



Ethernet Routing Switch

4800, 8800

Virtual Services Platform

4000, 7000, 8000, 9000

Engineering

> Shortest Path Bridging (802.1aq) Technical Configuration Guide

Extreme Networks

Document Date: December 2017

Document Number: NN48500-617

Document Version: 2.3

© 2017, Extreme Networks, Inc.

All Rights Reserved.

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](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](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/contact/> 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 permitted to use such Marks without prior written consent from Extreme Networks or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Extreme Networks or the applicable third party. Extreme Networks is a registered trademark of Extreme Networks, Inc.

All non-Extreme Networks trademarks are the property of their respective owners. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

For additional information on Extreme Networks trademarks, please see: <http://www.extremenetworks.com/company/legal/>

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 9000, ERS 4800, 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
- GRT : Global Route Table
- I-SID : 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
- 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. Changes text in various sections
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.

Table of Contents

Figures	11
Tables.....	12
1. Overview	14
1.1 Evolution of Ethernet Bridging.....	14
1.2 SPB Benefits	16
2. SPB Terminology	19
2.1 SPB	19
2.2 SPBM	19
2.3 IS-IS	19
2.4 B-VLAN	19
2.5 B-MAC.....	20
2.6 System ID.....	21
2.7 Backbone Edge Bridges (BEB) and Backbone Core Bridges (BCB).....	21
2.8 Connectivity Fault Management (CFM)	21
3. SPB Support Topologies.....	23
3.1 SPB L2 VSN.....	23
3.2 SPB L3 VSN.....	25
3.3 Inter VSN Routing	26
3.4 SPB IP Shortcuts	27
3.5 IP VPN Lite over SPB	28
4. UNI Types	29
4.1 L2VSN – C-VLAN UNI	29
4.2 L2VSN – Switched UNI	30
4.3 L2VSN – Transparent UNI	31
4.4 UNI Type – Example	32
5. Summary of SPB Features and Product Release Matrix.....	33
6. SPB Feature and License Matrix	34
7. Migration & Upgrades	35
7.1 Common Upgrade instructions.....	35
7.2 Upgrade from Pre-5.1 releases for the ERS 8800.....	35
7.3 Upgrade and SMLT Cluster	35
7.4 VSP 4000 and ERS 4850.....	35
7.5 VSP 7000	36
7.6 Activating SPB.....	38

7.7	Migrating traffic to SPB	40
7.8	Multicast	41
7.9	Migrating a VLAN to an L2 VSN.....	43
7.9.1	Migrating to Inter VSN Routing.....	43
7.10	VSP 9000 Notes.....	44
8.	Field Introduction & Support Specifications	45
8.1	Hardware and Deployment Specifications	45
8.2	Installation and Commissioning Specifications	45
8.3	Interoperability and Backwards / Forward Compatibility Specifications.....	45
9.	VSP 7000 – Fabric Interconnect.....	46
10.	ISIS Metrics - Optional	48
11.	SPB SMLT BEB Design Best Practices.....	49
11.1	SMLT BEB – C-VLAN Guidelines for L2VSN	49
11.2	SMLT BEB – ISIS Hello Timer Guidelines	50
11.3	SMLT BEB – RSMLT	51
11.4	SMLT BEB – VLACP Guidelines.....	52
11.5	SMLT BEB – VSP 7000 Guidelines	53
11.6	SMLT BEB – Virtual Inter-Switch Trunk (vIST)	53
11.7	SLPP Guard – ERS 4800.....	54
12.	SPB NNI SMLT – migrating existing SMLT network to SPB	55
13.	IS-IS TLV	59
14.	SPB Best Practices	60
15.	SPB Configuration.....	63
15.1	SPB Configuration.....	65
15.1.1	SPB and IS-IS Core Configuration.....	65
15.1.2	SPB NNI Interface Configuration	67
15.1.3	VSP 7000 – Fabric Interconnect Mesh	68
15.1.4	L2VSN Configuration	69
15.1.5	Inter-ISID Routing	70
15.1.6	L3VSN Configuration	72
15.1.7	IP Shortcuts	73
15.1.8	SPB Multicast Configuration	75
15.1.9	SMLT – Normal IST	77
15.1.10	Virtual IST	77
15.1.11	Connectivity Fault Management (CFM) Configuration	78
15.1.12	CFM Configuration Example – 7.1.1.x or higher	80
15.2	Using EDM	81

15.2.1	IS-IS and SPB Configuration	81
15.2.2	VSN Configuration	83
15.2.3	Connectivity Fault Management (CFM) Configuration – release 7.0 or 7.1.1.	86
16.	Configuration Examples	89
16.1	SPB – Core Setup	89
16.1.1	Configuration.....	91
16.1.2	Configuration using EDM – Using 8005 as an example.....	120
16.1.3	Verify Operations	128
16.2	SPB L2 VSN.....	153
16.2.1	VLAN and SMLT configuration.....	154
16.2.2	Layer 2 VSN configuration	156
16.2.3	Verify Operations	157
16.3	VSP 7000 & ERS 4800 – In-band Management via L2VSN.....	164
16.4	Multicast over L2VSN.....	167
16.4.1	Enable SPB Multicast – Global	168
16.4.2	Enable IGMP.....	168
16.4.3	Verify Operations	170
16.5	Inter VSN Routing	179
16.6	Inter-ISID Configuration	180
16.6.1	VRF configuration	180
16.6.2	Verification	181
16.7	SPB L3 VSN – SMLT	186
16.7.1	SPB IP Enable	187
16.7.2	VLAN Configuration	188
16.7.3	IPVPN Configuration.....	189
16.7.4	Enable L3VSN Configuration	190
16.7.5	Enable direct interface redistribution.....	192
16.7.6	Verify Operations	194
16.8	Extending L3VSN to the ERS 4800 via L2VSN	205
16.8.1	L2VSN Configuration	206
16.8.2	VRF Configuration	206
16.8.3	Verify Operations	207
16.9	Multicast over L3VSN.....	209
16.9.1	Enable SPB Multicast – Global	210
16.9.2	Enable Multicast VPN	210
16.9.3	Enable L3 SPB Multicast	210
16.9.4	Enable IGMP.....	210
16.9.5	Edge Switch	211

16.9.6	Verify Operations	212
16.10	SPB IP Shortcuts	220
16.10.1	IS-IS Layer 3 configuration	221
16.10.2	ECMP	225
16.10.3	Local VLAN configuration	225
16.10.4	Verify Operations	226
16.11	Multicast over IP Shortcuts	230
16.11.1	IP Shortcuts Multicast configuration	231
16.11.2	Enable IP Multicast at VLAN level	231
16.12	Verify Operations	232
16.12.1	Global Settings	232
16.12.2	Verify IGMP cache/group and senders	233
16.12.3	Verify SPB Multicast Routes	235
16.12.4	Verify multicast TLV's	236
16.12.5	Trace Multicast Routes	238
16.13	IPVPN-Lite L3 VPN over IS-IS	240
16.13.1	SMLT Cluster	242
16.13.2	8007 Configuration	247
16.13.3	Verify Operations	248
17.	Restrictions and Limitations	257
17.1	STP/RSTP/MSTP	257
17.2	SPB IS-IS	257
18.	Reference Documentation	258

Figures

Figure 1: SPBM Service Type Encapsulations	17
Figure 2: SPB L2 VSN	24
Figure 3: SPB L3 VSN	25
Figure 4: Inter VSN Routing.....	26
Figure 5: SPB IP Shortcuts	27
Figure 6: IP VPN Lite over SPB	28
Figure 7 – FI Rear Port Details	46
Figure 8: NNI - Triangle.....	55
Figure 9: NNI - SMLT Triangle A	55
Figure 10: NNI – SMLT Triangle B.....	56
Figure 11: NNI –Square A.....	56
Figure 12: NNI – Square B.....	56
Figure 13: NNI – SLT Square.....	56
Figure 14: NNI – SMLT Square	57
Figure 15: NNI – Full Mesh A.....	57
Figure 16: NNI – Full Mesh B.....	57
Figure 17: NNI – SMLT Full Mesh A	57
Figure 18: NNI – SMLT Full Mesh B	58

Tables

Table 1: IEEE Standards culminating with SPBM	15
---	----

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 clear 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. 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).

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 (I-SID) with a length of 24 bits. This I-SID can be used to identify any virtualized traffic across an 802.1ah encapsulated frame. In 802.1ah, these I-SIDs 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 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	2011	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 I-SID 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 I-SID 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 (I-SID) by local provisioning. Any BEB that has the same I-SID 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 (I-SID) by local provisioning. Any BEB that has the same I-SID 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 to 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 fault verification. (LBM)
- ▶ Linktrace messaging to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for 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 I-SIDs 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 I-SID number with an edge VLAN; creating an L3VSN is as simple as associating an I-SID number with a VRF and configuring the desired IS-IS IP route redistribution within the newly created L3VSN.

- ▶ Multicast

Multicast over SPB is supported on the ERS 8800 beginning in the 7.2 release, on the VSP 4000 in the 3.1 release, and the VSP 9000 in the 3.4 release. 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 additional 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 I-SIDs where the I-SID is simply assigned at the BEB to either a VLAN for virtualized layer 2 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.

2.5 B-MAC

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 SYS-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 (I-SID) 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 (I-SID) specific multicast SPTs based on which BEBs are configured with the corresponding I-SID. 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 I-SID 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>
```

SPBM UNICAST FIB ENTRY INFO					
DESTINATION	BVLAN	SYSID	HOST-NAME	OUTGOING	COST
ADDRESS	INTERFACE				
00:be:b0:00:00:02	40	00be.b000.0002	ERS-2	2/2	10
00:be:b0:00:00:02	41	00be.b000.0002	ERS-2	2/2	10
00:be:b0:00:00:30	40	00be.b000.0030	ERS-3	2/2	20
00:be:b1:00:03:04	40	00be.b000.0030	ERS-3	2/2	20
00:be:b0:00:00:30	41	00be.b000.0030	ERS-3	2/2	20
00:be:b0:00:00:40	40	00be.b000.0040	ERS-4	2/2	20
00:be:b0:00:00:40	41	00be.b000.0040	ERS-4	2/2	20
00: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>**

