



Configuring Fabric Layer 3 Services on VSP Operating System Software

Release 7.1 (VOSS)
9035544 Rev.01
August 2018

© 2017-2018, 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/software-licensing> 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> 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.

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/>

Contents

Chapter 1: Preface	6
Purpose.....	6
Training.....	6
Providing Feedback to Us.....	6
Getting Help.....	7
Extreme Networks Documentation.....	7
Subscribing to Service Notifications.....	8
Chapter 2: New in this Document	9
Notice about Feature Support.....	9
Chapter 3: SPBM and IS-IS configuration workflow	10
Chapter 4: IP Shortcuts configuration	11
IP Shortcuts configuration fundamentals.....	11
SPBM IP shortcuts.....	11
SPBM IPv6 Shortcuts.....	13
ECMP with IS-IS.....	16
IS-IS IP redistribution policies.....	19
IS-IS accept policies.....	26
IP Shortcuts configuration using the CLI.....	30
Configuring SPBM IPv4 Shortcuts.....	30
Configuring SPBM IPv6 Shortcuts.....	34
Configuring IS-IS Accept Policies.....	36
Viewing IS-IS accept policy information.....	46
IP Shortcuts configuration using EDM.....	49
Configuring a Circuitless IPv4 interface.....	49
Configuring a Circuitless IPv6 interface.....	49
Configuring IS-IS redistribution.....	50
Configuring SPBM IP shortcuts.....	51
Enabling IPv6 IS-IS redistribution.....	53
Applying IS-IS accept policies globally.....	54
Configuring an IS-IS accept policy for a specific advertising BEB.....	55
Configuring an IS-IS accept policy to apply for a specific I-SID.....	56
Configuring an IS-IS accept policy for a specific advertising BEB and I-SID.....	57
Configuring an I-SID list for an IS-IS accept policy.....	58
Configuring an IS-IS accept policy for a specific I-SID list.....	59
Configuring an IS-IS accept policy for a specific advertising BEB and I-SID-list.....	60
IP Shortcuts SPBM configuration example.....	61
Chapter 5: Layer 3 VSN configuration	65
Layer 3 VSN configuration fundamentals.....	65
SPBM Layer 3 VSN.....	65

Fabric Connect Service Types.....	67
ECMP with IS-IS.....	68
IS-IS redistribution policies.....	71
IS-IS accept policies.....	77
Enable/disable ICMP Response on VRFs/Layer 3 VSNs.....	82
Layer 3 VSN configuration using the CLI.....	82
Configuring SPBM IPv4 Layer 3 VSN.....	82
Configuring IS-IS Accept Policies.....	85
Configuring inter-VRF accept policies on VRFs.....	95
Configuring SPBM IPv6 Layer 3 VSN.....	98
Viewing IS-IS accept policy information.....	100
Displaying SPBM IPv6 unicast Forwarding Information Base.....	102
Layer 3 VSN configuration using EDM.....	103
Configuring SPBM IPv4 Layer 3 VSN.....	103
Configuring IS-IS redistribution.....	104
Applying IS-IS accept policies globally.....	106
Configuring an IS-IS accept policy for a specific advertising BEB.....	107
Configuring an IS-IS accept policy for a specific advertising BEB and I-SID.....	108
Configuring an I-SID list for an IS-IS accept policy.....	109
Configuring an IS-IS accept policy for a specific I-SID list.....	110
Configuring an IS-IS accept policy for a specific advertising BEB and I-SID-list.....	111
Configuring SPBM IPv6 Layer 3 VSN.....	112
Layer 3 VSN configuration example.....	113
Verifying Layer 3 VSN operation.....	117
Chapter 6: Layer 3 video surveillance basic configuration.....	121
Layer 3 Video Surveillance install script.....	121
Running the Layer 3 Video Surveillance install script.....	123
Appendix A: SPBM reference architectures.....	126
Reference architectures.....	126
Solution-specific reference architectures.....	135
Glossary.....	141

Chapter 1: Preface

Purpose

This document provides information on features in VSP Operating System Software (VOSS). VOSS runs on the following product families:

- Extreme Networks Virtual Services Platform 4000 Series
- Extreme Networks Virtual Services Platform 7200 Series
- Extreme Networks Virtual Services Platform 8000 Series (includes VSP 8200 and VSP 8400 Series)
- Extreme Networks Virtual Services Platform 8600

This document provides information and instructions to configure Fabric Layer 3 services on the switch.

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.

Training

Ongoing product training is available. For more information or to register, you can access the Web site at www.extremenetworks.com/education/.

Providing Feedback to Us

We are always striving to improve our documentation and help you work better, so we want to hear from you! We welcome all feedback but especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team about this document, please contact us using our short [online feedback form](#). You can also email us directly at internalinfodev@extremenetworks.com

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

- [GTAC \(Global Technical Assistance Center\) for Immediate Support](#)
 - Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact
 - Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- [GTAC Knowledge](#) – Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- [The Hub](#) – A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- [Support Portal](#) – Manage cases, downloads, service contracts, product licensing, and training and certifications.

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Extreme Networks Documentation

To find Extreme Networks product guides, visit our documentation pages at:

Current Product Documentation	www.extremenetworks.com/documentation/
Archived Documentation (for previous versions and legacy products)	www.extremenetworks.com/support/documentation-archives/
Release Notes	www.extremenetworks.com/support/release-notes

Open Source Declarations

Some software files have been licensed under certain open source licenses. More information is available at: www.extremenetworks.com/support/policies/software-licensing.

Subscribing to Service Notifications

Subscribe to receive an email notification for product and software release announcements, Vulnerability Notices, and Service Notifications.

About this task

You can modify your product selections at any time.

Procedure

1. In an Internet browser, go to <http://www.extremenetworks.com/support/service-notification-form/>.
2. Type your first and last name.
3. Type the name of your company.
4. Type your email address.
5. Type your job title.
6. Select the industry in which your company operates.
7. Confirm your geographic information is correct.
8. Select the products for which you would like to receive notifications.
9. Click **Submit**.

Chapter 2: New in this Document

The following sections detail what is new in *Configuring Fabric Layer 3 Services*.

New Content

[ECMP Impact on IS-IS Route Selection for Inter-VRF Routes with vIST](#) on page 18 is added to the document.

Notice about Feature Support

This document includes content for multiple hardware platforms across different software releases. As a result, the content can include features not supported by your hardware in the current software release.

If a documented command, parameter, tab, or field does not appear on your hardware, it is not supported.

For information about feature support, see *Release Notes*.

For information about physical hardware restrictions, see your hardware documentation.

Chapter 3: SPBM and IS-IS configuration workflow

The following section describes the generic work flow to configure SPBM and IS-IS infrastructure and services on your network.

 **Note:**

This section is an overview. For further details on the SPBM and IS-IS infrastructure and configuration, see the documents described in the Documentation sources section below.

1. Infrastructure configuration:

As a first step, you must configure your basic infrastructure for Shortest Path Bridging MAC (SPBM).

2. Services configuration:

After you complete the infrastructure configuration, you configure the appropriate services for your network to run on top of your base architecture. This includes:

- Layer 2 and Layer 3 VSNs
- IP Shortcuts
- Inter-VSN routing

3. Operations and Management:

To debug connectivity issues and isolate network faults in the SPBM network, you can use Connectivity Fault Management (CFM).

Chapter 4: IP Shortcuts configuration

This section provides concepts and procedures to configure IP Shortcuts.

IP Shortcuts configuration fundamentals

This section provides fundamental concepts for IP Shortcuts.

Fabric Connect supports both IPv4 Shortcuts and IPv6 Shortcuts. Because IPv6 Shortcuts depend on IPv4 Shortcuts, you should understand how IPv4 Shortcuts work (see [SPBM IP shortcuts](#) on page 11) before jumping to the IPv6 section.

SPBM IP shortcuts

In addition to Layer 2 virtualization, the SPBM model is extended to also support Routed SPBM, otherwise called SPBM IP Shortcuts.

Unlike Layer 2 VSN, with SPBM IP shortcuts, no I-SID configuration is required. Instead, SPBM nodes propagate Layer 3 reachability as “leaf” information in the IS-IS LSPs using Extended IP reachability TLVs (TLV 135), which contain routing information such as neighbors and locally configured subnets. SPBM nodes receiving the reachability information can use this information to populate the routes to the announcing nodes. All TLVs announced in the IS-IS LSPs are grafted onto the shortest path tree (SPT) as leaf nodes.

The following figure shows a network running SPBM IP shortcuts.




Figure 1: SPBM IP Shortcuts

In this example, BEB A receives a packet with a destination IP address in the subnet of VLAN 14 and knows to forward the packet to BEB D based on the IP route propagation within IS-IS. After a route lookup, BEB A knows that BEB D is the destination for the subnet and constructs a new B-MAC header with destination B-MAC: D. BCBs B and C need only perform normal Ethernet switching to forward the packet to BEB D. A route lookup is only required once, at the source BEB, to identify BEB D as the node that is closest to the destination subnet.

In contrast to IP routing or Multiprotocol Label Switching (MPLS), SPBM IP shortcuts provide a simpler method of forwarding IP packets in an Ethernet network using the preestablished Ethernet FIBs on the BEBs. SPBM allows a network to make the best use of routing and forwarding techniques, where only the BEBs perform an IP route lookup and all other nodes perform standard Ethernet switching based on the existing SPT. This allows for end to end IP-over-Ethernet forwarding without the need for ARP, flooding, or reverse learning.

In the above example, the SPBM nodes in the core that are not enabled with IP shortcuts can be involved in the forwarding of IP traffic. Since SPBM nodes only forward on the MAC addresses that comprise the B-MAC header, and since unknown TLVs in IS-IS are relayed to the next hop but ignored locally, SPBM nodes need not be aware of IP subnets to forward IP traffic.

With IP shortcuts, there is only one IP routing hop, as the SPBM backbone acts as a virtualized switching backplane.

The following figure shows a sample campus network implementing SPBM IP shortcuts.




Figure 2: SPBM IP shortcuts in a campus

To enable IP shortcuts on the BEBs, you must configure a circuitless IP address (loopback address) and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 135.

In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct, static, OSPF, RIP, or BGP routes into IS-IS. To advertise IPv6 routes from the BEBs into the SPBM network, you must enable route redistribution of IPv6 direct, IPv6 static, and OSPFv3 routes into IS-IS.

SPBM IPv6 Shortcuts

Both IPv4 and IPv6 Shortcuts use IS-IS as the Interior Gateway Protocol (IGP) and the link state packet (LSP) for reachability information. However, IPv4 Shortcuts use TLV 135 and IPv6 Shortcuts

use TLV 236. All TLVs announced in the IS-IS LSPs are grafted onto the shortest path tree (SPT) as leaf nodes. IS-IS transports the IPv6 reachability information to remote BEBs and uses the shortest path, calculated by SPBM, for data forwarding.

*** Note:**

You only configure the IPv6 address information on the edges. There is no IPv6 in the SPBM cloud.

IS-IS transports the IPv6 routes through TLV 236 in the LSP advertisements. These routes are installed in the Global Routing Table (GRT) with the node from which the LSPs carrying the IPv6 routes are received as the next hop.

IPv6 Shortcuts dependency on IPv4 Shortcuts

IPv6 Shortcuts function in a very similar manner to IPv4 Shortcuts and depends on IPv4 Shortcuts for some functions. For example, IPv6 Shortcuts use the BMAC (local and remote) information created by IPv4 Shortcuts.

! Important:

IPv4 Shortcuts must be enabled before you enable IPv6 Shortcuts.

An error is displayed if you try to enable IPv6 Shortcuts but do not have IPv4 Shortcuts already enabled.

IPv6 Shortcuts alone can be disabled while leaving IPv4 Shortcuts enabled. When IPv4 Shortcuts is disabled without disabling IPv6 Shortcuts disabled first, a warning or error message is displayed indicating that IPv6 should be disabled first.

Circuitless IPv6 (CLIPv6)

To enable IPv6 Shortcuts on the BEBs and to advertise the local BEB to other IS-IS nodes, you must configure a circuitless IPv6 address (loopback address) and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 236.

IPv6 Shortcuts support Circuitless IPv6 (CLIPv6), which ensures uninterrupted connectivity to the switch as long as there is an actual path to reach it. This route always exists and the circuit is always up because there is no physical attachment.

Migrating the GRT to IPv6 Shortcuts

Use the following steps to migrate the Global Router Table (GRT) to use IPv6 Shortcuts over the SPBM core:

- Identify the nodes that should be enabled with IPv6 Shortcuts. Apply these steps to all of these nodes.
- Activate and validate basic IPv6 Shortcuts. For information, see [SPBM IPv6 Shortcuts](#) on page 13.
- Configure IS-IS route preference to ensure that the IPv6 IGP protocol currently being used in the SPBM core is preferred over the IS-IS routes.

