSMLT

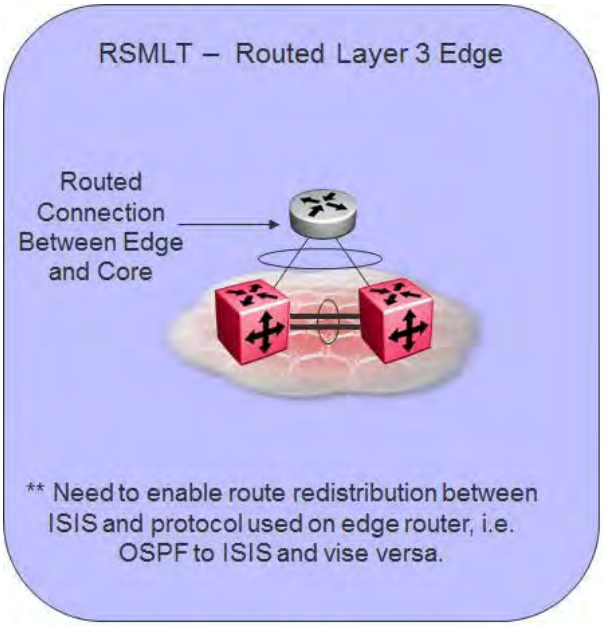

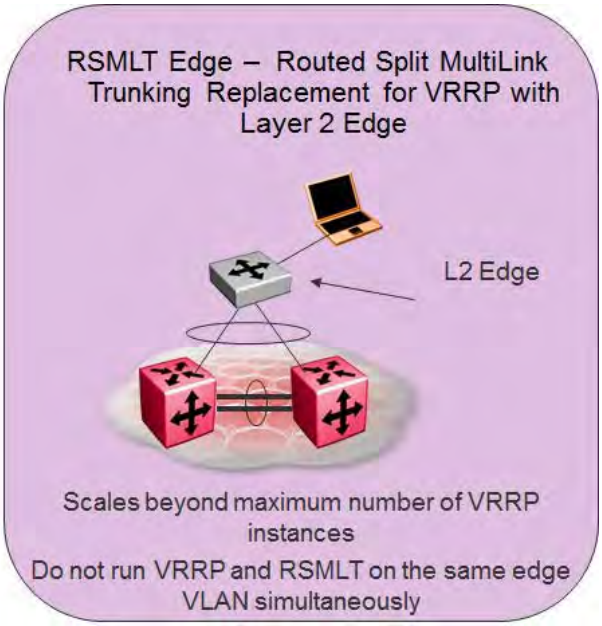

- Both RSMLT and RSMLT Edge is supported providing the SMLT cluster is either a VSP 9000 or ERS 8800 SMLT cluster or a VSP 4000/7200/8000 vIST SMLT cluster
- For RSMLT, if the OSPF network has multiple entry points via multiple SPB nodes, OSPF route policies must be configured on the SPB BEB switches to deny OSPF routes from each remote BEB entry point to prevent routing loops.

16.5SMLT BEB – VLACP Guidelines

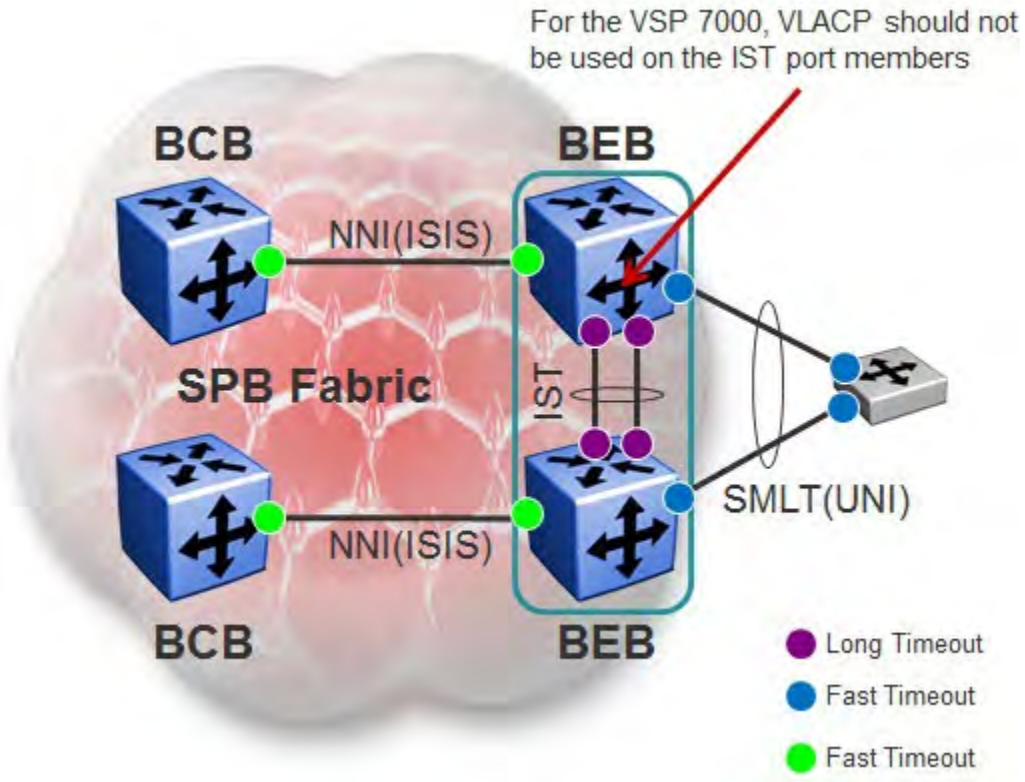

Connection Type	Fast Timer	Slow Timer	Timeout	Timeout Scale	ERS 8800 VSP 9000	VSP 7200 VSP 8000 VSP 4000 ERS 4800 ERS 5900	VSP 7000
IST – NNI ISIS	N/A	10000	Long	3	√	N/A	X
SMLT (UNI)	500ms	N/A	Short	5	√	√	√
NNI (ISIS)	500ms	N/A	Short	5	√	√	√

- Enable VLACP on all NNI ISIS enabled interfaces
- IST (which is now also an NNI connection) uses same VLACP slow timers
  - This does not apply to the VSP 7000 where VLACP should not be enabled on the IST port members
- Core facing NNI interfaces use same VLACP timers as SMLT UNI connections

16.6SMLT BEB – VSP 7000 Guidelines

For the VSP 7000, it is important to not enable the *filter-untagged-frame* option on the IST port members.

Prior to releases to 10.4.0, the default PVID of all IST ports must be the primary B-VLAN ID. This will happen automatically providing SPB is enabled first prior to enabling the IST. You can check the default PVID by entering the ACLI command *show vlan interface info <port list>.* To manually set the default PVID on the IST ports, use the ACLI command *vlan ports <port list> pvid <1-4096>.*

Starting with release 10.4.0, the PVID will automatically be set with the secondary BVLAN ID on the primary SMLT node and with the primary BLAN ID on the secondary SMLT node. This will also occur automatically when upgrading from release 10.3.x to 10.4.0. The output from the CLI *show vlan interface <port>* command will display either the primary or secondary PVID depending on which switch you are connected to.

Also, it is recommended to not enable VLACP on the IST and layer 3 routing is not supported with SPBM.

16.7SLPP Guard

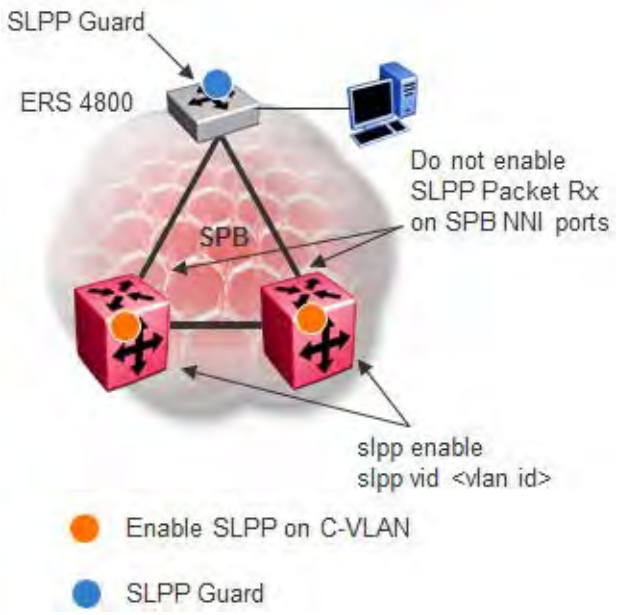

- SLPP Guard is supported on the ERS 3500, ERS 4000, ERS 5000, ERS 5900, and VSP 7000
- SLPP can enabled on the core bridges and in turn SLPP Guard can be enabled on the switch for local port loop detection
    - Because of this feature, SLPP can be enabled on the core SPB bridges and in turn allowing SLPP Guard to be enabled on the ERS 4800
- Only enable SLPP on the C-VLAN on the core SPB bridges
    - Do not enable SLPP Packet Rx on core NNI ports
        - Never want to take these ports down

17. SPB NNI SMLT – migrating existing SMLT network to SPB

When migrating from a legacy SMLT network to SPBM, under certain circumstances, you may have to change the MLT configuration as only one adjacency (port or MLT) is allowed between a pair of SPB switches. Please see the drawings shown below illustrating the various options. Please note this section does not apply to the ERS 4800 or ERS 5900 as SMLT is not supported on these products.

Please note the green links shown illustrate active links with IS-IS enabled where the link is either a physical port or MLT bundle. Most of the topologies only really apply when migrating to SPB with a SMLT cluster.

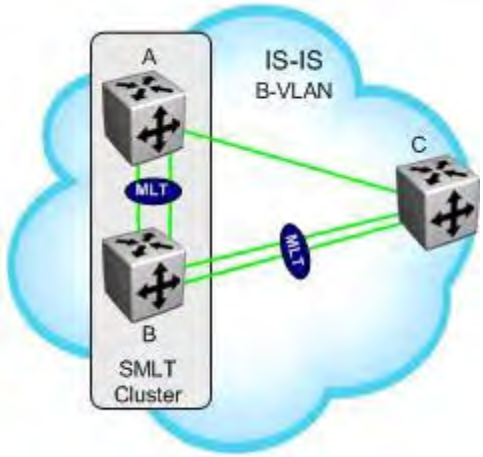

Figure 8: NNI - Triangle

In reference to switch C, it meets the requirement of only one link between a pair of SPB switches as it only has IS-IS enabled on the port to switch A and on the MLT bundle to switch B.

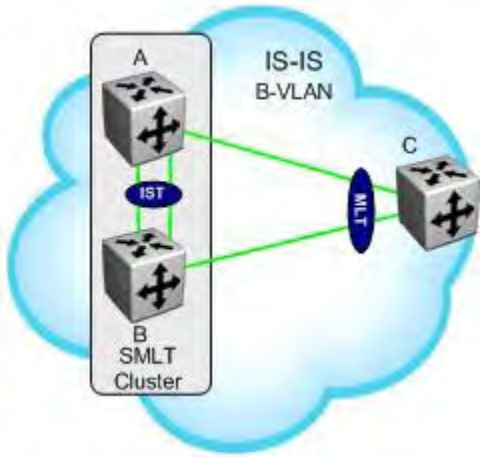

Figure 9: NNI - SMLT Triangle A

In reference to switch C, even though it has an MLT provisioned, IS-IS is provisioned on the physical ports to switch A and switch B. This type of configuration may show up when migrating to SPB where you may wish to not remove the MLT configuration. Please note that switch C can only be an ERS 8000 or VSP 9000.

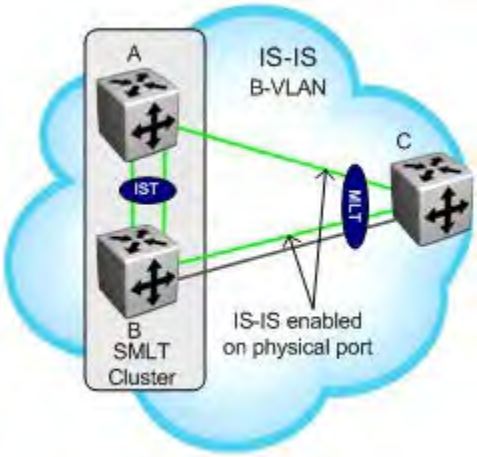

Figure 10: NNI – SMLT Triangle B

In reference to switch C, IS-IS cannot be enabled on the MLT bundle. If you wish to keep the MLT bundle, from switch C's perspective, enable IS-IS on the physical port to switch A and one of the physical ports to switch B. This applies when migrating from SMLT to SPB. If green field, then one should configure what is shown in figure 2. Please note that switch C can only be an ERS 8800 or VSP 9000

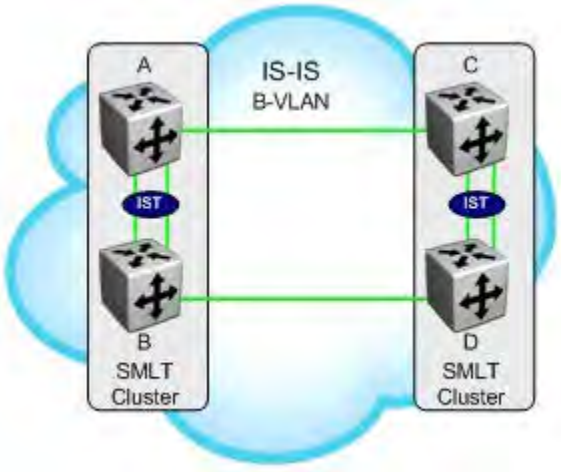

Figure 11: NNI –Square A

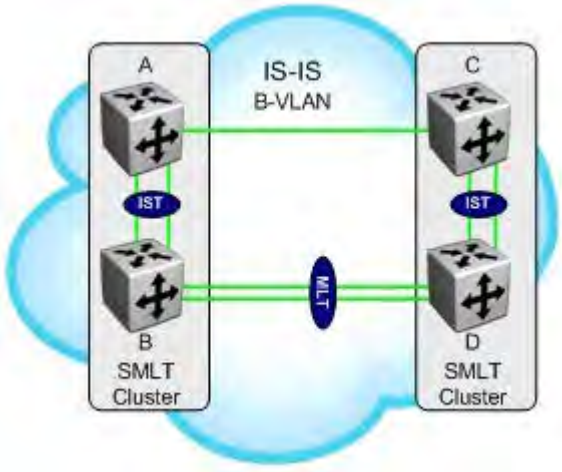

Figure 12: NNI – Square B

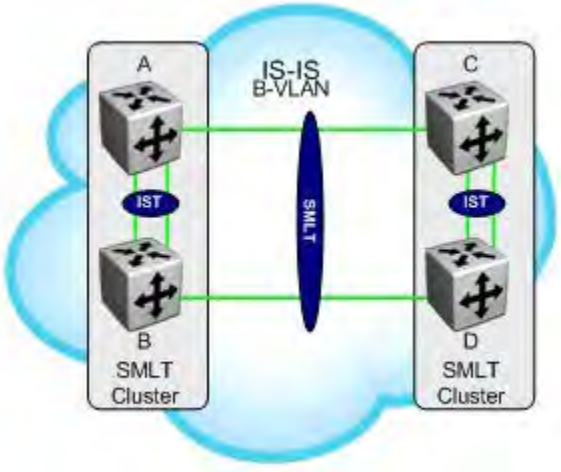

Figure 13: NNI – SLT Square

This diagram illustrates a likely scenario migrating from SMLT to SPB. The SMLT links could be made using regular MLT (with only one Ethernet port) or using SLT. In both cases, IS-IS should be enabled on the Ethernet port directly. You cannot enable IS- IS on the MLT (single port) if it has an SMLT ID. Please note that the switches can only be ERS 8800 or VSP 9000.

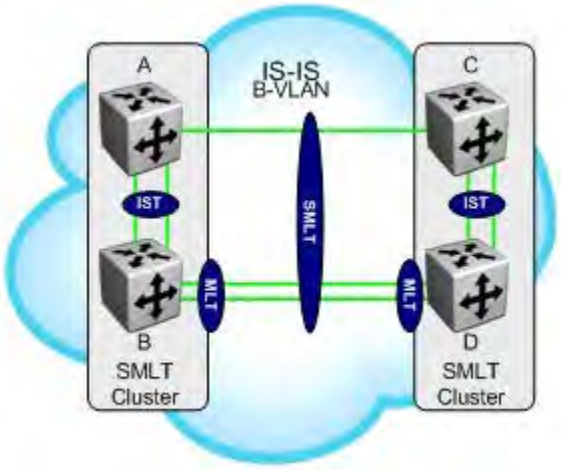

IS-IS is enabled on the link between nodes A and C. Between B and D, you cannot configure SPB on the MLT if it assigned with an SMLT ID. Once the SMLT ID is removed, then SPB can be enabled on the MLT.

Figure 14: NNI – SMLT Square

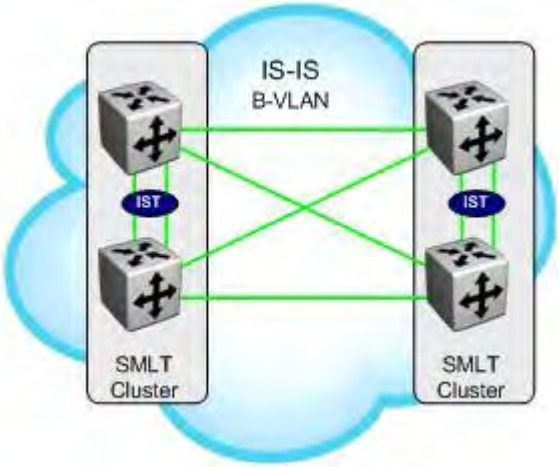

Figure 15: NNI – Full Mesh A

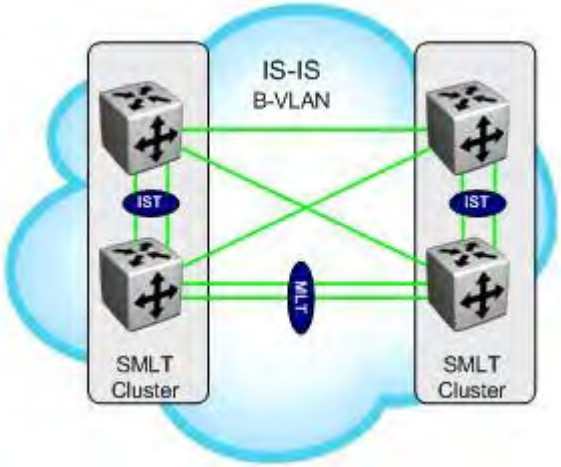

Figure 16: NNI – Full Mesh B

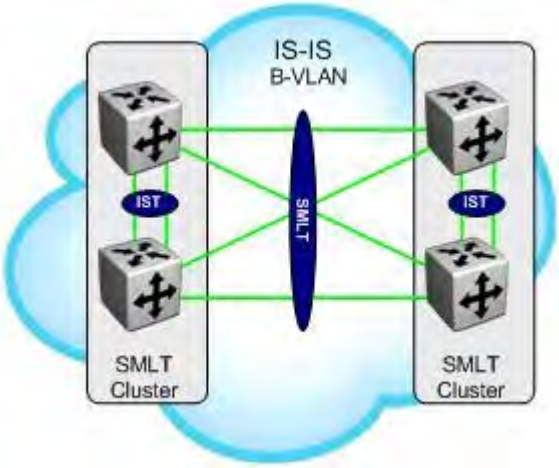

Figure 17: NNI – SMLT Full Mesh A

This diagram illustrates a common SMLT Full Mesh topology. Each switch has a local SMLT MLT defined with two Ethernet port members. When migrating this topology to SPB, IS-IS must not be enabled on the MLT instance, but, on the individual Ethernet ports which constitute it. Please note that the switches can only be ERS 8800 or VSP 9000.

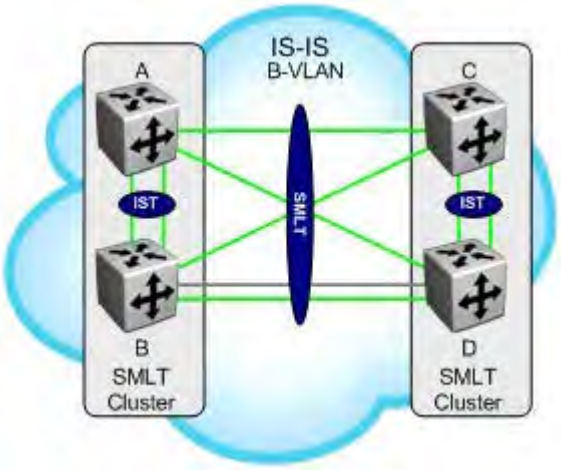

Figure 18: NNI – SMLT Full Mesh B

IS-IS should only be configured on one of the links between nodes B and D. Please note that the switches can only be ERS 8800 or VSP 9000.

18. IS-IS TLV

SPB uses IS-IS TLV (Type Length Value) and sub TLVs parameters to carry information in Link State Advertisements to other SPB enabled bridges including SPB services as shown in the table below

TLV	Description	Usage
1	Area Addresses	IS-IS area
3	End System Neighbors	B-MAC & SysName of itsself
22	Extended IS Reachability	IS-IS adjacencies Sub-TLV 29: Link Metric for SPBM alone
129	Protocol Supported	SPBM
135	TE IP Reachability	IP Reachability for IP Shortcuts in GRT
137	Host Name	ISIS router name
236	IPv6 Reachability	IP Reachability for IPv6 shortcuts in GRT
143	SPBM Instance & BVIDs	Sub-TLV 6: BVIDs to ECT algorithm Used in IS-IS Hellos only
144	SPBM Instance, Nick-name, BVLANs & ISIDs	Sub-TLV 1: SPBM Instance & Nick-name Sub-TLV 3: B-VLANs & L2VSN ISIDs
184	SPBM IPVPN Reachability	IP Reachability for L3 VSNs
185	SPBM ISID Constrained Source-Groups	IP Multicast stream availability for L2VSNs & L3VSNs
186	SPBM VRF-Opera/GRT Source-Groups	IP Multicast stream availability for GRT/VRF-0
236	IPv6 Reachability	IP Reachability for IPv6 Shortcuts in GRT

Table 11: ISIS TLV's

TLVs 1,3,22,129 & 135 are well known IS-IS TLVs which exited even before SPB was defined

TLVs 143 & 144 are new IS-IS TLVs defined for use by SPB

TLVs 184, 185 & 186 are new IS-IS TLVs defined in Extreme's IETF draft for SPB IP extensions

19. SPB Best Practices

The following are best practices when setting up SPB.

IS-IS

- Recommended to leave the IS-IS SYS-ID (B-MAC) with its default value to ensure no duplication in the network
    - If you do change manually the SYS-ID, please take the necessary steps to ensure there is no duplication in the network
        - In release 5.0 for the VSP 4000/7200/8000, duplicate SYS-ID detection is supported
- Create two B-VLANs to allow load distribution over both B-VLANs. Even if SMLT is not used, this is still good practice as adding a new B-VLAN to an existing configuration requires that IS-IS to be disabled therefore disrupting the network

SPB

- Use a different IS-IS Nick Name on each switch that is easily recognizable
- If IP is enabled, i.e. IP shortcuts, it is required that an IS-IS IP source address be added

IST

- If the nodes are to form an SMLT Cluster, the IST must be already up and running before enabling IS-IS on it on the VSP 9000 and ERS 8000
- On the VSP 7000, SPB should be first configured prior to enabling the IST

SMLT

- Each switch in the cluster must be configured to peer with its neighbor.
- A virtual B-MAC will be automatically created based on the lowest SYS-ID in the cluster plus one
    - The virtual B-MAC is used as the source B-MAC when forwarding traffic received from an SMLT/SLT UNI port into the SPB fabric. This allows reverse MAC learning on the remote BEBs to map the SMLT learnt customer MAC address to an SMLT cluster rather than to an individual BEB switch forming that cluster
    - If you choose to use the automatic created virtual B-MAC, careful consideration must be taken to ensure that the SYS-ID if configured on of the cluster switches is greater than one compared to its peer
    - If you have chosen to manually change the IS-IS SYS-ID (B-MAC), then you should do the same for the virtual B-MAC.

Please note, the virtual B-MAC or any System ID created should not conflict with any other System ID or virtual B-MAC in the network. In other words, please ensure there is no duplication anywhere in your network of System ID's and virtual B-MACs.

A safe practice, which is also future proof, would be to leave the lowest byte in the SYS-ID as all zeroes.

- There is a consistency check in place to ensure that L2 VSN VLANs cannot be added to the IST or to any IS-IS enabled interface on the ERS 8800; does not apply for the IST MLT on a VSP 9000
- L3 VSN VLANs must still be added to the IST

- A L3 VSN VLAN can also be a L2 VSN VLAN
    - For example, an ISID can be assigned to a VRF for L3 VSN. This does not restrict another ISID using a different value from the one assigned to the VRF to be assigned to VLANs within the VRF

ISIS Adjacency

Physical or MLT links between IS-IS switches

- Only a single point to point IS-IS adjacency is supported between a pair of IS-IS switches
    - For example, if there are two ports between a pair of IS-IS switches, IS-IS should only be configured on one of the two ports (if configured on both, only one of those links will form an IS-IS adjacency)
    - If a single MLT is configured between a pair of IS-IS switches, all ports (1-8) in the MLT will be utilized – note that the MLT must be configured first and then IS-IS can be enabled on the MLT

CFM

- If not using the simplified CFM configuration commands:
    - The Domain name must be same on all switches in a IS-IS area
    - The Maintenance Association must the same on all switches in a IS-IS area
        - Two Maintenance Associations should be created, one for each B-VLAN to allow CFM testing over both B-VLANs
    - The MIP can be configured the same on all switches in a IS-IS area or uniquely defined per switch
- The MEP id should be unique to every switch in the SPB network

The MIP must be configured at the same level as the MEP on all switches in the SPB network.

20. SPB Configuration

On the ERS 4800, ERS 8800, and VSP 7000, it is recommended to change the Spanning Tree mode to MSTP. By default, the VSP 4000, VSP 7200, VSP 8000, VSP 9000, and ERS 5900 support MSTP. This helps when using tools such as VLAN Manager in COM where the VLAN provisioned is broken down by Spanning Tree instance. To change the Spanning Tree mode to MSTP, please enter the following command:

ERS 8800

- 8800:5(config)#***boot config flags spanning-tree-mode mstp***

VSP 7000 & ERS 4800

- 7024XLS(config)#***spanning-tree mode mst***

Changing the Spanning Tree mode flag from default to MSTP on the ERS 8800 will result in a loss of configuration following the necessary reboot to activate the MSTP flag. This is because the syntax of certain commands in config.cfg (vlan creation & Spanning Tree port settings) changes in the two modes. It is therefore necessary to do a manual conversion of the config.cfg file (for example in a text editor using find & replace) in order to re-load the existing configuration file in MSTP mode.

20.1 SPB Configuration

20.1.1 ERS 8800 – Converting from CLI to ACLI

As the ERS 8000 supports CLI and ACLI, it is highly recommended to use ACLI as all other switches from Extreme only support ACLI. If you are presently using CLI, you can convert to ACLI using the following configuration.

```
ERS8800-1:5# copy /flash/config.cfg /flash/backup.cfg
ERS8800-1:5# save config file /flash/config_acli.cfg backup /flash/config.cfg mode acli
ERS8800-1:5# config boot flags acli true
ERS8800-1:5 config boot choice primary config-file /flash/config_acli.cfg
ERS8800-1:5# save boot
ERS8800-1:5# boot -y
```

20.1.2 SPB and IS-IS Core Configuration

SPB and IS-IS core configuration

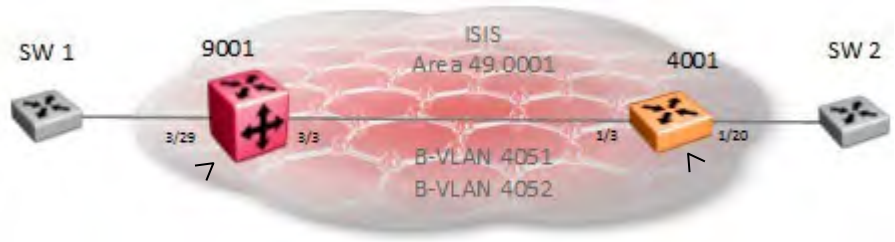


```
configure terminal
spbm
prompt 9001
router isis
  spbm 1
  spbm nick-name 0.90.01
  spbm b-vid 4051-4052 primary 4051
  manual-area 49.0001
exit
vlan create 4051 name BVLAN-1 type spbm-bvlan
vlan create 4052 name BVLAN-2 type spbm-bvlan
router isis enable
```

```
configure terminal
spbm
prompt 4001
router isis
  spbm 1
  spbm nick-name 0.40.01
  spbm b-vid 4051-4052 primary 4051
  manual-area 49.0001
exit
vlan create 4051 name BVLAN-1 type spbm-bvlan
vlan create 4052 name BVLAN-2 type spbm-bvlan
router isis enable
```

```
config terminal
spbm
prompt <word 0-255> **By default, becomes SPB System Name
router isis
       sys-name **Please see note above
       spbm 1
       system-id <xxxx.xxxx.xxxx – Optional, by default the base MAC is used>
```

```
       spbm 1 nick-name <x.xx.xx - 2.5 bytes>
       spbm 1 b-vid <prim vlan id,sec vlan id> primary <prim vlan id>
       manual-area <xx.xxxx.xxxx...xxxx - 1...13 bytes>
exit

vlan create <primary vlan-id> name "BVLAN-1" type spbm-bvlan
vlan create <secondary vlan-id> name "BVLAN-2" type spbm-bvlan
router isis enable
```

VSP 7000, ERS 4800, & ERS 5900

```
config terminal
snmp-server name <word 0-31> **By default, becomes SPB System Name
vlan create <primary vlan-id> name "BVLAN-1" type spbm-bvlan
vlan create <secondary vlan-id> name "BVLAN-2" type spbm-bvlan
spbm
router isis
       sys-name **Please see note above
       spbm 1
       system-id <xxxx.xxxx.xxxx – Optional, by default the base MAC is used>
       spbm 1 nick-name <x.xx.xx - 2.5 bytes>
       spbm 1 b-vid <prim vlan id,sec vlan id> primary <prim vlan id>
       manual-area <xx.xxxx.xxxx...xxxx - 1...13 bytes>
exit
router isis enable
```

Please note, if the IS-IS sys-name is not provisioned, by default, the global system name is used as the IS-IS sys-name. If you do wish to set the IS-IS sys-name, it can be set to a value different than global system name.

The primary and secondary VLAN provisioning must be the same on all SPB bridges, i.e. if VLAN 4051 is provisioned as the primary B-VLAN and VLAN 4052 is provisioned as the secondary B-VLAN, then this must be repeated on all SPB bridges.

On the VSP 7000, ERS 4800, and ERS 5900, the B-VLANs must be configured first prior to enabling SPB and ISIS.

By default, the SPB EtherType is set to 0x8100 on all Extreme switches when SPB is enabled. Please note this value is set on purpose to allow SPB to be transported across non-SPB networks, i.e. transparent VLAN service or a traditional Ethernet network. For SPB interoperability between different vendors, this value will have to be changed to the STP standard EtherType value of 0x88a8 unless this vendor also supports a SPB EtherType value of 0x8100.

The default switch behavior regarding System-Id is to use a MAC address within the MAC address range reserved for the switch. This ensures that there will be no de- stabilizing System-Id conflicts in the network. Extreme recommends the use of default System-Id values for this reason. To allow greater flexibility to customers, the use of configured System-Id values is also supported. When using configured System-Id

values it is very critical to ensure that each SPB enabled switch in the network uses a unique ISIS System-Id value.

If you do decide to change the System ID, it is recommended to set the locally administered bit. The second least significant bit of the most significant byte of the MAC address should be set to 1 to indicate the MAC address as locally administered. For more details, please go to http://en.wikipedia.org/wiki/MAC_address.

Please note that if you change the System ID after ISIS has already been enabled, you must also change the nickname. If you do not wish to change the nickname, you will still need to temporary change the nickname. After you change the System ID and temporarily change the nickname, enable ISIS, disable ISIS, change the nickname back to the original value, and then enable ISIS again.

Verify Operations:

```
show running-config module isis
show running-config module spbm
show isis
show isis spbm
show isis system-id
show isis nick-name
show isis manual-area
show isis spbm
show isis spbm nick-name
show isis spbm unicast-fib
show isis spbm unicast-fib vlan <BVLAN ID>
show isis spbm unicast-tree <BVLAN ID>
show isis lsdb
```

20.1.3 SPB NNI Interface Configuration

SPB and IS-IS core interface configuration

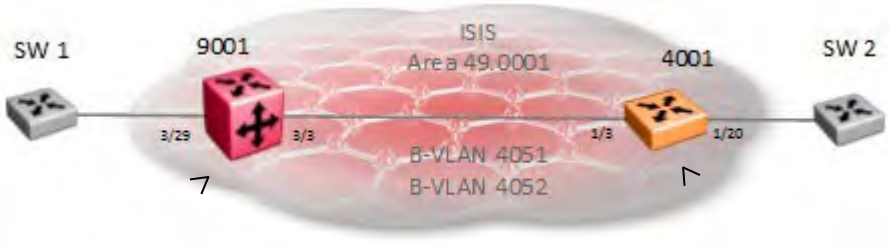


```
configure terminal
interface gigabitethernet 3/3
  no shutdown
  no spanning-tree mstp force-port-state enable
  isis
  isis spbm 1
  isis enable
exit
```

```
configure terminal
interface gigabitethernet 1/3
  no shutdown
  no spanning-tree mstp force-port-state enable
  isis
  isis spbm 1
  isis enable
exit
```

```
interface gigabitethernet <slot/port>
   isis
   isis spbm 1
   isis enable
exit
interface mlt <mlt id>
   isis
   isis spbm 1
   isis enable
exit
```

VSP 7000, ERS 4800, & ERS 5900:

```
interface ethernet <slot/port>
   isis
   isis spbm 1
   isis enable
exit
```

Please note that Spanning Tree should be disabled on all SPB NNI ports including all single ports or ports that are part of an MLT when the SBI NNI links are directly attached to another Extreme SPB switch. This does not apply to SMLT port members since SMLT disables Spanning Tree automatically.

As of release 10.3 for the VSP 7000 and 5.8 for the ERS 4800, the interface configuration changed from *interface fastEthernet <ports>* to *interface ethernet <ports>*

On the VSP 9000, VSP 4000, VSP 7200, and VSP 8000, by default all ports are administratively disabled.

On the ERS 4800 and VSP 7000, for all MLT's, ISIS is enabled at the port level, i.e. on each port that is a member of the MLT.

Verify Operations:

```
show isis interface
show isis adjacencies
show isis lsdb tlv 22 detail
show isis int-l1-cntl-pkts
```

20.1.4 CFM Configuration

CFM configuration

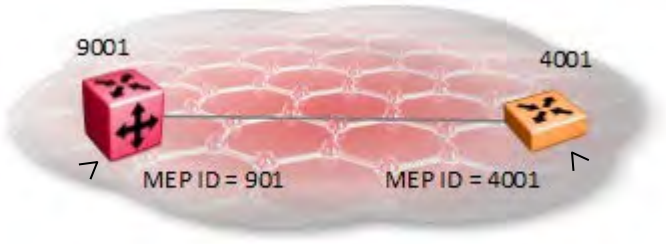


```
configure terminal
cfm spbm mepid 901
cfm spbm enable
```

```
configure terminal
cfm spbm mepid 4001
cfm spbm enable
```

```
config terminal
cfm spbm mepid <1-8191>
cfm spbm enable
```

By default, the CFM a Maintenance Domain name of *spbm* is used while two Maintenance Associations are created using the two B-VLAN IDs.

Verify Operations:

```
show cfm <maintenance-domain|maintenance-association|maintenance-endpoint>
l2 ping vlan <BVLAN ID> routernodename <system-name>
l2 ping vlan <BVLAN ID> mac <system-id>
l2 traceroute vlan <BVLAN ID> routernodename <system-name>
l2 traceroute vlan <BVLAN ID> mac <system-id>
```

20.1.5 VSP 7000 – Fabric Interconnect Mesh

20.1.5.1 Rear Port Mode

In the 10.2 release, the VSP 7000 can be configured in Fabric Interconnect Mesh (FI) mode by setting the rear-port mode to SPB. This allows the VSP 7000 to run SPB via the rear ports using stacking cables to connect to other VSP 7000s. In the 10.2.1 release, SMLT is supported allowing for either SPB or SMLT to operate via the rear port. In the 10.3 release, both SPB and SMLT are supported via the rear ports.

Please refer to the *Resilient Data Center Solutions Technical Configuration Guide* publication number NN48500-645 for more details.

ACLI - L2 VSN

```
config terminal
rear-port mode enable spb
Enabling rear port mode will disable Fabric Interconnect Stack operation.
Switch configuration will be reset to partial-defaults. Continue(yes/no)?yes
```

```
show rear-port mode
```

20.1.5.2 Rear Port Mode LACP Provisioning

By default when rear port mode is enabled, LACP is automatically enabled across all rear ports using a default LACP key of 4095. If you wish, you can change this value on one or more of the four rear ports. In SPB rear port mode, the port numbers for each rear port is as follows:

- FI Up (right) Bottom: Port 33
- FI Up (right) Top: Ports 34, 35, 36
- FI Down (left) Bottom: Port 37
- FI Down (left) Top: Ports 38, 39 (SPB) or ports 38, 39, 40 (Standard)

For example, to change the LACP on the *FI Up (right) Top* ports:

```
interface ethernet 34-36
  lacp key 4094
exit
show lacp aggr
show lacp port aggr <aggr id>
show lacp debug member 34-36
```

20.1.6 SMLT – Normal IST

Enabling IST

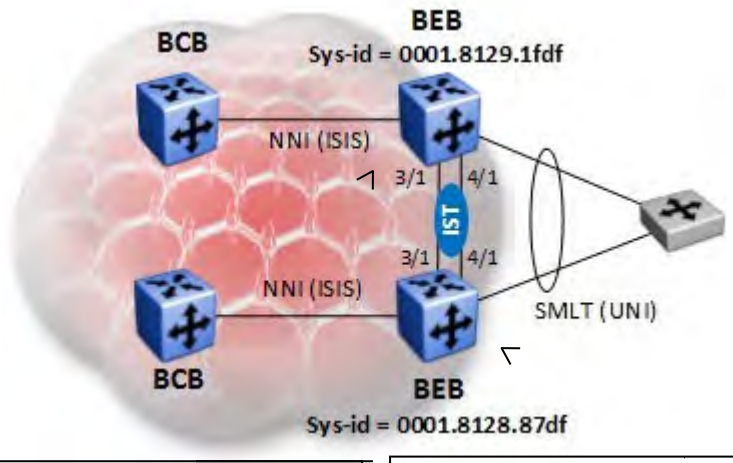


```
config terminal
vlan create 2 type port-mstprstp 0
interface vlan 2
   ip address 10.1.2.1 255.255.255.252
exit
mlt 1 enable name IST
mlt 1 member 3/1,4/1
mlt 1 vlan 2
interface mlt 1
   ist peer-ip 10.1.2.2 vlan 2
   ist enable
exit
router isis
  spbm 1 smlt-peer-system-id 0001.8128.87df
exit
```

```
config terminal
vlan create 2 type port-mstprstp 0
interface vlan 2
   ip address 10.1.2.2 255.255.255.252
exit
mlt 1 enable name IST
mlt 1 member 3/1,4/1
mlt 1 vlan 2
interface mlt 1
   ist peer-ip 10.1.2.1 vlan 2
   ist enable
exit
router isis
  spbm 1 smlt-peer-system-id 0001.8129.1fdf
exit
```

Enter the ALCI *show isis system-id* on each peer switch to get the System ID value for the *spbm 1 smlt-peer-system-id xxxx.xxxx.xxxx* command. Each switch in the cluster must peer with its peer's System ID value.

```
config terminal
vlan create x type port-mstprstp 0
interface vlan x
   ip address <ip address> <ip mask>
exit
mlt y enable name IST
mlt y member slot/port,slot/port
mlt y vlan x
interface mlt y
   ist peer-ip <ip address of peer> vlan x
   ist enable
exit
router isis
```

```
   spbm 1 smlt-virtual-bmac <xx:xx:xx:xx:xx:xx – Optional, by default the lowest B-MAC
   in the cluster is used plus one>
   spbm 1 smlt-peer-system-id <xxxx.xxxx.xxxx - system id of peer>
exit
```

Verify Operations:

```
show mlt
show ist mlt
show isis spbm
```

20.1.7 SMLT - Virtual IST (vIST)

The following shows how to provision the virtual IST. This feature will allow a SMLT cluster to have an IST between two cluster switches that does not require a physical connection between the cluster switches, i.e. an MLT with two or more ports. In practice it usually makes sense to use a direct connection between the vIST peers where this connection is now a regular NNI link and need not be a MLT.

Enabling vIST

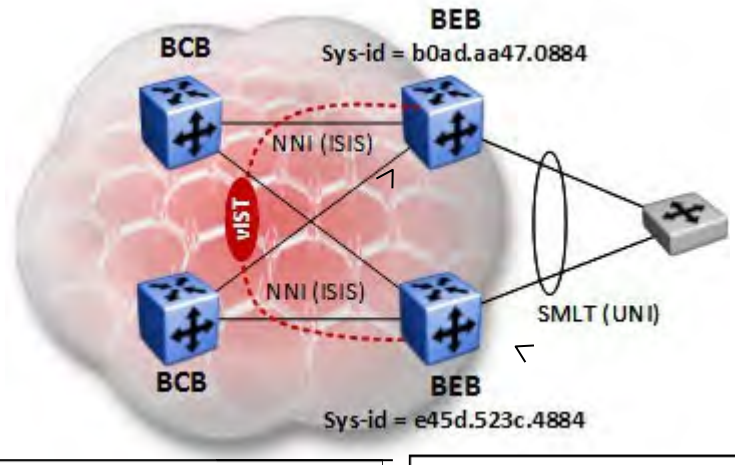


```
config terminal
router isis
   spbm 1 smlt-peer-system-id e45d.523c.4884
exit
vlan create 2 type port-mstprstp 0
vlan ISID 2 2002
interface vlan 2
   ip address 10.1.82.1 255.255.255.252
exit
virtual-ist peer-ip 10.1.82.2 vlan 2
```

```
config terminal
router isis
   spbm 1 smlt-peer-system-id b0ad.aa47.0884
exit
vlan create 2 type port-mstprstp 0
vlan ISID 2 2002
interface vlan 2
   ip address 10.1.82.2 255.255.255.252
exit
virtual-ist peer-ip 10.1.82.1 vlan 2
```

Enter the ACLI *show isis system-id* on each peer switch to get the System ID value for the *spbm 1 smlt-peer-system-id xxxx.xxxx.xxxx* command. Each switch in the cluster must peer with its peer's System ID value.

```
config terminal
router isis
   spbm 1 smlt-peer-system-id <xxxx.xxxx.xxxx - system id of peer>
   spbm 1 smlt-virtual-bmac <xx:xx:xx:xx:xx:xx - Optional, by default the lowest B-MAC
   in the cluster is used plus one>
exit
vlan create x type port-mstprstp 0
vlan ISID x <i-isid number>
interface vlan x
   ip address <ip address>/<mask>
exit
virtual-ist peer-ip <ip address of peer> vlan x
```

Verify Operations:

```
show virtual-ist
show isis spbm
```

Only C-VLANs (VLAN to which an ISID is assigned) can be assigned to SMLT ports with vIST. For each C-VLAN added, an ISID must be assigned to both SMLT cluster switches to allow learning between the two cluster switches including IP Shortcuts and L3VSN CVLANs.

BEB-1	BEB-2
IP Shortcuts Example – CVLAN Configuration	

`vlan create 1000 type port-mstprstp 0`
`vlan i-sid 1000 4001000`
`vlan mlt 1000 1`
`interface vlan 1000`
`  ip address 192.168.100.1 255.255.255.0`
`  ip rsmlt`
`  ip rsmlt holdup-timer 9999`
`exit`
	`vlan create 1000 type port-mstprstp 0`
	`vlan i-sid 1000 4001000`
	`vlan mlt 1000 1`
	`interface vlan 1000`
	`  ip address 192.168.100.2 255.255.255.0`
	`  ip rsmlt`
	`  ip rsmlt holdup-timer 9999`
	`exit`

L3VSN Example – CVLAN Configuration	

`ip vrf blue`
`vlan create 2255 type port-mstprstp 0`
`vlan i-sid 2255 5002255`
`vlan mlt 2255 1`
`interface Vlan 2255`
`  vrf blue`
`  ip address 192.168.100.1 255.255.255.0`
`  ip rsmlt`
`  ip rsmlt holdup-timer 9999`
`exit`
`router vrf blue`
`  ipvpn`
`  i-sid 3002255`
`  ipvpn enable`
`exit`
	`ip vrf blue`
	`vlan create 2255 type port-mstprstp 0`
	`vlan i-sid 2255 5002255`
	`vlan mlt 2255 1`
	`interface Vlan 2255`
	`  vrf blue`
	`  ip address 192.168.100.1 255.255.255.0`
	`  ip rsmlt`
	`  ip rsmlt holdup-timer 9999`
	`exit`
	`router vrf blue`
	`  ipvpn`
	`  i-sid 3002255`
	`  ipvpn enable`
	`exit`

20.1.8 L2VSN Configuration

L2 VSN

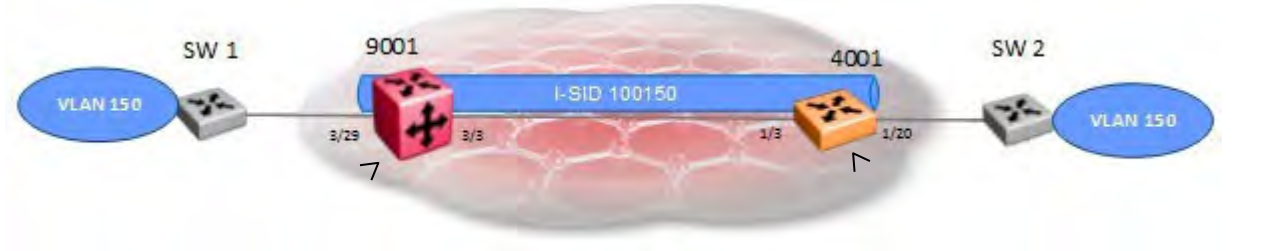


```
configure terminal
interface gigabitethernet 3/29
  no shutdown
  encapsulation dot1q
exit
vlan create 150 type port-mstprstp 0
vlan member 150 3/29
vlan i-sid 150 100150
vlan member remove 1 3/29
```

```
configure terminal
interface gigabitethernet 1/3
  no shutdown
  encapsulation dot1q
exit
vlan create 150 type port-mstprstp 0
vlan member 150 1/20
vlan i-sid 150 100150
vlan member remove 1 1/20
```

```
config terminal
vlan create <vlan id> type port-mstprstp 0
vlan members <vlan id> <slot/port>
vlan i-sid <vlan-id> <ISID: 0..16000000>
```

Although you can use any number from 1 to 16,777,215 as an ISID value, it is recommended not to use a value from 16,000,001 to 16,777,215. This range is used for Multicast over SPB.

Verify Operations:

```
show isis spbm i-sid all
show isis spbm i-sid all id <ISID value>
show isis spbm multicast-fib i-sid <ISID value>
show vlan i-sid
show isis lsdb tlv 144 detail
show isis lsdb sysid <system-id> tlv 144
All switches except VSP7000, ERS4800, and ERS 5900:
show vlan mac-address-entry <C-VLAN id>
show vlan remote-mac-table <C-VLAN id>
```

VSP7000, ERS4800, and ERS 5900:

```
show mac-address-table spbm <ISID value>
```

20.1.9 SwitchedUNI Configuration

Switched UNI

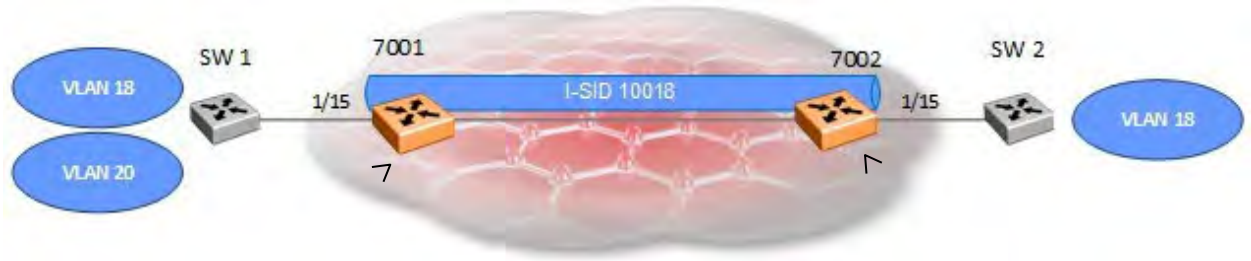


```
configure terminal
vlan ports 1/15 tagging tagall
vlan create 18 type spbm-switchedUni
vlan create 20 type spbm-switchedUni
i-sid 10018 vlan 18 port 1/15
i-sid 10018 vlan 20 port 1/15
vlan members remove 1 1/15
```

```
configure terminal
vlan ports 1/15 tagging tagall
vlan create 18 type spbm-switchedUni
i-sid 10018 vlan 18 port 1/15
vlan members remove 1 1/15
```

Please note that switches 7001 and 7002 are VSP7000 switches

```
config terminal
vlan create <vlan id> type spbm-switchedUni
i-sid <ISID: 0..16000000> vlan <vlan-id> port <port member>
```

Although you can use any number from 1 to 16,777,215 as an ISID value, it is recommended not to use a value from 16,000,001 to 16,777,215. This range is used for Multicast over SPB.

20.1.10Flex UNI Switched Configuration

Assuming we wish to forward tagged VLANs 10 and 11 plus untagged traffic between nodes 4002 and a vIST SMTL cluster made up of 8201 and 8202.

Switched UNI


```
configure terminal
vlan ports 1/5 tagging tagAll
vlan members remove 1 1/5
interface gigabitEthernet 1/5
  flex-uni enable
exit
i-sid 70010 elan
  c-vid 10 port 1/5
exit
i-sid 70011 elan
  c-vid 11 port 1/5
exit
i-sid 70012 elan
  untagged-traffic port 1/5
exit
```

```
configure terminal
mlt 9
mlt 9 encapsulation dot1q
mlt 9 member 2/2
interface mlt 9
  smlt
exit
vlan members remove 1 2/2
interface mlt 9
  flex-uni enable
exit
i-sid 70010 elan
  c-vid 10 mlt 9
exit
i-sid 70011 elan
  c-vid 11 mlt 9
exit
i-sid 70012 elan
  untagged-traffic mlt 9
exit
```

Please note that switches 8201 and 8202 are VSP 8000 switches while 4002 is a VSP 4000 switch.

Verify Operations

```
show mlt i-sid <mlt ID>
show interfaces gigabitEthernet i-sid <slot/port>
show isis spbm i-sid all id <i-sid id>
show isis spbm multicast-fib detail
show i-sid mac-address-entry <i-sid id>
```

Please note that MAC learning is done at an ISID level and not at a VLAN level. Also, the untagged traffic BPDU enable option cannot be enabled on the SMLT cluster; Spanning Tree is not supported.

To flush the MAC table, use the *i-sid mac-address-entry <i-sid id> flush* ACLI command.

20.1.11Transparent UNI Configuration

Transparent UNI

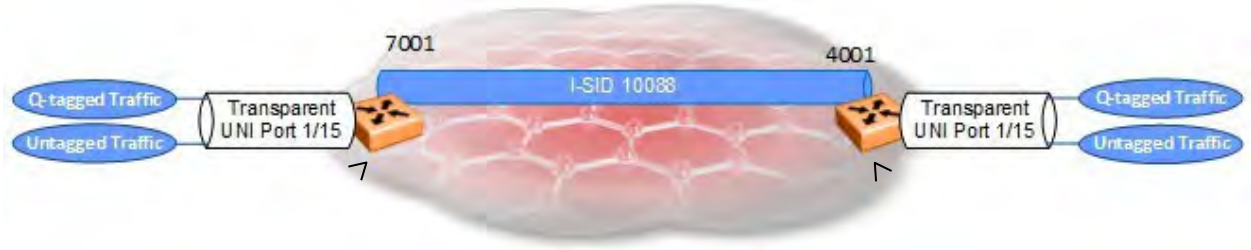


```
configure terminal
vlan member remove 1 1/15
i-sid 10088 port 1/15
```

```
configure terminal
i-sid 10088 elan-transparent
  port 1/15

  Adding Ports to Transparent UNI ISID
  removes it from all VLANS.
  Do you wish to continue (y/n) ? y
exit
```

Please note that 7001 is a VSP7000 switch while 4001 is a VSP4000 switch.

VSP 7000:

```
config terminal
i-sid <ISID: 0..16000000> port <port member>
```

VSP 4000, VSP 7200, and VSP 8000:

```
i-sid <ISID: 0..16000000> elan-transparent
  port <slot/port> | mlt <mlt-id>
exit
```

Although you can use any number from 1 to 16,777,215 as an ISID value, it is recommended not to use a value from 16,000,001 to 16,777,215. This range is used for Multicast over SPB.

Transparent UNI - SMLT

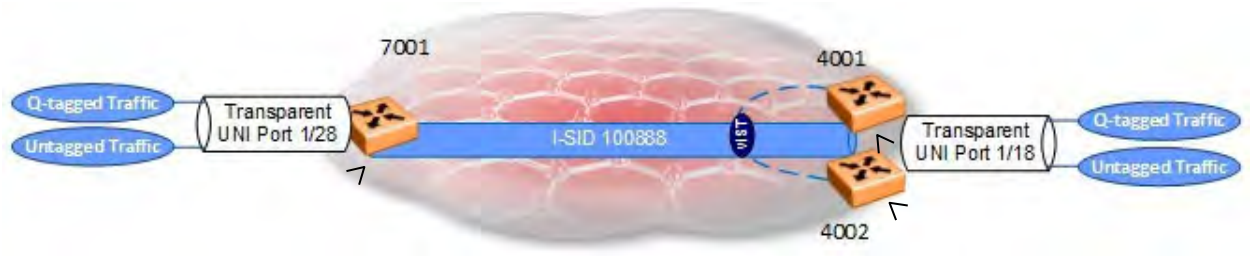


```
configure terminal
vlan member remove 1 1/28
i-sid 100888 port 1/28
```

```
configure terminal
mlt 6
mlt 6 member 1/18
interface mlt 6
  smlt
exit
i-sid 100888 elan-transparent
  mlt 6
  Adding MLT to Transparent UNI ISID removes
  it from all VLANS.
  Do you wish to continue (y/n)? y
exit
```

Please note that 7001 is a VSP7000 switch while 4001 and 4002 are VSP4000 switches using vIST.

VSP 7000:

```
config terminal
i-sid <ISID: 0..16000000> port <port member>
```

VSP 4000, VSP 7200, and VSP 8000:

```
mlt <mlt-id>
mlt <mlt-id> member <port members>
interface mlt <mlt-id>
  smlt
exit
i-sid <ISID: 0..16000000> elan-transparent
  mlt <mlt-id>
exit
```

20.1.12 Private VLAN (ETREE) Configuration

The Etree feature allows private VLANs to traverse the SPBM network for Layer 2 services.

Global Private VLAN Configuration

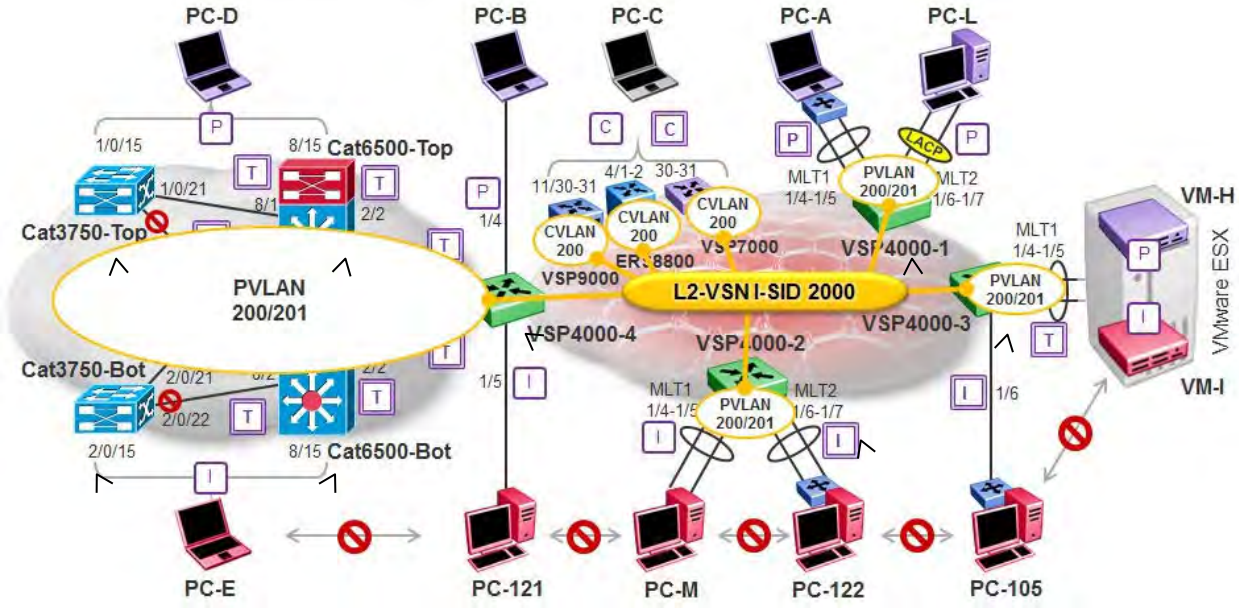


```
vlan 200
   private-vlan primary
   private-vlan association 201
exit
vlan 201
   private-vlan isolated
exit
```

```
configure terminal
vlan create 200 type pvlan-mstprstp 0 secondary 201
vlan i-sid 200 2000
```

VSP 4000, VSP 7200 & VSP 8000:

```
vlan create <vlan-id> type pvlan-mstprstp 0 secondary <secondary vlan-id>
vlan i-sid <vlan-id> <ISID: 0..16000000>
```

VMWare ESX Private VLAN Configuration

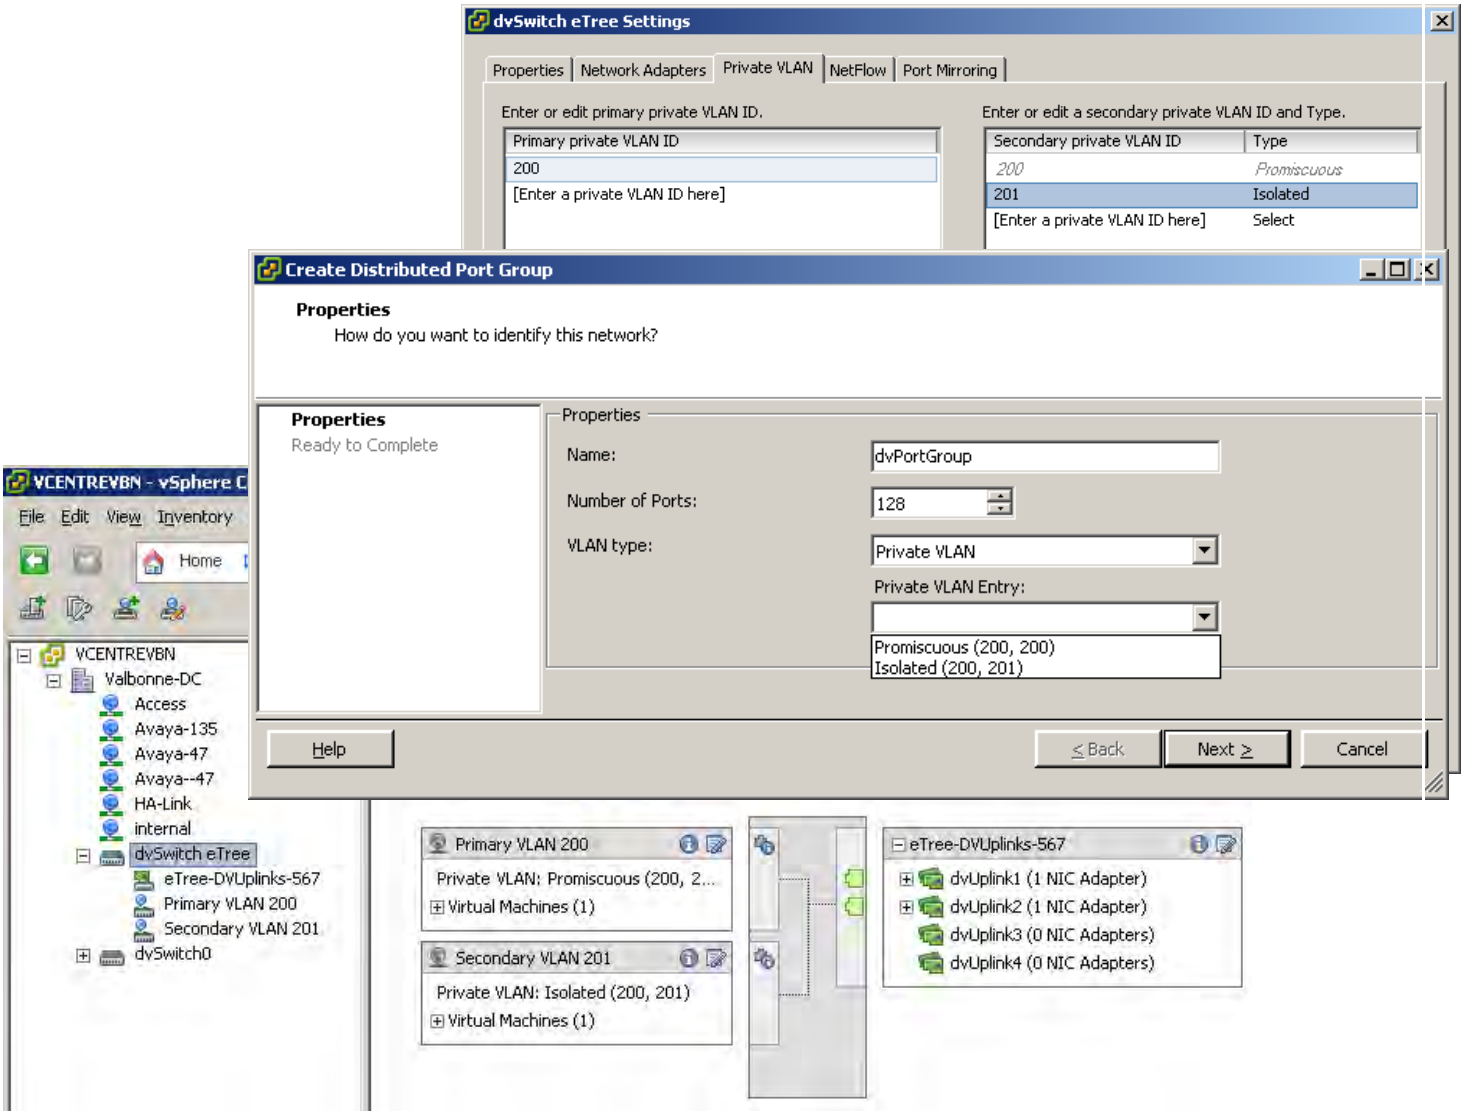

ESX PVLAN Config:

- Create new vSphere Distributed Switch...
- Set Private VLAN instances in tab shown
- Create two new Port Groups
    - one Primary using Promiscuous PVLAN
    - other Secondary using Isolated PVLAN id
- If the Extreme ethernet switch has an MLT configured for the ESX connections, remember to change the default ESX NIC Teaming mode
    - From: “Route based on originating virtual port”
    - To: “Route based on IP hash”

Private VLAN Trunk Configuration

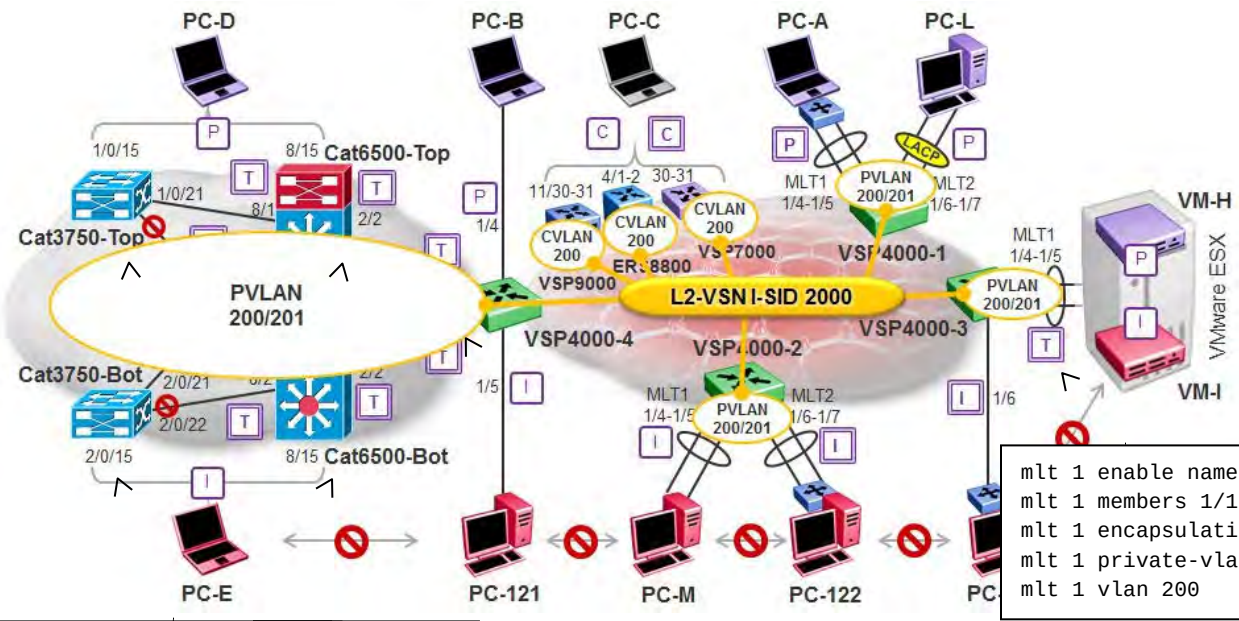


```
mlt 1 enable name ESX-Trunk
mlt 1 members 1/1-1/5
mlt 1 encapsulation dot1q
mlt 1 private-vlan trunk
mlt 1 vlan 200
```

```
interface gigabitEthernet <port>
  switchport
  switchport trunk encapsulation dot1q
  switchport mode trunk
exit
```

```
interface gigabitEthernet 1/47,1/48
  private-vlan trunk
exit
vlan members add 200 1/47,1/48
```

VSP 4000, VSP 7200 & VSP 8000:

Interface

```
interface gigabitEthernet <slot/port>
  private-vlan trunk
exit
vlan members add <vlan-id> <slot/port>
```

MLT

```
mlt <id> enable name <optional name>
mlt <id> member <slot/port>
mlt <id> encapsulation dot1q
mlt <id> private-vlan trunk
mlt <id> vlan <vlan-id>
```

Promiscuous Interface Configuration

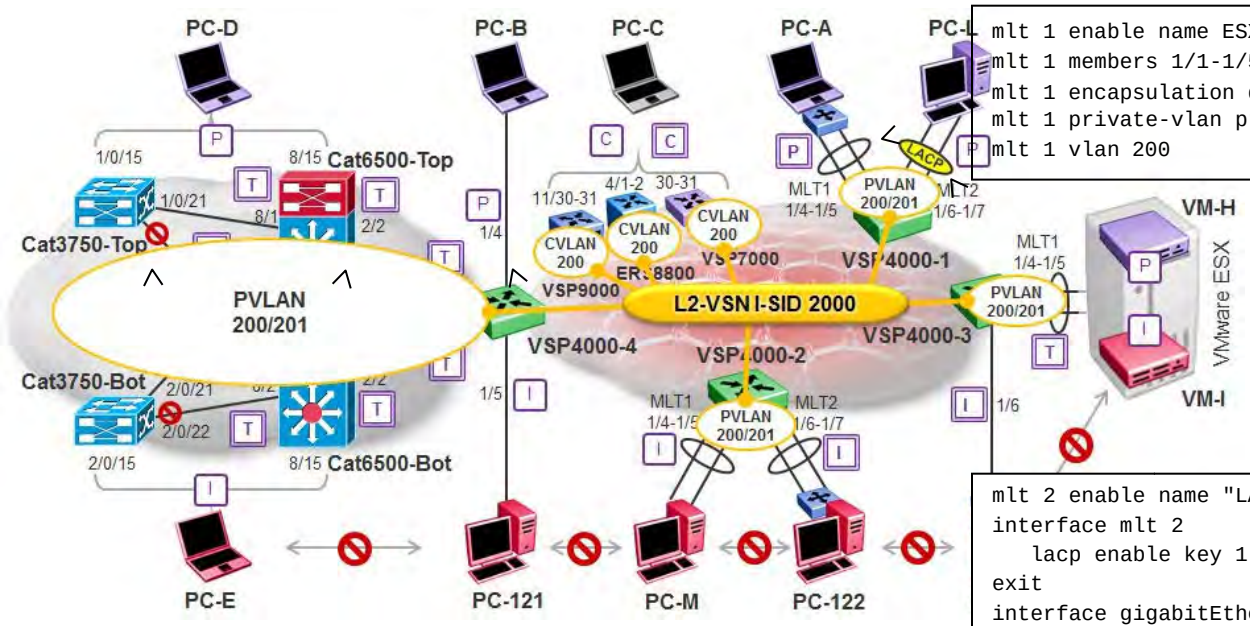


```
mlt 1 enable name ESX-Trunk
mlt 1 members 1/1-1/5
mlt 1 encapsulation dot1q
mlt 1 private-vlan promiscuous
mlt 1 vlan 200
```

```
interface gigabitEthernet <port>
  switchport
  switchport private-vlan mapping 200 201
  switchport mode private-vlan promiscuous
exit
```

```
interface gigabitEthernet 1/4
  private-vlan promiscuous
exit
vlan member add 200 1/4
```

```
mlt 2 enable name "LAG2-UnPr“
interface mlt 2
   lacp enable key 1
exit
interface gigabitEthernet
1/6,1/7
   private-vlan promiscuous
exit
vlan members add 200 1/6,1/7
interface gigabitEthernet
1/6,1/7
   lacp key 1
   lacp aggregation enable
   lacp timeout-time short
   lacp enable
exit
```

VSP 4000, VSP 7200 & VSP8000:

Interface

```
interface gigabitEthernet <slot/port>
  private-vlan promiscuous
exit
vlan members add <vlan-id> <slot/port>
```

MLT

```
mlt <id> enable name <optional name>
mlt <id> member <slot/port>
mlt <id> encapsulation dot1q
mlt <id> private-vlan promiscuous
mlt <id> vlan <vlan-id>
```

LACP

```
mlt 2 enable name <optional name>
interface mlt 2
   lacp enable key <key #>
exit
interface gigabitEthernet <slot/port>
   private-vlan promiscuous
exit
vlan members add <vlan-id> <slot/port>
interface gigabitEthernet <slot/port>
   lacp key <key 3>
   lacp aggregation enable
   lacp timeout-time short
```

```
    lacp enable
exit
```

VMWare ESX Promiscuous VM Configuration

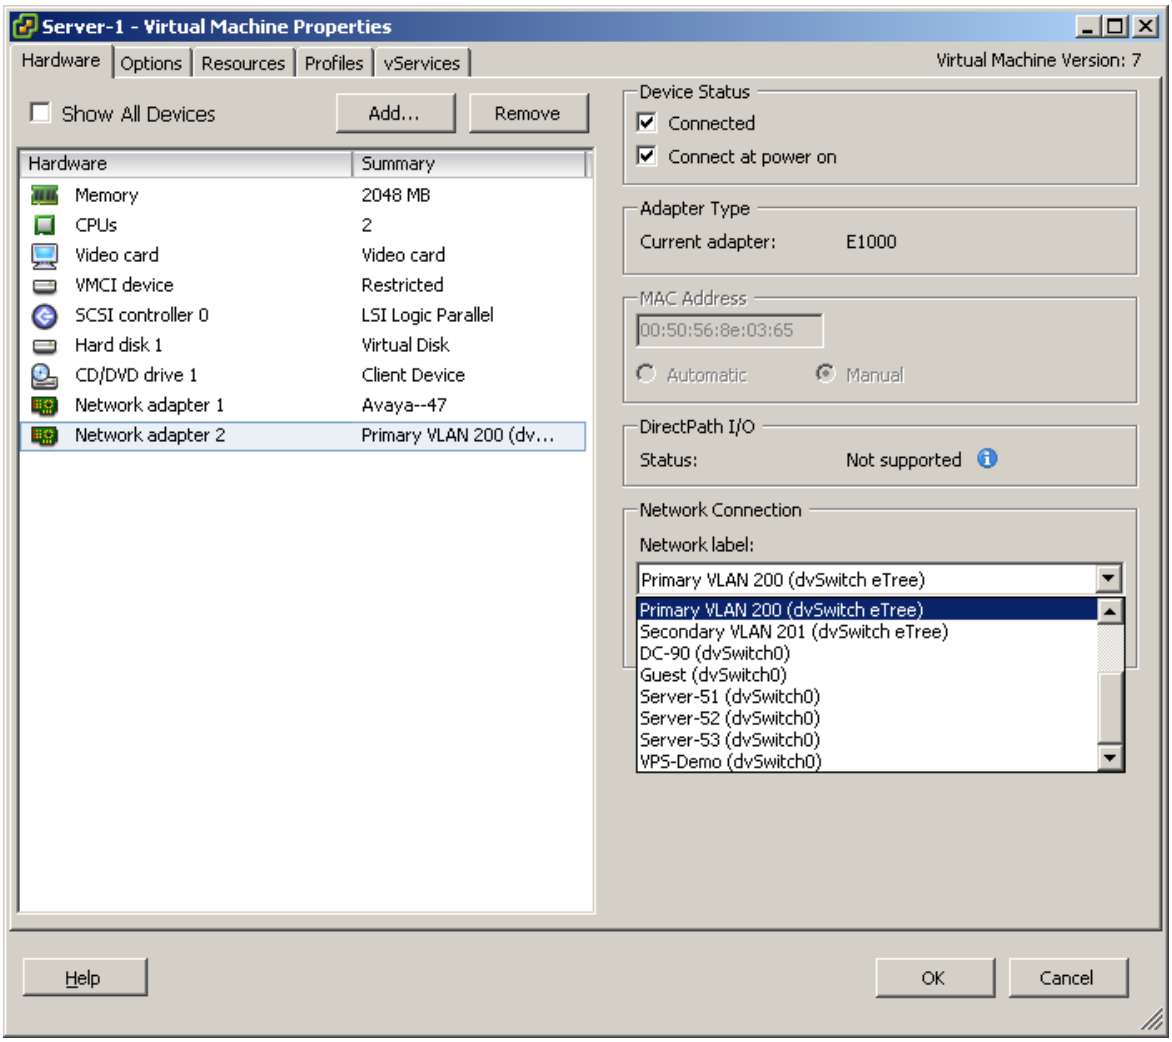

ESX Promiscuous VM Config:

- Assign Primary Port Group to VM

Isolated Interface Configuration

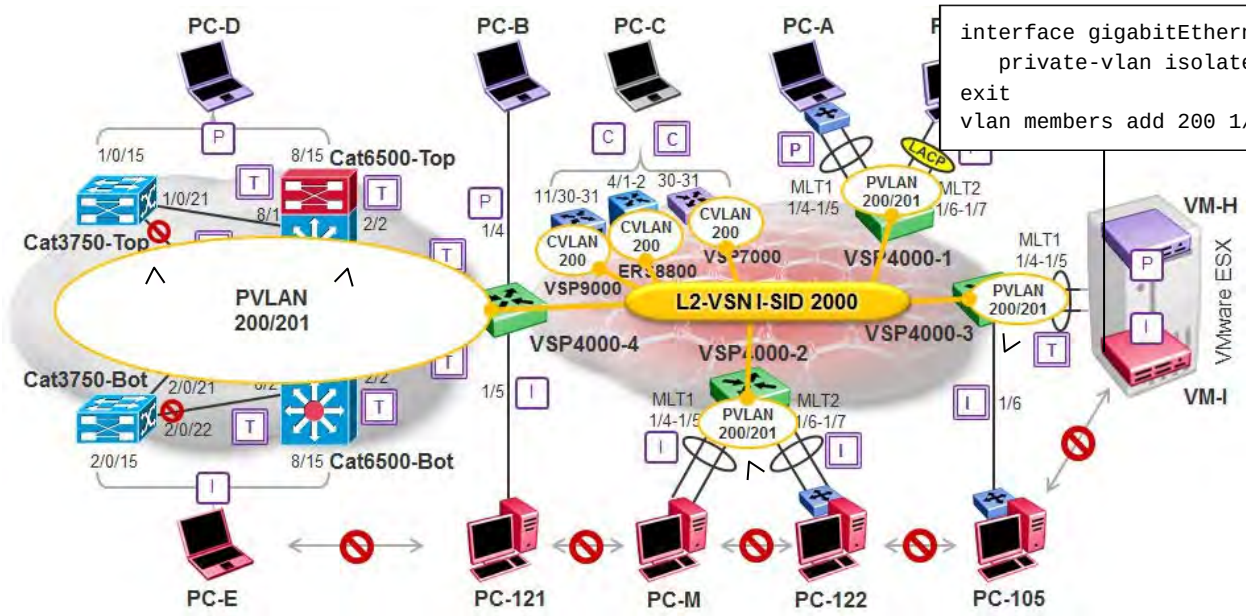


```
interface gigabitEthernet 1/6
   private-vlan isolated
exit
vlan members add 200 1/6
```

```
interface gigabitEthernet <port>
  switchport
  switchport private-vlan host-association 200 201
  switchport mode private-vlan host
exit
```

```
mlt 1 enable name "MLT1-UnIs"
mlt 1 member 1/4-1/5
mlt 1 private-vlan isolated
mlt 1 vlan 200
mlt 2 enable name "MLT2-TgIs"
mlt 2 member 1/6-1/7
mlt 2 encapsulation dot1q
mlt 2 private-vlan isolated
mlt 2 vlan 200
```

VSP 4000, VSP 7200 & VSP 8000:

Interface

```
interface gigabitEthernet <slot/port>
  private-vlan isolated
exit
vlan members add <vlan-id> <slot/port>
```

MLT

```
mlt <id> enable name <optional name>
mlt <id> member <slot/port>
mlt <id> encapsulation dot1q
mlt <id> private-vlan isolated
mlt <id> vlan <vlan-id>
```

VMWare ESX Isolated VM Configuration

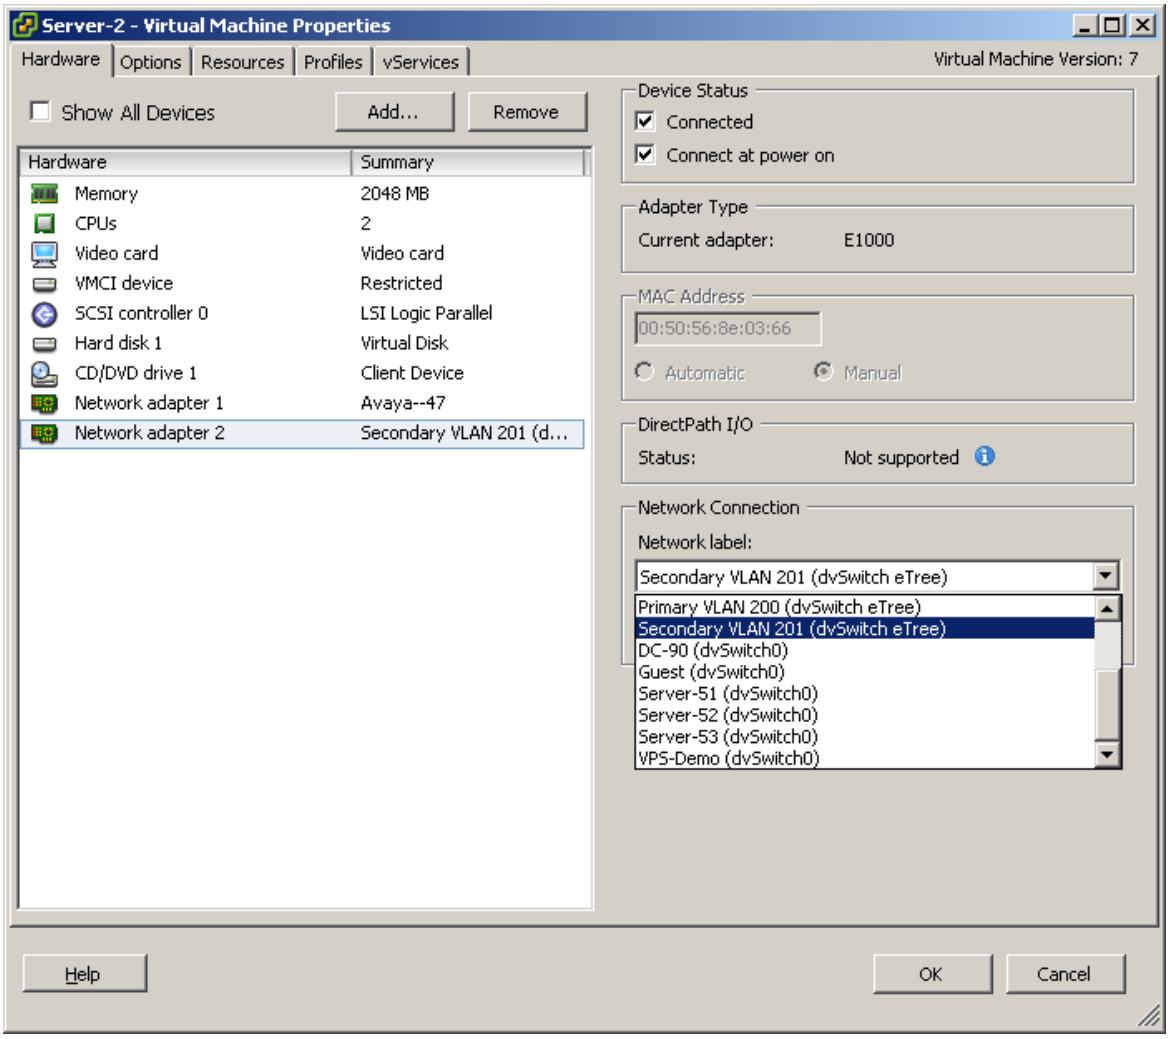

ESX Promiscuous VM Config:

- Assign Secondary Port Group to VM

20.1.13 L3VSN Configuration

L3 VSN with direct interface redistribution

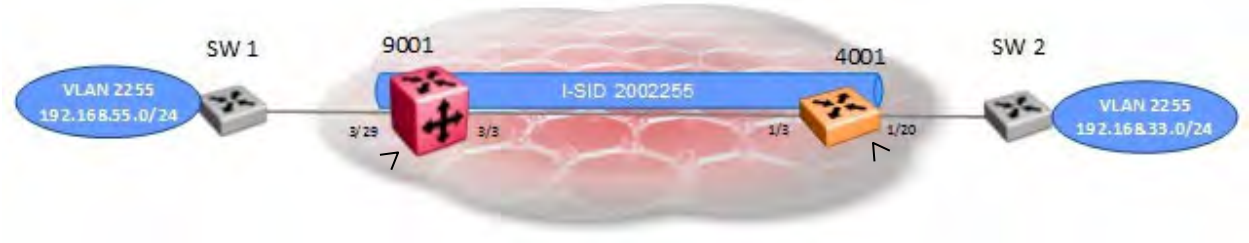


```
configure terminal
interface loopback 1
  ip address 1 10.1.90.1/255.255.255.255
exit
router isis
  ip-source-address 10.1.90.1
  spbm 1 ip enable
exit
ip vrf blue
interface gigabitethernet 3/29
  no shutdown
  encapsulation dot1q
exit
vlan create 2255 type port-mstprstp 0
vlan member 2255 3/29
interface vlan 2255
  vrf blue
  ip address 10.198.55.1 255.255.255.0
exit
router vrf blue
  ipvpn
  i-sid 2002255
  ipvpn enable
  isis redistribute direct
  isis redistribute direct enable
exit
isis apply redistribute direct vrf blue
vlan member remove 1 3/29
```

```
configure terminal
interface loopback 1
  ip address 1 10.1.40.1/255.255.255.255
exit
router isis
  ip-source-address 10.1.40.1
  spbm 1 ip enable
exit
ip vrf blue
interface gigabitethernet 1/20
  no shutdown
  encapsulation dot1q
exit
vlan create 2255 name type port-mstprstp 0
vlan members 2255 1/20
interface vlan 2255
  vrf blue
  ip address 10.198.33.1 255.255.255.0
exit
router vrf blue
  ipvpn
  i-sid 2002255
  ipvpn enable
  isis redistribute direct
  isis redistribute direct enable
exit
isis apply redistribute direct vrf blue
vlan member remove 1 1/20
```

```
config terminal
interface loopback <Interface id value; 1-256>
  ip address <Interface id value; 1-256>  <Ipv4address/mask>
exit
router isis
  ip-source-address <loopback ip>
  spbm 1 ip enable
exit
ip vrf <vrf-name>
vlan create <vlan id> type port-mstprstp 0
```

```
vlan members <vlan id> <slot/port>
interface vlan <vlan id>
  vrf <vrf-name>
  ip address <a.b.c.b mask>
exit
router vrf <vrf-name>
  ipvpn
  ISID <ISID: 0..16000000>
  ipvpn enable
  isis redistribute direct
  isis redistribute direct enable
exit
isis apply redistribute direct vrf <vrf-name>
```

Although you can use any number from 1 to 16,777,215 as an ISID value, it is recommended not to use a value from 16,000,001 to 16,777,215. This range is used for Multicast over SPB.

Although the above example only show direct interface redistribution into ISIS, other protocols such as BGP, OSPF, RIP, and Static can also be enabled.

Verify Operations:

```
show isis spbm (verify SPB IP is enabled globally)
show ip ipvpn
show ip interface vrf <vrf name>
show ip route vrf <vrf name>
ping <ip address> vrf <vrf name> source <local source ip address>
l2 ping ip-address <ip address> vrf <vrf name>
l2 traceroute ip-address <ip address> vrf <vrf name>
show isis lsdb tlv 184 detail
show isis lsdb sysid <system id> tlv 184 detail
```

20.1.14 L3VSN – leaking routes between VRF's

L3 VSN with two separate VRF's (red and blue) with one common shared VRF (shared). Allows routing from the red or blue VRF to the shared VRF and vise-versa, but, no routing between the red and blue VRF's. For all VRFs (red, blue, and shared), please see configuration step in the previous example titled *L3VSN Configuration*.

.


```
configure terminal
router vrf blue
  isis accept i-sid 2002290 enable
exit
isis apply accept vrf blue
router vrf red
  isis accept i-sid 2002290 enable
exit
isis apply accept vrf red
router vrf shared
  ip isid-list users list
2002255,2002256
  isis accept isid-list users enable
exit
isis apply accept vrf shared
```

```
configure terminal
router vrf blue
  isis accept i-sid 2002290 enable
exit
isis apply accept vrf blue
configure terminal
router vrf red
  isis accept i-sid 2002290 enable
exit
isis apply accept vrf red
```

Single ISID:

```
config terminal
router vrf <vrf-name>
  isis accept i-sid < ISID: 0..16000000> enable
exit
isis apply accept vrf <vrf-name>
```

ISID List:

```
config terminal
router vrf <vrf-name>
  ip isid-list <isid-list name> list <List of ISID values>
  isis accept isid-list <isid-list name> enable
exit
isis apply accept vrf <vrf-name>
```

20.1.15 IP Shortcuts

IP Shortcuts with direct interface redistribution

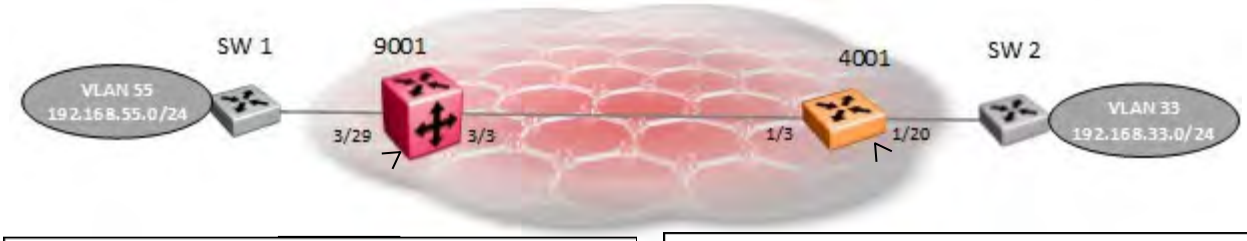


```
configure terminal
interface loopback 1
  ip address 1 10.1.90.1/255.255.255.255
exit
router isis
  ip-source-address 10.1.90.1
  spbm 1 ip enable
exit
interface gigabitethernet 3/29
  no shutdown
  encapsulation dot1q
exit
vlan create 55 type port-mstprstp 0
vlan member 55 3/29
interface vlan 55
  ip address 10.198.55.1 255.255.255.0
exit
router isis
  redistribute direct
  redistribute direct enable
exit
isis apply redistribute direct
vlan member remove 1 3/29
```

```
configure terminal
interface loopback 1
  ip address 1 10.1.40.1/255.255.255.255
exit
router isis
  ip-source-address 10.1.40.1
  spbm 1 ip enable
exit
interface gigabitethernet 1/20
  no shutdown
  encapsulation dot1q
exit
vlan create 33 name type port-mstprstp 0
vlan members 33 1/20
interface vlan 33
  ip address 10.198.33.1 255.255.255.0
exit
router isis
  redistribute direct
  redistribute direct enable
exit
isis apply redistribute direct
vlan member remove 1 1/20
```

```
config terminal
interface loopback < Interface id value; 1-256>
  ip address <Interface id value; 1-256>  <Ipv4address/mask>
exit
router isis
  ip-source-address <loopback ip>
  spbm 1 ip enable
exit
router isis
  redistribute direct
  redistribute direct enable
exit
isis apply redistribute direct
```

Although the above example only shows direct interface redistribution into ISIS, other protocols such as BGP, OSPF, RIP, and Static can also be redistributed.

Verify Operations:

```
show isis spbm (verify SPB IP is enabled globally)
show ip interface
show ip route
show isis spbm ip-unicast-fib
ping <ip address>
l2 ping <ip address>
l2 traceroute ip-address <ip address>
show isis lsdb tlv 135 detail
show isis lsdb sysid <system id> tlv 135 detail
```

20.1.16 IP Shortcut– Suppress IST Network

When IP Shortcuts is enabled with the option of redistribution of direct interfaces into ISIS, all local networks will be advertised including the IST or vIST network used for the IST VLAN. If you wish, you can suppress the IST network by simply adding a route map matching the IP subnet used for the IST or vIST.

Route-map to suppress the IST from being advertised if ISIS direct interface is enabled

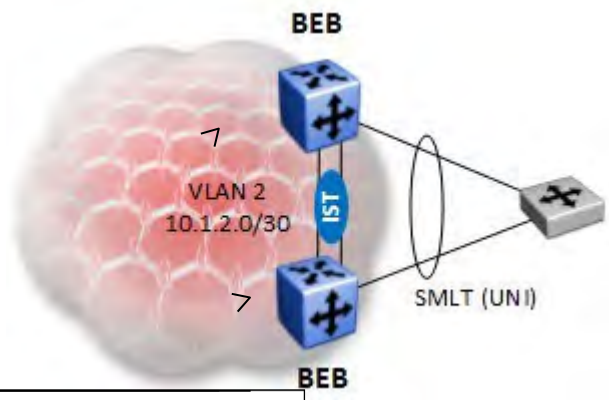


```
ip prefix-list "IST" 10.1.2.0/30
route-map "suppressIST" 1
  no permit
  enable
  match network "IST"
exit
route-map "suppressIST" 2
  enable
exit
router isis
  redistribute direct
  redistribute direct route-map "suppressIST"
  redistribute direct enable
exit
isis apply redistribute direct
```

20.1.17IP Shortcuts – leaking routes between GRT and VRF

IP Shortcuts and L3 VSN (VRF blue) where we wish to share the IP Shortcuts 192.168.55.0/24 subnet to the blue VRF.

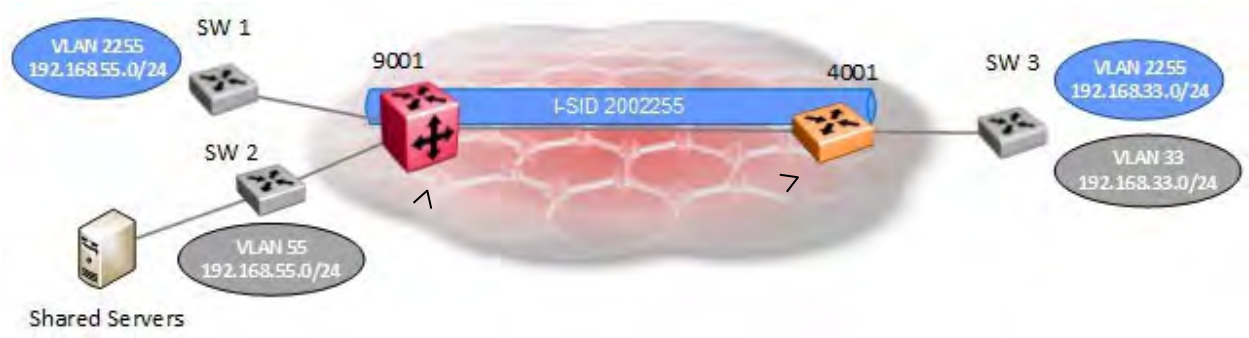


```
configure terminal
router vrf blue
ip prefix-list "shared" 192.168.55.0/24
route-map "shared_map" 1
  match network "shared"
  enable
exit
route-map "shared_map" 2
  no permit
  enable
exit
isis accept ISID 0 route-map "shared_map"
isis accept ISID 0 enable
exit
isis apply accept vrf blue
router isis
  accept ISID 2002255 enable
exit
isis apply accept
```

For the blue VRF used in this example, to accept GRT IP Shortcut routes, we need to specify an ISID of 0 for the ISIS accept policy.

20.1.18IP Shortcuts – redistribution of ISIS and OSPF

This section goes over several methods on how to redistribute ISIS into OSPF and vise-versa.

ERS 8000 – redistribution of ISIS into OSPF and vise-versa

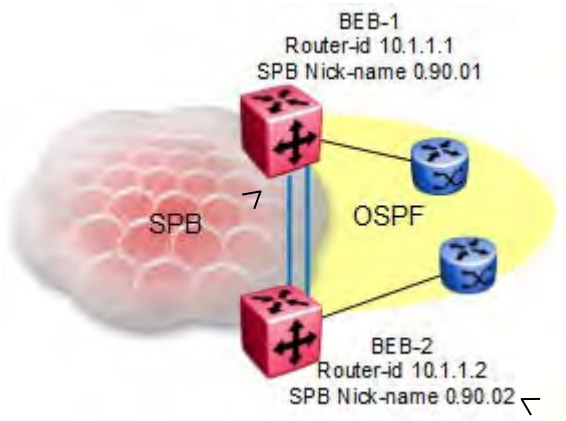


```
configure terminal
no ip alternative-route
ip route preference protocol spbm-level1 130
route-map "reject" 1
  no permit
  enable
exit
router ospf
  as-boundary-router enable
  accept adv-rtr 10.1.1.2 enable route-policy
"reject"
  redistribute isis
  redistribute isis enable
exit
ip ospf apply accept adv-rtr 10.1.1.2
router isis
  redistribute ospf
  redistribute ospf enable
exit
isis apply redistribute ospf
ip ospf apply redistribute isis
```

```
configure terminal
no ip alternative-route
ip route preference protocol spbm-level1 130
route-map "reject" 1
  no permit
  enable
exit
router ospf
  as-boundary-router enable
  accept adv-rtr 10.1.1.1 enable route-policy
"reject"
  redistribute isis
  redistribute isis enable
exit
ip ospf apply accept adv-rtr 10.1.1.1
router isis
  redistribute ospf
  redistribute ospf enable
exit
isis apply redistribute ospf
ip ospf apply redistribute isis
```

Please note the following:

- The ERS 8800 does not support ISIS accept policies
  - There is no way to prevent one border router from accepting ISIS routes from the other border router
  - The solution is to make OSPF preferred over ISIS by assigning ISIS a higher route preference over OSPF
    - We can now use OSPF accept policies to prevent one border router from accepting OSPF external routes from another

VSP 9000 or 4000/7200/8000 prior to release 5.0 – redistribution of ISIS into OSPF and vise-versa

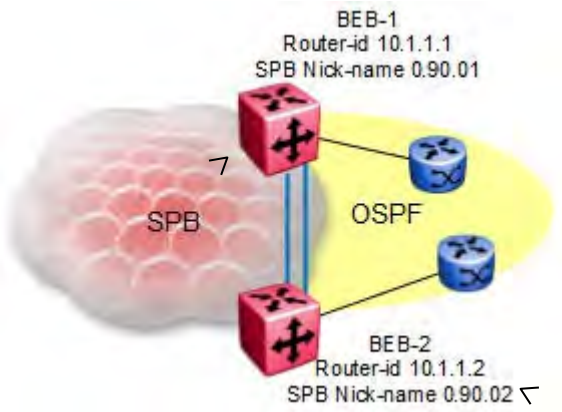


```
configure terminal
no ip alternative-route
route-map "reject" 1
  no permit
  enable
exit
router ospf
  as-boundary-router enable
  redistribute isis
  redistribute isis enable
exit
router isis
  accept adv-rtr 0.90.02
  accept adv-rtr 0.90.02 route-map "reject"
  accept adv-rtr 0.90.02 enable
  redistribute ospf
  redistribute ospf enable
exit
isis apply accept
isis apply redistribute ospf
ip ospf apply redistribute isis
```

```
configure terminal
no ip alternative-route
route-map "reject" 1
  no permit
  enable
exit
router ospf
  as-boundary-router enable
  redistribute isis
  redistribute isis enable
exit
router isis
  accept adv-rtr 0.90.01
  accept adv-rtr 0.90.01 route-map "reject"
  accept adv-rtr 0.90.01 enable
  redistribute ospf
  redistribute ospf enable
exit
isis apply accept
isis apply redistribute ospf
ip ospf apply redistribute isis
```

Please note the following:

- By default, when redistributing OSPF routes into ISIS, all ISIS switches with IP Shortcuts enabled will always use the internal metric with the SPBM cost
  - lower cost preferred, to the BEB nodes advertising the redistributed routes
  - If equal, use the prefix cost (external metric) with lower cost preferred and if prefix cost is the same, the routes are considered ECMP routes
- Instead of using the default internal metric, the ISIS metric type of external can be used as per the next example and is the preferred method for the VSP 4000/7200/8000 as of release 5.0

VSP 4000/7200/8000 release 5.0 or higher – redistribution of ISIS into OSPF and vise-versa using External Metrics

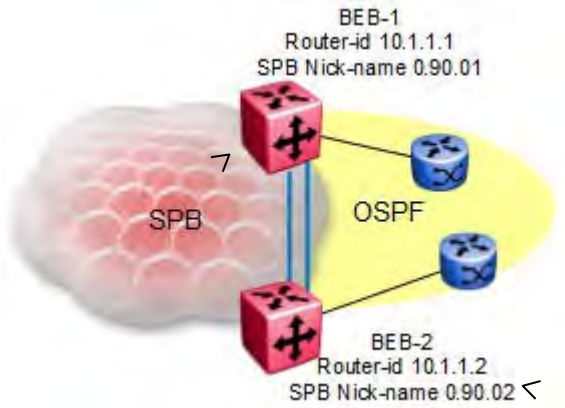


```
configure terminal
no ip alternative-route
route-map "reject" 1
  no permit
  match metric-type-isis external
  enable
exit
router ospf
  as-boundary-router enable
  redistribute isis
  redistribute isis enable
exit
router isis
  accept adv-rtr 0.90.02
  accept adv-rtr 0.90.02 route-map "reject"
  accept adv-rtr 0.90.02 enable
  redistribute ospf
  redistribute ospf metric-type external
  redistribute ospf enable
exit
isis apply accept
isis apply redistribute ospf
ip ospf apply redistribute isis
```

```
configure terminal
no ip alternative-route
route-map "reject" 1
  no permit
  match metric-type-isis external
  enable
exit
router ospf
  as-boundary-router enable
  redistribute isis
  redistribute isis enable
exit
router isis
  accept adv-rtr 0.90.01
  accept adv-rtr 0.90.01 route-map "reject"
  accept adv-rtr 0.90.01 enable
  redistribute ospf
  redistribute ospf metric-type external
  redistribute ospf enable
exit
isis apply accept
isis apply redistribute ospf
ip ospf apply redistribute isis
```

Please note the following:

- The ISIS metric type of external can be used instead of internal type
  - Prefix cost (external metric) with lowest cost will be preferred and if prefix cost is the same, the routes are considered ECMP routes
- A route-map can be used to specify metric type of none, internal or external; for example, a route-map can be used to selectively set specific routes as external
- When deciding which routes to add to the route table, a SPB router
  - Internal type routes
    - Always preferred over External type routes
    - Will always prefer the routes with the shortest path internal metric to the BEB node advertising them
    - Will only use the route external cost (prefix-cost) as a tie breaker

        - For route to go into ECMP, both the internal and external metric must be the same
    - External type routes
        - Will only consider the external route metric (prefix-cost)
        - For route to go into ECMP, only the external metric must be the same
- External metrics is supported as of release 5.0 for the VSP 4000, VSP 8000, and VSP 7200
    - Only IPv4 routes are supported in this release

VSP 4000/7200/8000 – Using a Route-Map to deny routes with External Metrics

```
configure terminal
route-map deny_ext 1
  no permit
  match metric-type-isis external
  enable
exit
router isis
  accept route-map deny_ext
exit
isis applyaccept
```

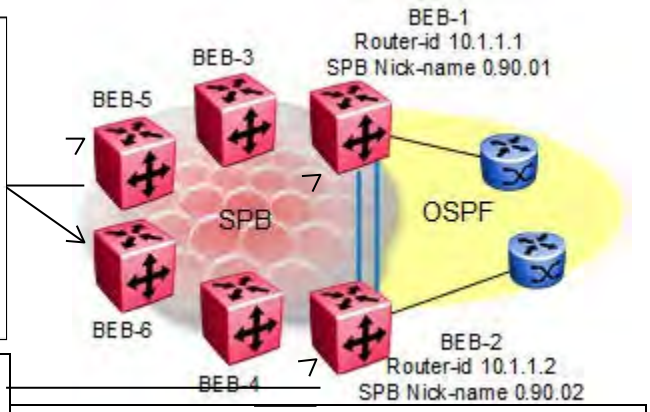


```
configure terminal
no ip alternative-route
route-map "reject" 1
  no permit
  enable
exit
router ospf
  as-boundary-router enable
  redistribute isis
  redistribute isis enable
exit
router isis
  accept adv-rtr 0.90.02
  accept adv-rtr 0.90.02 route-map "reject"
  accept adv-rtr 0.90.02 enable
  redistribute ospf
  redistribute ospf metric-type external
  redistribute ospf enable
exit
isis apply accept
isis apply redistribute ospf
ip ospf apply redistribute isis
```

```
configure terminal
no ip alternative-route
route-map "reject" 1
  no permit
  enable
exit
router ospf
  as-boundary-router enable
  redistribute isis
  redistribute isis enable
exit
router isis
  accept adv-rtr 0.90.01
  accept adv-rtr 0.90.01 route-map "reject"
  accept adv-rtr 0.90.01 enable
  redistribute ospf
  redistribute ospf metric-type external
  redistribute ospf enable
exit
isis apply accept
isis apply redistribute ospf
ip ospf apply redistribute isis
```

Please note the following:

- BEB-1 and BEB-2 are configured to redistribute OSPF with external metric type
- Only for BEB-5 and BEB-6, an accept policy is created to simply deny all external routes (OSPF routes) advertised from BEB-1 and BEB-2
    - End result, no external OSPF networks will be added to BEB-5 and BEB-6 route table
- BEB-3 and BEB-4 will add the OSPF networks into their respective route tables as ISIS routes

VSP 4000/7200/8000 – Change metric on BEB-1 to a higher value (worst metric) to use BEB-2 as default switch for all OSPF external routes

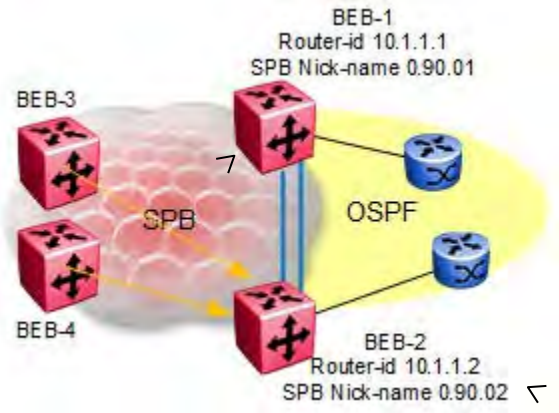


```
configure terminal
no ip alternative-route
route-map "reject" 1
  no permit
  enable
exit
router ospf
  as-boundary-router enable
  redistribute isis
  redistribute isis metric 200
  redistribute isis enable
exit
router isis
  accept adv-rtr 0.90.02
  accept adv-rtr 0.90.02 route-map "reject"
  accept adv-rtr 0.90.02 enable
  redistribute ospf
  redistribute ospf metric-type external
  redistribute ospf metric 200
  redistribute ospf enable
exit
isis apply accept
isis apply redistribute ospf
ip ospf apply redistribute isis
```

```
configure terminal
no ip alternative-route
route-map "reject" 1
  no permit
  enable
exit
router ospf
  as-boundary-router enable
  redistribute isis
  redistribute isis enable
exit
router isis
  accept adv-rtr 0.90.01
  accept adv-rtr 0.90.01 route-map "reject"
  accept adv-rtr 0.90.01 enable
  redistribute ospf
  redistribute ospf metric-type external
  redistribute ospf enable
exit
isis apply accept
isis apply redistribute ospf
ip ospf apply redistribute isis
```

Please note the following:

- The lowest external metric wins
  - As BEB-1 is setting the external metric for all OSPF routes redistributed into ISIS with a metric of 200 added, BEB-2 will be selected as it will have the best cost to all OSPF networks
  - As BEB-1 is advertising all ISIS routes into OSPF with a Type 2 metric of 200, all OSPF routers will select the router connected to BEB-2 as the preferred router
    - As alternative to setting the OSPF redistribution metric to 200 on BEB-1, we could have configured BEB-2 to redistribute ISIS using Type 1 external metric; OSPF External Type 1 is preferred over Type 2

20.1.19 Inter-VSN Routing

Inter-VSN Routing

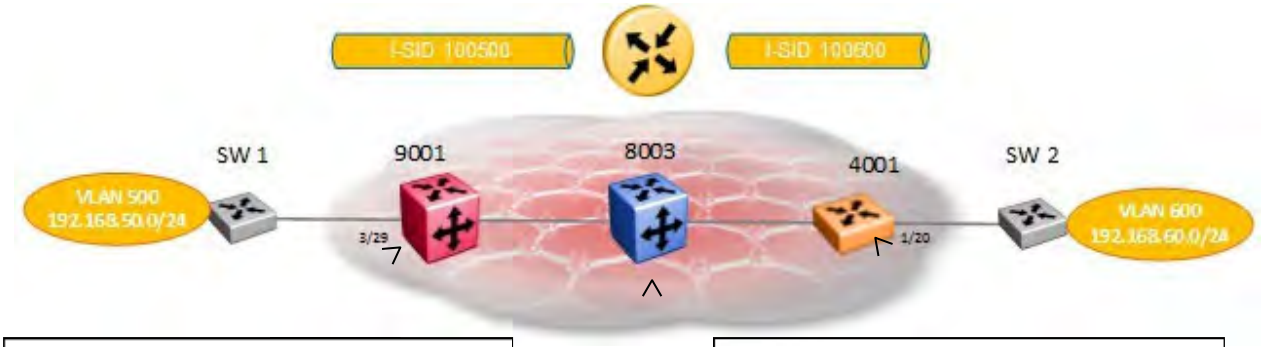


```
configure terminal
interface gigabitethernet 3/29
  no shutdown
  encapsulation dot1q
exit
vlan create 500 type port-mstprstp 0
vlan member 500 3/29
vlan i-sid 500 100500
vlan member remove 1 3/29
```

```
configure terminal
interface gigabitethernet 1/3
  no shutdown
  encapsulation dot1q
exit
vlan create 600 type port-mstprstp 0
vlan member 600 1/20
vlan i-sid 600 100600
vlan member remove 1 1/20
```

```
configure terminal
vlan create 500 type port-mstprstp 0
vlan create 600 type port-mstprstp 0
vlan i-sid 500 100500
vlan i-sid 600 100600
ip vrf intervlan
interface vlan 500
  vrf intervlan
  ip address 192.168.50.1 255.255.255.0
exit
interface vlan 600
  vrf intervlan
  ip address 192.168.60.1 255.255.255.0
exit
```

```
config terminal
vlan create <vlan id> type port-mstprstp 0
vlan i-sid <vlan-id> <ISID: 0..16000000>
ip vrf <vrf-name>
interface vlan <vlan id>
  vrf <vrf-name>
  ip address <a.b.c.b mask>
exit
```

Although you can use any number from 1 to 16,777,215 as an ISID value, it is recommended not to use a value from 16,000,001 to 16,777,215. This range is used for Multicast over SPB.

The VRF portion of the configuration can be added on any SPB switch in the network. For redundancy, the VRF portion of the configuration should be added on another SPB switch with VRRP Backup Master enabled.

For redundancy, it is recommended to enable Inter-ISID on another SPB switch in the network and enable VRRP with Backup Master.

```
config terminal
vlan create <vlan id> type port-mstprstp 0
vlan i-sid <vlan-id> <ISID: 0..16000000>
ip vrf <vrf-name>
interface vlan <vlan id>
  vrf <vrf-name>
  ip address <a.b.c.b mask>
  ip vrrp address <Vrid> <a.b.c.d>
  ip vrrp <1-255 - Vrid> backup-master enable
  ip vrrp <1-255 - Vrid> priority <1-255>
  ip vrrp <Vrid> enable
exit
```

20.1.20 IPv6 Shortcuts

IPv6 Shortcuts with direct interface redistribution

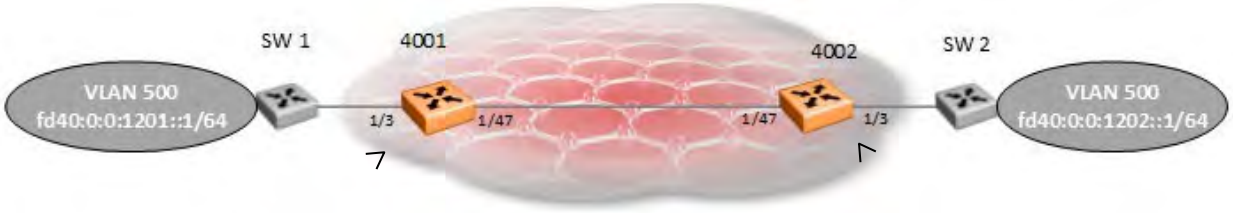


```
configure terminal
interface loopback 1
  ip address 1 10.4.4.1.1/255.255.255.255
  ipv6 interface address fd40::4:4:1/128
exit
router isis
  ip-source-address 10.1.90.1
  ipv6-source-address fd40::4:4:1
  spbm 1 ip enable
  spbm 1 ipv6 enable
exit
interface gigabitethernet 1/3
  no shutdown
  encapsulation dot1q
exit
vlan create 500 type port-mstprstp 0
vlan member 500 1/3
interface vlan 500
  ipv6 interface enable
  ipv6 interface address fd40:0:0:1201::1/64
exit
ipv6 forwarding
router isis
  ipv6 redistribute direct enable
exit
isis apply redistribute direct
vlan member remove 1 1/3
```

```
configure terminal
interface loopback 1
  ip address 1 10.4.4.2/255.255.255.255
  ipv6 interface address fd40::4:4:2/128
exit
router isis
  ip-source-address 10.1.40.1
  ipv6-source-address fd40::4:4:2
  spbm 1 ip enable
  spbm 1 ipv6 enable
exit
interface gigabitethernet 1/3
  no shutdown
  encapsulation dot1q
exit
vlan create 500 name type port-mstprstp 0
vlan members 500 1/3
interface vlan 500
  ipv6 interface enable
  ipv6 interface address fd40:0:0:1202::1/64
exit
ipv6forwarding
router isis
  ipv6 isis redistribute direct enable
exit
isis apply redistribute direct
vlan member remove 1 1/3
```

For IPv6 Shortcuts, IPv4 Shortcuts must be enabled with the addition of an IPv6 loopback address.

```
config terminal
interface loopback <Interface id value; 1-256>
  ip address <Interface id value; 1-256>  <Ipv4address/mask>
  ipv6 interface address <Ipv6address/prefix-len>
exit
router isis
  ip-source-address <loopback ip>
  ipv6-source-address <ipv6 loopback ip>
  spbm 1 ip enable
```

```
  spbm 1 ipv6 enable
exit
ipv6 forwarding
router isis
  redistribute direct
  redistribute direct enable
exit
isis apply redistribute direct
```

Although the above example only show direct interface redistribution into ISIS, other protocols such as OSPFv3 and Static can also be enabled.

Verify Operations:

```
show ipv6 interface
show ipv6 route
show isis spbm ipv6-unicast-fib
ping <ipv6 address>
l2 ping ip-address <ipv6 address>
l2 traceroute ip-address <ipv6 address>
show isis lsdb tlv 236 detail
show isis lsdb sysid <system id> tlv 236 detail
```

20.1.21 SPB Multicast Configuration

20.1.21.1 L2VSN Multicast

Enabling SPM Multicast

```
config terminal
router isis
   spbm 1 multicast enable
exit
interface vlan <vlan id>
   ip igmp snoop
   ip igmp snoop-querier-addr <ip addr>
   ip igmp ssm-snoop **If IGMPv3 is required
   ip igmp version 3 **If IGMPv3 is required
exit
```

ERS 4800 & ERS 5900:

```
config terminal
router isis
   spbm 1 multicast enable
exit
interface vlan <vlan id>
   ip igmp snooping
   ip igmp snoop-querier-addr <ip addr>
   ip igmp version 3 **If IGMPv3 is required
exit
```

For multicast over L2VSN's, please note if the SPB bridge is connected to an edge switch, it may be necessary to add an IGMP query address. If you omit adding a query address, the SPB bridge will send IGMP queries with a source address of 0.0.0.0. Depending on the edge switch model, it may not accept a query with a source address of 0.0.0.0. This is the case if using an Extreme stackable edge switch that supports IGMPv3.

Verify Operations:

```
show isis spbm (verify multicast is enabled)
show ip igmp interface
show isis spbm ip-multicast-route all
show isis spbm ip-multicast-route vsn-isid <ISID value>
show isis lsdb tlv 185 detail
show isis lsdb sysid <system id> tlv 185 detail
show ip igmp sender
```

The following commands can be performed on receiving BEB nodes:

```
show ip igmp cache
show ip igmp group
```

The following commands can be performed on sending BEB node:

```
l2 tracemroute source <ip addr of source> group <group ip address> vlan <C-VLAN id>
```

20.1.21.2 L3VSN Multicast

Enabling SPM Multicast

```
config terminal
router isis
   spbm 1 multicast enable
exit
router vrf <vrf name>
  mvpn enable
exit
interface vlan <vlan id>
   ip spb-multicast enable
   ip igmp version 3 **If IGMPv3 is required
exit
```

Verify Operations:

```
show isis spbm (verify multicast is enabled)
show ip igmp interface vrf <vrf name>
show isis spbm ip-multicast-route vrf <vrf name>
show isis spbm ip-multicast-route vrf <vrf name> group <group ip address>
show isis lsdb tlv 185 detail
show isis lsdb sysid <system id> tlv 185 detail
```

The following commands can be performed on the sending BEB node:

```
show ip igmp sender vrf <vrf name>
```

The following commands can be performed on receiving BEB nodes:

```
show ip igmp cache vrf <vrf name>
show ip igmp group vrf <vrf name>
```

20.1.21.3 IP Shortcuts Multicast

Enabling SPM Multicast

```
config terminal
router isis
   spbm 1 multicast enable
exit
interface vlan <vlan id>
   ip spb-multicast enable
   ip igmp version 3 **If IGMPv3 is required
exit
```

Verify Operations:

```
show isis spbm (verify multicast is enabled)
show ip igmp interface
show isis spbm ip-multicast-route
show isis spbm ip-multicast-route all
```

```
show isis spbm ip-multicast-route group <group ip address> source <source ip>
show isis lsdb tlv 186 detail
show isis lsdb sysid <system id> tlv 186 detail
```

The following commands can be performed on the sending BEB node:

```
show ip igmp sender
l2 tracemroute source <ip addr of source> group <group ip address>
```

The following commands can be performed on receiving BEB nodes:

```
show ip igmp cache
show ip igmp group
```

20.1.22 Multicast 239.255.255/24 – UPnP Filtering

Please be aware that if protocols such as Microsoft Universal Play and Play (UPnP) are enabled, multicast addresses in the 239.255.255/24 may be seen by SPB bridges depending if protocols such as UPnP is enabled on a Microsoft host. If you wish to deny the 239.255.255/24 address space, either an IGMP access list can be created at a VLAN level or an ACL can be created at a port level.

IGMP Access List – Assuming the local VLAN is using an IP subnet of 10.14.10.0/24.

VSP 9000 and ERS 8800:

```
ip prefix-list UPnP 239.255.255.0/24
interface vlan 10
ip igmp access-list "UPnP" 10.14.10.0/255.255.255.0 deny-both
exit
```

VSP 4000/7200/8000:

```
ip prefix-list UPnP 239.255.255.0/24
interface vlan 10
   ip igmp access-list "UPnP" deny-both
exit
```

```
show ip igmp access
================================================================================
                          IGMP Access - GlobalRouter
================================================================================
INTERFACE  GRP PREFIX       HOSTADDR         HOSTMASK         ACCESSMODE
 ------------------------------------------------------------------------------
Vlan10     UPnP             10.14.10.0       255.255.255.0    deny-both
```

ACL – VSP 4000/7200/8000/9000

```
filter acl 1 type inPort
filter acl 1 enable
filter acl port 1 <port list>
filter acl set 1 default-action permit
filter acl ace 1 1
filter acl ace action 1 1 deny
filter acl ace ethernet 1 1 ether-type eq 0x800
filter acl ace ip 1 1 dst-ip mask 239.255.255.0 8
filter acl ace 1 1 enable
```

ACL – ERS 8000

```
filter act 1
filter act 1 ip dstIp
filter act 1 ethernet etherType
filter apply act 1
filter acl 1 type inPort act 1
filter acl port 1 <port list>
filter acl set 1 default-action permit
filter acl ace 1 1
filter acl ace action 1 1 deny
filter acl ace ethernet 1 1 ether-type eq 0x800
filter acl ace ip 1 1 dst-ip eq 239.255.255.0-239.255.255.255
filter acl ace 1 1 enable
```

20.1.23 Connectivity Fault Management (CFM) Configuration

20.1.23.1 Manual CFM Configuration: Software releases 7.0 and 7.1 for the ERS 8800 and 3.3 for the VSP 9000

A Maintenance Domain (MD) up to 22 characters must be defined. To simplify the configuration when migrating to a future software release that support the simplified configuration for CFM, it is recommended to use a MD name of *spbm*. As two B-BVLANs are presently supported, a Maintenance Association (MA) for each B-VLAN must be defined if you wish to use CFM for testing on both of these B-BLANs. Assuming we have B-VLANs 4051 and 4052 defined, we will create two MA's with names of 4051 and 4052. If a Maintenance End Point (MEP) is defined, only a single value is supported for each MA.

CFM assuming MD = *spbm*, MA = *4051 & 4052*, and MEP = *2*

```
config terminal
cfm maintenance-domain spbm
cfm maintenance-association spbm 4051
cfm maintenance-association spbm 4052
cfm maintenance-endpoint spbm 4051 2 state enable
cfm maintenance-endpoint spbm 4052 2 state enable
vlan nodal-mep 4051 spbm 4051 2
vlan nodal-mep 4052 spbm 4052 2
```

20.1.23.2 Simplified CFM Configuration:

Starting in software release 7.1.1 for the ERS 8800, 3.4 for the VSP 9000, 10.2 for the VSP 7000, 5.7 for the ERS 4800, and 3.0 for the VSP 4000, CFM commands will automatically create a MEP and a MIP at a specific level for every SPB B-VLAN provisioned on the switch. Hence, you no longer have to configure explicit MEPs and MIPs and associated VLANs with MEPs and MIPs.

CFM – simplified configuration

```
config terminal
cfm cmac mepid <1-8191>
cfm cmac level <0-7>
cfm cmac enable
cfm spbm mepid <1-8191>
cfm spbm level <0-7>
cfm spbm enable
```

CMAC provisioning is only required on BEB where C-VLANs are terminated and is not supported at this time for the VSP 7000, VSP 8000, VSP 7200, ERS 5900, and ERS 4800.

Verify results using default values

```
ERS-8800:5# show cfm md info
================================================================================
                          Maintenance Domain
================================================================================
Domain Name              Domain Index    Level Domain Type
--------------------------------------------------------------------------------
cmac                     1               4     NODAL
spbm                     2               4     NODAL

 Total number of Maintenance Domain entries: 2.
ERS-8800:5# show cfm mep info
========================================================================
                          Maintenance Endpoint Config
========================================================================
DOMAIN                   ASSOCIATION             MEP  ADMIN
NAME                     NAME                    ID
------------------------------------------------------------------------
cmac                     1                       1    enable
cmac                     10                      1    enable
spbm                     4051                    1    enable
spbm                     4052                    1    enable
================================================================================
                          Maintenance Endpoint Service
================================================================================
DOMAIN_NAME              ASSN_NAME               MEP_ID TYPE    SERVICE_DESCRIPTION
--------------------------------------------------------------------------------
cmac                     1                       1      nodal   Vlan 1,    Level 4
cmac                     10                      1      nodal   Vlan 10,   Level 4
spbm                     4051                    1      nodal   Vlan 4051, Level 4
spbm                     4052                    1      nodal   Vlan 4052, Level 4
```

20.1.24CFM Configuration Example – 7.1.1.x or higher

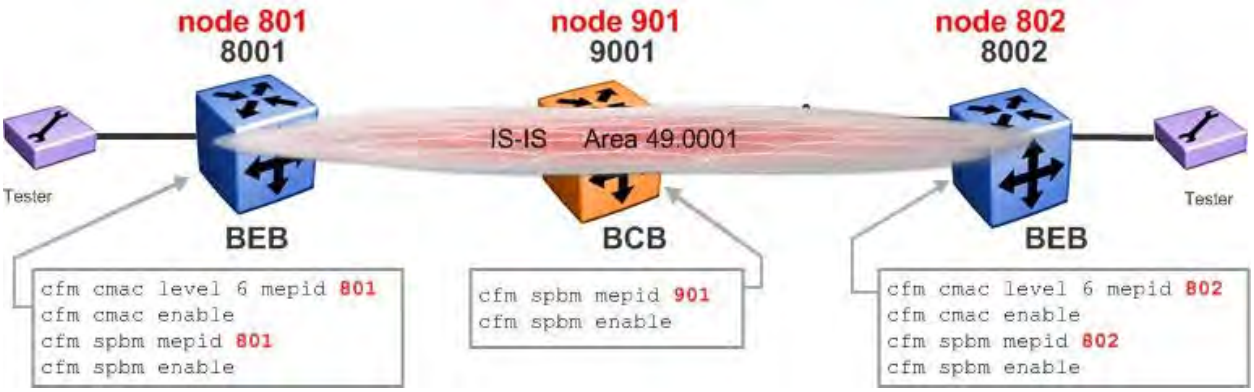

- This ensures that we get full OAM functionalities across:
  - SPB -> Backbone VLAN-ids (BVIDs) i.e. Infrastructure
  - CMAC -> Customer VLANs (CVLANs) i.e. Services
- If a node is acting as a BCB (i.e. it has no CVLANs) no point enabling CFM CMAC on it
- Use a higher level (6) on CMAC CFM
- Leave default level (4) on SPBM CFM

20.1.25Fabric Extend Configuration

Fabric Extend Configuration in VRF

```
config terminal
router isis
  ip-tunnel-source-address <tunnel IP source address> vrf <vrf name> ##see note below
exit
logical-intf isis <1-255> dest-ip <ISIS logical interface destination IP> name <word>
  isis
  isis spbm <1-100>
  isis spbm 1 l1-metric <1-16777215; optional if you wish to set the metric>
  isis enable
exit
```

Please note that the IP Tunnel Source Address must either be a brouter or a loopback address.

Fabric Extend Configuration in GRT

When the tunnel source address interface is configured in the GRT, a route-map must be created to match all the IP addresses used for the tunnel, i.e. any loopback, brouter, or VLAN addresses used for the tunnel and WAN addressing. This is important to prevent these networks from being advertised by the switch into ISIS for example when direct interface redistribution is enabled. For this reason, it is recommended to use the VRF method.

```
config terminal
router isis
  ip-tunnel-source-address <tunnel IP source address> ##see note below
exit
logical-intf isis <1-255> dest-ip <ISIS logical interface destination IP> name <word>
  isis
  isis spbm <1-100>
  isis spbm 1 l1-metric <1-16777215; optional if you wish to set the metric>
  isis enable
exit
```

Please note that the IP Tunnel Source Address must either be a brouter or a loopback address.

Fabric Extend Configuration in VID Tunneling

VID tunneling is supported over either a port or MLT using the configuration shown below. Please note that two VLAN Id's are required, one for each B-VLAN where the primary VID must be provisioned the same at both ends.

```
logical-intf isis <1-255> vid <vid 1,vid 2> primary <vid 1> port <slot|port>
logical-intf isis <1-255> vid <vid 1,vid 2> primary <vid 1> mlt <1-512>
```

20.1.26 ONA: Assigning a Static IP address to the Open Network Adapter

Assuming if we wish to add a static IP address to the ONA, please perform the following steps to factory default the ONA:

- Place PC to device side
- Power down the ONA – remove PoE on network side or remove DC adapter
- Hold down the mode switch located on the device side of the ONA, i.e. with a paper clip

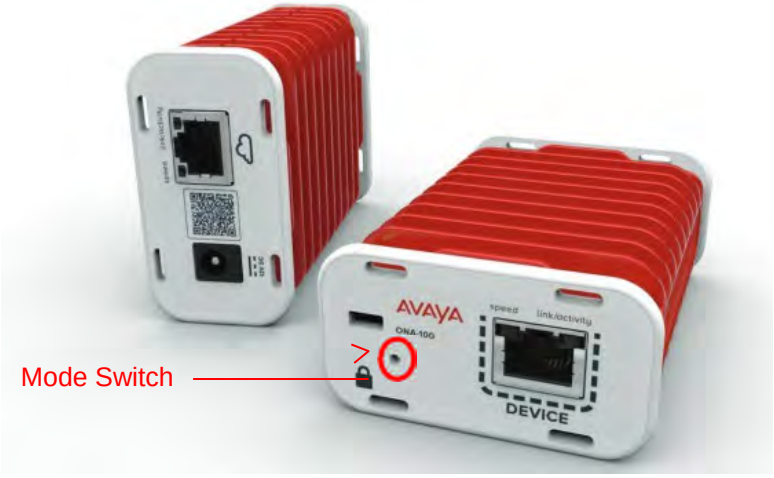

- While holding down the mode switch, power on the ONA via PoE on the WAN side of the ONA or using the power adapter
- After 5 seconds or longer, let go of mode switch
    - If you hold the mode switch for less than 5 seconds, it will simply reset the ONA
- PC should get an IP address in the 192.168.100.x/24 range from the supplied by the ONA

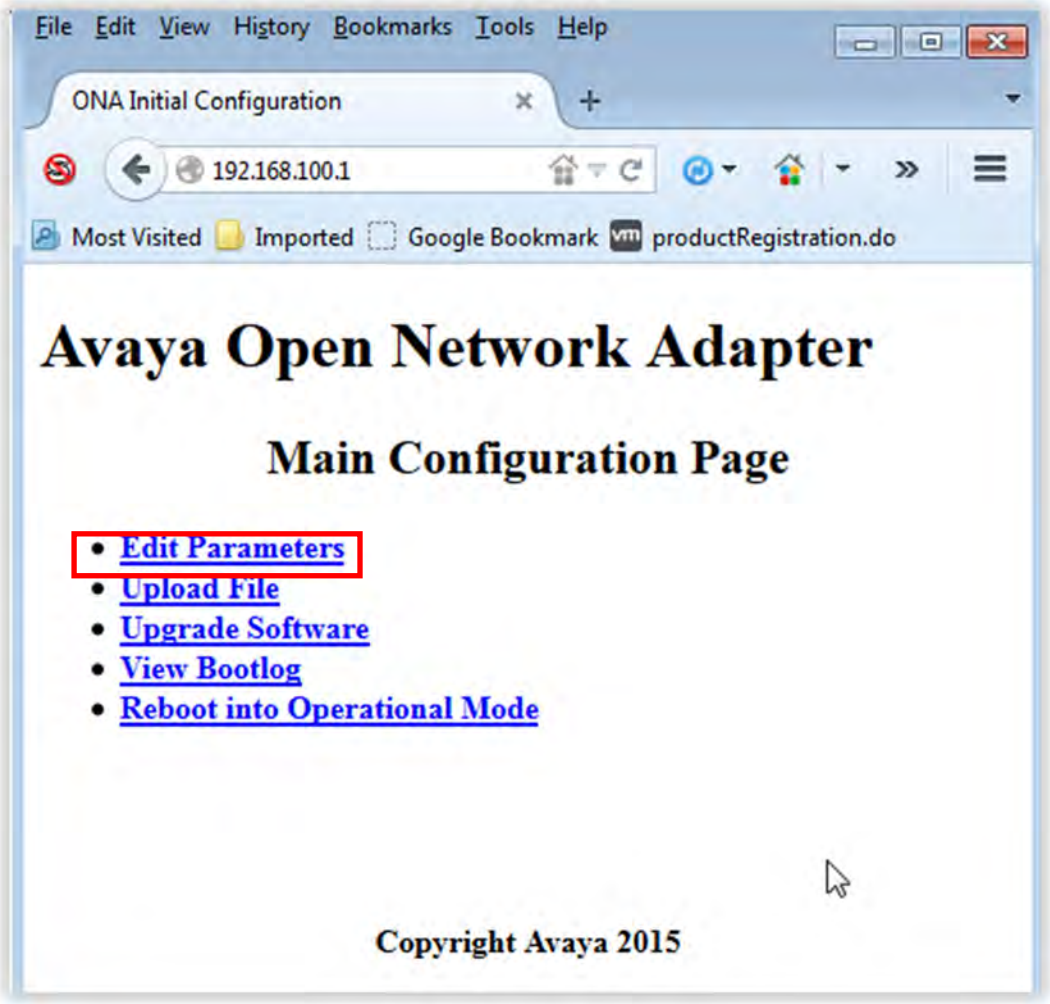

- Point browser to 192.168.100.1 and click on *Edit Parameters*
  - Set *Operational Mode* = *1* and at minimum, set the IP address, IP Mask, and Default Gateway
    - If using DHCP, only set the Operation Mode to 1
  - Click on *Save* when done and then click on *Return to Main Configuration Page* when the Extreme Open Network Adapter page is displayed
  - Either power off/on the ONA or click the *Reboot into Operational Mode* via the *Main Configuration Page*

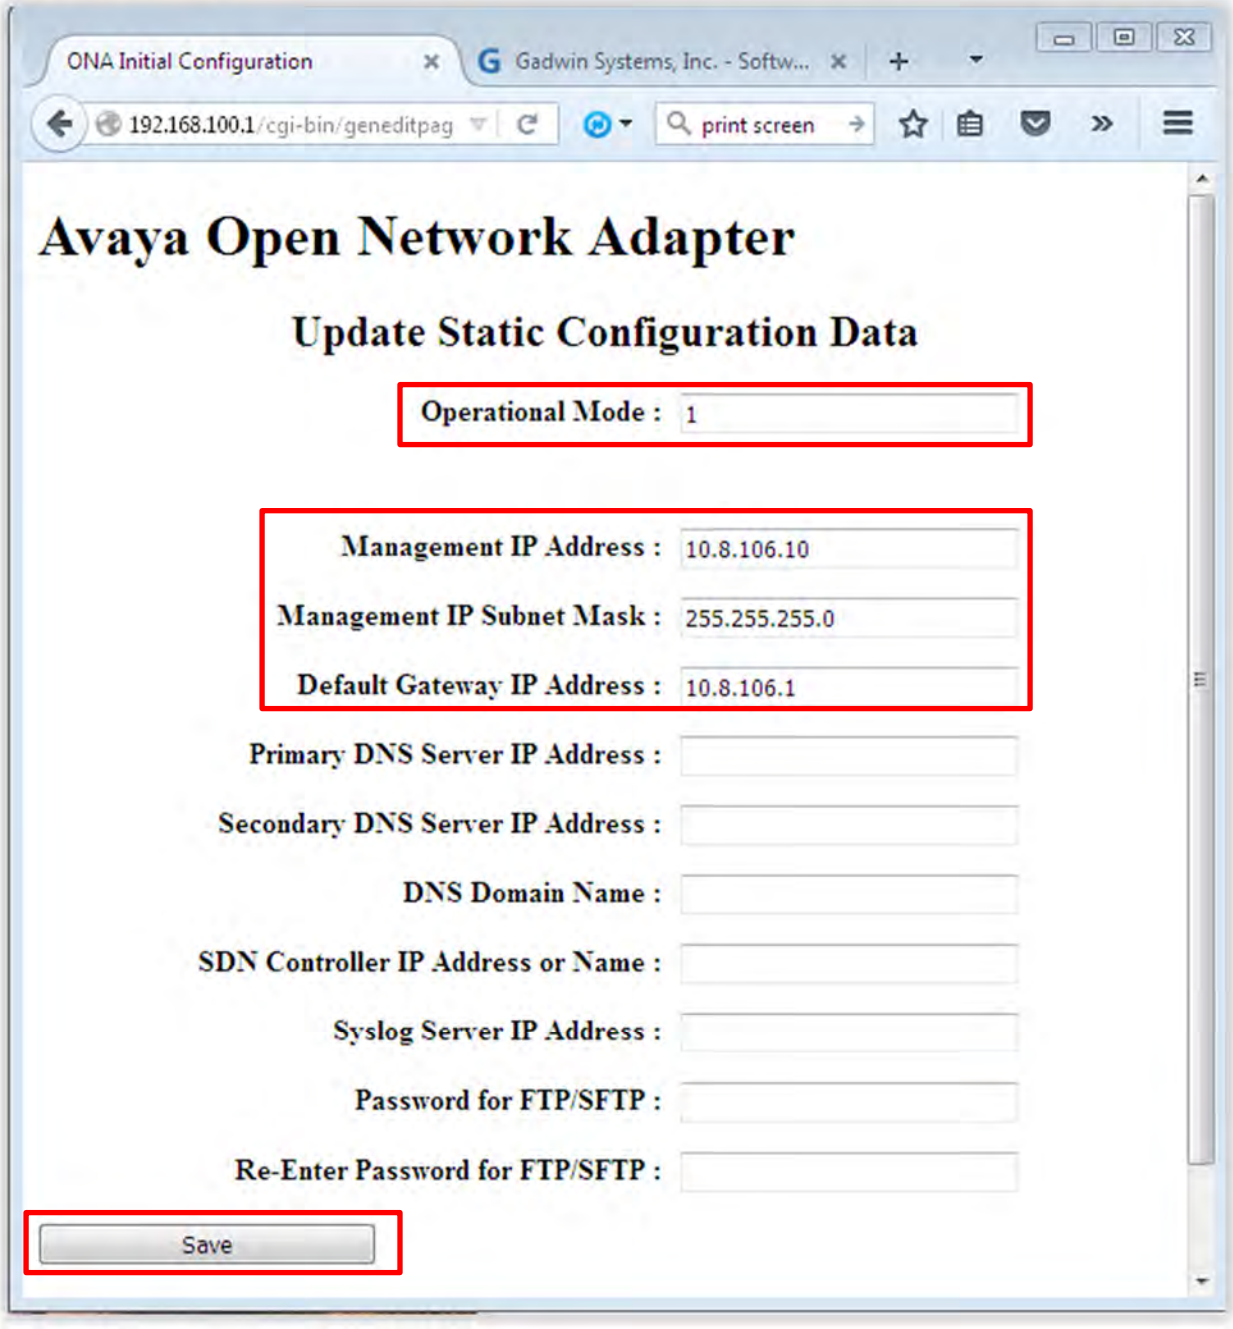

20.1.27 Fabric Extend over Routed Infrastructure using VRF to interconnect to routed network

Please begin by enable IP Shortcuts on all switches – see section above titled IP Shortcuts. On the VSP 4000 switches (4001 and 4002), we will create loopback 10 to be used for both the OSPF router-id and for the tunnel source IP address. On the VSP 8000 switch (8201), we will simple just use a brouter interface and use this IP address as the tunnel source. Switch 8201 in this example will use the loopback 10 interface created on the 4001 and 4002, as the tunnel termination point while nodes 4001 and 4002 will use the brouter interface on 8201 as the tunnel termination point.

Configure Core Networking - VRF

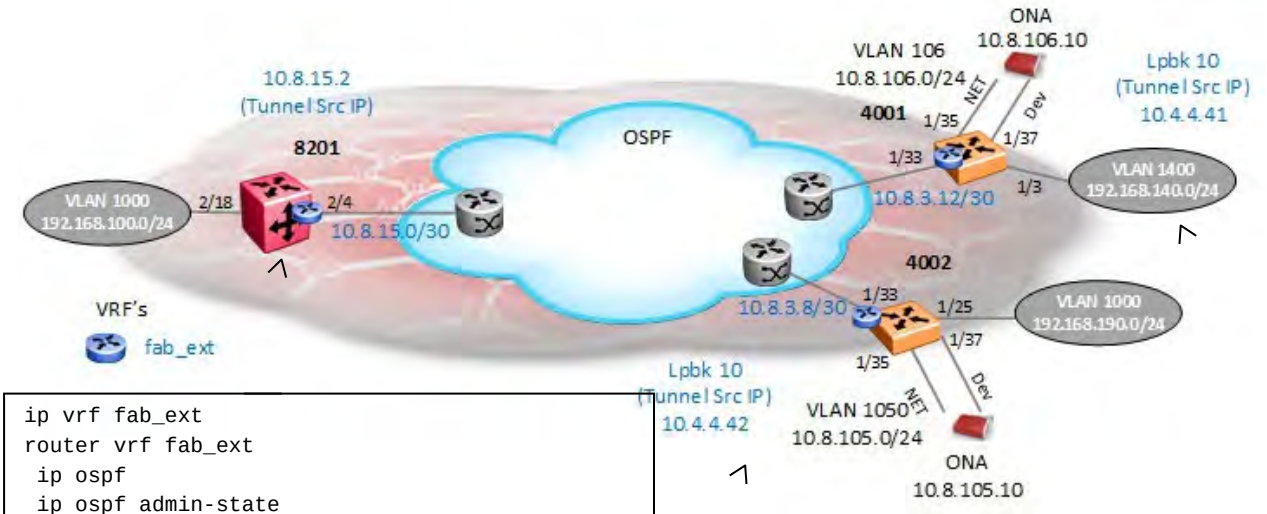


```
ip vrf fab_ext
router vrf fab_ext
 ip ospf
 ip ospf admin-state
exit
interface GigabitEthernet 2/4
  no shutdown
  vrf fab_ext
  brouter vlan 15 subnet 10.8.15.2/30
  no spanning-tree mstp force-port-state enable
  ip ospf enable
exit
```

```
ip vrf fab_ext
router vrf fab_ext
 ip ospf
 ip ospf admin-state
exit
vlan create 15 name "core_vlan15" type port-mstprstp 0
vlan members 15 1/33
interface Vlan 15
  vrf fab_ext
  ip address 10.8.3.14 255.255.255.252
  ip ospf enable
exit
interface GigabitEthernet 1/33
  no shutdown
  no spanning-tree mstp force-port-state enable
exit
interface loopback 10
  ip address 10 10.4.4.41/32 vrf fab_ext
  ip ospf 10 vrf fab_ext
exit
router vrf fab_ext
  ip ospf router-id 10.4.4.41
  ip ospf admin-state
exit
```

```
ip vrf fab_ext
router vrf fab_ext
 ip ospf
 ip ospf admin-state
exit
interface GigabitEthernet 1/33
  no shutdown
  vrf fab_ext
  brouter vlan 4 subnet 10.8.3.10/30
  no spanning-tree mstp force-port-state enable
  ip ospf enable
exit
interface loopback 10
  ip address 10 10.4.4.42/32 vrf fab_ext
  ip ospf 10 vrf fab_ext
exit
router vrf fab_ext
  ip ospf router-id 10.4.4.42
  ip ospf admin-state
exit
```

Verify Core Operations

Before beginning the Fabric Extend configuration, please ensure that nodes 8201, 4001, and 4002 are learning all the various remote networks via OSPF.

```
8201:1#show ip route -s 10.8.0.0/16 vrf fab_ext
================================================================================
                                  IP Route - VRF fab_ext
================================================================================
                                          NH                INTER
DST             MASK            NEXT      VRF/ISID   COST   FACE  PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.8.1.8        255.255.255.255 10.8.15.1   fab_ext    21     2/4   OSPF 0   IB   20
|
10.8.3.8        255.255.255.252 10.8.15.1   fab_ext    21     2/4   OSPF 0   IB   20
10.8.3.12       255.255.255.252 10.8.15.1   fab_ext    11     2/4   OSPF 0   IB   20
|
10.8.15.0       255.255.255.252 10.8.15.2    -         1      2/4   LOC  0   DB   0
10.8.100.0      255.255.255.0   10.8.15.1   fab_ext    11     2/4   OSPF 0   IB   20

8201:1#show ip route -s 10.4.4.0/24 vrf fab_ext
================================================================================
                                  IP Route - VRF fab_ext
================================================================================
                                        NH                INTER
DST             MASK            NEXT    VRF/ISID   COST   FACE   PROT AGE TYPE PRF
--------------------------------------------------------------------------------
10.4.4.41       255.255.255.255 10.8.15.1 fab_ext    21     2/4    OSPF 0   IB   20
10.4.4.42       255.255.255.255 10.8.15.1 fab_ext    31     2/4    OSPF 0   IB   20
```

Fabric Extend Configuration including ONA configuration for VSP 4000

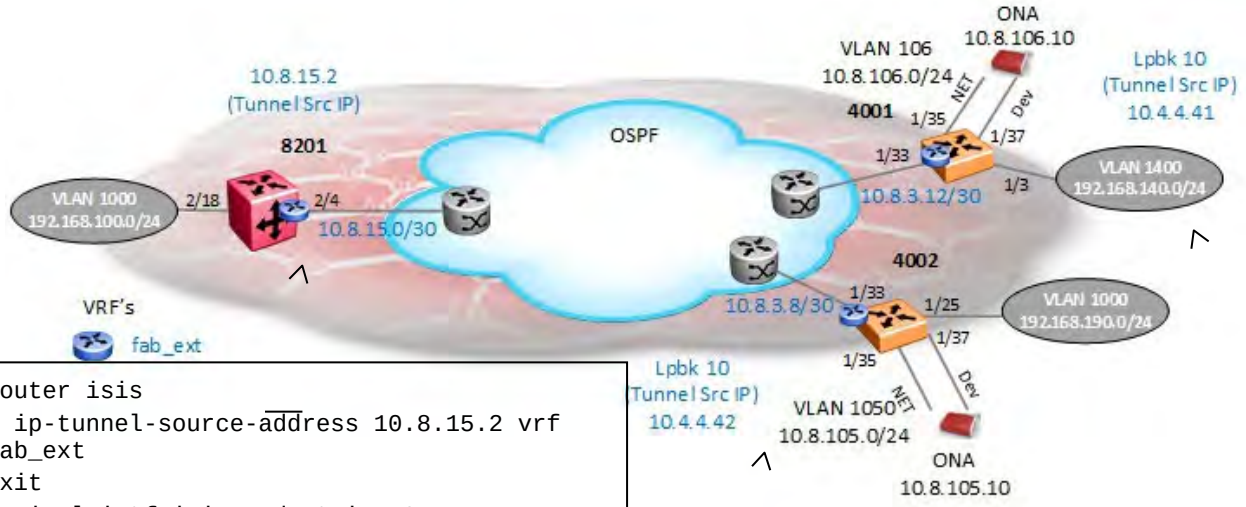


```
router isis
  ip-tunnel-source-address 10.8.15.2 vrf
fab_ext
exit
logical-intf isis 1 dest-ip 10.4.4.41 name
"tunnel_to_4001"
  isis
  isis spbm 1
  isis spbm 1 l1-metric 10000
  isis enable
exit
logical-intf isis 2 dest-ip 10.4.4.42 name
"tunnel_to_4002"
  isis
  isis spbm 1
  isis spbm 1 l1-metric 10000
  isis enable
exit
```

```
vlan create 106 name "ona_vlan106" type
port-mstprstp 0
vlan members 106 1/35
interface Vlan 106
  ip address 10.8.106.1 255.255.255.0
  ip dhcp-relay
exit
interface GigabitEthernet 1/35
  no shutdown
exit
interface GigabitEthernet 1/37
  no shutdown
exit
router isis
  ip-tunnel-source-address 10.4.4.41
port 1/37 vrf fab_ext
exit
logical-intf isis 1 dest-ip 10.8.15.2
name "tunnel_to_8201"
  isis
  isis spbm 1
  isis spbm 1 l1-metric 10000
  isis enable
exit
```

```
vlan create 1050 name "ona_vlan1050" type port-
mstprstp 0
vlan members 1050 1/35
interface Vlan 1050
  ip address 10.8.105.1 255.255.255.0
exit
interface GigabitEthernet 1/35
 no shutdown
exit
interface GigabitEthernet 1/37
  no shutdown
exit
router isis
  ip-tunnel-source-address 10.4.4.42 port 1/37 vrf
fab_ext
exit
```

```
logical-intf isis 2 dest-ip 10.8.15.2 name
"tunnel_to_8201"
  isis
  isis spbm 1
  isis spbm 1 l1-metric 10000
  isis enable
exit
```

Verify ISIS Interfaces

```
4002:1#show isis inter

================================================================================
                                ISIS Interfaces
================================================================================
IFIDX             TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
--------------------------------------------------------------------------------
tunnel_to_8201    pt-pt   Level 1   UP        UP         1      1       10000

--------------------------------------------------------------------------------
 1 out of 1 Total Num of ISIS interfaces
--------------------------------------------------------------------------------

4002:1#show isis logical-interface

================================================================================================================
                                            ISIS Logical Interfaces
================================================================================================================
IFIDX    NAME         ENCAP  L2_INFO                TUNNEL     L3_TUNNEL_NEXT_HOP_INFO
                      TYPE   PORT/MLT  VIDS(PRIMARY) DEST-IP    PORT/MLT   VLAN      VRF
----------------------------------------------------------------------------------------------------------------
2     tunnel_to_8201  IP      --        --           10.8.15.2  Port1/33   4         fab_ext

----------------------------------------------------------------------------------------------------------------
 1 out of 1 Total Num of Logical ISIS interfaces
----------------------------------------------------------------------------------------------------------------

4002:1#show isis adj

================================================================================
                                ISIS Adjacencies
================================================================================
INTERFACE         L STATE        UPTIME PRI HOLDTIME SYSID              HOST-NAME
--------------------------------------------------------------------------------
tunnel_to_8201    1 UP         00:12:15 127       20 b0ad.aa47.0884    8201

--------------------------------------------------------------------------------
 1 out of 1 interfaces have formed an adjacency
--------------------------------------------------------------------------------
```

Verify ONA Operations

```
4002:1#show khi fe-ona status

=========================================================
               ONA STATUS
=========================================================
ONA Device Status : UP
Running Release Name : VEGA1101.1.0.0.0int007
Last Image Upgrade Status : UPGRADE_SUCCESS
Last Image File Used For Upgrade : VEGA1101.1.0.0.0int007.tgz
-------------------------------------------------------------------------------

4002:1#show khi fe-ona detail

=========================================================
               ONA RUNTIME INFORMATION
=========================================================
ONA Port Number : 1/37
ONA Management Address : 10.8.105.10
Tunnel Source IP Address : 10.4.4.42
ONA LLDP Port Status : Enabled
ONA Device Port Status : UP
ONA Device Status : UP
MTU : 1950
ONA Network Port Number : 1/35
ONA Mac(ARP) Address : 10:cd:ae:69:b9:00
ONA Source VlanId : 1050
ONA Source Vlan IP : 10.8.105.1
ONA Gateway IP: 10.8.105.1
ONA Management IP Mask : 255.255.255.0
ONA Bootmode : 1
ONA Uptime : 93238 seconds
pbit-to-dscp-map p0=0 p1=0 p2=10 p3=18 p4=26 p5=34 p6=46 p7=46
-------------------------------------------------------------------------------------------
```

Please note the ONA will not come up until the *ip-tunnel-source-address* has been added with the corresponding device port and VRF name.