```
=====
```

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

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 - 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.

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 I-SIDs 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)

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 fault verification. (LBM)
- Linktrace messaging to a maintenance endpoint with intermediate points responding to indicate the path of the traffic within a domain for 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.
 - Hierarchical 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 CLI commands:
 - CLI
 - `lbn <mdName.maName.mepId.rmepMac>`
 - `ltm <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.
 - CLI

- l2ping <vlan.RouterNodeName | vlan.SystemIdMac | ipaddress>
- l2traceroute <vlan.RouterNodeName | vlan.SystemIdMac | ipaddress>
- l2tracetree <vlan.isid | vlan.isid.RouterNodeName | vlan.isid.SystemIdMac>
- 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 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 (I-SID) will be assigned on the BEB to each VLAN. All VLANs in the network that share the same I-SID 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 I-SID'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 (I-SID) will be assigned at a Virtual Router (VRF) level instead of at a VLAN level. All VRFs in the network that share the same I-SID will be able to participate in the same VSN.

In summary:

- One or more VRFs are created on the BEB switches with an assigned I-SID
 - All VRFs that share the same I-SID 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 I-SIDs. 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 can 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, it is 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. To forward the server traffic to the clients and vice-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 I-SID 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
 - o 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
 - o This will inject all direct (or OSPF, RIP, Static, BGP...) IP routes into IS-IS
 - o In a SMLT cluster, in the case of direct IP route redistribution, a route policy or route-map (CLI) 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

3.5 IP VPN Lite over SPB





Figure 6: IP VPN Lite over SPB

By using BGP IPVPNs, it is possible to provide hub and spoke configurations by manipulating the import and export Route Target (RT) values. This allows, for example, a server frame in a central site to have connectivity to all spokes, but, no connectivity between the spoke sites. BGP configuration is only required on the BEB sites where the backbone switches have no knowledge of any Layer 3 VPN IP addresses or routes.

Please note that IP VPN Lite over SPB is only supported on the ERS 8000 platform.


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 in VSP9000, ERS8800, VSP7000, VSP4000, ERS4800

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 I-SID
- Different VLAN-id on same or different ports can be assigned to same I-SID
 - o 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 I-SID
- Supported in VSP7000 release 10.2 and ERS4800 release 5.7

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
 - o VLACP/LACP PDUs are forwarded (VSP4000: unless configured on UNI port / MLT)
 - o 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 VSP7000 release 10.3 and VSP4000 release 3.1
- MLT Transparent UNI ports are supported (on VSP4000 even with LACP)
- Transparent UNIs should not be assigned to the same I-SID as Switched UNI or CVLAN UNIs

4.4 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 I-SIDs from different VLANs. The same VLAN can map to one I-SID on one port and to another I-SID on another.
	1	11	1000	
	1	12	2000	
	2	11	1000	
	2	12	3000	
C-VLAN	All	14	4000	Map entire VLAN to an I-SID. All member ports can send and receive traffic to / from I-SID.
	All	15	1000	
Transparent	5	All	5000	All traffic from the port that creates the transparent UNI goes to a single I-SID, regardless of VLAN.
	10	All	5000	

- I-SID 1000 will receive traffic from: Port 1 on VLAN 10, port 1 & port 2 on VLAN 11 and from all port members of VLAN 15.
- I-SID 2000 will receive traffic from: Port 1 on VLAN 12
- I-SID 3000 will receive traffic from: Port 2 on VLAN 12
- I-SID 4000 will receive traffic from: All member ports for VLAN 14.
- I-SID 5000 will receive traffic from: All traffic from ports port 5 & Port 10

5. Summary of SPB Features and Product Release Matrix

Capability Feature Matrix	ERS 8800	VSP 9000	VSP 8000	VSP 7000	VSP 4000	ERS 4800
L2 VSN	Y	Y	Y	Y	Y	Y
L2 VSN with Multicast (IGMP)	Y	Y	4.1 (1H15)	N		5.8.1 (CY 14)
L3 VSN	Y	Y	4.1 (1H15)	N	Y	N
L3 VSN with Multicast (IGMP)	Y	Y	4.1 (1H15)	N	Y	N
IP Shortcut Routing	Y	Y	Y	N	Y	N
IP Shortcut Routing with Multicast	Y	Y	4.1 (1H15)	N	Y	N
Inter-VSN Routing	Y	Y	Y	N	Y	N
IPVPN-Lite over SPB	Y	N	N	N	N	N
Enterprise Fabric & Switch Cluster Interoperability	Y	Y	Y	Y	4.1 (2H14)	N
Enterprise Fabric & Stackable Chassis Interoperability	N/A	N/A	N/A	Y	N	Y
Enterprise Fabric Connectivity Management (802.1ag)	Y	Y	Y	Y	Y	Y
CFM, L2 Ping, Traceroute, and Tracetree	Y	Y	Y	Y	Y	Y
BCB Mode (NNI-NNI)	Y	Y	Y	Y	Y	N
L2 Ping for Access VLAN (CVLAN)	Y	Y	4.1 (1H15)	10.4 (CY14)	Y	N
Switched UNI	N	5.0	5.0	Y	5.0	Y
Transparent UNI	N	N	N	Y	Y	N
ETREE	N	N	4.1 (1H15)	N	Y	N
vIST	N	N	Y	N	4.1 (2H14)	N

6. SPB Feature and License Matrix

Feature/Platform	ERS 4800	VSP 4000	VSP 7000	ERS 8800	VSP 9000	VSP 8000
L2 VSN	Base	Base	Base	Premier	Premier	Base
L3 VSN	N/A	Premier	N/A	Premier	Premier	TBD
IP-Shortcuts	N/A	Advanced	N/A	Premier	Premier	Base
Multicast L2 VSN	Base	Base	Base	Premier	Premier	TBD
Multicast L3 VSN	TBD	Premier	TBD	Premier	Premier	TBD
Multicast IP Shortcuts	TBD	Advanced	TBD	Premier	Premier	TBD
BCB Mode (NNI-NNI)	N/A	Base	Base	Premier	Premier	Base
Inter-ISID Routing	N/A	Advanced	N/A	Premier	Premier	Base
ETREE	TBD	Base	TBD	TBD	TBD	TBD
Switched UNI	Base	TBD	Base	N/A	TBD	TBD
Transparent UNI	TBD	Base	Base	TBD	Premier	TBD
VRF support	N/A	Premier	N/A	Premier	Premier	Base
Dual homing into a Fabric (SMLT Edge)	TBD	TBD	Base	Premier	Premier	Base

7. 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.

7.1 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.

7.2 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.

7.3 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.

7.4 VSP 4000 and ERS 4850

An ERS 4850 (Rev 10 or higher) can be converted to a VSP 4000 via a software conversion kit; a USB flash drive that contains the VSP 4000 run time software. The software conversion kit comes factory installed when ordering the VSP 4000 or can be included in the upgrade kit for an ERS 4850. Please note that only the ERS 4850 model, either an ERS 4850GTS or ERS 4850GTS-PWR+, can be converted. The VSP 4000 supports all the same SFP and SFP+ transceivers as the ERS 4850, but, does not support stacking. The VSP 4000 operating system is common with the VSP 9000 (release 3.3+) supporting the same CLI, full logging, KHI, and Flight Recorder.

7.5 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



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.

7.6 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
 - o 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.
 - o If the ERS 8800 is running in STG mode, verify that STG-63 is not current in use. SPB will use this STG.
 - o If running in MSTP mode – verify that MSTI-62 is not currently in use. SPB will use this MSTI.
- VSP 9000
 - o Activating SPB infrastructure reserves 100 multicast group ID (MGIDs) for SPB operation on the VSP 9000.
- VSP 9000, VSP 8000, VSP 7000, VSP 4000, ERS 4800, and ERS8800
 - o 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
 - o 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
 - o The IS-IS adjacencies will not come up if there is a discrepancy in the B-VLAN ids configuration between 2 nodes.
 - o The Primary VLAN IDs is also used on all IS-IS messages
- Enable SPB globally.
- Assign a unique nickname to each switch.
 - o An alarm will be logged if a duplicate nickname is provisioned in the network
- Assign a common Area ID
 - o 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.
 - o 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).

- o 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) 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 l2tracertout commands between all the switches in the network for both the B-VLANs.
- Verify that the path reported by the l2tracertout 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
 - o The C-VLAN is not added to the IST
 - o 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

7.7 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, and VSP 9000 release 3.4 add support for L2VSN, L3VSN and IP Shortcut. L2VSN multicast is not supported at this time for the VSP 8000, VSP 7000 and ERS 4800
 - o If a PIM router is connected to an SPB bridge, either use IGMP or static mroutes

The following traffic is not yet supported by SPB.

- IPv6 routed traffic (unicast or multicast) unless forward via L2VSN


Traffic which is not yet supported by SPB can continue to exist in parallel to SPB. For example, IPv6 traffic can be routed by OSPFv3 which is configured on SMLT VLANs which can remain independent from SPB running on the same physical infrastructure.

7.8 Multicast

Sender switch sends either TLV 185 (IP Shortcuts) or TLV 186 (L2 and L3 VSN) with Tx bit set and TLV 144 with Tx bit set

Source starts sending Multicast traffic for a particular group

Receiver sends IGMP host Membership report for group G (Join)



Multicast over SPB is supported in the 7.2 release for the ERS 8800, 3.1 for the VSP 4000, and the 3.4 release for the VSP 9000. 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
 - o Simple provision by enabling SPB multicast globally and IGMP snooping at L2VSN VLAN level
 - o Traffic does not cross L2VSN service boundary
 - o 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
 - o 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
 - o 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
 - o 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
 - o All or a subset of VLANs within a L3VSN can exchange IP multicast traffic between themselves
 - o Simply provision by enabling SPB multicast globally, enable MVPN on the VRF, and enable IP SPB Multicast on some or all the VLANs within the L3VSN

- Only those VLANs that have IP SPB Multicast enabled can pass multicast traffic
- o It is not a requirement to enable IP Shortcuts to support IP Multicast in the L3VSN
- o IPVPN creation and I-SID assignment for the IPVPN is required – but the IPVPN does not need to be enabled
- o 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
- o SMLT operation – does not apply to the VSP 4000 or ERS 4800
 - 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
 - o All or a subset of VLANs within GRT can exchange IP multicast traffic between themselves
 - o Simple provision by enabling SPB multicast globally and enabling IP SPB Multicast on some or all the GRT VLANs
 - Only those VLANs that have IP SPB Multicast enabled can pass multicast traffic
 - o It is not a requirement to enable IP Shortcuts to support IP multicast in the GRT using SPB
 - o 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

7.9 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> i-sid <isid value>`” CLI command. Use the same value of i-sid on all the switches. This step should restore service.
- SMLT deployment
 - o 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
 - o 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 I-SID is assigned to a VLAN.
 - o For a vIST, for the L2VSN C-VLAN, only the local SMLT ports are added. A separate L2VSN is used for the vIST VLAN.

7.9.1 Migrating to Inter VSN Routing

Inter VSN provides the ability to route traffic between extended VLANs where the VLANs have different I-SIDs. 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.

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

- 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.

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
 - o On all the switches where the VLAN is configured - remove all NNI ports
 - This will cause service interruption.
 - o On all the switches where the VLAN is configured – map the VLAN to an ISID. This will restore L2 connectivity (the `I2tracetree` 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.

- o 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.

7.10VSP 9000 Notes

- VSP 9000 supports SPB NNI Interfaces on the 9024XL, 9048XS, and 9012QQ cards.
- For L2VSN services on an IST switch
 - o 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).
 - o 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
 - o You will see the messages below during configuration.
 - CAUTION: Adding I-SID 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 I-SIDs 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

8. Field Introduction & Support Specifications

8.1 Hardware and Deployment Specifications

SPB is supported on ERS 8600/8800 family of switches that have the following hardware.

- Line Cards – R, RS, and 88xx modules
- CP – 8692 with Supermezz or 8895

There are no other special considerations for hardware other than the overall requirements that apply all the features in ERS7.0 and ERS 7.1 releases.

SPB is supported on the VSP 9000 where only the 9024XL supports SPB NNI Interfaces. Any of the other VSP 9000 modules can be used as UNI ports.

In regard to the ERS 4800 and VSP 4000, please see section 7.4 above.

8.2 Installation and Commissioning Specifications

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

8.3 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. In release 7.2 for the ERS 8800, only the new 802.1aq standard TLVs are supported.

9. 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 will be 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*.

10. 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



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.


11. SPB SMLT BEB Design Best Practices

11.1 SMLT BEB – C-VLAN Guidelines for L2VSN



- ▶ Customer VLAN (CVLAN) has I-SID 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)



11.2 SMLT BEB – ISIS Hello Timer Guidelines



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


- ▶ 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
 - $1 \times 27 = 9 \times 3 = 27$

11.3 SMLT BEB – RSMLT



- ▶ Both RSMLT and RSMLT Edge is supported providing the SMLT cluster is either a VSP 9000 or ERS 8800 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. At this time, ISIS route policies are not supported

11.4 SMLT BEB – VLACP Guidelines



Connection Type	Fast Timer	Slow Timer	Timeout	Timeout Scale	ERS 8800 VSP 9000	VSP 7000
IST – NNI ISIS	N/A	10000	Long	3	√	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

11.5 SMLT BEB – VSP 7000 Guidelines

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

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 CLI command `show vlan interface info <port list>`. To manually set the default PVID on the IST ports, use the CLI command `vlan ports <port list> pvid <1-4096>`.


Also, it is recommended to not enable VLACP on the IST.

11.6 SMLT 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, eliminates the single point of failure with a dedicated MLT. The vIST always up if there is SPBM connectivity between the vIST peers. Also, the vIST devices do not have to be the same type.

Like a traditional IST, the vIST still requires an IST VLAN with an IP address known as an IST VLAN. The difference is, the traditional IST requires that the IST VLAN is a member of the IST MLT whereas the vIST uses IST tunnel across an SPBM cloud where this tunnel is identified with an I-SID; the IST VLAN is assigned to an I-SID for identification. Both methods still require that you peer with the SPBM System-id of the peer node and both methods also require a SMLT virtual BMAC. Each C-VLAN, no matter the service type, will also require an C-VLAN to I-SID mapping for identification when using an vIST.

11.7 SLPP Guard – ERS 4800



- On the ERS 4800 only, SLPP can be enabled on the core bridges and in turn SLPP Guard can be enabled on the ERS 4800 for local port loop detection
 - The setting of the overload bit on the ERS 4800 allows it to operate as a stub node on the SPB network
 - This prevents traffic from one NNI port to be forwarded to another NNI port
 - 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

12. 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 VSP 4000 or ERS 4800 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 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 a 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.




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 could enable IS-IS on the MLT (single port), but, this would not be recommended




Figure 14: NNI – SMLT Square

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 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.



IS-IS should only be configured on one of the links between nodes B and D.

Figure 18: NNI – SMLT Full Mesh B

13. 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 itself
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
143	SPBM Instance & BVIDs	Sub-TLV 6: BVIDs to ECT algorithm Used in IS-IS Hellos only
144	SPBM Instance, Nick-name, BVLANS & I-SIDs	Sub-TLV 1: SPBM Instance & Nick-name Sub-TLV 3: B-VLANS & L2VSN I-SIDs
184	SPBM IPVPN Reachability	IP Reachability for L3 VSNs
185	SPBM I-SID Constrained Source-Groups	IP Multicast stream availability for L2VSNs & L3VSNs
186	SPBM VRF-0/GRT Source-Groups	IP Multicast stream availability for GRT/VRF-0

TLVs 1,3,22,129 & 135 are well known IS-IS TLVs which existed 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

14. SPB Best Practices

The following are best practices when setting up SPB.

IS-IS

- Recommended to change the IS-IS SYS-ID (B-MAC) with an easy to recognize address to easily identify a switch. This will help with troubleshooting to easily recognize source and destination addresses
 - If you leave the SYS-ID with its default value, safe practice as it ensures no duplication in the network, it may be difficult to recognize the source and destination B-MAC for troubleshooting purposes
 - If you do change manually the SYS-ID, please take the necessary steps to ensure there is no duplication in the network
- 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 I-SID can be assigned to a VRF for L3 VSN. This does not restrict another I-SID 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 – not 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 be 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.

Configuration

It is recommended to follow the SPB best practices as specified in the *SPB Deployment Considerations* document which can be found at <http://tools.ietf.org/html/draft-lapuh-spb-deployment-01>. The following is an example of recommended values:

- Spbm-id : 1
- BVID #1 and BVID #2 : 4051, 4052
- Nick-name : 0.01.<node-id>
- MEP-id : md.ma.<node-id>
- MD : spbm with level 4
- MA : 4051 & 4052
- MEP : <node-id>

-
- MIP : level 4
 - IS-IS Manual Area : 49.0000

15. SPB Configuration

For compatibility between the VSP 9000, VSP 4000, ERS 8800, ERS 4800, and VSP 7000, it is recommended to change the Spanning Tree mode to MSTP. By default, both the VSP 4000 and VSP 9000 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

CLI

- 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) to re-load the existing configuration file in MSTP mode.

It is recommended to follow the SPB best practices as recommended in the *SPB Deployment Considerations* document which can be found by going to <http://www.ietf.org/id/draft-lapuh-spb-deployment-01.txt>. In summary:

- It is recommended to define a global virtualization schema based on I-SIDs, and not tie VLAN ids directly to ISIDs ids in a 1 to 1 relationship throughout the network
- It is a good practice to manually configure System IDs and SPB Nicknames with a simple identification scheme coordinating the system ID numerically with the SPB Nickname for ease of troubleshooting
 - o For example, System IDs start with 0049.bb00.1000 for the first node, 0049.bb00.2000 for the second node and so on
 - 49 indicates a private address
 - the "00bb" indicates area "00bb"; 1000, 2000, etc., indicate the node number (1 through n)
 - These System IDs correspond to SPB Nicknames of 1.bb.10, 1.bb.20, 1.bb.30 for nodes 1, 2 and 3 respectively
 - o As an alternative, the System-ID could be constructed in such a way to identify node location such as the following:
 - 000z.0xxx.vyy0
 - Z = Core, Distribution, or Edge
 - Location xxx (000-FFF)
 - Bridge Mode v (4 = VSP 4000, 7 = VSP 7000, 8 = ERS 8800, 9 = VSP 9000)
 - Node identifier yy (00-FF)

- The 802.1aq standard defines up to 16 BVIDs where these BVIDs must be consistent across the SPB region



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.

15.1 SPB Configuration

15.1.1 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
system-id 0049.0090.0100
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
system-id 0049.0040.0100
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

```
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...xxx - 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 and ERS 4800, 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.

15.1.2 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
```

```
isis spbm 1
    isis enable
exit
interface mlt <mlt id>
    isis
    isis spbm 1
    isis enable
exit
```

ERS 4800 as of release 5.8 and VSP 7000 as of release 10.3:

```
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, 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.

15.1.3 VSP 7000 – Fabric Interconnect Mesh

15.1.3.1 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 is supported via the rear ports.

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

CLI - 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
```

15.1.3.2 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
```

15.1.4 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> <i-sid: 0..16000000>
```



Although you can use any number from 1 to 16,777,215 as an I-SID 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.

15.1.5 Inter-ISID Routing

L2 VSN



```
config terminal
vlan create <vlan id> type port-mstprstp 0
vlan i-sid <vlan-id> <i-sid: 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 I-SID 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-mstp 0
vlan i-sid <vlan-id> <i-sid: 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
```

15.1.6 L3VSN Configuration

L3 VSN with direct interface redistribution



```

config terminal
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
    i-sid <i-sid: 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 I-SID 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 shows direct interface redistribution into ISIS, other protocols such as BGP, OSPF, RIP, and Static can also be enabled.