- Enable redistribution of direct and static IPv6 routes into IS-IS.
- Create route policies to permit only IPv6 IGP routes from the access side of the SPBM network.
- Configure redistribution of routes from the IPv6 route table from each of the IPv6 IGP protocols into IS-IS along with the appropriate route policy.
- Use the `show isis spbm ipv6-unicast-fib` command to check the IS-IS LSDB, IS-IS routes, and to verify that all the desired IPv6 routes are now in IS-IS.
- Configure redistribution of IS-IS routes from the IPv6 route table into each of the IPv6 IGP protocols in use. This redistribution does not require a route policy since IS-IS is only supported in the SPBM core.
- Change IS-IS route-preference to ensure that IS-IS routes are preferred over other IPv6 IGP routes.
- Disable/delete old IPv6 IGP in the SPBM core.

! Important:

Use only one IPv6 routing protocol in the SPBM core to prevent the possibility of routing loops.

IPv6 Shortcut limitations and considerations

The following features are not supported:

- Disabling and enabling alternate routes for IPv6 routes
- Redistribution of RIP into IS-IS
- BGP+ and its redistribution into IS-IS
- 6-in-4 tunnels are not supported when the tunnel destination IP is reachable via IPv4 Shortcuts route.
- IS-IS accept policies are applicable to IPv4 Shortcut routes only, not IPv6 Shortcut routes.
- Route policies are not supported with redistribution of IPv6 routes.

Keep the following considerations in mind when configuring IPv6 Shortcuts:

- IPv4 Shortcuts must be enabled before enabling IPv6 Shortcuts.
- IPv6 Shortcuts support Circuitless IPv6 (CLIPv6) with the following limitations:
 - Stateless address autoconfiguration (SLAAC) is not supported on IPv6 CLIP interfaces.
 - IPv6 CLIP does not support link-local address configuration.
 - To configure an IPv6 address with a prefix length from 65 to 127 on a CLIP interface, you must enable the IPv6 mode flag.

*** Note:**

This limitation does not apply to VSP 4000 switches.

- Neighbor discovery (ND) does not run on an IPv6 CLIP interface. Therefore, the system does not detect when you configure a duplicate IPv6 address.

- Multiple IPv6 address configuration on an IPv6 CLIP interface is not supported.
- You can configure a maximum of 64 IPv6 CLIP interfaces.
- IPv6 CLIP interface is enabled by default and it cannot be disabled.
- IPv6 with vIST provides the same support as IPv4 with vIST.
- To help with debugging, CFM provides full support for both IPv4 and IPv6 addresses for the `12ping` and `12tracert` commands.

ECMP with IS-IS

The Equal Cost Multipath (ECMP) feature supports and complements the IS-IS protocol.

With ECMP, the switch can determine multiple equal-cost paths to the same destination prefix.

You can use multiple paths for load sharing of traffic. These multiple paths allow faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use your links between routers more efficiently when sending IP and IPv6 traffic. Equal Cost Multipath is formed using routes from the same protocol.

The number of multiple paths a switch can support differs by hardware platform. For more information about feature support, see .

ECMP within IS-IS routes

Equal Cost Multipath (ECMP) allows the device to determine up to eight equal cost paths to the same destination prefix. The maximum number of equal cost paths you can configure depends on the hardware platform. For more information, see *Release Notes*.

If the device learns the same route from multiple sources, the information is ECMP only if the routes:

- are from the same VSN
- have the same SPBM cost
- have the same prefix cost
- have the same IP route preference

Multiple BEBs can announce the same route, either because the Layer 2 LAN connects to multiple BEBs for redundancy, or because segments of the LAN are Layer 2 bridged. In Layer 2, if the device has to tie-break between multiple sources, the tie-breaking is based on cost and hop count.

In Layer 3, hop count is not used for tie-breaking. Instead, the device uses the following precedence rules to tie-break. In the following order, the device prefers:

1. Routes that do not include nodes with the overload bit set.

When a router node runs out of system resources (memory or CPU), it alerts the other routers in the network by setting the overload bit in its link-state packets (LSPs). When this bit is set, the node is not used for transit traffic but only for traffic packets destined to the node's directly connected networks and IP prefixes.

2. Local routes over remote routes.

If a route is learned locally, for example, through inter-VRF route leaking, it is most preferred.

3. Routes with the lowest route preference.

By default, IS-IS routes within the same VSN are added to the LSDB with a default preference of 7. Inter-VSN routes are added to the LSDB with a route preference of 200. You can however, change the route preference using IS-IS accept policies.

4. Metric type internal (type 1) over metric type external (type 2).
5. Routes with the lowest SPBM cost.
6. Routes with the lowest prefix cost.

If the metric type is internal, then the tie-break is on SPB cost first, and then on the prefix cost. Otherwise the tie-break is only on the prefix cost.

You can either change this using a route-map on the remote advertising node with the **redistribute** command, or using a route-map on the local node with the IS-IS accept policy.

7. Routes within a VSN with a lower Layer 3 VSN I-SID.

The device considers the Global Routing Table (GRT) to have an I-SID equal to zero.

When you use multiple B-VLANs in the SPBM core, multiple paths exist to reach a particular SPBM node, one on each B-VLAN; therefore, any IP prefix or IPv6 prefix that the device receives from a BEB results in multiple ECMP paths. These paths may or may not be physically diverse. SPBM supports up to two B-VLANs; a primary B-VLAN and a secondary B-VLAN.

If more ECMP paths are available than the configured number of paths, then the device adds the routes using the following order: The device selects all routes from the primary B-VLAN and orders the routes learned through that B-VLAN from lowest system ID to the highest IS-IS system ID, then the device moves on to select all routes from the secondary B-VLAN, ordering those routes from lowest IS-IS system ID to the highest IS-IS system ID until you reach the number of equal paths configured.

For example, consider an SPB core configured with two B-VLANs (primary B-VLAN 1000 and secondary B-VLAN 2000), and the device learns routes from two BEBs called BEB-A (with a lower IS-IS system ID) and BEB-B (with a higher IS-IS system ID), then the order in which the next-hop is chosen for those routes are as follows.

If a route is learned only from BEB-A with the maximum number of allowed ECMP paths configured as 8 (default), then the order in which the next-hop is chosen for that route is:

1. BEB-A B-VLAN 1000
2. BEB-A B-VLAN 2000

If routes are learned from both BEB-A and BEB-B with maximum number of allowed ECMP paths configured as 8 (default), then the order in which the next-hop is chosen for those routes are:

1. BEB-A B-VLAN 1000
2. BEB-B B-VLAN 1000
3. BEB-A B-VLAN 2000
4. BEB-B B-VLAN 2000

If ECMP is disabled, the maximum number of allowed ECMP paths is 1 and the device adds the route from the lowest system ID with the primary B-VLAN. In this example, the device adds BEB-A B-VLAN 1000.

*** Note:**

- ECMP is supported for IPv6 Shortcut routes.
- To add IS-IS equal cost paths in the routing table, you must enable IPv6 ECMP feature globally.

ECMP Impact on IS-IS Route Selection for Inter-VRF Routes with vIST

This section illustrates the impact ECMP can have on a configuration that implements user-defined VRFs in a vIST cluster and how to avoid incorrect route selection.

Understanding the Configuration

Imagine the following configuration:

- A vIST cluster exists with multiple VRF contexts.
- On both nodes, VRF A redistributes routes into IS-IS as external. VRF B uses an IS-IS accept policy to accept these routes.
- Each node learns three paths to the route:
 - The nodes learn one path using local inter-VRF redistribution.
 - The nodes learn the other two paths from the IST peer.
- The routes are treated as ECMP paths because the preference, metric-type, and metric are equal.

IS-IS sorts paths for the same route by source-BEB B-MAC and B-VLAN ID. The primary B-VLAN ID is first installed for each B-MAC, followed by the secondary B-VLAN ID for each B-MAC, as long as the ECMP max-path value is not reached. On the node with the lowest B-MAC, the first path listed is its own local inter-vrf route, while on the other node, the MIM path across the vIST is listed first.

If you disable ECMP, all but the first path is removed. Because IS-IS orders the paths by B-MAC, each node in the vIST cluster selects the same B-MAC as the nexthop. This configuration leads one of the nodes to select itself, the local inter-vrf route, while the other node selects the MIM path across the vIST to get to the inter-vrf route. This situation results in an incorrect route selection.

Avoiding Incorrect Route Selection

To avoid this situation, create a policy to prevent IS-IS from determining that the MIM path across the vIST and the local inter-VRF route are ECMP paths. Configure the local inter-VRF path as the preferred path, and the vIST path as the backup. The following list identifies way that you can accomplish this:

- Redistribute the VRF route into IS-IS using the internal metric-type. IS-IS will always select the local inter-VRF route. For more information about the metric type for IS-IS routes, see *Configuring Fabric Basics and Layer 2 Services*.
- If an IS-IS internal metric-type is not an option, configure an IS-IS accept policy to change the preference of inter-VRF routes learned from the IST peer. The local inter-VRF route is preferred over the inter-VRF routes learned from the IST peer.

IS-IS IP redistribution policies

When you connect an SPBM core using IP shortcuts to existing networks running a routing protocol such as OSPF or RIP, a redundant configuration requires two VSP switches:

- One router redistributes IP routes from Routing Information Protocol (RIP)/Open Shortest Path First (OSPF) into IS-IS (IP).
- The second router redistributes from IS-IS (IP) into RIP or OSPF.

The following figure illustrates this configuration.




Figure 3: Redundant OSPF or RIP network

In this scenario it is necessary to take extra care when redistributing through both VSP switches. By default the preference value for IP routes generated by SPBM-IP (IS-IS) is 7. This is a higher preference than OSPF (20 for intra-area, 25 for inter-area, 120 for ext type1, 125 for ext type2) or RIP (100).

! Important:

The lower numerical value determines the higher preference.

In the preceding diagram both nodes (VSP 8200G and VSP 8200D) have an OSPF or a RIP route to 192.168.10.0/24 with the next-hop to VSP 8200A.

As soon as the VSP 8200G node redistributes that IP route into IS-IS, the VSP 8200D node learns the same route through IS-IS from VSP 8200G. (The VSP8200G node already has the route through OSPF or RIP). Because IS-IS has a higher preference, VSP 8200D replaces its 192.168.10.0 OSPF route with an IS-IS one that points at VSP 8200G as the next-hop. The following figure illustrates this scenario.




Figure 4: Redistributing routes into IS-IS

Clearly this is undesirable and care needs to be taken to ensure that the two redistributing nodes (VSP 8200G and VSP 8200D) do not accept redistributed routes from each other. With IS-IS accept policies, you can associate an IS-IS accept policy on VSP 8200D to reject all redistributed IP routes received from VSP 8200G, and VSP 8200G to reject all redistribute IP routes from VSP 8200D.

*** Note:**

IS-IS accept policies do not apply to IPv6 shortcut routes and only apply to IPv4 shortcut routes.

An alternate way to solve the preceding problem with existing functionality is to reverse the problem by lowering the SPBM-IP (IS-IS) preference by configuring it to a value greater than RIP (100) or OSPF (20,25,120,125). For example, log on to Global Configuration mode and use the following command to configure a preference of 130:

```
ip route preference protocol spbm-level1 130
```

*** Note:**

For IPv6, the command is `ipv6 route preference protocol spbm-level1 130`

Now that the OSPF or RIP routes have a higher preference than SPBM-IP (IS-IS), the above problem is temporarily solved. However, the same issue resurfaces when the IS-IS IP routes are redistributed into OSPF or RIP in the reverse direction as shown in the following figure for OSPF:

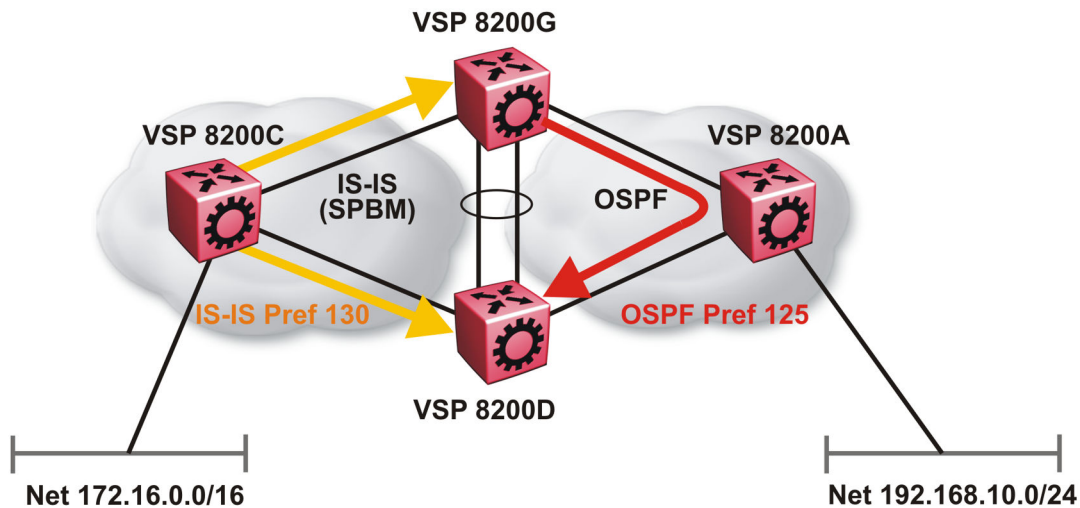


Figure 5: Redistributing routes into OSPF

In the preceding figure, both VSP 8200G and VSP 8200D have an IS-IS IP route for 172.16.0.0/16 with the next hop as VSP 8200C. As soon as the VSP 8200G redistributes the IS-IS route into OSPF, the VSP 8200D node learns that same route through OSPF from VSP 8200G. (The VSP 8200G node already has the route through IS-IS).