20.1.28 Fabric Extend over Routed Infrastructure using GRT to interconnect to routed network

Please note that the preferred method is using a VRF for the core routing as shown in the previous example. If you use the GRT method as shown below, please ensure you add a route policy to prevent all the core IP interfaces, local and loopback, from being advertised into ISIS. Please begin by enable IP Shortcuts on all switches – see section above titled IP Shortcuts.

Configure Core Networking - GRT

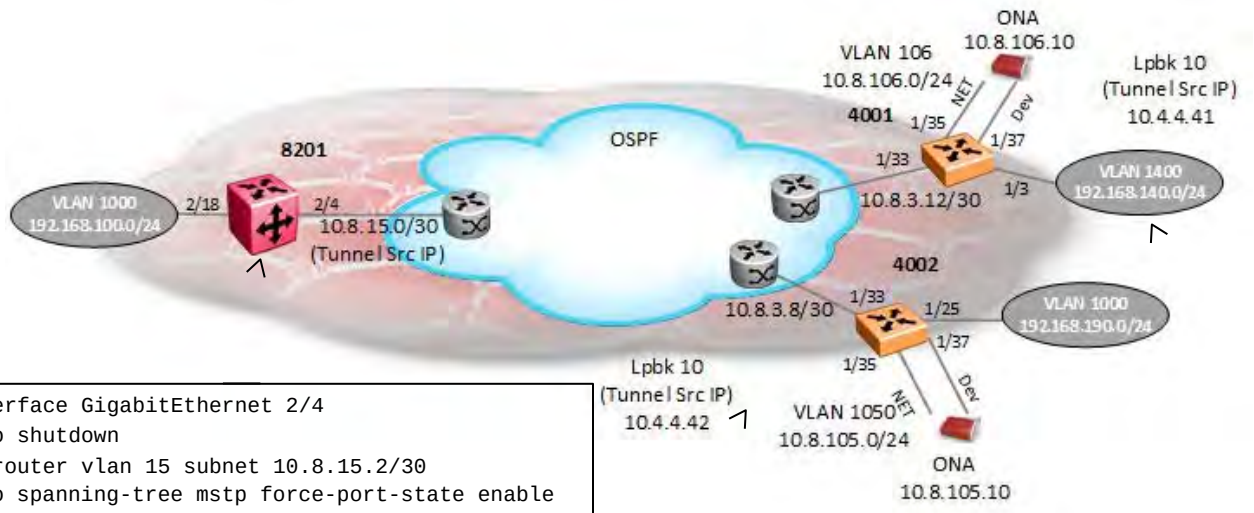


```
interface GigabitEthernet 2/4
  no shutdown
  brouter vlan 15 subnet 10.8.15.2/30
  no spanning-tree mstp force-port-state enable
  ip ospf enable
exit
router ospf enable
```

```
vlan create 15 name "core_vlan15" type port-mstprstp 0
vlan members 15 1/33
interface Vlan 15
  ip address 10.8.3.14 255.255.255.252
  ip ospf enable
exit
interface GigabitEthernet 1/33
  no shutdown
  no spanning-tree mstp force-port-state enable
exit
interface loopback 10
  ip address 10 10.4.4.41/32
  ip ospf 10
exit
router ospf enable
router ospf
  router-id 10.4.4.41
exit
```

```
interface GigabitEthernet 1/33
  no shutdown
  brouter vlan 4 subnet 10.8.3.10/30
  no spanning-tree mstp force-port-state enable
  ip ospf enable
exit
interface loopback 10
  ip address 10 10.4.4.42/32
  ip ospf 10
exit
router ospf enable
router ospf
  router-id 10.4.4.42
exit
```

Core GRT – Route Map/Route Policy

If using the GRT method for Fabric Extend, you have to add a route-map or route policy to prevent the local OSPF networks from also being advertised by ISIS. For this example, we will route a route-map for nodes 8201, 4001, and 4002.

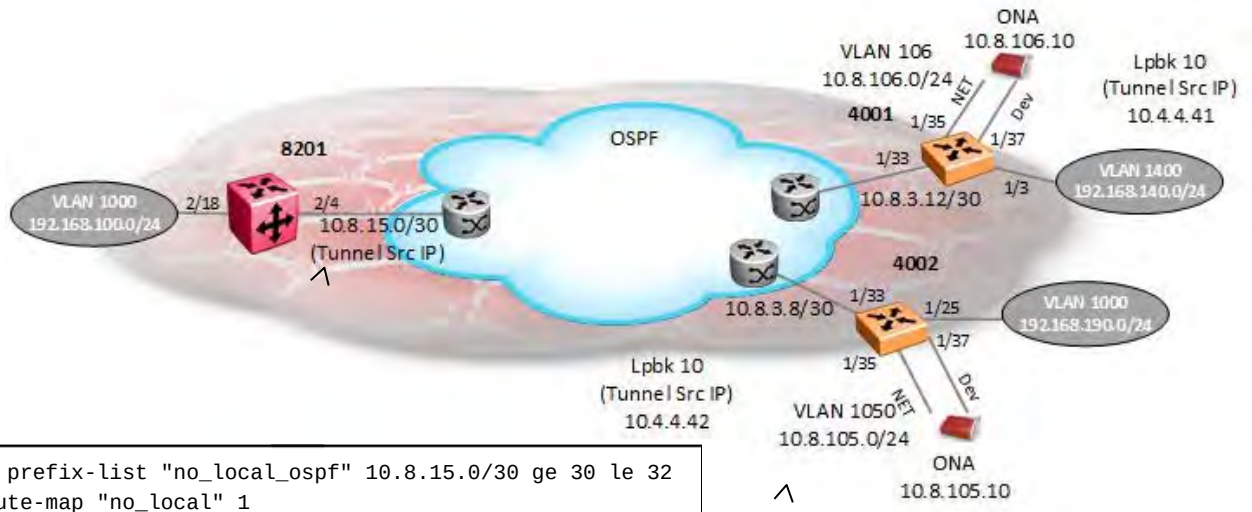


```
ip prefix-list "no_local_ospf" 10.8.15.0/30 ge 30 le 32
route-map "no_local" 1
  no permit
  enable
  match network "no_local_ospf"
  exit
route-map "no_loc" 65535
  permit
  enable
exit
router isis
    redistribute direct route-map “no_local”
exit
isis apply redistribute direct
```

```
ip prefix-list "no_local_ospf" 10.4.4.41/32
ip prefix-list "no_local_ospf" 10.8.3.12/30
ge 30 le 32
route-map "no_local" 1
  no permit
  enable
  match network "no_local_ospf"
exit
route-map "no_local" 65535
  permit
  enable
exit
router isis
  redistribute direct route-map "no_local"
exit
isis apply redistribute direct
```

```
ip prefix-list "no_local_ospf" 10.4.4.42/32
ip prefix-list "no_local_ospf" 10.8.3.8/30 ge 30 le 32
route-map "no_local" 1
  no permit
  enable
  match network "no_local_ospf"
exit
route-map "no_local" 65535
  permit
  enable
exit
router isis
  redistribute direct route-map "no_local"
exit
isis apply redistribute direct
```

Fabric Extend Configuration

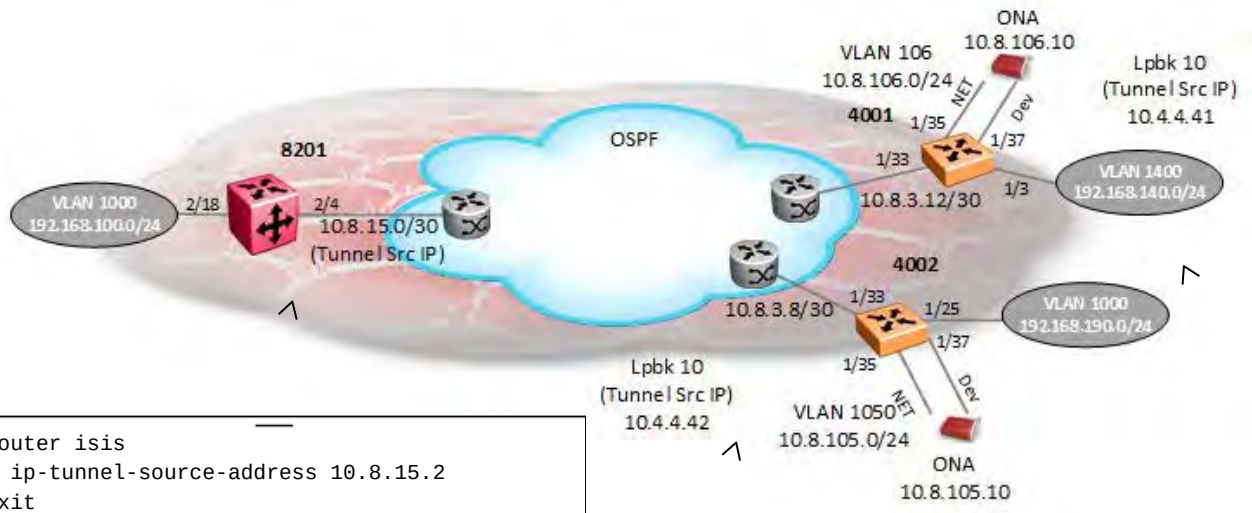


```
router isis
  ip-tunnel-source-address 10.8.15.2
exit
logical-intf isis 1 dest-ip 10.4.4.41 name
"tunnel_to_4001"
  isis
  isis spbm 1
  isis spbm 1 l1-metric 10000
  isis enable
exit
logical-intf isis 2 dest-ip 10.4.4.42 name
"tunnel_to_4002"
  isis
  isis spbm 1
  isis spbm 1 l1-metric 10000
  isis enable
exit
```

```
vlan create 1050 name "ona_vlan1050" type port-mstprstp 0
vlan members 1050 1/35
interface Vlan 1050
  ip address 10.8.105.1 255.255.255.0
exit
interface GigabitEthernet 1/35
 no shutdown
exit
interface GigabitEthernet 1/37
  no shutdown
exit
router isis
  ip-tunnel-source-address 10.4.4.42 port 1/37
exit
logical-intf isis 2 dest-ip 10.8.15.2 name
"tunnel_to_8201"
  isis
  isis spbm 1
  isis spbm 1 l1-metric 10000
  isis enable
exit
```

```
vlan create 106 name "ona_vlan106" type port-
mstprstp 0
vlan members 106 1/35
interface Vlan 106
  ip address 10.8.106.1 255.255.255.0
  ip dhcp-relay
exit
interface GigabitEthernet 1/35
  no shutdown
exit
interface GigabitEthernet 1/37
  no shutdown
exit
router isis
  ip-tunnel-source-address 10.4.4.41 port
1/37
exit
logical-intf isis 1 dest-ip 10.8.15.2 name
"tunnel_to_8201"
  isis
  isis spbm 1
  isis spbm 1 l1-metric 10000
  isis enable
exit
```

Verify the route-map is working

Verify that the OSPF networks are only learned by OSPF and not ISIS. In this example only the ONA networks (10.8.105.0/24 and 10.8.107.0/24) should be learned by ISIS.

```
8201:1#show ip route -s 10.8.0.0/16
================================================================================
                                  IP Route - GlobalRouter
================================================================================
                                                    NH                    INTER
DST             MASK            NEXT            VRF/ISID         COST    FACE     PROT AGE TYPE PRF
--------------------------------------------------------------------------------------------------
10.8.1.8        255.255.255.255 10.8.15.1       GlobalRouter     21      2/4      OSPF 0   IB   20
10.8.1.9        255.255.255.255 10.8.15.1       GlobalRouter     31      2/4      OSPF 0   IB   20
10.8.2.0        255.255.255.252 10.8.15.1       GlobalRouter     11      2/4      OSPF 0   IB   20
10.8.3.0        255.255.255.252 10.8.15.1       GlobalRouter     11      2/4      OSPF 0   IB   20
10.8.3.4        255.255.255.252 10.8.15.1       GlobalRouter     21      2/4      OSPF 0   IB   20
10.8.3.8        255.255.255.252 10.8.15.1       GlobalRouter     21      2/4      OSPF 0   IB   20
10.8.3.16       255.255.255.252 10.8.15.1       GlobalRouter     31      2/4      OSPF 0   IB   20
10.8.8.0        255.255.255.0   10.8.15.1       GlobalRouter     21      2/4      OSPF 0   IB   20
10.8.15.0       255.255.255.252 10.8.15.2       -                1       2/4      LOC  0   DB   0
10.8.21.0       255.255.255.0   10.8.15.1       GlobalRouter     31      2/4      OSPF 0   IB   20
10.8.100.0      255.255.255.0   10.8.15.1       GlobalRouter     11      2/4      OSPF 0   IB   20
10.8.105.0      255.255.255.0   4002            GlobalRouter     10000   3051     ISIS 0   IBS  7
10.8.106.0      255.255.255.0   4001            GlobalRouter     10000   3051     ISIS 0   IBS  7

8201:1#show ip route -s 10.4.4.0/24
================================================================================
                                  IP Route - GlobalRouter
================================================================================
                                                    NH                    INTER
DST             MASK            NEXT            VRF/ISID         COST    FACE     PROT AGE TYPE PRF
--------------------------------------------------------------------------------------------------
10.4.4.1        255.255.255.255 4001            GlobalRouter     10000   3051     ISIS 0   IBS  7
10.4.4.2        255.255.255.255 4002            GlobalRouter     10000   3051     ISIS 0   IBS  7
10.4.4.41       255.255.255.255 10.8.15.1       GlobalRouter     31      2/4      OSPF 0   IB   20
10.4.4.42       255.255.255.255 10.8.15.1       GlobalRouter     41      2/4      OSPF 0   IB   20
```

Verify ISIS Interfaces

```
4002:1#show isis inter

================================================================================
                                ISIS Interfaces
================================================================================
IFIDX             TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
--------------------------------------------------------------------------------
tunnel_to_8201    pt-pt   Level 1   UP        UP         1      1       10000

--------------------------------------------------------------------------------
 1 out of 1 Total Num of ISIS interfaces
--------------------------------------------------------------------------------

4002:1#show isis logical-interface

================================================================================================================
                                            ISIS Logical Interfaces
================================================================================================================
IFIDX    NAME         ENCAP  L2_INFO                TUNNEL     L3_TUNNEL_NEXT_HOP_INFO
                      TYPE   PORT/MLT  VIDS(PRIMARY) DEST-IP    PORT/MLT   VLAN      VRF
----------------------------------------------------------------------------------------------------------------
2     tunnel_to_8201  IP      --        --           10.8.15.2  Port1/33   4         fab_ext

----------------------------------------------------------------------------------------------------------------
 1 out of 1 Total Num of Logical ISIS interfaces
----------------------------------------------------------------------------------------------------------------

4002:1#show isis adj

================================================================================
                                ISIS Adjacencies
================================================================================
INTERFACE         L STATE        UPTIME PRI HOLDTIME SYSID              HOST-NAME
--------------------------------------------------------------------------------
tunnel_to_8201    1 UP         00:12:15 127       20 b0ad.aa47.0884    8201

--------------------------------------------------------------------------------
 1 out of 1 interfaces have formed an adjacency
--------------------------------------------------------------------------------
```

Verify ONA Operations

```
4002:1#show khi fe-ona status

=========================================================
               ONA STATUS
=========================================================
ONA Device Status : UP
Running Release Name : VEGA1101.1.0.0.0int007
Last Image Upgrade Status : UPGRADE_SUCCESS
Last Image File Used For Upgrade : VEGA1101.1.0.0.0int007.tgz
-------------------------------------------------------------------------------

4002:1#show khi fe-ona detail

=========================================================
               ONA RUNTIME INFORMATION
=========================================================
ONA Port Number : 1/37
ONA Management Address : 10.8.105.10
Tunnel Source IP Address : 10.4.4.42
ONA LLDP Port Status : Enabled
ONA Device Port Status : UP
ONA Device Status : UP
MTU : 1950
ONA Network Port Number : 1/35
ONA Mac(ARP) Address : 10:cd:ae:69:b9:00
ONA Source VlanId : 1050
ONA Source Vlan IP : 10.8.105.1
ONA Gateway IP: 10.8.105.1
ONA Management IP Mask : 255.255.255.0
ONA Bootmode : 1
ONA Uptime : 93238 seconds
pbit-to-dscp-map p0=0 p1=0 p2=10 p3=18 p4=26 p5=34 p6=46 p7=46
-------------------------------------------------------------------------------
```

Please note the ONA will not come up until the *ip-tunnel-source-address* has been added with the corresponding device port.

20.1.29 Fabric Extend over E-LAN/VPLS (L2) network using Layer 3 over Layer 2 tunneling using VSP 4000

Fabric Extend can be transported over a layer 2 core using VXLAN IP tunneling. For this example, we will assume there is no VLAN tagging to the core L2 network. For this simple setup, we can simply use a brouter port to the core and use the brouter IP address as the IP tunnel source address. If you require VLAN tagging to the core, please see the next section. Before you start the configuration below, enable IP Shortcuts as per the example above.

Configuration

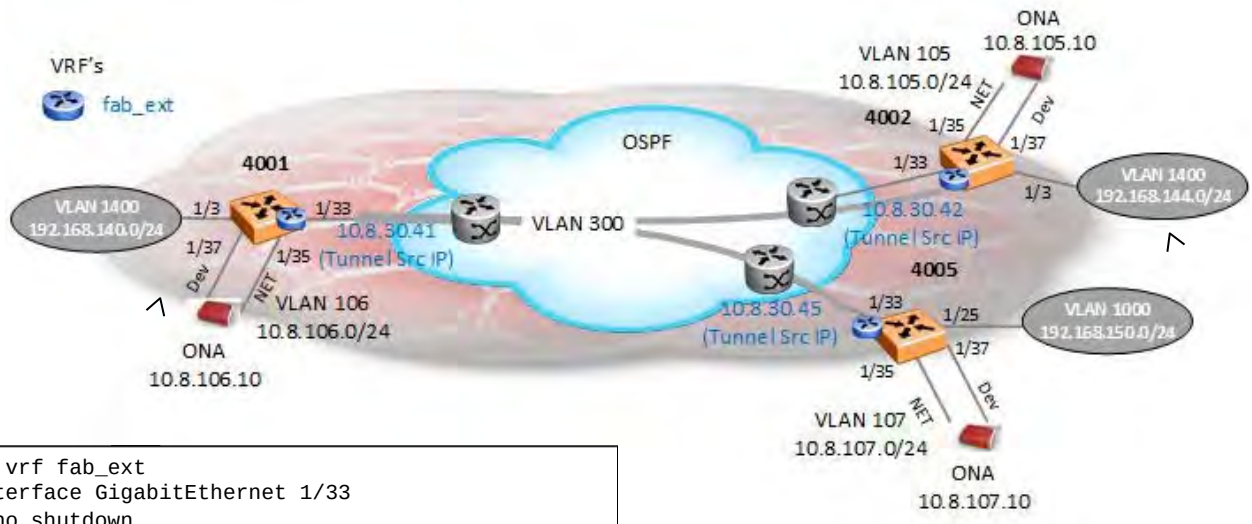


```
ip vrf fab_ext
interface GigabitEthernet 1/33
  no shutdown
  no spanning-tree mstp force-port-state enable
  vrf fab_ext
  brouter vlan 300 subnet 10.8.30.41/24
exit
vlan create 106 name "ona_vlan106" type port-
mstprstp 0
vlan members 106 1/35
interface Vlan 106
  ip address 10.8.106.1 255.255.255.0
exit
interface GigabitEthernet 1/37
  no shutdown
exit
router isis
  ip-tunnel-source-address 10.8.30.41 port 1/37
vrf fab_ext
exit
logical-intf isis 1 dest-ip 10.8.30.42 name
"tunnel_to_4002"show
  isis
  isis spbm 1
  isis enable
exit
logical-intf isis 1 dest-ip 10.8.30.45 name
"tunnel_to_4005"
  isis
  isis spbm 1
  isis enable
exit
```

```
ip vrf fab_ext
interface GigabitEthernet 1/33
  no shutdown
  no spanning-tree mstp force-port-state enable
  vrf fab_ext
  brouter vlan 300 subnet 10.8.30.42/24
exit
vlan create 105 name "ona_vlan105" type port-
mstprstp 0
vlan members 105 1/35
interface Vlan 105
  ip address 10.8.105.1 255.255.255.0
exit
interface GigabitEthernet 1/37
  no shutdown
exit
router isis
  ip-tunnel-source-address 10.8.30.42 port 1/37
vrf fab_ext
exit
logical-intf isis 1 dest-ip 10.8.30.41 name
"tunnel_to_4001"
  isis
  isis spbm 1
  isis enable
exit
```

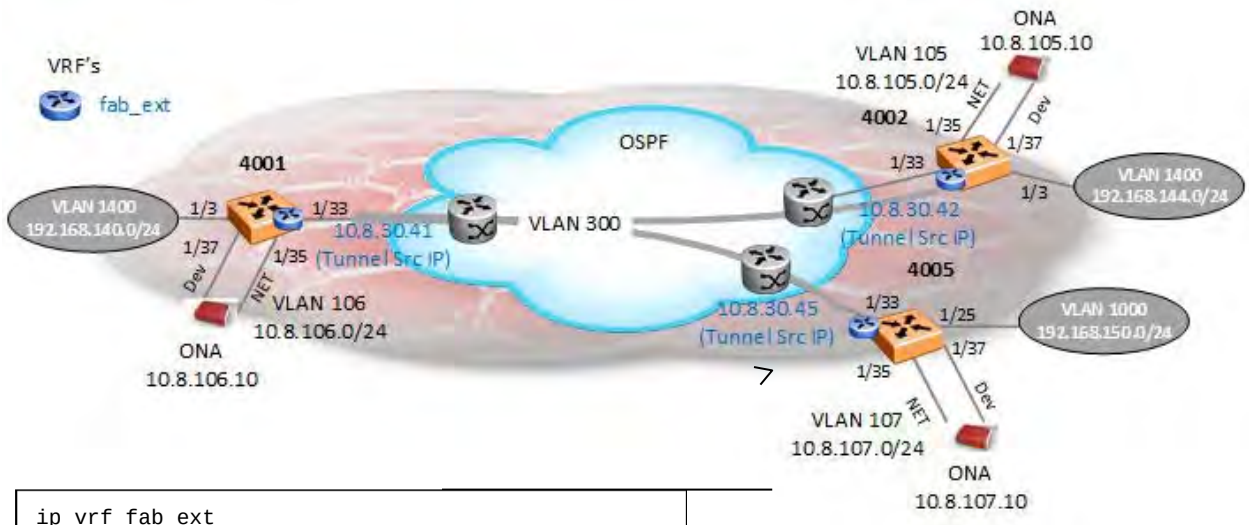


```
ip vrf fab_ext
interface GigabitEthernet 1/33
  no shutdown
  no spanning-tree mstp force-port-state enable
  vrf fab_ext
  brouter vlan 300 subnet 10.8.30.45/24
exit
vlan create 107 name "ona_vlan106" type port-
mstprstp 0
vlan members 107 1/35
interface Vlan 107
  ip address 10.8.107.1 255.255.255.0
exit
interface GigabitEthernet 1/37
  no shutdown
exit
router isis
  ip-tunnel-source-address 10.8.30.45 port 1/37
vrf fab_ext
exit
logical-intf isis 1 dest-ip 10.8.30.41 name
"tunnel_to_4002"
  isis
  isis spbm 1
  isis enable
exit
```

Verify VRF Operations

```
4001:1#show ip interface vrf fab_ext

================================================================================
                          IP Interface - VRF fab_ext
================================================================================
INTERFACE    IP             NET            BCASTADDR  REASM    VLAN  BROUTER
             ADDRESS        MASK           FORMAT     MAXSIZE  ID    PORT
 ------------------------------------------------------------------------------
Port1/33     10.8.30.41     255.255.255.0  ones       1500     300   true


All 1 out of 1 Total Num of IP interfaces displayed

4001:1#ping 10.8.30.42 vrf fab_ext
10.8.30.42 is alive

4001:1#ping 10.8.30.45 vrf fab_ext
10.8.30.45 is alive

4001:1#show ip arp vrf fab_ext

================================================================================================
                                  IP Arp - VRF fab_ext
================================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN  PORT                TYPE    TTL(10 Sec)
TUNNEL
 ------------------------------------------------------------------------------------------
10.8.30.1       b0:ad:aa:47:09:02  300   1/33                DYNAMIC 827
10.8.30.41      d4:ea:0e:10:e4:81  300   1/33                LOCAL   2160
10.8.30.42      a0:12:90:d3:ec:83  300   1/33                DYNAMIC 2154
10.8.30.45      d4:ea:0e:e0:98:80  300   1/33                DYNAMIC 2154
10.8.30.255     ff:ff:ff:ff:ff:ff  300   1/33                LOCAL   2160
```

Verify ONA Operations

```
4001:1#show khi fe-ona status

=========================================================
               ONA STATUS
=========================================================
ONA Device Status : UP
Running Release Name : VEGA1101_beta.1.0.0.0int014
Last Image Upgrade Status : UPGRADE_SUCCESS
Last Image File Used For Upgrade : VEGA1101_beta.1.0.0.0int014.tgz
-------------------------------------------------------------------------------

4001:1#show khi fe-ona detail

=========================================================
               ONA RUNTIME INFORMATION
=========================================================
ONA Port Number : 1/37
ONA Management Address : 10.8.106.10
Tunnel Source IP Address : 10.8.30.41
ONA LLDP Port Status : Enabled
ONA Device Port Status : UP
ONA Device Status : UP
MTU : 1950
ONA Network Port Number : 1/35
ONA Mac(ARP) Address : 10:cd:ae:69:b8:30
ONA Source VlanId : 106
ONA Source Vlan IP : 10.8.106.1
ONA Gateway IP: 10.8.106.1
ONA Management IP Mask : 255.255.255.0
ONA Bootmode : 1
ONA Uptime : 13365 seconds
pbit-to-dscp-map p0=0 p1=0 p2=10 p3=18 p4=26 p5=34 p6=46 p7=46
-------------------------------------------------------------------------------
```

Verify ISIS Interfaces

```
4001:1#show isis interface

=============================================================================
                              ISIS Interfaces
=============================================================================
IFIDX             TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
 ----------------------------------------------------------------------------
tunnel_to_4002    pt-pt   Level 1   UP        UP         1      1       20000
tunnel_to_4005    pt-pt   Level 1   UP        UP         1      1       20000
```

```
4001:1#show isis adjacencies

================================================================================
                                ISIS Adjacencies
================================================================================
INTERFACE         L STATE        UPTIME PRI  HOLDTIME SYSID             HOST-NAME
 ------------------------------------------------------------------------------
tunnel_to_4002    1 UP          03:42:50 127       23 a012.90d3.ec65    4002
tunnel_to_4005    1 UP          03:42:22 127       20 d4ea.0ee0.9865    4005

4001:1#show isis logical-interface

================================================================================================================
                                            ISIS Logical Interfaces
================================================================================================================
IFIDX    NAME             ENCAP   L2_INFO                  TUNNEL      l3_TUNNEL_NEXT_HOP_INFO
                          TYPE    PORT/MLT   VIDS(PRIMARY) DEST-IP     PORT/MLT   VLAN   VRF
----------------------------------------------------------------------------------------------------------------
1        tunnel_to_8201   IP      --         --            10.8.30.1   Port1/33   300    fab_ext
2        tunnel_to_4002   IP      --         --            10.8.30.42  Port1/33   300    fab_ext
3        tunnel_to_4005   IP      --         --            10.8.30.45  Port1/33   300    fab_ext
```

Verify Route Table

```
4001:1# show ip route
================================================================================================================
                                        IP Route - GlobalRouter
================================================================================================================
                                                NH                      INTER
DST             MASK            NEXT            VRF/ISID        COST    FACE     PROT AGE TYPE
PRF
----------------------------------------------------------------------------------------------------------------
10.4.4.1        255.255.255.255 10.4.4.1        -               1       0        LOC  0   DB   0
10.4.4.2        255.255.255.255 4002            GlobalRouter    20000   4051     ISIS 0   IBS  7
10.4.4.5        255.255.255.255 4005            GlobalRouter    20000   4051     ISIS 0   IBS  7
10.8.105.0      255.255.255.0   4002            GlobalRouter    20000   4051     ISIS 0   IBS  7
10.8.106.0      255.255.255.0   10.8.106.1      -               1       106      LOC  0   DB   0
10.8.107.0      255.255.255.0   4005            GlobalRouter    20000   4051     ISIS 0   IBS  7
192.168.100.0   255.255.255.0   8201            GlobalRouter    20000   4051     ISIS 0   IBS  7
192.168.140.0   255.255.255.0   192.168.140.1   -               1       1400     LOC  0   DB   0
192.168.144.0   255.255.255.0   4002            GlobalRouter    20000   4051     ISIS 0   IBS  7
192.168.150.0   255.255.255.0   4005            GlobalRouter    20000   4051     ISIS 0   IBS  7
```

20.1.30 Fabric Extend over E-LAN/VPLS (L2) network using Layer 3 over Layer 2 tunneling with VSP8000 or VSP7200

When using the VSP 8000 or VSP 7200, please note that they only support a single next hop IP address for all tunnels going through a single port. The ONA on the VSP 4000 does not have this restriction and as shown in the previous example, it can it used to provide tunnels to many remote sites over a flat layer 2 core. However, we can add a loopback interface to a second VRF and use this interface as the tunnel source IP address to overcome this limitation on the VSP 8000 or VSP 7200. To make this work, we need to create a hairpin between the two VRF's via an addition interface added to both VRF's with the same IP subnet in order for the core network to get to the tunnel network.

For this example, we will use a VSP 8000 at the main site with two VRFs named fab_ext and fab_lpbk. VRF fab_ext will be used in the core while vrf fab_lpbk will be used with a loopback interface for the tunnel source.

Configure Core Networking

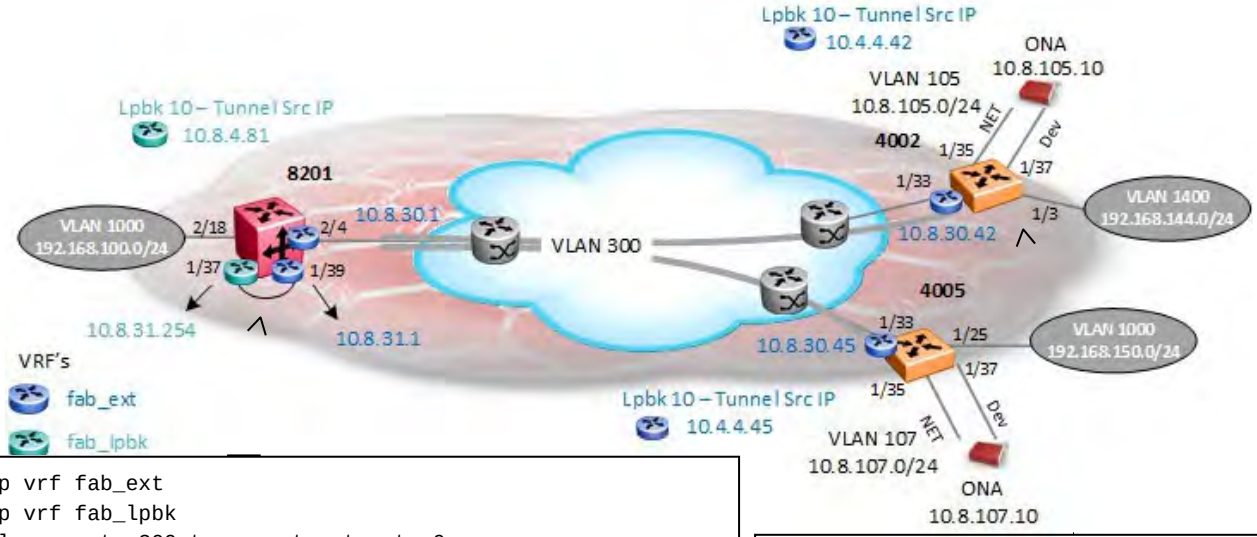


```
ip vrf fab_ext
ip vrf fab_lpbk
vlan create 300 type port-mstprstp 0
vlan members 300 2/4
interface Vlan 300
  vrf fab_ext
  ip address 10.8.30.1 255.255.255.0
exit
interface GigabitEthernet 1/37
  no shutdown
  vrf fab_lpbk
  brouter port 1/37 vlan 302 subnet 10.8.31.254/24
  no spanning-tree mstp force-port-state enable
exit
interface GigabitEthernet 1/39
  no shutdown
  vrf fab_ext
  brouter port 1/39 vlan 301 subnet 10.8.31.1/24
  no spanning-tree mstp force-port-state enable
exit
interface loopback 10
  ip address 10 10.8.4.81/32 vrf fab_lpbk
exit
router vrf fab_ext
  ip route 10.4.4.42 255.255.255.255 10.8.30.42 weight 10
  ip route 10.4.4.45 255.255.255.255 10.8.30.45 weight 10
  ip route 10.8.4.81 255.255.255.255 10.8.31.254 weight
10
exit
router vrf fab_lpbk
  ip route 10.4.4.42 255.255.255.255 10.8.31.1 weight 10
  ip route 10.4.4.45 255.255.255.255 10.8.31.1 weight 10
  ip route 10.8.30.0 255.255.255.0 10.8.31.1 weight 10
exit
```

```
ip vrf fab_ext
vlan create 300 type port-mstprstp 0
vlan members 300 1/33
interface Vlan 300
  vrf fab_ext
  ip address 10.8.30.42 255.255.255.0
exit
interface GigabitEthernet 1/33
  no shutdown
  no spanning-tree mstp force-port-state
enable
exit
interface loopback 10
  ip address 10 10.4.4.42/32 vrf fab_ext
exit
router vrf fab_ext
  ip route 10.8.4.81 255.255.255.255
10.8.30.1 weight 10
exit
```

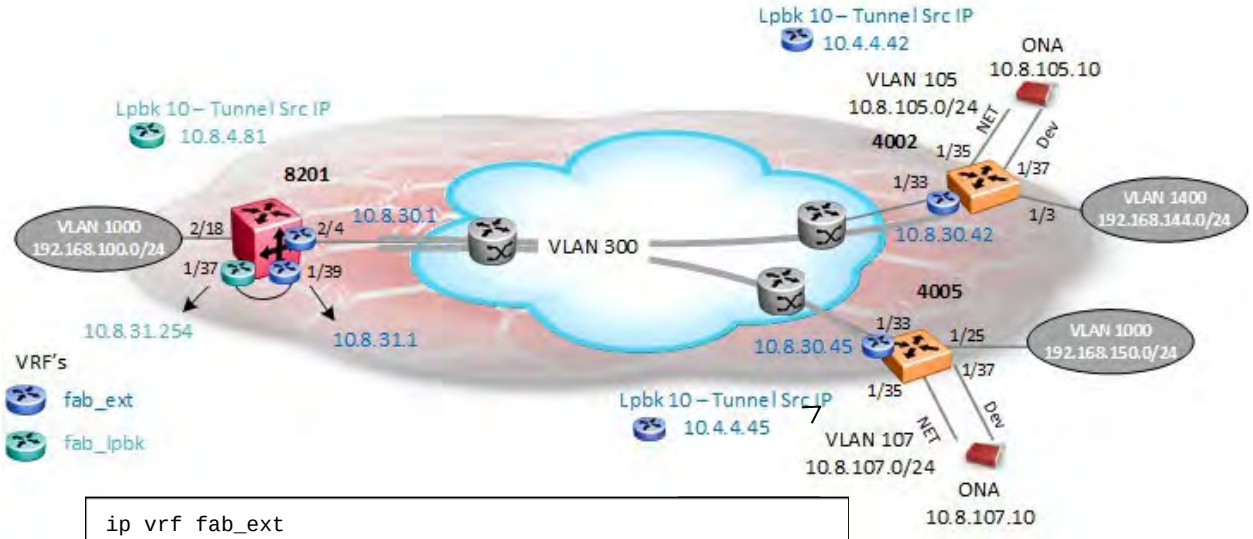


```
ip vrf fab_ext
vlan create 300 type port-mstprstp 0
vlan members 300 1/33
interface Vlan 300
  vrf fab_ext
  ip address 10.8.30.45 255.255.255.0
exit
interface GigabitEthernet 1/33
  no shutdown
  no spanning-tree mstp force-port-state enable
exit
interface loopback 10
  ip address 10 10.4.4.45/32 vrf fab_ext
exit
router vrf fab_ext
  ip route 10.8.4.81 255.255.255.255 10.8.30.1 weight 10
exit
```

In reference to the above drawing, if ports 1/37 and 1/39 are brouter ports where port 1/37 is in GRT and 1/39 is in a VRF, this will not work. If port 1/37 and port 1/39 are brouter ports and if port 1/37 is in one VRF and port 1/39 is in another VRF this solution will work. Also, if port 1/37 and port 1/39 are VLAN ports where the VLAN for port 1/37 is in GRT and the VLAN for port 1/39 is in a VRF, this solution will work.

Verify Routing

Before beginning the Fabric Extend configuration, please ensure that nodes 8201, 4002, and 4005 are able to route to each other, i.e. check the routing to the lpbk addresses used for the Fabric Extend tunnels.

```
4002:1#ping 10.8.30.1 vrf fab_ext
10.8.30.1 is alive

4002:1#ping 10.8.4.81 vrf fab_ext
10.8.4.81 is alive

8201:1#ping 10.8.30.42 vrf fab_ext
10.8.30.42 is alive

8201:1#ping 10.8.30.45 vrf fab_ext
10.8.30.45 is alive

8201:1#ping 10.8.30.42 vrf fab_lpbk source 10.8.4.81
10.8.30.42 is alive

8201:1#ping 10.8.30.45 vrf fab_lpbk source 10.8.4.81
10.8.30.45 is alive

8201:1#show ip arp vrf fab_ext

================================================================================
                                IP Arp - VRF fab_ext
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN  PORT                TYPE    TTL(10 Sec)
TUNNEL
--------------------------------------------------------------------------------
10.8.30.1        b0:ad:aa:47:09:02  300   -                   LOCAL   2160
10.8.30.42       a0:12:90:d3:ec:83  300   2/4                 DYNAMIC 2147
10.8.30.45       d4:ea:0e:e0:98:80  300   2/4                 DYNAMIC 2148
10.8.30.255      ff:ff:ff:ff:ff:ff  300   -                   LOCAL   2160
10.8.31.1        b0:ad:aa:47:09:03  301   1/39                LOCAL   2160
10.8.31.254      b0:ad:aa:47:09:04  301   1/39                DYNAMIC 1924
10.8.31.255      ff:ff:ff:ff:ff:ff  301   1/39                LOCAL   2160

8201:1#show ip arp vrf fab_lpbk

================================================================================
                                IP Arp - VRF fab_lpbk
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN  PORT                TYPE    TTL(10 Sec)
TUNNEL
--------------------------------------------------------------------------------
10.8.4.81        00:00:00:00:00:0a  -     -                   LOCAL   2160
10.8.31.1        b0:ad:aa:47:09:03  302   1/37                DYNAMIC 1915
10.8.31.254      b0:ad:aa:47:09:04  302   1/37                LOCAL   2160
10.8.31.255      ff:ff:ff:ff:ff:ff  302   1/37                LOCAL   2160
```

Fabric Extend and ONA Configuration

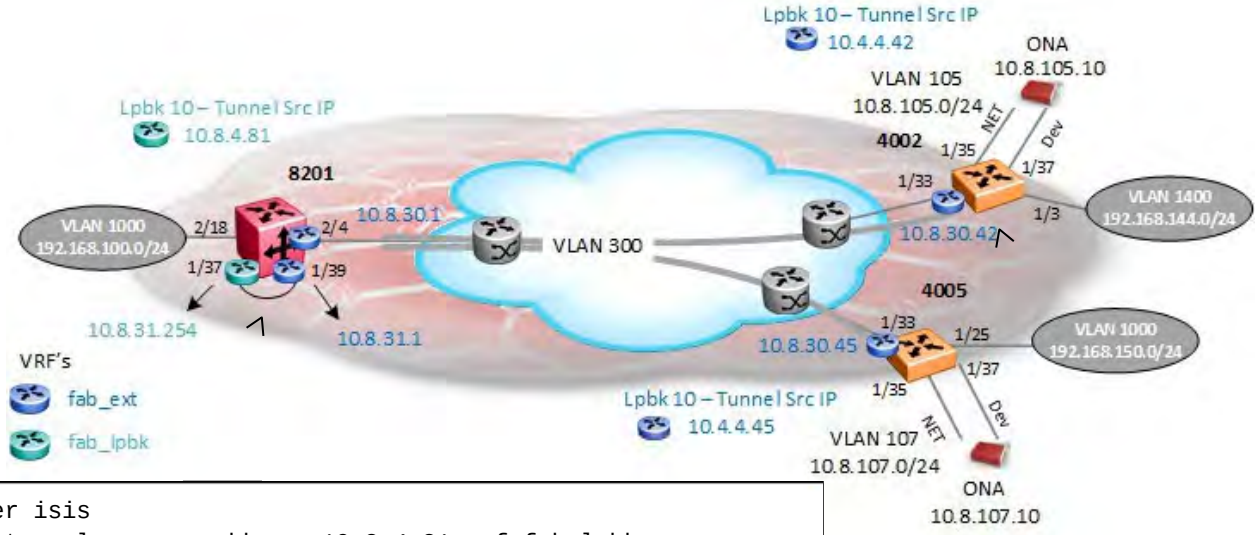


```
router isis
  ip-tunnel-source-address 10.8.4.81 vrf fab_lpbk
exit
logical-intf isis 1 dest-ip 10.4.4.42 name "tunnel_to_4002"
  isis
  isis spbm 1
  isis enable
exit
logical-intf isis 2 dest-ip 10.4.4.45 name "tunnel_to_4005"
  isis
  isis spbm 1
  isis enable
exit
```

```
vlan create 105 name "ona_vlan105" type port-mstprstp 0
vlan members 105 1/35
interface Vlan 105
  ip address 10.8.105.1 255.255.255.0
exit
interface GigabitEthernet 1/35
  no shutdown
exit
interface GigabitEthernet 1/37
  no shutdown
exit
router isis
  ip-tunnel-source-address 10.4.4.42 port 1/37 mtu 1950 vrf fab_ext
exit
logical-intf isis 1 dest-ip 10.4.4.81 name "tunnel_to_8201"
  isis
  isis spbm 1
  isis enable
exit
```

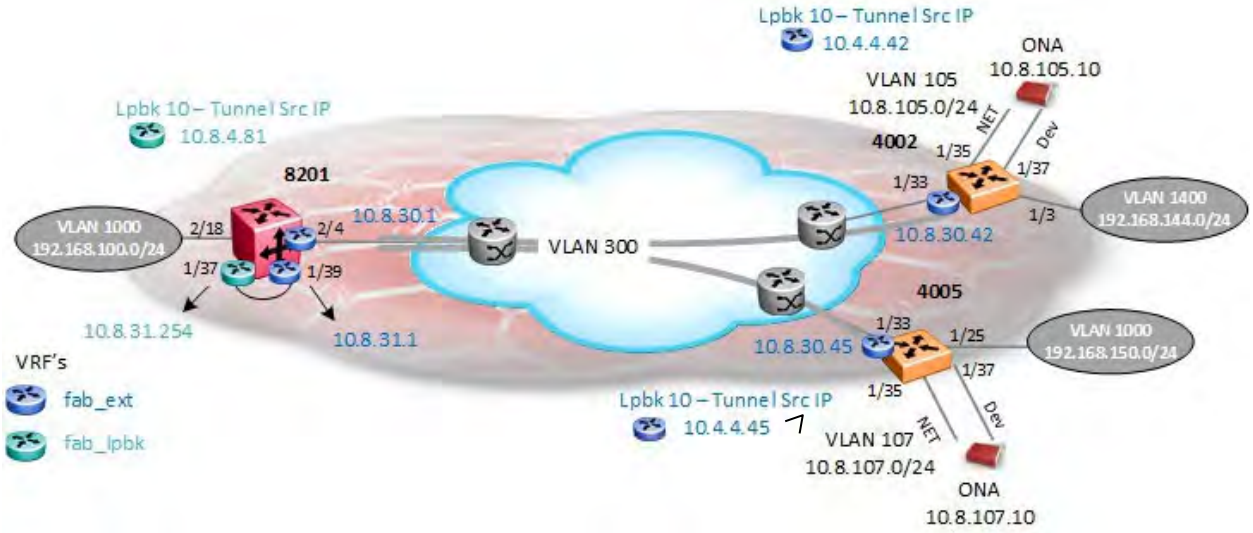


```
vlan create 107 name "ona_vlan107" type port-mstprstp 0
vlan members 107 1/35
interface Vlan 107
  ip address 10.8.105.1 255.255.255.0
exit
interface GigabitEthernet 1/35
  no shutdown
exit
interface GigabitEthernet 1/37
  no shutdown
exit
router isis
  ip-tunnel-source-address 10.4.4.45 port 1/37 mtu 1950 vrf fab_ext
exit
logical-intf isis 1 dest-ip 10.8.4.81 name "tunnel_to_8201"
  isis
  isis spbm 1
  isis enable
exit
```

Verify Operations

```
8201:1#show isis interface

================================================================================
                                ISIS Interfaces
================================================================================
IFIDX             TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
--------------------------------------------------------------------------------
Port1/17          pt-pt   Level 1   UP        UP         1      1       10
tunnel_to_4002    pt-pt   Level 1   UP        UP         1      1       20000
tunnel_to_4005    pt-pt   Level 1   UP        UP         1      1       20000

--------------------------------------------------------------------------------
 3 out of 3 Total Num of ISIS interfaces
--------------------------------------------------------------------------------

8201:1#show isis adjacencies

================================================================================
                                ISIS Adjacencies
================================================================================
INTERFACE         L STATE         UPTIME PRI  HOLDTIME SYSID             HOST-NAME
--------------------------------------------------------------------------------
Port1/17          1 UP      1d 01:39:27 127       20 a012.9002.d3df    7005
tunnel_to_4002    1 UP         01:00:40 127       24 a012.90d3.ec65    4002
tunnel_to_4005    1 UP         01:00:34 127       24 d4ea.0ee0.9865    4005

--------------------------------------------------------------------------------
 3 out of 3 interfaces have formed an adjacency
--------------------------------------------------------------------------------

8201:1#show isis logical-interface

================================================================================
                                                ISIS Logical Interfaces
================================================================================
IFIDX  NAME             ENCAP    L2_INFO                   TUNNEL     L3_TUNNEL_NEXT_HOP_INFO
                        TYPE     PORT/MLT   VIDS(PRIMARY)  DEST-IP    PORT/MLT   VLAN   VRF
--------------------------------------------------------------------------------
1      tunnel_to_4002   IP       --         --             10.4.4.42  Port1/37   302    fab_lpbk
2      tunnel_to_4005   IP       --         --             10.4.4.45  Port1/37   302    fab_lpbk

--------------------------------------------------------------------------------
 2 out of 2 Total Num of Logical ISIS interfaces
--------------------------------------------------------------------------------
```

```
8201:1#show ip route
================================================================================================
                                        IP Route - GlobalRouter
================================================================================================
                                                 NH                INTER
DST           MASK            NEXT          VRF/ISID      COST   FACE PROT AGE TYPE PRF
------------------------------------------------------------------------------------------------
10.1.1.81     255.255.255.255 10.1.1.81     -             1      0    LOC  0   DB   0
10.4.4.2      255.255.255.255 4002          GlobalRouter  20000  3051 ISIS 0   IBS  7
10.4.4.5      255.255.255.255 4005          GlobalRouter  20000  3051 ISIS 0   IBS  7
10.8.105.0    255.255.255.0   4002          GlobalRouter  20000  3051 ISIS 0   IBS  7
10.8.107.0    255.255.255.0   4005          GlobalRouter  20000  3051 ISIS 0   IBS  7
192.168.96.0  255.255.255.0   192.168.96.1  -             1      996  LOC  0   DB   0
192.168.100.0 255.255.255.0   192.168.100.1 -             1      1000 LOC  0   DB   0
192.168.144.0 255.255.255.0   4002          GlobalRouter  20000  3051 ISIS 0   IBS  7
192.168.150.0 255.255.255.0   4005          GlobalRouter  20000  3051 ISIS 0   IBS  7
```

20.1.31 Fabric Extend over E-LAN/VPLS (L2) network using VLAN Tunnels

Fabric Extend can also be transported over a layer 2 core without using VXLAN IP tunneling. In this case, logical ISIS interfaces are tunneling using B-VID translation where the two B-VLAN ID's are mapped to two different logical VLAN ID's. Please note that two core VLAN ID's are required for each point-to-point NNI connection between two BEB nodes. In in case, the ONA is not required with the VSP 4000, however, the third party core network must support jumbo frames.

Fabric Extend Configuration

```
logical-intf isis 1 vid 300-301 primary-vid 300 port 1/33 name "fe_to_8201"
isis
isis spbm 1
isis enable
exit

interface GigabitEthernet 1/33
  no shutdown
exit
```

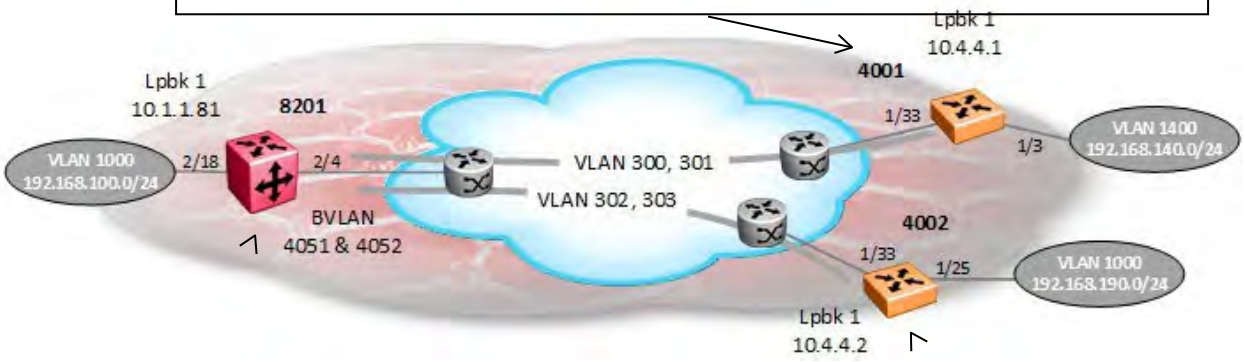


```
logical-intf isis 1 vid 300-301 primary-vid 300 port 2/4 name "fe_to_4001"
  isis
  isis spbm 1
  isis enable
exit
logical-intf isis 2 vid 302-303 primary-vid 302 port 2/4 name "fe_to_4002"
  isis
  isis spbm 1
  isis enable
exit

interface GigabitEthernet 2/4
  no shutdown
exit
```

```
logical-intf isis 1 vid 302-303 primary-vid 302 port 1/33 name "fe_to_8201"
  isis
  isis spbm 1
  isis enable
exit

interface GigabitEthernet 1/33
  no shutdown
exit
```

Verify ISIS Interfaces

```
8201:1#show isis logical-interface

================================================================================
                                    ISIS Logical Interfaces
================================================================================
IFIDX    NAME         ENCAP        L2_INFO                  TUNNEL    L3_TUNNEL_NEXT_HOP_INFO
                      TYPE         PORT/MLT   VIDS(PRIMARY)  DEST-IP   PORT/MLT   VLAN      VRF
--------------------------------------------------------------------------------------------------
1        fe_to_4001   L2-P2P-VID   Port2/4    300-301(300)   --        --         --        --
2        fe_to_4002   L2-P2P-VID   Port2/4    302-303(302)   --        --         --        --

8201:1#show isis interface

================================================================================
                              ISIS Interfaces
================================================================================
IFIDX            TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC
--------------------------------------------------------------------------------
fe_to_4001       pt-pt   Level 1   UP        UP         1      1       1000
fe_to_4002       pt-pt   Level 1   UP        UP         1      1       1000

8201:1#show isis adjacencies

================================================================================
                              ISIS Adjacencies
================================================================================
INTERFACE        L STATE         UPTIME PRI  HOLDTIME SYSID             HOST-NAME
--------------------------------------------------------------------------------
fe_to_4001       1 UP          17:54:30 127       23 d4ea.0e10.e465    4001
fe_to_4002       1 UP          18:53:10 127       23 a012.90d3.ec65    4002
```

Verify L2 Forwarding Table

```
8201:1#show isis spbm unicast-fib

================================================================================
                                SPBM UNICAST FIB ENTRY INFO
================================================================================
DESTINATION        BVLAN     SYSID            HOST-NAME         OUTGOING         COST
ADDRESS                                                        INTERFACE
--------------------------------------------------------------------------------
02:40:02:ff:ff:ff  3051      a012.90d3.ec65   4002              fe_to_4002       1000
a0:12:90:d3:ec:65  3051      a012.90d3.ec65   4002              fe_to_4002       1000
02:40:02:ff:ff:ff  3052      a012.90d3.ec65   4002              fe_to_4002       1000
a0:12:90:d3:ec:65  3052      a012.90d3.ec65   4002              fe_to_4002       1000
02:82:01:ff:ff:ff  3051      b0ad.aa47.0884   8201              cpp              0
b0:ad:aa:47:08:84  3051      b0ad.aa47.0884   8201              cpp              0
02:82:01:ff:ff:ff  3052      b0ad.aa47.0884   8201              cpp              0
b0:ad:aa:47:08:84  3052      b0ad.aa47.0884   8201              cpp              0
02:40:01:ff:ff:ff  3051      d4ea.0e10.e465   4001              fe_to_4001       1000
d4:ea:0e:10:e4:65  3051      d4ea.0e10.e465   4001              fe_to_4001       1000
02:40:01:ff:ff:ff  3052      d4ea.0e10.e465   4001              fe_to_4001       1000
d4:ea:0e:10:e4:65  3052      d4ea.0e10.e465   4001              fe_to_4001       1000
```

20.1.32 Fabric Attach Configuration

Fabric Attach Server – VSP 4000, VSP 8000 or VSP 7200

Fabric Attach (FA) is globally enabled by default and must be enabled on a port or MLT interface. By enable FA on a port or MLT, this will in turn enable LLDP on this interface and send out both FA TLVs and standard LLDP TLVs. Also, authentication is enabled by default.

```
config terminal
interface gigabitEthernet <slot/port>
  fa
  fa enable
exit
interface mlt <1-512>
  fa
  fa enable
exit
```

Fabric Attach / Zero Touch – FA Server

On the FA server, a FA port can be configured with the management VLAN using an ISID and VLAN ID value (c-vid). The c-vid is optional is and if not specified, then the management traffic will be untagged. The ISID is mandatory and is required for a network wide L2VSN. The FA server will announce this information in FA LLDP messages where only the management VLAN ID is announced as the ISID is not required. For untagged management, a VLAN ID of 4095 is announced.

```
config terminal
interface gigabitEthernet <slot/port>
  fa
  fa enable
  fa management i-sid <1-16777215> c-vid <1-4094>
exit
interface mlt <1-512>
  fa
  fa enable
  fa management i-sid <1-16777215> c-vid <1-4094>
exit
```

Fabric Attach Proxy or Proxy Standalone – ERS 4800 or ERS 5900

By default, Fabric Attach (FA) Proxy is globally enabled and all ports are enabled with LLDP. In regards to the ERS 4800, if it is running the standard image, you must disable FA message-authentication on the FA Server. FA message-authentication is only possible with the secure image for the ERS 4800. VLAN 1 membership must also be removed from all ports or at least the ports used for FA as FA will not otherwise remove it. If an MLT is used as the uplink to the FA server, then one must create the MLT and, only if the desired mode is Standalone Proxy, assign the MLT number as the FA trunk.

If the management VLAN is provisioned on the FA Server, this will be learned by the FA Proxy switch which in turn will attempt to get a management IP address via DHCP on it. The management VLAN ID will also be passed on to the FA Client, i.e. a WLAN 9100, via FA LLDP messages if used in the solution.

Fabric Attach– Using MLT Uplink

MLT configuration needs to be manually configured for example if using the FA Proxy switch to connect to a SMLT cluster.

```
config terminal
mlt <1-32> name <MLT name> member <list of ports> learning <disable|enable>
mlt <1-32> enable
vlan members remove 1 <list of ports>
```

Fabric Attach – Proxy Standalone

As mentioned above, FA Proxy is enabled by default. To configure the switch for FA Proxy Standalone, please enter the command shown below.

```
config terminal
fa standalone-proxy
fa uplink trunk <1-32>
```

In addition, we will also have to configure the management VLAN. The management IP address can be entered manually entered or DHCP can be used if enabled (*ip address source dhcp-when-needed*).

If Identity Engines is used, the policy created for a FA Proxy Standalone switch must contains the management VLAN ID and each VLAN all with an ISID value of 0. An actually ISID value is only required for a FA Proxy and FA Server configuration.

Fabric Attach– Using IDE Authentication Policy for WLAN 9100 FA Client Authentication

If the WLAN 9100 is deployed as a FA Client in the FA Proxy and FA Server solution, the FA Client needs to perform Non-EAP (NEAP) Multiple Host Single Authentication (MHSA) authentication using Identity Engines (IDE). A NEAP policy must be created on IDE which will push down all the required VLAN/ISID combinations including the management VLAN. Once a FA Client is authenticated, all VLAN mappings on this port will be cleared and the IDE policy VLAN/ISID and management VLAN settings will be used on this particular port.

If Wireless LAN Orchestration System (WOS) is used, please ensure that WOS Management VLAN setting in the profile used for the FA Client and the FA Management VLAN values using in the FA Server are the same. A WLAN 9100 access point will always use the WOS assigned management VLAN as the final setting.

Fabric Attach / Zero Touch – FA Proxy

FA Zero Touch eases the configuration process on FA-capable devices by automating specific configuration tasks required for FA functionality. For situations when you prefer or require manual configuration of the settings affected by Zero Touch, feature control is provided.

Fabric Attach must be enabled in order for Zero Touch to function.

When base Zero Touch functionality is enabled, FA Proxy and FA Client devices can acquire management VLAN data from the connected FA Server or FA Proxy and use it to facilitate manageability and network configuration. By default, base Zero Touch support is enabled. These options can be controlled by setting the appropriate zero touch options.

You can configure the following Zero Touch options on a FA Proxy switch:

```
4850GTS-PWR+(config)#fa zero-touch-options ?
  auto-port-mode-fa-client     Enable automated EAP client port configuration
  auto-pvid-mode-fa-client     Enable automatic client PVID/Mgmt VLAN update
  auto-trusted-mode-fa-client  Enable automatic trusted client connection
  ip-addr-dhcp                 Enable automated DHCP IP address acquisition
```

Automated FA Client Port Mode

When this option is enabled and FA Clients are present, the EAP settings for the interface on which the client is discovered, are automatically updated based on the FA Client type. If the FA Clients of the appropriate type are deemed no longer valid (when element aging causes the FA Client to be deleted from the discovered elements list), the EAP port settings revert to the previous state.

Automated configuration only applies to FA-enabled ports.

Two FA Client types are supported:

- Wireless Access Point Type 1 (direct network attachment)
- Wireless Access Point Type 2 (clients tunneled to controller)

If FA Standalone Proxy topology is in use, the EAP automated settings are updated only if the FA Client is a Wireless Access Point Type 1.

Automated PVID FA Client Port Mode

When this option is enabled, automatic port PVID and management VLAN membership updates are initiated based on the type of discovered FA Clients. This is applicable for FA Proxy, FA Server, and FA Standalone Proxy devices. Automated configuration is only applied to FA-enabled ports. Please note that VLAN 1 still have to be manually removed at this time.

QoS interface class data is updated based on the discovery and deletion (based on aging and port events) of the following FA Client types:

- Wireless Access Point Type 1 (direct network attachment)
- Wireless Access Point Type 2 (clients tunneled to controller)

Automated trusted FA Client connection

This option enables automatic trusted FA Client connection. When this option is enabled and FA Clients are present, the QoS settings for the interface on which the client is discovered are automatically updated to QoS 'Trusted'. If the FA Clients of the appropriate type are deemed no longer valid (when element aging causes the FA Client to be deleted from the discovered elements list), the QoS port settings revert to the previous state.

QoS interface class data is updated based on the discovery and deletion (based on aging) of the following FA Client types:

- Wireless Access Point Type 1 (direct network attachment)
- Wireless Access Point Type 2 (clients tunneled to controller)

Automated configuration only applies to FA-enabled ports.

IP Address DHCP

When this option is enabled, IP address source mode is updated on the FA Proxy device (receiver) to DHCP-When-Needed and initiates DHCP-based IP address acquisition if an IP address is not manually configured.

IP address source mode update only occurs during base Zero Touch processing when a new management VLAN is processed if this option is enabled.

20.1.33Identity Engines – Attribute Details

Outbound Attributes from IDE to FA Switch	
Fabric-Attach-VLAN-Create	This attribute behaves as Boolean 0 or not send = Switch will NOT create assigned VLAN if VLAN does not exist 1 = Switch will create assigned VLAN if the VLAN does not exit
Fabric-Attach-VLAN-ISID	– VLAN and ISID IDs are separated by ":" – Example: 20:20000 means VLAN=20 and ISID 20000 – VLAN = 1 to 4095 OR VLAN Name – ISID >0 and < max ISID
Fabric-Attach-VLAN-PVID	– PVID VLAN ID – VLAN = 1 to 4095

Inbound Attributes from FA Switch to IDE	
Fabric-Attach-Switch-Mode	0 or not send = Switch is assumed to have not concept of SPB/Fabric Attach (i.e. switch is in neither mode – not 1, 2, 3, 4, nor 5) 1 = Switch is FA Server with SPB Disabled and FA Enabled (Network Type 2) 2 = Switch is FA Server with SPB Enabled (Network Type 1a & 1b) 3 = Switch is FA Proxy connected to FA Server with SPB Disabled and FA Enabled (Network Type 2) 4 = Switch is FA Proxy connected to FA Server with SPB Enabled (Network Type 1a) 5 = Switch is FA Proxy Standalone (Network Type 3)
Fabric-Attach-Client-Id	MAC address of the FA Client taken from LLDP via FA Module: FA Element TLV > FA Element ID > System MAC Address
Fabric-Attach-Client-Type	FA Client Type taken from LLDP via FA Module: FA Element TLV > FA Element Type 1 = FA Element Type Other 2 = FA Server 3 = FA Proxy 4 = FA Server No Authentication 5 = FA Proxy No Authentication 6 = FA Client – Wireless AP Type 1 [clients direct network attachment] (e.g. AP 9100) 7 – FA Client – Wireless AP Type 2 [clients tunneled to controller]

20.1.34Fabric Attach Base Configuration – Adding a FA Proxy and FA Server

Fabric Attach Proxy Configuration – ERS 4800

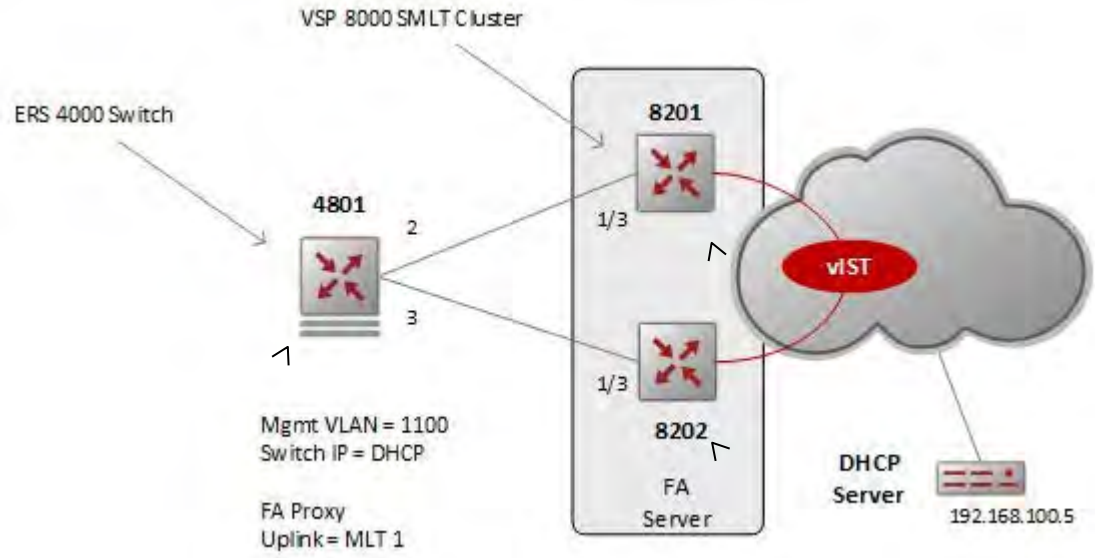


```
configure terminal
snmp-server name 4801
mlt 1 name fa_mlt1 member 2,3 learning disable
mlt 1 enable
vlan configcontrol automatic
vlan members remove 1 2,3
fa extended-logging
```

```
configure terminal
mlt 7 enable name mlt7_4801
mlt 7 encapsulation dot1q
mlt 7 member 1/3
interface mlt 7
  smlt
  fa
  fa enable
  fa management i-sid 5001100 c-vid 1100
exit
```

Please note that at this stage, the FA Proxy will not get an IP address until a platform VLAN has been added on the FA Server switches 8201 and 8202 – this will be the next step.

Please provision the FA Proxy switch 4801 first prior to connecting the uplink ports to the FA Server as the MLT has to be manually provisioned at this time. If this is not done, the FA Proxy will automatically learn the FA Server via one of the two ports and thereby adding the management VLAN on this port. On an ERS 4800, an MLT cannot be created if there is a mis- match in the VLAN ID's.

Verify Operations - LLDP

At this stage, as long as the FA Server has Fabric Attach enabled at a port or MLT level, the FA Proxy switch 4801 should have discovered the FA Server switches 8201 and 8202.

8201 & 8202

```
8202:1#show lldp port 1/3

================================================================================
                            LLDP Admin Port Status
================================================================================
--------------------------------------------------------------------------------------------------------------------
Port       AdminStatus  ConfigNotificationEnable
--------------------------------------------------------------------------------------------------------------------
1/3        txAndRx      disabled

8202:1#show lldp neighbor

================================================================================================
                                        LLDP Neighbor
================================================================================================

Port: 1/3       Index    : 3                  Time: 0 day(s), 02:53:49
                ChassisId: MAC Address        cc:f9:54:b4:ac:00
                PortId   : MAC Address        cc:f9:54:b4:ac:03
                SysName  : 4801
                SysCap   : Br / Br
                PortDescr: Port 3
                SysDescr : Ethernet Routing Switch 4850GTS-PWR+  HW:00
FW:5.8.0.1   SW:v5.9.0.005
```

4801

```
4801(config)#show lldp neighbor
----------------------------------------------------------------------------------------------------------------
                              LLDP neighbor
----------------------------------------------------------------------------------------------------------------
Port: 2     Index: 1                 Time: 0 days, 02:42:31
        ChassisId: MAC address       b0:ad:aa:47:08:00
        PortId:    MAC address       b0:ad:aa:47:08:02
        SysName:   8201
        SysCap:    rB / rB           (Supported/Enabled)
        PortDesc:  Extreme Virtual Services Platform 8284XSQ - GbicOther Port 1/3
        SysDescr:  VSP-8284XSQ (5.0.0.0) (PRIVATE)
----------------------------------------------------------------------------------------------------------------
Port: 3     Index: 3                 Time: 0 days, 02:42:58
        ChassisId: MAC address       e4:5d:52:3c:48:00
        PortId:    MAC address       e4:5d:52:3c:48:02
        SysName:   8202
        SysCap:    rB / rB           (Supported/Enabled)
        PortDesc:  Extreme Virtual Services Platform 8284XSQ - GbicOther Port 1/3
        SysDescr:  VSP-8284XSQ (5.0.0.0) (PRIVATE)
```

```
4801(config)#show lldp neighbor vendor-specific avaya fabric-attach
--------------------------------------------------------------------------------------------------------------
                      Neighbors LLDP info - Avaya FA TLVs
--------------------------------------------------------------------------------------------------------------
--------------------------------------------------------------------------------------------------------------
Port: 2
   Fabric Attach Data:
      Element Type: server
      Management VLAN: 1100
      System ID: b0:ad:aa:47:08:85:30:07:00:07
      Element State Flags (0x90):
          trafficTagged
          provisionModeSpbm
      Exported I-SID/VLAN Assignments:
         No I-SID/VLAN Assignments.
Port: 3
   Fabric Attach Data:
      Element Type: server
      Management VLAN: 1100
      System ID: b0:ad:aa:47:08:85:30:07:00:07
      Element State Flags (0x90):
          trafficTagged
          provisionModeSpbm
      Exported I-SID/VLAN Assignments:
         No I-SID/VLAN Assignments.
```

Verify Operations - FA default Agent Status

8201 & 8202

```
8202:1#show fa agent

================================================================================
                          Fabric Attach Configuration
================================================================================
                             FA Service : enabled
                        FA Element Type : server
                  FA Assignment Timeout : 240
                   FA Discovery Timeout : 240
                      FA Provision Mode : spbm
```

4801

```
4801(config)# show fa agent

Fabric Attach Service Status: Enabled
Fabric Attach Element Type: Proxy
Fabric Attach Zero Touch Status: Enabled
Fabric Attach Auto Provision Setting: Proxy
Fabric Attach Provision Mode: SPBM
Fabric Attach Client Proxy Status: Enabled
Fabric Attach Standalone Proxy Status: Disabled
Fabric Attach Agent Timeout: 75 seconds
```

```
Fabric Attach Extended Logging Status: Enabled
Fabric Attach Primary Server Id: b0:ad:aa:47:08:85:30:07:00:07 (SPBM)
Fabric Attach Primary Server Descr:
VSP-8284XSQ (5.0.0.0) (PRIVATE)
```

Verify Operations - FA Interface

8201 & 8202

```
8202:1#show fa interface

================================================================================
                          Fabric Attach Interfaces
================================================================================
INTERFACE    SERVER   MGMT     MGMT     MSG AUTH MSG AUTH
             STATUS   ISID     CVID     STATUS   KEY
--------------------------------------------------------------------------------
Mlt7         enabled  5001100  1100     enabled  ****
```

4801

```
4801(config)#show fa port-enable 2,3

Unit Port IfIndex Trunk Service Advertisement Authentication
____ ____ _______ _____ _____________________ ______________
1    2    2       1          Enabled             Enabled
1    3    3       1          Enabled             Enabled
```

Verify Operations - ELAN and ISID

8201 & 8202

```
8202:1#show i-sid 5001100

================================================================================
                                    Isid Info
================================================================================
ISID        ISID                  PORT              MLT               ORIGIN
ID          TYPE       VLANID     INTERFACES        INTERFACES
--------------------------------------------------------------------------------
5001100     ELAN       1100       -                 c1100:7           MANAGEMENT
                       1100       -                 8
```

Under the INTERFACES column, c1100:7 refers to C-VLAN 1100 and MLT 7 as used in this configuration example. Under the ORIGIN column, MANAGEMENT indicates the FA management VLAN has been provisioned on the FA Server switch.

Verify Operations - FA Elements

8201 & 8202

```
8201:1#show fa elements

================================================================================
                        Fabric Attach Discovery Elements
================================================================================
                         MGMT                                          ELEM ASGN
PORT    TYPE             VLAN STATE  SYSTEM ID                         AUTH AUTH
--------------------------------------------------------------------------------
1/3     proxy            1100 T / S  cc:f9:54:b4:ac:00:20:00:00:01    AP   AP

================================================================================
                      Fabric Attach Authentication Detail
================================================================================
        ELEM OPER                     ASGN OPER
PORT    AUTH STATUS                   AUTH STATUS
--------------------------------------------------------------------------------
1/3     successAuth                   successAuth

State Legend: (Tagging/AutoConfig)
T= Tagged,    U= Untagged,    D= Disabled,    S= Spbm,    V= Vlan,    I= Invalid

Auth Legend:
AP= Authentication Pass,AF= Authentication Fail,
NA= Not Authenticated,  N= None


--------------------------------------------------------------------------------
1 out of 1 Total Num of fabric attach discovery elements displayed
--------------------------------------------------------------------------------
```

4801:

```
4801(config)#show fa elements

Unit/   Element    Element           Element
Port     Type      Subtype            VLAN   Auth    System ID
---------- ---------- ------------------------- ---------- ------ -----------------------------------
MLT1    Server  Server (Auth)         1100    AP  b0:ad:aa:47:08:85:30:07:00:07

Legend:
Auth - AP=Authentication Pass, AF=Authentication Fail, NA=Not Authenticated
```

20.1.34.1 Fabric Attach – Adding a Platform VLAN on FA Server for Management VLAN

In order for the FA Proxy switch to get an IP address via DHCP, we will need to add a platform VLAN and add an IP address with DHCP. We will do this on the FA Server switches and enable the RSMLT Edge option assuming IP Shortcuts has already been enabled – please see the IP Shortcuts section for configuration details.