15.1.7 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
vrf blue
ip address 10.198.55.1 255.255.255.0
exit
router isis
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
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
vrf blue
ip address 10.198.33.1 255.255.255.0
exit
router isis
isis redistribute direct
isis redistribute direct enable
exit
isis apply redistribute direct vrf blue
vlan member remove 1 1/20

```

```

config terminal
interface loopback <1-256>
    ip address a.b.c.d mask
exit

```

```
router isis
  ip-source-address <loopback ip>
  spbm 1 ip enable
exit
ip ecmp
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 enabled.

15.1.8 SPB Multicast Configuration

15.1.8.1 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
```



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.

15.1.8.2 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
```

15.1.8.3 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
```

15.1.9 SMLT – Normal IST

Enabling IST

```
config terminal
mlt y enable name IST
mlt y member slot/port-slot/port
interface mlt y
    ist peer-ip <ip address of peer> vlan x
    ist enable
exit
vlan create x type port-mstprstp 0
vlan mlt x y
vlan members x slot/port-slot/port
interface vlan x
    ip address <ip address> <ip mask>
exit
router isis
    spbm 1 smlt-virtual-bmac xx:xx:xx:xx:xx:xx
    spbm 1 smlt-peer-system-id <xxxx.xxxx.xxxx - system id of peer>
exit
```

15.1.10 Virtual IST

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. This feature will be made for available, for example, on the VSP 8250 in the initial release.

Enabling vIST

```
config terminal
router isis
    spbm 1 spbm 1 smlt-peer-system-id <xxxx.xxxx.xxxx - system id of peer>
    spbm 1 spbm 1 smlt-virtual-bmac xxxx.xxxx.xxxx
vlan create x type port-mstprstp 0
vlan i-sid x <i-isid number>
interface vlan x
    ip address <ip address>/<mask>
    virtual-ist peer-ip <ip address of peer> vlan x
exit
```

15.1.11 Connectivity Fault Management (CFM) Configuration

15.1.11.1 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 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
```

15.1.11.2 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 or 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
```

15.1.12CFM Configuration Example – 7.1.1.x or higher



- This ensures that we get full OAM functionalities across:
 - o SPB -> Backbone VLAN-ids (BVIDs) i.e. Infrastructure
 - o 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



15.2 Using EDM

15.2.1 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




15.2.2 VSN Configuration

Extending a VLAN (L2VSN)

EDM


Configuration -> VLAN -> VLANs -> Advanced -> Select VLAN -> I-sid = <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




15.2.3 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.





16. Configuration Examples

16.1 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 & 9001	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



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.

16.1.1 Configuration

For this configuration example, all other switches are provisioned using CLI which is the default setting on all switches except for the ERS 8000.

16.1.1.1 Auto Save

On the VSP 7000 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 you make a change.

CLI

```
no autosave enable
```

To save the configuration, use either of the two commands

```
write memory
```

```
save configuration
```

16.1.1.2 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)

16.1.1.3 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

started configuring the switch in RSTP or MSTP mode.
The syntax used to create VLANs in any of these new
modes are NOT COMPATIBLE with the default mode (STP)

```
ERS-8606:5(config)#save boot
```

```
ERS-8606:5(config)#boot -y
```

16.1.1.4 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
```

16.1.1.5 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
```

16.1.1.6 Enable VLACP Globlaly

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
```

16.1.1.7 IST Configuration – SMLT Cluster switch 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



We will configure the vIST on the VSP 4000 latter one and the IST on VSP 7000 after we have provisioned SPBM.

VSP 9000 SMLT Cluster: Add IST VLAN 2 and add IP address

```
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 VLAN 2 and add IP address

```
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
```

16.1.1.8 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
	SPB System-ID	0049.0040.0100
	vIST Configuration	
	vIST VLAN	2
	IP address	10.4.2.1/30
	SMLT Peer System ID	0049.0040.0200
	SMLT Virtual BMAC	0049.0040.01ff
	I-SID	2002
	vIST peer	10.4.2.2
4002	SPB Nick Name	0.40.02
	SPB System-Name	4002
	SPB System-ID	0049.0040.0200
	vIST Configuration	
	vIST VLAN	2
	IP address	10.4.2.2/30

	SMLT Peer System ID	0049.0040.0100
	SMLT Virtual BMAC	0049.0040.01ff
	I-SID	2002
	vIST peer	10.4.2.1
4801	SPB Nick Name	0.48.01
	SPB System-Name	4801
	SPB System-ID	0049.0048.0100
7001	SPB Nick Name	0.70.01
	SPB System-Name	7001
	SPB System-ID	0049.0070.0100
7002	SPB Nick Name	0.70.02
	SPB System-Name	7002
	SPB System-ID	0049.0070.0200
7003	SPB Nick Name	0.70.03
	SPB System-Name	7003
	SPB System-ID	0049.0070.0300
7004	SPB Nick Name	0.70.04
	SPB System-Name	7004
	SPB System-ID	0049.0070.0400
8003	SPB Nick Name	0.80.03
	SPB System-Name	8003
	SPB System-ID	0049.0080.0300
8004	SPB Nick Name	0.80.04
	SPB System-Name	8004
	SPB System-ID	0049.0080.0400

8005	SPB Nick Name	0.80.05
	SPB System-Name	8005
	SPB System-ID	0049.0080.0500
	SMLT Virtual BMAC	00:49:00:08:05:ff
	SMLT Peer System-ID	0049.0080.0600
8006	SPB Nick Name	0.80.06
	SPB System-Name	8006
	SPB System-ID	0049.0080.0600
	SMLT Virtual BMAC	00:49:00:08:05:ff
	SMLT Peer System-ID	0049.0080.0500
8007	SPB Nick Name	0.80.07
	SPB System-Name	8007
	SPB System-ID	0049.0080.0700
9001	SPB Nick Name	0.90.01
	SPB System-Name	9001
	SPB System-ID	0049.0090.0100
	SMLT Virtual BMAC	00:49:00:90:01:ff
	SMLT Peer System-ID	0049.0090.0200
9002	SPB Nick Name	0.90.01
	SPB System-Name	9002
	SPB System-ID	0049.0090.0200
	SMLT Virtual BMAC	00:49:00:90:01:ff
	SMLT Peer System-ID	0049.0090.0100



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 if SPB is enabled prior to enabling the IST.

SPBM Configuration – VSP 4000

```
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-virtual-bmac 00:49:00:40:01:ff
4001:1(config-isis)#spbm 1 smlt-peer-system-id 0049.0040.0200
4001:1(config-isis)#system-id 0049.0040.0100
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 i-sid 2 2002
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
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)#system-id 0049.0040.0200
4002:1(config-isis)#spbm 1 smlt-peer-system-id 0049.0040.0100
4002:1(config-isis)#exit
4002:1(config)#interface vlan 2
4002:1(config-if)#interface ip address 10.4.2.2 255.255.255.252
4002:1(config-if)#exit
4002:1(config)#virtual-ist peer-ip 10.4.2.1 vlan 2
```

SPBM Configuration – VSP 7000

```
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 nick-name 0.70.01
7001(config-isis)#manual-area 49.0001
7001(config-isis)#system-id 0049.0070.0100
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)#system-id 0049.0070.0200
```

```
7003(config-isis)#spbm 1 nick-name 0.70.03
7003(config-isis)#system-id 0049.0070.0300
```

```
7004(config-isis)#spbm 1 nick-name 0.70.04
7004(config-isis)#system-id 0049.0070.0400
```

SPBM Configuration – ERS 8800

```
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)#system-id 0049.0080.0300
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
8004:5(config-isis)#system-id 0049.0080.0400
-----
8005:5(config-isis)#spbm 1 nick-name 0.80.05
8005:5(config-isis)#system-id 0049.0080.0500
8005:5(config-isis)#spbm 1 smlt-virtual-bmac 00:49:00:80:05:ff
8005:5(config-isis)#spbm 1 smlt-peer-system-id 0049.0080.0600
-----
8006:5(config-isis)#spbm 1 nick-name 0.80.06
8006:5(config-isis)#system-id 0049.0080.0600
8006:5(config-isis)#spbm 1 smlt-virtual-bmac 00:49:00:80:05:ff
8006:5(config-isis)#spbm 1 smlt-peer-system-id 0049.0080.0500
-----
8007:5(config-isis)#spbm 1 nick-name 0.80.07
8007:5(config-isis)#system-id 0049.0080.0700
```

SPBM Configuration – VSP 9000

```
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)#system-id 0049.0090.0100
9001:1(config-isis)#manual-area 49.0001
9001:1(config-isis)#spbm 1 smlt-virtual-bmac 00:49:00:90:01:ff
9001:1(config-isis)#spbm 1 smlt-peer-system-id 0049.0090.0200
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)#system-id 0049.0090.0200
9002:1(config-isis)#spbm 1 smlt-virtual-bmac 00:49:00:90:01:ff
9002:1(config-isis)#spbm 1 smlt-peer-system-id 0049.0090.0100
```

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)#system-id 0049.0048.0100
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.

16.1.1.9 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. Also, as MLT 1 is used for the IST, we will change the ISIS hello interval to 1 and hello multiplier to 27.

```
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)#isis l1-hello-interval 1
9001:1(config-mlt)#isis l1-hello-multiplier 27
9001:1(config-mlt)#exit
9001:1(config)#interface gigabitEthernet 3/3
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
```

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
```

16.1.1.10 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
```

16.1.1.11 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
```

```
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
```

16.1.1.12 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 happen 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
```

16.1.1.13 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
```

16.1.1.14 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)	807

	ID	
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
```

16.1.1.15 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, 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, these 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* paramter

VSP 4000, 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>
```

16.1.2 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

16.1.2.1 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




The screenshot shows the EDM configuration interface for a device named 8005 (vrf 0). The left sidebar displays a tree view of the configuration hierarchy, with 'IS-IS' selected under 'Configuration'. The main panel shows the 'IS-IS' configuration page with tabs for 'Protocol Summary', 'Globals', 'System Level', 'Interfaces', and 'Interfaces Le'. The 'Globals' tab is active, showing the following configuration parameters:

Parameter	Value	Range/Unit
AdminState:	☒ on ☐ off	
LevelType:	☒ level1 ☐ level1and2	
SystemId:	0049.0008.0500	
MaxLSPGenInt:	900	0..900 (secs)
CsnpInt:	10	1..600 (secs)
RxmtLspInt:	5	1..300 (secs)
PSNPInterval:	2	1..120 (secs)
SpfDelay:	100	0..5000 (1/1000 secs)
HostName:	8005	

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




16.1.2.2 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 be 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 4 – 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



16.1.3 Verify Operations

16.1.3.1 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 *0049.0080.0300* should be displayed.

8003:

```
=====
                        ISIS General Info
=====
                        AdminState : enabled
                        RouterType : Level 1
                        System ID : 0049.0080.0300
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
-----
49.0001.0049.0040.0100.00
```

4002:

```
=====
                        ISIS Net Info
=====
NET
-----
49.0001.0049.0040.0200.00
```

7001:

```
=====
                        ISIS Net Info
=====
NET
-----
49.0001.0049.0070.0100.00
```

7002:

```
=====
                        ISIS Net Info
```

=====

NET

49.0001.0049.0070.0200.00

7003:

=====

ISIS Net Info

=====

NET

49.0001.0049.0070.0300.00

7004:

=====

ISIS Net Info

=====

NET

49.0001.0049.0070.0400.00

9001:

=====

ISIS Net Info

=====

NET

49.0001.0049.0090.0100.00

9002:

=====

ISIS Net Info

```
=====
NET
-----
49.0001.0049.0090.0200.00
```

8003:

```
=====
ISIS Net Info
=====
NET
-----
49.0001.0049.0080.0300.00
```

8004:

```
=====
ISIS Net Info
=====
NET
-----
49.0001.0049.0080.0400.00
```

8005:

```
=====
ISIS Net Info
=====
NET
-----
49.0001.0049.0080.0500.00
```

8006:

```
=====
                        ISIS Net Info
=====
```

```
=====
NET
-----
49.0001.0049.0080.0600.00
```

8007:

```
=====
                        ISIS Net Info
=====
```

```
=====
NET
-----
49.0001.0049.0080.0700.00
```

4801:

```
=====
                        ISIS Net Info
=====
```

```
=====
NET
-----
49.0001.0049.0048.0100.00
```

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

16.1.3.2 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  HOLDDTIME SYSID      HOST-NAME
-----
Mlt1      1  UP          21:01:51 127      22 0049.0090.0100  9001
Port1/47  1  UP          1d 19:57:51 127      23 0049.0090.0200  9002

-----
2 out of 2 interfaces have formed an adjacency
-----
```

7001:

```
=====
                        ISIS Adjacencies
=====
INTERFACE L STATE      UPTIME PRI  HOLDDTIME SYSID      HOST-NAME
=====
Trunk: 64 1  UP        28d 13:12:00 127      17 0049.0070.0300  7003
Port: 3   1  UP         20:25:01 127      18 0049.0080.0500  8005
Trunk: 63 1  UP        14d 19:15:40 127      18 0049.0070.0100  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

16.1.3.3 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          00:49:00:90:01:ff    0049.0090.0200
```

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          00:49:00:08:05:ff    0049.0080.0600
```

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 4001 and 8005

4001:

```
=====
```

SPBM UNICAST FIB ENTRY INFO					
=====					
DESTINATION	BVLAN	SYSID	HOST-NAME	OUTGOING	COST
ADDRESS				INTERFACE	

00:49:00:08:03:00	4051	0049.0080.0300	8003	8001	20
00:49:00:08:03:00	4052	0049.0080.0300	8003	8001	20
00:49:00:08:04:00	4051	0049.0080.0400	8004	1/47	20
00:49:00:08:04:00	4052	0049.0080.0400	8004	1/47	20
00:49:00:08:05:00	4051	0049.0080.0500	8005	8001	30
00:49:00:08:05:ff	4051	0049.0080.0500	8005	8001	30
00:49:00:08:05:00	4052	0049.0080.0500	8005	8001	30
00:49:00:08:05:ff	4052	0049.0080.0500	8005	8001	30
00:49:00:08:05:ff	4051	0049.0080.0600	8006	1/47	30
00:49:00:08:06:00	4051	0049.0080.0600	8006	1/47	30
00:49:00:08:05:ff	4052	0049.0080.0600	8006	1/47	30
00:49:00:08:06:00	4052	0049.0080.0600	8006	1/47	30
00:49:00:08:07:00	4051	0049.0080.0700	8007	8001	30
00:49:00:08:07:00	4052	0049.0080.0700	8007	1/47	30
00:49:00:40:01:00	4051	0049.0040.0100	4001	cpp	0