Because OSPF has a higher preference, VSP 8200D replaces its 172.16.0.0/16 IS-IS route with an OSPF one. (Note that the 172.16.0.0/16 route will be redistributed into OSPF as an AS external route, hence with preference 120 or 125 depending on whether type1 or type2 was used). In this case, however, you can leverage OSPF Accept policies, which can be configured to prevent VSP 8200D from accepting any AS External (LSA5) routes from VSP 8200G and prevent VSP 8200G from accepting any AS External (LSA5) routes from VSP 8200D. The following is a sample configuration:

```
enable
configure terminal
route-map

IP ROUTE MAP CONFIGURATION - GlobalRouter

route-map "reject" 1
no permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

OSPF CONFIGURATION - GlobalRouter

router ospf enable

OSPF ACCEPT CONFIGURATION - GlobalRouter

router ospf
accept adv-rtr {A.B.C.D}
```

```
accept adv-rtr {A.B.C.D} enable route-map "reject"  
exit
```

*** Note:**

Disable alternative routes by issuing the command `no ip alternative-route` to avoid routing loops on the SMLT Backbone Edge Bridges (BEBs).

In the preceding figure, if VSP 8200A advertises 25000 OSPF routes to VSP 8200G and VSP 8200D, then both VSP 8200G and VSP 8200D install the 25000 routes as OSPF routes. Since VSP 8200D and VSP 8200G have OSPF to IS-IS redistribution enabled, they also learn these 25000 routes as IS-IS routes. IS-IS route preference is configured with a higher numerical value (130) than the OSPF route preference (125), so VSP 8200D and VSP 8200G keep IS-IS learned routes as alternative routes.

If VSP 8200A withdraws its 25000 OSPF routes, VSP 8200G and VSP 8200D remove the OSPF routes. While the OSPF routes are removed the routing tables of VSP 8200G and VSP 8200D activate the alternative IS-IS routes for the same prefix. Since VSP 8200G and VSP 8200D have IS-IS to OSPF redistribution enabled, VSP 8200A learns these routes as OSPF and this causes a routing loop. Use the `no ip alternative-route` command to disable alternative routes on VSP 8200G and VSP 8200D to avoid routing loops.

In the preceding figure, you leveraged OSPF Accept policies, which can be configured to prevent VSP 8200D from accepting any AS External (LSA5) routes from VSP 8200G and prevent VSP 8200G from accepting any AS External (LSA5) routes from VSP 8200D. In the case of a RIP access network, the preceding solution is not possible because RIP has no concept of external routes and no equivalent of accept policies. However, if you assume that a RIP network acts as an access network to an SPBM core, then it is sufficient to ensure that when IS-IS IP routes are redistributed into RIP they are aggregated into a single default route at the same time. The following figure and sample configuration example illustrates this scenario:




Figure 6: Redistributing routes into RIP

VSP 8200G

```

IP PREFIX LIST CONFIGURATION - GlobalRouter
ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION - GlobalRouter

route-map "inject-default" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "match-network" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "set-injectlist" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

RIP PORT CONFIGURATION

interface gigabitethernet 1/11
ip rip default-supply enable
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

```

IP Shortcuts configuration

```
router rip
redistribute isis
redistribute isis metric 1
redistribute isis route-map "inject-default"
redistribute isis enable
exit
```

IP REDISTRIBUTE APPLY CONFIGURATIONS

```
ip rip apply redistribute isis
```

VSP 8200A

RIP PORT CONFIGURATION

```
interface gigabitethernet 1/2
ip rip default-listen enable
exit
interface gigabitethernet 1/3
ip rip default-listen enable
exit
```

VSP 8200D

IP PREFIX LIST CONFIGURATION - GlobalRouter

```
ip prefix-list "default" 0.0.0.0/0 ge 0 le 32
```

IP ROUTE MAP CONFIGURATION - GlobalRouter

```
route-map "inject-default" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit
```

```
route-map "match-network" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit
```

```
route-map "set-injectlist" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit
```

RIP PORT CONFIGURATION

```
interface gigabitethernet 2/11
ip rip default-supply enable
exit
```

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

```
router rip
redistribute isis
redistribute isis metric 1
redistribute isis route-map "inject-default"
redistribute isis enable
exit
```

IP REDISTRIBUTE APPLY CONFIGURATIONS

```
ip rip apply redistribute isis
```


You can control the propagation of the default route on the RIP network so that both VSP 8200G and VSP 8200D supply the default route on their relevant interfaces, and not accept it on the same interfaces. Likewise, VSP 8200A will accept the default route on its interfaces to both VSP8200G and VSP8200D but it will not supply the default route back to them. This will prevent the default route advertised by VSP8200G from being installed by VSP8200D, and vice-versa.

The preceding example where IS-IS IP routes are aggregated into a single default route when redistributed into the RIP network also applies when redistributing IS-IS IP routes into OSPF if that OSPF network is an access network to an SPBM core. In this case use the following redistribution policy configuration as an example for injecting IS-IS IP routes into OSPF:

```
IP PREFIX LIST CONFIGURATION - GlobalRouter

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION - GlobalRouter

route-map "inject-default" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "match-network" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

route-map "set-injectlist" 1
permit
enable
match protocol local|static|rip|ospf|ebgp|ibgp|dvmrp|isis
exit

OSPF CONFIGURATION - GlobalRouter

router ospf enable
router ospf
as-boundary-router enable
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

router ospf
redistribute isis
redistribute isis route-map "inject-default"
redistribute isis enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip ospf apply redistribute isis
```

IS-IS accept policies

You can use Intermediate-System-to-Intermediate-System (IS-IS) accept policies to filter incoming IS-IS routes over the SPBM cloud and apply route policies to the incoming IS-IS routes. IS-IS accept policies enable the device to determine whether to add an incoming route to the routing table.

IS-IS accept policies and DvR

When you configure DvR in an SPB network, you can leverage IS-IS accept policies to control the DvR routes learned from the DvR backbone. The DvR backbone contains the master list of all the host routes learned from various DvR domains.

You can configure accept policies on a DvR Controller or a non-DvR BEB as a filter to determine which DvR host routes to accept into the routing table, from the DvR backbone. Accept policies apply to only those backbone (or inter-domain) host routes that are not part of the Controller's own DvR enabled subnets *and* do not have the same domain ID as that of the Controller.

For non-DvR BEBs, all the routes present in the backbone are learned, but you can still use the accept policies to filter specific routes.

For information on DvR, see *Configuring IPv4 Routing*.

IS-IS accept policy filters

You can filter traffic with IS-IS accept policies by:

- advertising BEB
- I-SID or I-SID list
- route-map
- backbone-route-map
- a combination of route-map and backbone-route-map

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

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

IS-IS accept policies can also apply route policies to determine what incoming traffic to accept into the routing table. With route policies the device can determine which routes to accept into the routing table based on the criteria you configure. You can match on the network or the route metric.

On DvR Controllers in a DvR domain, you can configure a backbone route policy to determine what host routes to accept from the DvR backbone, into the routing table. Also, just like on the route policy, you can configure match criteria, and set preferences on the backbone route policy.

To accept both IS-IS routes and host routes from the DvR backbone, you can configure both a route policy and a backbone route policy in the accept policy instance.

For more information on configuring route policies, see *Configuring IPv4 Routing*.

The following table describes IS-IS accept policy filters.

Filters into	Filter	Description
Global Routing Table (GRT)	accept route-map <i>WORD</i> <1-64>	By default, the device accepts all routes into the GRT and VRF routing table. This is the default accept policy.
	accept route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	This is the default accept policy with configuration to accept specific DvR host routes from the DvR backbone.
	accept adv-rtr <x.xx.xx> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB defined by the SPBM nickname. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept i-sid <1-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the I-SID, which represents a local or remote Layer 3 VSN. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept adv-rtr <x.xx.xx> i-sid <1-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the list of I-SIDs. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept adv-rtr <x.xx.xx> isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
Virtual Routing and Forwarding (VRF) routing table	isis accept adv-rtr <x.xx.xx> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB defined by the SPBM nickname.

Table continues...

Filters into	Filter	Description
		The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept i-sid <0-16777215> route-map WORD<1-64> backbone-route-map WORD<1-64>	The device filters based on the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept adv-rtr <x.xx.xx> i-sid <0-16777215> route-map WORD<1-64> backbone-route-map WORD<1-64>	The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept isid-list WORD<1-32> route-map WORD<1-64> backbone-route-map WORD<1-64>	The device filters based on the list of I-SIDs to which the IS-IS accept policy applies. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept adv-rtr <x.xx.xx> isid-list WORD<1-32> route-map WORD<1-64> backbone-route-map WORD<1-64>	The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept route-map WORD<1-64> route-map WORD<1-64> backbone-route-map WORD<1-64>	The device filters based on the route policy. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.

IS-IS accept policies for the GRT and VRFs

You can create an IS-IS accept policy for incoming routes for the Global Routing Table (GRT), which accepts routes into the routing table, or for a Virtual Routing and Forwarding (VRF) instance, which accepts incoming routes to the routing table of the VRF.

If you create an IS-IS accept policy on the switch for either the GRT or a VRF that operates at a global default level, the accept policy applies to all routes for all BEBs in the GRT or VRF.

If you create an IS-IS accept policy on the switch for a specific advertising BEB for either the GRT or a VRF, the IS-IS accept policy instance applies for that specific advertising BEB. If you use a more specific filter, the system gives preference to the specific filter over the global default level.

IS-IS accept policies for inter-VRF route redistribution

You can also use the filter mechanism for IS-IS accept policies to redistribute routes between different VRFs, or between a VRF and the GRT. For inter-VRF route redistribution, you match the filter based on the I-SID, which represents the Layer 3 VSN context.

You can apply the filter at the global default level, where the IS-IS accept policy applies to all routes for that I-SID from all BEBs, or at a specific advertising BEB level, where the filter only applies to a specific advertising BEB. The device gives preference to a specific filter for a specific advertising BEB over the global default filter.

For inter-VRF route redistribution, an I-SID value of 0 represents the GRT. For inter-VRF route redistribution between VRFs, the I-SID is the source VRF (or remote VRF).

IS-IS accept policy considerations

Consider the following when you configure IS-IS accept policies:

- The switch does not support IS-IS accept policies for IPv6 addresses.
- If a VRF uses a different protocol to redistribute routes from another VRF, the IS-IS accept policy feature cannot be used. You can only use the IS-IS accept policy for inter-VSN route redistribution between VRFs.

Precedence rules in the same VSN

The following precedence rules apply for IS-IS accept policies used in the same VSN:

- You can only apply one configured IS-IS accept policy for each route.
- You can apply either a default filter for all advertising BEBs or a filter for a specific advertising BEB.
- If you disable the accept filter, the system ignores the filter and the filter with the next highest precedence applies.
- The device prefers the `accept adv-rtr` filter, which filters based on a specific advertising BEB, over the default filter for all advertising BEBs.
- The device accepts all routes within the same VSN by default. You can apply a route policy to filter or change the characteristics of the route by metric or preference.
- The `i-sid` or `isid-list` filters are not valid for routes within the same VSN.

Precedence rules for inter-VSN route redistribution

The following precedence rules apply for IS-IS accept policies used for inter-VSN route redistribution:

- You can only apply one configured IS-IS accept policy for each route.
- You can apply filters at a global default level for all BEBs for a specific I-SID or I-SID list, or you can apply filters for a specific advertising BEB for a specific I-SID or I-SID list.
- If you disable the accept filter, the system ignores the filter and the filter with the next highest precedence applies.

- The device requires a specific filter to redistribute routes between VSNs through the use of the `i-sid` or `isid-list` filters.
- The `i-sid` filter takes precedence over the `isid-list` filter.
- The `adv-rtr` filter for a specific advertising BEB takes precedence over a filter with the same `i-sid` filter without the `adv-rtr` filter.
- The `i-sid` or `isid-list` filters only apply to routes for inter-VSN route redistribution.
- If multiple `isid-list` filters have the same I-SID within the list, the first on the list alphabetically has the higher precedence.