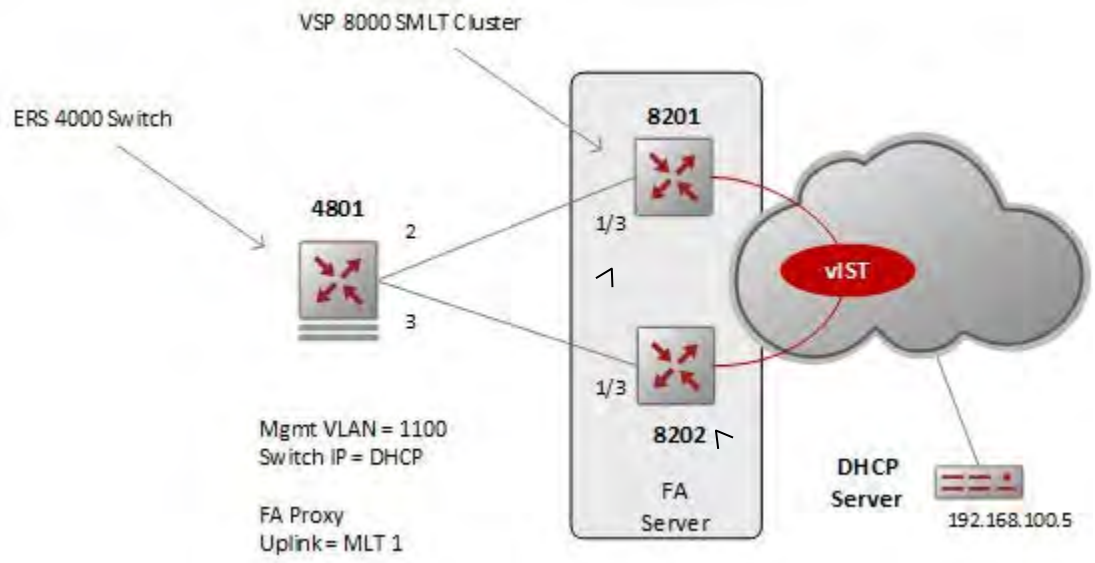


```
config terminal
vlan create 1100 name "fa_mgmt_vlan1100"
type port-mstprstp 0
vlan i-sid 1100 5001100
interface Vlan 1100
ip address 10.12.110.1 255.255.255.0
  ip dhcp-relay
  ip rsmlt
  ip rsmlt holdup-timer 9999
  ip dhcp-relay fwd-path 192.168.100.5
  ip dhcp-relay fwd-path 192.168.100.5
enable
  ip dhcp-relay fwd-path 192.168.100.5 mode
dhcp
exit
interface GigabitEthernet 1/3
  no shutdown
exit
ip rsmlt edge-support
```

```
configure terminal
vlan create 1100 name "fa_mgmt_vlan1100"
type port-mstprstp 0
vlan i-sid 1100 5001100
interface Vlan 1100
ip address 10.12.110.2 255.255.255.0
  ip dhcp-relay
  ip rsmlt
  ip rsmlt holdup-timer 9999
  ip dhcp-relay fwd-path 192.168.100.5
  ip dhcp-relay fwd-path 192.168.100.5
enable
  ip dhcp-relay fwd-path 192.168.100.5
mode dhcp
exit
interface GigabitEthernet 1/3
  no shutdown
exit
ip rsmlt edge-support
```

The platform VLAN ISID must match the ELAN ISID used as per the previous step.

Verify Operations - DHCP

At this stage, the Fabric Attach Proxy switch should get an IP address via DHCP plus all the Fabric Attach elements should be discovered.

```
4801(config)#show ip
Bootp/DHCP Mode: DHCP  When Needed

                          Configured        In Use        Last BootP/DHCP
                         -------------- -------------- --------------------
Stack IP Address:        192.168.1.2                   0.0.0.0
Switch IP Address:       192.168.1.1     10.12.110.10  10.12.110.10
Switch Subnet Mask:      255.255.255.0   255.255.255.0 255.255.255.0
Default Gateway:         0.0.0.0         10.12.110.1   10.12.110.1
```

20.1.34.2 Fabric Attach – Adding a L2VSN Service

Continuing from the base setup above, assuming we wish to add an L2VSN service from the FA Proxy switch for a local user to a remote SPBM node. This can be easily accomplished by simply adding a VLAN and ISID on the FA Proxy switch which in turn will be automatically added as an ELAN on the FA Server switches. All we then need to do is create an L2VSN with the same ISID on terminating SPBM switch.

For this example, SPBM node 7006 is a VSP 7000 switch

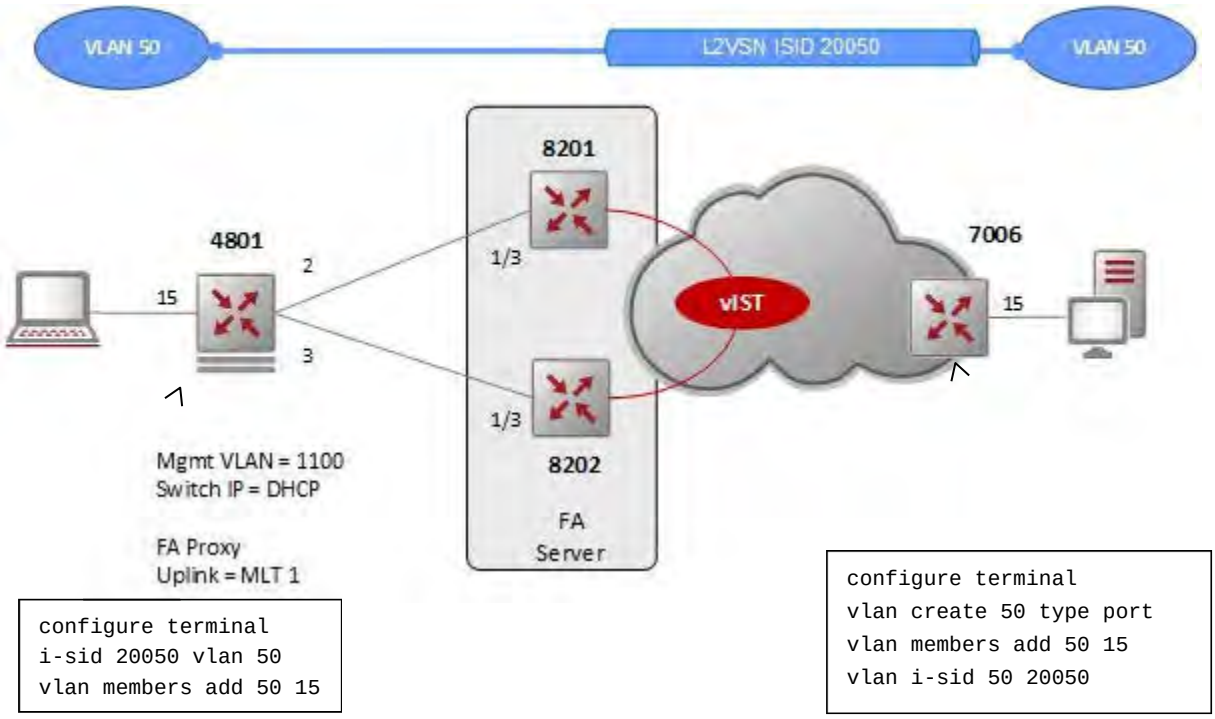

Verify Operations - FA Proxy switch 4801

```
4801#show fa i-sid 20050

 I-SID     VLAN      Source        Status
-----------  ------  ------------------  -----------
20050      50     Proxy           Active

Binding Count: 1
```

Verify the uplink MLT port members are added:

```
4801#show vlan id 50
Id    Name                    Type      Protocol          PID      Active IVL/SVL Mgmt
------ --------------------------- ----------- ---------------------- ---------- -------- ---------- ------
50    VLAN #50                Port      None              0x0000   Yes     IVL      No
        Port Members: 2-3,15
Total VLANs: 1
```

A log entry should also have been recorded provided the *fa extended-logging* option has been enabled:

```
4801#show logging sort-reverse
Type Time         Idx  Src Message
---- ------------ ---- --- -------
I    00:02:10:00  99    Fabric Attach: binding activation success (trunk 1 20050/50)
```

Verify Operations - FA Server switches 8201 and 8202

```
8201:1#show fa assignment

================================================================================
                        Fabric Attach Assignment Map
================================================================================
Interface  I-SID      Vlan       State      Origin
--------------------------------------------------------------------------------
1/3        20050      50         active     proxy
1/3        5001100    1100       active     proxy
1/3        5001110    1110       active     proxy
1/3        5001120    1120       active     proxy

8201:1#show i-sid 20050

================================================================================
                                    Isid Info
================================================================================
ISID        ISID                  PORT              MLT               ORIGIN
ID          TYPE       VLANID     INTERFACES        INTERFACES
--------------------------------------------------------------------------------
20050       ELAN       N/A        -                 c50:7             DISC_BOTH
```

To view MAC learning, use the *show i-sid mac-address-entry 20050* ALCI command as this is a ELAN service and not a platform VLAN where “20050” is the ISID used in this configuration example.

Verify Operations - SPBM switch 7006

```
7006#show isis spbm i-sid all id 20050
================================================================================
                            SPBM ISID INFO
================================================================================
ISID      SOURCE NAME    VLAN   SYSID            TYPE       HOST-NAME
================================================================================
20050     0.70.06        3052   a012.9004.2fdf   config     7006
20050     0.82.01        3051   b0ad.aa47.0884   discover   8201
20050     0.82.02        3052   e45d.523c.4884   discover   8202
```

20.1.34.3 Fabric Attach – Adding a L3VSN Service

Continuing from the base setup above, assuming we wish to add an L3VSN service. This can be easily accomplished by simply adding a VLAN and ISID on the FA Proxy switch which in turn will be automatically added as an ELAN on the FA Server switches. In this example, we will create a platform VLAN on the FA Server switches with RSMLT Edge support.

Fabric Attach Proxy Configuration

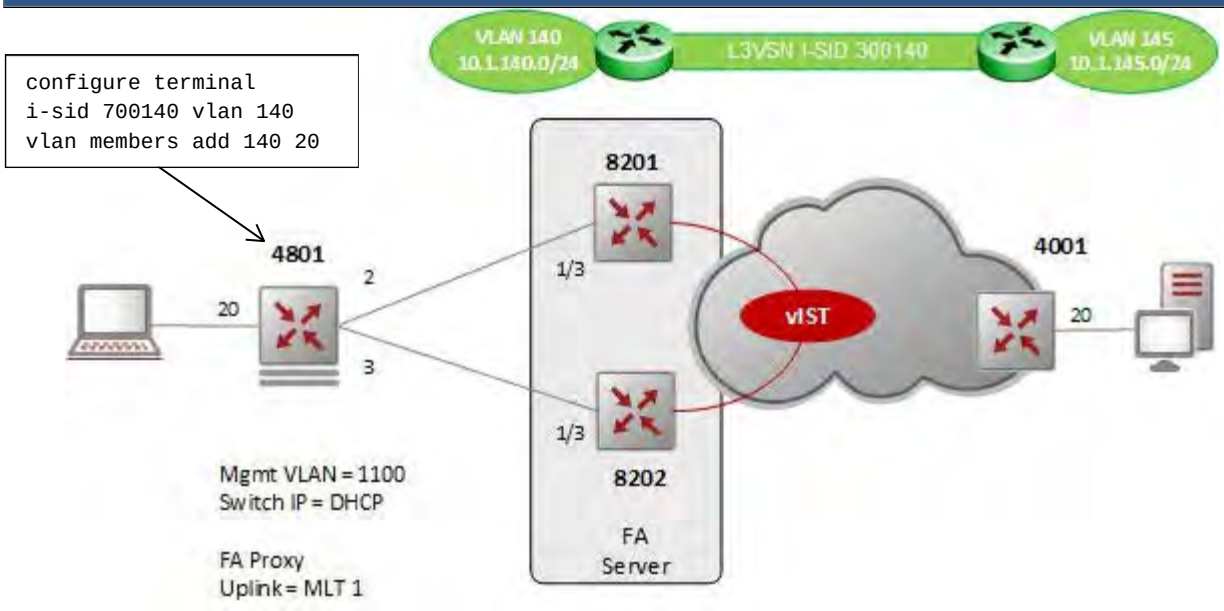

Verify Operations - ISID creation and state

```
4801#show fa i-sid 700140

 I-SID    VLAN     Source       Status
-------- ----  -------------  --------
700140    140   Proxy          Active

Binding Count: 1

4801(config)#show logging sort-reverse
Type Time                           Idx  Src Message
----- ---------------------------------- ----- ---- ----------
I    03:19:30:54                    130      Fabric Attach: binding activation success
(trunk 1 700140/140)
```

```
8201:1#show i-sid 700140
================================================================================================
                                        Isid Info
================================================================================================
ISID        ISID                    PORT                MLT                 ORIGIN
ID          TYPE       VLANID       INTERFACES          INTERFACES
------------------------------------------------------------------------------------------------
700140      ELAN       N/A          -                   c140:7              DISC_BOTH
```

Verify Operations – LLDP bindings

```
4801#show lldp neighbor vendor-specific avaya fabric-attach
------------------------------------------------------------------------------------------------------------
                      Neighbors LLDP info - Avaya FA TLVs
------------------------------------------------------------------------------------------------------------
------------------------------------------------------------------------------------------------------------
Port: 3
   Fabric Attach Data:
      Element Type: server
      Management VLAN: 1100
      System ID: b0:ad:aa:47:08:85:30:07:00:07
      Element State Flags (0x90):
          trafficTagged
          provisionModeSpbm
      Exported I-SID/VLAN Assignments:
         I-SID: 20050 VLAN: 50 Status: Active
         I-SID: 700140 VLAN: 140 Status: Active
Port: 2
   Fabric Attach Data:
      Element Type: server
      Management VLAN: 1100
      System ID: b0:ad:aa:47:08:85:30:07:00:07
      Element State Flags (0x90):
          trafficTagged
          provisionModeSpbm
      Exported I-SID/VLAN Assignments:
         I-SID: 20050 VLAN: 50 Status: Active
         I-SID: 700140 VLAN: 140 Status: Active
```

Verify Operations – FA assignments

```
8201:1#show fa assignment

================================================================================
                          Fabric Attach Assignment Map
================================================================================
Interface  I-SID      Vlan       State      Origin
--------------------------------------------------------------------------------
1/3        20050      50         active     proxy
1/3        700140     140        active     proxy
```

Fabric Attach Server Configuration – adding platform VLAN

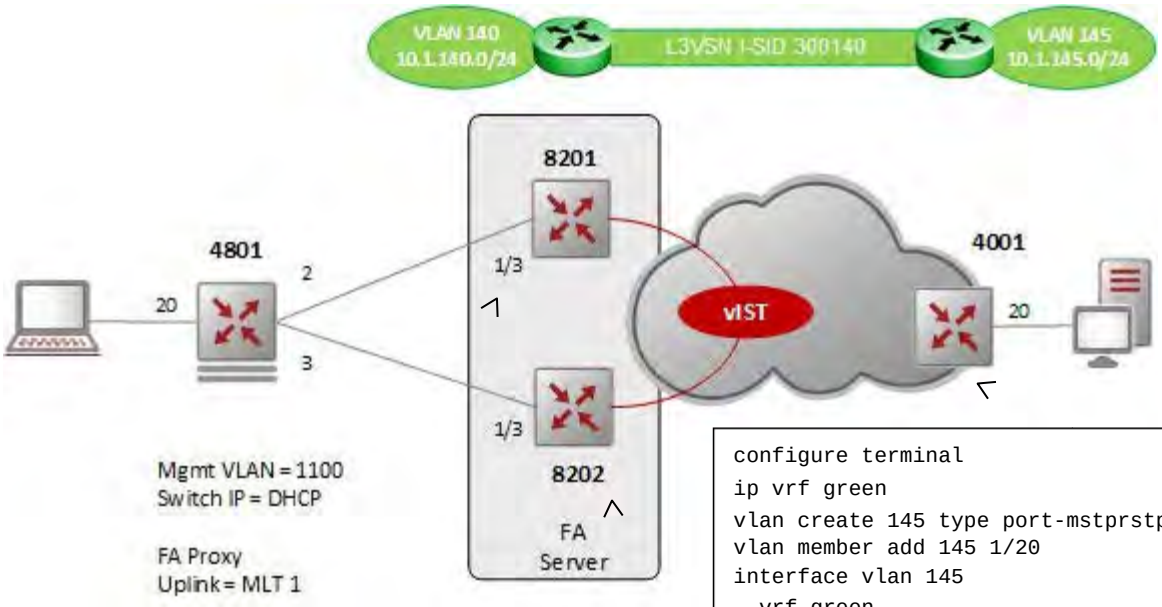


```
configure terminal
ip vrf green
vlan create 140 type port-mstprstp 0
vlan i-sid 140 700140
interface vlan 140
  vrf green
  ip address 10.1.140.1 255.255.255.0
  ip rsmlt
  ip rsmlt holdup-timer 9999
exit
router vrf green
  ipvpn
  i-sid 300140
  ipvpn enable
  isis redistribute direct
  isis redistribute direct enable
exit
isis apply redistribute direct vrf green
```

```
configure terminal
ip vrf green
vlan create 145 type port-mstprstp 0
vlan member add 145 1/20
interface vlan 145
  vrf green
  ip address 10.1.145.1 255.255.255.0
exit
router vrf green
  ipvpn
  i-sid 300140
  ipvpn enable
  isis redistribute direct
  isis redistribute direct enable
exit
isis apply redistribute direct vrf green
interface gigabitEthernet 1/20
  no shutdown
exit
```

```
configure terminal
ip vrf green
vlan create 140 type port-mstprstp 0
vlan i-sid 140 700140
interface vlan 140
  vrf green
  ip address 10.1.140.2 255.255.255.0
  ip rsmlt
  ip rsmlt holdup-timer 9999
exit
router vrf green
  ipvpn
  i-sid 300140
  ipvpn enable
  isis redistribute direct
  isis redistribute direct enable
exit
isis apply redistribute direct vrf green
```

Verify Operations – IP routing

VSP Cluster Switches

```
show ip route vrf green
show ip rsmlt vrf green
show ip arp vrf green
ping <remote IP> vrf green source <local IP>
```

20.1.34.4 Fabric Attach - Adding a WLAN 9100 FA Client with EAP Device authentication via Identity Engines

Continuing from the base setup above, assuming we wish to add an FA client. For this example, we will add a FA WLAN 9100 AP client and use Identity Engines to authenticate the AP and push down all the VLAN and ISID combinations required.

Fabric Attach Proxy Configuration – ERS 4800

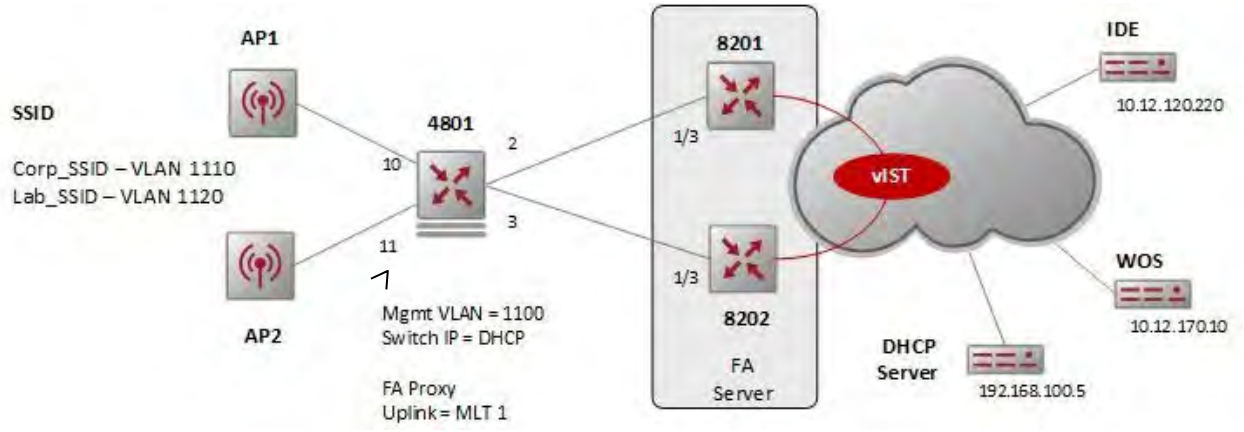


```
configure terminal
fa zero-touch-options auto-port-mode-fa-client
radius server host 10.12.120.220 key acct-enable
  Enter key: ******
  Confirm key: ******
```

The fa zero-touch-option setting of auto-port-mode-fa-client will automatically enable all the various EAP MHSA settings and enable EAP on the Fabric Attach (FA) client ports when Fabric Attached discovers a FA client. Once the FA client is discovered and authenticated against Identify Engines (IDE), IDE will overwrite the port VLAN port membership.

Once the FA WLAN 9100 AP Client is authenticated using EAP MAC authentication on the ERS 4800 FA Proxy switch, all VLAN information on the port will be lost. Hence, the policy used on Identity Engines will need to provision the management VLAN PVID in addition to the management VLAN and ISID combination. This is in addition to enabling FA VLAN create and adding all the user VLANs and ISID combinations required by the WLAN 9100 AP. Please see the Outbound Attributes chart in the section titled “Identity Engines – Attribute Details”.

Verify Operations

At this point, we should be able to discover the FA Client assuming a WLAN 9100 AP is attached to port 11.

```
4801(config)#show fa elements

Unit/  Element    Element           Element
Port    Type      Subtype            VLAN   Auth   System ID
------- ------- ------------------------- ------- ---------
MLT1   Server  Server (Auth)       1100    AP  b0:ad:aa:47:08:85:30:07:00:07
1/11   Client  Wireless AP(Type 1)  0       AP  64:a7:dd:03:39:03:00:00:00:01

Legend:
Auth - AP=Authentication Pass, AF=Authentication Fail, NA=Not Authenticated

4801#show lldp port 11 neighbor vendor-specific avaya fabric-attach
--------------------------------------------------------------------------------------------
                    Neighbors LLDP info - Avaya FA TLVs
--------------------------------------------------------------------------------------------
--------------------------------------------------------------------------------------------
Port: 11
   Fabric Attach Data:
      Element Type: client
      Management VLAN: 0
      System ID: 64:a7:dd:03:39:03:00:00:00:01
      Element State Flags (0x60):
          trafficTaggedAndUntagged
          provisionModeDisabled
      Exported I-SID/VLAN Assignments:

4801(config)#show vlan interface info 11
      Filter     Filter
     Untagged Unregistered
Port  Frames     Frames     PVID PRI    Tagging     Name
----- ---------- ---------- ---------- ------- ---------------------------- --------------------
11    No       Yes          1    0    UntagPvidOnly Port 11
```

Please note that if Identity Engines is not reachable or has not been provisioned, the PVID on the access port to the wireless AP will remain in the default VLAN. Also, if the WLAN 9100 AP is in Factory Default, it will have no VLANs thus the management VLAN will be 0.

```
4801#show logging sort-reverse
Type Time                           Idx   Src Message
----- ---------------------------------- ----- ---- ---------
I    03:20:34:08                    151       Fabric Attach: device discovered (Authentication
Pass - port 11)
I    03:20:31:45                    146       PoE Port Detection Status:  Unit 1 / Port 11
Status: Delivering Power
```

Adding Platform VLANs

On the FA Server switches, assuming IP Shortcuts has been enabled, we will add the the C-VLAN's used by the FA Proxy and FA Client assuming the Management VLAN has already been configured - see section titled "*Fabric Attach – Adding a Platform VLAN on FA Server for Management VLAN*".

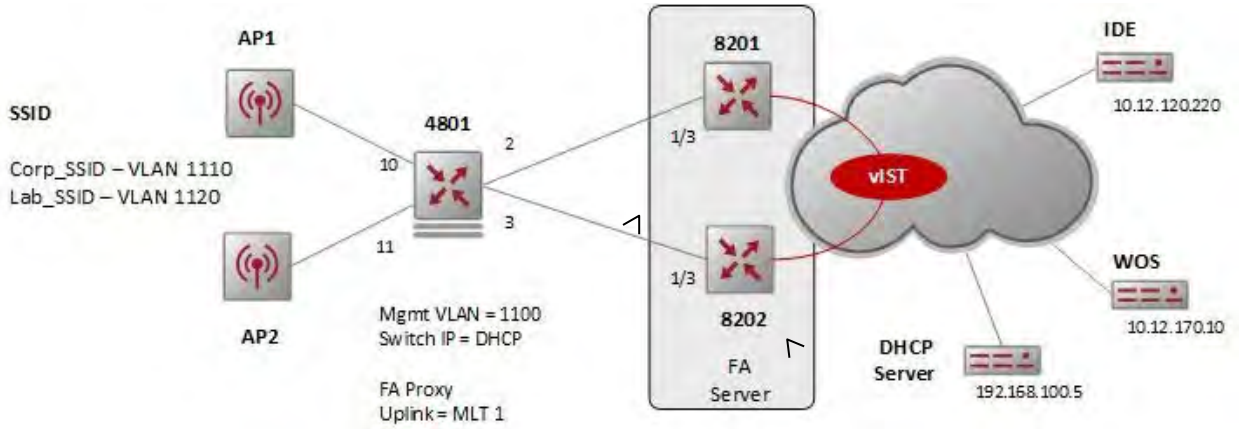


```
vlan create 1110 name "Corp_vlan1110" type
port-mstprstp 0
vlan i-sid 1110 5001110
interface Vlan 1110
  ip address 10.12.111.1 255.255.255.0
  ip dhcp-relay
  ip rsmlt
  ip rsmlt holdup-timer 9999
  ip dhcp-relay fwd-path 192.168.100.5
  ip dhcp-relay fwd-path 192.168.100.5 enable
  ip dhcp-relay fwd-path 192.168.100.5 mode
dhcp
exit
vlan create 1120 name "Lab_vlan1120" type
port-mstprstp 0
vlan i-sid 1120 5001120
interface Vlan 1120
  ip address 10.12.112.1 255.255.255.0
  ip dhcp-relay
  ip rsmlt
  ip rsmlt holdup-timer 9999
  ip dhcp-relay fwd-path 192.168.100.5
  ip dhcp-relay fwd-path 192.168.100.5 enable
  ip dhcp-relay fwd-path 192.168.100.5 mode
dhcp
exit
```

```
vlan create 1110 name "Corp_vlan1110" type
port-mstprstp 0
vlan i-sid 1110 5001110
interface Vlan 1110
  ip address 10.12.111.2 255.255.255.0
  ip dhcp-relay
  ip rsmlt
  ip rsmlt holdup-timer 9999
  ip dhcp-relay fwd-path 192.168.100.5
  ip dhcp-relay fwd-path 192.168.100.5 enable
  ip dhcp-relay fwd-path 192.168.100.5 mode
dhcp
exit
vlan create 1120 name "Lab_vlan1120" type
port-mstprstp 0
vlan i-sid 1120 5001120
interface Vlan 1120
  ip address 10.12.112.2 255.255.255.0
  ip dhcp-relay
  ip rsmlt
  ip rsmlt holdup-timer 9999
  ip dhcp-relay fwd-path 192.168.100.5
  ip dhcp-relay fwd-path 192.168.100.5 enable
  ip dhcp-relay fwd-path 192.168.100.5 mode
dhcp
exit
```

Identity Engines Setup – Add an internal group for the AP's

We will create a new group to be used for the IDE policy as a container for all the AP MAC address that will be authenticated when an AP connects to the Fabric Attach Proxy switch.

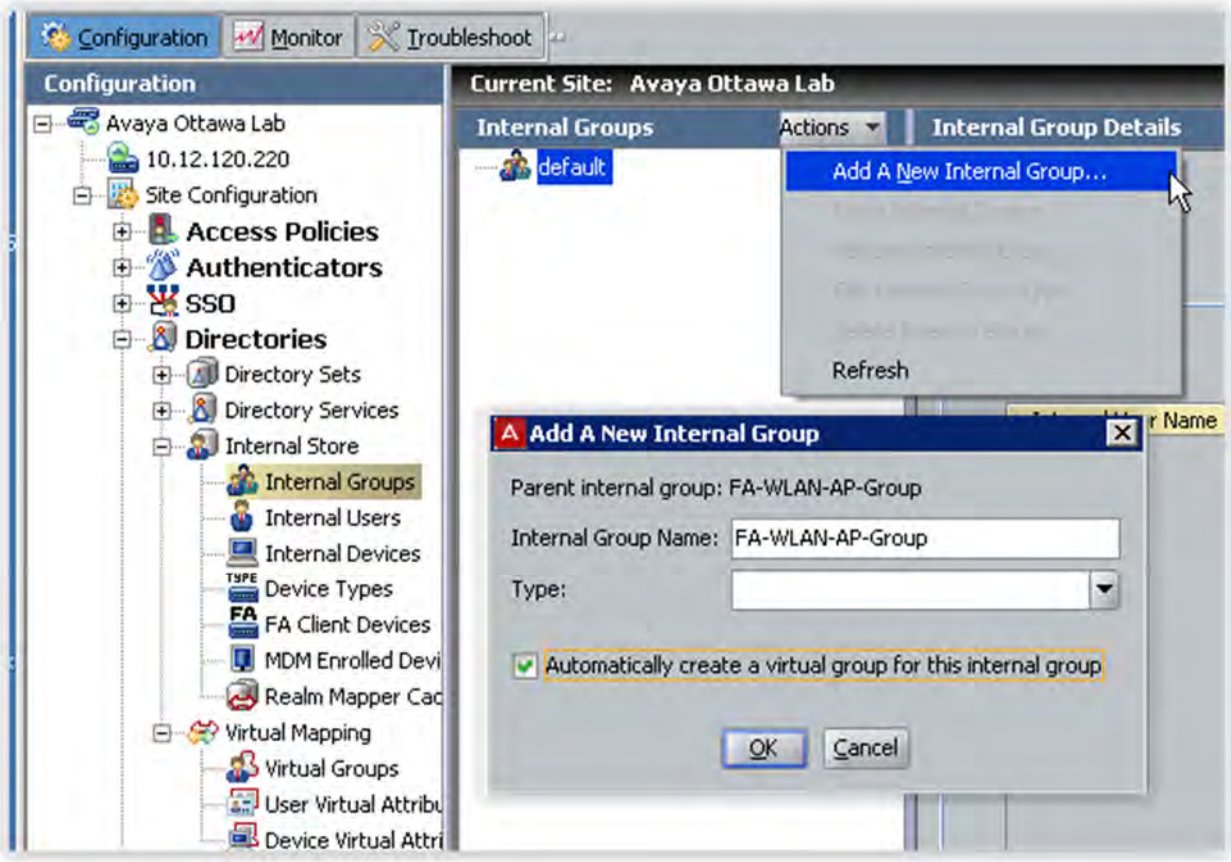

Identity Engines Setup – Set the RADIUS outbound values to be used to provision all the VLANs on the Fabric Attach Proxy switch

Please note we need to add the management VLAN, VLAN create, and all the necessary VLANs that will be used on the wireless AP's. For a Fabric Attach Proxy setup, the ISID value must be entered for all user VLANs including the management VLAN, i.e. 1100:5001100 for VLAN 1100 using ISID 5001100.

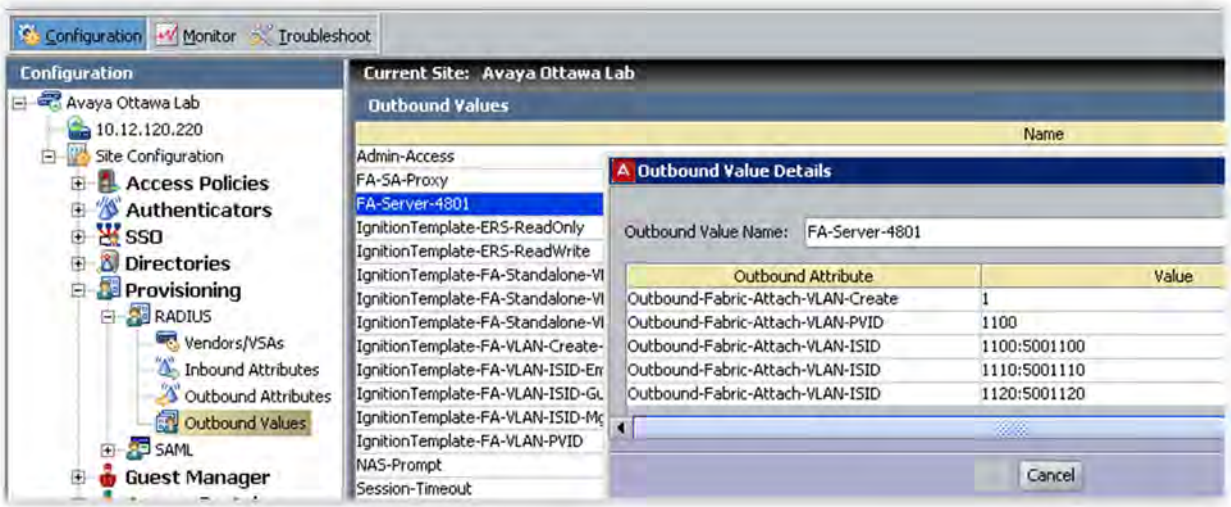

Identity Engines Setup – Set the RADIUS policy to be used to authenticate the MAC addresses for the wireless AP's

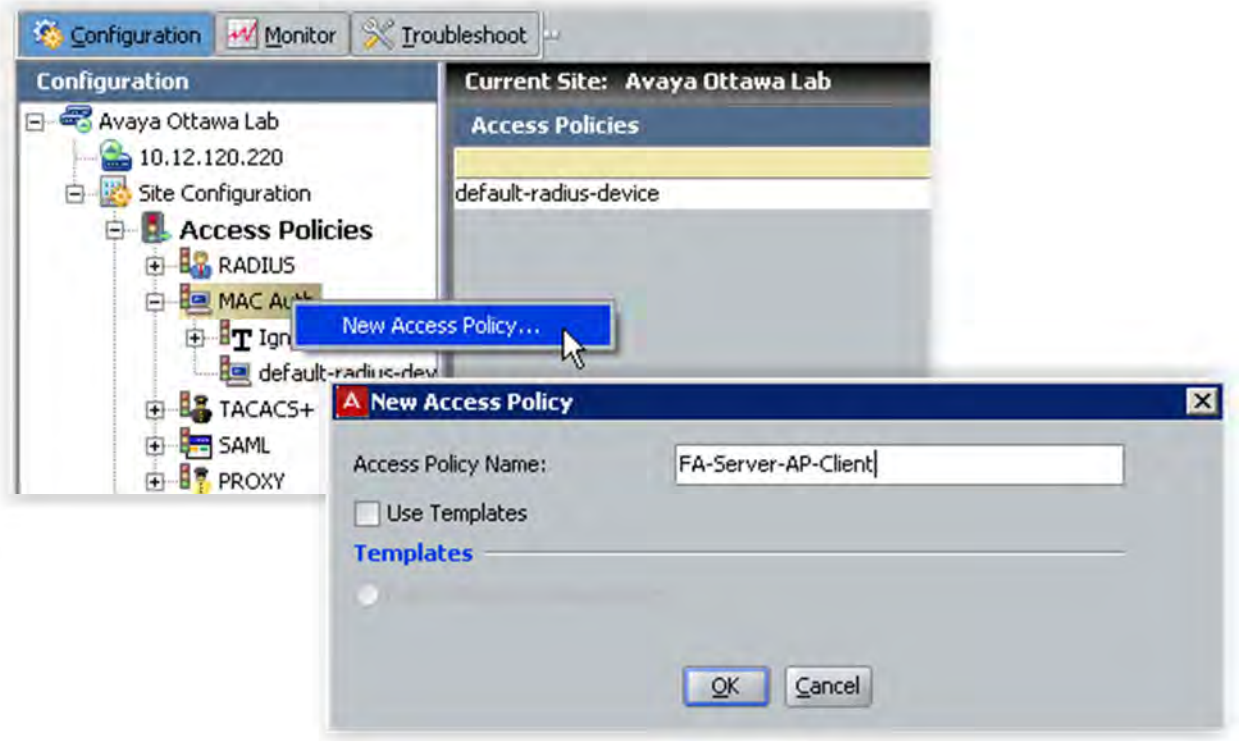

The FA Proxy switch zero touch option of *auto-port-mode-fa-client* must be enabled for the above policy to work. The policy as shown above is recommended to provide the most secure method for authenticating the WLAN 9100 AP where we are checking for a FA Client Type = 6 (FA Client – Wireless AP Type 1) in addition to the device address.

Identity Engines Setup – Add switch 4801 as an Fabric Attach Proxy authenticator

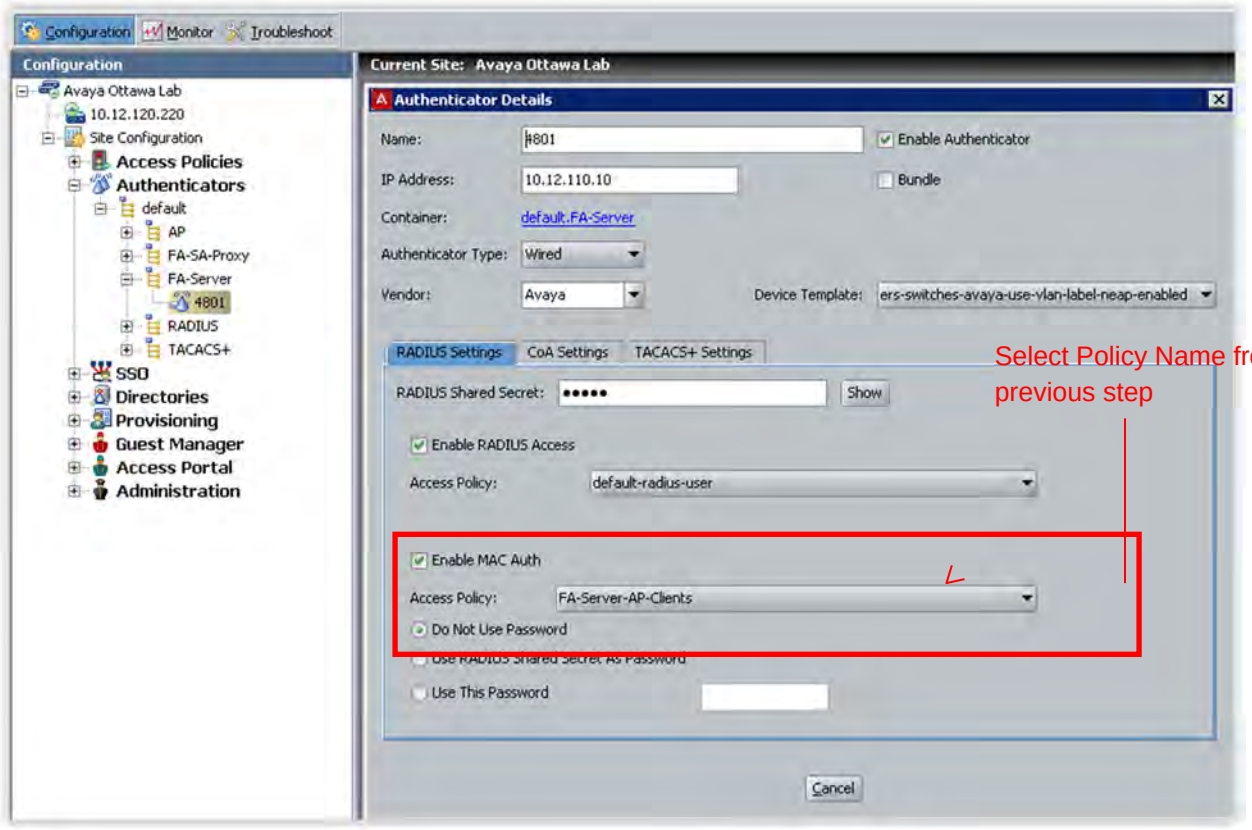

Identity Engines Setup – Adding Fabric Attach Client AP to the FA-WLAN-AP-Group

The first time the Fabric Attach client AP is connected to the ERS 4800 switch, it will fail device authentication and you will need to add the AP device MAC to the internal store. To do this, via Ignition Dashboard, click on the site name and go to *RADIUS AAA Summary -> Failed tab ->* double-click the MAC entry, and add the MAC to the group container created above

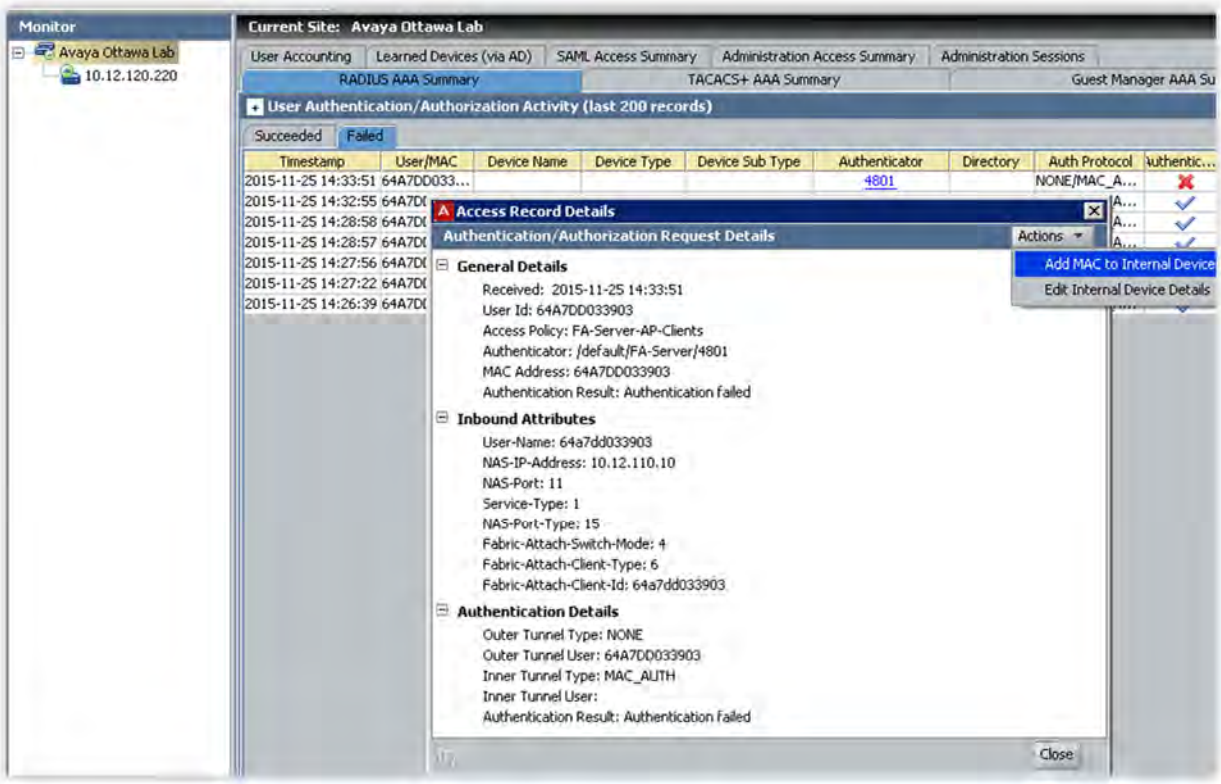

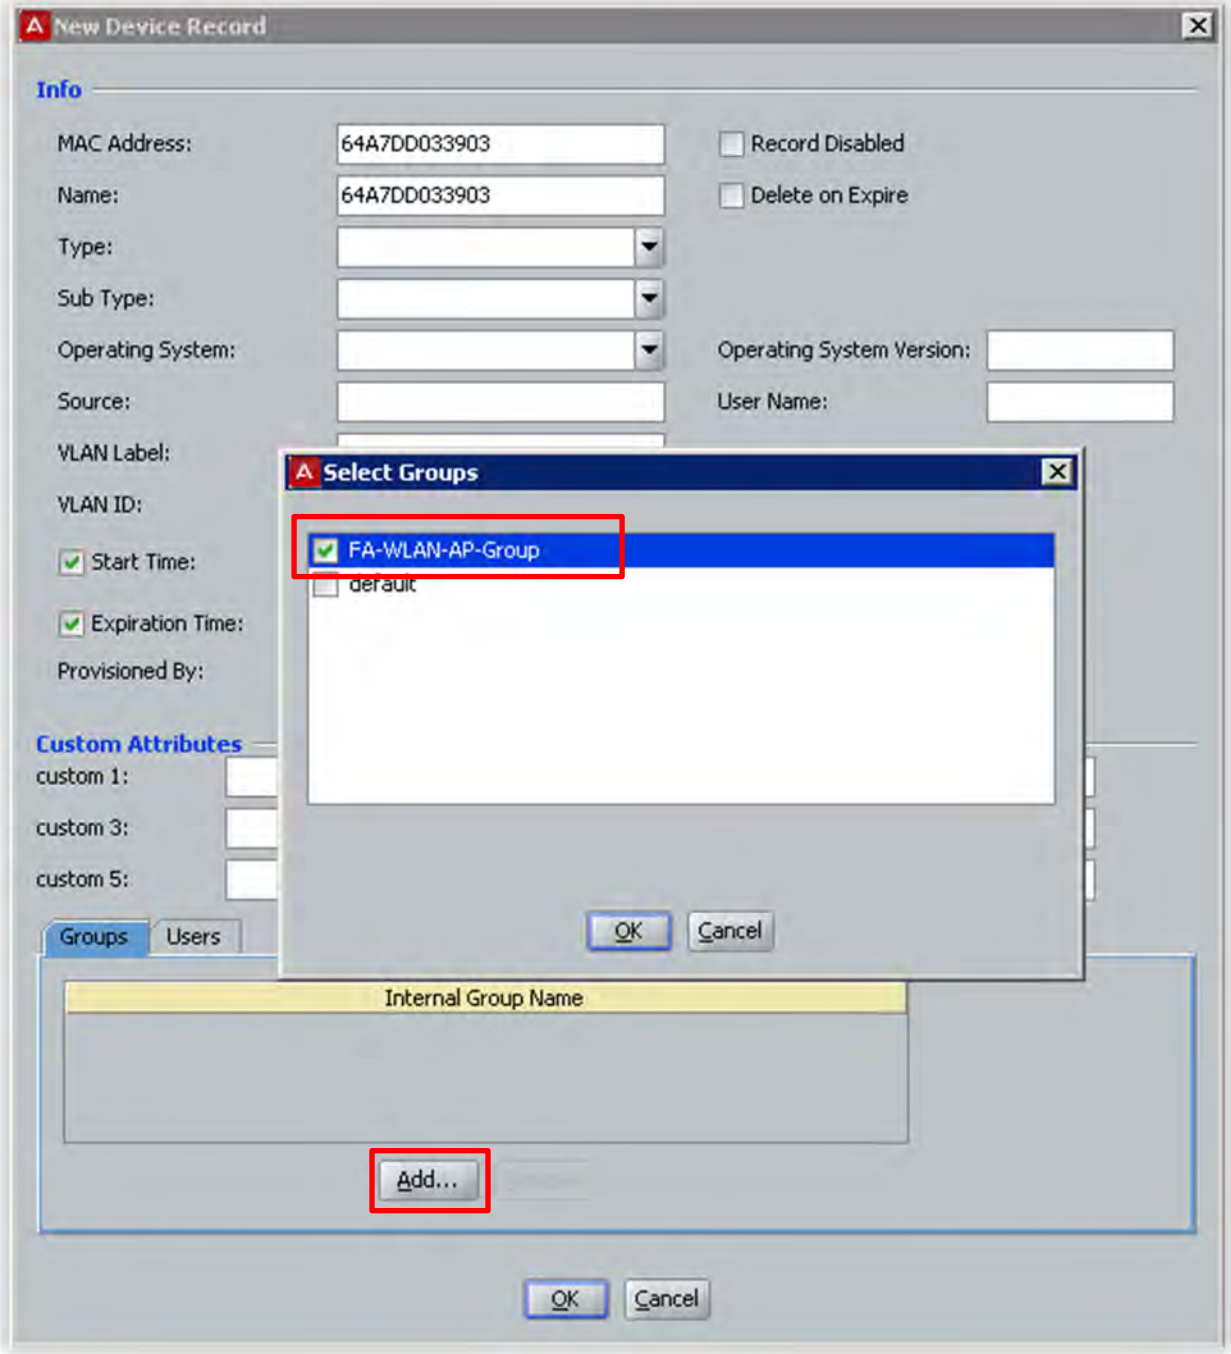

Verify Operations – FA Proxy VLAN

After IDE is configured with a policy and the AP is authenticated:

```
4801#show eapol multihost non-eap-mac status 11
Port Client MAC Address State                                          Vid  Pri
----- ------------------- ---------------------------------------------- ------ -----
11   64:A7:DD:03:39:03  Authenticated By RADIUS                        N/A  N/A
Total number of authenticated clients: 1
```

Note the default PVID is now the management PVID of 1100 and the C-VLAN PVIDs of 1110 and 1120 have been added port 11 as per the policy on Identity Engines.

```
4801#show vlan interface info 11
      Filter     Filter
     Untagged Unregistered
Port  Frames     Frames    PVID PRI    Tagging    Name
----- ----------- ---------------- ------ ----- ----------------- ----------------------
11   No        Yes          1100 0   UntagPvidOnly Port 11

4801#show vlan interface vids 11
Port VLAN VLAN Name          VLAN VLAN Name          VLAN VLAN Name
----- ------ ---------------------- ------ ---------------------- ------ ----------------------
11   1100 VLAN #1100         1110 VLAN #1110         1120 VLAN #1120
-----------------------------------------------------------------------
```

As mentioned above, the first time an AP is connected to the FA Proxy switch, it's MAC address needs to be added to the Identities Engines internal group and the AP will need to be authenticated again. This can be accomplished by either clearing the MAC address or shutting down and then bringing back up PoE power at a port level.

To clear the MAC address entries at a port level:

```
4801(config)#interface ethernet 1/11
4801(config-if)#clear mac-address-table
4801(config-if)#exit
```

To shutdown and bring back PoE power:

```
4801(config)#interface ethernet 1/11
4801(config-if)#poe poe-shutdown
4801(config-if)#no poe-shutdown
4801(config-if)#exit
```

Verify Operations – IDE Monitor

Via IDE, go to *Monitor -> <Name of your IDE Server> -> Access* to look at the record details.

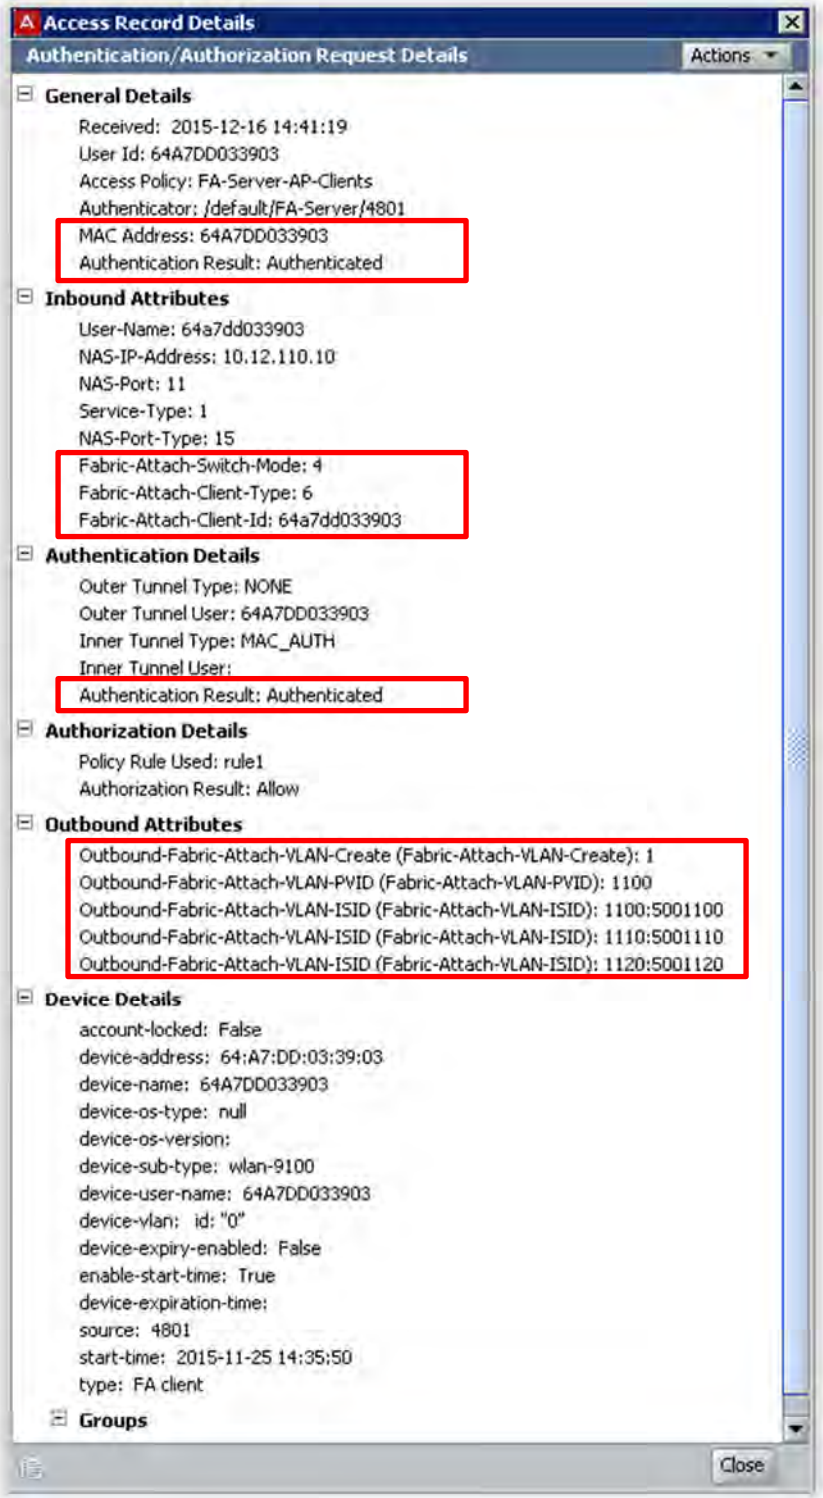

Verify Operations – FA Server

After IDE configured with a policy and the AP being authenticated, the VLAN and ISID assignment should be learned on the FA server.

```
8201:1(config)#show fa assignment
================================================================================
                          Fabric Attach Assignment Map
================================================================================
Interface  I-SID      Vlan       State      Origin
 ------------------------------------------------------------------------------
1/3        20050      50         active     proxy
1/3        700140     140        active     proxy
1/3        5001100    1100       active     proxy
1/3        5001110    1110       active     proxy
1/3        5001120    1120       active     proxy
1/3        5001150    1150       active     proxy

--------------------------------------------------------------------------------------------------------------
 6 out of 6 Total Num of fabric attach assignment mappings displayed
--------------------------------------------------------------------------------------------------------------
```

Verify Operations – Verify EAP Configuration

The FA zero touch setting of *auto-port-mode-fa-client* will automatically enable the following EAP settings:

```
4801#show eapol port 11
EAP Administrative State                          :  Enabled
Protocol Version                                  :  2
Port-mirroring on EAP ports                       :  Disabled
EAP User Based Policies                           :  Disabled
EAP User Based Policies Filter On MAC Addresses : Disabled
Port:  11
    Admin Status          :  Auto
    Authorized            :  Yes
    Admin Directions      :  Both
    Oper  Directions      :  Both
    ReAuth Enable         :  No
    ReAuth Period         :  3600
    Quiet Period : 60 Supplicant
    Timeout : 30 Server Timeout
    :  30
    Max Requests    :  2 Dynamic
    RADIUS Server :  No

4801#show eapol multihost
Allow Local Non-EAP Clients                          :  Disabled
Non-EAP RADIUS Authentication                        :  Enabled
Non-EAP AutoLearned After Single Authent (MHSA)      :  Enabled
Non-EAP DHCP Phone Authentication                    :  Disabled
EAPoL Request Packet Generation Mode                 :  Multicast
```

```
EAP RADIUS Assigned VLANs                              :  Disabled
Non-EAP RADIUS Assigned VLANs                          :  Disabled
Non-EAP RADIUS Password Attribute Format               :  MACAddr
Non-EAP User Based Policies                            :  Disabled
Non-EAP User Based Policies Filter On MAC Addresses :  Disabled
EAP Protocol                                           :  Enabled
Use Most Recent RADIUS Assigned VLAN                   :  Disabled
Non-EAP ReAuthentication                               :  Disabled
Block Different RADIUS Assigned VLAN Authentication : Disabled
Dummy ADAC Radius Requests                             :  Disabled
ADAC Non-EAP Phone Authentication                      :  Disabled
Fail Open VLAN                                         :  Disabled
Fail Open VLAN ID                                      :  1
Fail Open VLAN Continuity Mode                         :  Disabled

4801#show eapol multihost interface 11
Port:  11
    MultiHost Status                                    :  Enabled
    Total Maximum Number of Clients                     :  1
    Maximum Number of  EAP Clients                      :  1
    Maximum Number of Non-EAP Clients                   :  1
    Allow Local Non-EAP Clients                         :  Disabled
    Non-EAP RADIUS Authentication                       :  Enabled
    Non-EAP AutoLearned After Single Auth (MHSA)        :  Enabled
    Non-EAP DHCP Phone Authentication                   :  Disabled
    EAPoL Request Packet Generation Mode                :  Multicast
    EAP RADIUS Assigned VLANs                           :  Disabled
    Non-EAP RADIUS Assigned VLANs                       :  Disabled
    EAP Protocol                                        :  Enabled
    Use Most Recent RADIUS Assigned VLAN                :  Disabled
    Block Different RADIUS Assigned VLAN Authentication : Disabled
    ADAC Non-EAP Phone Authentication                   : Disabled
    MHSA No limit Non-EAP Authentication                :  Enabled
```

WAP 9132 Verify Operations - Assuming a WAP 9132 is connected to port 11 on the Proxy switch and using SSH to connect the WAP 9132

```
A17142803390B# show fabric-attach status

Fabric Attach Status
----------------------------
State                          enabled
Element Type                   FA Client - Wireless Access Point Type 1
Element State                  untagged
Management VLAN                0
Element  Gig1  MAC  Address  64:a7:dd:03:39:03
Element  Gig2  MAC  Address  64:a7:dd:03:39:04
Message Auth Key               set

A17142803390B# show fabric-attach elements

Fabric Attach Elements

Interface     Element IP     Element Type   Mgmt VLAN  Element MAC        Last Updated
------------ ------------  ------------  ---------  -----------------  ------------
gig1          10.12.110.10   FA Proxy         1100     cc:f9:54:b3:d4:00       13

A17142803390B# show lldp neighbors

LLDP List

Hostname          IP Address      Model        Interface      VLAN  Capabilities
-------------------------------------------------------------- ------ ----------------
4801             10.12.110.10               Port 11             Bridge Router
```

20.1.34.5 Fabric Attach – Changing the FA authentication key

By default, the FA WLAN 9100 client, FA proxy, and FA server all ship with a pre-defined secret FA key. The only exception is with the non-secure image for an ERS 4800 switch.

To disable FA message authentication, enter the following command.

ERS 4800

```
interface ethernet <port>
  no fa message-authentication
exit
```

VSP 4000/7200/8000

```
a) Port Level:
interface gigabitEthernet <slot/port>
  no fa message-authentication
exit

b) MLT level:
interface mlt <1-512>>
  no fa message-authentication
exit
```

To change the FA authentication key:

ERS 4800

```
a) At Port Level:
interface ethernet <port>
  fa authentication-key
  Enter authentication key (length - 1..32): <word>
  Confirm authentication key (length - 1..32): <word>
exit

b) At Global Level (applied to all ports):
fa authentication-key
Enter authentication key (length - 1..32): <word>
  Confirm authentication key (length - 1..32): <word>
```

VSP 4000/7200/8000

```
a) Port level:
interface gigabitEthernet <slot/port>
  fa authentication-key <word>
exit
b) MLT level:
interface mlt <1-512>
  fa authentication-key <word>
exit
```

20.1.35Fabric Attach Proxy Standalone

Assuming we are using an ERS 4800 for the Fabric Attach Proxy Standalone switch connected to a ERS 8800 IST core. In regards to the core switches, we need to configure C-VLANs 1100, 1110, and 1120 with tagged trunk links to the ERS 4800. In regards to the Wireless Access Points, by default, LLDP and Fabric attach is enabled.

Fabric Attach Proxy Standalone Configuration - ERS 4800

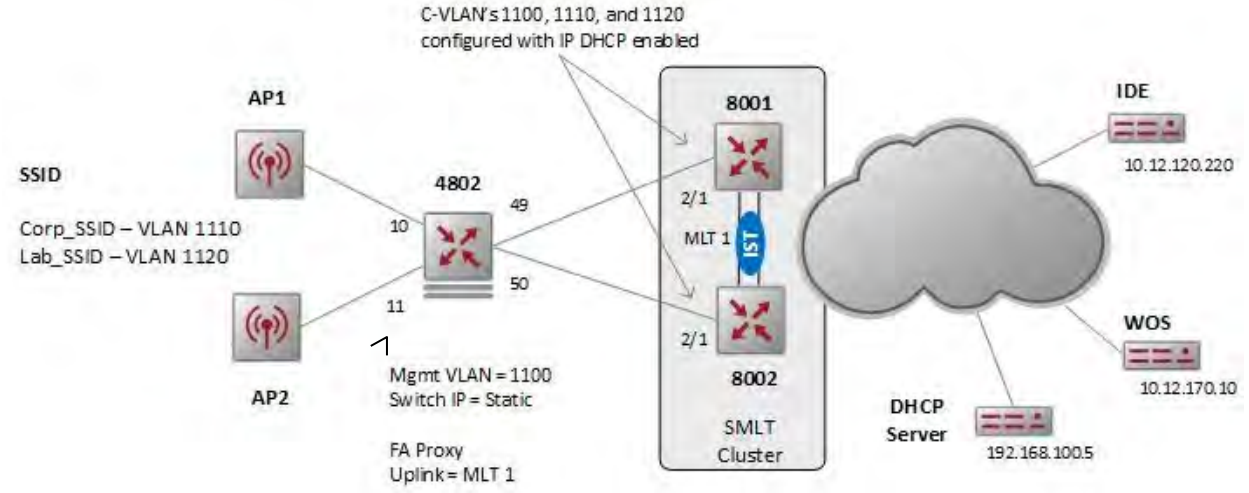


```
configure terminal
snmp-server name 4802
vlan create 1100 name mgmt_vlan1100 type port
vlan configcontrol automatic
vlan ports 49-50 tagging tagAll
vlan members 1100 49-50
vlan members remove 1 49-50
vlan mgmt 1100
mlt 1 member 49-50 learning disable
mlt 1 enable
ip address switch 10.12.110.254 netmask 255.255.255.0
ip default-gateway 10.12.110.1
radius server host 10.12.120.220 key acct-enable
  Enter key: ******
  Confirm key: ******
fa standalone-proxy
fa uplink trunk 1
fa extended-logging
fa zero-touch-option auto-port-mode-fa-client
```

The fa zero-touch-option setting of auto-port-mode-fa-client will automatically enable all the various EAP MHSA settings and enable EAP on the Fabric Attach (FA) client ports when Fabric Attached discovers a FA client. Once the FA client is discovered and authenticated against Identify Engines (IDE), IDE will overwrite the port VLAN port membership.

Fabric Attach Proxy Standalone Configuration - ERS 8800 SMLT Cluster

Assuming IP Shortcuts has been enabled on the ERS 8800 cluster, we need to add the management VLAN and all other C-VLANs used by the FA Proxy Standalone switch and FA Client. The configuration shown below simply shows the management VLAN configuration (VLAN 1100) with DHCP enabled. The other two C-VLAN configuration will be similar.

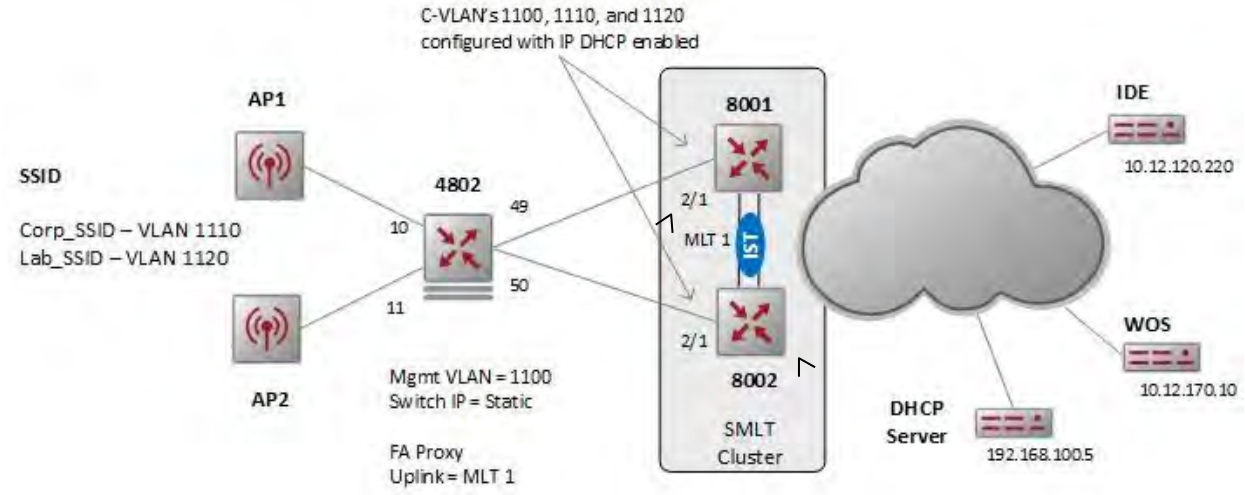


```
configure terminal
vlan ports 2/1 tagging tagAll
vlan create 1100 name "fa_mgmt_vlan1100" type
port-mstprstp 0
vlan mlt 1100 1
vlan members 1100 2/1
vlan members remove 1 2/1
interface GigabitEthernet 2/1
  smlt 132
exit
interface Vlan 1100
  ip address 10.12.110.1 255.255.255.0
  ip dhcp-relay
  ip rsmlt
  ip rsmlt holdup-timer 9999
  ip dhcp-relay fwd-path 192.168.100.5
  ip dhcp-relay fwd-path 192.168.100.5 enable
  ip dhcp-relay fwd-path 192.168.100.5 mode dhcp
exit
ip rsmlt edge-support
```

```
configure terminal
vlan ports 2/1 tagging tagAll
vlan create 1100 name "fa_mgmt_vlan1100" type
port-mstprstp 0
vlan mlt 1100 1
vlan members 1100 2/1
vlan members remove 1 2/1
interface GigabitEthernet 2/1
  smlt 132
exit
exit interface Vlan 1100
  ip address 10.12.110.2 255.255.255.0
  ip dhcp-relay
  ip rsmlt
  ip rsmlt holdup-timer 9999
  ip dhcp-relay fwd-path 192.168.100.5
  ip dhcp-relay fwd-path 192.168.100.5 enable
  ip dhcp-relay fwd-path 192.168.100.5 mode dhcp
exit
ip rsmlt edge-support
```

Identity Engines Setup – Add an internal group for the AP's

We will create a new group to be used for the IDE policy as a container for all the AP MAC address that will be authenticated when an AP connects to the Fabric Attach Proxy Standalone switch.

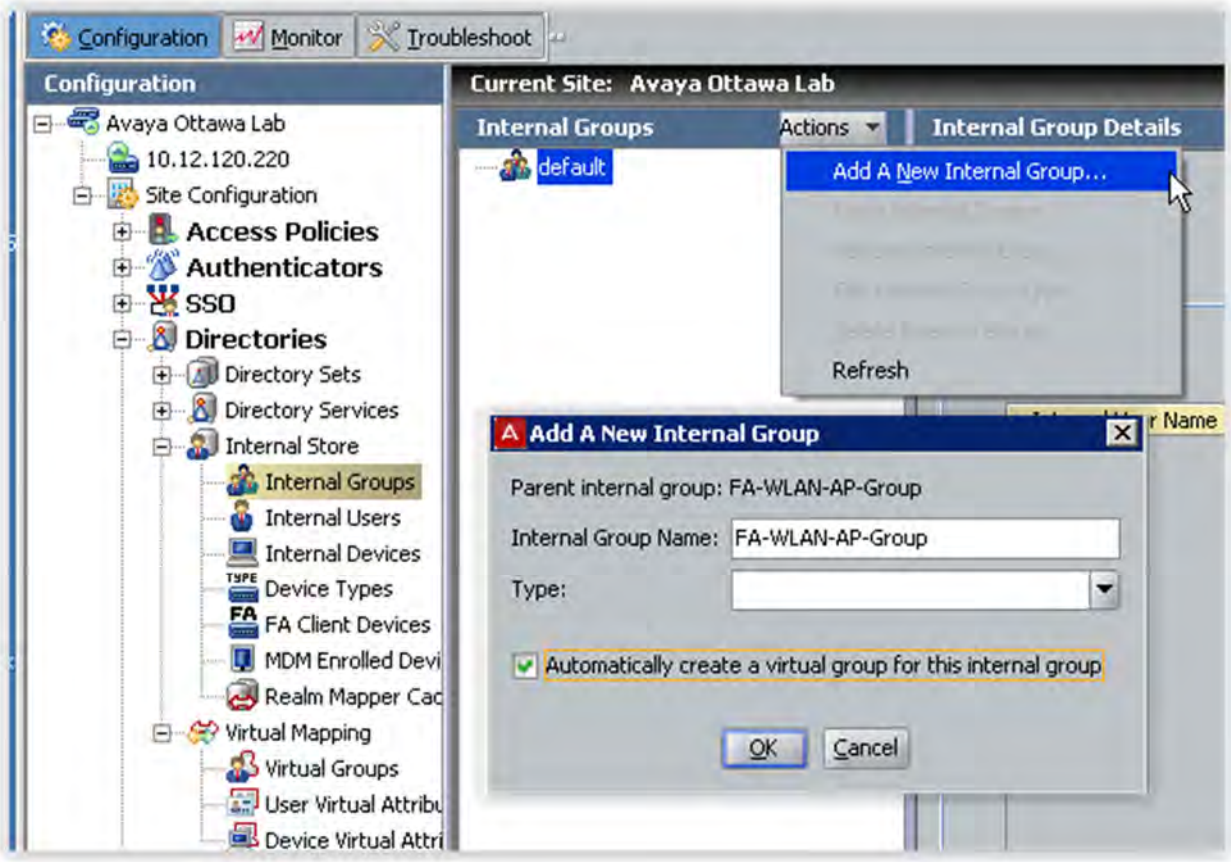

Identity Engines Setup – Set the RADIUS outbound values to be used to provision all the VLANs on the Fabric Attach Proxy Standalone switch

Please note we need to add the management VLAN, VLAN create, and all the necessary VLANs that will be used on the wireless AP's. For a Fabric Attach Standalone Proxy setup, the VLAN-ISID value must be entered with an ISID value of 0 for all user VLANs including the management VLAN, i.e. 1100:0 for VLAN 1100.

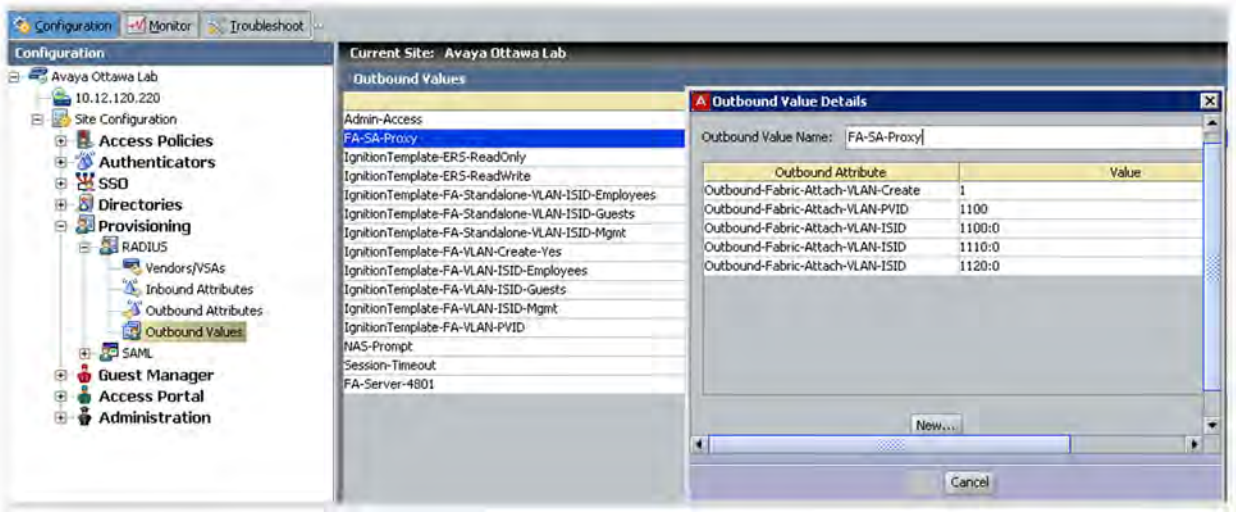

Identity Engines Setup – Set the RADIUS policy to be used to authenticate the MAC addresses for the wireless AP's

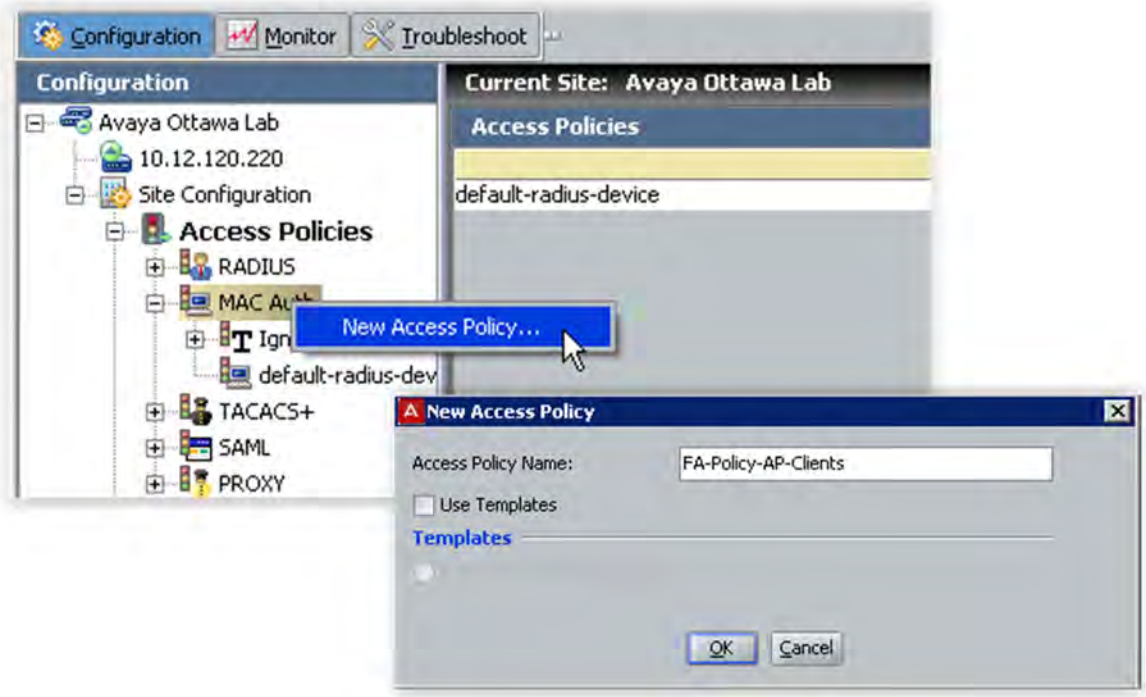

The FA Proxy switch zero touch option of *auto-port-mode-fa-client* must be enabled for the above policy to work. The policy as shown above is recommended to provide the most secure method for authenticating the WLAN 9100 AP where we are checking for a FA Client Type = 6 (FA Client – Wireless AP Type 1) in addition to the device address.