00:49:00:40:01:00	4052	0049.0040.0100	4001	cpp	0
00:49:00:40:02:00	4051	0049.0040.0200	4002	8001	20
00:49:00:40:02:00	4052	0049.0040.0200	4002	1/47	20
00:49:00:70:01:00	4051	0049.0070.0100	7001	8001	50
00:49:00:70:01:00	4052	0049.0070.0100	7001	8001	50
00:49:00:70:02:00	4051	0049.0070.0200	7002	8001	40
00:49:00:70:02:00	4052	0049.0070.0200	7002	8001	40
00:49:00:70:03:00	4051	0049.0070.0300	7003	1/47	40
00:49:00:70:03:00	4052	0049.0070.0300	7003	1/47	40
00:49:00:70:04:00	4051	0049.0070.0400	7004	1/47	50
00:49:00:70:04:00	4052	0049.0070.0400	7004	1/47	50
00:49:00:90:01:00	4051	0049.0090.0100	9001	8001	10
00:49:00:90:01:ff	4051	0049.0090.0100	9001	8001	10
00:49:00:90:01:00	4052	0049.0090.0100	9001	8001	10
00:49:00:90:01:ff	4052	0049.0090.0100	9001	8001	10
00:49:00:90:01:ff	4051	0049.0090.0200	9002	1/47	10
00:49:00:90:02:00	4051	0049.0090.0200	9002	1/47	10
00:49:00:90:01:ff	4052	0049.0090.0200	9002	1/47	10
00:49:00:90:02:00	4052	0049.0090.0200	9002	1/47	10

 Total number of SPBM UNICAST FIB entries 34

8007:

=====					
SPBM UNICAST FIB ENTRY INFO					
=====					
DESTINATION ADDRESS	BVLAN	SYSID	HOST-NAME	OUTGOING INTERFACE	COST

00:49:00:08:03:00	4051	0049.0080.0300	8003	3/27	10
00:49:00:08:03:00	4052	0049.0080.0300	8003	3/27	10
00:49:00:08:04:00	4051	0049.0080.0400	8004	3/28	10
00:49:00:08:04:00	4052	0049.0080.0400	8004	3/28	10
00:49:00:08:05:00	4051	0049.0080.0500	8005	3/27	20
00:49:00:08:05:ff	4051	0049.0080.0500	8005	3/27	20
00:49:00:08:05:00	4052	0049.0080.0500	8005	3/27	20
00:49:00:08:05:ff	4052	0049.0080.0500	8005	3/27	20
00:49:00:08:05:ff	4051	0049.0080.0600	8006	3/28	20
00:49:00:08:06:00	4051	0049.0080.0600	8006	3/28	20
00:49:00:08:05:ff	4052	0049.0080.0600	8006	3/28	20
00:49:00:08:06:00	4052	0049.0080.0600	8006	3/28	20
00:49:00:08:07:00	4051	0049.0080.0700	8007	cpp	0

```

00:49:00:08:07:00 4052 0049.0080.0700 8007          cpp          0
00:49:00:40:01:00 4051 0049.0040.0100 4001          3/27         30
00:49:00:40:01:00 4052 0049.0040.0100 4001          3/28         30
00:49:00:40:02:00 4051 0049.0040.0200 4002          3/27         30
00:49:00:40:02:00 4052 0049.0040.0200 4002          3/28         30
00:49:00:70:01:00 4051 0049.0070.0100 7001          3/27         40
00:49:00:70:01:00 4052 0049.0070.0100 7001          3/27         40
00:49:00:70:02:00 4051 0049.0070.0200 7002          3/27         30
00:49:00:70:02:00 4052 0049.0070.0200 7002          3/27         30
00:49:00:70:03:00 4051 0049.0070.0300 7003          3/28         30
00:49:00:70:03:00 4052 0049.0070.0300 7003          3/28         30
00:49:00:70:04:00 4051 0049.0070.0400 7004          3/28         40
00:49:00:70:04:00 4052 0049.0070.0400 7004          3/28         40
00:49:00:90:01:00 4051 0049.0090.0100 9001          3/27         20
00:49:00:90:01:ff 4051 0049.0090.0100 9001          3/27         20
00:49:00:90:01:00 4052 0049.0090.0100 9001          3/27         20
00:49:00:90:01:ff 4052 0049.0090.0100 9001          3/27         20
00:49:00:90:01:ff 4051 0049.0090.0200 9002          3/28         20
00:49:00:90:02:00 4051 0049.0090.0200 9002          3/28         20
00:49:00:90:01:ff 4052 0049.0090.0200 9002          3/28         20
00:49:00:90:02:00 4052 0049.0090.0200 9002          3/28         20

```

Total number of SPBM UNICAST FIB entries 34

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.

16.1.3.4 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

0049.0080.0300.00-00	1	748	0x1545	0xbc63	8003
0049.0080.0400.00-00	1	1068	0x1191	0xbc2	8004
0049.0080.0500.00-00	1	978	0xae9	0x8e34	8005
0049.0080.0600.00-00	1	930	0x1555	0x90c3	8006
0049.0080.0700.00-00	1	1090	0x23c	0x1dcb	8007
0049.0040.0100.00-00	1	1175	0x144a	0xabc	4001
0049.0040.0100.00-01	1	1175	0x11d9	0x7421	4001
0049.0040.0100.00-02	1	1175	0x1085	0x9890	4001
0049.0040.0100.00-03	1	1175	0xafd	0xa41	4001
0049.0040.0200.00-00	1	355	0x11f8	0xfb1c	4002
0049.0040.0200.00-01	1	355	0x11cd	0xb1ec	4002
0049.0040.0200.00-02	1	355	0x1073	0xb77f	4002
0049.0040.0200.00-03	1	355	0xaf6	0x3619	4002
0049.0070.0100.00-00	1	945	0x76c	0x68d7	7001
0049.0070.0200.00-00	1	1072	0x1d1f	0xfaff	7002
0049.0070.0300.00-00	1	1168	0x1d1c	0x89ae	7003
0049.0070.0400.00-00	1	335	0x1d1b	0xf8a3	7004
0049.0090.0100.00-00	1	647	0xc1	0x3177	9001
0049.0090.0200.00-00	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
0049.0080.0300.00-00	1	832	0x1545	0xbc63	8003
0049.0080.0400.00-00	1	1142	0x1191	0xbc2	8004
0049.0080.0500.00-00	1	1056	0xae9	0x8e34	8005
0049.0080.0600.00-00	1	1009	0x1555	0x90c3	8006
0049.0080.0700.00-00	1	1165	0x23c	0x1dcb	8007
0049.0040.0100.00-00	1	371	0x1449	0xcbb	4001
0049.0040.0100.00-01	1	371	0x11d8	0x7620	4001
0049.0040.0100.00-02	1	371	0x1084	0x9a8f	4001
0049.0040.0100.00-03	1	371	0xafc	0xc40	4001
0049.0040.0200.00-00	1	451	0x11f8	0xfb1c	4002
0049.0040.0200.00-01	1	451	0x11cd	0xb1ec	4002
0049.0040.0200.00-02	1	451	0x1073	0xb77f	4002
0049.0040.0200.00-03	1	451	0xaf6	0x3619	4002
0049.0070.0100.00-00	1	1024	0x76c	0x68d7	7001
0049.0070.0200.00-00	1	1147	0x1d1f	0xfaff	7002
0049.0070.0300.00-00	1	357	0x1d1b	0x8bad	7003
0049.0070.0400.00-00	1	431	0x1d1b	0xf8a3	7004
0049.0090.0100.00-00	1	736	0xc1	0x3177	9001
0049.0090.0200.00-00	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

16.1.3.5 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(SPB 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: 0049.0080.0300.00-00      SeqNum: 0x00001546      Lifetime:   604
      Chksum: 0xba64  PDU Length: 225
      Host_name: 8003
      Attributes:      IS-Type 1
TLV:137 Host_name: 8003


Level-1 LspID: 0049.0080.0400.00-00      SeqNum: 0x00001192      Lifetime:   924
      Chksum: 0x9c3   PDU Length: 173
```

Host_name: 8004
Attributes: IS-Type 1
TLV:137 Host_name: 8004

Level-1 LspID: 0049.0080.0500.00-00 SeqNum: 0x00000aea Lifetime: 834
Chksum: 0x8c35 PDU Length: 841
Host_name: 8005
Attributes: IS-Type 1
TLV:137 Host_name: 8005

Level-1 LspID: 0049.0080.0600.00-00 SeqNum: 0x00001556 Lifetime: 788
Chksum: 0x8ec4 PDU Length: 818
Host_name: 8006
Attributes: IS-Type 1
TLV:137 Host_name: 8006

Level-1 LspID: 0049.0080.0700.00-00 SeqNum: 0x0000023d Lifetime: 949
Chksum: 0x1bcc PDU Length: 466
Host_name: 8007
Attributes: IS-Type 1
TLV:137 Host_name: 8007

Level-1 LspID: 0049.0040.0100.00-00 SeqNum: 0x0000144b Lifetime: 1005
Chksum: 0x8bd PDU Length: 124
Host_name: 4001
Attributes: IS-Type 1
TLV:137 Host_name: 4001

Level-1 LspID: 0049.0040.0200.00-00 SeqNum: 0x000011fa Lifetime: 1085
Chksum: 0xf71e PDU Length: 124
Host_name: 4002
Attributes: IS-Type 1
TLV:137 Host_name: 4002

Level-1 LspID: 0049.0070.0100.00-00 SeqNum: 0x0000076d Lifetime: 784
Chksum: 0x66d8 PDU Length: 124
Host_name: 7001
Attributes: IS-Type 1
TLV:137 Host_name: 7001

Level-1 LspID: 0049.0070.0200.00-00 SeqNum: 0x00001d20 Lifetime: 911
Chksum: 0xf801 PDU Length: 195
Host_name: 7002
Attributes: IS-Type 1
TLV:137 Host_name: 7002

Level-1 LspID: 0049.0070.0300.00-00 SeqNum: 0x00001d1d Lifetime: 1007
Chksum: 0x87af PDU Length: 179
Host_name: 7003
Attributes: IS-Type 1
TLV:137 Host_name: 7003

Level-1 LspID: 0049.0070.0400.00-00 SeqNum: 0x00001d1d Lifetime: 1082
Chksum: 0xf4a5 PDU Length: 176
Host_name: 7004
Attributes: IS-Type 1
TLV:137 Host_name: 7004

Level-1 LspID: 0049.0090.0100.00-00 SeqNum: 0x000000c2 Lifetime: 506
Chksum: 0x2f78 PDU Length: 829
Host_name: 9001
Attributes: IS-Type 1
TLV:137 Host_name: 9001

Level-1 LspID: 0049.0090.0200.00-00 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 0049.0090.0100.00-00 tlv 22 detail*

```
=====
                        ISIS LSDB (DETAIL)
=====
-----

Level-1 LspID: 0049.0090.0100.00-00      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
                0049.0080.0300.00 (8003)      Metric:10
                  SPBM Sub TLV:
                    port id: 194 num_port 1
                    Metric: 10
                0049.0090.0200.00 (9002)      Metric:10
                  SPBM Sub TLV:
                    port id: 6144 num_port 1
                    Metric: 10
                0049.0040.0200.00 (4002)      Metric:10
                  SPBM Sub TLV:
                    port id: 220 num_port 1
                    Metric: 10
                0049.0040.0100.00 (4001)      Metric:10
                  SPBM Sub TLV:
                    port id: 6151 num_port 1
                    Metric: 10
```

16.1.3.6 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 NAME	ASSOCIATION NAME	MEP ID	ADMIN
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, Level 4
spbm	4052	401	nodal	Vlan 4052, Level 4
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

16.1.3.7 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

```
----00:49:00:08:07:00    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

```
----00:49:00:08:07:00    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 (00:49:00:70:01:00),  vlan 4051
0    4001                (00:49:00:40:01:00)
1    9001                (00:49:00:90:01:00)
2    8003                (00:49:00:08:03:00)
3    8005                (00:49:00:08:05:00)
4    7001                (00:49:00:70:01:00)
```

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 00:49:00:08:05:ff
```

Please wait for l2traceroute to complete or press any key to abort

```
l2traceroute to (00:49:00:08:05:ff),  vlan 4051
0    4001                (00:49:00:40:01:00)
1    9001                (00:49:00:90:01:00)
2    8003                (00:49:00:08:03:00)
3    8005                (00:49:00:08:05:00)
```

```
4001:1#l2 traceroute vlan 4052 mac 00:49:00:08:05:ff
```


Please wait for l2traceroute to complete or press any key to abort

```
l2traceroute to (00:49:00:08:05:ff),  vlan 4052
0    4001                (00:49:00:40:01:00)
1    9002                (00:49:00:90:02:00)
2    8004                (00:49:00:08:04:00)
3    8006                (00:49:00:08:06:00)
```

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%

16.2 SPB L2 VSN



For this example, we will configure the following:

- L2 VSN VLANs
 - o VLAN ID = 1155 configured on switches 4001, 9001, 9002, 7001, and 7002 using I-SID = 1001155
 - The L2VSN is provisioned on SMLT cluster switches 9001 & 9002 for edge switch Switch-1
 - The L2VSN is provisioned on SMLT cluster switches 7001 & 7002 for edge server
 - o VLAN ID = 2256 configured on switches 8005, 8006 and 8007 using I-SID = 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 16.1.

16.2.1 VLAN and 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 and SLT 129 on the SMLT cluster 8005 & 8006.

VSP 4000 Switches

4001:

```
4001:1(config)#vlan create 1155 name VSN-Blue type port-mstprstp 0
4001:1(config)#vlan members add 1155 1/10
```

8005 & 8006 SMLT Cluster Switches – SLT on port 2/21 using SLT id 129

8005 & 8006: Same configuration on both switches

```
8005:5(config)#vlan create 2256 name VSN-Red type port-mstprstp 0
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 add 2256 2/21
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
```

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 – SMLT on port 3/25 using SMLT id 2

9001 & 9002: Same configuration on both switches

```
9001:1(config)#vlan create 1155 name VSN-Blue type port-mstprstp 0
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 2
9001:1(config-mlt)#exit
9001:1(config)#vlan members remove 1 3/25
9001:1(config)#vlan mlt 1155 2
9001:1(config)#vlan mlt 1155 1
```

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)#slpp packet-rx-threshold 5
9001:1(config-if)#vlacp timeout short
9001:1(config-if)#vlacp timeout-scale 5
9001:1(config-if)#vlacp fast-periodic-time 500
9001:1(config-if)#vlacp 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 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  
7001(config)#interface ethernet 25  
7001(config-if)#smlt 65  
7001(config-if)#exit
```

16.2.2 Layer 2 VSN configuration

VSP 4000 Switches

4001:

```
4001:1(config)#vlan i-sid 1155 1001155
```

VSP 7000 Switches

7001 & 7002: Same configuration on both switches

```
7001(config)#vlan i-sid 1155 1001155
```

SMLT Cluster Switches – 8005 & 8006

8005 & 8006: Same configuration on both switches

```
8005:5(config)#vlan i-sid 2256 1002256
```

ERS 8800 Switch - 8007

```
8007:5(config)#vlan i-sid 2256 1002256
```

VSP 9000 Switches

9001 & 9002: Same configuration on 9001 and 9002

```
9001:1(config)# vlan i-sid 1155 1001155
```