Route preference

The relative value of the route preference among different protocols determines which protocol the device prefers. If multiple protocols are in the routing table, the device prefers the route with the lower value. You can change the value at the protocol level, and you can also change the preference of incoming ISIS routes using the route-map with the ISIS Accept policy filter.

Route metric

Use route-map to change the metric of a route when you accept a remote IS-IS route with IS-IS accept policies.

You can use route-map to change the metric of a route when you redistribute the route from another protocol to IS-IS through the route redistribution mechanism.

You can also configure the route metric with the base `redistribute` command without the use of route-map.

For more information on the configuration of route-map, see *Configuring IPv4 Routing*.

IP Shortcuts configuration using the CLI

This section provides procedures to configure IP Shortcuts using the CLI.

Configuring SPBM IPv4 Shortcuts

In addition to Layer 2 virtualization, the SPBM model is extended to also support Routed SPBM, otherwise called SPBM IP Shortcuts.

SPBM allows a network to make the best use of routing and forwarding techniques, where only the BEBs perform an IP route lookup and all other nodes perform standard Ethernet switching based on the existing shortest path tree. This allows for end to end IP-over-Ethernet forwarding without the need for ARP, flooding, or reverse learning.

To enable IP shortcuts on the BEBs, you must configure a circuitless IP (CLIP) address (loopback address), and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 135. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct and static routes into IS-IS.

*** Note:**

The loopback address on each switch or BEB must all be in different subnets to ensure connectivity between them. To do this, use a 32-bit mask with the CLIP address.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Before redistributing routes into IS-IS, you must create the Customer VLANs, add slots/ports, and add the IP addresses and network masks.

Procedure

1. Enter Loopback Interface Configuration mode


```
enable
configure terminal
interface Loopback <1-256>
```
2. Configure a CLIP interface to use as the source address for SPBM IP shortcuts:


```
ip address [<1-256>] <A.B.C.D/X>
```
3. Exit the Loopback Interface Configuration mode to Global Configuration mode:


```
exit
```
4. Log on to IS-IS Router Configuration mode:


```
router isis
```
5. Specify the CLIP interface as the source address for SPBM IP shortcuts:


```
ip-source-address <A.B.C.D>
```
6. Configure SPBM IP shortcuts:


```
spbm <1-100> ip enable
```
7. Display the status of SPBM IP shortcuts on the switch:


```
show isis spbm
```
8. Identify routes on the local switch to be announced into the SPBM network:


```
redistribute {bgp | direct | ospf | rip | static}
```
9. Enable routes to be announced into the SPBM network


```
redistribute {bgp | direct | ospf | rip | static} enable
```
10. If you want to delete the configuration, use the no option:


```
no redistribute {bgp | direct | ospf | rip | static}
no redistribute {bgp | direct | ospf | rip | static} enable
```
11. Exit to Global Configuration mode:

```
exit
```

12. Apply the configured redistribution:

```
isis apply redistribute {bgp | direct | ospf | rip | static | vrf  
WORD<1-16>}
```

Example

```
Switch:1> enable  
Switch:1# configure terminal  
Switch:1(config)# interface loopback 1  
Switch:1(config-if)# ip address 192.0.2.2/8  
Switch:1(config-if)# exit  
Switch:1(config)# router isis  
Switch:1(config-isis)# ip-source-address 192.0.2.2  
Switch:1(config-isis)# spbm 1 ip enable  
Switch:1(config-isis)# show isis spbm
```

```
show isis spbm
```

```
=====
```

ISIS SPBM Info							
SPBM INSTANCE	B-VID	PRIMARY VLAN	NICK NAME	LSDB TRAP	IP	IPV6	MULTICAST
1	4086-4087	4086	3.03.01	disable	enable	enable	disable

```
=====
```

ISIS SPBM SMLT Info			
SPBM INSTANCE	SMLT-SPLIT-BEB	SMLT-VIRTUAL-BMAC	SMLT-PEER-SYSTEM-ID
1	primary	00:00:03:03:03:03	0000.0303.0302

```
-----
```

Total Num of SPBM instances: 1

```
-----
```

```
Switch:1(config-isis)# redistribute rip  
Switch:1(config-isis)# redistribute rip enable  
Switch:1(config-isis)# exit  
Switch:1(config)# isis apply redistribute rip
```

Variable definitions

Use the data in the following table to use the **ip address** command.

Variable	Value
<1–256>	Specifies an interface ID value. This value is optional.
<A.B.C.D/X>	Specifies an IP address and subnet mask. Use the no option to delete the specified IP address.
<A.B.C.D>	Specifies an IP address. Use the no option to delete the specified IP address.

Use the data in the following table to use the **ip-source-address** command.

Variable	Value
<A.B.C.D>	Specifies the CLIP interface to use as the source address for SPBM IP shortcuts.

Use the data in the following table to use the **spbm** command.

Variable	Value
<1–100> ip enable	Enables or disables SPBM IP shortcut state. The default is disabled. Use the no or default options to disable SPBM IP shortcuts.

Use the data in the following table to use the **redistribute** command.

Variable	Value
{bgp direct ospf rip static}	Specifies the protocol.
enable	Enables the redistribution of the specified protocol into the SPBM network. The default is disabled. Use the no option to disable the redistribution.
metric <0–65535>	Configures the metric (cost) to apply to redistributed routes. The default is 1.
metric-type {external internal}	Configures the type of route to import into the protocol. The default is internal.
route-map WORD<0–64>	Configures the route policy to apply to redistributed routes. Type a name between 0 to 64 characters in length.
subnets {allow suppress}	Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose suppress to advertise subnets aggregated to their classful subnet. Choose allow to advertise the subnets individually with the learned or configured mask of the subnet. The default is allow.

Use the data in the following table to use the **isis apply redistribute** command.

Variable	Value
{bgp direct ospf rip static}	Specifies the protocol.

Configuring SPBM IPv6 Shortcuts

! Important:

You must enable IPv4 Shortcuts before you enable IPv6 Shortcuts because IPv6 Shortcuts depend on IPv4 Shortcuts for some functions.

Configuring IPv6 Shortcuts is essentially the same as the IPv4 procedure except you use the following IPv6 commands instead of their IPv4 equivalents:

- Use `ipv6 interface address` to create a CLIPv6 interface with an IPv6 address.
- Use `ipv6 ipv6-source-address` to specify the CLIPv6 interface as the source address for IPv6 Shortcuts.
- Use `spbm ipv6 enable` to enable IPv6 Shortcuts.
- Use `ipv6 redistribute {direct | static | ospf } enable` to control the redistribution of GRT IPv6 routes into the SPBM IS-IS domain.
- Use `ipv6 route preference protocol spbm-level1` to change route preference values for IPv6 Shortcut routes learned through IS-IS.

To enable IPv6 Shortcuts on the BEBs, you must configure a circuitless IPv6 (CLIPv6) address (loopback address), and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 236. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct and static routes into IS-IS.

* Note:

The loopback address on each switch or BEB must all be in different subnets to ensure connectivity between them. To do this, use a 32-bit mask with the CLIP address, and the CLIPv6 address prefix must be 128.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Before redistributing routes into IS-IS, you must create the Customer VLANs, add slots/ports, and add the IPv6 addresses and network masks.

Procedure

1. Enter Loopback Interface Configuration mode

```
enable
configure terminal
interface Loopback <1-256>
```

2. Configure a CLIPv6 interface to use as the source address for SPBM IPv6 Shortcuts:

```
ipv6 interface address WORD<0-255>
```

3. Exit the Loopback Interface Configuration mode to Global Configuration mode:

```
exit
```

4. Log on to IS-IS Router Configuration mode:

```
router isis
```

5. Specify the CLIPv6 interface as the source address for SPBM IPv6 Shortcuts:

```
ipv6-source-address WORD<0-46>
```

6. Enable SPBM IPv6 Shortcuts:

```
spbm <1-100> ipv6 enable
```

7. Display the status of SPBM IPv6 Shortcuts on the switch:

```
show isis spbm
```

8. Identify IPv6 routes on the local switch to be announced into the SPBM network.

- a. To redistribute routes for directly connected subnets:

```
ipv6 redistribute direct enable
```

- b. To redistribute static routes into IS-IS:

```
ipv6 redistribute static enable
```

- c. To redistribute OSPFv3 routes into IS-IS, enter IS-IS Router Configuration mode and then enable IS-IS redistribution:

```
router isis
```

```
ipv6 redistribute ospf enable
```

9. If you want to delete the configuration, use the no option:

```
no redistribute {direct | static | ospf} enable
```

10. Exit to Global Configuration mode:

```
exit
```

11. **(Optional)** Change route preference values for IPv6 Shortcut routes learned through IS-IS:

```
ipv6 route preference protocol spbm-level1 <0-255>
```

12. Apply the configured redistribution:

```
isis apply redistribute {bgp | direct | ospf | rip | static | vrf}
```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface loopback 123
Switch:1(config-if)#ipv6 interface address 123::1/128
Switch:1(config-if)#exit
```

IP Shortcuts configuration

```
Switch:1(config)#router isis
Switch:1(config-isis)#ipv6 ipv6-source-address <non-link-local ipv6-address>
Switch:1(config-isis)#spbm 1 ipv6 enable
Switch:1(config-isis)#show isis spbm
=====
                        ISIS SPBM Info
=====
SPBM      B-VID      PRIMARY   NICK      LSDB      IP        IPV6      MULTICAST
INSTANCE  VLAN
-----
1         4086-4087  4086     3.03.01  disable  enable  enable  disable
=====
                        ISIS SPBM SMLT Info
=====
SPBM      SMLT-SPLIT-BEB      SMLT-VIRTUAL-BMAC      SMLT-PEER-SYSTEM-ID
INSTANCE
-----
1         primary             00:00:03:03:03:03      0000.0303.0302
-----
Total Num of SPBM instances: 1
=====
```

Variable definitions

Use the data in the following table to use the IPv6 Shortcuts commands.

Variable	Value
ipv6-source-address <i>WORD</i> <0-46>	Specifies the source IPv6 address for locally generated IPv6 packets whose egress port is an SPBM NNI port. The <i>WORD</i> <0-46> value must be a locally configured loopback IPv6 address (CLIPv6). Use the no option to delete the specified IPv6 address.
spbm<1-100> ipv6 enable	Enables or disables SPBM IPv6 Shortcuts. The default is disabled. Use the no or default options to disable SPBM IPv6 Shortcuts.
ipv6 route preference protocol spbm-level1 <0-255>	Sets the route preference value for IPv6 Shortcut routes learned through IS-IS. The default preference is 7.
ipv6 redistribute {direct static ospf} enable	Specifies the GRT IPv6 route that you want to redistribute into the SPBM IS-IS domain. The default is disabled. Use the no option to disable the redistribution.

Configuring IS-IS Accept Policies

Use the following procedure to create and enable IS-IS accept policies to apply to routes from all Backbone Edge Bridges (BEBs) or to all routes from a specific BEB.

Use IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud. Accept policies apply to incoming traffic and determine whether to add the route to the routing table.

If DvR is enabled on your switch, and the switch is either a DvR Controller or a non-DvR BEB within the domain, you can configure IS-IS accept policies to accept specific host routes from the DvR backbone. For information on DvR, see *Configuring IPv4 Routing*.

IS-IS accept policies are disabled by default.

*** Note:**

- The `isis apply accept [vrf WORD<1-16>]` command can disrupt traffic and cause temporary traffic loss. After you apply `isis apply accept [vrf <1-16>]`, the command reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. You should make all the relevant accept policy changes, and then apply `isis apply accept [vrf WORD<1-16>]` at the end.
- If the route policy changes, you must reapply the IS-IS accept policy, unless the IS-IS accept policy was the last sequence in the configuration.
- The `isis apply accept [vrf WORD<1-16>]` command is not saved in the configuration file. If you use a saved configuration file for IS-IS accept policy configuration, you must apply the `isis apply accept [vrf WORD<1-16>]` command at the end.
- The number of unique Layer 3 VSN I-SIDs used on a BEB is limited to the number of VRFs supported on the switch. This includes the I-SID values used for Layer 3 VSNs and the I-SID values specified for the ISIS accept policy filters, which can be configured using the `ip isid-list [ISID#], accept i-sid <value>`, or `accept adv-rtr <isis nn> i-sid <value>` commands.

The switch supports 24 VRFs by default, so, in a default configuration, you cannot create an `ip isid-list` or accept policy with more than 24 unique I-SID entries. However, the configured VRFs take up an entry, so the formula to calculate the limit is: [24 VRF Limit – (currently configured VRFs)]. This gives the number of unique I-SIDs that can be used directly in the IS-IS accept policy filters, which you implement with the `ip isid-list` or `accept policy` command. The I-SIDs used for Layer 3 VSNs can be reused in IS-IS accept policy filters without affecting the limit.

If you increase the VRF scaling, you can create more Layer 3 VSNs. For more information about how to increase the number of supported VRFs, see *Configuring IPv4 Routing*. The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see *Release Notes*.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.
- Ensure that DvR is enabled on the switch before you configure an IS-IS accept policy with a backbone route policy, to accept host routes from the DvR backbone.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. **(Optional)** If you want to accept routes from a variety of I-SIDs, create an I-SID list before you create an IS-IS accept policy for the I-SID list:

```
ip isid-list WORD<1-32> [<1-16777215>][list WORD<1-1024>]
```

3. **(Optional)** Delete an I-SID list:

```
no ip isid-list WORD<1-32> [<1-16777215>][list WORD<1-1024>]
```

4. Enter IS-IS Router Configuration mode:

```
enable
configure terminal
router isis
```

Configure IS-IS accept policies with a route policy or a backbone route policy or a combination of both, to determine which routes the IS-IS accept policy applies to.

Configure one of the following types of IS-IS accept policies.

- **An IS-IS accept policy with only the route policy:**

The IS-IS routes are selectively accepted based on the route policy. Since the backbone route policy is not configured, all host routes from the DvR backbone are *denied*.

If you do not configure a route policy, by default, all IS-IS routes are *accepted*.

- **An IS-IS accept policy with only the backbone route policy:**

The DvR host routes from the DvR backbone are selectively accepted based on the backbone route policy. Since the route policy is not configured, all IS-IS host routes are accepted.

If you do not configure a backbone route policy, all host routes from the DvR backbone are *denied*.

- **An IS-IS accept policy with both route policy and backbone route policy:**

IS-IS routes are selectively accepted based on the route policy and host routes from the DvR backbone are selectively accepted based on the backbone route policy.

5. Configure an IS-IS accept policy instance with a route policy.

Use one of the following options:

- a. Create an IS-IS accept policy instance to apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>]
```

- b. Create an IS-IS accept policy instance to apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> [i-sid <1-16777215>][isid-list WORD
<1-32>]
```

- c. **(Optional)** Delete an IS-IS accept policy instance:

```
no accept [adv-rtr <x.xx.xx>][i-sid <1-16777215>][isid-list WORD
<1-32>]
```

- d. Specify an IS-IS route policy to apply to routes from all BEBs:

```
accept route-map WORD<1-64>
```

- e. Specify an IS-IS route policy to apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx>[route-map WORD<1-64>]
```

- f. **(Optional)** Delete an IS-IS route policy:

```
no accept [adv-rtr <x.xx.xx>] [route-map]
```

- g. Enable an IS-IS route accept instance:

```
accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-
list WORD<1-32>]
```

- h. **(Optional)** Disable an IS-IS route accept instance:

```
no accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-
list WORD<1-32>]
```

6. Configure an IS-IS accept policy instance with a backbone route policy to accept host routes from the DvR backbone:

*** Note:**

IS-IS accept policies typically apply to all IS-IS routes. However, to accept DvR host routes from the DvR backbone, you *must* explicitly configure the IS-IS accept policy with a backbone route policy.

Use one of the following options:

- a. Create the default IS-IS accept policy instance to accept host routes from the DvR backbone:

```
accept backbone-route-map WORD <1-64>
```

- b. **(Optional)** Delete the default IS-IS accept policy instance with backbone route policy configuration:

```
no accept backbone-route-map
```

- c. Create an IS-IS accept policy instance to accept host routes from the DvR backbone, and apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>] backbone-
route-map WORD<1-64>
```

- d. **(Optional)** Delete an IS-IS accept policy instance with backbone route policy configuration, which applies to all BEBs for a specific I-SID or I-SID list:

```
no accept [i-sid <1-16777215>][isid-list WORD <1-32>] backbone-
route-map
```

- e. Create an IS-IS accept policy instance to accept host routes from the DvR backbone and apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> backbone-route-map WORD <1-64>
```

- f. **(Optional)** Delete an IS-IS accept policy instance with backbone route policy configuration, which applies to a specific advertising BEB

```
no accept adv-rtr <x.xx.xx> backbone-route-map
```

7. Configure an IS-IS accept policy with both route policy and backbone route policy, to selectively accept IS-IS routes as well as host routes from the DvR backbone.

- a. Create the default IS-IS accept policy instance with a route policy to accept IS-IS routes and a backbone route policy to accept host routes from the DvR backbone:

```
accept route-map WORD<1-32> backbone-route-map WORD <1-64>
```

- b. **(Optional)** Delete the default IS-IS accept policy with route policy and backbone route policy configuration:

```
no accept route-map backbone-route-map
```

- c. Create an accept policy instance to selectively accept IS-IS routes and host routes from the DvR backbone, and apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>] route-map
WORD<1-32> backbone-route-map WORD<1-64>
```

- d. **(Optional)** Delete an accept policy instance with route policy and backbone route policy configuration, which applies to all BEBs for a specific I-SID or I-SID list:

```
no accept [i-sid <1-16777215>][isid-list WORD <1-32>] route-map
backbone-route-map
```

- e. Create an IS-IS accept policy instance to selectively accept IS-IS routes and host routes from the DvR backbone, and apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> route-map WORD<1-32> backbone-route-map
WORD <1-64>
```

- f. **(Optional)** Delete an IS-IS accept policy instance with route policy and backbone route policy configuration, which applies to a specific advertising BEB:

```
no accept adv-rtr <x.xx.xx> route-map backbone-route-map
```

8. Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

```
isis apply accept [vrf WORD <1-16>]
```

9. Exit IS-IS Router Configuration mode:

```
exit
```

You are in Global Configuration mode.

Example

Configure an I-SID based IS-IS accept policy with the route policy `test`:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#route-map test 1
Switch:1(route-map)#enable
Switch:1(route-map)#exit

Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 101
Switch:1(config-isis)#accept i-sid 101 route-map test
Switch:1(config-isis)#accept i-sid 101 enable
Switch:1#exit
Switch:1(config)#isis apply accept
```

Configuration of IS-IS accept policy to accept host routes from the DvR backbone**Example 1:**

To accept host routes from the DvR backbone, you must configure a backbone route policy and apply it to the IS-IS accept policy.

1. Configure a route policy for DvR:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch:1(config)#route-map dvrmap1 1
Switch:1(route-map)#enable
```

2. Configure an IS-IS accept policy for I-SID 10, and apply the route policy as a backbone route policy:

```
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 backbone-route-map dvrmap1
Switch:1(config-isis)#accept i-sid 10 enable
Switch:1(config-isis)#exit
```

OR

Configure the default accept policy for IS-IS and DvR, and apply the route policy as a backbone route policy:

```
Switch:1(config)#route-map isismap1 1
Switch:1(route-map)#enable
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept route-map isismap1 backbone-route-map dvrmap1
```

3. Apply the IS-IS accept policy:

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
Switch:1(config)#exit
```

4. Verify the configuration:

```
Switch:1#show ip isis accept
```

```
=====
Isis Accept - GlobalRouter
=====
```

ADV_RTR	I-SID	ISID-LIST	ENABLE	POLICY	BACKBONE POLICY
-	10	-	TRUE		dvrmap1
-	-	-		isismap1	dvrmap1

2 out of 2 Total Num of Isis Accept Policies displayed

Example 2:

Configure an IS-IS accept policy for I-SID 10 that accepts DvR host routes in a subnet, for example, subnet 126.1.1.0/24.

1. Configure an IP prefix list:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list listPrefix 126.1.1.0/24
```

2. Create the route policy dvrmap2 to match the IP prefix list:

```
Switch:1(config)#route-map dvrmap2 1
Switch:1(route-map)#match network listPrefix
Switch:1(route-map)#enable
```

3. Create an IS-IS accept policy with I-SID 10 and apply the route policy as a backbone route policy:

```
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 backbone-route-map dvrmap2
Switch:1(config-isis)#accept i-sid 10 enable
```

4. Apply the IS-IS accept policy:

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
```

The above command causes IS-IS to accept all routes with I-SID 10. To deny IS-IS routes and accept only DvR host routes, you can configure an additional IS-IS route policy as follows:

```
Switch:1(config)#route-map isismap2 1
Switch:1(route-map)#no permit
Switch:1(route-map)#enable

Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 route-map isismap2 backbone-route-map dvrmap2
Switch:1(config-isis)#accept i-sid 10 enable
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
```

5. Verify the configuration:

```
Switch:1(config)#exit
Switch:1#show ip isis accept

=====
Isis Accept - GlobalRouter
=====

ADV_RTR  I-SID  ISID-LIST  ENABLE POLICY  BACKBONE POLICY
-----
-        10     -          TRUE  isismap2     dvrmap2
```

```
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configuration of IS-IS accept policies for a specific VRF instance

Example 1:

Configure IS-IS accept policies to accept host routes from the DvR backbone, for a specific VRF instance.

1. In the VRF green context, configure the route policy `dvrmap3` for DvR:

```
Switch:1(config)#router vrf green
Switch:1(router-vrf)#route-map dvrmap3 1
Switch:1(router-vrf-routemap)#enable
```

2. Use one of the following options to configure an IS-IS accept policy, and apply the route policy as a backbone route policy:

Configure an IS-IS accept policy for a specific advertising BEB with nickname `1.11.11`:

```
Switch:1(router-vrf-routemap)#isis accept adv-rtr 1.11.11 backbone-route-map
dvrmap3
Switch:1(router-vrf-routemap)#exit
Switch:1(router-vrf)#isis accept adv-rtr 1.11.11 enable
```

```
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
POLICY
-----
1.11.11  -         -                        TRUE                      dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

```
Switch:1(config)#show ip isis accept vrfids 2
```

```
=====
                        Isis Accept - VRF green
=====
ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
POLICY
-----
1.11.11  -         -                        TRUE                      dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure an accept policy for I-SID 10:

```
Switch:1(router-vrf)#isis accept i-sid 10 backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
POLICY
-----
-         10       -                        TRUE                      dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure an accept policy for the I-SID list listisids:

```
Switch:1(router-vrf)#isis accept isid-list listisids backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
```

ADV_RTR	I-SID	ISID-LIST	ENABLE POLICY	BACKBONE POLICY
-	10	listisids	TRUE	dvrmap3

1 out of 1 Total Num of Isis Accept Policies displayed

Configure the default accept policy for IS-IS and DvR:

```
Switch:1(router-vrf)#route-map isismap3 1
Switch:1(router-vrf-routemap)#
Switch:1(router-vrf-routemap)#enable
Switch:1(router-vrf-routemap)#
Switch:1(router-vrf-routemap)#isis accept route-map isismap3 backbone-route-map
dvrmap3
Switch:1(router-vrf)#
```

```
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
```

ADV_RTR	I-SID	ISID-LIST	ENABLE POLICY	BACKBONE POLICY
-	-	-	TRUE isismap3	dvrmap3

1 out of 1 Total Num of Isis Accept Policies displayed

Configure the default accept policy for DvR:

```
Switch:1(router-vrf)#isis accept backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
```

ADV_RTR	I-SID	ISID-LIST	ENABLE POLICY	BACKBONE POLICY
-	-	-	TRUE	dvrmap3

1 out of 1 Total Num of Isis Accept Policies displayed

Example 2:

Configure an accept policy for I-SID 10 that accepts DvR host routes in a subnet, for example, subnet 126.1.1.0/24.

1. Configure an IP prefix list:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list listPrefix 126.1.1.0/24
```

2. For a specific VRF instance, create a route policy to match the IP prefix list:

```
Switch:1(config)#router vrf green
Switch:1(router-vrf)#route-map dvrmap4 1
Switch:1(router-vrf-routemap)#match network listPrefix
Switch:1(router-vrf-routemap)#enable
Switch:1(router-vrf-routemap)#exit
Switch:1(router-vrf)#
```

3. Create an IS-IS accept policy with I-SID 10, and apply the route policy as the backbone route policy:

```
Switch:1(router-vrf)#accept i-sid 10 backbone-route-map dvrmap4
Switch:1(router-vrf)#accept i-sid 10 enable
```

4. Apply the IS-IS accept policy:

```
Switch:1(router-vrf)#exit
Switch:1(config)#isis apply accept
```

5. Verify the configuration:

```
Switch:1(config)#exit
Switch:1(router-vrf)#show ip isis accept vrf green

=====
Isis Accept - VRF green
=====
ADV_RTR   I-SID     ISID-LIST           ENABLE POLICY       BACKBONE
POLICY
-----
-         -         -                   TRUE                 dvrmap4

1 out of 1 Total Num of Isis Accept Policies displayed
```

Variable definitions

Use the data in the following table to use the `ip isid-list` command.