Identity Engines Setup – Add switch 4802 as an Fabric Attach Proxy Standalone authenticator

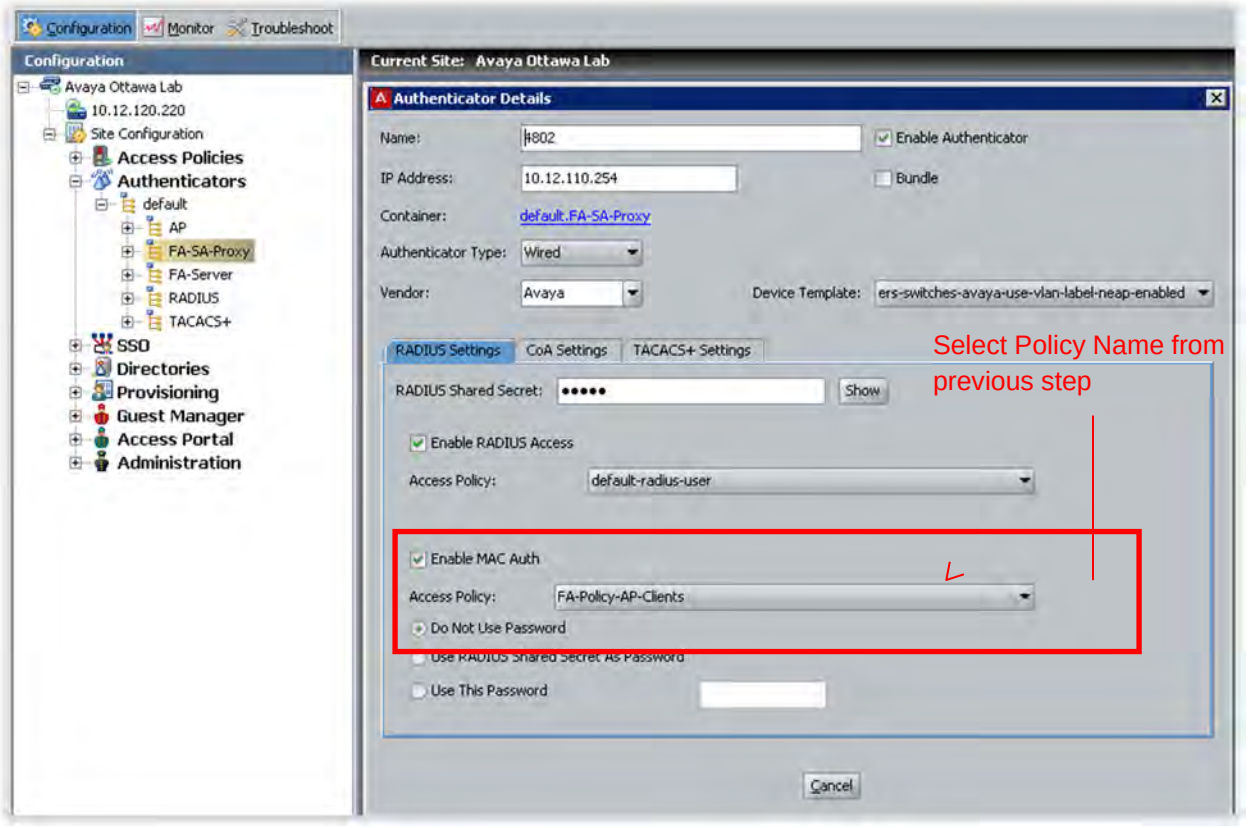

Identity Engines Setup – Adding Fabric Attach Client AP to the FA-WLAN-AP-Group

The first time the Fabric Attach client AP is connected to the ERS 4800 switch, it will fail device authentication and you will need to add the AP device MAC to the internal store. To do this, via Ignition Dashboard, click on the site name and go to *RADIUS AAA Summary -> Failed tab ->* double-click the MAC entry, and add the MAC to the group container created above

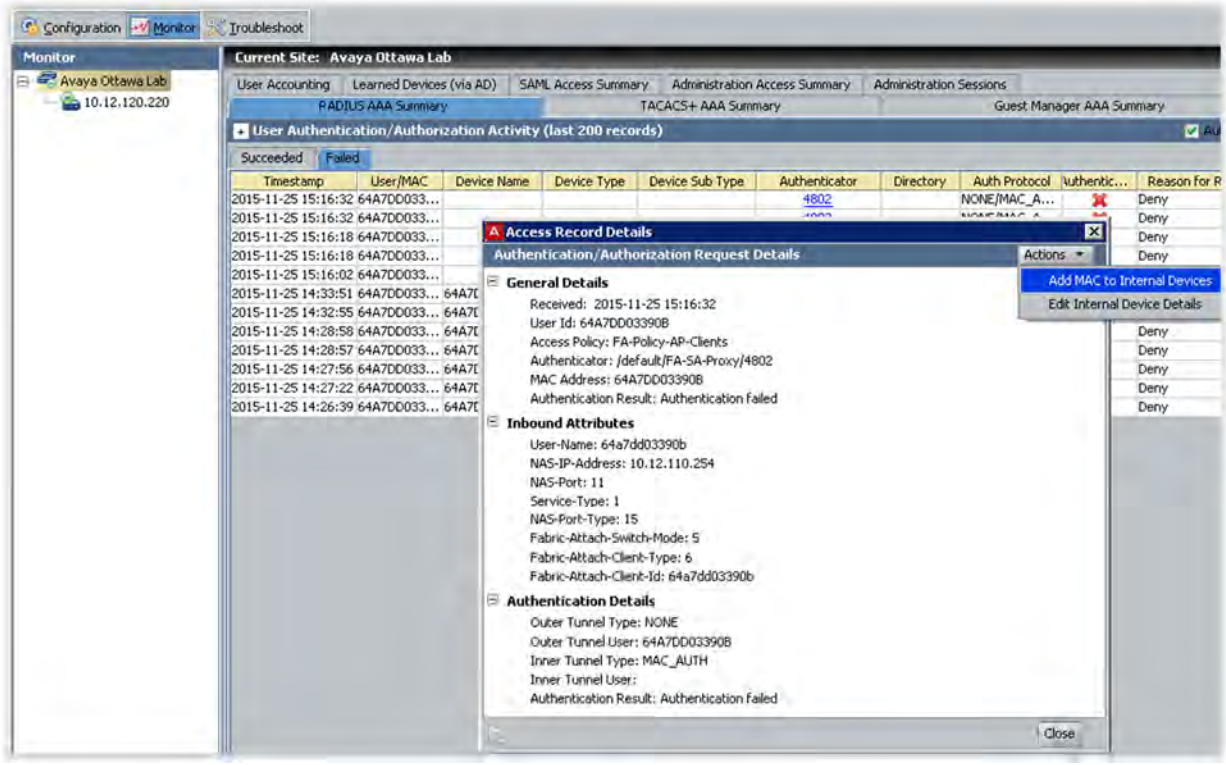

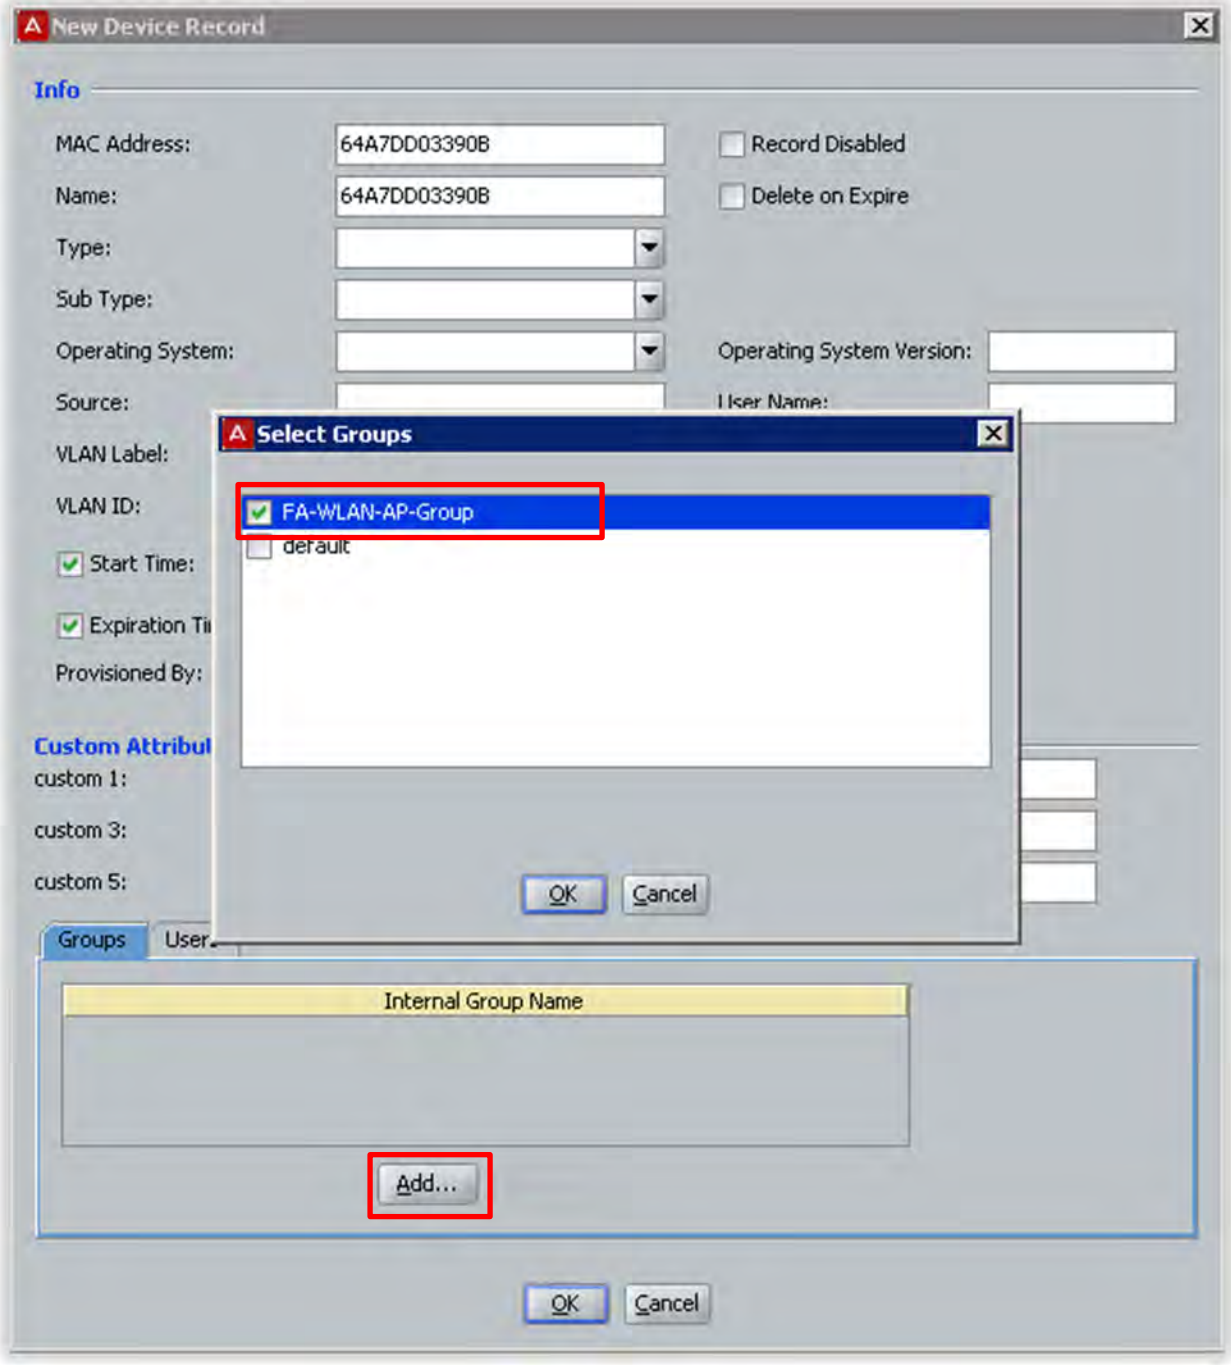

Verify Operations – FA Elements

```
4802#show fa elements

Unit/   Element    Element           Element
Port     Type      Subtype            VLAN    Auth    System ID
-------- ---------- -------------------------- ---------- ------ ----------------------------------------
1/11    Client  Wireless AP (Type 1)  0       AP  64:a7:dd:03:39:0b:00:00:00:01

Legend:
Auth - AP=Authentication Pass, AF=Authentication Fail, NA=Not Authenticated

4802#show logging sort-reverse
Type Time                Idx Src Message
---- ---------------------------- ------------------
I    00:04:15:44         168      Trap:  lldpRemTableChange Deletes = 8
|
I    00:04:14:44         163      Fabric Attach: binding activation success (ifc 11 0/1120)
I    00:04:14:44         162      Fabric Attach: binding activation success (trunk 1 0/1120)
I    00:04:14:44         161      Fabric Attach: binding activation success (ifc 11 0/1110)
I    00:04:14:44         160      Fabric Attach: binding activation success (trunk 1 0/1110)
I    00:04:14:44         159      Fabric Attach: binding activation success (ifc 11 0/1100)
I    00:04:14:44         158      Fabric Attach: binding activation success (trunk 1 0/1100)
I    00:04:14:44         38       Fabric Attach: device discovered (Authentication Pass - port
11)
|
I    00:04:06:06         153      Fabric Attach: device discovered (Auth Pass – port 11)
I    00:04:06:06         152      Trap:  lldpXMedTopologyChangeDetected,    Subtype = 7 Class = 1

4802#show lldp port 11 neighbor detail
------------------------------------------------------------------------------------------------------------
                              LLDP neighbor
------------------------------------------------------------------------------------------------------------
Port: 11Index: 2                      Time: 0 days, 00:32:22
        ChassisId: Locally assigned   A17142803390B
        PortId:    MAC address        64:a7:dd:03:39:03
        SysName:   A17142803390B
        SysCap:    rWB / rB           (Supported/Enabled)
        PortDesc:  gig1
        SysDescr:
Extreme WAP9132, 1.0GB (400MHz), Version: Extreme AOS 7.2.8 (Jul 17 2015), Build: 5
494

 PVID: 0
 PPVID Supported: none
 PPVID Enabled: none
 VLAN Name List: none

 Dot3-MAC/PHY Auto-neg: supported/enabled       OperMAUtype:  100BaseTXFD
 LinkAggr:  aggregatable/not aggregated         AggrPortID:   0

 PMD auto-neg:          10Base(T, TFD), 100Base(TX, TXFD), 1000Base(TFD)
```

```
 MED-Capabilities: CNLDI / C          (Supported/Current)
 MED-Device type:  Endpoint Class 1
-------------------------------------------------------------------------------------------------------
Sys capability: O-Other; R-Repeater; B-Bridge; W-WLAN accesspoint; r-Router;
T-Telephone; D-DOCSIS cable device; S-Station only.
Total neighbors: 1
Med Capabilities-C: N-Network Policy; L-Location Identification; I-Inventory;
S-Extended Power via MDI - PSE; D-Extended Power via MDI - PD.
```

Verify Operations – FA Proxy Stanalone VLAN

After IDE is configured with a policy and the AP is authenticated:

```
4802#show eapol multihost non-eap-mac status 11
Port Client MAC Address State                                          Vid  Pri
------ ------------------------ ------------------------------------------------------------ ------ -----
11   64:A7:DD:03:39:03  Authenticated By RADIUS                        N/A  N/A
Total number of authenticated clients: 1
```

Note the default PVID is now the management PVID of 1100 and the C-VLAN PVIDs of 1110 and 1120 have been added port 11 as per the policy on Identity Engines.

```
4802#show vlan interface info 11
      Filter     Filter
     Untagged Unregistered
Port  Frames     Frames    PVID PRI    Tagging    Name
------ ----------- ---------------- ------ ----- ------------------ ----------------------
11   No        Yes          1100 0   UntagPvidOnly Port 11

4802#show vlan interface vids 11
Port VLAN VLAN Name          VLAN VLAN Name          VLAN VLAN Name
------ ------ ---------------------- ------ ---------------------- ------ ----------------------
11   1100 VLAN #1100         1110 VLAN #1110         1120 VLAN #1120
-----------------------------------------------------------------------
```

As mentioned above, the first time an AP is connected to the FA Proxy switch, it's MAC address needs to be added to the Identities Engines internal group and the AP will need to be authenticated again. This can be accomplished by either clearing the MAC address or shutting down and then bringing back up PoE power at a port level.

To clear the MAC address entries at a port level:

```
4802(config)#interface ethernet 1/11
4802(config-if)#clear mac-address-table
4802(config-if)#exit
```

To shutdown and bring back PoE power:

```
4802(config)#interface ethernet 1/11
4802(config-if)#poe poe-shutdown
4802(config-if)#no poe-shutdown
4802(config-if)#exit
```

Verify Operations – IDE Monitor

Via IDE, go to *Monitor -> <Name of your IDE Server> -> Access* to look at the record details.

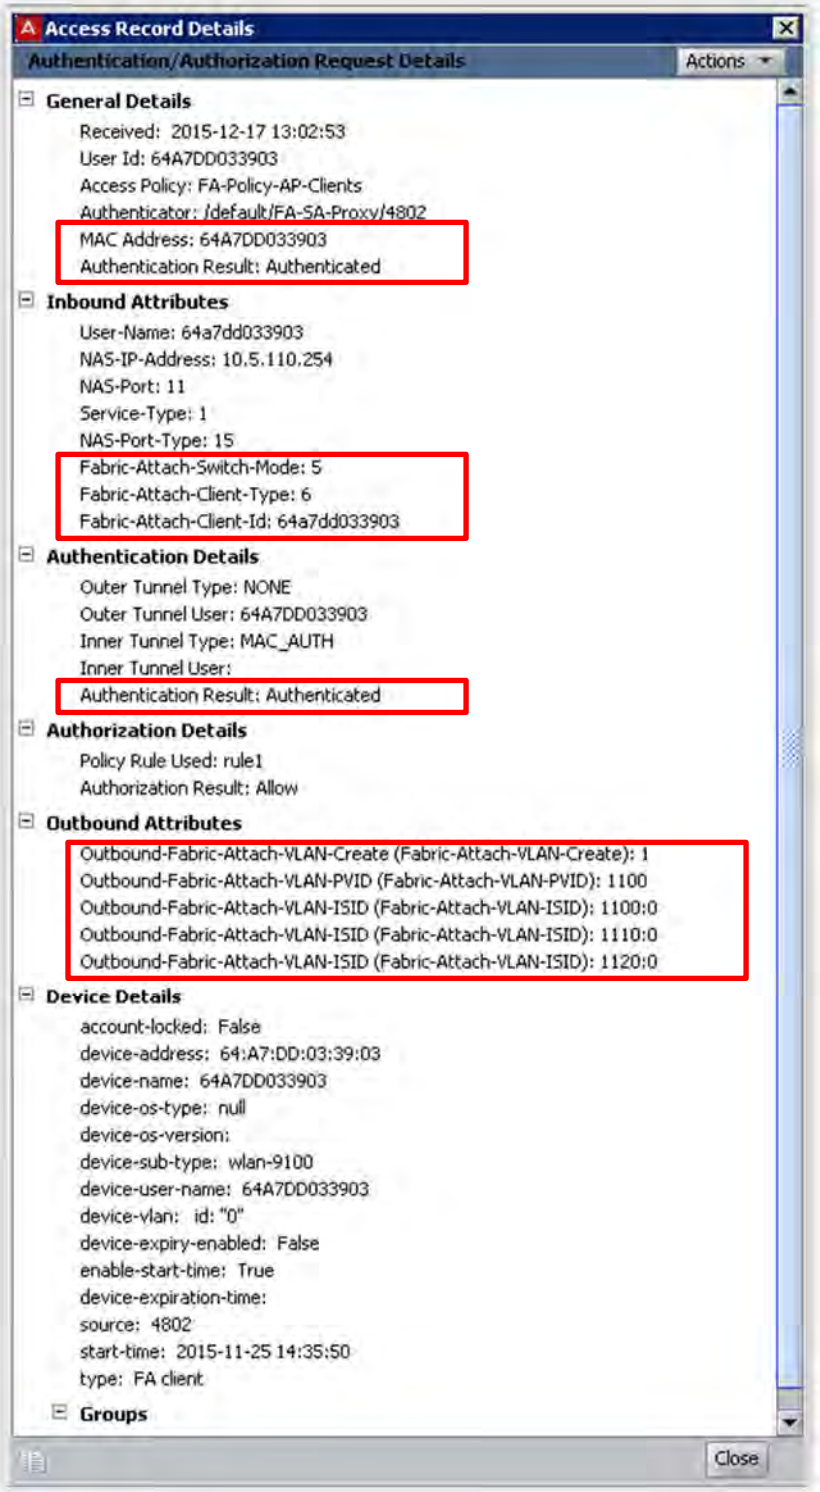

Verify Operations - EAP Setting on ERS 4800 assuming an AP has successfully authenticated on port 11

The FA zero touch setting of *auto-port-mode-fa-client* will automatically enable the following EAP settings:

```
4802#show eapol multihost non-eap-mac status
Port Client MAC Address State                                          Vid  Pri
------ ------------------------ ------------------------------------------------------------ ------ -----
11   64:A7:DD:03:39:0B  Authenticated By RADIUS                        N/A  N/A
Total number of authenticated clients: 1

4802#show eapol port 11
EAP Administrative State                        :  Enabled
Protocol Version                                :  2
Port-mirroring on EAP ports                     :  Disabled
EAP User Based Policies                         :  Disabled
EAP User Based Policies Filter On MAC Addresses : Disabled
Port:  11
    Admin Status          :  Auto
    Authorized            :  Yes
    Admin Directions      :  Both
    Oper  Directions      :  Both
    ReAuth Enable         :  No
    ReAuth Period         :  3600
    Quiet Period : 60 Supplicant
    Timeout : 30 Server Timeout
    :  30
    Max Requests    :  2 Dynamic
    RADIUS Server :  No

4802#show eapol multihost
Allow Local Non-EAP Clients                          :  Disabled
Non-EAP RADIUS Authentication                        :  Enabled
Non-EAP AutoLearned After Single Authent (MHSA)      :  Enabled
Non-EAP DHCP Phone Authentication                    :  Disabled
EAPoL Request Packet Generation Mode                 :  Multicast
EAP RADIUS Assigned VLANs                            :  Disabled
Non-EAP RADIUS Assigned VLANs                        :  Disabled
Non-EAP RADIUS Password Attribute Format             :  MACAddr
Non-EAP User Based Policies                          :  Disabled
Non-EAP User Based Policies Filter On MAC Addresses  :  Disabled
EAP Protocol                                         :  Enabled
Use Most Recent RADIUS Assigned VLAN                 :  Disabled
Non-EAP ReAuthentication                             :  Disabled
Block Different RADIUS Assigned VLAN Authentication  : Disabled
Dummy ADAC Radius Requests                           :  Disabled
ADAC Non-EAP Phone Authentication                    :  Disabled
Fail Open VLAN                                       :  Disabled
Fail Open VLAN ID                                    :  1
Fail Open VLAN Continuity Mode                       :  Disabled
```

```
4802#show eapol multihost interface 11
Port:  11
    MultiHost Status                                        :  Enabled
    Total Maximum Number of Clients                         :  1
    Maximum Number of  EAP Clients                          :  1
    Maximum Number of Non-EAP Clients                       :  1
    Allow Local Non-EAP Clients                             :  Disabled
    Non-EAP RADIUS Authentication                           :  Enabled
    Non-EAP AutoLearned After Single Auth (MHSA)            :  Enabled
    Non-EAP DHCP Phone Authentication                       :  Disabled
    EAPoL Request Packet Generation Mode                    :  Multicast
    EAP RADIUS Assigned VLANs                               :  Disabled
    Non-EAP RADIUS Assigned VLANs                           :  Disabled
    EAP Protocol                                            :  Enabled
    Use Most Recent RADIUS Assigned VLAN                    :  Disabled
    Block Different RADIUS Assigned VLAN Authentication     : Disabled
    ADAC Non-EAP Phone Authentication                       : Disabled
    MHSA No limit Non-EAP Authentication                    :  Enabled
```

WAP 9132 Verify Operations - Assuming a WAP 9132 is connected to port 11 on the Proxy Standalone switch and using SSH to connect the WAP 9132

```
A17142803390B# show fabric-attach status

Fabric Attach Status
---------------------------
State                         enabled
Element Type                  FA Client - Wireless Access Point Type 1
Element State                 untagged
Management VLAN               0
Element  Gig1  MAC  Address  64:a7:dd:03:39:03
Element  Gig2  MAC  Address  64:a7:dd:03:39:04
Message Auth Key              set

A17142803390B# show fabric-attach elements

Fabric Attach Elements

Interface     Element IP     Element Type   Mgmt VLAN  Element MAC        Last Updated
------------ ------------  ------------  ---------  -----------------  ------------
gig1          10.5.110.254   FA Proxy         1100     cc:f9:54:b4:ac:00       15

A17142803390B# show lldp neighbors

LLDP List

Hostname          IP Address      Model            Interface      VLAN  Capabilities
-------------------------------------------------- ------------ ------ ----------------
4802              10.12.110.254                    Port 11              Bridge Router
```

20.2 Using EDM

20.2.1 IS-IS and SPB Configuration

SPB and IS-IS core configuration

EDM

a) Configuration -> IS-IS -> SPBM –> Globals -> GlobalEnable = enable -> Apply

b) Configuration -> VLAN -> VLANs -> Basic -> Insert -> add Id, provide a Name if you wish and Type = spbm-bvlan -> Insert

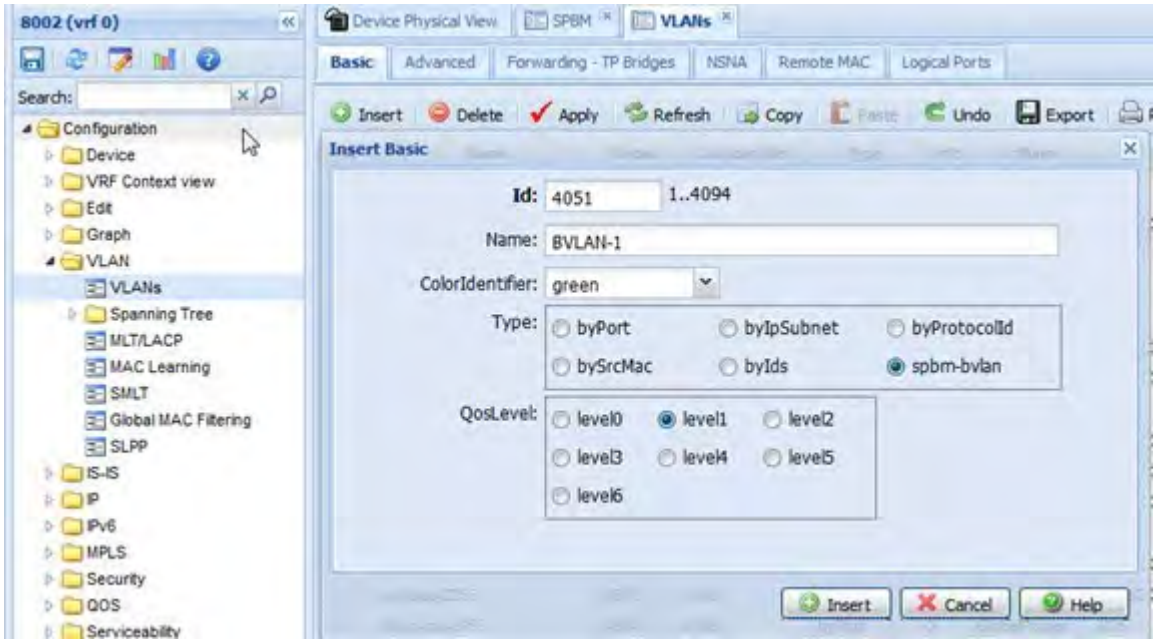

c) Configuration -> IS-IS -> SPBM –> SPBM -> Insert -> add Id, Node Nick Name, Primary VLAN (used with SMLT configurations), B-VLAN ID’s, and click on *Insert* when done (configuration shown below is for an SMLT setup to B-VLANs 4051 and 4052)

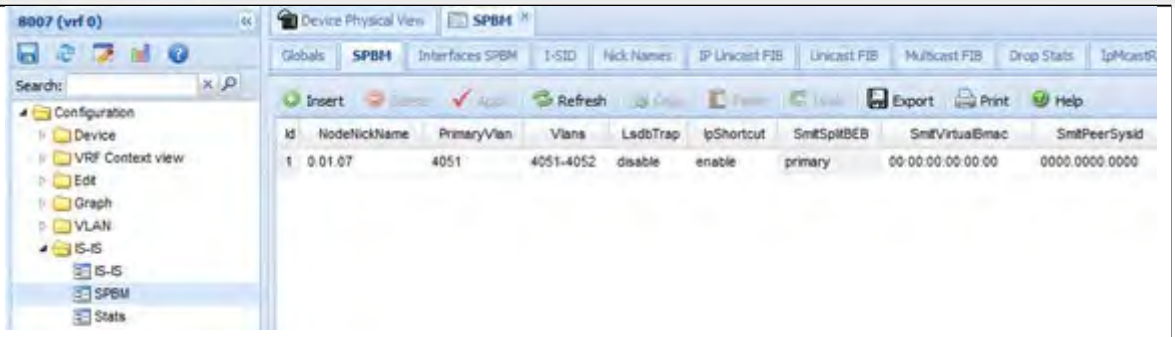

d) Configuration -> IS-IS -> IS-IS -> Interfaces -> Insert -> enter index number, select *Port* or *Mlt*, then AdminStatus = *off* (enable once SPBM is enabled in next step)

e) Configuration -> IS-IS -> IS-IS -> Interfaces -> <select index number from previous set> -> SPBM -> Insert -> enter SPBM id and state = enable -> Insert

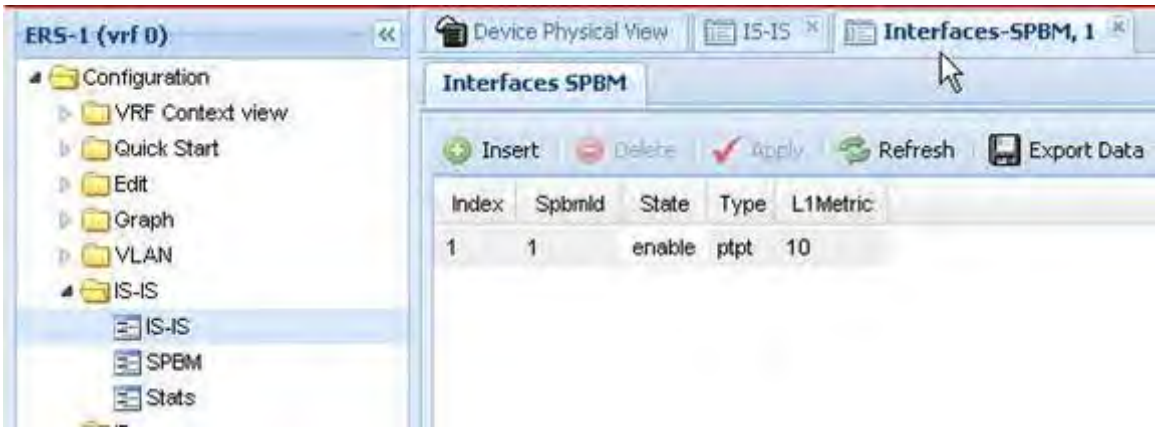

f) Configuration -> IS-IS -> IS-IS -> Interfaces -> <select index number from previous set> -> AdminState = on -> Apply

g) Configuration -> IS-IS -> IS-IS -> Manual Area -> Insert -> AreaAddr =

<manual area id in format of xx.yyyy>

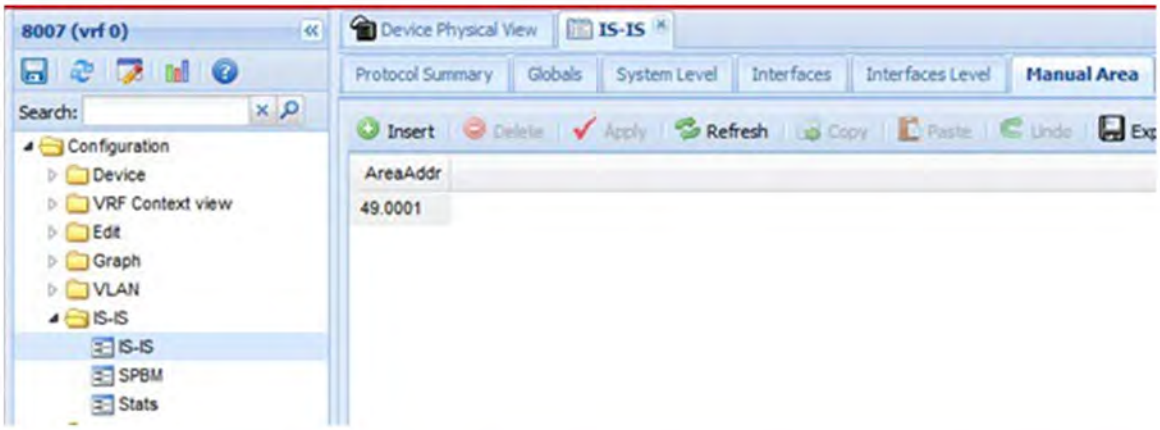

h) Configuration -> IS-IS -> IS-IS -> Globals -> AdminState = on -> Apply

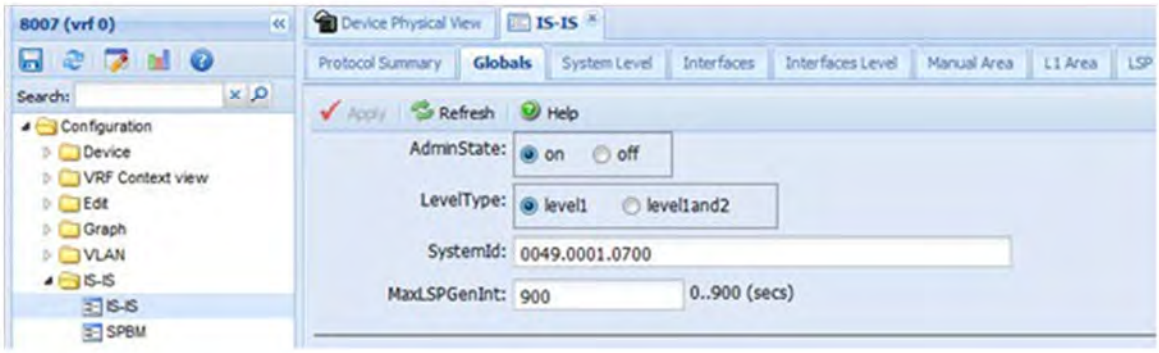

20.2.2 VSN Configuration

Extending a VLAN (L2VSN)

EDM

`Configuration -> VLAN -> VLANs –> Advanced -> Select VLAN -> ISID = <0..16777215> -> Apply`

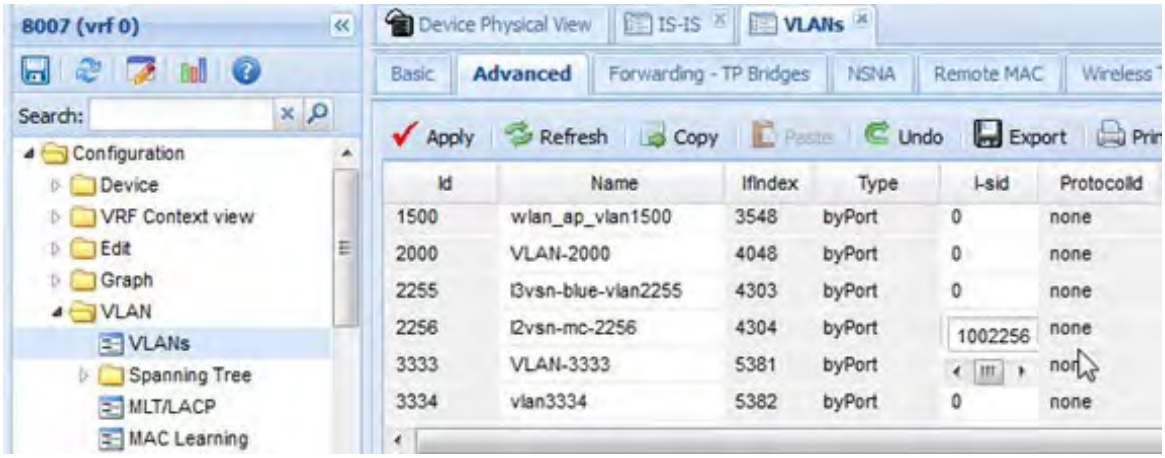

Extending a VLAN (L3VSN)

EDM

a) Configuration -> IP -> VRF –> Insert -> Enter ID, VRF name, any other options -> Insert

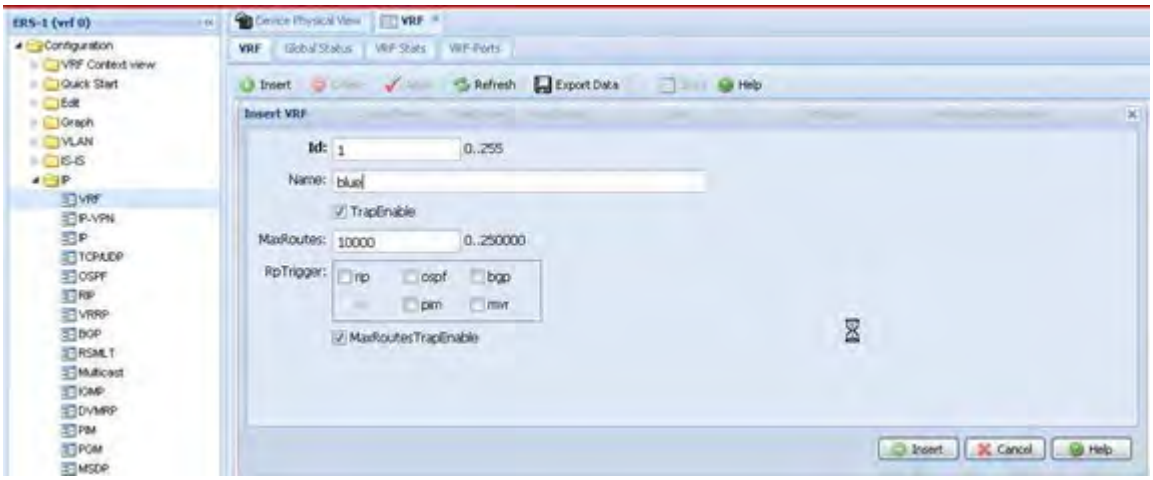

b) Configuration -> IP -> IP-VPN –> VPN -> Insert -> Select VRF ID -> Insert

c) Configuration -> IP -> IP-VPN –> VPN -> <select VrdId> -> IsidNumber = 0..16777215> -> Enable = true -> Apply

20.2.3 Connectivity Fault Management (CFM) Configuration – release 7.0 or 7.1.1.

Add Maintenance Domain (string up to 22 characters), Maintenance Association (string up to 22 characters), and maintenance end point (id from 1 to 8191). There may only be one MEP per SPB VLAN in the 7.1 release and CFM is only supported on SPB VLANs. When assigning a Maintenance Intermediate Point (MIP) level to an SPB VLAN, the value may be 0 to 7; there is only one MIP supported per SPB VLAN in the 7.1 release. It is recommended that MEP and MIP use the same level. The MEP level is configured under the Maintenance Domain of a given MEP

EDM

a) Configuration -> Edit -> Diagnostics –> CFM -> MD -> Insert -> enter Index ID, Name, Level -> Insert

b) Configuration -> Edit -> Diagnostics –> CFM -> MD -> select MD instance -> MaintenanceAssociation -> Insert -> Enter MA index number and MA name -> Insert. Repeat for each B-VLAN, i.e. Association Index = 1 for B-VLAN 4051, and Association Index = 2 for B-VLAN 4052

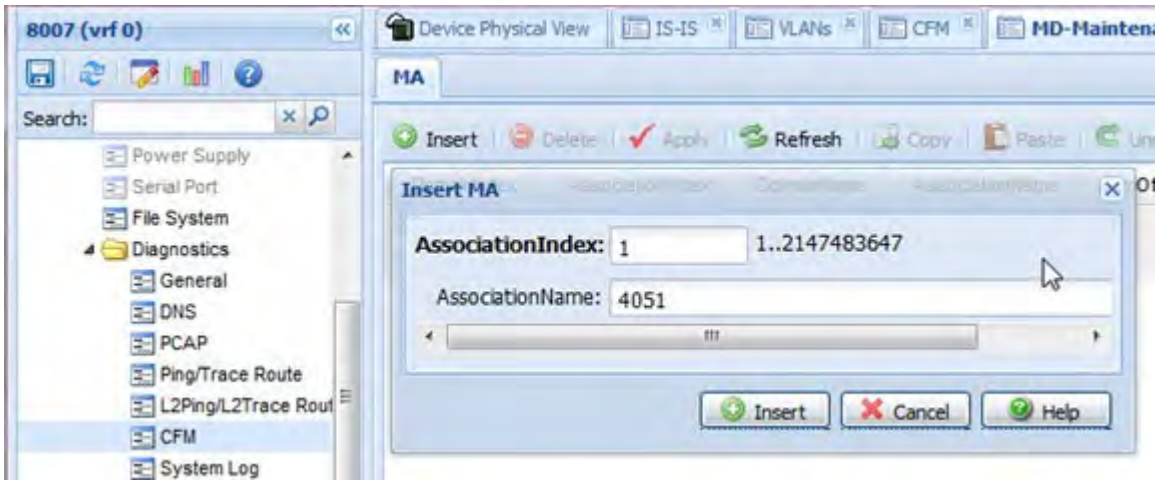

c) Configuration -> Edit -> Diagnostics –> CFM -> MD -> select MD instance -> MaintenanceAssociation -> select MA index -> MaintenanceEndPoint -> Insert -> enter id, AdminState = enable -> Insert. Repeat for each B-VLAN. Please keep note of MA Id used as this will be required for next step

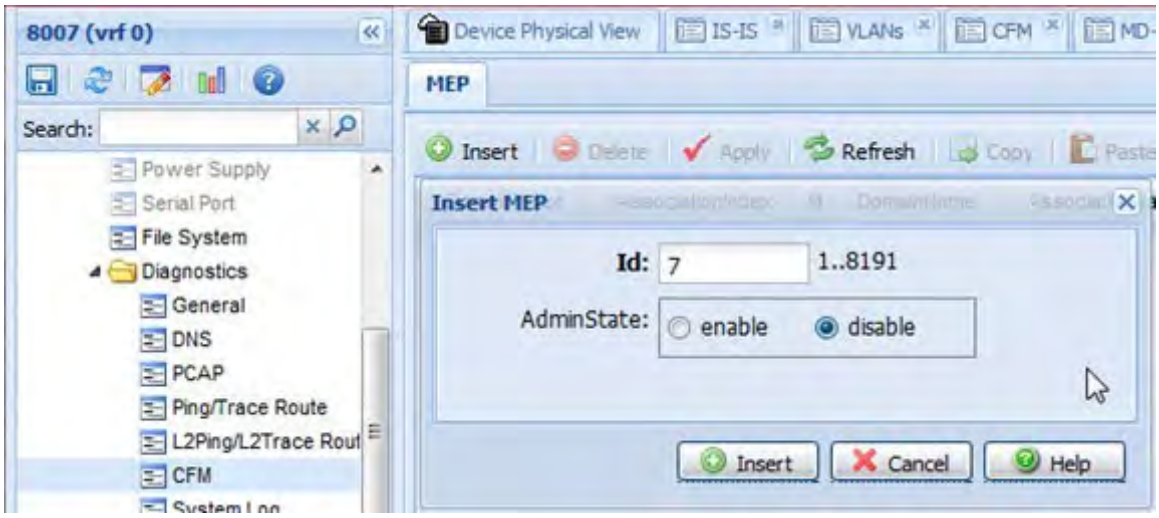

d) Configuration -> VLAN -> VLANs -> Advanced -> select B-VLAN -> Nodal -> NodalMepList = <md string>.<ma string>.<mep id>. Repeat for each B-VLAN.

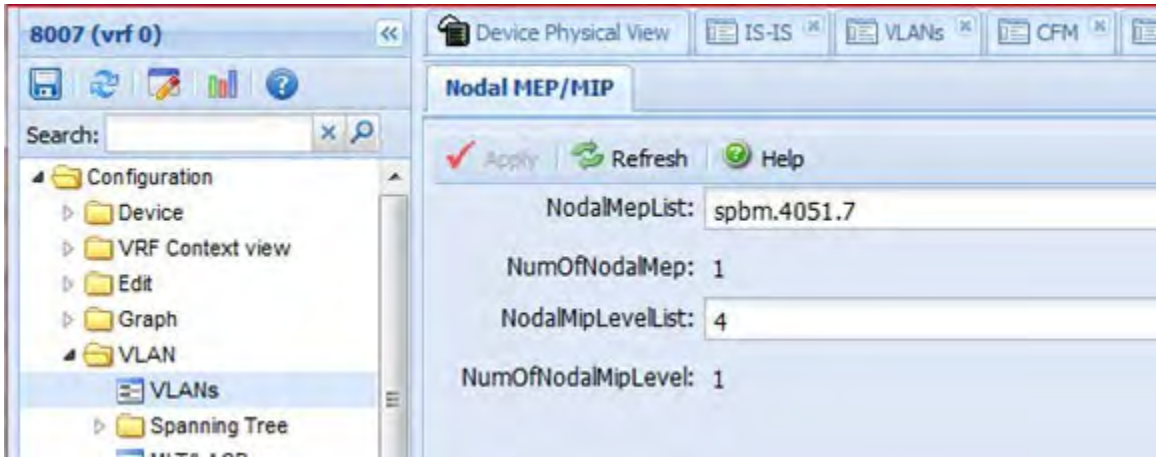

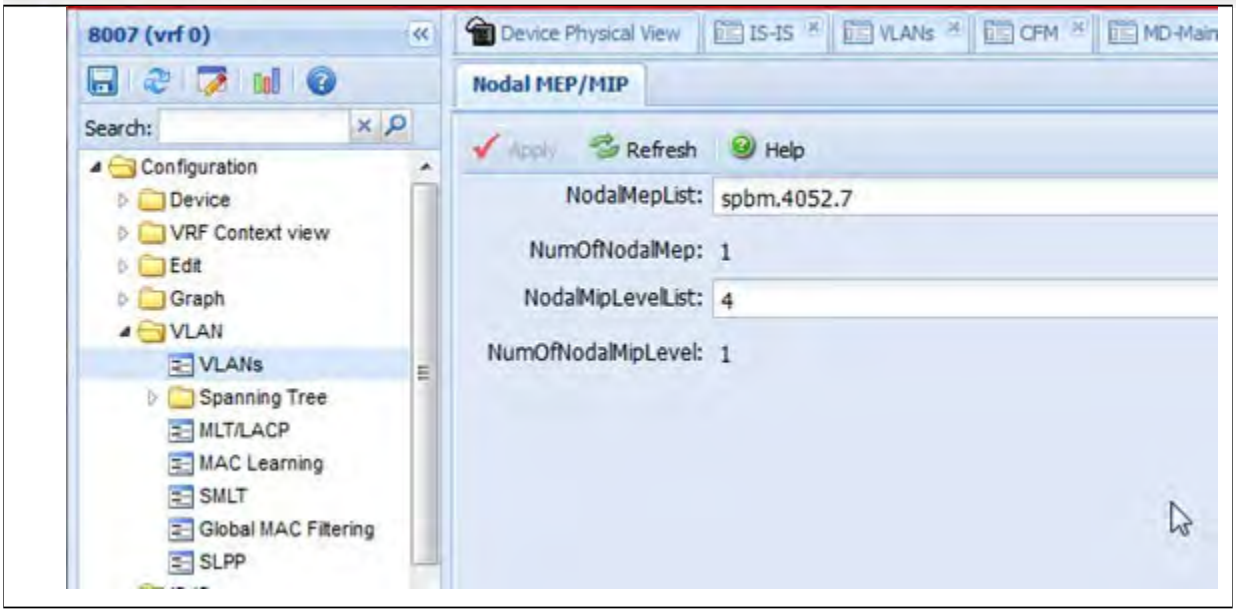

21. VLAN and ISID Restrictions using TACACS+ via Identity Engines

For security concerns, customers may wish to restrict users from only entering specific VLAN and ISID combinations. For example, for building x, an administrator wishes to only allow a local user to add VLANs 2000-2399 and only use I-SIDs 2002000-2002399. Regular expressions via Identity Engines TACACS+ Device Command Sets can be used to restrict specific ranges.

On a VSP 7000, ERS 5900, and ERS 4800 supports up 15 different TACACS+ levels are supported. For each level, we can restrict what commands are allowed and or denied and also allow regular expressions to restrict a command to a specific range. Please see the Management Access Security TCG, publication number NN48500-594 for more details on how to configure TACACS+ and setting up IDE.

The VSP 4000/7200/8200 and VSP 9000 support up to 6 levels as per the table below. Please see the Management Access Security TCG, publication number NN48500-650 for more details on how to configure TACACS+ and setting up IDE.

VSP 4000/7200/8000/9000 TACACS+ Access Levels

Access Level	Privilege Level
None	0 and 7 to 14
Read only	1
Layer 1 read write	2
Layer 2 read write	3
Layer 3 read write	4
Read write	5
Read write all	6
Read write all	15

VSP 4000/7200/8000 Enhanced Security TACACS+ Attributes

Access Level	VSA Attribute 26 – Vendor Identifier 1584 Type 192 value
None-Access	0, 4, 5, 7 to 14
Auditor	1
Security	2
Operator	3
Privilege	N/A – Not allowed by TACACS+
Admin	6
Admin	15

21.1TACACS+ Switch Configuration

VSP 7000, ERS 4800/5900 TACACS+ Configuration

```
tacacs server host <primary TACACS+ ip address> key
  Enter key: *****
  Confirm key: *****
tacacs authorization enable
tacacs accounting enable
tacacs authorization level all
cli password telnet tacacs
cli password serial tacacs
```

Please note SNMP and WEB access will be disabled one TACACS+ is enabled. Both options can be re-enabled again after initially enabled TACACS+.

VSP 4000/7200/8000/9000 TACACS+ Configuration

Option 1: without a source IP address defined:

```
tacacs server host <primary TACACS+ ip address> key <word, 0-128>
tacacs protocol enable
tacacs accounting enable cli
tacacs authorization enable
tacacs authorization level all
```

Option 2: without a source IP address defined, i.e. loopback interface:

```
interface loopback <1-256>
  ip address <1-256> <ip address/mask>
exit
tacacs server host <primary TACACS+ ip address> key <word, 0-128> source <ip address>
source-ip-interface enable
tacacs protocol enable
tacacs accounting enable cli
tacacs authorization enable
tacacs authorization level all
```

21.2TACACS+ Configuration – Identity Engines

Assuming we wish to restrict user123 to only allowing the following

- VLAN 2000-2399
- ISID's 2002000-2002399
- Deny all other VLAN and ISID combinations

IDE - Add a new device command set by going to *Configuration -> Site Configuration -> Access Policies ->TACACS+ -> Device Command Sets* and click on *New*

Via the *New Device Command Set* window, enter a name (***level5_set1*** as used in this example) and click on Add for each ACLI command set:

- For all the normal commands, via the *Device Command* window, select *Simple Command using Keywords and Arguments* and *Allow*
- For the command with ranges, via the *Device Command* window, select *Allow* first via the *Simple Command using Keywords and Arguments* tab and then click on the *Advanced Command Matching the Regular Expression* tab to add the regular expression

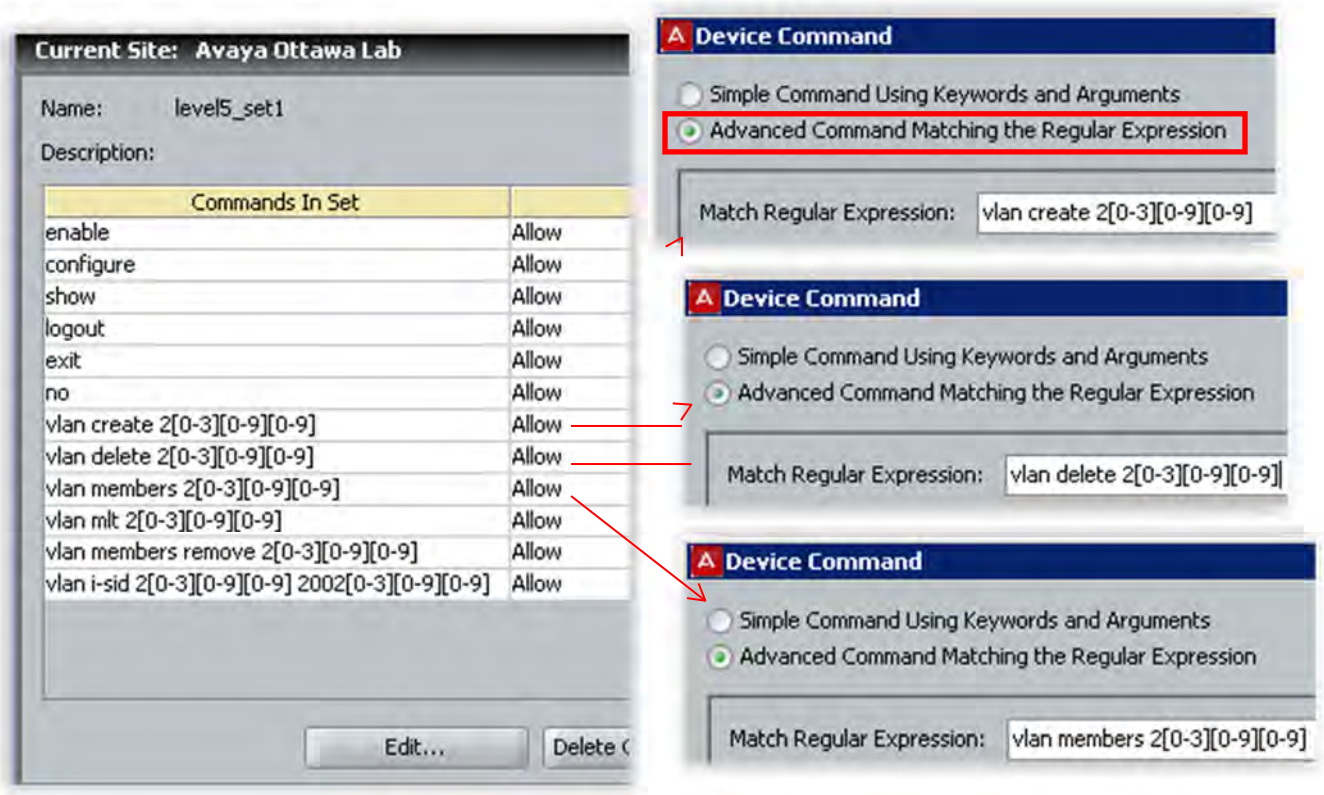

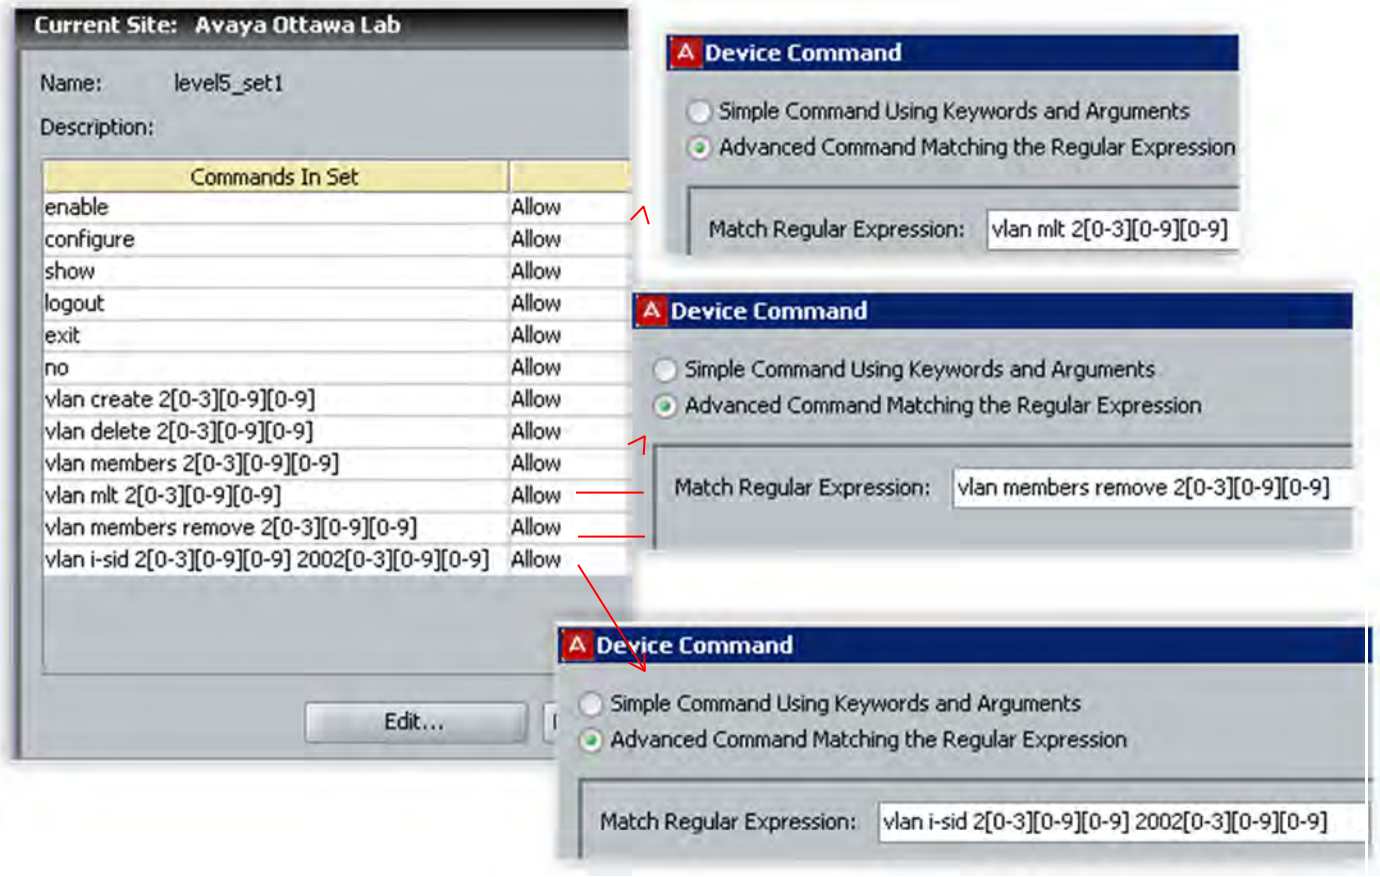

IDE – For the TACACS+ Policy, make sure you select the above device command set for the appropriate user via *Configuration -> Site Configuration -> Access Policies ->TACACS+ -> <policy name>*

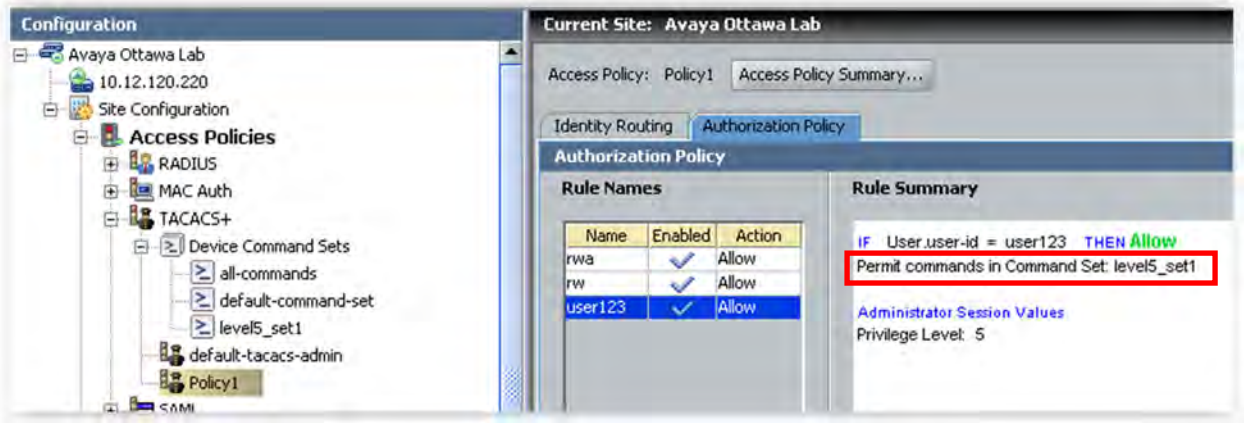

22. Configuration Examples

22.1 SPB – Core Setup

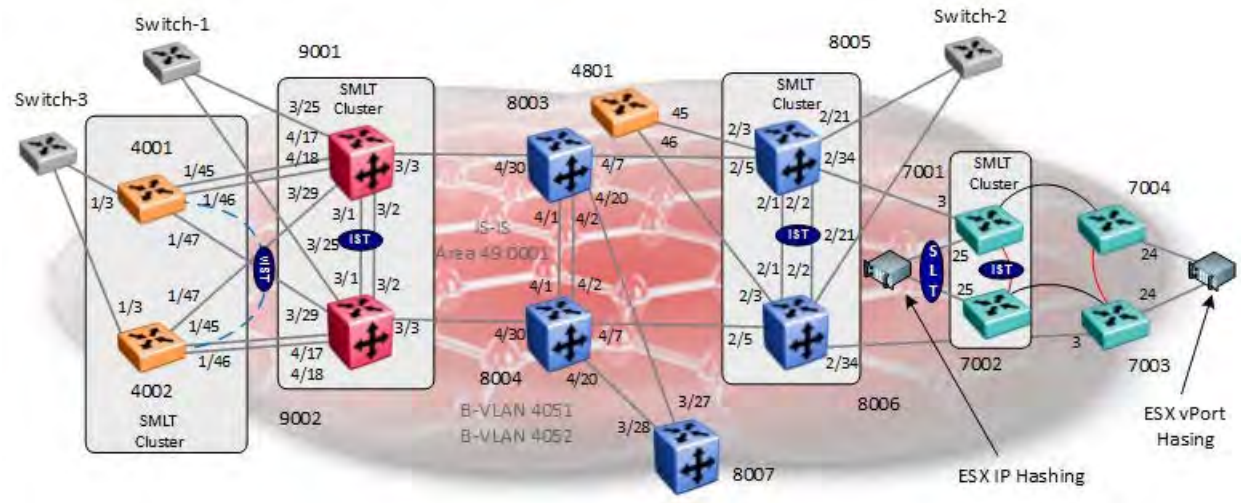

For this configuration example, we will show how to provision SPB on the following platforms:

- Common SBP Settings

Switch	Parameter	Value
All switches	B-VLANs	4051, 4052 where 4051 is the primary B-VLAN
	VLAN Names	BVLAN-1 and BVLAN-2
	IS-IS Area	49.0001
	IS-IS	Enable
	SPBM	Enable, using instance 1

- Unique SPB Settings

Platform	System Name	Nick Name	CFM MEPID
VSP 4000	400x	0.40.0x	40x
VSP 7000	700x	0.70.0x	700x
VSP 9000	900x	0.90.0x	90x
ERS 4800	480x	0.48.0x	480x
ERS 8000	800x	0.80.0x	80x

- SMLT IST Settings

SMLT Cluster	VLAN	VLAN Members	Subnet	VLACP
9001 & 9002	2	3/1 & 3/2	10.5.2.0/30	Yes – Long timeout
8005 & 8006	2	2/1 & 2/2	10.2.1.0/30	Yes – Long timeout
7001 & 7002	2	38-39	10.70.2.0/30	No

- SMLT vIST Settings

SMLT Cluster	VLAN	Subnet
4001 & 4002	2	10.4.2.0/30

For compatibility between the VSP 9000 and VSP 4000 with the ERS 8800 and VSP 7000, it is recommended to change the Spanning Tree mode to MSTP on the ERS 8800 and VSP 7000. This allows, for example, COM's VLAN Wizard to dynamically add VLANs between the ERS 8800 and VSP 9000 as the Wizard adds VLANs by Spanning Tree instance.

22.1.1 Configuration

For this configuration example, all switches are provisioned using ACLI which is the default setting on all switches with the exception of the ERS 8000.

22.1.1.1 Configuration Mode

ACLI

```
config terminal
```

The ERS 8800 supports either CLI or ACLI. The VSP 4000, 7000, ERS 4800, and VSP 9000 only support ACLI. On an ERS 8800 switch, to change from CLI to ACLI, enter the CLI command *config boot flags acli true* and *save boot*. Prior to making the change, you should convert the configuration file to ACLI by entering the command *save config file /flash/config_acli.cfg backup /flash/config.cfg mode acli*. To change from ACLI to CLI, enter the ACLI command *no boot config flags acli* and *save boot*. Prior to making the change, you should convert the configuration file to CLI by entering the command *save config file /flash/config_cli.cfg backup /flash/config.cfg mode cli*.

22.1.1.2 Auto Save

On the VSP 7000, ERS 5900, and ERS 4000 platforms, auto-save the configuration is enabled by default. If you wish, you can disable this feature and then manually save the configuration each time you make a change.

ACLI

```
no autosave enable
```

To save the configuration, use either of these commands

```
write memory
```

```
save configuration
```

```
copy config nvram
```

22.1.1.3 VSP 7000 – Rear Port Mode

Switch	Parameter	Value
Rear Port		
7001, 7002, 7003, 7004	Rear port mode	Enabled & SPB

For this example, the VSP 7000 is configured in Fabric Interconnect (FI) mode. Hence, we will change the rear-port mode to SPB.

Enable rear-port mode on switches 7001, 7002, 7003, and 7004

7001, 7002, 7003 & 7004:

```
rear-port mode enable spb
Enabling rear port mode will disable Fabric Interconnect Stack operation.
Switch configuration will be reset to partial-defaults. Continue(yes/no)?y
--------------------------------------------------------------------------------
show rear-port mode
Rear Port Mode:                Enabled SPB (Loopback Port Reserved)
Rear Port Operational State:   Operational SPB (Loopback Port Reserved)
```

22.1.1.4 Option: Change Spanning Tree mode to MSTP

For the ERS 8800, ERS 4800, and VSP 7000, we will change the Spanning Tree mode to MSTP. This is the default setting on the VSP 4000 and VSP 9000. When using tools such as VLAN Manager in COM, it is recommended to change the Spanning Tree mode to MSTP.

VSP 7000 & ERS 4800 Option – change spanning mode to MSTP on ERS 8000, VSP 7000, and ERS 4800 switches

4801, 7001, 7002, 7003 & 7004:

```
spanning-tree mode mst
New operational mode MSTP will take effect upon reset
7024XLS(config)#boot
Reboot the unit(s) (y/n) ? yRebooting . . .
show spanning-tree mode
Current STP Operation Mode: MSTP
```

ERS 8800 Option – change spanning mode to MSTP on switches 8003, 8004 ,8005 ,8006, and 8007

8003, 8004, 8005, 8006 & 8007:

```
ERS-8606:5(config)# boot config flags spanning-tree-mode mstp
Warning: Please save boot configuration and reboot the switch
         for this to take effect.
Warning: Please carefully save your configuration files before
```

```
         starting configuring the switch in RSTP or MSTP mode.
         The syntax used to create VLANs in any of these new
         modes is NOT COMPATIBLE with the default mode (STP)
```

```
ERS-8606:5(config)#save boot
```

```
ERS-8606:5(config)#boot –y
```

22.1.1.5 System Name

VSP 4000 Switches - Configure system name

prompt <4001|4002>

VSP 7000 Switches - Configure system name

snmp-server name <7001|7002|7003|7004>

ERS 8800 Switches - Configure system name

prompt <8003|8004|8005|8006|8007>

VSP 9000 Switches - Configure system name

```
prompt <9001|9002>
```

ERS 4800 Switches - Configure system name

```
snmp-server name 4801
```

22.1.1.6 Option – Configure out-of-band management interface

As an option on the ERS 8000, VSP 7000, and VSP 9000, an out-of-band management interface can be configured.

VSP 7000 Switches – Add out-of-band configuration

```
ip mgmt address <switch|stack> <ip address> netmask <subnet mask>
Either add a default gateway or static route(s)
ip mgmt default-gateway <gateway IP>
or
ip mgmt route <destination IP> <destination subnet mask> <gateway IP>
```

```
show mgmt-port status
show ip mgmt switch
show ip mgmt route
```

ERS 8000 Switches - Add out-of-band configuration

```
boot config net mgmt ip <ip address>/<subnet mask> cpu-slot <cpu slot number>
boot config net mgmt route add <ip address>/<subnet mask> <gateway IP>
save boot
```

```
show boot config net
```

Up to 5 static routes can be configured and no out-of-band default route is supported.

VSP 9000 Switches – Add out-of-band configuration

```
interface mgmtEthernet <slot|port>
ip address <ip address> <subnet mask>
exit
```

As an option, a management virtual IP address can be configured valid for both CPU's when two are used

```
sys mgmt-virtual-ip <ip address>/<subnet mask>
router vrf MgmtRouter
ip route <destination IP> <destination subnet mask> <gateway IP> weight <1-65535>
exit
```

```
show interfaces mgmtEthernet
show interfaces mgmtEthernet <config-L1|error|statistics>
show interfaces mgmtEthernet <config-L1|error|statistics> <slot|port>
show interfaces mgmtEthernet
show ip route vrf MgmtRouter
```

22.1.1.7 Enable VLACP Globally

VSP 4000, VSP 9000, and ERS 8000 Switches – Enable VLACP globally

```
vlacp enable
```

VSP 7000 & ERS 4800 Switches - Enable VLACP globally

```
vlacp enable
vlacp macaddress 180.c200.f
```

22.1.1.8 IST Configuration – SMLT Cluster switch 4001 & 4002, 9001 & 9002 and 8005 & 8006

Switch	Feature	Parameter	Value
9001, 9002 8005, 8006	IST	MLT ID	1
		VLAN	2
	VLACP (IST port members)	Timers	Long (slow)
		Time-out Scale	3
		VLACP MAC	01:80:c2:00:00:0f
		Slow periodic time	10000
9001	IST VLAN	IP address	10.5.2.1/30
		Ports	3/1,3/2
9002	IST VLAN	IP address	10.5.2.2/30
		Ports	3/1,3/2
8005	IST VLAN	IP address	10.2.1.1/30
		Ports	2/1,2/2
8006	IST VLAN	IP address	10.2.1.2/30
		Ports	2/1,2/2
4001	IST VLAN	VLAN ID	2
		IP address	10.4.2.1/30
4002	IST VLAN	VLAN ID	2
		IP address	10.4.2.12/30

For vIST on the VSP 4000, we only need to create the IST VLAN and add an IP address for now. For the VSP 7000, we will create the IST after we have provisioned SPBM.