16.2.3 Verify Operations

16.2.3.1 Verify IS-IS I-SID

Step 1 – Show IS-IS I-SID

```
show isis spbm i-sid all
show isis spbm i-sid all <id|nick-name|vlan>
```

EDM

Configuration -> IS-IS -> SPBM -> I-SID

Results:

4001:

```
=====
                        SPBM ISID INFO
=====
ISID      SOURCE NAME  VLAN  SYSID      TYPE      HOST_NAME
-----
1001155  0.40.01      4051  0049.0040.0100  config    4001
1001155  0.70.01      4051  0049.0070.0100  discover  7001
1001155  0.70.02      4051  0049.0070.0200  discover  7002
1001155  0.90.01      4051  0049.0090.0100  discover  9001
1001155  0.90.02      4052  0049.0090.0200  discover  9002
```

7001: (7002 will be the same)

```
=====
                        SPBM ISID INFO
=====
ISID      SOURCE NAME  VLAN  SYSID      TYPE      HOST-NAME
-----
1001155  0.40.01      4051  0049.0040.0100  discover  4001
1001155  0.70.01      4051  0049.0070.0100  config    7001
1001155  0.70.02      4051  0049.0070.0200  discover  7002
1001155  0.90.01      4051  0049.0090.0100  discover  9001
1001155  0.90.02      4052  0049.0090.0200  discover  9002
```

9001: (9002 will be the same)

```
=====
                        SPBM ISID INFO
=====
ISID      SOURCE NAME  VLAN  SYSID                TYPE      HOST_NAME
-----
1001155  0.40.01             4051  0049.0040.0100      discover  4001
1001155  0.70.01             4051  0049.0070.0100      discover  7001
1001155  0.70.02             4051  0049.0070.0200      discover  7002
1001155  0.90.01             4051  0049.0090.0100      config    9001
1001155  0.90.02             4052  0049.0090.0200      discover  9002
```

8005: (8006 will be the same)

```
=====
                        SPBM ISID INFO
=====
ISID      SOURCE NAME  VLAN  SYSID                TYPE      HOST_NAME
-----
1002256  0.80.05             4051  0049.0080.0500      config    8005
1002256  0.80.06             4052  0049.0080.0600      discover  8006
1002256  0.80.07             4052  0049.0080.0700      discover  8007
```

8007:

```
=====
                        SPBM ISID INFO
=====
ISID      SOURCE NAME  VLAN  SYSID                TYPE      HOST_NAME
-----
1002256  0.80.05             4051  0049.0080.0500      discover  8005
1002256  0.80.06             4052  0049.0080.0600      discover  8006
1002256  0.80.07             4052  0049.0080.0700      config    8007
```

On each switch, verify the following:

Option	Verify
ISID TYPE	For switches 4001, 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.

16.2.3.2 Show IS-IS LSP Details pertaining to I-SIDs 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 I-SID

```
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 0049.0070.0100.00-00 tlv 144 sub-tlv 3 detail
```

```
=====
                        ISIS LSDB (DETAIL)
=====
-----
```

```
Level-1 LspID: 0049.0070.0100.00-00      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: 00-49-00-70-01-00
```

```
    BVID:4051
```

```
    Number of ISID's:1
```

```
        1001155(Both)
```

```
    Instance: 0
```

```
    Metric: 0
```

```
    B-MAC: 00-49-00-70-01-00
```

```
    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 0049.0070.0100.00-00 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 00-49-00-70-01-00 . 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, I-SID 1002256 should be displayed

16.2.3.3 Unknown unicast or multicast/broadcast traffic

The multicast addresses are built out of two pieces. Each SPB node must be configured with a unique Nick-name 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 I-SID id converted to hex forming the Multicast MAC address.

For example, in reference to 8005:

- Nickname = 0.80.05
- I-SID = 1002256 (0x0f:4b:10)
- Multicast address = 03:08:05:0f:4b:10 for I-SID 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 i-sid 1002256
```

SPBM MULTICAST FIB ENTRY INFO					
MCAST DA	ISID	BVLAN	SYSID	HOST-NAME	OUTGOING-INTERFACES
03:08:05:0f:4b:10	1002256	4051	0049.0080.0500	8005	2/5, IST
03:08:06:0f:4b:10	1002256	4052	0049.0080.0600	8006	
03:08:07:0f:4b:10	1002256	4052	0049.0080.0700	8007	

On each switch, verify the following information:

Option	Verify
MCAST DA	Verify that the correct multicast address for each switch I-SID 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 I-SID 1002256 • 8005: 03:08:05:0f:4b:10 • 8006: 03:08:06:0f:4b:10 • 8007: 03:08:07:0f:4b:10

16.2.3.4 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 i-sid <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	Port-1/10	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	00:49:00:70:02:00	4051	7002	9001	false
1155	learned	00:0c:29:9b:a8:31	00:49:00:90:01:ff	4051	9001	9001	false
1155	learned	00:18:71:ea:31:bb	00:49:00:70:01:00	4051	7001	9001	false

7001#*show mac-address-table spbm i-sid 1001155*

Mac Address Table Aging Time: 300

Learning Enabled Ports ALL

Number of addresses: 4

MAC Address	I-SID	Source	Vid	BVid	Dest-MAC	Dest-Sys-Name
00-0C-29-35-62-A4	1001155	Trunk 31		4051	00-49-00-70-02-00	7002
00-0C-29-9B-A8-31	1001155	Trunk 31		4051	00-49-00-90-01-FF	9001
00-0C-29-D6-81-E5	1001155	Trunk 31		4051	00-49-00-40-01-00	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 I-SID-1001155 for MAC addressed from I-SID 1001155 and I-SID-1002256 for MAC addressed from I-SID 1002256
DEST-MAC	For remote entries, the remote B-MAC address of the SPB switch should be shown as the remote destination MAC.

16.3 VSP 7000 & ERS 4800 – In-band Management via L2VSN

An L2VSN 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 16.10.

Switch	Parameter	Value
L2VSN – for in-band management		
8005, 8006	Mgmt VLAN	101
7001, 7002, 7003, 7004 4801	I-SID	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 i-sid 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 i-sid 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 i-sid 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
```

16.4 Multicast over L2VSN



Continuing from example used in Section 16.2, we will simply enable multicast support for L2VSN i-sid 1002256.

16.4.1 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
```

16.4.2 Enable IGMP

16.4.2.1 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.

16.4.2.2 Edge Switch

Assuming the edge switch is an Extreme stackable switch with the latest firmware, enable IGMP snoop and proxy.

Extreme Stackable Switches

CLI

```
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
```

16.4.3 Verify Operations

16.4.3.1 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
```

16.4.3.2 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
STATICIP
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
```

16.4.3.3 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 / i-sid 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            160000002  4051   8005
10.5.41.20            232.2.2.2            160000002  4052   8006
```

16.4.3.4 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-3 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: 0049.0080.0500.00-00      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: 0049.0080.0600.00-00      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 0049.0080.0700.00-00 tlv 144 sub-tlv 3 detail
```

```
=====
                        ISIS LSDB (DETAIL)
=====
-----
Level-1 LspID: 0049.0080.0700.00-00      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: 00-49-00-08-07-00
      BVID:4051
      Number of ISID's:1

                        1002256(Rx)

      Instance: 0
      Metric: 0
      B-MAC: 00-49-00-08-07-00
      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)

16.4.3.5 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

I-SID : 16000003

```
=====
1   8005           00:49:00:08:05:00 -> 8003           00:49:00:08:03:00
2   8003           00:49:00:08:03:00 -> 8007           00:49:00:08:07:00
```

```
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

I-SID : 16000003

```
=====
1   8006           00:49:00:08:06:00 -> 8004           00:49:00:08:04:00
2   8004           00:49:00:08:04:00 -> 8007           00:49:00:08:07:00
```

16.5 Inter VSN Routing



Continuing from configuration example 15.2 (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 15.2
- On SPB bridges 8005 and 8006, we will add the following:
 - o 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

16.6 Inter-ISID Configuration

In addition to the configuration to the configuration used in 13.2, we will add the following configuration.

16.6.1 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
```

16.6.2 Verification

16.6.2.1 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
=====
```

DST	MASK	NEXT	NH VRF	INTER					
				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
=====
```

DST	MASK	NEXT	NH VRF	INTER					
				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	PRI0	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	PRI0	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	I-SID-1002256	DYNAMIC	2118
10.5.41.3	00:1e:1f:48:f2:2d	2256	I-SID-1002256	DYNAMIC	2118
10.5.40.3	00:1e:1f:48:f2:24	1155	I-SID-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	I-SID-1001155	DYNAMIC	2154
10.5.40.51	00:0c:29:35:62:a4	1155	I-SID-1001155	DYNAMIC	2160
10.5.40.5	00:0c:29:d6:81:e5	1155	I-SID-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	I-SID-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	I-SID-1001155	DYNAMIC	2124
10.5.40.51	00:0c:29:35:62:a4	1155	I-SID-1001155	DYNAMIC	2143
10.5.40.5	00:0c:29:d6:81:e5	1155	I-SID-1001155	DYNAMIC	2159

16.6.2.2 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	I-SID-1001155	false 1 false
1155	learned	00:0c:29:9b:a8:31	I-SID-1001155	false 1 false
1155	learned	00:0c:29:d6:81:e5	I-SID-1001155	false 1 false
1155	learned	00:18:71:ea:31:bb	I-SID-1001155	false 1 false
1155	learned	00:1e:1f:48:f2:24	I-SID-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	I-SID-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	I-SID-1002256	false	1	true
2256	self	00:24:43:b4:e2:2d	Port-cpp	false	1	false

16.7SPB L3 VSN – SMLT



For this example, we will configure the SMLT switch cluster with the following:

- SPB IP
 - o SPB IP parameter must be enabled BEB switches 4002, 8005, 8006, and 8007
 - o An IS-IS source IP address must be configured (loopback/circuitless IP address)
- VRF's - BEB Nodes only
 - o VRF Red
 - VLAN ID = 2254 configured on switches 4002 and 8007
 - Assign I-SID 2002254 to VRF Red
 - o VRF Blue
 - VLAN ID = 2255 configured on switches 4002, 8005, and 8006
 - VRF Blue configured on SMLT cluster 8005 & 8006 for users off Switch-2
 - Assign I-SID 2002255 to VRF Blue

This example is a continuation from the base setup used in Section 16.1.

16.7.1 SPB IP Enable

16.7.1.1 IS-IS Layer 3 configuration

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
```

16.7.1.2 VRF Configuration

4002

```
4002:1(config)#ip vrf blue
4002:1(config)#ip vrf red
```

8005 & 8006

8005 & 8006: Same configuration on both switches

```
8005:5(config)#ip vrf blue
8005:5(config)#ip vrf red
```

8007

```
8007:5(config)#ip vrf red
```

16.7.2 VLAN Configuration

4002

```
4002:1(config)#vlan create 2254 name vsnred-2254 type port-mstprstp 0
4002:1(config)#vlan create 2255 name vsnblue-2255 type port-mstprstp 0
4002:1(config)#vlan members add 2254 1/11
4002:1(config)#vlan members add 2255 1/12
```

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 2254 name "vsnred-2254" type port-mstprstp 0
8005:5(config)#vlan create 2255 name "vsnblue-2255" type port-mstprstp 0
8005:5(config)#vlan ports 2/21 tagging tagAll
8005:5(config)#vlan members add 2255 2/21
8005:5(config)#vlan members add 2256 2/21
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)#slpp packet-rx
8005:5(config-if)#slpp packet-rx-threshold 5
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
8005:5(config)#slpp enable
8005:5(config)#slpp vid 2254,2255
```

8006 will have the same SMLT best practices configuration as 8005 except for the one item shown below.

```
8006:5(config-if)#slpp packet-rx-threshold 50
```

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
```

16.7.3 IPVPN Configuration

4002 - Add IP address and VRF to VLANs 2254 & 2255

```
4002:1(config)#interface vlan 2254
4002:1(config-if)#vrf red
4002:1(config-if)#ip address 192.168.15.1 255.255.255.0
4002:1(config-if)#exit
4002:1(config)#interface vlan 2255
4002:1(config-if)#vrf blue
4002:1(config-if)#ip address 192.168.5.1 255.255.255.0
4002:1(config-if)#exit
```

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 rsmt
8005:5(config-if)#ip rsmt holdup-timer 9999
8005:5(config-if)#exit
8005:5(config)#interface vlan 2254
8005:5(config-if)#vrf red
8005:5(config-if)#ip address 192.168.3.1 255.255.255.0
8005:5(config-if)#ip rsmt
8005:5(config-if)#ip rsmt holdup-timer 9999
8005:5(config-if)#exit
8005:5(config)#ip rsmt 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
8006:5(config)#interface vlan 2254
8006:5(config-if)#ip address 192.168.3.2 255.255.255.0
```

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)#ip rsmt
8007:5(config-if)#ip rsmt holdup-timer 9999
8007:5(config-if)#exit
```

16.7.4 Enable L3VSN Configuration

4002 - Enable L3 IPVPN and i-sid to VRF red and blue

```
4002:1(config)#router vrf red
4002:1(router-vrf)#ipvpn
4002:1(router-vrf)#i-sid 2002254
4002:1(router-vrf)#ipvpn enable
4002:1(router-vrf)#exit
4002:1(config)#router vrf blue
4002:1(router-vrf)#ipvpn
4002:1(router-vrf)#i-sid 2002255
4002:1(router-vrf)#ipvpn enable
4002:1(router-vrf)#exit
```

8005 and 8006 - Enable L3 IPVPN and i-sid to VRF blue

8005 & 8006: Same configuration on both switches

```
8005:5(config)#router vrf red  
8005:5(router-vrf)#ipvpn  
8005:5(router-vrf)#i-sid 2002254  
8005:5(router-vrf)#ipvpn enable  
8005:5(router-vrf)#exit  
8005:5(config)#router vrf blue  
8005:5(router-vrf)#ipvpn  
8005:5(router-vrf)#i-sid 2002255  
8005:5(router-vrf)#ipvpn enable  
8005:5(router-vrf)#exit
```

8007 - Enable L3 IPVPN and i-sid to VRF red

```
8007:5(config)#router vrf red  
8007:5(router-vrf)#ipvpn  
8007:5(router-vrf)#i-sid 2002254  
8007:5(router-vrf)#ipvpn enable  
8007:5(router-vrf)#exit
```

16.7.5 Enable direct interface redistribution

4002 - Redistribute IP Networks via IS-IS – Direct Interfaces

```
4002:1(config)#router vrf red
4002:1(router-vrf)#isis redistribute direct
WARNING: Routes will not be injected until apply command is issued after enable
command
4002:1(router-vrf)#isis redistribute direct enable
4002:1(router-vrf)#exit
4002:1(config)#router vrf blue
4002:1(router-vrf)#isis redistribute direct
WARNING: Routes will not be injected until apply command is issued after enable
command
4002:1(router-vrf)#isis redistribute direct enable
4002:1(router-vrf)#exit
4002:1(config)#isis apply redistribute direct vrf red
4002:1(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 red
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)#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
```