Variable	Value
<code>WORD<1-32></code>	Creates a name for your I-SID list.
<code><1-16777215></code>	Specifies an I-SID number.
<code>list WORD<1-1024></code>	Specifies a list of I-SID values. For example, in the format 1,3,5,8-10.

Use the data in the following table to use the `accept` command.

Variable	Value
<code>adv-rtr <x.xx.xx></code>	Specifies the SPBM nickname for each advertising BEB to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter.
<code>backbone-route-map WORD<1-64></code>	Specifies the DvR backbone route map.

Table continues...

Variable	Value
enable	Enables an IS-IS accept policy.
i-sid <1-16777215>	Specifies an I-SID number to represent a local or remote Layer 3 VSN to which the IS-IS accept policy applies. Use the parameter to apply a filter for routes from a specific I-SID that represents the remote VSN. Based on the routing policy the system applies, the system can redistribute the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).
isid-list WORD<1-32>	Specifies the I-SID list name that represents the local or remote Layer 3 VSNs to which the IS-IS accept policy applies. Use the parameter to apply a default filter for all routes from a specific I-SID that represents the remote VSN. Based on the routing policy the system applies, the system redistributes the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).
route-map WORD<1-64>	Specifies a route policy by name. You must configure the route policy earlier in a separate procedure.

Use the data in the following table to use the **isis apply accept** command.

Variable	Value
vrf WORD<1-16>	Specifies a specific VRF instance.

Viewing IS-IS accept policy information

Use the following procedure to view IS-IS accept policy information on the switch.

Procedure

1. Display IS-IS accept policy information:

```
show ip isis accept [vrf WORD<1-16>][vrfids WORD<0-512>]
```

2. Display I-SID list information:

```
show ip isid-list [vrf WORD<1-16>][vrfids WORD<0-512>][WORD<1-32>]
```

3. Display route information:

```
show ip route [vrf WORD<1-16>]
```

The NH VRF/ISID column displays the I-SID for inter-Virtual Services Network (VSN) routes redistributed with IS-IS accept policies, only if the I-SID redistributed does not have an IP VSN associated with it. If an IP VSN exists for that I-SID, the VRF name displays. If the I-SID is 0, the column represents and displays as the GlobalRouter.

The existing IS-IS routes for Layer 3 VSNs continue to display as the VRF name of the IP VSN.

4. Display the SPBM IP unicast Forwarding Information Base (FIB):

```
show isis spbm ip-unicast-fib [all] [id <1-16777215>][spbm-nh-as-
mac]
```

Example

View IS-IS accept policy information:

```
Switch:1#show ip route vrf test
=====
                        IP Route - VRF test
=====
DST          MASK          NEXT          NH          INTER
VRF/ISID    COST  FACE  PROT  AGE  TYPE  PRF
-----
1.1.1.5     255.255.255.255  1.1.1.5  GlobalRouter  0    0    ISIS  0    IB    200
1.1.1.13    255.255.255.255  Switch13  GRT          10   1000 ISIS  0    IBSV  7
1.1.1.200   255.255.255.255  Switch200 GRT          10   1000 ISIS  0    IBSV  7
5.7.1.0     255.255.255.0    5.7.1.1  -            1    7    LOC   0    DB    0
13.7.1.0    255.255.255.0    Switch13  GlobalRouter  10   1000 ISIS  0    IBSV  7
100.0.0.0   255.255.255.0    100.0.0.1 GlobalRouter  0    100  ISIS  0    IB    200
111.1.1.0   255.255.255.0    111.1.1.1 hub        0    111  ISIS  0    IB    200

Switch:1(config)#show isis spbm ip-unicast-fib
=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF  VRF  DEST  Destination  NH  BEB  OUTGOING  SPBM  PREFIX  IP ROUTE
ISID ISID  ID     NH BEB  VLAN  INTERFACE  COST  COST  PREFERENCE
-----
GRT  -    101   1.1.1.13/32  Switch13  1000  1/7      10    44     7
GRT  -    101   1.1.1.13/32  Switch13  1001  1/7      10    44     7

-----
Total number of SPBM IP-UNICAST FIB entries 2
-----

Switch:1(config)#show ip isid-list test
=====
                        IP ISID LIST
=====
List Name          I-SID          VRF
-----
test               1              GlobalRouter
                  3              GlobalRouter
                  4              GlobalRouter
                  5              GlobalRouter
                  10             GlobalRouter
                  22             GlobalRouter
```

All 6 out of 6 Total Num of Isid Lists displayed

```
Switch:1(router-vrf)#show ip isid-list vrf red
```

```
=====
                        IP ISID LIST red
=====
List Name                I-SID                VRF
-----
test1                    11                   1
                        12                   1
                        13                   1
                        14                   1
                        15                   1
```

Variable definitions

Use the data in the following table to use the `show ip isis accept` command.

Variable	Value
vrf WORD<1-16>	Displays I-SID list information for a particular VRF by name.
vrfids WORD<0-512>	Displays I-SID list information for a particular VRF ID.

Use the data in the following table to use the `show ip isid-list` command.

Variable	Value
vrf WORD<1-16>	Displays I-SID list information for a particular VRF by name.
vrfids WORD<0-512>	Displays I-SID list information for a particular VRF ID.
WORD<1-32>	Displays I-SID list information for a particular I-SID list name.

Use the data in the following table to use the `show ip route` command.

Variable	Value
vrf WORD<1-16>	Displays I-SID list information for a particular VRF by name.

Use the data in the following table to use the `show isis spbm ip-unicast-fib` command.

Variable	Value
all	Displays all IS-IS SPBM IP unicast Forwarding Information Base (FIB) information.
id <1-16777215>	Displays IS-IS SPBM IP unicast FIB information by I-SID ID.
spbm-nh-as-mac	Displays the next hop B-MAC of the IP unicast FIB entry.

IP Shortcuts configuration using EDM

This section provides procedures to configure IP Shortcuts using Enterprise Device Manager (EDM).

Configuring a Circuitless IPv4 interface

About this task

You can use a circuitless IPv4 (CLIPv4) interface to provide uninterrupted connectivity to your system.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IP**.
3. Click the **Circuitless IP** tab.
4. Click **Insert**.
5. In the **Interface** field, assign a CLIP interface number.
6. Enter the IP address.
7. Enter the network mask.
8. Click **Insert**.
9. To delete a CLIP interface, select the interface and click **Delete**.

Circuitless IP field descriptions

Use the data in the following table to use the **Circuitless IP** tab.

Name	Description
Interface	Specifies the number assigned to the interface.
Ip Address	Specifies the IP address of the CLIP.
Net Mask	Specifies the network mask.

Configuring a Circuitless IPv6 interface

Before you begin

Change the VRF instance as required to configure a Circuitless IPv6 interface on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

About this task

You can use a circuitless IPv6 (CLIPv6) interface to provide uninterrupted connectivity to your system.

Procedure

1. In the navigation pane, expand the **Configuration > IPv6** folders.
2. Click **IPv6**.
3. Click the **Circuitless IP** tab.
4. Click **Insert**.
5. In the **Interface** field, assign a CLIP interface number.
6. Type the IPv6 address and prefix length.

Circuitless IPv6 field descriptions

Use the data in the following table to use the **Circuitless IPv6** tab.

Name	Description
Interface	Specifies the interface to which this entry applies.
Addr	Specifies the IPv6 address to which this entry applies.
AddrLen	Specifies the prefix length value for this address. You cannot change the address length after you create it. You must provide this value to create an entry in this table.

Configuring IS-IS redistribution

Use this procedure to configure IS-IS redistribution. In the Virtual Routing and Forwarding (VRF), just like in the Global Router, the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP, within the context of a VRF. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

The VRF specific routes are transported in TLV 184 with the I-SID assigned to the VPNs. After extracting the IP VPN IP reachability information, the routes are installed in the route tables of the appropriate VRFs based on the I-SID association.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Redistribute** tab.
4. Click **Insert**.
5. Complete the fields as required.
6. Click **Insert**.

IS-IS Redistribute field descriptions

Use the data in the following table to configure the **IS-IS Redistribute** tab.

Name	Description
DstVrflid	Specifies the destination Virtual Routing and Forwarding (VRF) ID used in the redistribution.
Protocol	Specifies the protocols that receive the redistributed routes.
SrcVrflid	Specifies the source VRF ID used in the redistribution. For IS-IS, the source VRF ID must be the same as the destination VRF ID.
RouteSource	Specifies the source protocol for the route redistribution entry.
Enable	Enables or disables a redistribution entry. The default is disable.
RoutePolicy	Specifies the route policy to be used for the detailed redistribution of external routes from a specified source into the IS-IS domain.
Metric	Specifies the metric for the redistributed route. The value can be a range between 0 to 65535. The default value is 0. Use a value that is consistent with the destination protocol.
MetricType	Specifies the metric type. Specifies a type1 or a type2 metric. For metric type1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type2, the cost of the external routes is equal to the external cost alone. The default is type2.
Subnets	Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose suppress to advertise subnets aggregated to their classful subnet. Choose allow to advertise the subnets individually with the learned or configured mask of the subnet. The default is allow.

Configuring SPBM IP shortcuts

In addition to Layer 2 virtualization, the SPBM model is extended to also support Routed SPBM, otherwise called SPBM IP Shortcuts.

SPBM allows a network to make the best use of routing and forwarding techniques, where only the BEBs perform an IP route lookup and all other nodes perform standard Ethernet switching based on the existing shortest path tree. This allows for end to end IP-over-Ethernet forwarding without the need for ARP, flooding, or reverse learning.

To enable IP shortcuts on the BEBs, you must configure a circuitless IP address (loopback address) and specify this address as the IS-IS source address. This source address is automatically advertised into IS-IS using TLV 135. In addition, to advertise routes from the BEBs into the SPBM network, you must enable route redistribution of direct and static routes into IS-IS.

After you have configured the SPBM infrastructure, you can enable SPBM IP shortcuts to advertise IP routes across the SPBM network using the following procedure.

Before you begin

- You must configure the required SPBM and IS-IS infrastructure, which includes the creation of SPBM B-VLANs.
- Before redistributing routes into IS-IS, you must create the Customer VLANs, add slots/ports, and add the IP addresses and network masks.
- You must configure a circuitless IP (CLIP) interface:
 - To configure an IPv4 CLIP interface, see [Configuring a Circuitless IPv4 interface](#) on page 49
 - To configure an IPv6 CLIP interface, see [Configuring a Circuitless IPv6 interface](#) on page 49

Procedure

1. In the navigation pane, expand the **Configuration > IS-IS** folders.
2. Click **IS-IS**.
3. From the **Globals** tab, in the **IpSourceAddress** field, specify the CLIP interface to use as the source address for SBPM IP Shortcuts.

 **Note:**

For IPv6 Shortcuts, select **ipv6** in the **Ipv6SourceAddressType** field, and then use the **Ipv6SourceAddress** field to specify the CLIPv6 interface to use as the source address for SBPM IPv6 Shortcuts.

4. Click **Apply**.
5. In the navigation pane, expand the **Configuration > IS-IS > SPBM** folders.
6. Click the **SPBM** tab.
7. In the **IpShortcut** field select **enable**.

 **Note:**

For IPv6 Shortcuts, select **enable** in the **Ipv6Shortcut** field.

8. Click **Apply**.
9. In the navigation pane, expand the **Configuration > IP** folders.
10. Click **Policy**.
11. Click the **Route Redistribution** tab.
12. Click **Insert** to identify routes on the local switch to be announced into the SPBM network.
13. Using the fields provided, specify the source protocols to redistribute into IS-IS. In the **Protocol** field, ensure to specify **isis** as the destination protocol.
14. Click **Insert**.

Enabling IPv6 IS-IS redistribution

Enable redistribution to announce routes of a certain source protocol type into the IPv6 IS-IS domain.

You can redistribute directly connected routes, IPv6 static routes, IPv6 BGP routes, OSPFv3 routes, and RIPng routes into IPv6 IS-IS.

Before you begin

Change the VRF instance as required to enable IPv6 IS-IS redistribution on a specific VRF instance. Not all parameters are configurable on non-default VRFs.

Procedure

1. In the navigation pane, expand the **Configuration > IPv6** folders.
2. Click **IS-IS**.
3. Click the **Redistribute** tab.
4. For the type of route source, double-click the cell in the **Enable** column to change the value.
5. Select **enable**.
6. Click **Apply**.

Redistribute field descriptions

Use the data in the following table to use the **Redistribute** tab.