VSP 9000 SMLT Cluster: Add IST MLT, VLAN 2 with IP address, and enable VLACP

```
9001:1(config)#vlan create 2 name "IST_vlan2" type port-mstprstp 0
9001:1(config)#mlt 1
9001:1(config)#mlt 1 name IST
9001:1(config)#mlt 1 member 3/1,3/2
9001:1(config)#mlt 1 encapsulation dot1q
9001:1(config)#vlan mlt 2 1
9001:1(config)#interface vlan 2
9001:1(config-if)#ip address 10.5.2.1 255.255.255.252
9001:1(config-if)#exit
9001:1(config)#interface mlt 1
9001:1(config-mlt)#ist peer-ip 10.5.2.2 vlan 2
9001:1(config-mlt)#ist enable
9001:1(config-mlt)#exit
9001:1(config)#interface gigabitEthernet 3/1,3/2
9001:1(config-if)#vlacp slow-periodic-time 10000 funcmac-addr 01:80:c2:00:00:0f
9001:1(config-if)#vlacp enable
9001:1(config-if)#exit
9001:1(config)#vlacp enable
```

For 9002, use the same configuration as above except for the items shown below

---

```
9002:1(config)#interface vlan 2
9002:1(config-if)#ip address 10.5.2.2 255.255.255.252
9002:1(config-if)#exit
9002:1(config)#interface mlt 1
9002:1(config-mlt)#ist peer-ip 10.5.2.1 vlan 2
9002:1(config-mlt)#ist enable
9002:1(config-mlt)#exit
```

ERS 8800 SMLT Cluster: Add IST MLT, VLAN 2 with IP address, and enable VLACP

```
8005:5(config)#vlan create 2 name "IST_VLAN" type port-mstprstp 0
8005:5(config)#mlt 1
8005:5(config)#mlt 1 name IST
8005:5(config)#mlt 1 member 2/1,2/2
8005:5(config)#mlt 1 encapsulation dot1q
8005:5(config)#vlan 2 mlt 1
8005:5(config)#interface vlan 2
8005:5(config-if)#ip create 10.2.1.1 255.255.255.0
8005:5(config-if)#exit
8005:5(config)#interface mlt 1
8005:5(config-mlt)#ist peer-ip 10.2.1.2 vlan 2
8005:5(config-mlt)#ist enable
8005:5(config-mlt)#exit
8005:5(config)#interface gigabitEthernet 2/1,2/2
8005:5(config-if)#vlacp slow-periodic-time 10000 funcmac-addr 01:80:c2:00:00:0f
8005:5(config-if)#vlacp enable
8005:5(config-if)#exit
8005:5(config)#vlacp enable
```

For 8006, use the same configuration as above except for the items shown below

---

```
8006:5(config)#interface vlan 2
8006:5(config-if)#ip create 10.2.1.2 255.255.255.0
8006:5(config-if)#exit
8006:5(config)#interface mlt 1
8006:5(config-mlt)#ist peer-ip 10.2.1.1 vlan 2
8006:5(config-mlt)#ist enable
8006:5(config-mlt)#exit
```

VSP 4000 SMLT Cluster: Add IST VLAN 2 with IP address

4001:

```
4001:1(config)#interface vlan 2
4001:1(config-if)#interface ip address 10.4.2.1 255.255.255.252
4001:1(config-if)#exit
```

4002:

```
4001:1(config)#interface vlan 2
4001:1(config-if)#interface ip address 10.4.2.2 255.255.255.252
4001:1(config-if)#exit
```

22.1.1.9 IS-IS and SPB Global Configuration

Switch	Parameter	Value
SPB		
All switches	B-VLANs	4051, 4052 where 4051 is the primary B-VLAN
	VLAN Names	BVLAN-1 and BVLAN-2
	IS-IS Area	49.0001
	IS-IS	Enable
	SPBM	Enable, using instance 1
4001	SPB Nick Name	0.40.01
	SPB System-Name	4001
	vIST Configuration	
	SMLT Peer System ID	a012.90d3.ec65 (System ID of 4002)
	ISID	2002
	vIST peer	10.4.2.2
4002	SPB Nick Name	0.40.02
	SPB System-Name	4002
	vIST Configuration	
	SMLT Peer System ID	d4ea.0e10.e465 (System ID of 4001)
	ISID	2002
	vIST peer	10.4.2.1
4801	SPB Nick Name	0.48.01
	SPB System-Name	4801
7001	SPB Nick Name	0.70.01
	SPB System-Name	7001

	SMLT Peer System-ID	3cb1.5bff.5fdf (System ID of 7002)
7002	SPB Nick Name	0.70.02
	SPB System-Name	7002
	SMLT Peer System-ID	fca8.41f6.37df (System ID of 7001)
7003	SPB Nick Name	0.70.03
	SPB System-Name	7003
7004	SPB Nick Name	0.70.04
	SPB System-Name	7004
8003	SPB Nick Name	0.80.03
	SPB System-Name	8003
8004	SPB Nick Name	0.80.04
	SPB System-Name	8004
8005	SPB Nick Name	0.80.05
	SPB System-Name	8005
	SMLT Peer System-ID	001e.1f48.f3df (System ID of 8006)
8006	SPB Nick Name	0.80.06
	SPB System-Name	8006
	SMLT Peer System-ID	0024.43b4.e3df (System ID of 8005)
8007	SPB Nick Name	0.80.07
	SPB System-Name	8007
9001	SPB Nick Name	0.90.01
	SPB System-Name	9001
	SMLT Peer System-ID	d4ea.0efd.e4ac (System ID of 9002)

9002	SPB Nick Name	0.90.01
	SPB System-Name	9002
	SMLT Peer System-ID	d4ea.0efd.e3df (System ID of 9001)

Please note for the VSP7000, it is recommended to provision SPB first prior to enabling the IST. The default PVID on all IST ports must be the primary B-VLAN ID. This will happen automatically as long as SPB is enabled prior to enabling the IST.

For the SMLT cluster switches, use the ACLI show *isis system-id command* on the peer cluster switch to get the System ID value.

SPBM Configuration – VSP 4000: Use the *show isis system-id* command on the peer cluster switch to get the smlt-peer-system-id value

```
4001:1(config)#spbm
4001:1(config)#router isis
4001:1(config-isis)#spbm 1
4001:1(config-isis)#spbm 1 nick-name 0.40.01
4001:1(config-isis)#spbm 1 b-vid 4051-4052 primary 4051
4001:1(config-isis)#spbm 1 smlt-peer-system-id a012.90d3.ec65
4001:1(config-isis)#manual-area 49.0001
4001:1(config-isis)#exit
4001:1(config)#vlan create 4051 name BVLAN-1 type spbm-bvlan
4001:1(config)#vlan create 4052 name BVLAN-2 type spbm-bvlan
4001:1(config)#vlan create 2 name "vlan2_IST" type port-mstprstp 0
4001:1(config)#vlan ISID 2 2002
4001:1(config)#virtual-ist peer-ip 10.4.2.2 vlan 2
4001:1(config)#router isis enable
--------------------------------------------------------------------------------
For 4002, use the same configuration as above except for the items shown below
--------------------------------------------------------------------------------
4002:1(config)#router isis
4002:1(config-isis)#spbm 1 nick-name 0.40.02
4002:1(config-isis)#spbm 1 smlt-peer-system-id d4ea.0e10.e465
4002:1(config-isis)#exit
4002:1(config)#virtual-ist peer-ip 10.4.2.1 vlan 2
```

SPBM Configuration – VSP 7000: Use the *show isis system-id* command on the peer cluster switch to get the smlt-peer-system-id value for the SMLT cluster switches

```
7001(config)#vlan configcontrol automatic
7001(config)#vlan create 4051 name BVLAN-1 type spbm-bvlan
7001(config)#vlan create 4052 name BVLAN-2 type spbm-bvlan
7001(config)#spbm
7001(config)#router isis
7001(config-isis)#spbm 1
7001(config-isis)#spbm 1 b-vid 4051-4052 primary 4051
7001(config-isis)#spbm 1 smlt-peer-system-id 3cb1.5bff.5fdf
7001(config-isis)#spbm 1 nick-name 0.70.01
7001(config-isis)#manual-area 49.0001
7001(config-isis)#sys-name 7001
7001(config-isis)#exit
7001(config)#router isis enable
```

For switches 7002, 7003, and 7004, use the same configuration as above except for the items shown below

---

```
7002(config-isis)#spbm 1 nick-name 0.70.02
7002(config-isis)#spbm 1 smlt-peer-system-id fca8.41f6.37df
```

```
7003(config-isis)#spbm 1 nick-name 0.70.03
```

```
7004(config-isis)#spbm 1 nick-name 0.70.04
```

SPBM Configuration – ERS 8800: Use the *show isis system-id* command on the peer cluster switch to get the smlt-peer-system-id value for the SMLT cluster switches

```
8003:5(config)#spbm
8003:5(config)#router isis
8003:5(config-isis)#spbm 1
8003:5(config-isis)#spbm 1 nick-name 0.80.03
8003:5(config-isis)#spbm 1 b-vid 4051-4052 primary 4051
8003:5(config-isis)#manual-area 49.0001
8003:5(config-isis)#exit
8003:5(config)#vlan create 4051 name BVLAN-1 type spbm-bvlan
8003:5(config)#vlan create 4052 name BVLAN-2 type spbm-bvlan
8003:5(config)#router isis enable
```

For switches 8004, 8005, 8006, and 8007, use the same configuration as above except for the items shown below. Note that bridges 8005 and 8006 also has the additional configuration to support SPB over SMLT.

```
8004:5(config-isis)#spbm 1 nick-name 0.80.04
```

```
8005:5(config-isis)#spbm 1 nick-name 0.80.05
8005:5(config-isis)#spbm 1 smlt-peer-system-id 001e.1f48.f3df
```

```
8006:5(config-isis)#spbm 1 nick-name 0.80.06
8006:5(config-isis)#spbm 1 smlt-peer-system-id 0024.43b4.e3df
```

```
8007:5(config-isis)#spbm 1 nick-name 0.80.07
```

SPBM Configuration – VSP 9000: Use the *show isis system-id* command on the peer cluster switch to get the smlt-peer-system-id value

```
9001:1(config)#spbm
9001:1(config)#router isis
9001:1(config-isis)#spbm 1
9001:1(config-isis)#spbm 1 nick-name 0.90.01
9001:1(config-isis)#spbm 1 b-vid 4051-4052 primary 4051
9001:1(config-isis)#manual-area 49.0001
9001:1(config-isis)#spbm 1 smlt-peer-system-id d4ea.0efd.e4ac
9001:1(config-isis)#exit
9001:1(config)#vlan create 4051 name BVLAN-1 type spbm-bvlan
9001:1(config)#vlan create 4052 name BVLAN-2 type spbm-bvlan spbm
9001:1(config)#router isis enable
```

For 9002, use the same configuration as above except for the items shown below

---

```
9002:1(config-isis)#spbm 1 nick-name 0.90.02
9002:1(config-isis)#spbm 1 smlt-peer-system-id d4ea.0efd.e3df
```

SPBM Configuration – ERS 4800

```
4801(config)#vlan configcontrol automatic
4801(config)#vlan create 4051 name BVLAN-1 type spbm-bvlan
4801(config)#vlan create 4052 name BVLAN-2 type spbm-bvlan
4801(config)#spbm
4801(config)#router isis
4801(config-isis)#spbm 1
4801(config-isis)#spbm 1 b-vid 4051-4052 primary 4051
4801(config-isis)#spbm 1 nick-name 0.48.01
4801(config-isis)#manual-area 49.0001
4801(config-isis)#sys-name 4801
4801(config-isis)#exit
4801(config)#router isis enable
```

SPB must be globally enabled first prior to adding SPB VLANs. If you create any SPB VLANs prior to globally enabling SPB, all SPB VLAN must be deleted. Also note that for the VSP 7000 as of release 10.2, the two B-VLANs must first be created prior to adding the B-VLANs to the SPB configuration.

22.1.1.10 IS-IS SPB Interface Configuration

Please note that Spanning Tree should be disabled on all SPB NNI interfaces that are not configured as SMLT ports. SMLT by default will disable Spanning Tree.

VSP 4000 - SPB Interface Configuration

4001 & 4002: Same configuration on both switches

```
4001:1(config)#mlt 1 enable name 9001
4001:1(config)#mlt 1 member 1/45-1/46
4001:1(config)#mlt 1 encapsulation dot1q
4001:1(config)#interface mlt 1
4001:1(config-mlt)#isis
4001:1(config-mlt)#isis spbm 1
4001:1(config-mlt)#isis enable
4001:1(config-mlt)#exit
4001:1(config)#interface gigabitEthernet 1/47
4001:1(config-if)#isis
4001:1(config-if)#isis spbm 1
4001:1(config-if)#isis enable
4001:1(config-if)#exit
4001:1(config)#interface gigabitEthernet 1/45-1/47
4001:1(config-if)#no spanning-tree mstp force-port-state enable
Disabling CIST would also disable all other MST instances.
Are you sure you want to continue (y/n) ? y
4001:1(config-if)#exit
```

VSP 7000 - SPB Interface Configuration

Rear ports – for this configuration example, we are provisioning only the FI Red (ports 38 & 39) and Black FI (ports 34, 35, & 36) rear ports. We simply just need to select one of the rear ports and enable SPB as all ports are using the same LACP LAG. On switches 7001 & 7002, we are also configuring the FI Red ports as an SMLT IST interface, hence, we will need to disable LACP as an IST interface does not support LACP – we will perform this step latter in this configuration example. Note if you select all ports, i.e. 34-39, and enable SPBM, this will work, but, you will simply get an error message stating IS-IS is already enabled on port 39, 35, and 36 – simply just ignore this error message.

7003 and 7004: Same configuration on both switches

```
7003(config)#interface ethernet 34,38
7003(config-if)#isis
7003(config-if)#isis spbm 1
7003(config-if)#isis enable
7003(config-if)#spanning-tree mstp learning disable
```

7001 & 7003: Front Ports – Same configuration on both switches

```
7001(config)#interface ethernet 3,34
7001(config-if)#isis
7001(config-if)#isis spbm 1
```

```
7001(config-if)#isis enable
7001(config-if)#spanning-tree mstp learning disable
7001(config-if)#exit
```

ERS 8800 - SPB Interface Configuration

8003 & 8004: Same configuration on both switches

```
8003:5(config)#mlt 1 enable name isis_mlt_1
8003:5(config)#mlt 1 member 4/1-4/2
8003:5(config)#mlt 1 encapsulation dot1q
8003:5(config)#interface mlt 1
8003:5(config-mlt)#isis
8003:5(config-mlt)#isis spbm 1
8003:5(config-mlt)#isis enable
8003:5(config-mlt)#exit
8003:5(config)#interface GigabitEthernet 4/1-4/2
8003:5(config-if)#no spanning-tree mstp force-port-state enable
Disabling CIST would also disable all other MST instances.
Are you sure you want to continue (y/n) ? y
8003:5(config-if)#exit
8003:5(config)#interface GigabitEthernet 4/7,4/20,4/30
8003:5(config-if)#isis
8003:5(config-if)#isis spbm 1
8003:5(config-if)#isis enable
8003:5(config-if)#no spanning-tree mstp force-port-state enable
Disabling CIST would also disable all other MST instances.
Are you sure you want to continue (y/n) ? y
8003:5(config-if)#exit
```

8005 & 8006: Same configuration on both switches. Also, as MLT 1 is used for the IST, we will change the ISIS hello interval to 1 and hello multiplier to 27.

```
8005:5(config)#interface mlt 1
8005:5(config-mlt)#isis
8005:5(config-mlt)#isis spbm 1
8005:5(config-mlt)#isis enable
8005:5(config-mlt)#isis l1-hello-interval 1
8005:5(config-mlt)#isis l1-hello-multiplier 27
8005:5(config-mlt)#exit
8005:5(config)#interface GigabitEthernet 2/5,2/34
8005:5(config-if)#isis
8005:5(config-if)#isis spbm 1
```

```
8005:5(config-if)#isis enable
8005:5(config-if)#no spanning-tree mstp  force-port-state enable
Disabling CIST would also disable all other MST instances.
Are you sure you want to continue (y/n) ? y
8005:5(config-if)#exit
8007:
8007:5(config)#interface GigabitEthernet 3/27,3/28
8007:5(config-if)#isis
8007:5(config-if)#isis spbm 1
8007:5(config-if)#isis enable
8007:5(config-if)#no spanning-tree mstp  force-port-state enable
Disabling CIST would also disable all other MST instances.
Are you sure you want to continue (y/n) ? y
8007:5(config-if)#exit
```

VSP 9000 - SPB Interface Configuration

```
9001 & 9002: Same configuration on both switches.
9001:1(config)#interface mlt 1
9001:1(config-mlt)#isis
9001:1(config-mlt)#isis spbm 1
9001:1(config-mlt)#isis enable
9001:1(config-mlt)#exit
9001:1(config)#interface gigabitEthernet 3/3,3/29
9001:1(config-if)#isis
9001:1(config-if)#isis spbm 1
9001:1(config-if)#isis enable
9001:1(config-if)#no spanning-tree mstp force-port-state enable
Disabling CIST would also disable all other MST instances.
Are you sure you want to continue (y/n) ? y
9001:1(config-if)#exit
9001:1(config)#mlt 8 enable
9001:5(config)#mlt 8 member 4/17,4/18
9001:5(config)#mlt 8 encapsulation dot1q
9001:5(config)#interface mlt 8
9001:5(config-mlt)#isis
9001:5(config-mlt)#isis spbm 1
9001:5(config-mlt)#isis enable
9001:5(config-mlt)#exit
9001:5(config)#interface gigabitEthernet 4/17,4/18
```

```
9001:5(config-if)#no shutdown
9001:5(config-if)#no spanning-tree mstp force-port-state enable
Disabling CIST would also disable all other MST instances.
Are you sure you want to continue (y/n) ? y
9001:5(config-if)#exit
```

ERS 4800 - SPB Interface Configuration

```
4801(config)#interface ethernet 45,46
4801(config-if)#isis
4801(config-if)#isis spbm 1
4801(config-if)#isis enable
4801(config-if)#spanning-tree mstp learning disable
```

22.1.1.11 Remove default VLAN from all SPB ports

Note this section only applies to the ERS 4800 and VSP 7000.

ERS 4800 - Remove default VLAN from ISIS port members

```
4801(config)#vlan members remove 1 45,46
```

VSP 7000 - Remove default VLAN from ISIS port members

```
7001 & 7003: Same configuration on both switches
7001(config)#vlan members remove 1 3,34-39
```

22.1.1.12 Other best practice items – VLACP and discard untagged frames

For added protection and faster link failure detection, it is recommended to also enable VLACP on all IS-IS ports. VLACP is already enabled on the IST port member so the rest of this configuration covers the IS-IS ports.

VSP 4000 - Interface Configuration

```
4001 & 4002: Same configuration on both switches
4001:1(config)#interface gigabitEthernet 1/45-1/47
4001:1(config-if)#untagged-frames-discard
4001:1(config-if)#vlacp fast-periodic-time 500 timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
4001:1(config-if)#vlacp enable
4001:1(config-if)#exit
```

VSP 7000 - Interface Configuration

```
7001 & 7003: Same configuration on both switches
```

```
7001(config)#vlan ports 3 filter-untagged-frame enable
7001(config)#interface ethernet 3
7001(config-if)#vlacp timeout short
7001(config-if)#vlacp timeout-scale 5
7001(config-if)#vlacp enable
7001(config-if)#exit
```

ERS 8800 - Interface Configuration

8003 & 8004: Same configuration on both switches

```
8003:5(config)#interface gigabitEthernet 4/1,4/2,4/7,4/20,4/30
8003:5(config-if)#untagged-frames-discard
8003:5(config-if)#vlacp fast-periodic-time 500 timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
8003:5(config-if)#vlacp enable
8003:5(config-if)#exit
```

8005 & 8006: Same configuration on both switches

```
8005:5(config)#interface gigabitEthernet 2/5,2/34
8005:5(config-if)#untagged-frames-discard
8005:5(config-if)#vlacp fast-periodic-time 500 timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
8005:5(config-if)#vlacp enable
8005:5(config-if)#exit
```

8007:

```
8007:5(config)#interface gigabitEthernet 3/27,3/28
8007:5(config-if)#untagged-frames-discard
8007:5(config-if)#vlacp fast-periodic-time 500 timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
8007:5(config-if)#vlacp enable
8007:5(config-if)#exit
```

VSP 9000 - Interface Configuration

9001 & 9002: Same configuration on both switches

```
9001:1(config)#interface GigabitEthernet 3/3,3/29,4/17,4/18
9001:1(config-if)#untagged-frames-discard
9001:1(config-if)#vlacp fast-periodic-time 500 timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
9001:1(config-if)#vlacp enable
9001:1(config-if)#exit
```

ERS 4800 - Interface Configuration

```
4801(config)#vlan ports 45,46 filter-untagged-frame enable
```

```
4801(config)#interface ethernet 3
4801(config-if)#vlacp timeout short
4801(config-if)#vlacp timeout-scale 5
4801(config-if)#vlacp enable
4801(config-if)#exit
```

22.1.1.13 IST Configuration – SMLT Cluster switch 7001 & 7002

The following port based VLANs will be configured on the SMLT Switch cluster

Switch	Feature	Parameter	Value
7001 & 7002	IST	MLT ID	1
		VLAN	2
		VLAN Port Members	38 & 39
7001 & 7002	LACP	Aggregation	Disable on ports 38-39
		Mode	Off on ports 38-39
7001	IST VLAN	IP address	10.70.2.5/30
		Ports	38,39
7002	IST VLAN	IP address	10.70.2.6/30
		Ports	38,39

For the VSP 7000, it is important to not enable the *filter-untagged-frame* option on the IST port members. Also, the default PVID of all IST ports must be the primary B-VLAN ID; for this example, this will be B-VLAN ID 4051. This will happened automatically providing SPB is enable first prior to enabling the IST.

Also, it is recommended to not enable VLACP on the IST.

Please note that Spanning Tree should be disabled on all SPB NNI interfaces that are not configured as SMLT ports. SMLT by default will disable Spanning Tree.

Since we will be adding an IST interface via the red rear ports, ports 38 & 39, we will have to disable LACP on these ports and add an MLT.

VSP 7000 – Disable LACP on ports 38 & 39

7001 & 7002: Same configuration on both switches

```
7001(config)#interface ethernet 38,39
7001(config-if)#no lacp aggregation enable
7001(config-if)#lacp mode off
7001(config-if)#exit
```

Prior to enabling the IST, LACP must be disabled on the rear port member that are being used for the IST

VSP 7000 – Create MLT to be used by IST

7001 & 7002: Same configuration on both switches

```
7001(config)#mlt 1 name IST enable member 38,39 learning disable
```

Verify MLT configuration

```
7001(config)#show mlt 1
Id Name                 Members                  Bpdu    Mode            Status  Type
-- -------------------- ------------------------ ------- --------------- ------- -------
1  IST                  38-39                    All     Basic           Enabled Trunk
```

VSP 7000 – SPB Interface Configuration

7001 & 7002: Same configuration on both switches

```
7001(config)#interface ethernet 38,39
7001(config-if)#isis
7001(config-if)#isis spbm 1
7001(config-if)#isis enable
7001(config-if)#spanning-tree mstp learning disable
7001(config-if)#exit
```

VSP 7000 – Remove default VLAN from SPB ports

7001 & 7002: Same configuration on both switches

```
7001(config)#vlan members remove 1 38,39
```

VSP 7000 – Add IST VLAN 2 and IP address

7001 & 7002: Same configuration on both switches

```
7001(config)#vlan create 2 name ist type port
7001(config)#vlan members 2 38,39
7001(config)#vlan members remove 1 38,39
```

IP configuration on 7001

```
7001(config)#interface vlan 2
7001(config-if)#ip address 10.70.2.5 255.255.255.252
7001(config-if)#exit
```

IP configuration on 7002

```
7002(config)#interface vlan 2
7002(config-if)#ip address 10.70.2.6 255.255.255.252
7002(config-if)#exit
```

VSP 7000 – Create IST

```
7001(config)#interface mlt 1
7001(config-if)#ist peer-ip 10.70.2.6 vlan 2
7001(config-if)#ist enable
7002(config-if)#exit
```

```
7002(config)#interface mlt 1
7002(config-if)#ist peer-ip 10.70.2.5 vlan 2
7002(config-if)#ist enable
7002(config-if)#exit
```

Verify IST Operation assuming the SMLT cluster peer is also configured

```
7001(config)#show ist
MLT ID Enabled Running Master Peer IP Address Vlan ID
--------- ---------- ---------- --------- --------------------- ----------
1      YES     YES     NO     10.70.2.2       2
7001(config)#show smlt
================================================================================
                                        MLT SMLT Info
================================================================================
MLT   SMLT     ADMIN    CURRENT
ID    ID       TYPE     TYPE
--------------------------------------------------------------------------------
1               ist     ist
```

Verify the default VLAN is now the primary B-VLAN ID. Also, make sure the Filter Untagged Frames option is disabled.

```
7001(config)#show vlan interface info 38,39
      Filter     Filter
     Untagged Unregistered
Port  Frames     Frames    PVID PRI    Tagging     Name
----- ---------- ---------- ------ ----- ---------- ----------
38    No       Yes          4051 0   TagAll        Port 38
39    No       Yes          4051 0   TagAll        Port 39
```

22.1.1.14 ISIS L1-metric – Optional

As an option, we can change the default metric on the FI rear ports and SPB front ports to reflect the actual port speeds.

VSP 7000 – ISIS L1 Metric

Switches 7001 and 7003: Same configuration on both switches

```
7001(config)#interface ethernet 3
7001(config-if)#isis spbm 1 l1-metric 200
7001(config-if)#exit
7001(config)#interface ethernet 34
7001(config-if)#isis spbm 1 l1-metric 17
7001(config-if)#exit
7001(config)#interface ethernet 38
7001(config-if)#isis spbm 1 l1-metric 25
7001(config-if)#exit
```

Switches 7002 and 7004: Same configuration on both switches

```
7002(config)#interface ethernet 34
7002(config-if)#isis spbm 1 l1-metric 17
7002(config-if)#exit
7002(config)#interface ethernet 38
7002(config-if)#isis spbm 1 l1-metric 25
7002(config-if)#exit
```

22.1.1.15 Connectivity Fault Management (CFM) Configuration

Switch	Parameter	Value
CFM		
All bridges	CFM	Enabled
	*Maintenance Domain Name	spbm
	*Maintenance Association Name	4051
	*Maintenance Association Name	4052
4001	Maintenance End Point (MEP) ID	401
4002	Maintenance End Point (MEP) ID	402
4801	Maintenance End Point (MEP) ID	4801
7001	Maintenance End Point (MEP) ID	7001
7002	Maintenance End Point (MEP) ID	7002
7003	Maintenance End Point (MEP) ID	7003
7004	Maintenance End Point (MEP) ID	7004
8003	Maintenance End Point (MEP) ID	803
8004	Maintenance End Point (MEP) ID	804
8005	Maintenance End Point (MEP) ID	805
8006	Maintenance End Point (MEP) ID	806
8007	Maintenance End Point (MEP) ID	807
9001	Maintenance End Point (MEP) ID	901
9002	Maintenance End Point (MEP) ID	902

* Default values on all switches

VSP 4000 - CFM Configuration

4001:

```
4001:1(config)#cfm spbm mepid 401
4001:1(config)#cfm spbm enable
```

4002:

```
4002:1(config)#cfm spbm mepid 402
4002:1(config)#cfm spbm enable
```

VSP 7000 – CFM Configuration

7001:

```
7001(config)#cfm spbm mepid 7001
7001(config)#cfm spbm enable
```

7002:

```
7002(config)#cfm spbm mepid 7002
7002(config)#cfm spbm enable
```

7003:

```
7003(config)#cfm spbm mepid 7003
7003(config)#cfm spbm enable
```

7004:

```
7004(config)#cfm spbm mepid 7004
7004(config)#cfm spbm enable
```

ERS 8800 - CFM Configuration

8003:

```
8003:5(config)#cfm spbm mepid 803
8003:5(config)#cfm spbm enable
```

8004:

```
8004:5(config)#cfm spbm mepid 804
8004:5(config)#cfm spbm enable
```

8005:

```
8005:5(config)#cfm spbm mepid 805
8005:5(config)#cfm spbm enable
8005:5(config)#cfm cmac mepid 805
8005:5(config)#cfm cmac enable
```

8006:

```
8006:5(config)#cfm spbm mepid 806
8006:5(config)#cfm spbm enable
8006:5(config)#cfm cmac mepid 806
```

```
8006:5(config)#cfm cmac enable
```

8007:

```
8006:5(config)#cfm spbm mepid 807
8006:5(config)#cfm spbm enable
8006:5(config)#cfm cmac mepid 807
8006:5(config)#cfm cmac enable
```

VSP 9000 - CFM Configuration assuming 3.4 or higher is used

9001:

```
9001:1(config)#cfm spbm mepid 901
9001:1(config)#cfm spbm enable
9001:1(config)#cfm cmac mepid 901
9001:1(config)# cfm cmac enable
```

9002:

```
9002:1(config)# cfm spbm mepid 902
9002:1(config)#cfm spbm enable
9002:1(config)#cfm cmac mepid 902
9002:1(config)# cfm cmac enable
```

22.1.1.16 QoS

QoS by default is enabled on all NNI interfaces. Depending on the switch, QoS may still have to be enabled on the UNI interface or filters must be used to provide end-to-end QoS.

On the VSP 4000, VSP 8000, VSP 7200, and VSP 9000, the interface level parameters *802.1p-override disable*, *enable-diffserv enable* and no *access-diffserv enable* are the default settings. On an UNI interface, this has the overall result of honoring p-bits for bridge traffic and DSCP values for routed traffic. Note that on the ERS 8000, the diffserv settings are disabled by default; the *enable-diffserv* parameter must be enabled for the ERS 8000 to behave the same as the VSP 9000 and VSP 4000. Note that with these settings, on any untagged L2 port, i.e. a port member of a C-VLAN used for an L2VSN, as there is no p-bit to determine the QoS level, either the port or VLAN QoS level determines the QoS classification. To be safe, it is recommended to enable the *802.1p-override* parameter. This has the net effect of honoring the DSCP value for L2 traffic so it makes no difference if the ingress port is tagged or untagged.

VSP 4000, VSP 9000, and ERS 8800 – QoS Configuration

VSP 4000 & VSP 9000: All C-VLAN port members – L2 and L3

```
interface gigabitEthernet <slot/port>
qos 802.1p-override enable
```

ERS 8000: All C-VLAN port members – L2 and L3

```
interface gigabitEthernet <slot/port>
qos 802.1p-override enable
enable-diffserv enable
```

If you do not wish to trust the incoming traffic, i.e. remark all traffic to Best Effort and use ACL's to remark traffic, you need to enable the *access-diffserv* parameter

VSP 4000, VSP 8000, VSP 7200, VSP 9000, and ERS 8000:

```
interface gigabitEthernet <slot/port>
access-diffserv enable
```

On the VSP 7000, by default, all ports are members of the default interface group *allQoSPolicyIfcs* has an interface class of trusted resulting in all traffic being trusted. This results in honoring the DSCP value and updating the 802.1 p-bit value based on the DSCP mapping table. If you wish to not trust the incoming traffic and use Traffic Profiles to remark traffic, you need to create an interface group of *untrusted*.

```
qos if-group name <word> class untrusted
qos if-assign port <port list> name <word>
```

22.1.2 Configuration using EDM – Using 8005 as an example

If using EDM to config SPB, please follow the steps shown below. The following configuration is in reference to 9002 and assumes the base configuration has been configured – i.e. VLAN and SMLT configuration

22.1.2.1 SPB and B-VLAN Configuration

8005 - Step 1 – Via EDM, go to Configuration -> IS-IS -> SPBM -> Global and enable SPBM globally

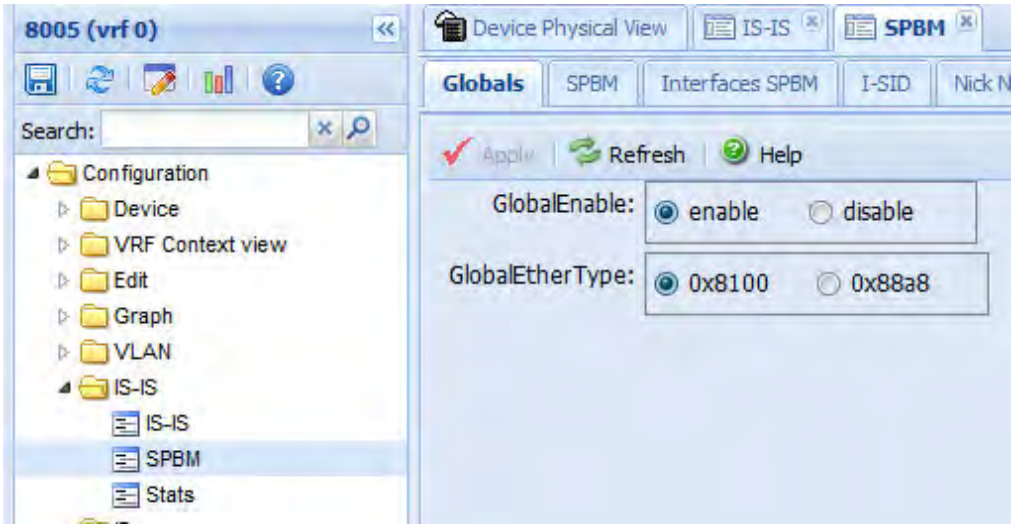

8005: Step 2 – Via EDM, go to Configuration -> VLAN -> VLANs -> Basic -> Insert to add primary B-VLAN 4051 (make sure to select Type: spbm-bvlan)

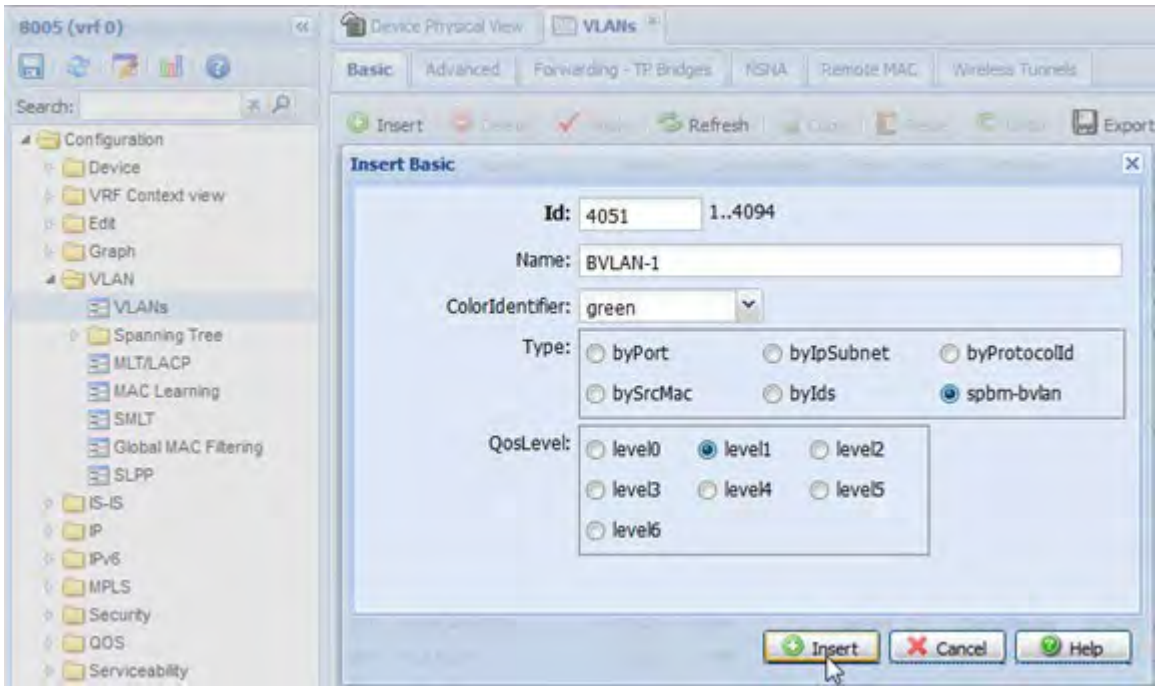

8005: Step 3 – Via EDM, go to Configuration -> VLAN -> VLANs -> Basic -> Insert to add secondary B-VLAN 4052 (make sure to select Type: spbm-bvlan)

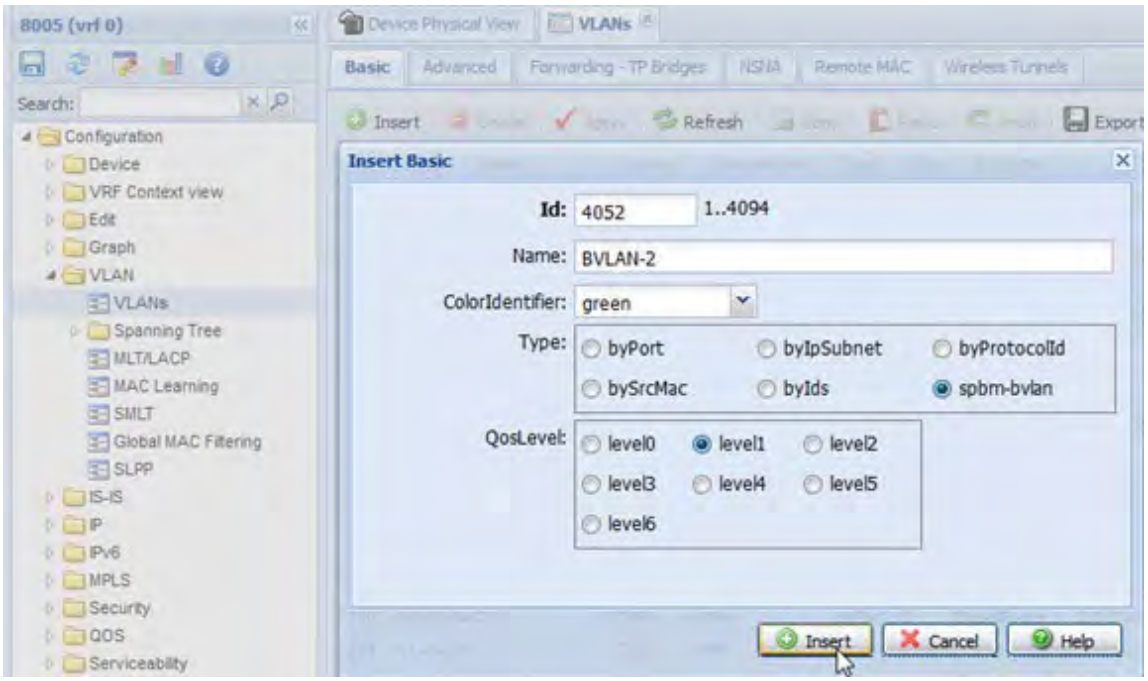

8005: Step 4 – Via EDM, go to Configuration -> IS-IS -> IS-IS -> Global, add the SPBM System ID and set the Admin State to enable

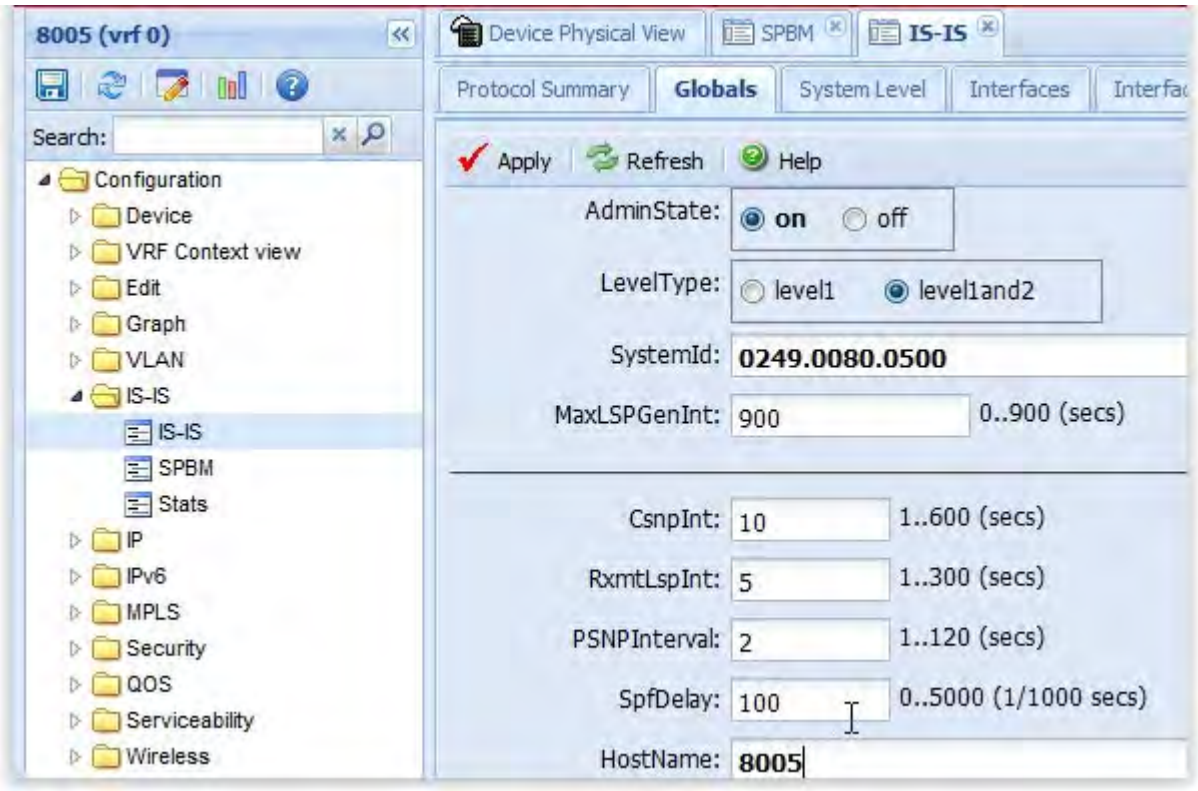

8005: Step 5 – Via EDM, go to Configuration -> IS-IS -> SPBM -> SPBM, add the SPBM node nickname, primary VLAN, and both primary and secondary VLANs as ERS-3 is part of an SMLT cluster

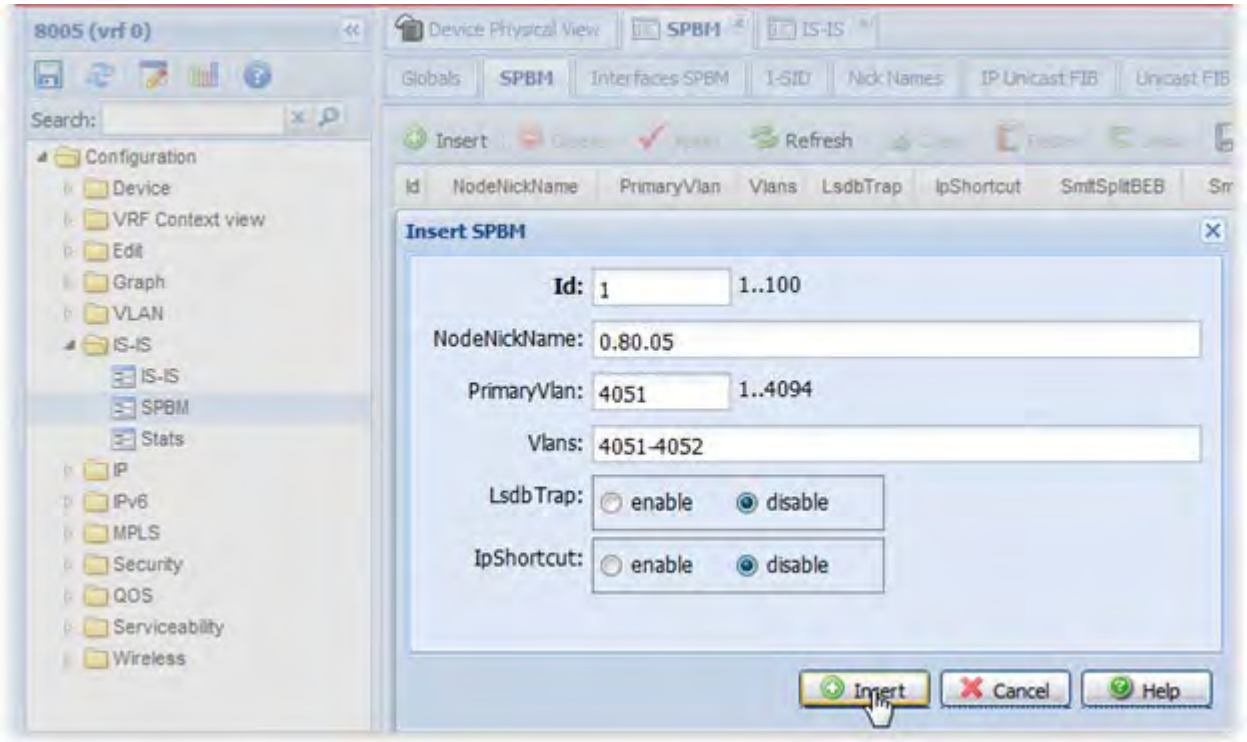

8005: Step 6 – Via EDM, go to Configuration -> IS-IS -> IS-IS -> Manual Area to add the IS-IS area which in our example is area 49.0001

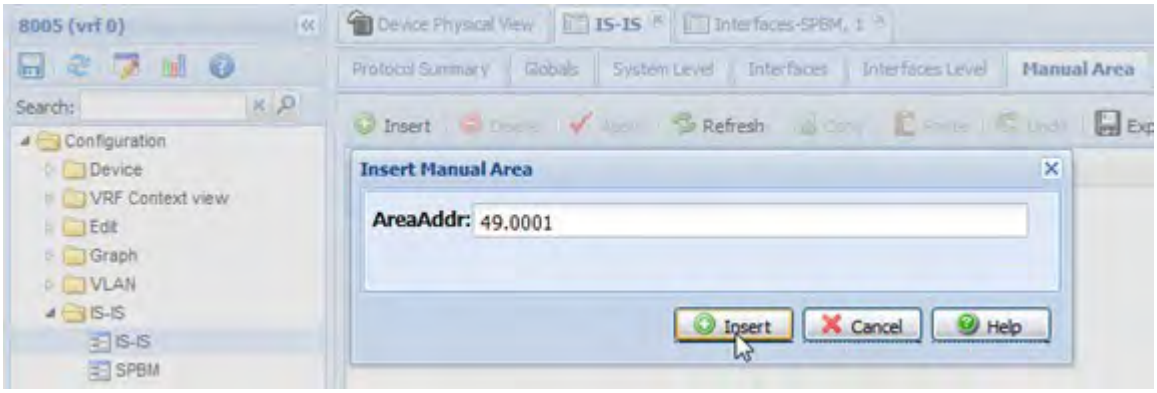

8005: Step 7 – Via EDM, go to Configuration -> IS-IS -> SPBM -> SPBM and change the and SMLT peer B-MAC (001e.1f48.f3df) – use the ACLI *show isis system-id* command on the peer cluster switch to get the System ID value

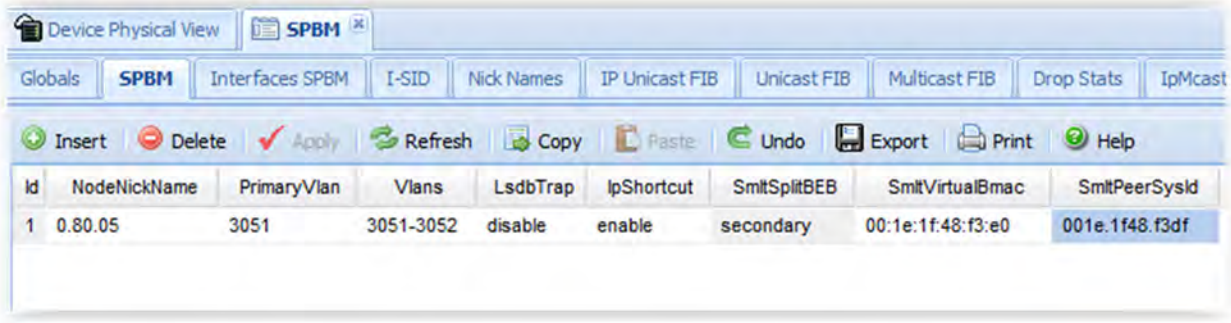

22.1.2.2 IS-IS and SPB Configuration

8005: Step 1 – Via EDM, go to Configuration -> IS-IS -> IS-IS -> Interfaces to add IS-IS on all appropriate interfaces; in regards to 8005, this will the IST interface, port 2/5 and 2/34. Do not enable IS-IS (AdminState = off) until SPBM is enabled on the interface

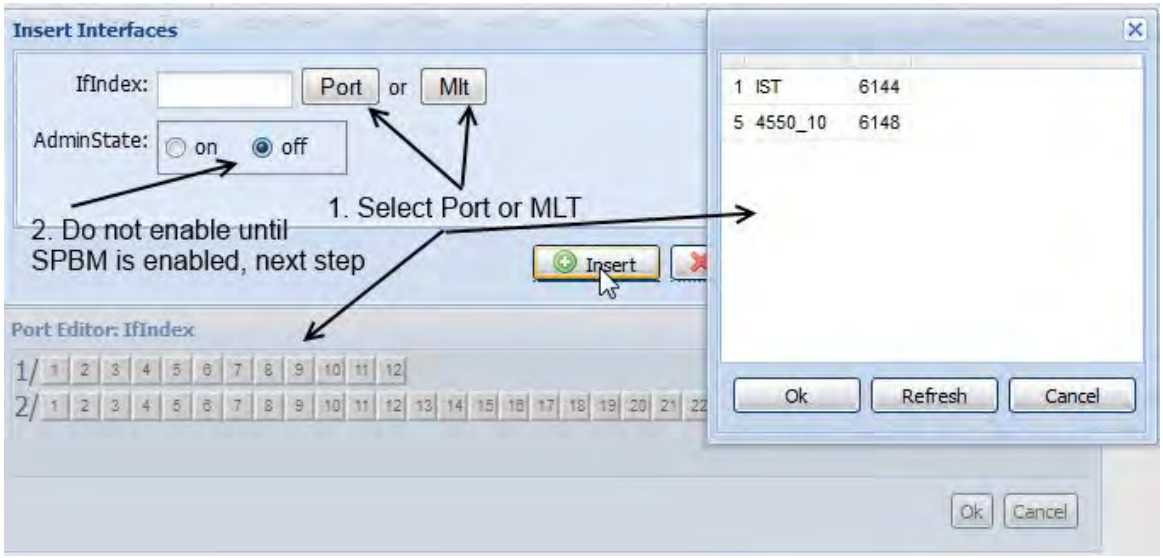

8005: Step 2 – Via EDM, go to Configuration -> IS-IS -> IS-IS -> Interfaces, select interface and then click on SPBM

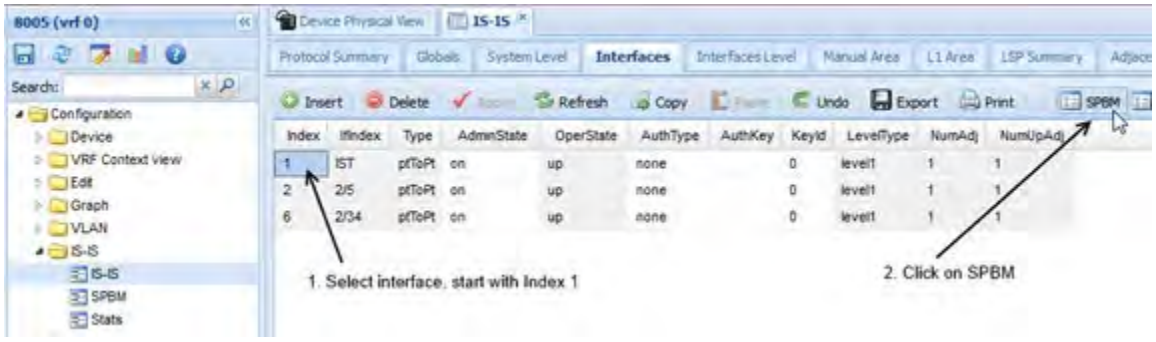

8005: Step 3 – Via SBPM windows, select SPBM Id of 1 and enable SPBM

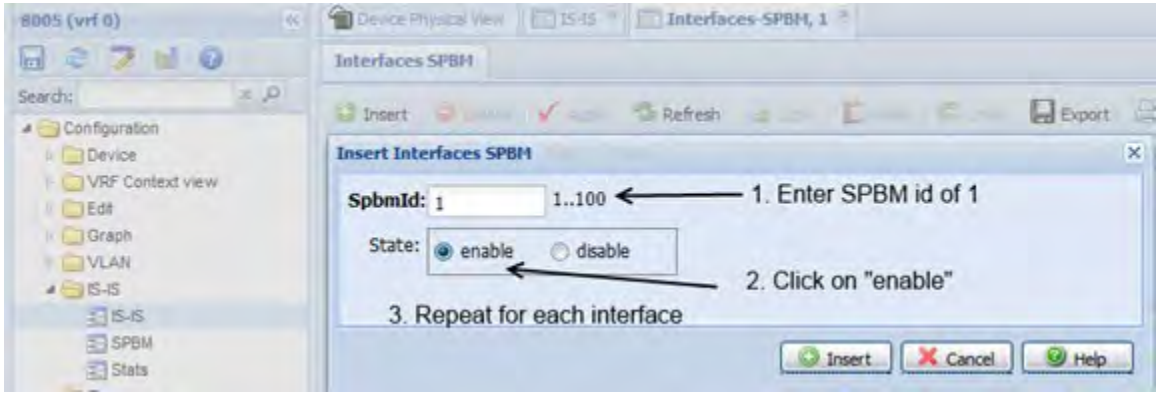

8005: Step 4 – Via EDM, go back to Configuration -> IS-IS -> IS-IS -> Interface and enable IS-IS on each interface

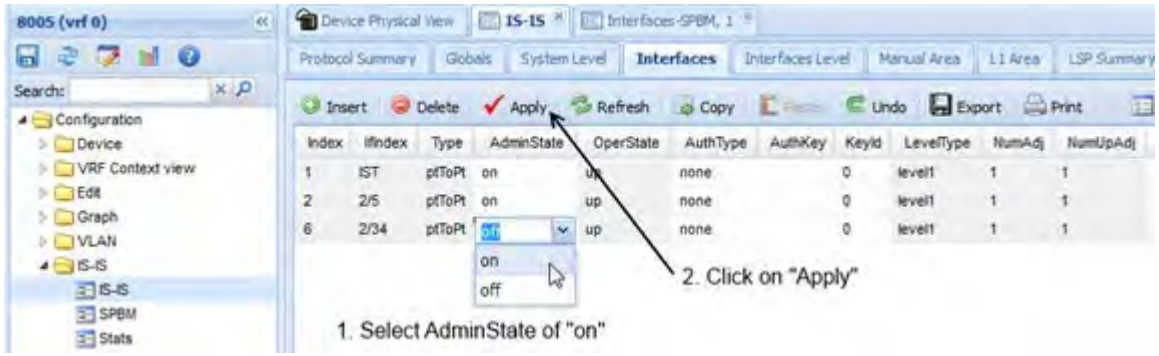

22.1.3 Verify Operations

22.1.3.1 Global Settings

Step 1 – Verify IS-IS global settings:

```
8800:5#show isis
```

Results: Example from 8003. Admin state should show *enabled* and in our case the configured B-MAC address of *0080.2dbe.23df* should be displayed.

8003:

```
================================================================================
                              ISIS General Info
================================================================================
                           AdminState : enabled
                           RouterType : Level 1
                            System ID : 0080.2dbe.23df
                  Max LSP Gen Interval : 900
                               Metric : wide
                 Overload-on-startup : 20
                             Overload : false
                       Csnp Interval : 10
                       PSNP Interval : 2
                    Rxmt LSP Interval : 5
                            spf-delay : 100
                          Router Name : 8003
                    ip source-address :
                    Num of Interfaces : 4
                Num of Area Addresses : 1
```

Step 2 – Verify IS-IS network information

```
show isis net
```

Results: From all switches

4001:

```
================================================================================

                                ISIS Net Info

================================================================================

NET

--------------------------------------------------------------------------------------------------------

49.0001.d4ea.0e10.e465.00
```

4002:

```
================================================================================

                                ISIS Net Info

================================================================================

NET

--------------------------------------------------------------------------------------------------------

49.0001.a012.90d3.ec65.00
```

7001:

```
================================================================================

                                ISIS Net Info

================================================================================

NET

--------------------------------------------------------------------------------------------------------

49.0001.fca8.41f6.37df.00
```

7002:

```
================================================================================

                                ISIS Net Info
```

```
================================================================================

NET

--------------------------------------------------------------------------------

49.0001.3cb1.5bff.5fdf.00


7003:

================================================================================

                                ISIS Net Info

================================================================================

NET

--------------------------------------------------------------------------------

49.0001.7030.1823.7fdf.00


7004:

================================================================================

                                ISIS Net Info

================================================================================

NET

--------------------------------------------------------------------------------

49.0001.7030.1823.9bdf.00


9001:

================================================================================

                                ISIS Net Info

================================================================================

NET

--------------------------------------------------------------------------------

49.0001.d4ea.0efd.e3df.00


9002:

================================================================================
```

```
                                    ISIS Net Info
================================================================================
NET
--------------------------------------------------------------------------------
49.0001.d4ea.0efd.e4ac.00
```

8003:

```
================================================================================
                                    ISIS Net Info
================================================================================
NET
--------------------------------------------------------------------------------
49.0001.0080.2dbe.23df.00
```

8004:

```
================================================================================
                                    ISIS Net Info
================================================================================
NET
--------------------------------------------------------------------------------
49.0001.00e0.7bbc.23df.00
```

8005:

```
================================================================================
                                    ISIS Net Info
================================================================================
NET
--------------------------------------------------------------------------------
49.0001.0024.43b4.e3df.00
```

8006:

```
================================================================================

                                ISIS Net Info

================================================================================

NET

--------------------------------------------------------------------------------------------------------------

49.0001.001e.1f48.f3df.00


8007:

================================================================================

                                ISIS Net Info

================================================================================

NET

--------------------------------------------------------------------------------------------------------------

49.0001.00e0.7bb3.07df.00


4801:

================================================================================

                                ISIS Net Info

================================================================================

NET

--------------------------------------------------------------------------------------------------------------

49.0001.ccf9.54b4.ac21.00
```

On each switch, verify the following:

Option	Verify
System ID	This is the unique MAC address which will be used by SPB to build adjacencies and forwarding.
NET	Should be displayed as ***49.0001.<system id>.00*** where 49.0001 is the IS-IS area ID

22.1.3.2 Verify IS-IS Interface and Adjacencies

Step 1 – Verify IS-IS interfaces:

```
show isis interface
```

Results: From switch 4001 and 7002

4001:

```
================================================================================
                              ISIS Interfaces
================================================================================
IFIDX      TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC

--------------------------------------------------------------------------------
Mlt1       pt-pt   Level 1   UP        UP         1      1       10
Port1/47   pt-pt   Level 1   UP        UP         1      1       10
```

7001:

```
================================================================================
                              ISIS Interfaces
================================================================================
IFIDX      TYPE    LEVEL     OP-STATE  ADM-STATE  ADJ    UP-ADJ  SPBM-L1-METRIC

--------------------------------------------------------------------------------
Trunk: 64  pt-pt   Level 1   UP        UP          1     1       10
Port: 3    pt-pt   Level 1   UP        UP          1     1       10
Trunk: 63  pt-pt   Level 1   UP        UP          1     1       10

7001#show mlt
Id Name             Members               Bpdu   Mode           Status  Type
-- ---------------- --------------------- ------ -------------- ------- ------
63 Trunk #31        38-39                 Single DynLag/Basic   Enabled Trunk
64 Trunk #32        34-36                 Single DynLag/Basic   Enabled Trunk
```

Step 2 – Verify IS-IS adjacencies

```
show isis adjacencies
```

Results: From switch 4001 and 7002

4001:

```
================================================================================
                                ISIS Adjacencies
================================================================================
INTERFACE L STATE        UPTIME PRI HOLDTIME SYSID             HOST-NAME
--------------------------------------------------------------------------------
Mlt1      1 UP         21:01:51 127       22 d4ea.0efd.e3df    9001
Port1/47  1 UP       1d 19:57:51 127       23 d4ea.0efd.e4ac    9002

--------------------------------------------------------------------------------
 2 out of 2 interfaces have formed an adjacency
--------------------------------------------------------------------------------
```

7001:

```
================================================================================
                                ISIS Adjacencies
================================================================================
INTERFACE L STATE        UPTIME PRI HOLDTIME SYSID             HOST-NAME

================================================================================
Trunk: 64 1 UP     28d 13:12:00 127       17  7030.1823.7fdf   7003
Port: 3   1 UP          20:25:01 127       18  0024.43b4.e3df   8005
Trunk: 63 1 UP     14d 19:15:40 127       18  fca8.41f6.37df   7001
--------------------------------------------------------------------------------
 3 interfaces have formed an adjacency
--------------------------------------------------------------------------------
```

On each switch, verify the following:

Option	Verify
IS-IS Interface	
TYPE	The value displayed should be ***pt-pt*** which indicates Point to Point
OP-STATE ADM-STATE	The value displayed should be ***UP*** which indicates that IS-IS have been configured and is operational for the interface index shown
IS-IS Adjacencies	
STATE HOST-NAME	Should be displayed as ***UP*** indicating there is an adjacency with its neighbor as shown via ***HOST-NAME***

22.1.3.3 Verify IS-IS SPB Information

Step 1 – Verify IS-IS interfaces

```
show isis spbm
```

Results: From switch 4001, 7002, 9001, 8003, 8005, and 8007

4001:

```
================================================================================
                                ISIS SPBM Info
================================================================================
SPBM        B-VID       PRIMARY     NICK      LSDB      IP        MULTICAST
INSTANCE                VLAN        NAME      TRAP
--------------------------------------------------------------------------------
1           4051-4052   4051        0.40.01   disable   enable    disable
```

7002:

```
============================================================
                   ISIS SPBM Info
============================================================
SPBM       B-VID          PRIMARY      NICK             LSDB
INSTANCE                  VLAN         NAME             TRAP
============================================================
1          4051-4052      4051         0.70.02          FALSE
```

9001:

```
================================================================================
                                ISIS SPBM Info
================================================================================
SPBM        B-VID       PRIMARY     NICK      LSDB      IP        MULTICAST
INSTANCE                VLAN        NAME      TRAP
--------------------------------------------------------------------------------
1           4051-4052   4051        0.90.01   disable   enable    enable

================================================================================
                             ISIS SPBM SMLT Info
================================================================================
SPBM        SMLT-SPLIT-BEB       SMLT-VIRTUAL-BMAC     SMLT-PEER-SYSTEM-ID
INSTANCE
```

```
-------------------------------------------------------------------------------------------------------------
1          primary              d4:ea:0e:fd:e3:e0    d4ea.0efd.e4ac
```

8003:

```
===============================================================================
                                ISIS SPBM Info
===============================================================================================
SPBM       B-VID      PRIMARY    NICK     LSDB     IP       MULTICAST
INSTANCE              VLAN       NAME     TRAP
-------------------------------------------------------------------------------------------------------------
1          4051-4052  4051       0.80.03  disable  enable   disable

===============================================================================
                             ISIS SPBM SMLT Info
===============================================================================================
SPBM       SMLT-SPLIT-BEB        SMLT-VIRTUAL-BMAC     SMLT-PEER-SYSTEM-ID
INSTANCE
 -----------------------------------------------------------------------------
1          primary               00:00:00:00:00:00
```

8005:

```
===============================================================================
                                ISIS SPBM Info
===============================================================================================
SPBM       B-VID      PRIMARY    NICK     LSDB     IP       MULTICAST
INSTANCE              VLAN       NAME     TRAP
-------------------------------------------------------------------------------------------------------------
1          4051-4052  4051       0.80.05  disable  enable   enable

===============================================================================
                             ISIS SPBM SMLT Info
===============================================================================================
SPBM       SMLT-SPLIT-BEB       SMLT-VIRTUAL-BMAC     SMLT-PEER-SYSTEM-ID
INSTANCE
 -----------------------------------------------------------------------------
1          primary              00:1e:1f:48:f3:e0    001e.1f48.f3df
```

```
8007:

================================================================================
                                ISIS SPBM Info
================================================================================
SPBM        B-VID       PRIMARY     NICK      LSDB      IP        MULTICAST
INSTANCE                VLAN        NAME      TRAP
--------------------------------------------------------------------------------
1           4051-4052   4051        0.80.07   disable   enable    enable

================================================================================
                              ISIS SPBM SMLT Info
================================================================================
SPBM        SMLT-SPLIT-BEB       SMLT-VIRTUAL-BMAC      SMLT-PEER-SYSTEM-ID
INSTANCE
--------------------------------------------------------------------------------
   1           primary               00:00:00:00:00:00
```

Step 2 – show isis spbm unicast-fib

```
show isis spbm unicast-fib
```

Results: From switch 8007

```
8007:

================================================================================
                         SPBM UNICAST FIB ENTRY INFO
================================================================================
DESTINATION         BVLAN  SYSID            HOST-NAME          OUTGOING      COST
ADDRESS                                                       INTERFACE
--------------------------------------------------------------------------------
00:80:2d:be:23:df   4051   0080.2dbe.23df   8003               3/27          10
02:80:03:ff:ff:ff   4051   0080.2dbe.23df   8003               3/27          10
00:80:2d:be:23:df   4052   0080.2dbe.23df   8003               3/27          10
02:80:03:ff:ff:ff   4052   0080.2dbe.23df   8003               3/27          10
00:e0:7b:b3:07:df   4051   00e0.7bb3.07df   8007               cpp           0
02:80:07:ff:ff:ff   4051   00e0.7bb3.07df   8007               cpp           0
00:e0:7b:b3:07:df   4052   00e0.7bb3.07df   8007               cpp           0
02:80:07:ff:ff:ff   4052   00e0.7bb3.07df   8007               cpp           0
00:e0:7b:bc:23:df   4051   00e0.7bbc.23df   8004               3/28          10
02:80:04:ff:ff:ff   4051   00e0.7bbc.23df   8004               3/28          10
00:e0:7b:bc:23:df   4052   00e0.7bbc.23df   8004               3/28          10
02:80:04:ff:ff:ff   4052   00e0.7bbc.23df   8004               3/28          10
```

02:40:02:ff:ff:ff	4051	a012.90d3.ec65	4002	3/27	30
a0:12:90:d3:ec:65	4051	a012.90d3.ec65	4002	3/27	30
a0:12:90:d3:ec:66	4051	a012.90d3.ec65	4002	3/27	30
a0:12:90:d3:ec:65	4052	a012.90d3.ec65	4002	3/28	30
a0:12:90:d3:ec:66	4052	a012.90d3.ec65	4002	3/28	30
a0:12:90:d3:ec:66	4051	d4ea.0e10.e465	4001	3/27	30
d4:ea:0e:10:e4:65	4051	d4ea.0e10.e465	4001	3/27	30
02:40:01:ff:ff:ff	4052	d4ea.0e10.e465	4001	3/28	30
a0:12:90:d3:ec:66	4052	d4ea.0e10.e465	4001	3/28	30
d4:ea:0e:10:e4:65	4052	d4ea.0e10.e465	4001	3/28	30
02:80:05:ff:ff:ff	4051	b0ad.aa47.0884	8005	3/27	20
00:24:43:b4:e3:df	4051	0024.43b4.e3df	8005	3/27	20
00:1e:1f:48:f3:e0	4051	0024.43b4.e3df	8005	3/27	20
00:24:43:b4:e3:df	4052	0024.43b4.e3df	8005	3/27	20
00:1e:1f:48:f3:e0	4052	0024.43b4.e3df	8005	3/27	20
02:80:06:ff:ff:ff	3052	e45d.523c.4884	8006	3/28	20
00:1e:1f:48:f3:e0	4051	001e.1f48.f3df	8006	3/28	20
00:1e:1f:48:f3:df	4051	001e.1f48.f3df	8006	3/28	20
00:1e:1f:48:f3:e0	4052	001e.1f48.f3df	8006	3/28	20
00:1e:1f:48:f3:df	4052	001e.1f48.f3df	8006	3/28	20
fc:a8:41:f6:37:df	4051	fca8.41f6.37df	7001	3/27	220
3c:b1:5b:ff:5f:e0	4052	fca8.41f6.37df	7001	3/28	220
fc:a8:41:f6:37:df	4052	fca8.41f6.37df	7001	3/28	220
3c:b1:5b:ff:5f:df	4051	3cb1.5bff.5fdf	7002	3/27	237
3c:b1:5b:ff:5f:e0	4051	3cb1.5bff.5fdf	7002	3/27	237
3c:b1:5b:ff:5f:df	4052	3cb1.5bff.5fdf	7002	3/28	237
70:30:18:23:7f:df	4051	7030.1823.7fdf	7003	3/28	30
70:30:18:23:7f:df	4052	7030.1823.7fdf	7003	3/28	30
70:30:18:23:9b:df	4051	7030.1823.9bdf	7004	3/28	40
70:30:18:23:9b:df	4052	7030.1823.9bdf	7004	3/28	40
02:90:01:ff:ff:ff	4051	d4ea.0efd.e3df	9001	3/27	20
d4:ea:0e:fd:e3:df	4051	d4ea.0efd.e3df	9001	3/27	20
d4:ea:0e:fd:e3:e0	4051	d4ea.0efd.e3df	9001	3/27	20
02:90:01:ff:ff:ff	4052	d4ea.0efd.e3df	9001	3/27	20
d4:ea:0e:fd:e3:df	4052	d4ea.0efd.e3df	9001	3/27	20
d4:ea:0e:fd:e3:e0	4052	d4ea.0efd.e3df	9001	3/27	20
02:90:02:ff:ff:ff	4051	d4ea.0efd.e4ac	9002	3/28	20
d4:ea:0e:fd:e3:e0	4051	d4ea.0efd.e4ac	9002	3/28	20
d4:ea:0e:fd:e4:ac	4051	d4ea.0efd.e4ac	9002	3/28	20
02:90:02:ff:ff:ff	4052	d4ea.0efd.e4ac	9002	3/28	20
d4:ea:0e:fd:e3:e0	4052	d4ea.0efd.e4ac	9002	3/28	20
d4:ea:0e:fd:e4:ac	4052	d4ea.0efd.e4ac	9002	3/28	20

On each switch, verify the following:

Option	Verify
IS-IS SPB	
B-VID PRIMARY VLAN	The B-VLAN is displayed should be ***4051*** and ***4052*** where the primary B-VLAN should be ***4051***
NICK NAME	The value displayed should be as follows per this configuration example: • 4001: ***0.40.01*** • 4001: ***0.40.02*** • 7001: ***0.70.01*** • 7002: ***0.70.02*** • 7003: ***0.70.03*** • 7004: ***0.70.04*** • 9001: ***0.90.01*** • 9002: ***0.90.02*** • 8003: ***0.80.03*** • 8004: ***0.80.04*** • 8005: ***0.80.05*** • 8006: ***0.80.06*** • 8007: ***0.80.07***
SPB Unicast FIB	
FIB ENTRY	For each host, there should be a destination forwarding entry via both B-VLANs. Note that the default metric is 10 for all links.