16.7.6 Verify Operations

16.7.6.1 Verify RSMLT Information

8005 & 8006 - Verify RSMLT is up and operational for both VRF instances

```
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

16.7.6.2 Verify IS-IS I-SID

Show IS-IS I-SID pertaining to each vrf instance

```
show ip ipvpn
```

Results:

4002:

```
4002:1#show ip ipvpn
```

```
VRF Name      : blue
Ivpn-state    : enabled
I-sid         : 2002255
```

```
VRF Name      : red
Ivpn-state    : enabled
I-sid         : 2002254
```

8005 & 8006:

```
8005:5#show ip vrf ipvpn
```

```
VRF Name      : blue
Ivpn-state    : enabled
I-sid         : 2002255
```

```
VRF Name      : red
Ivpn-state    : enabled
I-sid         : 2002254
```

8007:

```
8007:5#show ip vrf ipvpn
```

```
VRF Name      : red
Ivpn-state    : enabled
I-sid         : 2002254
```

On each switch, verify the following:

Option	Verify
VRF Name Ivpn-state I-sid	For the VRF Name of blue , the Ivpn-state should display enabled with an I-sid value of 2002255 . For the VRF Name of red , the Ivpn-state should display enabled with an I-sid value of 2002254 .

16.7.6.3 Show IS-IS SPB IP Unicast Forwarding database

Show IS-IS SPB IP Unicast FIB using i-sid 2002254 and 2002255 as used in this

configuration example

```
show isis spbm ip-unicast-fib id <i-sid id>
```

Results: Example from 4002

4002:

```
4002:1(config)#show isis spbm ip-unicast-fib id 2002255
```

```
=====
```

SPBM IP-UNICAST FIB ENTRY INFO							
=====							
VRF	ISID	Destination	NH BEB	OUTGOING VLAN	SPBM INTERFACE	COST	PREFIX COST

blue	2002255	192.168.33.0/24	8005	4051	1/47	30	1
blue	2002255	192.168.33.0/24	8005	4052	1/47	30	1
blue	2002255	192.168.33.0/24	8006	4051	8001	30	1
blue	2002255	192.168.33.0/24	8006	4052	8001	30	1

```
4002:1(config)#show isis spbm ip-unicast-fib id 2002254
```

```
=====
```

SPBM IP-UNICAST FIB ENTRY INFO							
=====							
VRF	ISID	Destination	NH BEB	OUTGOING VLAN	SPBM INTERFACE	COST	PREFIX COST

red	2002254	192.168.3.0/24	8005	4051	1/47	30	1
red	2002254	192.168.3.0/24	8005	4052	1/47	30	1
red	2002254	192.168.3.0/24	8006	4051	8001	30	1
red	2002254	192.168.3.0/24	8006	4052	8001	30	1
red	2002254	192.168.133.0/24	8007	4051	1/47	30	1
red	2002254	192.168.133.0/24	8007	4052	8001	30	1

16.7.6.4 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: 0049.0080.0500.00-00   SeqNum: 0x00000cdf   Lifetime:   606
      Chksum: 0xbcc6  PDU Length: 810
      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
      Vrf ISID:2002254
            Metric:1      Prefix Length:24
            IP Address: 192.168.3.0

Level-1 LspID: 0049.0080.0600.00-00   SeqNum: 0x0000174a   Lifetime:   498
      Chksum: 0x52c4  PDU Length: 787
      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
      Vrf ISID:2002254
```

```

Metric:1          Prefix Length:24
IP Address: 192.168.3.0

Level-1 LspID: 0049.0080.0700.00-00      SeqNum: 0x00000439      Lifetime: 639
      Chksum: 0x2b9a  PDU Length: 425
      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

Level-1 LspID: 0049.0040.0200.00-03      SeqNum: 0x0000001d      Lifetime: 1173
      Chksum: 0x9aea  PDU Length: 103
      Host_name: 4002
      Attributes:      IS-Type 1
TLV:184 SPBM IPVPN Reachability:
      Vrf ISID:2002254
            Metric:1          Prefix Length:24
            IP Address: 192.168.15.0

      Vrf ISID:2002255
            Metric:1          Prefix Length:24
            IP Address: 192.168.5.0

```

In reference to 4002, verify the following:

Option	Verify
Level 1 TLV:184	As this example, in reference to 4002, for the blue vrf I-SID 2002255, we should learn routes 192.168.33.0/24 from 8005 and 8006 and for the red vrf I-SID 2002254 , we should learn routes 192.168.3.0/24 from 8005 and 8006 and route 192.168.133.0/24 from 8007 .

16.7.6.5 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 blue
```

```
=====
```

IP Route - VRF blue

```
=====
```

			NH	INTER					
DST	MASK	NEXT	VRF	COST	FACE	PROT	AGE	TYPE	PRF
10.5.40.0	255.255.255.0	10.5.40.1	-	1	40	LOC	0	DB	0
192.168.5.0	255.255.255.0	192.168.5.1	-	1	2255	LOC	0	DB	0
192.168.33.0	255.255.255.0	8005	GlobalRouter	30	4051	ISIS	0	IBSV	7

3 out of 3 Total Num of Route Entries, 3 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 red*

```
=====
                        IP Route - VRF red
=====
DST                MASK                NEXT                NH                INTER
VRF                COST FACE    PROT AGE  TYPE PRF
-----
192.168.3.0        255.255.255.0    8005                GlobalRouter      30   4051  ISIS 0   IBSV 7
192.168.15.0       255.255.255.0    192.168.15.1       -                 1    2254  LOC  0   DB   0
192.168.133.0      255.255.255.0    8007                GlobalRouter      30   4051  ISIS 0   IBSV 7
```

3 out of 3 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

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 address for SPBM routes is the remote BEB MAC address

16.7.6.6 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 192.168.15.1
192.168.133.1 is alive
4002:1#ping 192.168.3.1 vrf red source 192.168.15.1
192.168.3.1 is alive
4002:1#ping 192.168.33.1 vrf blue source 192.168.5.1
192.168.33.1 is alive
4002:1#ping 192.168.33.2 vrf blue source 192.168.5.1
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 192.168.15.1
traceroute to 192.168.133.1, 30 hops max, 56 byte packets (vrf red)
 1  192.168.133.1  1.852 ms 2.410 ms 1.929 ms
4002:1#traceroute 192.168.33.1 vrf blue source 192.168.5.1
traceroute to 192.168.33.1, 30 hops max, 56 byte packets (vrf blue)
 1  192.168.33.1  2.154 ms 2.618 ms 2. 12 ms
```

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

10.5.40.1	d4:ea:0e:15:30:86	40	-	LOCAL	2160	
10.5.40.5	00:18:71:ea:31:bb	40	1/13	DYNAMIC	2016	
10.5.40.255	ff:ff:ff:ff:ff:ff	40	-	LOCAL	2160	
192.168.5.1	d4:ea:0e:15:30:85	2255	-	LOCAL	2160	
192.168.5.51	00:0c:29:9b:a8:31	2255	1/12	DYNAMIC	2120	
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	00:0c:29:9b:a8:31	Port-1/12	false	-
2255	self	d4:ea:0e:15:30:85	Port-cpp	false	-

16.8 Extending L3VSN to the ERS 4800 via L2VSN



Continuing from the L3VSN example in Section 16.7, we will extend the blue vrf to an ERS 4800 switch by adding a L2VSN between SPB bridges 8005, 8006 and 4801 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 I-SID 1002558 to local VLAN 2558 on SPB bridges 8005, 8006, and 8007
 - On bridges 8005 and 8006
 - Add VLAN 2558 to the blue vrf configured in Section 16.7
 - For VLAN 2558, add IP subnet 192.168.58.0/24 with a VRRP virtual IP address of 192.168.58.1 and VRRP Backup Master enabled

16.8.1 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 i-sid 2258 1002258
```

4801 – Assuming we are using local ports 3-11

```
4801(config)#vlan create 2558 type port
```

```
4801(config)#vlan configcontrol automatic
```

```
4801(config)#vlan members add 2558 3-11
```

```
4801(config)#i-sid 1002258 vlan 2558
```

16.8.2 VRF Configuration

8005 and 8006

8005:

```
8005:5(config)#interface vlan 2558
```

```
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 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
```

16.8.3 Verify Operations

16.8.3.1 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	PRI0	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	PRI0	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	up	200 (NO)

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 , its 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.

16.8.3.2 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
=====
```

DST	MASK	NEXT	NH VRF	COST	INTER FACE	PROT	AGE	TYPE	PRF
10.5.40.0	255.255.255.0	10.5.40.1	-	1	40	LOC	0	DB	0
192.168.5.0	255.255.255.0	192.168.5.1	-	1	2255	LOC	0	DB	0
192.168.33.0	255.255.255.0	8005	GlobalRouter	30	4051	ISIS	0	IBSV	7
192.168.58.0	255.255.255.0	8005	GlobalRouter	30	4051	ISIS	0	IBSV	7

16.9 Multicast over L3VSN



Continuing from example used in Section 16.7, we will simply enable multicast support for L3VSN I-SID 2002254 (red vrf) between SPB bridges 4001, 8005, 8006 and 8007.

16.9.1 Enable SPB Multicast – Global

16.9.1.1 IS-IS Layer 3 configuration

4001, 8005, 8006, and 8007: Enable SPB Multicast, global

4001, 8005, 8006 and 8007: Same configuration on all switches

```
4001:1(config)#router isis  
4001:1(config)#spbm 1 multicast enable  
4001:1(config)#exit
```

16.9.2 Enable Multicast VPN

4001, 8005, 8006, and 8007: Enable multicast VPN

4001, 8005, 8006 and 8007: Same configuration on all switches

```
4001:1(config)#router vrf red  
4001:1(router-vrf)#mvpn enable  
4001:1(router-vrf)#exit
```

16.9.3 Enable L3 SPB Multicast

4001, 8005, 8006, and 8007: Enable L3 SPB multicast at VLAN level

4001, 8005, 8006 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
```

16.9.4 Enable IGMP

16.9.4.1 Enable IGMPv2 at VLAN level

Default setting, no configuration required

16.9.4.2 Enable IGMPv3 at VLAN level

4001, 8005, 8006, and 8007: Enable IGMPv3, i.e. on VLAN 2254

4001, 8005, 8006 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
```

16.9.5 Edge Switch

Assuming the edge switch is an Extreme stackable switch with the latest firmware, enable IGMP snoop and proxy.

Switch-2 & Switch-3: 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
```

16.9.6 Verify Operations

16.9.6.1 Global Settings

Verify SPB multicast is enabled

```
show isis spbm multicast
```

Results: From 8005 & 8006

```
8005 8006> show isis spbm multicast
```

Response from 8005:

```
multicast : enable
fwd-cache-timeout : 210
```

Response from 8006:

```
multicast : enable
fwd-cache-timeout : 210
```

16.9.6.2 Verify IGMP interfaces

Verify IGMP interfaces

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

Results: From 8005 & 8006

```
8005 8006# show ip igmp interface vrf red
```

Response from 8005:

```
=====
IGMP Interface - VRF red
=====
      QUERY      OPER      QUERY  WRONG      LASTMEM
IF   INTVL STATUS VERS.  VERS  QUERIER  MAXRSPT  QUERY  JOINS  ROBUST  QUERY  MODE
-----
V2254 125   active  3     3    192.168.3.1 100     0     18     2     10    routed-spb
```

Response from 8006:

```
=====
IGMP Interface - VRF red
=====
      QUERY      OPER      QUERY  WRONG      LASTMEM
IF   INTVL STATUS VERS.  VERS  QUERIER  MAXRSPT  QUERY  JOINS  ROBUST  QUERY  MODE
-----
```

V2254 125 active 3 3 192.168.3.2 100 0 18 2 10 routed-spb

16.9.6.3 Verify IGMP cache/group and senders

Assuming the multicast sender connect to Switch-2 off SPB bridges 8005 & 8006 is sending a multicast stream of 232.1.1.1 with a receiver connected to Switch-3 off SPB bridge 8007.

Step 1 - Verify IGMP cache / group

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

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

Results:

```
8007:5#show ip igmp cache vrf red
```

```
=====
                        IGMP Cache - VRF red
=====
GRPADDR      INTERFACE  LASTREPORTER  EXPIRYTIME      VERSION1HOSTTIMER  TYPE
STATICIP
ORTS
-----
232.1.1.1      Vlan2254    192.168.133.10  0day,00h:03m:21s  0day,00h:00m:00s  DYNAMIC NULL
```

```
8007:5#show ip igmp group vrf red
```

```
=====
                        IGMP Group - VRF red
=====
GRPADDR      INPORT      MEMBER      EXPIRATION TYPE
-----
232.1.1.1      V2254-4/35  192.168.133.10  191      Dynamic
```

1 out of 1 group Receivers displayed

Step 2 - Verify IGMP sender

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

Results: From 8005 & 8006

```
8005 8006> show ip igmp sender vrf red
```

Response from 8005:

```
=====
                                IGMP Sender - VRF red
=====
                                PORT/
GRPADDR      IFINDEX  MEMBER      MLT      STATE
-----
232.1.1.1    Vlan 2254  192.168.3.10  2/21    NOTFILTERED
```

Response from 8006:

```
=====
                                IGMP Sender - VRF red
=====
                                PORT/
GRPADDR      IFINDEX  MEMBER      MLT      STATE
-----
232.1.1.1    Vlan 2254  192.168.3.10  2/21    NOTFILTERED
```

16.9.6.4 Verify SPB Multicast Routes

Assuming the multicast sender connect to switch 8007 is sending four multicast streams in the range from 239.10.10.10 to 239.10.10.13 while both receivers join all groups.

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 8007

```
8007:5#show isis spbm ip-multicast-route vrf red detail
```

```
=====
                        SPBM IP-MULTICAST ROUTE INFO - VRF NAME: red, VSN-ISID: 2002254
=====
```

Source	Group	Data ISID	BVLAN	NNI	Rcvrs	UNI Rcvrs	Source-BEB
192.168.3.10	232.1.1.1	16000005	4051	-		V2254:4/35	8005
192.168.3.10	232.1.1.1	16000005	4052	-		V2254:4/35	8006

16.9.6.5 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 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 8005 and 8006 perspective taken from 8007

```
8007:5#show isis lsdb tlv 185 lspid 0049.0080.0500.00-00 detail
```

```
=====
                        ISIS LSDB (DETAIL)
=====
-----

Level-1 LspID: 0049.0080.0500.00-00      SeqNum: 0x00000d37      Lifetime: 782
      Chksum: 0xd91d  PDU Length: 889
      Host_name: 8005
      Attributes:      IS-Type 1
TLV:185 SPBM IPVPN :
      VSN ISID:2002254
      BVID      :4051
      Metric:0
      IP Source Address: 192.168.3.10
      Group Address   : 232.1.1.1
      Data ISID       : 16000005
      TX              : 1
```