Name	Description
DstVrflid	Shows the ID of the destination virtual router and forwarder (VRF).
Protocol	Shows the routing protocol that receives the external routing information. In this case, the routing protocol is IPv6 IS-IS.
SrcVrflid	Shows the ID of the source VRF.
RouteSource	Shows the source protocol from which to receive routes to insert into the IPv6 IS-IS domain. <ul style="list-style-type: none"> • For Global Router — all routing protocols can be redistributed into IS-IS. • For VRFs — only direct and static routes can be redistributed into IS-IS.
Enable	Configures the status of route redistribution. The default is disable.

Applying IS-IS accept policies globally

Apply IS-IS accept policies globally. Use IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud. Accept policies apply to incoming traffic and determine whether to add the route to the routing table.

After you apply the IS-IS accept filters, the device removes and re-adds all routes with updated filters.

IS-IS accept policies are disabled by default.

* Note:

- After you apply IS-IS accept policies globally the application can disrupt traffic and cause temporary traffic loss. After you configure the IS-IS accept policies value to **Apply**, the device reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. You should make all the relevant accept policy changes, and then apply IS-IS accept policies globally at the end.
- If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- Ensure the IP IS-IS filter exists.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Global** tab.
4. Select a name from the list or enter name in the **DefaultPolicyName** field to specify the route policy name for the default filter.
5. Select **Apply** to apply the default policy.

Accept Global field descriptions

Use the data in the following table to configure the **Accept Global** tab.

Name	Description
DefaultPolicyName	Specifies the route policy name for the default filter.
DefaultBackbonePolicyName	Specifies the backbone host route policy name for the default filter.

Table continues...

Name	Description
Apply	Applies the default policy when you configure the field to apply. The device only activates the default policy if the route map (the default policy name) has a value. If you do not select apply, the device takes no action. The GRT always returns no action.

Configuring an IS-IS accept policy for a specific advertising BEB

Configure an IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB). Specify the SPBM nickname and the IS-IS accept policy name to allow you to apply the IS-IS accept policy.

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.

* **Note:**

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Nick Name** tab.
4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.
8. Click **Insert**.

Accept Nick Name field descriptions

Use the data in the following table to configure the **Accept Nick Name** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter. The value is 2.5 bytes in the format <x.xx.xx>.
Enable	Enables or disables the SPBM nickname advertising router entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies a route policy. You must configure a policy earlier in a separate procedure.
BackbonePolicyName	Specifies the route policy for the backbone routes. You must configure a policy earlier in a separate procedure.

Configuring an IS-IS accept policy to apply for a specific I-SID

Configure an IS-IS accept policy for a specific I-SID number to represent a local or remote Layer 3 VSN, which allows the system to redistribute the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Isid** tab.
4. Click **Insert**.
5. In the **Isid** field, specify the SPBM nickname.
6. Select enable in the **Enable** check box to enable the filter.

7. In the **PolicyName** field, specify the route-map name.
8. Click **Insert**.

Accept Isid field descriptions

Use the data in the following table to configure the **Accept Isid** tab.

Name	Description
Isid	Configures a specific I-SID number to represent a local or remote Layer 3 VSN to which the IS-IS accept policy applies. Based on the routing policy the system applies, the system redistributes the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).
Enable	Enables or disables the I-SID entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies the route map name. You must configure a policy earlier in a separate procedure.
BackbonePolicyName	Specifies the backbone route map name. You must configure a policy earlier in a separate procedure.

Configuring an IS-IS accept policy for a specific advertising BEB and I-SID

Configures a specific advertising Backbone Edge Bridge (BEB) with a specific I-SID to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB.

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Nick-Name Isid** tab.

4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Isid** field, specify an I-SID number.
7. Select enable in the **Enable** check box to enable the filter.
8. In the **PolicyName** field, specify the route-map name.
9. Click **Insert**.

Accept Nick-Name Isid descriptions

Use the data in the following table to configure the **Accept Nick-Name Isid** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The value is 2.5 bytes in the format <x.xx.xx>.
Isid	Specifies an I-SID used to filter. The value 0 is used for the Global Router.
Enable	Enables or disables the I-SID entry. The default is disabled.
PolicyName	Specifies the route policy name. You must configure a policy earlier in a separate procedure.
BackBonePolicyName	Specifies the backbone route policy name. You must configure a policy earlier in a separate procedure.

Configuring an I-SID list for an IS-IS accept policy

Configures a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IS-IS accept policy applies. After you create the list of I-SID numbers, you must then create, configure, and enable the IS-IS accept policy.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Isid-List** tab.
4. Click **Insert**.
5. In the **Name** field, specify a name for the I-SID list.

6. Select **Isid** or **Isid-List**.
7. Specify an I-SID number or a list of I-SID numbers.
8. Click **Insert**.

Isid-List field descriptions

Use the data in the following table to configure the **Isid-List** tab.

Name	Description
Name	Specifies the name of the I-SID list.
Isid or Isid-List	Specifies that you either want to add a particular I-SID or a list of I-SID numbers.
Isid	Specifies a particular I-SID number or a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).

Configuring an IS-IS accept policy for a specific I-SID list

Configure an IS-IS accept policy for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.

Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Isid-List** tab.
4. Click **Insert**.
5. In the **Name** field, specify the I-SID list name.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.

- Click **Insert**.

Accept Isid-List field descriptions

Use the data in the following table to configure **Accept Isid-List** tab.

Name	Description
Name	Specifies the name of I-SID list.
Enable	Enables or disables the I-SID list entry. The value must be enabled to filter. The default is disabled.
PolicyName	Specifies the route policy name.
BackBonePolicyName	Specifies the backbone route policy name.

Configuring an IS-IS accept policy for a specific advertising BEB and I-SID-list

Configure an IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB) for a specific I-SID list to represent local or remote Layer 3 VSNS, which allows the system to redistribute the remote VSNS to the VSN where you applied the filter.

* Note:

If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

About this task

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Nick-Name Isid-List** tab.
4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Name** field, specify an I-SID list name.
7. Select enable in the **Enable** check box to enable the filter.

8. In the **PolicyName** field, specify the route-map name.
9. Click **Insert**.

Accept Nick–Name Isid-List field descriptions

Use the data in the following table to configure the **Accept Nick-Name Isid-List** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter is present the device applies the specific filter. The value is 2.5 bytes in the format <x.xx.xx>.
Name	Specifies the name of the I-SID list used to filter.
Enable	Enables or disables the SPBM nickname advertising router entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies a route policy name.
BackBonePolicyName	Specifies a backbone route policy name.

IP Shortcuts SPBM configuration example

The following figure shows a sample IP Shortcuts over SPBM deployment.

Examples and network illustrations in this document may illustrate only one of the supported platforms. Unless otherwise noted, the concept illustrated applies to all supported platforms.




Figure 7: SPBM IP Shortcuts

The following sections show the steps required to configure the SPBM IP Shortcuts parameters in this example. You must first configure basic SPBM and IS-IS infrastructure. For more information, see *Configuring Fabric Basics and Layer 2 Services*.

Note the following:

- IP IS-IS redistribution needs to be configured to inject IP shortcuts routes into IS-IS. The one exception is the circuitless IP address configured as the IS-IS ip-source-address. This address is automatically advertised without the need for a redistribution rule.
- In the displayed configuration, only direct routes are injected (the same configuration is possible for static routes). To inject IPv6 routes, you must enable route redistribution of IPv6 direct, IPv6 static, and OSPFv3 routes into IS-IS.
- No IP address needs to be configured on SwitchG.

The following sections show the steps required to configure the SPBM IP Shortcuts parameters in this example.

SwitchC

```
CIRCUITLESS INTERFACE CONFIGURATION - GlobalRouter
```

```
interface loopback 1
ip address 1 10.0.0.1/255.255.255.255
exit
```

```
ISIS CONFIGURATION
```

```
router isis
ip-source-address 10.0.0.1
```

```
ISIS SPBM CONFIGURATION
```

```
spbm 1 ip enable
exit
```

```
VLAN CONFIGURATION
```

```
vlan create 13 type port-mstprstp 0
vlan members 13 1/2 portmember
interface Vlan 13
ip address 10.0.13.1 255.255.255.0
exit
```

```
IP REDISTRIBUTION CONFIGURATION - GlobalRouter
```

```
router isis
redistribute direct
redistribute direct metric 1
redistribute direct enable
exit
```

```
IP REDISTRIBUTE APPLY CONFIGURATIONS
```

```
isis apply redistribute direct
```

SwitchD

```
CIRCUITLESS INTERFACE CONFIGURATION - GlobalRouter
```

```
interface loopback 1
ip address 1 10.0.0.2/255.255.255.255
exit
```

```
ISIS CONFIGURATION
```

```
router isis
ip-source-address 10.0.0.2
```

```

ISIS SPBM CONFIGURATION

spbm 1 ip enable
exit

VLAN CONFIGURATION

vlan create 14 type port-mstprstp 0
vlan member add 14 1/2
interface Vlan 14
ip address 10.0.14.1 255.255.255.0
exit

IP REDISTRIBUTION CONFIGURATION - GlobalRouter

router isis
redistribute direct
redistribute direct metric 1
redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct

```

Verifying operation — SwitchC

```

SwitchC:1# show isis spbm ip-unicast-fib
=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF      ISID      Destination      NH BEB      VLAN      OUTGOING      SPBM PREFIX
                                  INTERFAC     COST COST
-----
GRT      -          10.0.0.2/32     SwitchD     4000      1/30         20    1
GRT      -          10.0.14.1/24   SwitchD     4000      1/30         20    1
=====

Total number of SPBM IP-UNICAST FIB entries 2
=====

SwitchC:1# show ip route
=====
                        IP Route - GlobalRouter
=====
DST          MASK          NEXT          VRF      NH      INTER     PROT AGE  TYPE  PRF
          COST      FACE
-----
10.0.0.1     255.255.255.255  10.0.0.1     -        1       0        LOC   0    DB   0
10.0.0.2     255.255.255.255  SwitchD      Glob~   20     4000     ISIS  0    IBS  7
10.0.13.1    255.255.255.0   10.0.13.1    -        1       13       LOC   0    DB   0
10.0.14.1    255.255.255.0   SwitchD      Glob~   20     4000     ISIS  0    IBS  7

4 out of 4 Total Num of Route Entries, 4 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

```

Verifying operation — SwitchD

SwitchD:1# show isis spbm ip-unicast-fib

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

VRF	ISID	Destination	NH BEB	VLAN	OUTGOING INTERFACE	SPBM COST	PREFIX COST
GRT	-	10.0.0.1/32	SwitchC	4000	1/20	20	1
GRT	-	10.0.13.1/24	SwitchC	4000	1/20	20	1

```
-----
Total number of SPBM IP-UNICAST FIB entries 2
-----
```

SwitchD:1# show ip route

```
=====
                        IP Route - GlobalRouter
=====
```

DST	MASK	NEXT	VRF	NH COST	INTER FACE	PROT	AGE	TYPE	PRF
10.0.0.1	255.255.255.255	SwitchC	Glob~	20	4000	ISIS	0	IBS	7
10.0.0.2	255.255.255.255	10.0.0.2	-	1	0	LOC	0	DB	0
10.0.13.1	255.255.255.0	SwitchC	Glob~	20	4000	ISIS	0	IBS	7
10.0.14.1	255.255.255.0	10.0.14.1	-	1	14	LOC	0	DB	0

4 out of 4 Total Num of Route Entries, 4 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

Chapter 5: Layer 3 VSN configuration

The following sections provide concepts and procedures to configure Layer 3 VSN.

Layer 3 VSN configuration fundamentals

This section provides fundamental concepts on Layer 3 VSN.

SPBM Layer 3 VSN

The SPBM Layer 3 VSN feature is a mechanism to provide IP connectivity over SPBM for VRFs. SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF.




Figure 8: SPBM Layer 3 VSN

In the preceding figure, the BEBs are connected over the SPBM cloud running IS-IS. VRF red and green are configured on the BEBs. VRF red on BEB A has to send and receive routes from VRF red on BEB D. Similar operations are required for VRF green on BEB A and BEB D.

IS-IS TLV 184 is used to advertise SPBM Layer 3 VSN route information across the SPBM cloud. To associate advertised routes with the appropriate VRF, each VRF is associated with an I-SID. All VRFs in the network that share the same I-SID participate in the same VSN.

*** Note:**

IPv4 Layer 3 VSN and IPv6 Layer 3 VSN co-exist and share the same I-SID. You need to configure I-SID only once. The advantage of having two separate VPNs, one for IPv4 and one for IPv6 is because it gives user an option to enable them separately.

In this example, I-SID 101 is associated with VRF green and I-SID 102 is associated with VRF red. The I-SID is used to tie the advertised routes to a particular VRF. This identifier has to be the same on all edge nodes for a particular VRF, and has to be unique across all the VRFs on the same node

When IS-IS receives an update from an edge node, it looks for the Layer 3 VSN TLV, and if one exists, it looks at the I-SID identifier. If that identifier is mapped to a local VRF, extracts the IPv4 or IPv6 routes and add them to the RTM of that VRF.