22.1.3.4 Verify IS-IS Link-State Database

Step 1 – Show IS-IS LSDB

```
show isis lsdb
```

Results: From switches 4001 and 9001

4001:

```
================================================================================
                                ISIS LSDB
================================================================================
LSP ID                    LEVEL    LIFETIME  SEQNUM    CHKSUM   HOST-NAME
--------------------------------------------------------------------------------
0080.2dbe.23df.00-00      1        748       0x1545    0xbc63   8003
00e0.7bbc.23df.00-00      1        1068      0x1191    0xbc2    8004
0024.43b4.e3df.00-00      1        978       0xae9     0x8e34   8005
001e.1f48.f3df.00-00      1        930       0x1555    0x90c3   8006
00e0.7bb3.07df.00-00      1        1090      0x23c     0x1dcb   8007
```

```
d4ea.0e10.e465.00-00      1        1175     0x144a    0xabc    4001
d4ea.0e10.e465.00-01      1        1175     0x11d9    0x7421   4001
d4ea.0e10.e465.00-02      1        1175     0x1085    0x9890   4001
d4ea.0e10.e465.00-03      1        1175     0xafd     0xa41    4001
a012.90d3.ec65.00-00      1        355      0x11f8    0xfb1c   4002
a012.90d3.ec65.00-01      1        355      0x11cd    0xb1ec   4002
a012.90d3.ec65.00-02      1        355      0x1073    0xb77f   4002
a012.90d3.ec65.00-03      1        355      0xaf6     0x3619   4002
fca8.41f6.37df.00-00      1        945      0x76c     0x68d7   7001
3cb1.5bff.5fdf.00-00      1        1072     0x1d1f    0xfaff   7002
7030.1823.7fdf.00-00      1        1168     0x1d1c    0x89ae   7003
7030.1823.9bdf.00-00      1        335      0x1d1b    0xf8a3   7004
d4ea.0efd.e3df.00-00      1        647      0xc1      0x3177   9001
d4ea.0efd.e4ac.00-00      1        650      0xbd      0x9a1a   9002
--------------------------------------------------------------------------------
Level-1 : 19 out of 19 Total Num of LSP Entries
Level-2 : 0 out of 0 Total Num of LSP Entries
--------------------------------------------------------------------------------
```

9001:

```
================================================================================
                                ISIS LSDB
================================================================================
LSP ID                    LEVEL    LIFETIME  SEQNUM    CHKSUM   HOST-NAME
 ------------------------------------------------------------------------
0080.2dbe.23df.00-00      1        832      0x1545    0xbc63   8003
00e0.7bbc.23df.00-00      1        1142     0x1191    0xbc2    8004
0024.43b4.e3df.00-00      1        1056     0xae9     0x8e34   8005
001e.1f48.f3df.00-00      1        1009     0x1555    0x90c3   8006
00e0.7bb3.07df.00-00      1        1165     0x23c     0x1dcb   8007
d4ea.0e10.e465.00-00      1        371      0x1449    0xcbb    4001
d4ea.0e10.e465.00-01      1        371      0x11d8    0x7620   4001
d4ea.0e10.e465.00-02      1        371      0x1084    0x9a8f   4001
d4ea.0e10.e465.00-03      1        371      0xafc     0xc40    4001
a012.90d3.ec65.00-00      1        451      0x11f8    0xfb1c   4002
a012.90d3.ec65.00-01      1        451      0x11cd    0xb1ec   4002
a012.90d3.ec65.00-02      1        451      0x1073    0xb77f   4002
a012.90d3.ec65.00-03      1        451      0xaf6     0x3619   4002
fca8.41f6.37df.00-00      1        1024     0x76c     0x68d7   7001
3cb1.5bff.5fdf.00-00      1        1147     0x1d1f    0xfaff   7002
7030.1823.7fdf.00-00      1        357      0x1d1b    0x8bad   7003
```

```
7030.1823.9bdf.00-00      1        431      0x1d1b   0xf8a3   7004
d4ea.0efd.e3df.00-00      1        736      0xc1     0x3177   9001
d4ea.0efd.e4ac.00-00      1        737      0xbd     0x9a1a   9002
--------------------------------------------------------------------------------
Level-1 : 19 out of 19 Total Num of LSP Entries
Level-2 : 0 out of 0 Total Num of LSP Entries
--------------------------------------------------------------------------------
```

On each switch, verify the following:

Option	Verify
LSP ID HOST-NAME	For each switch, the LSDB table should have a LSP ID entry for each neighbor including its own LSP ID for a total of seven entries

22.1.3.5 Verify IS-IS LSP Details

Step 1 – Show IS-IS LSDB details

```
show isis lsdb?
  detail show isis lsdb detailed information
  level   show isis lsdb information by level
  local   show isis local lsdb information
  lspid   show isis lsdb information by lspid
  sysid   show isis lsdb information by system-id
  tlv     show isis lsdb by tlv type
  <cr>

show isis lsdb tlv ?
   <1-186>  Enter tlv type: 1(Area Addresses), 3(End System Neighbors), 5(Prefix
            Neighbors), 22(TE Neighbors), 128(IP Addresses), 129(Protocol
            Supported), 135(TE IP Reachability), 137(Host Name), 144(Multi
            Topology), 180(SPBM Instance), 183(ISID), 184(IPVPN
            Reachability),185(IPVPN Multicast), 186 (IPMC Multicast)
```

Results: From 4001

4001: Example showing SPB Host names

```
4001:1#show isis lsdb tlv 137 detail

================================================================================
                                ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: 0080.2dbe.23df.00-00     SeqNum: 0x00001546      Lifetime:    604
        Chksum: 0xba64  PDU Length: 225
        Host_name: 8003
        Attributes:     IS-Type 1
TLV:137 Host_name: 8003


Level-1 LspID: 00e0.7bbc.23df.00-00     SeqNum: 0x00001192      Lifetime:    924
        Chksum: 0x9c3   PDU Length: 173
```

```
        Host_name: 8004
        Attributes:     IS-Type 1
TLV:137 Host_name: 8004


Level-1 LspID: 0024.43b4.e3df.00-00     SeqNum: 0x00000aea      Lifetime:   834
        Chksum: 0x8c35  PDU Length: 841
        Host_name: 8005
        Attributes:     IS-Type 1
TLV:137 Host_name: 8005


Level-1 LspID: 001e.1f48.f3df.00-00     SeqNum: 0x00001556      Lifetime:   788
        Chksum: 0x8ec4  PDU Length: 818
        Host_name: 8006
        Attributes:     IS-Type 1
TLV:137 Host_name: 8006


Level-1 LspID: 00e0.7bb3.07df.00-00     SeqNum: 0x0000023d      Lifetime:   949
        Chksum: 0x1bcc  PDU Length: 466
        Host_name: 8007
        Attributes:     IS-Type 1
TLV:137 Host_name: 8007


Level-1 LspID: d4ea.0e10.e465.00-00     SeqNum: 0x0000144b      Lifetime:  1005
        Chksum: 0x8bd   PDU Length: 124
        Host_name: 4001
        Attributes:     IS-Type 1
TLV:137 Host_name: 4001


Level-1 LspID: a012.90d3.ec65.00-00     SeqNum: 0x000011fa      Lifetime:  1085
        Chksum: 0xf71e  PDU Length: 124
        Host_name: 4002
        Attributes:     IS-Type 1
TLV:137 Host_name: 4002
```

```
Level-1 LspID: fca8.41f6.37df.00-00     SeqNum: 0x0000076d      Lifetime:   784
        Chksum: 0x66d8  PDU Length: 124
        Host_name: 7001
        Attributes:     IS-Type 1
TLV:137 Host_name: 7001


Level-1 LspID: 3cb1.5bff.5fdf.00-00     SeqNum: 0x00001d20      Lifetime:   911
        Chksum: 0xf801  PDU Length: 195
        Host_name: 7002
        Attributes:     IS-Type 1
TLV:137 Host_name: 7002


Level-1 LspID: 7030.1823.7fdf.00-00     SeqNum: 0x00001d1d      Lifetime:  1007
        Chksum: 0x87af  PDU Length: 179
        Host_name: 7003
        Attributes:     IS-Type 1
TLV:137 Host_name: 7003


Level-1 LspID: 7030.1823.9bdf.00-00     SeqNum: 0x00001d1d      Lifetime:  1082
        Chksum: 0xf4a5  PDU Length: 176
        Host_name: 7004
        Attributes:     IS-Type 1
TLV:137 Host_name: 7004


Level-1 LspID: d4ea.0efd.e3df.00-00     SeqNum: 0x000000c2      Lifetime:   506
        Chksum: 0x2f78  PDU Length: 829
        Host_name: 9001
        Attributes:     IS-Type 1
TLV:137 Host_name: 9001


Level-1 LspID: d4ea.0efd.e4ac.00-00     SeqNum: 0x000000be      Lifetime:   508
        Chksum: 0x981b  PDU Length: 796
        Host_name: 9002
```

```
        Attributes:     IS-Type 1
TLV:137 Host_name: 9002
```

4001: Example, to view ISIS adjacencies in reference to SPB bridge 9001

```
4001:1#show isis lsdb lspid d4ea.0efd.e3df.00-00 tlv 22 detail

================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: d4ea.0efd.e3df.00-00     SeqNum: 0x000000c2      Lifetime:   773
        Chksum: 0x2f78  PDU Length: 829
        Host_name: 9001
        Attributes:     IS-Type 1
TLV:22  Extended IS reachability:
        Adjacencies: 4
        TE Neighbors: 4
                0080.2dbe.23df.00 (8003)        Metric:10
                        SPBM Sub TLV:
                                port id: 194 num_port 1
                                Metric: 10
                d4ea.0efd.e4ac.00 (9002)        Metric:10
                        SPBM Sub TLV:
                                port id: 6144 num_port 1
                                Metric: 10
                a012.90d3.ec65.00 (4002)        Metric:10
                        SPBM Sub TLV:

                                port id: 220 num_port 1
                                Metric: 10
                d4ea.0e10.e465.00 (4001)        Metric:10
                        SPBM Sub TLV:
                                port id: 6151 num_port 1
                                Metric: 10
```

22.1.3.6 Verify CFM Configuration

Step 1 – Verify CFM Maintenance Domain

```
show cfm maintenance-domain
```

Results: The following is shown from 8003 perspective which should be the same on all switches

4001:

```
================================================================================
                          Maintenance Domain
================================================================================
Domain Name              Domain Index    Level Domain Type
--------------------------------------------------------------------------------
spbm                     1               4     NODAL

 Total number of Maintenance Domain entries: 1.
```

Step 2 – Verify CFM Maintenance Association Configuration and Status

```
show cfm maintenance-association
```

Results: The following is shown from 4001 perspective which should be the same on all switches

4001:

```
================================================================================
                        Maintenance Association Status
================================================================================
Domain Name             Assn Name               Domain Idx  Assn Idx
--------------------------------------------------------------------------------
spbm                    4051                    1           1
spbm                    4052                    1           2

 Total number of Maintenance Association entries: 2.

================================================================================
                        Maintenance Association config
================================================================================
Domain Name             Assn Name
--------------------------------------------------------------------------------
spbm                    4051
spbm                    4052

 Total number of MA entries: 2.
```

Step 3 – Verify CFM Maintenance Endpoint Configuration and Status

```
show cfm maintenance-endpoint
```

Results: The following is shown from 8003 perspective; the information should be the same on all switches except for the MEP ID (1 for 9001, 2 for 9002, 3 for 8003, 4 for 8004, 5 for 8005, 6 for 8006, 7 for 8007)

4001:

```
================================================================================
                      Maintenance Endpoint Config
================================================================================
DOMAIN                ASSOCIATION           MEP  ADMIN
NAME                  NAME                  ID
--------------------------------------------------------------------------------
spbm                  4051                  401   enable
spbm                  4052                  401   enable
 Total number of MEP entries: 2.
================================================================================
                      Maintenance Endpoint Service
================================================================================
DOMAIN_NAME           ASSN_NAME             MEP_ID TYPE    SERVICE_DESCRIPTION
--------------------------------------------------------------------------------
spbm                  4051                  401    nodal   Vlan 4051, Level4
spbm                  4052                  401    nodal   Vlan 4052, Level4

 Total number of MEP entries: 2.
```

On 4001 as used in this example, verify the following information:

Option	Verify
DOMAIN NAME	Should be displayed with a name of ***spbm*** as configured in this example
Assn Name ASSOCIATION NAME	Should be displayed with a name of ***4051*** and ***4052*** as configured in this example
SERVICE_DESCRIPTION	Should be displayed as ***Vlan 4051 & Vlan 4052, Level 4*** if CFM is operational and configured correctly where Level 4 is the default CFM level

22.1.3.7 Use CFM Command to verify operations

Step 1 – Use L2 ping command to verify network connectivity to neighbors. The neighbor format is BVID.Remote Router Name for CLI

```
l2 ping vlan <vlan id> routernodename <Router Node Name>
```

Results: The following is shown from 9001 perspective pinging switch 4001

4001:

```
4001:1#l2 ping vlan 4051 routernodename 8007

 Please wait for l2ping to complete or press any key to abort

----00:e0:7b:b3:07:df    L2 PING Statistics---- 0(64) bytes of data
1 packets transmitted, 1 packets received,    0.00% packet loss
  round-trip (us)        min/max/ave/stdv = 4479/4479/4479.00/ 0.00

4001:1#l2 ping vlan 4052 routernodename 8007

 Please wait for l2ping to complete or press any key to abort

----00:e0:7b:b3:07:df    L2 PING Statistics---- 0(64) bytes of data
1 packets transmitted, 1 packets received,    0.00% packet loss
  round-trip (us)        min/max/ave/stdv = 2996/2996/2996.00/ 0.00
```

Step 2 – Use L2 traceroute command to verify network route to neighbors

```
l2 traceroute vlan <vlan id> routernodename <Router Node Name>
```

Results: The following is shown from 4001 perspective to switch 7001

4001:

```
4001:1#l2 traceroute vlan 4051 routernodename 7001

 Please wait for l2traceroute to complete or press any key to abort

l2traceroute to 7001  (3c:b1:5b:ff:5f:e0),  vlan 4051
0    4001                (d4:ea:0e:10:e4:65)
1    9001                (d4:ea:0e:fd:e3:df)
2    8003                (00:80:2d:be:23:df)
3    8005                (00:24:43:b4:e3:df)
4    7001                (3c:b1:5b:ff:5f:e0)
```

Step 3 – Use L2 traceroute command to verify network route to neighbors; for example, diverse route to a SMLT virtual B-MAC

```
l2 traceroute vlan <vlan id> mac <Mac>
```

Results: The following is shown from 4001 perspective to SMLT virtual B-MAC of SMLT cluster 8005 & 8006

4001:

```
4001:1#l2 traceroute vlan 4051 mac 00:1e:1f:48:f3:e0
 Please wait for l2traceroute to complete or press any key to abort
l2traceroute to   (00:1e:1f:48:f3:e0),  vlan 4051
0    4001                (d4:ea:0e:10:e4:65)
1    9001                (d4:ea:0e:fd:e3:df)
2    8003                (00:80:2d:be:23:df)
3    8005                (00:24:43:b4:e3:df)

4001:1#l2 traceroute vlan 4052 mac 00:1e:1f:48:f3:e0
 Please wait for l2traceroute to complete or press any key to abort
l2traceroute to   (00:1e:1f:48:f3:e0),  vlan 4052
0    4001                (d4:ea:0e:10:e4:65)
1    9002                (d4:ea:0e:fd:e4:ac )
2    8004                (00:e0:7b:bc:23:df)
3    8006                (00:1e:1f:48:f3:df)
```

Verify the following information:

Option	Verify
L2 PING Statistics	If everything has been configured correctly and during normal operations, the packets received should display ***0.00% packet loss***
LBMs lost	If everything has been configured correctly and during normal operations, the LBMs loss should display ***0.00%***

22.1.3.8 Verify vIST

Step 1 – Verify the vIST is up and running

show virtual-ist

Results:

4001:

```
================================================================================
                                  IST Info
================================================================================
PEER-IP              VLAN     ENABLE    IST
ADDRESS              ID       IST       STATUS
--------------------------------------------------------------------------------
10.4.2.2             2        true      up


NEGOTIATED                                MASTER/
DIALECT      IST STATE                    SLAVE
--------------------------------------------------------------------------------
v4.0         Up                           Slave
```

Step 2 – Verify the vIST statistics

show virtual-ist stat

Results:

4001:

```
================================================================================
                           IST Message Statistics
================================================================================
PROTOCOL MESSAGE            COUNT
```

```
Ist Down                 :  5
Hello Sent               :  94208
Hello Recv               :  94099
Learn MAC Address Sent   :  10663
Learn MAC Address Recv   :  4596
MAC Address AgeOut Sent  :  1286
MAC Address AgeOut Recv  :  248
MAC Address Expired Sent :  0
MAC Address Expired Recv :  0
Delete Mac Address Sent  :  2
Delete Mac Address Recv  :  0
Smlt Down Sent           :  0
Smlt Down Recv           :  1
Smlt Up Sent             :  8
Smlt Up Recv             :  6
Send MAC Address Sent    :  4
Send MAC Address Recv    :  4
IGMP Sent                :  3
IGMP Recv                :  0
Port Down Sent           :  0
Port Down Recv           :  0
Request MAC Table Sent   :  4
Request MAC Table Recv   :  5
Unknown Msg Type Recv    :  0
Mlt Table Sync Req Sent  :  0
Mlt Table Sync Req Recv  :  0
Mlt Table Sync Sent      :  4
Mlt Table Sync Recv      :  4
Port Update Sent         :  0
Port Update Recv         :  0
Entry Update Sent        :  0
Entry Update Recv        :  0
Dialect Negotiate Sent   :  6
Dialect Negotiate Recv   :  4
Update Response Sent     :  0
Update Response Recv     :  0
Transaction Que HiWaterM :  0
Poll Count Hi Water Mark :  61
```

22.2 SMLT Configuration

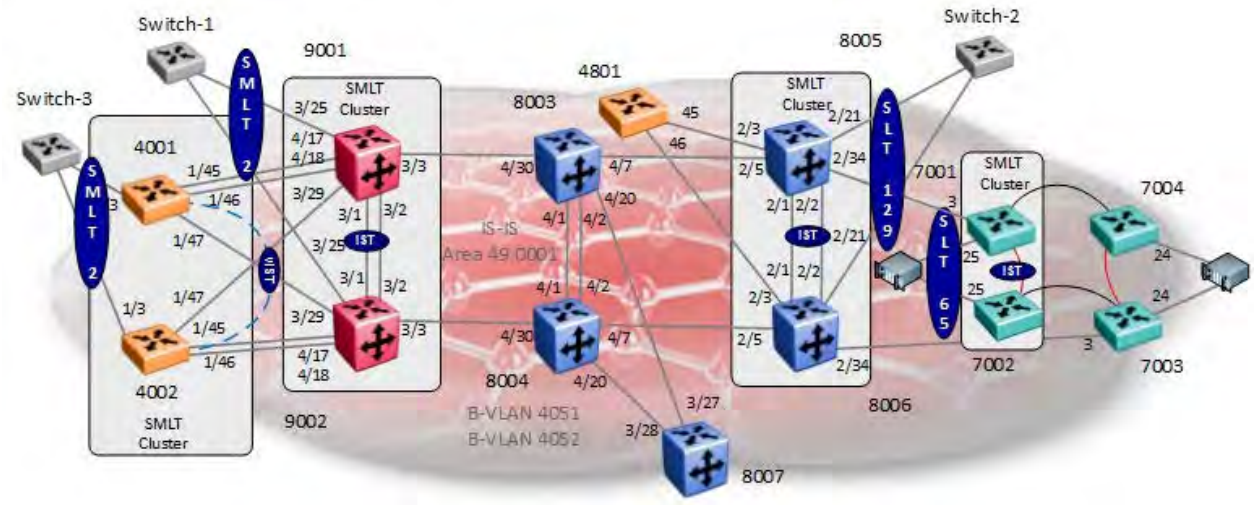

Assuming the edge switches are Extreme stackable switches, we will also enable VLACP, VLAN tagging, SLPP, and untagged frames discard as per the SMLT best practices. For this example, we will create SMLT id 2 on the SMLT cluster 9001 & 9002, SMLT id 2 on the SMLT cluster 4001 & 4001, SLT 129 on the SMLT cluster 8005 & 8006, and SLT 65 on the SMLT cluster 7001 & 7002.

4001 & 4002 SMLT Cluster Switches – SMLT on port 1/3

4001 & 4002: Same configuration on both switches

```
4001:1(config)#mlt 2 enable
4001:1(config)#mlt 2 member 1/3
4001:1(config)#mlt 2 encapsulation dot1q
4001:1(config)#interface mlt
4001:1(config-mlt)#smlt
4001:1(config-mlt)#exit
4001:1(config)#vlan members remove 1 1/3
4001:1(config)#interface gigabitEthernet 1/3
4001:1(config-if)#no shutdown
```

```
As per SMLT best practices, we will also enable VLACP, untagged frames discard,
and SLPP. Note that VLACP will have to also be enabled on Switch-3.
```

```
4001:1(config)#interface GigabitEthernet 1/3
4001:1(config-if)#untagged-frames-discard
4001:1(config-if)#vlacp fast-periodic-time 500 timeout short timeout-scale 5
funcmac-addr 01:80:c2:00:00:0f
4001:1(config-if)#vlacp enable
4001:1(config-if)#slpp
```

```
4001:1(config-if)#slpp packet-rx packet-rx-threshold 5
4001:1(config-if)#exit
4001:1(config)#slpp enable
4001:1(config)#slpp vid 1155
```

For 4002, use the same configuration as above except for the items shown below

---

```
4002:1(config-if)#slpp packet-rx packet-rx-threshold 50
```

8005 & 8006 SMLT Cluster Switches – SLT on port 2/21 using SLT id 129

8005 & 8006: Same configuration on both switches

```
8005:5(config)#interface GigabitEthernet 2/21
8005:5(config-if)#encapsulation dot1q
8005:5(config-if)#smlt 129
8005:5(config-if)#exit
8005:5(config)#vlan members remove 1 2/21
```

As per SMLT best practices, we will also enable VLACP, untagged frames discard, and SLPP. Note that VLACP will have to also be enabled on Switch-2.

```
8005:5(config)#interface GigabitEthernet 2/21
8005:5(config-if)#untagged-frames-discard
8005:5(config-if)#vlacp fast-periodic-time 500 timeout short timeout-scale 5
funcmac-addr 01:80:c2:00:00:0f
8005:5(config-if)#vlacp enable
8005:5(config-if)#slpp
8005:5(config-if)#slpp packet-rx packet-rx-threshold 5
8005:5(config-if)#exit
8005:5(config)#slpp enable
8005:5(config)#slpp vid 2256
```

For 8006, use the same configuration as above except for the items shown below

---

```
8006:5(config-if)#slpp packet-rx packet-rx-threshold 50
```

9001 & 9002 SMLT Cluster Switches – SMLT on port 3/25

9001 & 9002: Same configuration on both switches

```
9001:1(config)#mlt 2 enable
9001:1(config)#mlt 2 member 3/25
9001:1(config)#mlt 2 encapsulation dot1q
9001:1(config-mlt)#interface mlt 2
9001:1(config-mlt)#smlt
9001:1(config-mlt)#exit
9001:1(config)#vlan members remove 1 3/25
9001:1(config)#interface gigabitEthernet 3/25
9001:1(config-if)#no shutdown
```

As per SMLT best practices, we will also enable VLACP, untagged frames discard, and SLPP. Note that VLACP will have to also be enabled on Switch-1.

```
9001:1(config)#interface gigabitEthernet 3/25
9001:1(config)#untagged-frames-discard
9001:1(config)#slpp packet-rx
9001:1(config-if)#vlacp fast-periodic-time 500 timeout short timeout-scale 5 funcmac-addr 01:80:c2:00:00:0f
9001:1(config-if)#vlacp enable
9001:1(config-if)#exit
9001:1(config)#slpp enable
9001:1(config)#slpp vid 1155
```

For 9002, use the same configuration as above except for the items shown below

---

```
9002:1(config-if)#slpp packet-rx packet-rx-threshold 50
```

7001 & 7002 SMLT Cluster Switches - SLT on port 25 using SLT id 65 assuming the server uses an tagged port

7001 & 7002: Same configuration on both switches

```
7001(config)#vlan ports 25 tag tagall
7001(config)#vlan configcontrol automatic
7001(config)#vlan members remove 1 25
7001(config)#interface ethernet 25
7001(config-if)#smlt 65
7001(config-if)#exit
```

22.2.1 Verify Operations

Assuming a local VLAN have been provisioned on the SMLT cluster switches and also on the edge switch or server (next step for example via the L2VSN configuration), the SMLT state should be up and operational.

22.2.1.1 Verify SMLT

Step 1 – Show SMLT state

```
show smlt mlt
show smlt gigabitethernet (on ERS 8000 to display SLT)
show smlt ethernet (on VSP 7000 to display SLT)
EDM
Configuration -> IS-IS -> SPBM -> ISID
```

Results:

```
4001: (4002 will be the same)
================================================================================
                                Mlt SMLT Info
================================================================================
MLT    ADMIN    CURRENT
ID     TYPE     TYPE
--------------------------------------------------------------------------------
2      smlt     smlt

9001: (9002 will be the same)
================================================================================
                                Mlt SMLT Info
================================================================================
MLT    ADMIN    CURRENT
ID     TYPE     TYPE
--------------------------------------------------------------------------------
2      smlt     smlt

8005: (8006 will be the same)
================================================================================
                                  SMLT Info
================================================================================
PORT   SMLT     ADMIN    CURRENT
NUM    ID       TYPE     TYPE
```

```
--------------------------------------------------------------------------------------------------------------
2/21  129      smlt     smlt

7001: (7002 will be the same)
==============================================================================
                                        SLT Info
==============================================================================
PORT  SMLT     ADMIN    CURRENT
NUM   ID       TYPE     TYPE
------------------------------------------------------------------------------------------------------------
25      65      slt     slt
```

On each switch, verify the following:

Option	Verify
CURRENT TYPE	Depending on the switch, either ***SMLT*** or ***SLT*** should be displayed indicating that the local VLAN, port members, and SMLT ID are configured on both switches and the correct VLAN and ports are also configured on the edge switch

22.3 SPB L2 VSN Configuration

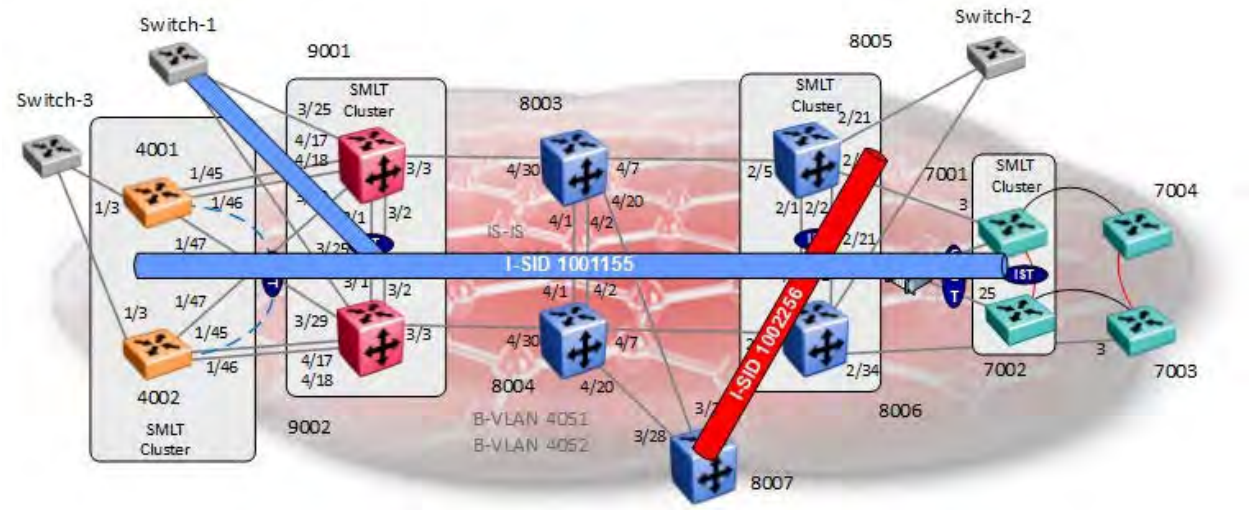

For this example, we will configure the following:

- L2 VSN VLANs
  - VLAN ID = 1155 configured on SMLT cluster switches 4001 & 4002, 9001 & 9002, and 7001 & 7002 using ISID = 1001155
    - The L2VSN is provisioned on SMLT cluster switches 9001 & 9002 for edge switch Switch-1
    - The L2VSN is provisioned on SMLT cluster switches 4001 & 4002 for edge switch Switch-3
    - The L2VSN is provisioned on SMLT cluster switches 7001 & 7002 for edge server
  - VLAN ID = 2256 configured on switches 8005, 8006 and 8007 using ISID = 1002256
    - The L2VSN is provisioned on SMLT cluster switches 8005 & 8006 for edge switch Switch-2

This example is a continuation from the base setup used in Section 22.1 and 22.2 for the SMLT configuration.

22.3.1 VLAN configuration

4001 & 4002 SMLT Cluster Switches

4001 & 4002: Same configuration on both switches

```
4001:1(config)#vlan create 1155 type port-mstprstp 0
4001:1(config)#vlan mlt 1155 2
```

8005 & 8006 SMLT Cluster Switches

8005 & 8006: Same configuration on both switches

```
8005:5(config)#vlan create 2256 name VSN-Red type port-mstprstp 0
8005:5(config)#vlan members add 2256 2/21
```

8007

```
8007:5(config)#vlan create 2256 name VSN-Red type port-mstprstp 0
8007:5(config)#vlan ports 4/25 tagging tagall
8007:5(config)#vlan members add 2256 4/35
8007:5(config)#vlan members remove 1 4/35
```

9001 & 9002 SMLT Cluster Switches

9001 & 9002: Same configuration on both switches

```
9001:1(config)#vlan create 1155 name VSN-Blue type port-mstprstp 0
9001:1(config)#vlan mlt 1155 2
9001:1(config)#vlan mlt 1155 1
```

7001 & 7002 SMLT Cluster Switches - SLT on port 25 using SLT id 65 assuming the server uses an untagged port

7001 & 7002: Same configuration on both switches

```
7001(config)#vlan create 1155 name VSN-Blue type port
7001(config)#vlan configcontrol automatic
7001(config)#vlan members add 1155 10
```

22.3.2 Layer 2 VSN configuration

VSP 4000 Switches

4001 & 4002: Same configuration on both switches

```
4001:1(config)#vlan ISID 1155 1001155
```

VSP 7000 Switches

7001 & 7002: Same configuration on both switches

```
7001(config)#vlan ISID 1155 1001155
```

SMLT Cluster Switches – 8005 & 8006

8005 & 8006: Same configuration on both switches

```
8005:5(config)#vlan ISID 2256 1002256
```

ERS 8800 Switch - 8007

```
8007:5(config)#vlan ISID 2256 1002256
```

VSP 9000 Switches

9001 & 9002: Same configuration on 9001 and 9002

```
9001:1(config)# vlan ISID 1155 1001155
```

22.3.2.1 Verify IS-IS ISID

Step 1 – Show IS-IS ISID

```
show isis spbm ISID all
show isis spbm ISID all <id|nick-name|vlan>
```

EDM

Configuration -> IS-IS -> SPBM -> ISID

Results:

4001: (4002 will be the same)

```
================================================================================
                                SPBM ISID INFO
================================================================================
ISID      SOURCE NAME    VLAN    SYSID                TYPE         HOST_NAME
--------------------------------------------------------------------------------------------------------------
1001155 0.40.01         4051    d4ea.0e10.e465       config       4001
1001155 0.40.02         4052    a012.90d3.ec65       discover     4002
1001155 0.70.01         4051    fca8.41f6.37df       discover     7001
1001155 0.70.02         4052    3cb1.5bff.5fdf       discover     7002
1001155 0.90.01         4051    d4ea.0efd.e3df       discover     9001
1001155 0.90.02         4052    d4ea.0efd.e4ac       discover     9002
```

8005: (8006 will be the same)

```
================================================================================
                                SPBM ISID INFO
================================================================================
ISID    SOURCE NAME     VLAN    SYSID                TYPE         HOST_NAME
--------------------------------------------------------------------------------------------------------------
1002256   0.80.05        4051    0024.43b4.e3df       config       8005
1002256   0.80.06        4052    001e.1f48.f3df       discover     8006
1002256   0.80.07        4052    00e0.7bb3.07df       discover     8007
```

8007:

```
================================================================================
                                SPBM ISID INFO
================================================================================
ISID    SOURCE NAME     VLAN    SYSID                TYPE         HOST_NAME
--------------------------------------------------------------------------------------------------------------
1002256 0.80.05         4051    0024.43b4.e3df       discover     8005
```

```
1002256 0.80.06         4052    001e.1f48.f3df       discover      8006
1002256 0.80.07         4052    00e0.7bb3.07df       config        8007
```

On each switch, verify the following:

Option	Verify
ISID TYPE	For switches 4001, 4002, 7001, 7002, 9001, and 9002, for example, in reference to ISID ***1001155***, TYPE should show ***config*** in reference its own SYSID and ***discover*** to each neighbor. For switches 8005, 8006 and 8007, for example, in reference to ISID ***1002256***, TYPE should show ***config*** in reference its own SYSID and ***discover*** to each neighbor.

22.3.2.2 Show IS-IS LSP Details pertaining to ISIDs provisioned

In an IS-IS network, each IS router advertises one or more IS-IS Link State Protocol Data Units (LSPs) with routing information. Within each LSP, there is a fixed header and a number of TLVs with encoded information. The following command is used to show details of a LSP in detail to a specific neighbor displaying the encoded information in the TLVs.

Step 1 – Show IS-IS ISID

```
show isis lsdb lspid <is-is system id>.00-00 detail
show isis lsdb lspid <is-is system id>.00-00 tlv 144 sub-tlv 3 detail
```

Results: From 9001 for perspective for 7001

9001:

```
9001:1#show isis lsdb lspid fca8.41f6.37df.00-00 tlv 144 sub-tlv 3 detail
================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: fca8.41f6.37df.00-00     SeqNum: 0x000007d7      Lifetime:   406
        Chksum: 0x986d  PDU Length: 160
        Host_name: 7001
        Attributes:     IS-Type 1
TLV:144 SUB-TLV 3       ISID:
                Instance: 0
                Metric: 0
                B-MAC: fc-a8-41-f6-37-df
                BVID:4051
                Number of ISID's:1
```

```
                1001155(Both)


        Instance: 0
        Metric: 0
        B-MAC: 00-e0-7b-b3-07-df
        BVID:4052

        Number of ISID's:1
                1001155(Rx)
```

In reference to 7001 as used in this example from 9001, verify the following:

Option	Verify
Level 1	As this example is in reference 7001, IS-IS LDP ID of ***fca8.41f6.37df.00-00*** should be displayed with its Host Name of ***7001***.
TLV:144 Sub-tlv 3	TLV 144 sub-tlv 3 is the SPBM ISID TLV. The SPBM instance is set to ***0*** indicating only one instance is supported today. The BMAC entry is the advertising MAC address which for this example should be the BMAC of 8005 displayed as ***fc-a8-41-f6-37-df***. For BVID, VLAN IDs of ***4051*** and ***4052*** should be displayed as these are the two VLANs used in this configuration with BVID 4051 being the primary. For each BVID, ISID ***1001155*** should be displayed

22.3.2.3 Unknown unicast or multicast/broadcast traffic

The multicast addresses are built out of two pieces. Each SPB node must be configured with a unique Nickname that is carried in the IS-IS link state database and is used to form the first portion of the multicast MAC address (with the multicast bit set: multicast address is Nickname & “3”). The second portion is the ISID id converted to hex forming the Multicast MAC address.

For example, in reference to 8005:

- Nickname = 0.80.05
- ISID = 1002256 (0x0f:4b:10)
- Multicast address = 03:08:05:0f:4b:10 for ISID 1002256

Step 1 – Display Multicast address used for unknown unicast or multicast/broadcast traffic

```
show isis spbm multicast-fib
```

EDM

```
Configuration -> IS-IS -> SPBM -> Multicast FIB
```

Results: The following is shown from 8005

```
8005:3#show isis spbm multicast-fib ISID 1002256
================================================================================
                        SPBM MULTICAST FIB ENTRY INFO
================================================================================
MCAST DA         ISID     BVLAN  SYSID            HOST-NAME      OUTGOING-INTERFACES
--------------------------------------------------------------------------------
03:08:05:0f:4b:10  1002256  4051  0024.43b4.e3df   8005           2/5,IST,2/21
03:08:06:0f:4b:10  1002256  4052  001e.1f48.f3df   8006
03:08:07:0f:4b:10  1002256  4052  00e0.7bb3.07df   8007
```

On each switch, verify the following information:

Option	Verify
MCAST DA	Verify that the correct multicast address for each switch ISID 10011155 • 4001: ***03:04:01:0f:46:c3*** • 9001: ***03:09:01:0f:46:c3*** • 9002: ***03:09:02:0f:46:c3*** • 7001: ***03:07:01:0f:46:c3*** • 7002: ***03:07:02:0f:46:c3*** ISID 1002256 • 8005: ***03:08:05:0f:4b:10*** • 8006: ***03:08:06:0f:4b:10*** • 8007: ***03:08:07:0f:4b:10***

22.3.2.4 MAC Address Table

Step 1 – Display MAC address table, local or remote

ERS 8800 & VSP 9000:

```
show vlan mac-address-entry <vlan id>
show vlan mac-address-entry 2256
```

EDM

```
Configuration -> VLAN -> VLANs -> Forwarding
```

VSP 7000:

```
show mac-address-table spbm
show mac-address-table spbm ISID <1-16777215>
```

Results: The following is shown from 4001 and 7001

```
4001:1#show vlan mac-address-entry 1155

================================================================================
                                   Vlan Fdb
================================================================================
VLAN            MAC                             SMLT
ID   STATUS     ADDRESS            INTERFACE  REMOTE TUNNEL
1155 learned     00:0c:29:35:62:a4  9001       false  7002
1155 learned     00:0c:29:9b:a8:31  9001       false  9001
1155 learned     00:0c:29:d6:81:e5  MLT-2      false  -
1155 learned     00:18:71:ea:31:bb  9001       false  7001

4001:1#show vlan remote-mac-table 1155

================================================================================
                                Vlan Remote Mac Table
================================================================================
VLAN STATUS  MAC-ADDRESS        DEST-MAC           BVLAN DEST-SYSNAME      PORTS    SMLTREMOTE
--------------------------------------------------------------------------------------------
1155 learned 00:0c:29:35:62:a4  3c:b1:5b:ff:5f:df  4051  7002              9001     false
1155 learned 00:0c:29:9b:a8:31  d4:ea:0e:fd:e3:df  4051  9001              9001     false
1155 learned 00:18:71:ea:31:bb  3c:b1:5b:ff:5f:e0  4051  7001              9001     false

7001#show mac-address-table spbm ISID 1001155
Mac Address Table Aging Time: 300
```

```
Learning Enabled Ports ALL
Number of addresses: 4

MAC Address       ISID    Source    Vid  BVid Dest-MAC          Dest-Sys-Name
----------------------- ----------- ------------ ------ ------ ----------------------- ------------------
00-0C-29-35-62-A4 1001155  Trunk 31       4051 3C-B1-5B-FF-5F-DF 7002
00-0C-29-9B-A8-31 1001155  Trunk 31       4051 D4-EA-0E-FD-E3-DF 9001
00-0C-29-D6-81-E5 1001155  Trunk 31       4051 D4-EA-0E-10-E4-65 4001
00-18-71-EA-31-BB 1001155  Port 10   1155
```

In reference to each switch, verify the following information:

Option	Verify
MAC Address INTERFACE	The MAC address displayed will vary depending on the MAC address of the end-user device. The interface should display ***ISID-1001155*** for MAC addressed from ISID 1001155 and ***ISID-1002256*** for MAC addressed from ISID 1002256
DEST-MAC	For remote entries, the remote B-MAC address of the SPB switch should be shown as the remote destination MAC.

22.4 VSP 7000 & ERS 4800 – In-band Management via L2VSN

An L2VLSN can be created to provide in-band management for the VSP 7000 and ERS 4800. For example, let's assume we wish to use the 10.12.11/0/24 subnet to manage the VSP 7000 and ERS 4800. On bridges 8005 and 8006, we will enable VRRP with backup-master to provide routing to the rest of the network. We will also have to enable IP Shortcuts on both 8005 and 8006 – please see section 17.10.

Switch	Parameter	Value
L2VSN – for in-band management		
8005, 8006	Mgmt VLAN	101
7001, 7002, 7003, 7004 4801	ISID	1000101
8005	IP Address	10.12.11.2/24
8006	IP Address	10.12.11.3/24
7001	IP address	10.12.11.11/24
7002	IP address	10.12.11.12/24
7003	IP address	10.12.11.13/24
7004	IP address	10.12.11.14/24
4801	IP address	10.12.11.15/24
IP Configuration – 8005 and 8006		
8005 8006	VRRP ID	11
	VRRP VIP	10.12.11.1
	Backup Master	Enable
8005	VRRP Priority	150

Please note, for the VSP 7000, if you also use the out-of-management management interface, you cannot have two default gateways – that is one for the in-band and another for the out-of-band management interfaces. If you also use the out-of-band management interface, please use static routes and use a default route on the in-band interface.

VSP 7000: Add in-band L2VSN and IP address

```
7001(config)#vlan create 101 name mgmt-101 type port
7001(config)#vlan mgmt 101
7001(config)#ip address 10.12.11.11 netmask 255.255.255.0 default-gateway 10.12.11.1
7001(config)#vlan ISID 101 1000101
```

For switches 7002, 7003, and 7004, use the same configuration as above except for the items shown below

```
7002(config)#ip address 10.12.11.12 netmask 255.255.255.0 default-gateway 10.12.11.1
7003(config)#ip address 10.12.11.13 netmask 255.255.255.0 default-gateway 10.12.11.1
7004(config)#ip address 10.12.11.14 netmask 255.255.255.0 default-gateway 10.12.11.1
```

```
7004(config)#show ip
Bootp/DHCP Mode: Disabled
                          Configured        In Use        Last BootP/DHCP
                        ---------------- ---------------- ----------------
Stack IP Address:       0.0.0.0                           0.0.0.0
Switch IP Address:      10.12.11.14       10.12.11.14     0.0.0.0
Switch Subnet Mask:     255.255.255.0     255.255.255.0   0.0.0.0
Mgmt Stack IP Address:  0.0.0.0
Mgmt Switch IP Address: 10.136.56.54      10.136.56.54
Mgmt Subnet Mask:       255.255.255.0     255.255.255.0
Mgmt Def Gateway:       0.0.0.0
Default Gateway:        10.12.11.1        10.12.11.1      0.0.0.0
```

ERS 4800: Add in-band L2VSN and IP address

```
4801(config)#vlan create 101 name mgmt-101 type port
4801(config)#vlan mgmt 101
4801(config)#ip address 10.12.11.15 netmask 255.255.255.0 default-gateway 10.12.11.1
4801(config)#vlan ISID 101 1000101
```

```
4801(config)#show ip
Bootp/DHCP Mode: Disabled

                          Configured        In Use        Last BootP/DHCP
                        ------------------- ------------------- ---------------------------
Stack IP Address:       0.0.0.0                         0.0.0.0
Switch IP Address:      10.12.11.15     10.12.11.15     0.0.0.0
Switch Subnet Mask:     255.255.255.0   255.255.255.0   0.0.0.0
Default Gateway:        10.12.11.1      10.12.11.1      0.0.0.0
```

8005 & 8006: Add in-band L2VSN and IP address

```
8005:5(config)#vlan create 101 type port-mstprstp 0
8005:5(config)#vlan ISID 101 1000101
8005:5(config)#interface Vlan 101
8005:5(config-if)#ip address 10.12.11.2 255.255.255.0
8005:5(config-if)#ip vrrp address 11 10.12.11.1
8005:5(config-if)#ip vrrp 11 backup-master enable priority 150
8005:5(config-if)#ip vrrp 11 enable
8005:5(config-if)#exit
```

For 8006, use the same configuration as above except for the items shown below

```
8005:5(config-if)#ip address 10.12.11.3 255.255.255.0
8005:5(config-if)#ip vrrp 11 backup-master enable
```

22.5 Multicast over L2VSN

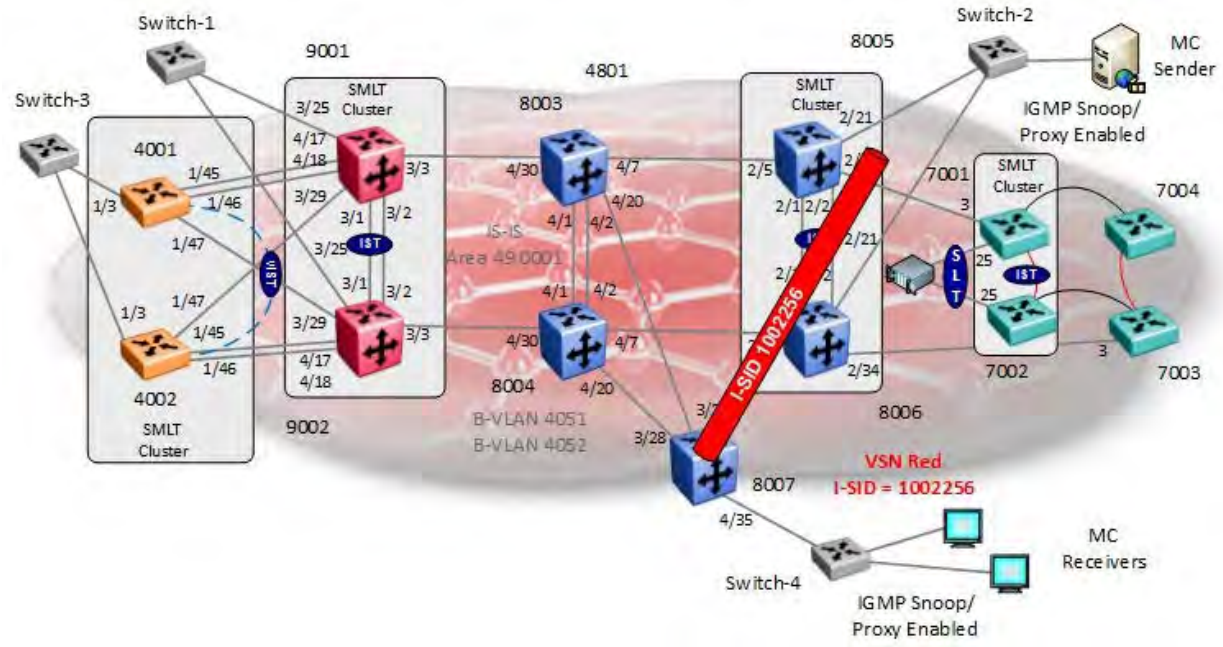

Continuing from example used in Section 17.2, we will simple enable multicast support for L2VSN ISID 1002256.

22.5.1 Enable SPB Multicast – Global

ERS 8800 Switches

8005, 8006 & 8007: Same configuration on all switches

```
8005:5(config)#router isis
8005:5(config-isis)#spbm 1 multicast enable
8005:5(config-isis)#exit
```

22.5.2 Enable IGMP

22.5.2.1 Enable IGMPv2 at VLAN level

ERS 8800 Switches

8005 & 8006: Same configuration on both switches

```
8005:5(config)#interface vlan 2256
8005:5(config-if)#ip igmp proxy
8005:5(config-if)#ip igmp snooping
8005:5(config-if)#ip igmp snoop-querier-addr 192.168.156.1
8005:5(config-if)#exit
## If IGMPv3 is used:
8005:5(config-if)#ip igmp ssm-snoop
8005:5(config-if)#ip igmp version 3
```

8007:

```
8007:5(config)#interface vlan 2256
8007:5(config-if)#ip igmp proxy
8007:5(config-if)#ip igmp snooping
8007:5(config-if)#ip igmp snoop-querier-addr 192.168.56.1
8007:5(config-if)#exit
## If IGMPv3 is used:
8007:5(config-if)#ip igmp ssm-snoop
8007:5(config-if)#ip igmp version 3
```

Please note, if the ERS 8800 is connected to an edge switch, it may be necessary to add an IGMP query address. If you omit adding a query address, the ERS 8800 will send IGMP queries with a source address of 0.0.0.0. Depending on the edge switch model, it may not accept a query with a source address of 0.0.0.0.

22.5.2.2 Edge Switch

Assuming the edge switch is an Extreme stackable switch with the latest firmware, enable IGMP snoop and proxy.

Extreme Stackable Switches

```
ACLI
ERS-Stackable(config)#interface vlan 2256
ERS-Stackable(config-if)#ip igmp snoop
ERS-Stackable(config-if)#ip igmp proxy
## If IGMPv3 is used:
ERS-Stackable(config-if)#ip igmp version 3
```

22.5.3 Verify Operations

22.5.3.1 Global Settings

Step 1 – Verify SPB multicast is enabled

```
show isis spbm multicast
```

Results:

8007: (8005 and 8006 should be the same)

```
================================================================================
                                ISIS SPBM Info
================================================================================
SPBM        B-VID       PRIMARY     NICK      LSDB      IP        MULTICAST
INSTANCE                VLAN        NAME      TRAP
--------------------------------------------------------------------------------
1           4051-4052   4051        0.80.07   disable   enable    enable
```

Step 2 – Verify IGMP interfaces

```
show ip igmp interface
```

Results:

8007: (results from 8005 and 8006 will be same except for the querier address)

```
================================================================================
                        IGMP Interface - GlobalRouter
================================================================================
       QUERY               OPER             QUERY    WRONG             LASTMEM
IF     INTVL STATUS VERS.  VERS QUERIER      MAXRSPT QUERY JOINS ROBUST QUERY MODE
--------------------------------------------------------------------------------
V2256  125    active 2     2    192.168.56.1 100     0     12    2      10    snoop-spb
```

22.5.3.2 Verify IGMP cache/group and senders

Assuming the multicast sender is using IGMPv3 (source IP 10.5.41.20@232.2.2.2) connect to Switch-2 off SPB bridges 8005 & 8006 with a receiver (10.5.41.10) connected to Switch-3 off SPB bridges 8007.

Step 1 – Verify SPB multicast is enabled

```
show ip igmp cache
show ip igmp group
```

Results:

8007:

```
8007:3#show ip igmp cache
================================================================================
                          IGMP Cache - GlobalRouter
================================================================================
GRPADDR         INTERFACE  LASTREPORTER    EXPIRYTIME        VERSION1HOSTTIMER TYPE
STATICP
ORTS
--------------------------------------------------------------------------------------------------------
232.2.2.2       Vlan2256   10.7.30.5       0day,00h:04m:08s  0day,00h:00m:00s    DYNAMIC NULL


8007:3#show ip igmp group
================================================================================
                          IGMP Group - GlobalRouter
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------------------------------
232.2.2.2       V2256-4/35      10.5.41.10      54         Dynamic
```

Step 2 – Verify IGMP sender

```
show ip igmp sender
```

Results:

8007:

```
8007:3#show ip igmp sender
================================================================================
                          IGMP Sender - GlobalRouter
================================================================================
                                        PORT/
GRPADDR         IFINDEX    MEMBER       MLT        STATE
--------------------------------------------------------------------------------------------------------
232.2.2.2       Vlan 2256  10.5.41.20   spb        NOTFILTERED
```

22.5.3.3 Verify SPB Multicast Routes

Step 1 – Verify all SPB multicast routes

```
show isis spbm ip-multicast-route all
```

Results:

8007:

```
8007:3#show isis spbm ip-multicast-route all
================================================================================================
                                        SPBM IP-MULTICAST ROUTE INFO ALL
================================================================================================
Type     VrfName     Vlan  Source          Group      VSN-ISID    Data ISID   BVLAN Source-
BEB
                     Id
------------------------------------------------------------------------------------------------
snoop    GRT         2256  10.5.41.20      232.2.2.2  1002256     16000002    4051  8005
snoop    GRT         2256  10.5.41.20      232.2.2.2  1002256     16000002    4052  8006
------------------------------------------------------------------------------------------------
```

Step 2 – Verify SPB multicast routes pertaining to VLAN 2256 / ISID 1002256

```
show isis spbm ip-multicast-route vlan 2256
show isis spbm ip-multicast-route vsn-isid 1002256
```

Results:

8007:

```
8007:3#show isis spbm ip-multicast-route vlan 2256
================================================================================================
                      SPBM IP-MULTICAST ROUTE INFO - VLAN ID : 2256, VSN-ISID : 1002256
================================================================================================
Source          Group           Data ISID  BVLAN Source-BEB
------------------------------------------------------------------------------------------------
10.5.41.20      232.2.2.2       16000002   4051  8005
10.5.41.20      232.2.2.2       16000002   4052  8006
```

22.5.3.4 Verify multicast TLV's

Assuming the multicast sender is using IGMPv3 (source IP 10.5.41.20@232.2.2.2) connect to Switch-2 off SPB bridges 8005 & 8006 with a receiver (10.5.41.10) connected to Switch-4 off SPB bridges 8007. TLV 185 in relationship to bridges 8005 and 8006 should have the Tx bit set and also send TLV 144 with the Tx bit set. Each multicast group should have its own unique data ISID with a value of 1600000x. The receiver switches (8007) should have TLV 144 with the Rx bit set.

Step 1 – Verify IP multicast source, group addresses, and Tx bit set on the BEB bridge where the multicast source is located via TLV 185

```
show isis lsdb tlv 185 detail
```

Results:

8005:

```
8005:3#show isis lsdb tlv 185 detail
================================================================================
                                ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: 0024.43b4.e3df.00-00     SeqNum: 0x00000bb5      Lifetime:   375
        Chksum: 0xb302  PDU Length: 889
        Host_name: 8005
        Attributes:     IS-Type 1
TLV:185 SPBM IPVPN :
                VSN ISID:1002256
                BVID    :4051
                        Metric:0
                        IP Source Address: 10.5.41.20
                        Group Address    : 232.2.2.2
                        Data ISID        : 16000002
                        TX               : 1


Level-1 LspID: 001e.1f48.f3df.00-00     SeqNum: 0x00001621      Lifetime:   662
        Chksum: 0x4cf8  PDU Length: 866
        Host_name: 8006
        Attributes:     IS-Type 1
TLV:185 SPBM IPVPN :
                VSN ISID:1002256
                BVID     4052
```

```
                        Metric:0
                        IP Source Address: 10.5.41.20
                        Group Address    : 232.2.2.2
                        Data ISID        : 16000002
                        TX               : 1
```

Step 2 – Verify on the BEB bridges where the multicast receivers are located via TLV 144, the Rx bit is set with a B-MAC of 03-08-05-00-00-00 & 03-08-06-00-00-00 (03 indicated multicast while 08-05 & 08-06 are the Nick Names of BEB bridges 8005 & 8006 with the multicast source). On the BEB bridges where the source is located, the Tx bit should be set

```
show isis lsdb tlv 144 detail
show isis lsdb lspid tlv 144 sub-tlv 3 detail
show isis lsdb lspid <lsp id> tlv 144 detail
show isis lsdb lspid <lsp id> tlv 144 sub-tlv 3 detail
```

Results:

```
8005:3#show isis lsdb lspid 00e0.7bb3.07df.00-00 tlv 144 sub-tlv 3 detail
================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------
Level-1 LspID: 00e0.7bb3.07df.00-00     SeqNum: 0x00000311      Lifetime:  1032
        Chksum: 0xf022  PDU Length: 502
        Host_name: 8007
        Attributes:     IS-Type 1
TLV:144 SUB-TLV 3       ISID:
                Instance: 0
                Metric: 0
                B-MAC: 00-e0-7b-b3-07-df
                BVID:4051
                Number of ISID's:1


                        1002256(Rx)


                Instance: 0
                Metric: 0
                B-MAC: 00-e0-7b-b3-07-df
                BVID:4052
```

```
Number of ISID's:1

        1002256(Both)

Instance: 0
Metric: 0
B-MAC: 03-08-05-00-00-00
BVID:4051
Number of ISID's:1
        16000002(Rx)

Instance: 0
Metric: 0
B-MAC: 03-08-06-00-00-00
BVID:4052
Number of ISID's:1
        16000002(Rx)
```

22.5.3.5 Trace Multicast Routes

On the switch where the multicast sender is located, in our example this would be switch 8007, you can trace the multicast route by specifying the source, group, and VLAN.

Step 1 – Verify all SPB multicast routes

```
l2 tracemroute source <source address> group <group address> vlan <C-VLAN id>
```

Results: Since the multicast source is via bridges 8005 & 8006, we will use the following command to view the multicast route for group address 232.2.2.2

```
8005:3#l2 tracemroute source 10.5.41.20 group 232.2.2.2 vlan 2256
 Please wait for l2tracemroute to complete or press any key to abort
Source : 10.5.41.20
Group  : 232.2.2.2
VLAN   : 2256
BMAC   : 03:08:05:f4:24:03
B-VLAN : 4051
ISID   : 16000003
================================================================================
1   8005              00:24:43:b4:e3:df -> 8003              00:80:2d:be:23:df
2   8003              00:80:2d:be:23:df -> 8007              00:e0:7b:b3:07:df

8006:3#l2 tracemroute source 10.5.41.20 group 232.2.2.2 vlan 2256
Please wait for l2tracemroute to complete or press any key to abort
Source : 10.5.41.20
Group  : 232.2.2.2
VLAN   : 2256
BMAC   : 03:08:06:f4:24:03
B-VLAN : 4052
ISID   : 16000003
================================================================================
1   8006              00:1e:1f:48:f3:df -> 8004              00:e0:7b:bc:23:df
2   8004              00:e0:7b:bc:23:df -> 8007              00:e0:7b:b3:07:df
```

22.6 Inter VSN Routing

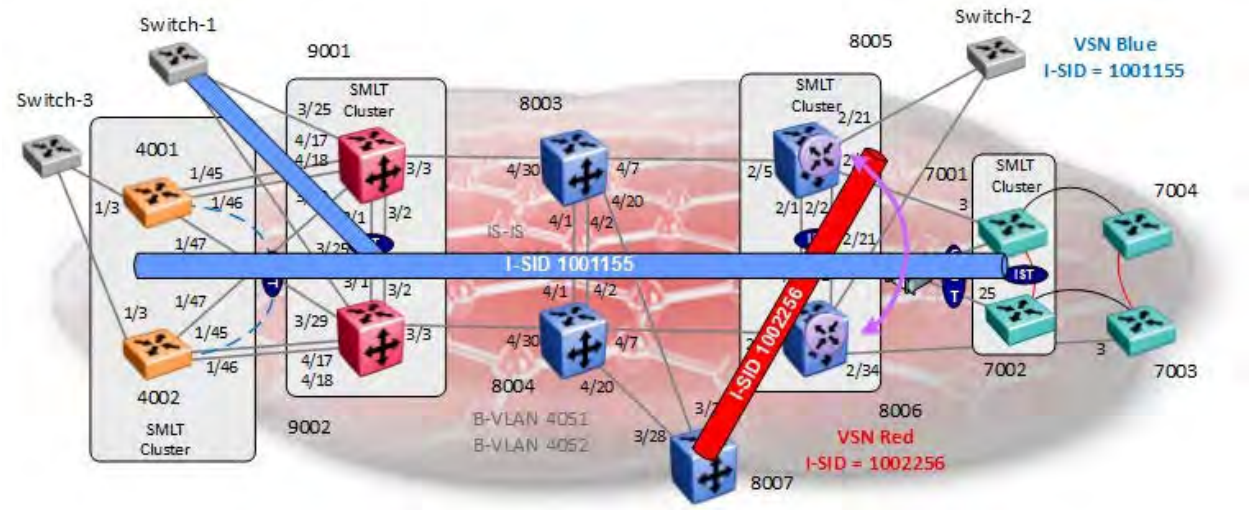

Continuing from configuration example 22.3 (L2VSN), assuming we wish to route between the Layer 2 Red and Blue Layer 2 VSNs. This can be accomplished by creating a VRF instance and adding the appropriate VLANs. For redundancy purposes, we can also create a VRF between two SPB bridges and run VRRP between for redundancy. We will enable Inter VSN routing by adding a VRF instance and then adding the Blue VSN and Red VSN on 8005 and 8006 and run VRRP between them. The end result will allow user or servers to forward traffic between Red and Blue VSNs.

In summary, we will configure the following:

- Use the base configuration from configuration example 22.1 and 22..2
- On SPB bridges 8005 and 8006, we will add the following:
    - A VRF instance named *inter-isid* with the following
        - 8005
            - Add an IP address of 10.5.40.2 to VLAN 1155 with a VRRP VIP of 10.5.40.1 and VRRP backup master enabled
            - Add an IP address of 10.5.41.2/24 to VLAN 2256 with a VRRP VIP of 10.5.41.1 and VRRP backup master enabled
        - 8006
            - Add an IP address of 10.5.40.3/24 to VLAN 1155 with a VRRP VIP of 10.5.40.1 and VRRP backup master enabled
            - Add an IP address of 10.5.41.3/24 to VLAN 2256 with a VRRP VIP of 10.5.41.1 and VRRP backup master enabled

22.7 Inter-ISID Configuration

In addition to the configuration to the configuration used in 22.3, we will add the following configuration.

22.7.1 VRF configuration

8005 & 8006 – Create VRF and add IP addressing to VLANs 1155 and 2256, enable VRRP with backup master, and make 8005 VRRP master for VLAN 1155

8005:

```
8005:5(config)#ip vrf inter-isid
8005:5(config)#interface vlan 1155
8005:5(config-if)#vrf inter-isid
8005:5(config-if)#ip address 10.5.40.2 255.255.255.0
8005:5(config-if)#ip vrrp address 10.5.40.1
8005:5(config-if)#ip vrrp 55 backup-master enable priority 150
8005:5(config-if)#ip vrrp 55 enable
8005:5(config-if)#exit
8005:5(config)#interface vlan 2256
8005:5(config-if)#vrf inter-isid
8005:5(config-if)#ip address 10.5.41.2 255.255.255.0
8005:5(config-if)#ip vrrp address 10.5.41.1
8005:5(config-if)#ip vrrp 56 backup-master enable
8005:5(config-if)#ip vrrp 56 enable
8005:5(config-if)#exit
```

For 8006, use the same configuration as above except for the items shown below

---

```
8005:5(config)#interface vlan 1155
8005:5(config-if)#ip address 10.5.40.3 255.255.255.0
8005:5(config-if)#ip vrrp 55 backup-master enable
|
8005:5(config)#interface vlan 2256
8005:5(config-if)#ip address 10.5.41.3 255.255.255.0
8005:5(config-if)#ip vrrp 56 backup-master enable priority 150
```

22.7.2 Verification

22.7.2.1 IP Route and ARP Table

Step 1 – Verify route table for VRF inter-isid

```
show ip route vrf inter-isid
```

Results:

```
8005 8006#show ip route vrf inter-isid
Response from 8005:

================================================================================================
                                IP Route - VRF inter-isid
================================================================================================
                                                 NH                INTER
 DST              MASK             NEXT           VRF          COST FACE PROT AGE TYPE PRF
------------------------------------------------------------------------------------------------
10.5.40.0        255.255.255.0    10.5.40.2       -           1    1155  LOC  0   DB   0
10.5.41.0        255.255.255.0    10.5.41.2       -           1    2256  LOC  0   DB   0

2 out of 2 Total Num of Route Entries, 2 Total Num of Dest Networks displayed.
------------------------------------------------------------------------------------------------
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed

Response from 8006:
================================================================================================
                                IP Route - VRF inter-isid
================================================================================================
                                                 NH                INTER
DST              MASK             NEXT           VRF          COST FACE PROT AGE TYPE PRF
------------------------------------------------------------------------------------------------
10.5.40.0        255.255.255.0    10.5.40.3       -           1    1155  LOC  0   DB   0
10.5.41.0        255.255.255.0    10.5.41.3       -           1    2256  LOC  0   DB   0

2 out of 2 Total Num of Route Entries, 2 Total Num of Dest Networks displayed.
------------------------------------------------------------------------------------------------
TYPE Legend:
```

```
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed
```

Step 2 – Verify VRRP operations

show ip vrrp vrf inter-isid

Results:

```
8005 8006# show ip vrrp vrf inter-isid
Response from 8005:
================================================================================
                          VRRP Info - VRF inter-isid
================================================================================

VRID  P/V   IP              MAC               STATE    CONTROL  PRIO  ADV
 ------------------------------------------------------------------------------
55    1155  10.5.40.1       00:00:5e:00:01:37  Master   Enabled  150   1
56    2256  10.5.41.1       00:00:5e:00:01:38  Back Up  Enabled  100   1

VRID  P/V   MASTER          UP TIME               HLD DWN  CRITICAL IP (ENABLED)
 ------------------------------------------------------------------------------
55    1155  10.5.40.2       0 day(s), 00:07:56    0        0.0.0.0         (No)
56    2256  10.5.41.3       0 day(s), 00:06:34    0        0.0.0.0         (No)

VRID  P/V   BACKUP MASTER   BACKUP MASTER STATE   FAST ADV (ENABLED)
--------------------------------------------------------------------------------------------------------
55    1155  enable          down                   200       (NO)
56    2256  enable          up                     200       (NO)


Response from 8006:
================================================================================
                          VRRP Info - VRF inter-isid
================================================================================

VRID  P/V   IP              MAC               STATE    CONTROL  PRIO  ADV
--------------------------------------------------------------------------------------------------------
55    1155  10.5.40.1       00:00:5e:00:01:37  Back Up  Enabled  100   1
```

```
56    2256  10.5.41.1       00:00:5e:00:01:38  Master    Enabled  150   1

VRID  P/V   MASTER          UP TIME               HLD DWN  CRITICAL IP (ENABLED)
-------------------------------------------------------------------------------------------------------
55    1155  10.5.40.2       0 day(s), 00:07:57    0        0.0.0.0          (No)
56    2256  10.5.41.3       0 day(s), 00:06:35    0        0.0.0.0          (No)

VRID  P/V   BACKUP MASTER   BACKUP MASTER STATE   FAST ADV (ENABLED)
-------------------------------------------------------------------------------------------------------
55    1155  enable          up                    200       (NO)
56    2256  enable          down                  200       (NO)
```

Step 3 – Verify ARP table

```
show ip arp vrf inter-isid
```

Results:

```
8005 8006#show ip arp vrf inter-isid
Response from 8005:
================================================================================
                            IP Arp - VRF inter-isid
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN    PORT       TYPE    TTL(10 Sec)
-------------------------------------------------------------------------------------------------------
10.5.40.2        00:24:43:b4:e2:25  1155    -        LOCAL   2160
10.5.40.255      ff:ff:ff:ff:ff:ff  1155    -        LOCAL   2160
10.5.41.2        00:24:43:b4:e2:2d  2256    -        LOCAL   2160
10.5.41.255      ff:ff:ff:ff:ff:ff  2256    -        LOCAL   2160
10.5.40.1        00:00:5e:00:01:37  1155    -        LOCAL   2160
10.5.41.1        00:00:5e:00:01:38  2256    -        LOCAL   2160
10.5.41.10       00:0c:29:26:b5:af  2256  ISID-1002256  DYNAMIC 2118
10.5.41.3        00:1e:1f:48:f2:2d  2256  ISID-1002256  DYNAMIC 2118
10.5.40.3        00:1e:1f:48:f2:24  1155  ISID-1001155  DYNAMIC 2118
10.5.41.20       00:0c:29:d9:96:59  2256    2/21      DYNAMIC 2121
10.5.40.52       00:0c:29:9b:a8:31  1155  ISID-1001155  DYNAMIC 2154
10.5.40.51       00:0c:29:35:62:a4  1155  ISID-1001155  DYNAMIC 2160
10.5.40.5        00:0c:29:d6:81:e5  1155  ISID-1001155  DYNAMIC 2159
Response from 8006:
================================================================================
                            IP Arp - VRF inter-isid
```

```
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN    PORT       TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
10.5.40.3        00:1e:1f:48:f2:24  1155    -        LOCAL   2160
10.5.40.255      ff:ff:ff:ff:ff:ff  1155    -        LOCAL   2160
10.5.41.3        00:1e:1f:48:f2:2d  2256    -        LOCAL   2160
10.5.41.255      ff:ff:ff:ff:ff:ff  2256    -        LOCAL   2160
10.5.40.1        00:00:5e:00:01:37  1155    -        LOCAL   2160
10.5.41.1        00:00:5e:00:01:38  2256    -        LOCAL   2160
10.5.41.10       00:0c:29:26:b5:af  2256  ISID-1002256  DYNAMIC 2160
10.5.41.20       00:0c:29:d9:96:59  2256    2/21      DYNAMIC 2160
10.5.40.52       00:0c:29:9b:a8:31  1155  ISID-1001155  DYNAMIC 2124
10.5.40.51       00:0c:29:35:62:a4  1155  ISID-1001155  DYNAMIC 2143
10.5.40.5        00:0c:29:d6:81:e5  1155  ISID-1001155  DYNAMIC 2159
```

22.7.2.2 MAC Address Table

Step 1 – Verify MAC table for VRF inter-isid

```
show vlan mac-address-entry <vlan id>
```

Results:

```
8005#show vlan mac-address-entry 1155
================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                          QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE        MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
1155 self        00:00:5e:00:01:37  Port-cpp          false   1      false
1155 learned     00:0c:29:35:62:a4  ISID-1001155    false   1      false
1155 learned     00:0c:29:9b:a8:31  ISID-1001155    false   1      false
1155 learned     00:0c:29:d6:81:e5  ISID-1001155    false   1      false
1155 learned     00:18:71:ea:31:bb  ISID-1001155    false   1      false
1155 learned     00:1e:1f:48:f2:24  ISID-1001155    false   1      true
1155 self        00:24:43:b4:e2:25  Port-cpp          false   1      false
8005#show vlan mac-address-entry 2256
================================================================================
                                  Vlan Fdb
================================================================================
VLAN            MAC                                          QOS    SMLT
ID   STATUS     ADDRESS            INTERFACE        MONITOR LEVEL  REMOTE
--------------------------------------------------------------------------------
```

```
2256 self        00:00:5e:00:01:38  Port-cpp          false   1      false
2256 learned     00:0c:29:26:b5:af  ISID-1002256    false   1      false
2256 learned     00:0c:29:d9:96:59  Port-2/21         false   1      false
2256 learned     00:1e:1f:48:f2:2d  ISID-1002256    false   1      true
2256 self        00:24:43:b4:e2:2d  Port-cpp          false   1      false
```

22.8 SPB L3 VSN – SMLT

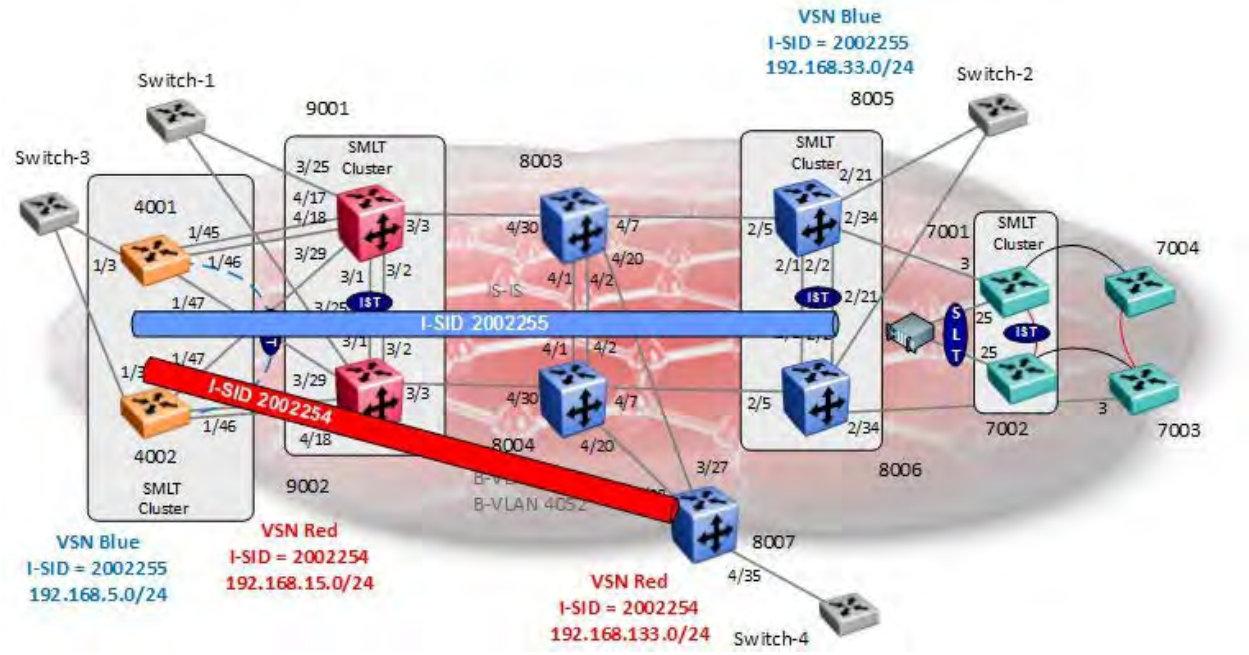

For this example, we will configure the SMLT switch cluster with the following:

- SPB IP
  - SPB IP parameter must be enabled BEB switches SMLT cluster switches 4001 & 4002, 8005 & 8006, and standalone switch 8007
  - An IS-IS source IP address must be configured (loopback/curcuitless IP address)
- VRF's - BEB Nodes only
  - VRF Red
    - VLAN ID = 2254 configured on switches 4001, 4002 and 8007
    - Assign ISID 2002254 to VRF Red
    - On the vIST cluster 4001 and 4002, we will also need to assign an ISID to the VLAN where we will use ISID 3002254
  - VRF Blue
    - VLAN ID = 2255 configured on switches 4001, 4002, 8005, and 8006
    - Assign ISID 2002255 to VRF Blue
    - On the vIST cluster 4001 and 4002, we will also need to assign an ISID to the VLAN where we will use ISID 3002255

This example is a continuation from the base setup used in Section 22.1 and 22.2 for the SMLT configuration.