```
8007:5#show isis lsdb tlv 185 lspid 0049.0080.0600.00-00 detail
```

```
=====
                        ISIS LSDB (DETAIL)
=====
-----

Level-1 LspID: 0049.0080.0600.00-00      SeqNum: 0x000017a2      Lifetime: 813
      Chksum: 0x2364  PDU Length: 866
      Host_name: 8006
      Attributes:      IS-Type 1
TLV:185 SPBM IPVPN :
```



```
VSN ISID:2002254
BVID      :4052
Metric:0
IP Source Address: 192.168.3.10
Group Address      : 232.1.1.1
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 8007

```
8007:5#show isis lsdb tlv 144 sub-tlv 3 lspid 0049.0080.0700.00-00 detail
```

```
=====
                        ISIS LSDB (DETAIL)
=====
-----

Level-1 LspID: 0049.0080.0700.00-00      SeqNum: 0x00000493      Lifetime: 747
      Chksum: 0x75f1  PDU Length: 461
      Host_name: 8007
      Attributes:      IS-Type 1
TLV:144 SUB-TLV 3      ISID:
      |
      |
      Instance: 0
      Metric: 0
      B-MAC: 03-08-05-00-00-00
      BVID:4051
      Number of ISID's:1
              16000005(Rx)

      Instance: 0
      Metric: 0
      B-MAC: 03-08-06-00-00-00
      BVID:4052
```

Number of ISID's:1

16000005(Rx)

8007# show isis lsdb tlv 144 sub-tlv 3 lspid 0049.0080.0500.00-00 detail

```
=====
                        ISIS LSDB (DETAIL)
=====
-----
```

Level-1 LspID: 0049.0080.0500.00-00 SeqNum: 0x00000d3a Lifetime: 1143

Chksum: 0x7a77 PDU Length: 889

Host_name: 8005

Attributes: IS-Type 1

TLV:144 SUB-TLV 3 ISID:

|

|

Instance: 0

Metric: 0

B-MAC: 03-00-00-00-00-00

BVID:4051

Number of ISID's:1

16000005(Tx)

8007# **show isis lsdb tlv 144 sub-tlv 3 lspid 0049.0080.0600.00-00 detail**

```
=====
                        ISIS LSDB (DETAIL)
=====
-----
```

Level-1 LspID: 0049.0080.0600.00-00 SeqNum: 0x000017a5 Lifetime: 1136

Chksum: 0x95ec PDU Length: 866

Host_name: 8006

Attributes: IS-Type 1

TLV:144 SUB-TLV 3 ISID:

|

|

Instance: 0

Metric: 0

B-MAC: 03-00-00-00-00-00

BVID:4052

Number of ISID's:1

16000005(Tx)

16.9.6.6 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 8005 & 8006, we will use the following command to view the multicast route for group address 239.1.1.55l2

```
8005 8006> l2 tracemroute source 192.168.3.10 group 232.1.1.1 vlan 2254 vrf red
```

Response from 8005:

Please wait for l2tracemroute to complete or press any key to abort

Source : 192.168.3.10

Group : 232.1.1.1

VRF : red ID 2

BMAC : 03:08:05:f4:24:06

B-VLAN : 4051

I-SID : 16000005

```
=====
1  8005          00:49:00:08:05:00 -> 8003          00:49:00:08:03:00
2  8003          00:49:00:08:03:00 -> 8007          00:49:00:08:07:00
```

Response from 8006:

Please wait for l2tracemroute to complete or press any key to abort

Source : 192.168.3.10

Group : 232.1.1.1

VRF : red ID 2


BMAC : 03:08:06:f4:24:06

B-VLAN : 4052

I-SID : 16000006

```
=====
1  8006          00:49:00:08:06:00 -> 8004          00:49:00:08:04:00
2  8004          00:49:00:08:04:00 -> 8007          00:49:00:08:07:00
```

16.10 SPB IP Shortcuts



- SPB IP
 - o SPB IP parameter must be enabled on BEB bridges 4001, 4002, 9001, 9002, 8005, 8006, and 8007
 - o An IS-IS source IP address must be configured (loopback/circuitless IP address)
- IP Configuration
 - o CLIP/Loopback #1 as shown in the above diagram
 - o Local VLAN and IP addressing as shown in the above diagram
 - o 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 16.1.

16.10.1 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
```

16.10.1.1 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 and enable SPB IP

4001 and 4002: Same configuration on both switches

```
4001:1(config)# router isis
4001:1(config-isis)#redistribute direct
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
```

16.10.2 ECMP

Enable ECMP using the following command

- `ip ecmp`

16.10.3 Local VLAN configuration

There are no special configuration requirements for the local VLAN provisioning. For the SMLT cluster configuration, you can enable RSMLT Edge or VRRP with backup master using the SMLT best practices. All that is required is adding an IP address to the VLAN itself with no routing protocol. Please refer to the *Switch Clustering using SMLT with ERS* Technical Configuration Guide for more information, publication number *NN48500-518*.

16.10.4 Verify Operations

16.10.4.1 Verify IP Route Table

Verify IP Routes

show ip route

Results: From bridge 4001

4001:

IP Route - GlobalRouter									
DST	MASK	NEXT	NH	INTER					
			VRF	COST	FACE	PROT	AGE	TYPE	PRF
10.1.1.1	255.255.255.255	9001	GlobalRouter	10	4051	ISIS	0		IBSE
10.1.1.1	255.255.255.255	9001	GlobalRouter	10	4052	ISIS	0		IBSE
10.1.1.2	255.255.255.255	9002	GlobalRouter	10	4051	ISIS	0		IBSE
10.1.1.2	255.255.255.255	9002	GlobalRouter	10	4052	ISIS	0		IBSE
10.1.1.5	255.255.255.255	8005	GlobalRouter	30	4051	ISIS	0		IBSE
10.1.1.5	255.255.255.255	8005	GlobalRouter	30	4052	ISIS	0		IBSE
10.1.1.6	255.255.255.255	8006	GlobalRouter	30	4051	ISIS	0		IBSE
10.1.1.6	255.255.255.255	8006	GlobalRouter	30	4052	ISIS	0		IBSE
10.1.1.7	255.255.255.255	8007	GlobalRouter	30	4051	ISIS	0		IBSE
10.1.1.7	255.255.255.255	8007	GlobalRouter	30	4052	ISIS	0		IBSE
10.4.4.1	255.255.255.255	10.4.4.1	-	1	0	LOC	0		DB
10.4.4.2	255.255.255.255	4002	GlobalRouter	20	4051	ISIS	0		IBSE
10.4.4.2	255.255.255.255	4002	GlobalRouter	20	4052	ISIS	0		IBSE
10.5.120.0	255.255.255.0	9001	GlobalRouter	10	4051	ISIS	0		IBSE
10.5.120.0	255.255.255.0	9002	GlobalRouter	10	4051	ISIS	0		IBSE
10.5.120.0	255.255.255.0	9001	GlobalRouter	10	4052	ISIS	0		IBSE
10.5.120.0	255.255.255.0	9002	GlobalRouter	10	4052	ISIS	0		IBSE
10.7.101.0	255.255.255.0	8007	GlobalRouter	30	4051	ISIS	0		IBSE
10.7.101.0	255.255.255.0	8007	GlobalRouter	30	4052	ISIS	0		IBSE
10.12.100.0	255.255.255.0	8005	GlobalRouter	30	4051	ISIS	0		IBSE
10.12.100.0	255.255.255.0	8006	GlobalRouter	30	4051	ISIS	0		IBSE
10.12.100.0	255.255.255.0	8005	GlobalRouter	30	4052	ISIS	0		IBSE
10.12.100.0	255.255.255.0	8006	GlobalRouter	30	4052	ISIS	0		IBSE

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 *show ip route spbm-nh-as-mac*.

16.10.4.2 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							
=====							
VRF	ISID	Destination	NH BEB	OUTGOING VLAN	SPBM INTERFACE	PREFIX 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

16.10.4.3 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
```

Results: From bridge 8007


4001:

```
4001:1#show isis lsdb lspid 0049.0080.0700.00-00 tlv 135 detail
```

```
=====
                        ISIS LSDB (DETAIL)
=====
-----

Level-1 LspID: 0049.0080.0700.00-00      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
```

16.11 Multicast over IP Shortcuts



Continuing from example used in Section 16.10, we will simply enable multicast support for IP Shortcuts on all SPB bridges.

16.11.1 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
```

16.11.2 Enable IP Multicast at VLAN level

Enable IP multicast at VLAN level

4001, 4002, 9001, 9002, 8005, 8006, and 8007: Same configuration on both 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
```

16.12 Verify Operations

16.12.1 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	

16.12.2 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 receiver 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 STATICPORTS	INTERFACE	LASTREPORTER	EXPIRYTIME	VERSION1HOSTTIMER	TYPE
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

16.12.3 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
-----
```

16.12.4 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 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: 0049.0080.0500.00-00    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: 0049.0080.0600.00-00    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 0049.0080.0500.00-00 tlv 144 detail
```

```
=====
                        ISIS LSDB (DETAIL)
=====
-----

Level-1 LspID: 0049.0080.0500.00-00      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 0049.0080.0600.00-00 tlv 144 detail
```

```
=====
                        ISIS LSDB (DETAIL)
=====
-----

Level-1 LspID: 0049.0080.0600.00-00      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 0049.0080.0700.00-00 tlv 144 detail*

```

=====
                        ISIS LSDB (DETAIL)
=====
-----

```

```

Level-1 LspID: 0049.0080.0700.00-00      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)

```

16.12.5 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

I-SID : 16000008

```
=====
1  8005          00:49:00:08:05:00 -> 8003          00:49:00:08:03:00
2  8003          00:49:00:08:03:00 -> 8007          00:49:00:08:07:00
```

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

I-SID : 16000009

```
=====
1  8006          00:49:00:08:06:00 -> 8004          00:49:00:08:04:00
2  8004          00:49:00:08:04:00 -> 8007          00:49:00:08:07:00
```

16.13 IPVPN-Lite L3 VPN over IS-IS



For this configuration example, we will show how to configure IPVPN Lite on top of SPB allowing support for both leaking routes between VRF's and providing hub-and-spoke operation. For this example, we will simply configure IPVPN Lite using VRF green.

Note, for this example, SPB bridges 8001 and 8002 are ERS 8800 switches. IPVPN-Lite over ISIS is only supported on the ERS 8800.

- SPB and IS-IS configuration
 - o Please see previous example in Section 16.1
- IPVPN Lite Configuration
 - o Add a Circuitless/Loopback IP address for iBGP peering using ID #1
 - 10.1.1.1/32 on 8001, 10.1.1.2/32 on 8002, and 10.1.1.7/32 on 8007
 - Enable IPVPN Lite
 - o Add a Circuitless/Loopback IP address for IPVPN Lite using subnet 172.16.x.254/24 using ID #4
 - 172.16.1.254 on 8001, 172.16.2.254/24 on 8002, and 172.16.7.254/24 on 8007
 - o BGP Configuration
 - AS = 65000
 - iBGP peering between 8001, 8002, and 8007 using CLIP/Loopback #1 where Clip/Loopback #1 (also the BGP router-id)
 - Enable IPVPN Lite capability to each iBGP peer
 - o VPN Lite Configuration
 - Use VRF green id #4 with IPVPN enabled

- Add a Route Distinguisher (RD) where the ID is derived from the CLIP/Loopback address and vrf ID <CLIP/Loopback:vrf id>
 - 172.16.1.1:4 on 8001, 172.16.2.1:4 on 8002, and 172.16.7.1:4 on 8007
- Add a Route Target (RT) where the ID is derived from the autonomous system ID
 - 65000:60004 on all switches 8001, 8002, and 8007

This example is a continuation from the base setup used in Section 16.1.



Please note that multicast is not supported on IPVPN-Lite L3 VPN. Also, the VSP 9000 does not support this feature.

16.13.1 SMLT Cluster

16.13.1.1 Add Circuitless IP addresses

As we have already added CLIP addresses 1 and 2 in the previous example, we will create CLIP 1 to be used for iBGP peering and CLIP 4 to be used for IPVPN-Lite. Overall, we will configure the following:

- Loopback 1 with IP address of 10.1.1.1/32 which will be used for iBGP peering on switch 8001 and CLIP 1 with IP address of 10.1.1.2/32 which will be used for iBGP peering on switch 8002
- CLIP 4 with IP subnet of 172.16.1.254/24 on 8001 and CLIP 4 with IP address of 172.16.2.254/24 on 8002 where an IP address will be used in this range for the IPVPN Lite RD value

8001: Step 1 – Add loopback 1 to be used for iBGP peering

```
CLI
8001:1(config)#interface loopback 1
8001:1(config-if)#ip address 10.1.1.1/32
8001:1(config-if)#exit
```

8002: Step 1 – Add loopback 1 to be used for iBGP peering

```
CLI
8002:1(config)#interface loopback 1
8002:1(config-if)#ip address 10.1.1.2/32
8002:1(config-if)#exit
```

8001: Step 2 – Add loopback 4 and enable IPVPN-Lite

```
CLI
8001:1(config)#interface loopback 4
8001:1(config-if)#ip address 172.16.1.254/24
8001:1(config-if)#ip ipvpn-lite-capability enable
8001:1(config-if)#exit
```

8002: Step 2 – Add loopback 4 and enable IPVPN-Lite

```
CLI
8002:1(config)#interface loopback 4
8002:1(config-if)#ip address 172.16.2.254/24
8002:1(config-if)#ip ipvpn-lite-capability enable
8002:1(config-if)#exit
```

16.13.1.2 Enable IP Routing over IS-IS

8001: Step 1 – Add loopback 1 to be used for iBGP peering

```
CLI
8001:1(config)#router isis
8001:1(config-isis)#ip-source-address 10.1.1.1
8001:1(config-isis)#spbm 1 ip enable
8001:1(config-isis)#exit
```

8002: Step 1 – Add loopback 1 to be used for iBGP peering

```
CLI
8002:1(config)#router isis
8002:1(config-isis)#ip-source-address 10.1.1.2
8002:1(config-isis)#spbm 1 ip enable
8002:1(config-isis)#exit
```

16.13.1.3 Enable BGP and IP-VPN Lite

8001: Step 1 – Add BGP router-id. Note that the BGP router-id is derived from the OSPF router-id

```
CLI
8001:1(config)#router ospf
8001:1(config-ospf)#router-id 10.1.1.1
8001:1(config-ospf)#exit
```

8002: Step 1 – Add BGP router-id. Note that the BGP router-id is derived from the OSPF router-id

```
CLI
8002:1(config)#router ospf
8002:1(config-ospf)#router-id 10.1.1.2
8002:1(config-ospf)#exit
```

8001: Step 2 – Add BGP global settings