22.8.1 SPB IP Enable

22.8.1.1 IS-IS Layer 3 configuration

4001

```
4001:1(config)#interface loopback 1
4001:1(config-if)#ip address 1 10.4.4.1/255.255.255.255
4001:1(config-if)#exit
4001:1(config)#router isis
4001:1(config-isis)#ip-source-address 10.4.4.1
4001:1(config-isis)#spbm 1 ip enable
4001:1(config-isis)#exit
```

4002

```
4002:1(config)#interface loopback 1
4002:1(config-if)#ip address 1 10.4.4.2/255.255.255.255
4002:1(config-if)#exit
4002:1(config)#router isis
4002:1(config-isis)#ip-source-address 10.4.4.2
4002:1(config-isis)#spbm 1 ip enable
4002:1(config-isis)#exit
```

8005

```
8005:5(config)#interface loopback 1
8005:5(config-if)#ip address 1 10.1.1.5/255.255.255.255
8005:5(config-if)#exit
8005:5(config)#router isis
8005:5(config-isis)#ip-source-address 10.1.1.5
8005:5(config-isis)#spbm 1 ip enable
8005:5(config-isis)#exit
```

8006

```
8006:5(config)#interface loopback 1
8006:5(config-if)#ip address 1 10.1.1.6/255.255.255.255
8006:5(config-if)#exit
8006:5(config)#router isis
8006:5(config-isis)#ip-source-address 10.1.1.6
8006:5(config-isis)#spbm 1 ip enable
8006:5(config-isis)#exit
```

8007

```
8007:5(config)#interface loopback 1
8007:5(config-if)#ip address 1 10.1.1.7/255.255.255.255
8007:5(config-if)#exit
8007:5(config)#router isis
8007:5(config-isis)#ip-source-address 10.1.1.7
8007:5(config-isis)#spbm 1 ip enable
8007:5(config-isis)#exit
```

22.8.1.2 VRF Configuration

4001 and 4002

4001 & 4002: Same configuration on both switches

```
4001:1(config)#ip vrf blue
4001:1(config)#ip vrf red
```

8005 & 8006

8005 & 8006: Same configuration on both switches

```
8005:5(config)#ip vrf blue
```

8007

```
8007:5(config)#ip vrf red
```

22.8.2 VLAN Configuration

4001 and 4001 – Add VLAN and ISID identifier for the vIST

4001 & 4002: Same configuration on both switches assuming we are using the SMLT configuration via MLT 2

```
4001:1(config)#vlan create 2254 name vsnred-2254 type port-mstprstp 0
4001:1(config)#vlan ISID 2254 3002254
4001:1(config)#vlan create 2255 name vsnblue-2255 type port-mstprstp 0
4001:1(config)#vlan ISID 2255 3002255
4001:1(config)#vlan mlt 2254 2
4001:1(config)#vlan mlt 2255 2
```

8005 and 8006

8005 & 8006: Same configuration on both switches assuming we are using SLT 129 via port 2/21 and MLT 1 for the IST

```
8005:5(config)#vlan create 2255 name "vsnblue-2255" type port-mstprstp 0
8005:5(config)# vlan members add 2255 2/21
```

8007

```
8007:5(config)#vlan create 2254 name "vsnred-2254" type port-mstprstp 0
8007:5(config)#vlan ports 4/35 tagging tagAll
8007:5(config)# vlan members add 2254 4/35
8007:5(config)# vlan members remove 1 4/35
```

22.8.3 IPVPN Configuration

4001 and 4002 - Add IP address and VRF to VLANs 2254 & 2255

4001:

```
4001:1(config)#interface vlan 2254
4001:1(config-if)#vrf red
4001:1(config-if)#ip address 192.168.15.1 255.255.255.0
4001:1(config-if)#ip rsmlt
4001:1(config-if)#ip rsmlt holdup-timer 9999
4001:1(config-if)#exit
4001:1(config)#interface vlan 2255
4001:1(config-if)#vrf blue
4001:1(config-if)#ip address 192.168.5.1 255.255.255.0
4001:1(config-if)#ip rsmlt
4001:1(config-if)#ip rsmlt holdup-timer 9999
4001:1(config-if)#exit
4001:1(config)#ip rsmlt edge-support
```

For 4002, use the same configuration as above except for the items shown below

---

```
4002:1(config)#interface vlan 2254
4002:1(config-if)#ip address 192.168.15.2 255.255.255.0
4002:1(config)#interface vlan 2255
4002:1(config-if)#ip address 192.168.5.2 255.255.255.0
```

8005 and 8006 - Add IP address and VRF to VLANs 2254 and 2555, enable RSMLT Edge by setting the holdup timer to infinity (9999), and enable RSMT edge support globally

8005:

```
8005:5(config)#interface vlan 2255
8005:5(config-if)#vrf blue
8005:5(config-if)#ip address 192.168.33.1 255.255.255.0
8005:5(config-if)#ip rsmlt
8005:5(config-if)#ip rsmlt holdup-timer 9999
8005:5(config-if)#exit
8005:5(config)#ip rsmlt edge-support
```

For 8006, use the same configuration as above except for the items shown below

---

```
8006:5(config)#interface vlan 2255
8006:5(config-if)#ip address 192.168.33.2 255.255.255.0
```

8007 - Add IP address and VRF to VLAN 2254

```
8007:5(config)#interface vlan 2254
8007:5(config-if)#vrf red
8007:5(config-if)#ip address 192.168.133.1 255.255.255.0
8007:5(config-if)#exit
```

22.8.4 Enable L3VSN Configuration

4001 and 4002 - Enable L3 IPVPN and ISID to VRF red and blue

4001 & 4002: Same configuration on both switches

```
4001:1(config)#router vrf red
4001:1(router-vrf)#ipvpn
4001:1(router-vrf)#ISID 2002254
4001:1(router-vrf)#ipvpn enable
4001:1(router-vrf)#exit
4001:1(config)#router vrf blue
4001:1(router-vrf)#ipvpn
4001:1(router-vrf)#ISID 2002255
4001:1(router-vrf)#ipvpn enable
4001:1(router-vrf)#exit
```

8005 and 8006 - Enable L3 IPVPN and ISID to VRF blue

8005 & 8006: Same configuration on both switches

```
8005:5(config)#router vrf blue
8005:5(router-vrf)#ipvpn
8005:5(router-vrf)#ISID 2002255
8005:5(router-vrf)#ipvpn enable
8005:5(router-vrf)#exit
```

8007 - Enable L3 IPVPN and ISID to VRF red

```
8007:5(config)#router vrf red
8007:5(router-vrf)#ipvpn
8007:5(router-vrf)#ISID 2002254
8007:5(router-vrf)#ipvpn enable
8007:5(router-vrf)#exit
```

22.8.5 Enable direct interface redistribution

4001 and 4002 - Redistribute IP Networks via IS-IS – Direct Interfaces

4001 & 4002: Same configuration on both switches

```
4001:1(config)#router vrf red
4001:1(router-vrf)#isis redistribute direct
WARNING: Routes will not be injected until apply command is issued after enable command
4001:1(router-vrf)#isis redistribute direct enable
4001:1(router-vrf)#exit
4001:1(config)#router vrf blue
4001:1(router-vrf)#isis redistribute direct
WARNING: Routes will not be injected until apply command is issued after enable command
4001:1(router-vrf)#isis redistribute direct enable
4001:1(router-vrf)#exit
4001:1(config)#isis apply redistribute direct vrf red
4001:1(config)#isis apply redistribute direct vrf blue
```

8005 and 8006 - Redistribute IP Networks via IS-IS – Direct Interfaces

8005 & 8006: Same configuration on both

```
8005:5(config)#router vrf blue
8005:5(router-vrf)#isis redistribute direct
WARNING: Routes will not be injected until apply command is issued after enable command
8005:5(router-vrf)#isis redistribute direct enable
8005:5(router-vrf)#exit
8005:5(config)#isis apply redistribute direct vrf red
8005:5(config)#isis apply redistribute direct vrf blue
```

8007 - Redistribute IP Networks via IS-IS – Direct Interfaces

```
8007:5(config)#router vrf red
8007:5(router-vrf)#isis redistribute direct
WARNING: Routes will not be injected until apply command is issued after enable command
8007:5(router-vrf)#isis redistribute direct enable
8007:5(router-vrf)#exit
8007:5(config)#isis apply redistribute direct vrf red
```

22.8.6 Verify Operations

22.8.6.1 Verify RSMLT Information

4001 & 4002 and 8005 & 8006 - Verify RSMLT is up and operational for both VRF instances

```
show ip rsmlt vrf red
show ip rsmlt vrf blue
```

Results:

```
8005 8006#show ip rsmlt vrf blue
Response from 8005:
================================================================================
                        Ip Rsmlt Local Info - VRF blue
================================================================================
VID   IP              MAC               ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------
2255  192.168.33.1    00:24:43:b4:e2:23 Enable  Up    60     infinity
VID   SMLT ID                              SLT ID
--------------------------------------------------------------------------------
2255     5                                  129
VID   IPv6            MAC               ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------
VID   SMLT ID                              SLT ID
--------------------------------------------------------------------------------
================================================================================
                        Ip Rsmlt Peer Info - VRF blue
================================================================================
VID   IP              MAC               ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------
2255  192.168.33.2    00:1e:1f:48:f2:22 Enable  Up    60     infinity
VID   HDT REMAIN  HUT REMAIN  SMLT ID                              SLT ID

--------------------------------------------------------------------------------
2255  60          infinity      5                                  129
VID   IPv6            MAC               ADMIN   OPER  HDTMR  HUTMR
--------------------------------------------------------------------------------
VID   HDT REMAIN  HUT REMAIN  SMLT ID                              SLT ID
--------------------------------------------------------------------------------

Response from 8006:
================================================================================
                        Ip Rsmlt Local Info - VRF blue
```

```
================================================================================
VID    IP              MAC               ADMIN    OPER   HDTMR  HUTMR
--------------------------------------------------------------------------------
2255   192.168.33.2    00:1e:1f:48:f2:22 Enable   Up     60     infinity
VID    SMLT ID                              SLT ID
--------------------------------------------------------------------------------
2255     5                                    129
VID    IPv6            MAC               ADMIN    OPER   HDTMR  HUTMR
--------------------------------------------------------------------------------
VID    SMLT ID                              SLT ID
--------------------------------------------------------------------------------
================================================================================
                        Ip Rsmlt Peer Info - VRF blue
================================================================================
VID    IP              MAC               ADMIN    OPER   HDTMR  HUTMR
--------------------------------------------------------------------------------
2255   192.168.33.1    00:24:43:b4:e2:23 Enable   Up     60     infinity
VID    HDT REMAIN  HUT REMAIN  SMLT ID                                SLT ID
--------------------------------------------------------------------------------
2255   60          infinity     5                                     129
VID    IPv6            MAC               ADMIN    OPER   HDTMR  HUTMR
--------------------------------------------------------------------------------
VID    HDT REMAIN  HUT REMAIN  SMLT ID                                SLT ID
--------------------------------------------------------------------------------
```

On each SMLT cluster switch, verify the following:

Option	Verify
VID	The value displayed should be ***2254*** for vrf red and ***2255*** for vrf blue as per the VLAN ID used in this configuration example.
ADMIN	The value displayed should be ***Enable*** which indicates that RSMLT has been enabled for this interface
OPER	Should be displayed as ***Up*** indicating that RSMLT is operational
HUTMR	Should be displayed as ***infinity*** as we configured this interface as a RSMLT Edge interface with a holdup-timer timer value of 9999
SLT ID	The value displayed should be ***129*** for SMLT cluster 8005 & 8005 as per the SMLT ID's used in this example

22.8.6.2 Verify IS-IS ISID

Show IS-IS ISID pertaining to the vIST

```
show isis spbm ISID all id <3002254|3002255>
```

Results:

4001 & 4002:

```
4001:1#show isis spbm ISID all id 3002254

================================================================================
                                SPBM ISID INFO
================================================================================
ISID      SOURCE NAME    VLAN    SYSID                TYPE         HOST_NAME
 ------------------------------------------------------------------------------
3002254 0.40.01         4051    d4ea.0e10.e465       config       4001
3002254 0.40.02         4052    a012.90d3.ec65       discover     4002

4001:1#show isis spbm ISID all id 3002255

================================================================================
                                SPBM ISID INFO
================================================================================
ISID      SOURCE NAME    VLAN    SYSID                TYPE         HOST_NAME
 ------------------------------------------------------------------------------
3002255 0.40.01         4051    d4ea.0e10.e465       config       4001
3002255 0.40.02         4052    a012.90d3.ec65       discover     4002
```

Show IS-IS ISID pertaining to each vrf instance

```
show ip ipvpn
```

Results:

4001 & 4002:

```
4001:1#show ip ipvpn

          VRF Name             : blue
          Ipvpn-state          : enabled
          ISID                : 2002255

          VRF Name             : red
          Ipvpn-state          : enabled
          ISID                : 2002254
```

8005 & 8006:

```
8005:5#show ip vrf ipvpn
          VRF Name              : blue
          Ipvpn-state           : enabled
          ISID                 : 2002255
```

8007:

```
8007:5#show ip vrf ipvpn
          VRF Name              : red
          Ipvpn-state           : enabled
          ISID                 : 2002254
```

On each switch, verify the following:

Option	Verify
VRF Name Ipvpn-state ISID	For the VRF Name of ***blue***, the Ipvpn-state should display ***enabled*** with an ISID value of ***2002255***. For the VRF Name of ***red***, the Ipvpn-state should display ***enabled*** with an ISID value of ***2002254***.

22.8.6.3 Show IS-IS SPB IP Unicast Forwarding database

Show IS-IS SPB IP Unicast FIB using ISID 2002254 and 2002255 as used in this configuration example

```
show isis spbm ip-unicast-fib id <ISID id>
```

Results: Example from 4002

4002:

```
4002:1(config)#show isis spbm ip-unicast-fib id 2002255
================================================================================================
                                    SPBM IP-UNICAST FIB ENTRY INFO
================================================================================================
          VRF     DEST                                      OUTGOING  SPBM    PREFIX  IP ROUTE
VRF       ISID    ISID    Destination       NH BEB     VLAN INTERFACE COST    COST    PREFERENCE
------------------------------------------------------------------------------------------------
blue      2002255 2002255 192.168.5.0/24    4001       4051 1/47      20       1       7
blue      2002255 2002255 192.168.5.0/24    4001       4052 9001      20       1       7
blue      2002255 2002255 192.168.33.0/24   8005       4051 1/47      30       1      7
blue      2002255 2002255 192.168.33.0/24   8005       4052 1/47      30       1       7
blue      2002255 2002255 192.168.33.0/24   8006       4051 9001      30       1       7
blue      2002255 2002255 192.168.33.0/24   8006       4052 9001      30       1      7

4002:1(config)#show isis spbm ip-unicast-fib id 2002254
================================================================================================
                                    SPBM IP-UNICAST FIB ENTRY INFO
================================================================================================
          VRF     DEST                                      OUTGOING  SPBM    PREFIX  IP ROUTE
VRF       ISID    ISID    Destination       NH BEB     VLAN INTERFACE COST    COST    PREFERENCE
------------------------------------------------------------------------------------------------
red       2002255 2002255 192.168.15.0/24    4001       4051 1/47      20       1       7
red       2002255 2002255 192.168.15.0/24    4001       4052 9001      20       1       7
red       2002254 2002255 192.168.133.0/24   8007       4051 1/47      30       1       7
red       2002254 2002255 192.168.133.0/24   8007       4052 9001      30       1       7
```

22.8.6.4 Show IS-IS LSP Details

In a IS-IS network, each IS router advertises one or more IS-IS Link State Protocol Data Units (LSPs) with routing information. Within each LSP, there is a fixed header and a number of TLVs with encoded information. The following command is used to show details of a LSP in detail to a specific neighbor displaying the encoded information in the TLVs.

Show IS-IS LSP details

```
show isis lsdb tlv 184 detail
show isis lsdb lspid <is-is system id>.00-00 tlv 184 detail
```

Results: Example from 4002

4002:

```
4002:1#show isis lsdb tlv 184 detail

================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------
Level-1 LspID: d4ea.0e10.e465.00-03     SeqNum: 0x00000009      Lifetime: 1027
        Chksum: 0xbdc0 PDU Length: 89
        Host_name: 4001
        Attributes:     IS-Type 1
TLV:184 SPBM IPVPN Reachability:
                Vrf ISID:2002255
                        Metric:1        Prefix Length:24
                        IP Address: 192.168.5.0

                Vrf ISID:2002254
                        Metric:1        Prefix Length:24
                        IP Address: 192.168.15.0


Level-1 LspID: a012.90d3.ec65.00-03     SeqNum: 0x0000000a      Lifetime: 1181
        Chksum: 0xe990 PDU Length: 89
        Host_name: 4002
        Attributes:     IS-Type 1
TLV:184 SPBM IPVPN Reachability:
                Vrf ISID:2002255
                        Metric:1        Prefix Length:24
                        IP Address: 192.168.5.0
```

```
                Vrf ISID:2002254
                        Metric:1        Prefix Length:24
                        IP Address: 192.168.15.0


Level-1 LspID: 0024.43b4.e3df.00-00     SeqNum: 0x000023c1      Lifetime: 1182
        Chksum: 0x56dd PDU Length: 772
        Host_name: 8005
        Attributes:     IS-Type 1
TLV:184 SPBM IPVPN Reachability:
                Vrf ISID:2002255
                        Metric:1        Prefix Length:24
                        IP Address: 192.168.33.0


Level-1 LspID: 001e.1f48.f3df.00-00     SeqNum: 0x000023b6      Lifetime: 1183
        Chksum: 0xee6   PDU Length: 797
        Host_name: 8006
        Attributes:     IS-Type 1
TLV:184 SPBM IPVPN Reachability:
                Vrf ISID:2002255
                        Metric:1        Prefix Length:24
                        IP Address: 192.168.33.0


Level-1 LspID: 00e0.7bb3.07df.00-00     SeqNum: 0x000000d4      Lifetime:   921
        Chksum: 0x8c2d PDU Length: 395
        Host_name: 8007
        Attributes:     IS-Type 1
TLV:184 SPBM IPVPN Reachability:
                Vrf ISID:2002254
                        Metric:1        Prefix Length:24
                        IP Address: 192.168.133.0
```

In reference to 4002, verify the following:

Option	Verify
Level 1 TLV:184	As this example, in reference to 4002, for the blue vrf ISID 2002255, we should learns routes***192.168.33.0/24*** from ***4001*** and ***192.168.33.0/24*** from ***8005*** and ***8006*** and for the red vrf ISID ***2002254***, we should learn routes ***192.168.15.0/24*** from ***4001*** and ***192.168.133.0/24*** from ***8007***.

22.8.6.5 IP Route Table

Use the following command to display the routes for each VRF instance

Display IP route table for each VRF instance

```
show ip route vrf blue
show ip route vrf red
```

Results: Example from 4002

4002:

```
4002:1#show ip route vrf red
================================================================================================
                                        IP Route - VRF red
================================================================================================
                                                    NH                    INTER
DST             MASK            NEXT                VRF/ISID         COST FACE PROT AGE TYPE
PRF
------------------------------------------------------------------------------------------------
192.168.15.0    255.255.255.0   192.168.15.2        -                1    2254  LOC  0   DB   0
192.168.133.0   255.255.255.0   8007                red              30   4051  ISIS 0   IBSV 7

2 out of 2 Total Num of Route Entries, 2 Total Num of Dest Networks displayed.
------------------------------------------------------------------------------------------------
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed

4002:1#show ip route vrf blue
================================================================================================
                                        IP Route - VRF blue
================================================================================================
                                                    NH                    INTER
DST             MASK            NEXT                VRF/ISID         COST FACE PROT AGE TYPE
PRF
------------------------------------------------------------------------------------------------
192.168.5.0     255.255.255.0   192.168.5.2         -                1    2255  LOC  0   DB   0
192.168.33.0    255.255.255.0   8005                blue             30   4051  ISIS 0   IBSV 7

2 out of 2 Total Num of Route Entries, 2 Total Num of Dest Networks displayed.
```

In reference to each switch, verify the following information:

Option	Verify
Next PROT TYPE	All local interfaces should display ***LOC*** whereas all learned routes should display ***ISIS*** with the appropriate next-hop address and type of ***IBSV***. The next hop ISIS System-name should be displayed for remote networks

22.8.6.6 Verify VRF L3 operations

Step 1 - Use ping command to verify network connectivity to neighbors

```
ping <host> vrf <value> source <source ip>
```

Results: Example from 4002

4002:

```
4002:1#ping 192.168.133.1 vrf red source 192.168.15.2
192.168.133.1 is alive
4002:1#ping 192.168.33.1 vrf blue source 192.168.5.2
192.168.33.1 is alive
4002:1#ping 192.168.33.2 vrf blue source 192.168.5.2
192.168.33.2 is alive
```

Step 2 - Use traceroute command to verify network connectivity to neighbors

```
traceroute <host> vrf <value> source <source ip>
```

Results: Example from 4002

4002:

```
4002:1#traceroute 192.168.133.1 vrf red source 192.168.15.2
traceroute to 192.168.133.1, 30 hops max, 56 byte packets (vrf red)
 1  192.168.133.1  1.997 ms 2.773 ms 1.808 ms
4002:1#traceroute 192.168.33.1 vrf blue source 192.168.5.2
traceroute to 192.168.33.1, 30 hops max, 56 byte packets (vrf blue)
 1  192.168.33.1  1.847 ms 2.424 ms 1.706 ms
```

Step 3 - Use l2 traceroute command to verify network connectivity to neighbors

```
L2 traceroute ip-address <host> vrf <value>
```

Results: Example from 4002

4002:

```
4002:1#l2 traceroute ip-address 192.168.133.1 vrf red

 Please wait for l2trace to complete or press any key to abort

L2 Trace  Statistics : IP 192.168.133.1, paths found 1
================================================================
8007  (00:e0:7b:b3:07:df),  vlan 4051

0    4002                (a0:12:90:d3:ec:65)
1    9001                (d4:ea:0e:fd:e3:df)
2    8003                (00:80:2d:be:23:df)
3    8007                (00:e0:7b:b3:07:df)

4002:1#l2 traceroute ip-address 192.168.33.1 vrf blue

 Please wait for l2trace to complete or press any key to abort

L2 Trace  Statistics : IP 192.168.33.1, paths found 1
================================================================
8005  (00:24:43:b4:e3:df ),  vlan 4051

0    4002               (a0:12:90:d3:ec:65)
1    9001               (d4:ea:0e:fd:e3:df)
2    8003               (00:80:2d:be:23:df)
3    8005               (00:24:43:b4:e3:df )
```

Step 3 - Verify ARP and local MAC entry for local hosts

```
show ip arp vrf <vrf name>
show vlan mac-address-entry <vlan id>
```

Results: Example from 4002 for vrf blue

4002:

```
4002:1#show ip arp vrf blue
================================================================================================
                                    IP Arp - VRF blue
================================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN  PORT                TYPE    TTL(10 Sec) TUNNEL
------------------------------------------------------------------------------------------------
192.168.5.1     d4:ea:0e:10:e4:8b  2255  9001                DYNAMIC 1711  4001
192.168.5.2     d4:ea:0e:15:30:88  2255  -                   LOCAL   2160
192.168.5.255   ff:ff:ff:ff:ff:ff  2255  -                   LOCAL   2160

4002:1#show vlan mac-address-entry 2255
================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                            SMLT
ID   STATUS     ADDRESS            INTERFACE   REMOTE TUNNEL
--------------------------------------------------------------------------------
2255 learned    d4:ea:0e:10:e4:8b  9001        true   4001
2255 self       d4:ea:0e:15:30:88  Port-cpp    false  -
```

22.9 Extending L3VSN to the VSP 7000 Cluster via L2VSN

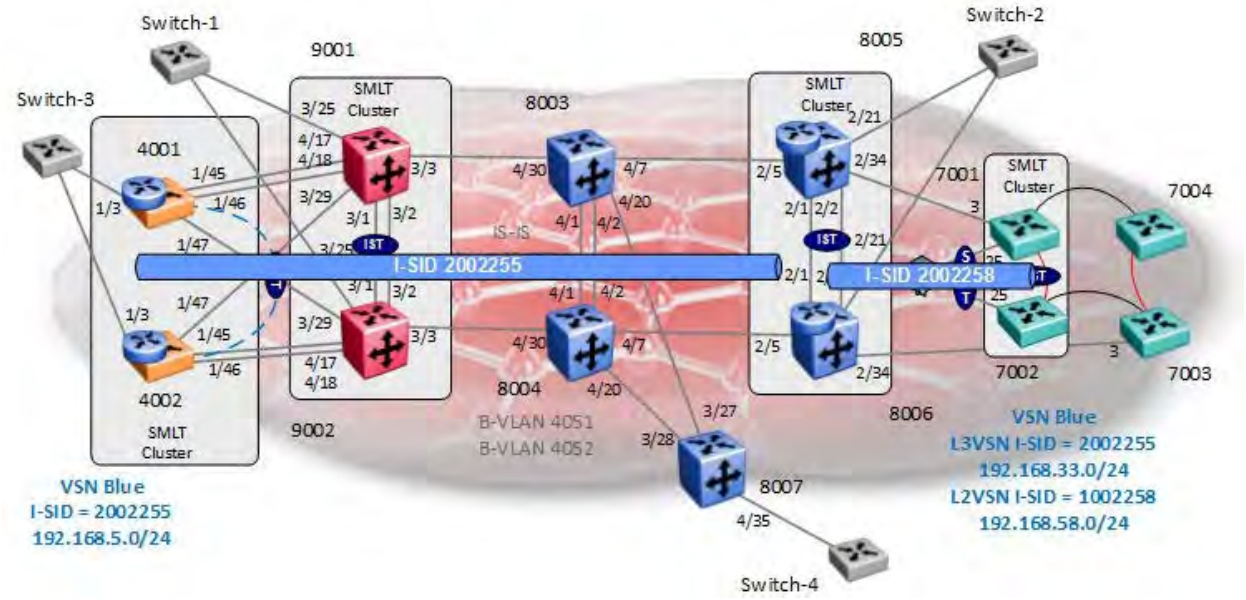

Continuing from the L3VSN example in Section 22.8, we will extend the blue vrf to the VSP 7000 7001 & 7002 SMLT cluster by adding a L2VSN between SPB bridges 8005 & 8006 and 7001 & 7002 and then adding the L2VSN VLAN provisioned on SPB bridges 8005 and 8006 to the blue vrf. For redundancy, we will also enable VRRP with Backup Master on 8005 & 8006.

In summary, we will configure the following:

L2VSN

- Assign ISID 1002558 to local VLAN 2558 on SPB bridges 8005 & 8006, and 7001 & 7002
  - On bridges 8005 and 8006
    - Add VLAN 2258 to the blue vrf configured in Section 22.8
    - For VLAN 2258, add IP subnet 192.168.58.0/24 with a VRRP virtual IP address of 192.168.58.1 and VRRP Backup Master enabled

22.9.1 L2VSN Configuration

8005 and 8006

8005 & 8006: Same configuration on both switches

```
8005:5(config)#vlan create 2258 type port-mstprstp 0
8005:5(config)#vlan ISID 2258 1002258
```

7001 and 7002 – Assuming we are using local ports 11. Please see section 17.2 for the SMLT configuration

7001 & 7002: Same configuration on both switches

```
7001(config)#vlan create 2258 type port
7001(config)#vlan configcontrol automatic
7001(config)#vlan members add 2258 25
7001(config)#vlan ISID 2258 1002258
```

22.9.2 VRF Configuration

8005 and 8006

8005:

```
8005:5(config)#interface vlan 2258
8005:5(config-if)#vrf blue
8005:5(config-if)#ip address 192.168.58.2 255.255.255.0
8005:5(config-if)#ip vrrp address 58 192.168.58.1
8005:5(config-if)#ip vrrp 58 backup-master enable
8005:5(config-if)#ip vrrp 58 enable
8005:5(config-if)#exit
```

8006 will have the same configuration except for the items shown below assuming also that we wish to make 8006 the VRRP master

---

```
8006:5(config-if)#vrf blue
8006:5(config-if)#ip address 192.168.58.3 255.255.255.0
8006:5(config-if)#ip vrrp 58 backup-master enable priority 150
```

22.9.3 Verify Operations

22.9.3.1 Verify L2VSN Operations

Verify L2VSN is configured and discovered

```
show isis spbm ISID all id 1002258
```

Results: Example from switch 7001

```
7001#show isis spbm ISID all id 1002258
================================================================================
                              SPBM ISID INFO
================================================================================
ISID      SOURCE NAME   VLAN   SYSID            TYPE       HOST-NAME
================================================================================
1002258   0.70.01       4051   fca8.41f6.37df   config     7001
1002258   0.70.02       4052   3cb1.5bff.5fdf   discover   7002
1002258   0.80.05       4051   0024.43b4.e3df   discover   8005
1002258   0.80.06       4052   001e.1f48.f3df   discover   8006
```

22.9.3.2 Verify VRRP Operations

8005 & 8006 - Verify RSMLT is up and operational for both VRF instances

```
show ip vrrp address vrid <1-255> vrf <name>
```

Results:

```
8005 8006#show ip vrrp address vrid 58 vrf blue
Response from 8005:
================================================================================
                              VRRP Info - VRF blue
================================================================================

VRID  P/V   IP              MAC               STATE    CONTROL PRIO  ADV
--------------------------------------------------------------------------------------------------------------
58    2558  192.168.58.1    00:00:5e:00:01:3a  Back Up  Enabled  100   1

VRID  P/V   MASTER          UP TIME               HLD DWN CRITICAL IP (ENABLED)
--------------------------------------------------------------------------------------------------------------
58    2558  192.168.58.3    0 day(s), 01:43:08    0        0.0.0.0         (No)

VRID  P/V   BACKUP MASTER   BACKUP MASTER STATE   FAST ADV (ENABLED)
```

```
-------------------------------------------------------------------------------------------------------
58    2558  enable          up                    200       (NO)


Response from 8006:
===============================================================================
                              VRRP Info - VRF blue
===============================================================================

VRID  P/V   IP              MAC                STATE    CONTROL PRIO  ADV
-------------------------------------------------------------------------------------------------------
58    2558  192.168.58.1    00:00:5e:00:01:3a  Master   Enabled  150   1

VRID  P/V   MASTER          UP TIME               HLD DWN CRITICAL IP (ENABLED)
-------------------------------------------------------------------------------------------------------
58    2558  192.168.58.3    0 day(s), 01:43:23    0        0.0.0.0         (No)

VRID  P/V   BACKUP MASTER   BACKUP MASTER STATE   FAST ADV (ENABLED)
-------------------------------------------------------------------------------------------------------
58    2558  enable          down                  200       (NO)
```

On each 8005 and 8006, verify the following:

Option	Verify
IP	Under IP, the VRRP address of ***192.168.58.1*** should be shown.
Master	As bridge 8005 has been configured with a VRRP priority of ***150***, it's interface IP address of ***192.168.58.3*** should be shown as master.
State Backup Master State	Bridge 8005 should have a state of ***Backup*** while 8006 should have a state of ***Master*** as it has the higher VRRP priority under ***State***. Likewise, under ***Backup Master State***, 8005 should display ***up*** while 8006 should display ***down***.
Backup Master	Both 8005 and 8006 should display ***enable*** to indicate that VRRP Backup Master has been enabled.

22.9.3.3 IP Route Table

Use the following command to display the routes for each VRF instance

Display IP route table for each VRF instance

```
show ip route vrf blue
```

Results: Example from 4002 where the 192.168.58.0/24 should be populated

4002:

```
4002:1#show ip route vrf blue

=================================================================================================
                                        IP Route - VRF blue
=================================================================================================
                                                    NH                    INTER
DST             MASK            NEXT                VRF/ISID        COST FACE PROT AGE TYPE
PRF
-------------------------------------------------------------------------------------------------
192.168.5.0     255.255.255.0   192.168.5.2         -               1    2255  LOC  0   DB   0
192.168.33.0    255.255.255.0   8005                blue            30   4051  ISIS 0   IBSV 7
192.168.58.0    255.255.255.0   8005                blue            30   4051  ISIS 0   IBSV 7
```

22.10 Multicast over L3VSN

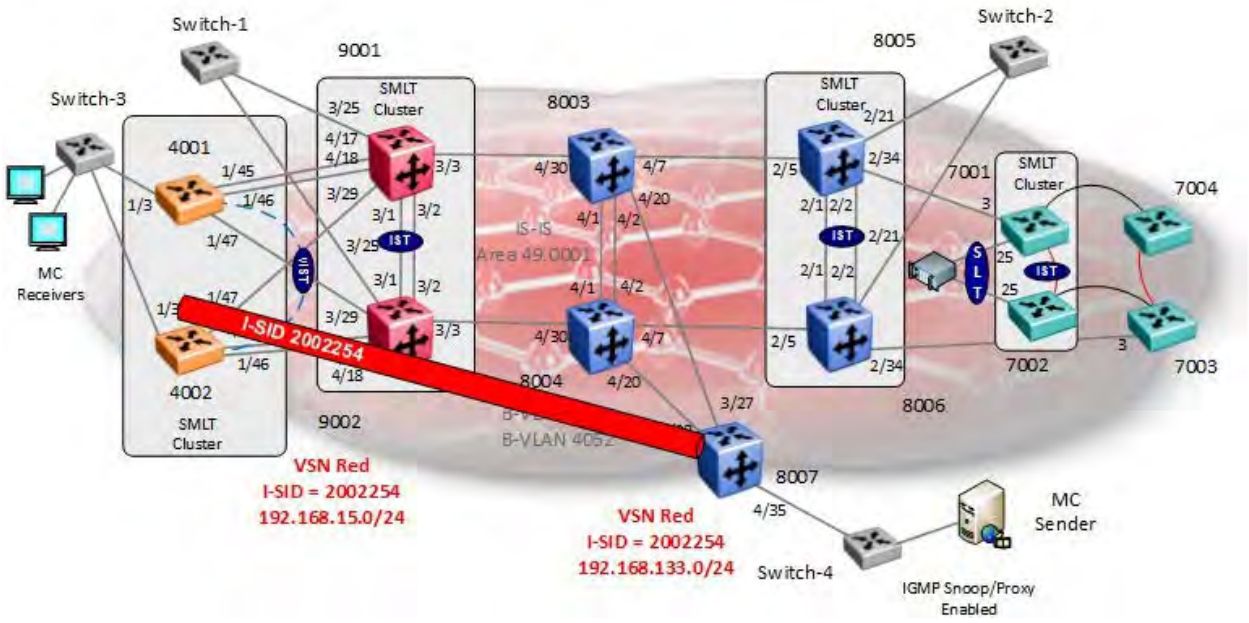

Continuing form example used in Section 22.8, we will simple enable multicast support for L3VSN ISID 2002254 (red vrf) between SPB bridges 4001, 4002, and 8007.

22.10.1Enable SPB Multicast – Global

22.10.1.1 IS-IS Layer 3 configuration

4001, 4002 and 8007: Enable SPB Multicast, global

4001, 4002, and 8007: Same configuration on all switches

```
4001:1(config)#router isis
4001:1(config)#spbm 1 multicast enable
4001:1(config)#exit
```

22.10.2 Enable Multicast VPN

4001, 4002 and 8007: Enable multicast VPN

4001, 4002, and 8007: Same configuration on all switches

```
4001:1(config)#router vrf red
4001:1(router-vrf)#mvpn enable
4001:1(router-vrf)#exit
```

22.10.3 Enable L3 SPB Multicast

4001, 4002, and 8007: Enable L3 SPB multicast at VLAN level

4001, 4002, and 8007: Same configuration on all switches

```
4001:1(config)#interface vlan 2254
4001:1(config-if)#ip spb-multicast enable
4001:1(config-if)#exit
```

22.10.4 Enable IGMP

22.10.4.1 Enable IGMPv2 at VLAN level

Default setting, no configuration required

22.10.4.2 Enable IGMPv3 at VLAN level

4001, 4002, and 8007: Enable IGMPv3, i.e. on VLAN 2254

4001, 4002, and 8007: Same configuration on all switches

```
4001:1(config)#interface vlan 2254
4001:1(config-if)#ip igmp compatibility-mode
4001:1(config-if)#ip igmp version 3
```

22.10.5Edge Switch

Assuming the edge switch is an Extreme stackable switch with the latest firmware, enable IGMP snoop and proxy.

Switch-3 & Switch-4: Enable IGMPv3, i.e. on VLAN 2254

```
ERS-Stackable(config)#interface vlan 2254
ERS-Stackable(config-if)#ip igmp snoop
ERS-Stackable(config-if)#ip igmp proxy
## If IGMPv3 is used:
ERS-Stackable(config-if)#ip igmp version 3
```

22.10.6Verify Operations

22.10.6.1 Global Settings

Verify SPB multicast is enabled

```
show isis spbm multicast
```

Results: From 4001 & 4002

```
4001 4002# show isis spbm multicast
Response from 4001:
                            multicast : enable
                      fwd-cache-timeout : 210


Response from 4002:
                            multicast : enable
                      fwd-cache-timeout : 210
```

22.10.6.2 Verify IGMP interfaces

Verify IGMP interfaces

```
show ip igmp interface vrf <vrf name>
```

Results: From 4001 & 4002

```
4001 4002# show ip igmp interface vrf red
Response from 4001:

================================================================================
                            Igmp Interface - VRF red
================================================================================
       QUERY              OPER            QUERY    WRONG            LASTMEM
IF     INTVL STATUS VERS. VERS QUERIER    MAXRSPT QUERY JOINS ROBUST QUERY MODE
--------------------------------------------------------------------------------------------------------
V2254  125    active 2     2    192.168.15.1 100     0     0     2      10    routed-spb

1 out of 1 entries displayed


Response from 4002:
```

```
================================================================================
                          Igmp Interface - VRF red
================================================================================
       QUERY              OPER            QUERY   WRONG            LASTMEM
IF     INTVL STATUS VERS. VERS QUERIER    MAXRSPT QUERY JOINS ROBUST QUERY MODE
--------------------------------------------------------------------------------------------------------
V2254  125    active 2     2   192.168.15.2 100     0     0     2      10    routed-spb

1 out of 1 entries displayed
```

22.10.6.3 Verify IGMP cache/group and senders

Assuming the multicast sender connect to Switch-4 off SPB bridges 8007 is sending a multicast stream of 232.1.1.1 with a receiver connected to Switch-3 off SPB SMLT cluster bridges 4001 & 4002 .

Step 1 - Verify IGMP cache / group

```
show ip igmp cache vrf <vrf name>
show ip igmp group vrf <vrf name>
```

Results:

```
4001 4002# show ip igmp cache vrf red
Response from 4001:
================================================================================================

================================================================================
                            IGMP Cache - VRF red
================================================================================
================================================================================================
GRPADDR         INTERFACE LASTREPORTER    EXPIRATION              V1HOSTTIMER TYPE    STATICPORTS
------------------------------------------------------------------------------------------------
239.1.1.1       Vlan2254  192.168.15.100  0 day(s), 00h:03m:46s   0           DYNAMIC NULL


Response from 4002:
================================================================================================

================================================================================
                            IGMP Cache - VRF red
================================================================================
================================================================================================
GRPADDR         INTERFACE LASTREPORTER    EXPIRATION              V1HOSTTIMER TYPE    STATICPORTS
```

```
239.1.1.1       Vlan2254  192.168.15.100  0 day(s), 00h:03m:47s   0           DYNAMIC NULL

4001 4002# show ip igmp group vrf red
Response from 4001:

================================================================================
                              Igmp Group - VRF red
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
239.1.1.1       V2254-1/3       192.168.15.100  138        Dynamic


Response from 4002:

================================================================================
                              Igmp Group - VRF red
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------
239.1.1.1       V2254-1/3       192.168.15.100  139        Dynamic
```

Step 2 - Verify IGMP sender

```
show ip igmp sender vrf <vrf name>
```

Results: From 8007

```
8007:5#show ip igmp sender vrf red
================================================================================
                            IGMP Sender - VRF red
================================================================================
                                          PORT/
GRPADDR         IFINDEX    MEMBER         MLT        STATE
--------------------------------------------------------------------------------
239.1.1.1       Vlan 2254  192.168.133.100 4/35       NOTFILTERED
```

22.10.6.4 Verify SPB Multicast Routes

Verify all SPB multicast routes

```
show isis spbm ip-multicast-route vrf <vrf name>
show isis spbm ip-multicast-route vrf <vrf name> detail
show isis spbm ip-multicast-route vrf <vrf name> group <IP addr>
show isis spbm ip-multicast-route vrf <vrf name> group <IP addr> detail
show isis spbm ip-multicast-route vrf <vrf name> group <IP addr> source <ip>
show isis spbm ip-multicast-route vrf <vrf name> group <IP addr> source <ip>
detail
```

Results: From 4001 & 4002

```
4001 4002# show isis spbm ip-multicast-route vrf red detail
Response from 4001:

================================================================================================
                    SPBM IP-MULTICAST ROUTE INFO - VRF NAME : red, VSN-ISID : 2002254
================================================================================================
Source          Group           Data ISID  BVLAN NNI Rcv        UNI Rcvrs         Source-BEB
------------------------------------------------------------------------------------------------
192.168.133.100 239.1.1.1       16000005   4051  -               V2254:1/3         8007

Response from 4002i:

================================================================================================
SPBM IP-MULTICAST ROUTE INFO - VRF NAME : red, VSN-ISID : 2002254
================================================================================================
Source          Group           Data ISID  BVLAN NNI Rcv        UNI Rcvrs         Source-BEB
------------------------------------------------------------------------------------------------
192.168.133.100 239.1.1.1       16000005   4051  -               V2254:1/3         8007
```

22.10.6.5 Verify multicast TLV's

Assuming we have a sender via switch 8007 and receivers via the two SMLT clusters. TLV 185 in relationship to switch 8007 should have the Tx bit set and also send TLV 144 with the Tx bit set. Each multicast group should have its own unique data ISID with a value of 1600000x. The receiver switches (9001, 9002, 8005, and 8006) should have TLV 144 with the Rx bit set.

Step 1 - Verify IP multicast source, group addresses, and Tx bit set on the BEB bridge where the multicast source is located via TLV 185

```
show isis lsdb tlv 185 detail
```

Results: From 4001 perspective taken from 8007

```
4001:1#show isis lsdb sysid 00e0.7bb3.07df tlv 185 detail

================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: 00e0.7bb3.07df.00-00     SeqNum: 0x000002de      Lifetime:   981
        Chksum: 0x5251 PDU Length: 443
        Host_name: 8007
        Attributes:     IS-Type 1
TLV:185 SPBM IPVPN :
                VSN ISID:2002254
                BVID     4051
                        Metric:0 q
                        IP Source Address: 192.168.133.100
                        Group Address    : 239.1.1.1
                        Data ISID        : 16000005
                        TX               : 1
```

Step 2 - Verify on the BEB bridges where the multicast receivers are located via TLV 144, the Rx bit is set with a B-MAC of 03-01-07-00-00-00 (03 indicated multicast while 01-07 is the Nick Name of BEB bridge 8007 with the multicast source). On the BEB bridges where the source is located, the Tx bit should be set

```
show isis lsdb tlv 144 detail
show isis lsdb lspid tlv 144 sub-tlv 3 detail
show isis lsdb lspid <lsp id> tlv 144 detail
show isis lsdb lspid <lsp id> tlv 144 sub-tlv 3 detail
```

Results: Receiver is via SPB bridge 4001

```
4001:1#show isis lsdb sysid d4ea.0e10.e465 tlv 144 sub-tlv 3 detail
================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------------
|
|
Level-1 LspID: d4ea.0e10.e465.00-05     SeqNum: 0x0000000a      Lifetime: 1121
        Chksum: 0xf575 PDU Length: 49
        Host_name: 4001
        Attributes:     IS-Type 1
TLV:144 SUB-TLV 3       ISID:
                Instance: 0
                Metric: 0
                B-MAC: 03-80-07-00-00-00
                BVID:4051
                Number of ISID's:1
                        16000005(Rx)
```

22.10.6.6 Trace Multicast Routes

On the switch where the multicast sender is located, in our example this would be switch 8007, you can trace the multicast route by specifying the source, group, and VLAN.

Verify all SPB multicast routes

```
l2 tracemroute source <source address> group <group address> vlan <C-VLAN id>
vrf <vrf name>
```

Results: Since the multicast source is via switch 8007, we will use the following command to view the multicast route for group address 239.1.1.1

```
8007:5#l2 tracemroute source 192.168.133.100 group 239.1.1.1 vrf red

Please wait for l2tracemroute to complete or press any key to abort

Source : 192.168.133.100
Group : 239.1.1.1
VRF    : red ID 1
BMAC   : 03:80:07:f4:24:05
B-VLAN : 4051
ISID : 16000005

================================================================================
1   8007             00:e0:7b:b3:07:df -> 8003             00:80:2d:be:23:df
2   8003             00:80:2d:be:23:df -> 9001             00:01:81:29:1f:df
3   9001             00:01:81:29:1f:df   -> 4001            d4:ea:0e:10:e4:65
3   9001             00:01:81:29:1f:df   -> 4002            a0:12:90:d3:ec:65
```

22.11 SPB IP Shortcuts

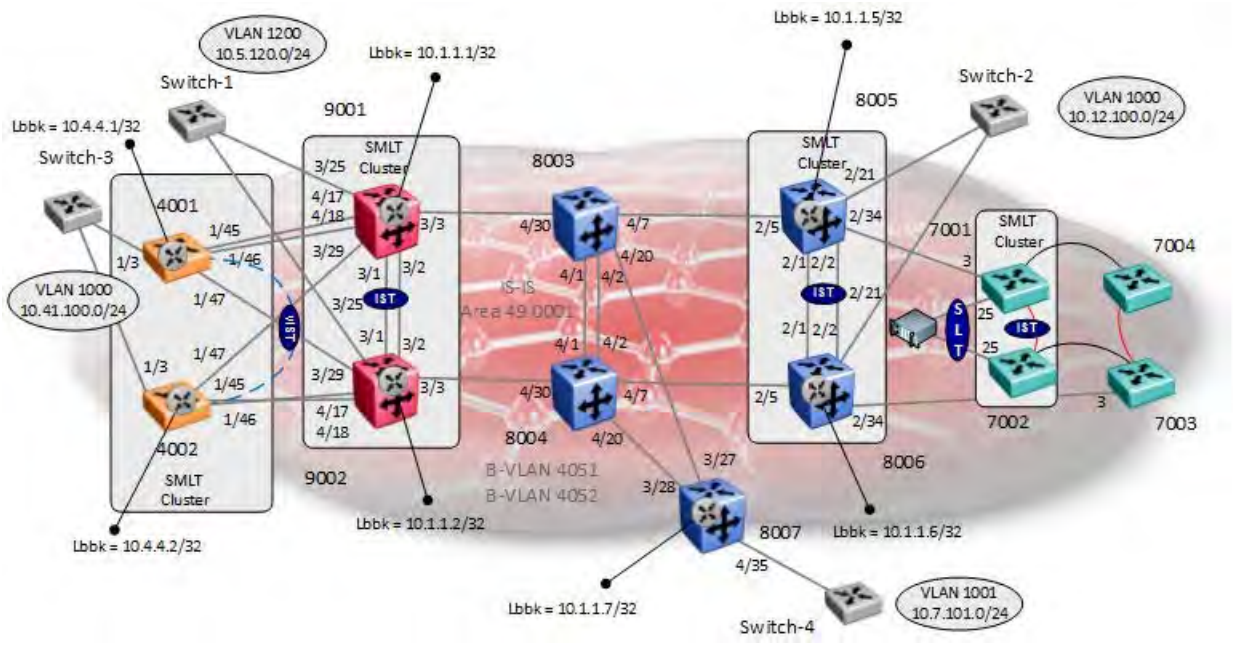

- SPB IP
  - SPB IP parameter must be enabled on BEB bridges 4001, 4002, 9001, 9002, 8005, 8006, and 8007
  - An IS-IS source IP address must be configured (loopback/curcuitless IP address)
- IP Configuration
  - CLIP/Loopback #1 as shown in the above diagram
  - Local VLAN and IP addressing as shown in the above diagram
  - Redistribution of direct interfaces to IS-IS (SPB) on each BEB bridge
    - Please note, on the SMLT cluster, a route policy must be create to deny the IST subnet as by default, all local interfaced will be redistributed into IS-IS unless if you wish to distribute the IST network

This example is a continuation from the base setup used in Section 22.1 and SMLT configuration in Section 22.2.

22.11.1 IS-IS Layer 3 configuration

VSP 4000 Switches - Create Loopback IP address for the IS-IS source address and enable SPB IP

4001:

```
4001:1(config)#interface loopback 1
4001:1(config-if)#ip address 10.4.4.1/32
4001:1(config-if)#exit
4001:1(config)#router isis
4001:1(config-isis)#ip-source-address 10.4.4.1
4001:1(config-isis)#spbm 1 ip enable
4001:1(config-isis)#exit
```

4002:

```
4002:1(config)#interface loopback 1
4002:1(config-if)#ip address 10.4.4.2/32
4002:1(config-if)#exit
4002:1(config)#router isis
4002:1(config-isis)#ip-source-address 10.4.4.2
4002:1(config-isis)#spbm 1 ip enable
4002:1(config-isis)#exit
```

VSP 9000 Switches - Create Loopback IP address for the IS-IS source address and enable SPB IP

9001:

```
9001:1(config)#interface loopback 1
9001:1(config-if)#ip address 10.1.1.1/32
9001:1(config-if)#exit
9001:1(config)#router isis
9001:1(config-isis)#ip-source-address 10.1.1.1
9001:1(config-isis)#spbm 1 ip enable
9001:1(config-isis)#exit
```

9002:

```
9002:1(config)#interface loopback 1
9002:1(config-if)#ip address 10.1.1.2/32
9002:1(config-if)#exit
9002:1(config)#router isis
9002:1(config-isis)#ip-source-address 10.1.1.2
9002:1(config-isis)#spbm 1 ip enable
9002:1(config-isis)#exit
```

ERS 8800 Switches - Create Loopback IP address for the IS-IS source address and enable SPB IP

8005:

```
8005:5(config)#interface loopback 1
8005:5(config-if)#ip address 10.1.1.5/32
8005:5(config-if)#exit
8005:5(config)#router isis
8005:5(config-isis)#ip-source-address 10.1.1.5
8005:5(config-isis)#spbm 1 ip enable
8005:5(config-isis)#exit
```

8006:

```
8006:5(config)#interface loopback 1
8006:5(config-if)#ip address 10.1.1.6/32
8006:5(config-if)#exit
8006:5(config)#router isis
8006:5(config-isis)#ip-source-address 10.1.1.6
8006:5(config-isis)#spbm 1 ip enable
8006:5(config-isis)#exit
```

8007:

```
8007:5(config)#interface loopback 1
8007:5(config-if)#ip address 10.1.1.7/32
8007:5(config-if)#exit
8007:5(config)#router isis
8007:5(config-isis)#ip-source-address 10.1.1.7
8007:5(config-isis)#spbm 1 ip enable
8007:5(config-isis)#exit
```

22.11.1.1 Redistribute direct interfaces

For the SMLT cluster switches, we will also add a policy to suppress the IST interface

VSP 4000 Switches - Create Loopback IP address for the IS-IS source address, enable SPB IP, and create route-map to suppress the IST network

4001 and 4002: Same configuration on both switches

```
4001:1(config)#ip prefix-list IST 10.4.2.0/30
4001:1(config)#route-map suppressIST 1
4001:1(route-map)#no permit
4001:1(route-map)#enable
4001:1(route-map)#match network "IST"
4001:1(route-map)#exit
4001:1(config)#route-map suppressIST 2
4001:1(route-map)#enable
4001:1(route-map)#match protocol local
4001:1(route-map)#exit
4001:1(config)#router isis
4001:1(config-isis)#redistribute direct
4001:1(config-isis)#redistribute direct route-map suppressIST
4001:1(config-isis)#redistribute direct enable
4001:1(config-isis)#exit
4001:1(config)#isis apply redistribute direct
```

VSP 9000 Switches - Create Loopback IP address for the IS-IS source address, enable SPB IP, and create route-map to suppress the IST network

9001 and 9002: Same configuration on both switches

```
9001:1(config)#ip prefix-list IST 10.5.2.0/30
9001:1(config)#route-map suppressIST 1
9001:1(route-map)#enable
9001:1(route-map)#match network IST
9001:1(route-map)#exit
9001:1(config)#route-map suppressIST 1 deny
9001:1(config)#route-map suppressIST 2
9001:1(route-map)#enable
9001:1(route-map)#match protocol local
9001:1(route-map)#exit
9001:1(config)#router isis
9001:1(config-isis)#redistribute direct
9001:1(config-isis)#redistribute direct route-map suppressIST
```

```
9001:1(config-isis)#redistribute direct enable
9001:1(config-isis)#exit
9001:1(config)#isis apply redistribute direct
```

8005 and 8006 - Create Loopback IP address for the IS-IS source address, enable SPB IP, and create route policy to suppress the IST network

8005 and 8006: Same configuration on both switches

```
8005:5(config)#ip prefix-list IST 10.2.1.0/30
8005:5(config)#route-map suppressIST 1
8005:5(route-map)#no permit
8005:5(route-map)#enable
8005:5(route-map)#match network IST
8005:5(route-map)#exit
8005:5(config)#route-map suppressIST 2
8005:5(route-map)#enable
8005:5(route-map)#match protocol local
8005:5(route-map)#exit
8005:5(config)#router isis
8005:5(config-isis)#redistribute direct
8005:5(config-isis)#redistribute direct route-map suppressIST
8005:5(config-isis)#redistribute direct enable
8005:5(config-isis)#exit
8005:5(config)#isis apply redistribute direct
```

8007 - Create Loopback IP address for the IS-IS source address and enable SPB IP

8007:

```
8007:5(config)#router isis
8007:5(config-isis)#redistribute direct
8007:5(config-isis)#redistribute direct
8007:5(config-isis)#redistribute direct enable
8007:5(config-isis)#exit
8007:5(config)#isis apply redistribute direct
```

22.11.2 ECMP

Enable ECMP using the following command

- `ip ecmp`

22.11.3 Local VLAN configuration

VSP 4000 Switches - Create local VLAN, add ISID for the vIST, add IP address to VLAN, and ether enable RSMLT Edge or VRRP with Backup Master. For this example, we will enable RSMLT Edge

4001:

```
4001:1(config)#vlan create 1000 type port-mstprstp 0
4001:1(config)#vlan ISID 1000 3001000
4001:1(config)#vlan mlt 1000 2
4001:1(config)#interface vlan 1000
4001:1(config-if)#ip address 10.41.100.1 255.255.255.0
4001:1(config-if)#ip rsmlt
4001:1(config-if)#ip rsmlt holdup-timer 9999
4001:1(config-if)#exit
4001:1(config)#ip rsmlt edge-support
```

4002: Same configuration as 4001 except for the IP address

```
4002:1(config-if)#ip address 10.41.100.2 255.255.255.0
```

8005 & 8006 SMLT Cluster Switches

8005:

```
8005:5(config)#vlan create 1000 type port-mstprstp 0
8005:5(config)#vlan members add 1000 2/21
8005:5(config)#interface vlan 1000
8005:5(config-if)#ip address 10.12.100.1 255.255.255.0
8005:5(config-if)#ip rsmlt
8005:5(config-if)#ip rsmlt holdup-timer 9999
8005:5(config-if)#exit
8005:5(config)#ip rsmlt edge-support
```

8006: Same configuration as 8005 except for the IP address

```
8005:5(config-if)#ip address 10.12.100.1 255.255.255.0
```

9001 & 9002 SMLT Cluster Switches

9001 & 9002: Same configuration on both switches

```
9001:1(config)#vlan create 1200 type port-mstprstp 0
9001:1(config)#vlan mlt 1000 2
9001:1(config)#vlan mlt 1000 1
9001:1(config)#interface vlan 1000
9001:1(config-if)#ip address 10.12.100.1 255.255.255.0
9001:1(config-if)#ip rsmlt
9001:1(config-if)#ip rsmlt holdup-timer 9999
9001:1(config-if)#exit
9001:1(config)#ip rsmlt edge-support
```

9002: Same configuration as 9001 except for the IP address

```
9001:1(config-if)#ip address 10.5.120.1 255.255.255.0
```

8007

```
8007:5(config)#vlan create 1001 type port-mstprstp 0
8007:5(config)#vlan members add 1001 4/35
8005:5(config)#interface vlan 1001
8005:5(config-if)#ip address 10.7.101.1 255.255.255.0
8005:5(config-if)#exit
```

22.11.4 Verify Operations

22.11.4.1 Verify IP Route Table

Verify IP Routes

```
show ip route
```

Results: From bridge 4001

4001:

```
================================================================================
                          IP Route - GlobalRouter
================================================================================
                                                    NH          INTER
DST              MASK            NEXT               VRF    COST FACE PROT AGE TYPE PRF
--------------------------------------------------------------------------------------------------------
10.1.1.1         255.255.255.255 9001        GlobalRouter     10   4051  ISIS 0   IBSE 7
10.1.1.1         255.255.255.255 9001        GlobalRouter     10   4052  ISIS 0   IBSE 7
10.1.1.2         255.255.255.255 9002        GlobalRouter     10   4051  ISIS 0   IBSE 7
10.1.1.2         255.255.255.255 9002        GlobalRouter     10   4052  ISIS 0   IBSE 7
10.1.1.5         255.255.255.255 8005        GlobalRouter     30   4051  ISIS 0   IBSE 7
10.1.1.5         255.255.255.255 8005        GlobalRouter     30   4052  ISIS 0   IBSE 7
10.1.1.6         255.255.255.255 8006        GlobalRouter     30   4051  ISIS 0   IBSE 7
10.1.1.6         255.255.255.255 8006        GlobalRouter     30   4052  ISIS 0   IBSE 7
10.1.1.7         255.255.255.255 8007        GlobalRouter     30   4051  ISIS 0   IBSE 7
10.1.1.7         255.255.255.255 8007        GlobalRouter     30   4052  ISIS 0   IBSE 7
10.4.4.1         255.255.255.255 10.4.4.1    -                1    0     LOC  0   DB   0
10.4.4.2         255.255.255.255 4002        GlobalRouter     20   4051  ISIS 0   IBSE 7
10.4.4.2         255.255.255.255 4002        GlobalRouter     20   4052  ISIS 0   IBSE 7
10.5.120.0       255.255.255.0   9001        GlobalRouter     10   4051  ISIS 0   IBSE 7
10.5.120.0       255.255.255.0   9002        GlobalRouter     10   4051  ISIS 0   IBSE 7
10.5.120.0       255.255.255.0   9001        GlobalRouter     10   4052  ISIS 0   IBSE 7
10.5.120.0       255.255.255.0   9002        GlobalRouter     10   4052  ISIS 0   IBSE 7
10.7.101.0       255.255.255.0   8007        GlobalRouter     30   4051  ISIS 0   IBSE 7
10.7.101.0       255.255.255.0   8007        GlobalRouter     30   4052  ISIS 0   IBSE 7
10.12.100.0      255.255.255.0   8005        GlobalRouter     30   4051  ISIS 0   IBSE 7
10.12.100.0      255.255.255.0   8006        GlobalRouter     30   4051  ISIS 0   IBSE 7
10.12.100.0      255.255.255.0   8005        GlobalRouter     30   4052  ISIS 0   IBSE 7
10.12.100.0      255.255.255.0   8006        GlobalRouter     30   4052  ISIS 0   IBSE 7
10.41.100.0      255.255.255.0   10.41.100.1  -               1    1000  LOC  0   DB   0


--------------------------------------------------------------------------------------------------------
```

```
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route
PROTOCOL Legend:
v=Inter-VRF route redistributed
```

To display the B-MAC for the attribute “NEXT”, enter the CLI command *show ip route info spbm-nh-as-mac* or ACLI command *show ip route spbm-nh-as-mac*.

22.11.4.2 Verify IS-IS SPB IP Unicast FIB

Verify IP Routes from remote BEBs

```
show isis spbm ip-unicast-fib
```

Results: From bridge 4001

4001:

```
================================================================================
                        SPBM IP-UNICAST FIB ENTRY INFO
================================================================================
                                                          OUTGOING  SPBM       PREFIX
VRF       ISID    Destination       NH BEB            VLAN INTERFACE COST       COST
--------------------------------------------------------------------------------------
GRT       -       10.1.1.1/32       9001              4051 9001      10         1
GRT       -       10.1.1.1/32       9001              4052 9001      10         1
GRT       -       10.1.1.2/32       9002              4051 1/47      10         1
GRT       -       10.1.1.2/32       9002              4052 1/47      10         1
GRT       -       10.1.1.5/32       8005              4051 9001      30         1
GRT       -       10.1.1.5/32       8005              4052 9001      30         1
GRT       -       10.1.1.6/32       8006              4051 1/47      30         1
GRT       -       10.1.1.6/32       8006              4052 1/47      30         1
GRT       -       10.1.1.7/32       8007              4051 9001      30         1
GRT       -       10.1.1.7/32       8007              4052 1/47      30         1
GRT       -       10.4.4.2/32       4002              4051 9001      20         1
GRT       -       10.4.4.2/32       4002              4052 1/47      20         1
GRT       -       10.5.120.0/24     9001              4051 9001      10         1
GRT       -       10.5.120.0/24     9001              4052 9001      10         1
GRT       -       10.5.120.0/24     9002              4051 1/47      10         1
GRT       -       10.5.120.0/24     9002              4052 1/47      10         1
GRT       -       10.7.101.0/24     8007              4051 9001      30         1
```

GRT	-	10.7.101.0/24	8007	4052 1/47	30	1
GRT	-	10.12.100.0/24	8005	4051 9001	30	1
GRT	-	10.12.100.0/24	8005	4052 9001	30	1
GRT	-	10.12.100.0/24	8006	4051 1/47	30	1
GRT	-	10.12.100.0/24	8006	4052 1/47	30	1
GRT	-	10.41.100.0/24	4002	4051 9001	20	1
GRT	-	10.41.100.0/24	4002	4052 1/47	20	1

22.11.4.3 Verify IS-IS Extended IP Reachability TLV (135)

IS-IS uses TLV 135 for extended IP reachability. You can view TLV 135 details by issuing the command shown below.

Verify TLV 135 details

```
show isis lsdb tlv 135 detail
show isis lsdb lspid <isis system id>.00-00 tlv 135 detail
show isis lsdb sysid <isis system id> tlv 135 detail
```

Results: From bridge 8007

4001:

```
4001:1#show isis lsdb lspid 00e0.7bb3.07df.00-00 tlv 135 detail

================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: 00e0.7bb3.07df.00-00     SeqNum: 0x000004f9      Lifetime: 1123
        Chksum: 0xa95b PDU Length: 425
        Host_name: 8007
        Attributes:     IS-Type 1
TLV:135 TE IP Reachability: 19
                Metric: 1       Prefix Length: 32
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 10.1.1.7
                Metric: 1       Prefix Length: 24
                UP/Down Bit: FALSE              Sub TLV Bit: FALSE
                IP Address: 10.7.101.0
```

22.12 Multicast over IP Shortcuts

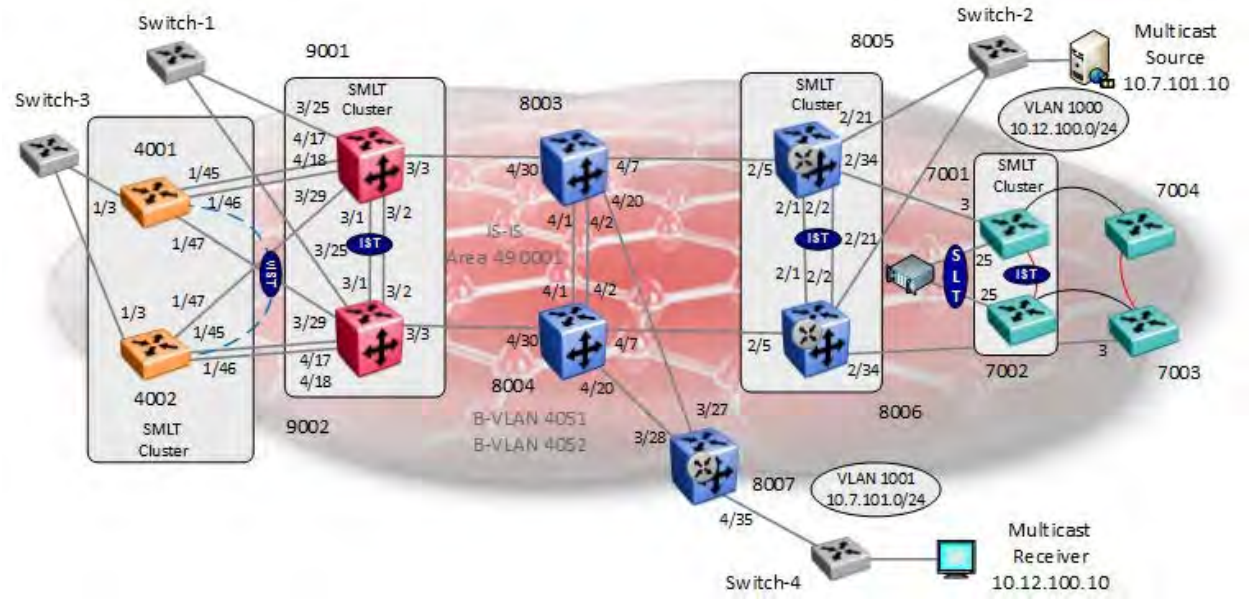

Continuing form example used in Section 22.11, we will simple enable multicast support for IP Shortcuts on all SPB bridges.

22.12.1 IP Shortcuts Multicast configuration

Enable IP multicast globally

4001, 4002, 9001, 9002, 8005, 8006, and 8007: Same configuration on all switches

```
8005:5(config)#router isis
8005:5(config-isis)#spbm 1 multicast enable
8005:5(config-isis)#exit
```

22.12.2 Enable IP Multicast at VLAN level

Enable IP multicast at VLAN level

4001, 4002, 9001, 9002, 8005, 8006, and 8007: Same configuration on all switches

```
8005:5(config)#interface vlan 1000
8005:5(config-isis)#ip spb-multicast enable
8005:5(config-isis)#exit
```

```
Enable IGMPv3 if used, default is IGMPv2
```

```
8005:5(config)#interface vlan 1000
8005:5(config-if)#ip igmp version 3
```

Enable IP multicast at VLAN level

4001, 4002, 9001, 9002, 8005, 8006, and 8007: Same configuration on both switches

```
8007:5(config)#interface vlan 1001
8007:5(config-isis)#ip spb-multicast enable
8007:5(config-isis)#exit
```

```
Enable IGMPv3 if used, default is IGMPv2
```

```
8007:5(config)#interface vlan 1001
8007:5(config-if)#ip igmp version 3
```

22.13 Verify Operations

22.13.1 Global Settings

Verify SPB multicast is enabled

```
show isis spbm multicast
```

Results: From bridge 8007

8007:

```
================================================================================
                                ISIS SPBM Info
================================================================================
SPBM       B-VID      PRIMARY    NICK     LSDB     IP       MULTICAST
INSTANCE              VLAN       NAME     TRAP
--------------------------------------------------------------------------------
1          4051-4052  4051       0.80.07 disable enable  enable

================================================================================
                              ISIS SPBM SMLT Info
================================================================================
SPBM       SMLT-SPLIT-BEB       SMLT-VIRTUAL-BMAC     SMLT-PEER-SYSTEM-ID
INSTANCE
--------------------------------------------------------------------------------
1          primary              00:00:00:00:00:00
```

22.13.2 Verify IGMP cache/group and senders

Assuming the multicast sender connect to Switch-2 (via 8005 and 8006) is sending a multicast stream using a group address of 232.1.1.1 while a receivers off Switch-3 joins this group.

Step 1 - Verify IGMP cache / group

```
show ip igmp cache
show ip igmp group
```

Results: From bridge 8007

8007:

```
8007:5#show ip igmp cache
================================================================================
                          IGMP Cache - GlobalRouter
================================================================================
GRPADDR         INTERFACE  LASTREPORTER    EXPIRYTIME        VERSION1HOSTTIMER TYPE
STATICPORTS
--------------------------------------------------------------------------------------------------------
232.1.1.1       Vlan1001   10.7.101.10     0day,00h:03m:41s  0day,00h:00m:00s    DYNAMIC NULL


1 out of 1 entries displayed


8007:5#show ip igmp group
================================================================================
                          IGMP Group - GlobalRouter
================================================================================
GRPADDR         INPORT          MEMBER          EXPIRATION TYPE
--------------------------------------------------------------------------------------------------------
232.1.1.1       V1001-4/35      10.7.101.10     217        Dynamic


1 out of 1 group Receivers displayed
```

Step 2 - Verify IGMP sender

```
show ip sender
```

Results: From 8005 and 8006 – the SPB bridge where the sender is located

```
8005 8006> show ip igmp sender
Response from 8005:
================================================================================
                           IGMP Sender - GlobalRouter
================================================================================
                                          PORT/
GRPADDR         IFINDEX    MEMBER         MLT        STATE
--------------------------------------------------------------------------------
232.1.1.1       Vlan 1000  10.12.100.10   2/21       NOTFILTERED

1 out of 1 entries displayed

Response from 8006:
================================================================================
                           IGMP Sender - GlobalRouter
================================================================================
                                          PORT/
GRPADDR         IFINDEX    MEMBER         MLT        STATE
--------------------------------------------------------------------------------
232.1.1.1       Vlan 1000  10.12.100.10   2/21       NOTFILTERED

1 out of 1 entries displayed
```

22.13.3 Verify SPB Multicast Routes

Assuming the multicast sender connect to switch 8007 is sending multicast stream 232.1.1.1 while the receivers joins this group.

Verify IGMP cache / group

```
show isis spbm ip-multicast-route
show isis spbm ip-multicast-route all
show isis spbm ip-multicast-route info detail
show isis spbm ip-multicast-route info group <IP addr>
show isis spbm ip-multicast-route info group <IP addr> detail
show isis spbm ip-multicast-route info group <IP addr> source <ip>
show isis spbm ip-multicast-route info group <IP addr> source <ip> detail
```

Results: From bridge 8007

8007:

```
================================================================================================
                                        SPBM IP-MULTICAST ROUTE INFO
================================================================================================
Source          Group           Data ISID  BVLAN Source-BEB
------------------------------------------------------------------------------------------------
10.12.100.10    232.1.1.1       16000008   4051  8005
10.12.100.10    232.1.1.1       16000009   4052  8006

------------------------------------------------------------------------------------------------
 Total Number of SPBM IP MULTICAST ROUTE Entries: 2
------------------------------------------------------------------------------------------------
```

22.13.4 Verify multicast TLV's

Assuming we have a sender via the SMLT cluster 8005 and 8006 and a receiver via 8007. TLV 186 in relationship to switch 8005 & 8006 should have the Tx bit set and also send TLV 144 with the Tx bit set. Each multicast group should have its own unique data ISID with a value of 1600000x. The receiver bridges (8007) should have TLV 144 with the Rx bit set.

Step 1 - Verify IP multicast source, group addresses, and Tx bit set on the BEB bridge where the multicast source is located via TLV 186.

```
show isis lsdb tlv 186 detail
```

Results: From bridge 8007

```
8007:5#show isis lsdb tlv 186 detail
================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: 0024.43b4.e3df.00-00     SeqNum: 0x00000da4      Lifetime:   422
        Chksum: 0xfb52 PDU Length: 903
        Host_name: 8005
        Attributes:     IS-Type 1
TLV:186 SPBM IP Multicast:
                GRT ISID
                        Metric:0
                        IP Source Address: 10.12.100.10
                        Group Address    : 232.1.1.1
                        Data ISID        : 16000008
                        BVID             : 4051
                        TX               : 1
                        Route Type       : Internal


Level-1 LspID: 001e.1f48.f3df.00-00     SeqNum: 0x00001810      Lifetime:   646
        Chksum: 0x8556 PDU Length: 880
        Host_name: 8006
        Attributes:     IS-Type 1
                GRT ISID
                        Metric:0
                        IP Source Address: 10.12.100.10
                        Group Address    : 232.1.1.1
                        Data ISID        : 16000009
```

```
                        BVID            : 4052
                        TX              : 1
                        Route Type      : Internal
```

Step 2 – Verify on the BEB bridges where the multicast receivers are located via TLV 144, the Rx bit is set with a B-MAC of 03-08-05-00-00-00 and 03-08-06-00-00-00 (03 indicated multicast while 08-05 is the Nick Name of BEB bridge 8005 and 08-06 is the Nick Name of the BEB bridge 8006 with the multicast source). On the BEB bridges where the source is located, the Tx bit should be set

```
show isis lsdb tlv 144 detail
show isis lsdb lspid tlv 144 sub-tlv 3 detail
show isis lsdb lspid <lsp id> tlv 144 detail
show isis lsdb lspid <lsp id> tlv 144 sub-tlv 3 detail
```

Results: From bridge 8007

```
8007:5#show isis lsdb lspid 0024.43b4.e3df.00-00 tlv 144 detail
================================================================================
                               ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: 0024.43b4.e3df.00-00     SeqNum: 0x00000da5      Lifetime:   950
        Chksum: 0xf953 PDU Length: 903
        Host_name: 8005
        Attributes:     IS-Type 1
                Instance: 0
                Metric: 0
                B-MAC: 03-00-00-00-00-00
                BVID:4051
                Number of ISID's:1
                        16000008(Tx)
8007> show isis lsdb lspid 001e.1f48.f3df.00-00 tlv 144 detail
================================================================================
                               ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: 001e.1f48.f3df.00-00     SeqNum: 0x00001811      Lifetime: 1124
        Chksum: 0x8357 PDU Length: 880
        Host_name: 8006
        Attributes:     IS-Type 1
```

```
                Instance: 0
                Metric: 0
                B-MAC: 03-00-00-00-00-00
                BVID:4052
                Number of ISID's:1
                        16000009(Tx)

8007:5#show isis lsdb lspid 00e0.7bb3.07df.00-00 tlv 144 detail
================================================================================
                              ISIS LSDB (DETAIL)
================================================================================
--------------------------------------------------------------------------------------------------------------

Level-1 LspID: 00e0.7bb3.07df.00-00     SeqNum: 0x00000504      Lifetime: 1136
        Chksum: 0x4ca2 PDU Length: 461
        Host_name: 8007
        Attributes:     IS-Type 1
TLV:144 SUB-TLV 1       SPBM INSTANCE:
                Instance: 0
                Metric: 0
                B-MAC: 03-08-05-00-00-00
                BVID:4051
                Number of ISID's:1
                        16000008(Rx)


                Instance: 0
                Metric: 0
                B-MAC: 03-08-06-00-00-00
                BVID:4052
                Number of ISID's:1
                        16000009(Rx)
```

22.13.5 Trace Multicast Routes

On the switch where the multicast sender is located, in our example this would be switch 8005 and 8006, you can trace the multicast route by specifying the source, group, and VLAN.

Verify all SPB multicast routes

```
l2 tracemroute source <source address> group <group address> vlan <C-VLAN id>
```

Results: From bridge 8007

```
8005 8006> l2 tracemroute source 10.12.100.10 group 232.1.1.1
Response from 8005:
 Please wait for l2tracemroute to complete or press any key to abort
Source : 10.12.100.10
Group  : 232.1.1.1
VRF    : GRT  ID 0
BMAC   : 03:08:05:f4:24:08
B-VLAN : 4051
ISID   : 16000008
================================================================================
1   8005             00:24:43:b4:e3:df -> 8003             00:80:2d:be:23:df
2   8003             00:80:2d:be:23:df -> 8007             00:e0:7b:b3:07:df

Response from 8006:
 Please wait for l2tracemroute to complete or press any key to abort
Source : 10.12.100.10
Group  : 232.1.1.1
VRF    : GRT  ID 0
BMAC   : 03:08:06:f4:24:09
B-VLAN : 4052
ISID   : 16000009
================================================================================
1   8006             00:1e:1f:48:f3:df -> 8004             00:e0:7b:bc:23:df
2   8004             00:e0:7b:bc:23:df -> 8007             00:e0:7b:b3:07:df
```

23. Restrictions and Limitations

23.1 STP/RSTP/MSTP

- SPB is not supported in RSTP mode
- C-VLAN level loop across SPB NNI ports can't be detected and need to be solved at provisional level.
- SPB NNI ports are not part of L2VSN C-VLAN and BPDU are not transmitted over the SPB tunnel. SPB can only guarantee loop-free topologies consisting of the NNI ports.
- SPB uses STG 63/MSTI 62 internally so 62/63 can't be used by other VLAN/MSTI. If STG/MSTI 62/63 is used in the configuration on non-SPB customer network, then STG/MSTI 61 is used internally.
- When an ISIS interface is created, MSTP is automatically disabled for MSTI-62 allowing traffic for the B-VLAN to be forwarded and not blocked. This also allows traffic for other VLANs other than the BVLANs to co-exist with SPBM on the same interface.
- SPB B-VLANs need to be configured on all bridges as well in the same MSTP region. This is required by MSTP itself to generate the correct digest if MSTP is in use and configured with Regional settings. In MSTP mode, when a C-VLAN is created on the BEB, make sure the same VLAN is created on all switches in the same MSTP region to have correct digest.

23.2 SPB IS-IS

- IP IS-IS

IP over IS-IS is not supported. IS-IS protocol is only to facilitate SPB.

- Level 1 IS-IS Only

SPB only use level 1 IS-IS. Level 2 IS-IS is currently not supported.

- Wide Metric Only

IS-IS standard defines wide (32bit) metric and narrow (8 bits) metrics. Only wide metric is supported.

- IS-IS HA – ERS 8800 and VSP 9000

SPB support full HA (High Availability). SPB and IS-IS configuration and dynamic information (adjacencies, LSPs etc.) are all HA synced to the standby CPU to ensure seamless switchover.

Switching between the CPUs is very quick - there is a sub-second second gap between the active CPU down and the standby CPU up.

To avoid IS-IS adjacencies bounce during switchover, the default hello interval value of 9 seconds and hello multiple of 3 are good for most normal configurations. They may need to be increased depending on overall system load.

- IS-IS sys-name

By default, the IS-IS sys-name is derived from the global system name setting. If you do set the IS-IS sys-name parameter, please ensure that a different value from the global system name is used.

24. Reference Documentation

Document Title	Publication Number	Description
Release Notes for VOSS	Various	Release Notes for VSP Operating System Software. Refer to the Release Notes for the specific software version.