```
CLI
8001:1(config)#router bgp
8001:1(router-bgp)#no auto-summary
8001:1(router-bgp)#no synchronization
8001:1(router-bgp)#quick-start enable
```

<pre> 8001:1(router-bgp)#enable 8001:1(router-bgp)#exit 8001:1(config)#router bgp 65000 enable </pre>
8002: Step 2 – Add BGP global settings
<pre> CLI 8002:1(config)#router bgp 8002:1(router-bgp)#no auto-summary 8002:1(router-bgp)#no synchronization 8002:1(router-bgp)#quick-start enable 8002:1(router-bgp)#enable 8002:1(router-bgp)#exit 8002:1(config)#router bgp 65000 enable </pre>
8001: Step 3 – Add BGP peers and enable IPVPN-Lite to 8002 peer
<pre> CLI 8001:1(config)#router bgp 8001:1(router-bgp)#neighbor 10.1.1.7 8001:1(router-bgp)#no neighbor 10.1.1.7 enable 8001:1(router-bgp)#neighbor 10.1.1.7 remote-as 65000 8001:1(router-bgp)#neighbor 10.1.1.7 update-source 10.1.1.1 8001:1(router-bgp)#neighbor 10.1.1.7 address-family vpnv4 enable 8001:1(router-bgp)#neighbor 10.1.1.7 ipvpn-lite-capability enable 8001:1(router-bgp)#neighbor 10.1.1.7 enable 8001:1(router-bgp)#exit </pre>
8002: Step 3 – Add BGP peers and enable IPVPN-Lite to 8001 peer
<pre> CLI 8002:1(config)#router bgp 8002:1(router-bgp)#neighbor 10.1.1.7 8002:1(router-bgp)#no neighbor 10.1.1.7 enable 8002:1(router-bgp)#neighbor 10.1.1.7 remote-as 65000 8002:1(router-bgp)#neighbor 10.1.1.7 update-source 10.1.1.2 8002:1(router-bgp)#neighbor 10.1.1.7 address-family vpnv4 enable 8002:1(router-bgp)#neighbor 10.1.1.7 ipvpn-lite-capability enable 8002:1(router-bgp)#neighbor 10.1.1.7 enable 8002:1(router-bgp)#exit </pre>

16.13.1.4 L3VSN Configuration

8001: Step 1 – VRF Green configuration

```

CLI
8001:1(config)#ip vrf green
8001:1(config)#vlan create 1010 name VRF-Green type port-mstprstp 0
8001:1(config)#vlan members add 1010 3/5
8001:1(config)#interface vlan 1010
8001:1(config-if)#vrf green
8001:1(config-if)#ip address 10.1.110.1 255.255.255.0
8001:1(config-if)#ip rsmt
8001:1(config-if)#ip rsmt holdup-timer 9999
8001:1(config-if)#exit
8001:1(config)#vlan mlt 1010 1
  
```

8002: Step 1 – VRF Green configuration

```

CLI
8002:1(config)#ip vrf green
8002:1(config)#vlan create 1010 name VRF-Green type port-mstprstp 0
8002:1(config)#vlan members add 1010 3/5
8002:1(config)#interface vlan 1010
8002:1(config-if)#vrf green
8002:1(config-if)#ip address 10.1.110.2 255.255.255.0
8002:1(config-if)#ip rsmt
8002:1(config-if)#ip rsmt holdup-timer 9999
8002:1(config-if)#exit
8002:1(config)#vlan mlt 1010 1
  
```

8001: Step 2 – Enable IPVPN-Lite on VRF Green

```

CLI
8001:1(config)#router vrf green
8001:1(router-vrf)#ipvpn
8001:1(router-vrf)#rd 172.16.1.254 4
8001:1(router-vrf)#route-target both 65000 60004
8001:1(router-vrf)#no ip bgp auto-summary
8001:1(router-vrf)#ipvpn enable
8001:1(router-vrf)#exit
  
```

8002: Step 2 – Enable IPVPN-Lite on VRF Green

```

CLI
  
```

```
8002:1(config)#router vrf green  
8002:1(router-vrf)#ipvpn  
8002:1(router-vrf)#rd 172.16.2.254 4  
8002:1(router-vrf)#route-target both 65000 60004  
8002:1(router-vrf)#no ip bgp auto-summary  
8002:1(router-vrf)#ipvpn enable  
8002:1(router-vrf)#exit
```

16.13.28007 Configuration

16.13.2.1 Add Circuitless IP addresses

As we have already added CLIP addresses 1 and 2 in the previous example, we will create CLIP 1 to be used for iBGP peering and CLIP 4 to be used for IPVPN-Lite. Overall, we will configure the following:

- CLIP 1 with IP address of 10.1.1.7/32 which will be used for iBGP peering
- CLIP 4 with IP address of 172.16.7.254/24 where an IP address will be used in this range for the IPVPN Lite RD value

8007: Step 1 – Add CLIP 1 to be used for iBGP peering

CLI

```
8007:5# config ip circuitless-ip-int 1 create 10.1.1.7/32
```

8007: Step 2 – Add CLIP 4 and enable IPVPN-Lite

CLI

```
8007:5# config ip circuitless-ip-int 4 create 172.16.7.254/24
```

```
8007:5# config ip circuitless-ip-int 4 ipvpn-lite-capability enable
```

16.13.2.2 Enable BGP and IP-VPN Lite

8007: Step 1 – Add BGP router-id. Note that the BGP router-id is derived from the OSPF router-id

CLI

```
8007:5# config ip ospf router-id 10.1.1.7
```

8007: Step 2 – Add BGP global settings

CLI

```
8007:5# config ip bgp auto-summary disable
```

```
8007:5# config ip bgp synchronization disable
```

```
8007:5# config ip bgp local-as 65000
```

```
8007:5# config ip bgp aggregation disable
```

```
8007:5# config ip bgp enable
```

```
8007:5# config ip bgp quick-start enable
```

8007: Step 3 – Add BGP peers and enable IPVPN-Lite on each peer

CLI

```
8007:5# config ip bgp neighbor 10.1.1.1 create
```

```
8007:5# config ip bgp neighbor 10.1.1.2 create
```

```
8007:5# config ip bgp neighbor 10.1.1.1 remote-as 65000
```

```
8007:5# config ip bgp neighbor 10.1.1.2 remote-as 65000
8007:5# config ip bgp neighbor 10.1.1.1 update-source-interface 10.1.1.7 add
8007:5# config ip bgp neighbor 10.1.1.2 update-source-interface 10.1.1.7 add
8007:5# config ip bgp neighbor 10.1.1.1 address-family vpnv4 enable
8007:5# config ip bgp neighbor 10.1.1.2 address-family vpnv4 enable
8007:5# config ip bgp neighbor 10.1.1.1 ipvpn-lite-capability enable
8007:5# config ip bgp neighbor 10.1.1.2 ipvpn-lite-capability enable
8007:5# config ip bgp neighbor 10.1.1.1 admin-state enable
8007:5# config ip bgp neighbor 10.1.1.2 admin-state enable
```

16.13.2.3 Enable IP Routing over IS-IS

8007: Step 1 – Add CLIP 1 to be used for iBGP peering

```
CLI
8007:5# config isis ip source-address 10.1.1.7
8007:5# config isis spbm 1 ip enable
```

16.13.2.4 L3VSN Configuration

8007: Step 1 – VRF Green configuration

```
CLI
8007:5# config ip vrf green create
8007:5# config vlan 1010 create byport-mstprstp 0 name VPN-Green
8007:5# config vlan 1010 port add 4/35
8007:5# config vlan 1010 vrf green
8007:5# config vlan 1010 ip create 10.7.110.1/24
```

8007: Step 2 – Enable IPVPN-Lite on VRF Green

```
CLI
8007:5# config ip vrf green ipvpn create
8007:5# config ip vrf green ipvpn rd 172.16.7.1:4
8007:5# config ip vrf green ipvpn rt add both 65000:60004
8007:5# config ip vrf green ipvpn enable
8007:5# config ip vrf green bgp auto-summary disable
```

16.13.3 Verify Operations

16.13.3.1 IP Routing over IS-IS

Step 1 – Display SPB IP unicast forwarding database

CLI

show isis spbm ip-unicast-fib

Results: The following is from 8001 and 8007 perspective

8001:1# ***show isis spbm ip-unicast-fib***

```
=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
```

				OUTGOING	SPBM	PREFIX
VRF	ISID	Destination	NH BEB	VLAN	INTERFACE	COST
GRT	-	10.1.1.2/32	8002	4051	IST	10
GRT	-	10.1.1.2/32	8002	4052	IST	10
GRT	-	10.1.1.7/32	8007	4051	3/3	20
GRT	-	10.1.1.7/32	8007	4052	3/3	20
GRT	-	172.16.2.0/24	8002	4051	IST	10
GRT	-	172.16.2.0/24	8002	4052	IST	10
GRT	-	172.16.7.0/24	8007	4051	3/3	20
GRT	-	172.16.7.0/24	8007	4052	3/3	20

8007:5# ***show isis spbm ip-unicast-fib***

```
=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
```

				OUTGOING	SPBM	PREFIX
VRF	ISID	Destination	NH BEB	VLAN	INTERFACE	COST
GRT	-	10.1.1.1/32	8001	4051	3/27	20

GRT	-	10.1.1.1/32	8001	4052 3/27	20	1
GRT	-	10.1.1.2/32	8002	4051 3/28	20	1
GRT	-	10.1.1.2/32	8002	4052 3/28	20	1
GRT	-	172.16.1.0/24	8001	4051 3/27	20	1
GRT	-	172.16.1.0/24	8001	4052 3/27	20	1
GRT	-	172.16.2.0/24	8002	4051 3/28	20	1
GRT	-	172.16.2.0/24	8002	4052 3/28	20	1

Step 2 – Display IP route table

CLI

show ip route info

show ip route

Results: The following is from 8001 and 8007 perspective

8001:1# *show ip route info*

IP Route - GlobalRouter										
DST	MASK	NEXT	NH	INTER			AGE	TYPE	PRF	
			VRF	COST	FACE	PROT				
10.1.1.1	255.255.255.255	10.1.1.1	-	1	0	LOC	0	DB	0	
10.1.1.2	255.255.255.255	8002	GlobalRout~	10	4051	ISIS	0	IBSE	7	
10.1.1.2	255.255.255.255	8002	GlobalRout~	10	4052	ISIS	0	IBSE	7	
10.1.1.7	255.255.255.255	8007	GlobalRout~	20	4051	ISIS	0	IBSE	7	
10.1.1.7	255.255.255.255	8007	GlobalRout~	20	4052	ISIS	0	IBSE	7	
172.16.1.0	255.255.255.0	172.16.1.254	-	1	0	LOC	0	DB	0	
172.16.2.0	255.255.255.0	8002	GlobalRout~	10	4051	ISIS	0	IBSE	7	
172.16.2.0	255.255.255.0	8002	GlobalRout~	10	4052	ISIS	0	IBSE	7	
172.16.7.0	255.255.255.0	8007	GlobalRout~	20	4051	ISIS	0	IBSE	7	
172.16.7.0	255.255.255.0	8007	GlobalRout~	20	4052	ISIS	0	IBSE	7	

8007:5# *show ip route info*

IP Route - GlobalRouter										
			NH	INTER						
DST	MASK	NEXT	VRF	COST	FACE	PROT	AGE	TYPE	PRF	

10.1.1.1	255.255.255.255	8001	GlobalRout~	20	4051	ISIS	0	IBSE	7	
10.1.1.1	255.255.255.255	8001	GlobalRout~	20	4052	ISIS	0	IBSE	7	

10.1.1.2	255.255.255.255	8002	GlobalRout~	20	4051	ISIS	0	IBSE	7
10.1.1.2	255.255.255.255	8002	GlobalRout~	20	4052	ISIS	0	IBSE	7
172.16.1.0	255.255.255.0	8001	GlobalRout~	20	4051	ISIS	0	IBSE	7
172.16.1.0	255.255.255.0	8001	GlobalRout~	20	4052	ISIS	0	IBSE	7
172.16.2.0	255.255.255.0	8002	GlobalRout~	20	4051	ISIS	0	IBSE	7
172.16.2.0	255.255.255.0	8002	GlobalRout~	20	4052	ISIS	0	IBSE	7
172.16.7.0	255.255.255.0	172.16.7.254	-	1	0	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

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 IBS . The next hop address for SPBM routes is the remote BEB MAC address.

16.13.3.2 BGP Operation

Use the following command to ensure BGP peering between neighbors. Note that the SMLT cluster only requires peering to ERS-1 and not between the cluster switches.

Step 1 – Ensure BGP peering between neighbors

CLI
show ip bgp sum

Results: The following is from 8001 and 8007

```
8001:1# show ip bgp sum

=====

                        BGP Summary - GlobalRouter

=====

                        BGP version - 4
                        local-as - 65000
                        Identifier - 10.1.1.1
                        Decision state - Idle

                        The total number of routes is 0


BGP NEIGHBOR INFO :

      NEIGHBOR      RMTAS      STATE      HLDTM      KPALV      HLDCFG      KPCFG      WGHT      CONRTY      ADVINT
-----
10.1.1.7           65000      Established 180         60         180         60         100         120         5
```

Total bgp neighbors: 1

8007:5# *show ip bgp sum*

=====

BGP Summary - GlobalRouter

=====

BGP version - 4

local-as - 65000

Identifier - 10.1.1.7

Decision state - Idle

The total number of routes is 0

BGP NEIGHBOR INFO :

NEIGHBOR	RMTAS	STATE	HLDTM	KPALV	HLDCFG	KPCFG	WGHT	CONRTY	ADVINT

10.1.1.2	65000	Established	180	60	180	60	100	120	5
10.1.1.1	65000	Established	180	60	180	60	100	120	5

Total bgp neighbors: 2

16.13.3.3 IP Route Table

Use the following command to display the routes for VRF Green.

Step 1 – Display IP route table for each VRF instance

CLI

show ip route info vrf green

show ip route vrf green

Results: The following is from 8001 and 8007 for VRF green

8001:1# *show ip route vrf green*

```
=====
```

IP Route - VRF green									
=====									
			NH	INTER					
DST	MASK	NEXT	VRF	COST	FACE	PROT	AGE	TYPE	PRF

10.1.110.0	255.255.255.0	10.1.110.1	-	1	1010	LOC	0	DB	0
10.7.110.0	255.255.255.0	8007	GlobalRout~	0	4051	BGP	0	IBSV	175

8007:5# *show ip route info vrf green*

```
=====
```

IP Route - VRF green									
=====									
			NH	INTER					
DST	MASK	NEXT	VRF	COST	FACE	PROT	AGE	TYPE	PRF

10.1.110.0	255.255.255.0	8001	GlobalRout~	0	4051	BGP	0	IBSV	175
10.7.110.0	255.255.255.0	10.7.110.1	-	1	1010	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

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 BGP with the appropriate next-hop address and type of IBSV . The next hop address for SPBM routes is the remote BEB MAC address.

17. Restrictions and Limitations

17.1 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/MSTI 63 internally and can not be used by other VLAN/MSTI. If STG 63 is used in the configuration on non-SPB customer network, then STG 62 is used internally.
- 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. 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.

17.2 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 not currently 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

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.

Since ERS 8600/8800 HA framework, 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.

18. Reference Documentation

Document Title	Publication Number	Description