With SPBM Layer 3 VSN, the packet forwarding works in a similar fashion as the IP Shortcuts on the Global Router, with the difference that the encapsulation includes the I-SID to identify the VRF that the packet belongs to. The following figure shows the packet forwarding for VRF red.

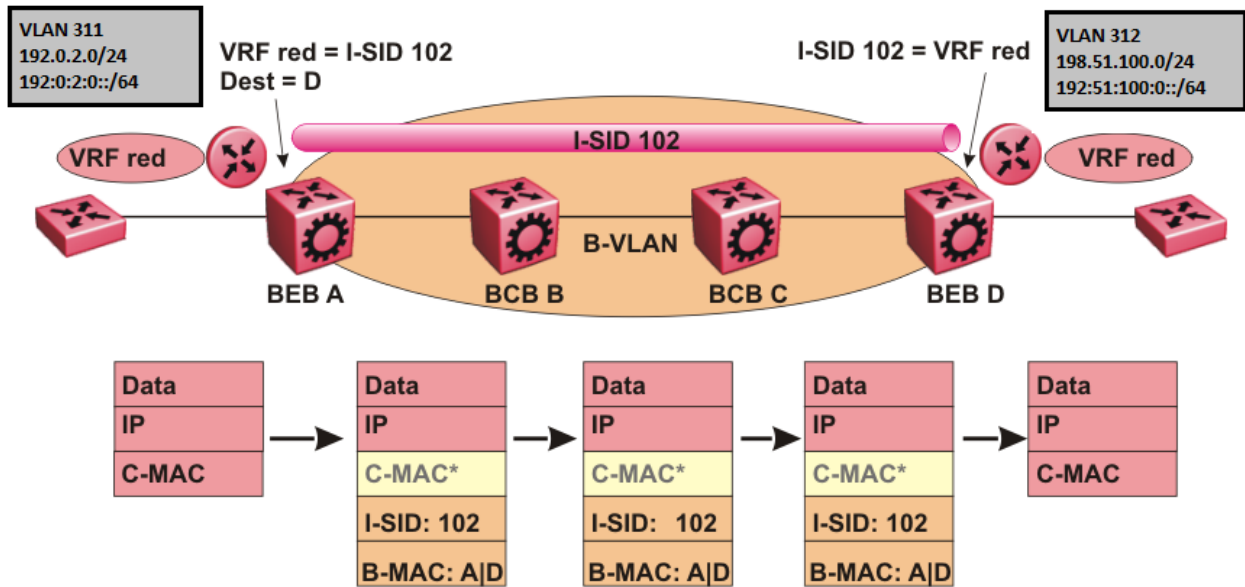


Figure 9: Packet forwarding in SPBM Layer 3 VSN

When BEB A receives traffic from VRF red that must be forwarded to the far-end location, it performs a lookup and determines that VRF red is associated with I-SID 102 and that BEB D is the destination for I-SID 102. BEB A then encapsulates the IP data into a new B-MAC header, using destination B-MAC: D.

*** Note:**

With SPBM Layer 3 VSN, the C-MAC header is all null. This header does not have any significance in the backbone. It is included to maintain the same 802.1ah format for ease of implementation.

At BEB D, the node strips off the B-MAC encapsulation, and performs a lookup to determine the destination for traffic with I-SID 102. After identifying the destination as VRF red, the node forwards the packet to the destination VRF.

*** Note:**

IPv4 Layer 3 VSN and IPv6 Layer 3 VSN co-exist and share the same I-SID. The advantage of having two separate VPNs, one for IPv4 and one for IPv6 is because it gives user an option to enable them separately.

IPv6 Layer 3 VSN limitations and considerations

Keep the following considerations in mind when configuring the IPv6 Layer 3 VSN :

- You can enable IPv6 Layer3 VSN only when `spbm boot` config flag is true.
- IPv4 Shortcuts and IPv6 Shortcuts must be enabled.
- IPv6 route redistribution is available only for static and direct routes.

The following features are not supported:

- Inter VRF route redistribution into ISIS.
- IPv6 Accept Policies.

Fabric Connect Service Types

The Fabric Connect technology delivers Layer 2 and Layer 3 virtualization. These virtualized Layer 2 and Layer 3 instances are referred to as Virtual Service Networks (VSNs). A Service Identifier (I-SID) is used to uniquely distinguish these service instances network-wide, and a User Network Interface (UNI) is the boundary or demarcation point between the “service layer” of traditional networks, that is VLANs and VRFs, and the Fabric Connect “service layer”, that is Layer 2 & Layer 3 VSNs.

- Layer 2 VSNs are virtual broadcast domains interconnecting UNI members that share the same Layer 2 VSN I-SID. MAC learning/aging is applied to all Layer 2 VSNs.
- Layer 3 VSNs are virtual routed Layer 3 networks (Layer 3 VPN) leveraging IS-IS as the routing protocol between VRFs that share the same Layer 3 VSN I-SID.

Fabric Connect uses the User-Network-Interface (UNI) to denote the capabilities and attributes of the service interfaces. Fabric connect devices support the following UNI types:

- *VLAN UNI (C-VLAN)* — a device-specific VLAN-ID maps to a Layer 2 VSN I-SID – all device physical ports that are associated with the VLAN are therefore associated with the UNI.
- *Flex UNI* — it has the following sub-types:
 - *Switched UNI* — a VLAN-ID and a given port (VID, port) maps to a Layer 2 VSN I-SID. With this UNI type, VLAN-IDs can be reused on other ports and therefore mapped to different I-SIDs.
 - *Transparent Port UNI* — a physical port maps to a Layer 2 VSN I-SID (all traffic through that port, 802.1Q tagged or untagged, ingress and egress is mapped to the I-SID). Note: All VLANs on a Transparent Port UNI interface now share the same single MAC learning table of the Transparent Port UNI I-SID.
- *E-Tree UNI* — it extends Private VLANs beyond one Switch to form a network-wide E-Tree service infrastructure. An E-Tree UNI is a Layer 2 VSN where broadcast traffic flows from Hub sites to Spokes sites, and from Spokes to Hubs, but not between Spoke sites. E-Tree Hubs can be formed with any VLAN UNI, while E-Tree Spokes must be configured as Private VLAN UNIs.

- *Layer 3 VSN UNI* — a device-specific VRF maps to an I-SID, and the control plane exchanges the Layer 3 routes belonging to the same I-SID. All VRFs in a network sharing the same Layer 3 I-SID effectively form a Layer 3 VPN. Layer 3 VSNs can be configured to simultaneously support both IP Unicast and IP Multicast.

For more information on Layer 3 VSN, see *Configuring Fabric Layer 3 Services*.

ECMP with IS-IS

The Equal Cost Multipath (ECMP) feature supports and complements the IS-IS protocol.

With ECMP, the switch can determine multiple equal-cost paths to the same destination prefix.

You can use multiple paths for load sharing of traffic. These multiple paths allow faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use your links between routers more efficiently when sending IP and IPv6 traffic. Equal Cost Multipath is formed using routes from the same protocol.

The number of multiple paths a switch can support differs by hardware platform. For more information about feature support, see .

ECMP within IS-IS routes

Equal Cost Multipath (ECMP) allows the device to determine up to eight equal cost paths to the same destination prefix. The maximum number of equal cost paths you can configure depends on the hardware platform. For more information, see *Release Notes*.

If the device learns the same route from multiple sources, the information is ECMP only if the routes:

- are from the same VSN
- have the same SPBM cost
- have the same prefix cost
- have the same IP route preference

Multiple BEBs can announce the same route, either because the Layer 2 LAN connects to multiple BEBs for redundancy, or because segments of the LAN are Layer 2 bridged. In Layer 2, if the device has to tie-break between multiple sources, the tie-breaking is based on cost and hop count.

In Layer 3, hop count is not used for tie-breaking. Instead, the device uses the following precedence rules to tie-break. In the following order, the device prefers:

1. Routes that do not include nodes with the overload bit set.

When a router node runs out of system resources (memory or CPU), it alerts the other routers in the network by setting the overload bit in its link-state packets (LSPs). When this bit is set, the node is not used for transit traffic but only for traffic packets destined to the node's directly connected networks and IP prefixes.

2. Local routes over remote routes.

If a route is learned locally, for example, through inter-VRF route leaking, it is most preferred.

3. Routes with the lowest route preference.

By default, IS-IS routes within the same VSN are added to the LSDB with a default preference of 7. Inter-VSN routes are added to the LSDB with a route preference of 200. You can however, change the route preference using IS-IS accept policies.

4. Metric type internal (type 1) over metric type external (type 2).

5. Routes with the lowest SPBM cost.

6. Routes with the lowest prefix cost.

If the metric type is internal, then the tie-break is on SPB cost first, and then on the prefix cost. Otherwise the tie-break is only on the prefix cost.

You can either change this using a route-map on the remote advertising node with the **redistribute** command, or using a route-map on the local node with the IS-IS accept policy.

7. Routes within a VSN with a lower Layer 3 VSN I-SID.

The device considers the Global Routing Table (GRT) to have an I-SID equal to zero.

When you use multiple B-VLANs in the SPBM core, multiple paths exist to reach a particular SPBM node, one on each B-VLAN; therefore, any IP prefix or IPv6 prefix that the device receives from a BEB results in multiple ECMP paths. These paths may or may not be physically diverse. SPBM supports up to two B-VLANs; a primary B-VLAN and a secondary B-VLAN.

If more ECMP paths are available than the configured number of paths, then the device adds the routes using the following order: The device selects all routes from the primary B-VLAN and orders the routes learned through that B-VLAN from lowest system ID to the highest IS-IS system ID, then the device moves on to select all routes from the secondary B-VLAN, ordering those routes from lowest IS-IS system ID to the highest IS-IS system ID until you reach the number of equal paths configured.

For example, consider an SPB core configured with two B-VLANs (primary B-VLAN 1000 and secondary B-VLAN 2000), and the device learns routes from two BEBs called BEB-A (with a lower IS-IS system ID) and BEB-B (with a higher IS-IS system ID, then the order in which the next-hop is chosen for those routes are as follows.

If a route is learned only from BEB-A with the maximum number of allowed ECMP paths configured as 8 (default), then the order in which the next-hop is chosen for that route is:

1. BEB-A B-VLAN 1000
2. BEB-A B-VLAN 2000

If routes are learned from both BEB-A and BEB-B with maximum number of allowed ECMP paths configured as 8 (default), then the order in which the next-hop is chosen for those routes are:

1. BEB-A B-VLAN 1000
2. BEB-B B-VLAN 1000
3. BEB-A B-VLAN 2000
4. BEB-B B-VLAN 2000

If ECMP is disabled, the maximum number of allowed ECMP paths is 1 and the device adds the route from the lowest system ID with the primary B-VLAN. In this example, the device adds BEB-A B-VLAN 1000.

*** Note:**

- ECMP is supported for IPv6 Shortcut routes.
- To add IS-IS equal cost paths in the routing table, you must enable IPv6 ECMP feature globally.

ECMP Impact on IS-IS Route Selection for Inter-VRF Routes with vIST

This section illustrates the impact ECMP can have on a configuration that implements user-defined VRFs in a vIST cluster and how to avoid incorrect route selection.

Understanding the Configuration

Imagine the following configuration:

- A vIST cluster exists with multiple VRF contexts.
- On both nodes, VRF A redistributes routes into IS-IS as external. VRF B uses an IS-IS accept policy to accept these routes.
- Each node learns three paths to the route:
 - The nodes learn one path using local inter-VRF redistribution.
 - The nodes learn the other two paths from the IST peer.
- The routes are treated as ECMP paths because the preference, metric-type, and metric are equal.

IS-IS sorts paths for the same route by source-BEB B-MAC and B-VLAN ID. The primary B-VLAN ID is first installed for each B-MAC, followed by the secondary B-VLAN ID for each B-MAC, as long as the ECMP max-path value is not reached. On the node with the lowest B-MAC, the first path listed is its own local inter-vrf route, while on the other node, the MIM path across the vIST is listed first.

If you disable ECMP, all but the first path is removed. Because IS-IS orders the paths by B-MAC, each node in the vIST cluster selects the same B-MAC as the nexthop. This configuration leads one of the nodes to select itself, the local inter-vrf route, while the other node selects the MIM path across the vIST to get to the inter-vrf route. This situation results in an incorrect route selection.

Avoiding Incorrect Route Selection

To avoid this situation, create a policy to prevent IS-IS from determining that the MIM path across the vIST and the local inter-VRF route are ECMP paths. Configure the local inter-VRF path as the preferred path, and the vIST path as the backup. The following list identifies way that you can accomplish this:

- Redistribute the VRF route into IS-IS using the internal metric-type. IS-IS will always select the local inter-VRF route. For more information about the metric type for IS-IS routes, see *Configuring Fabric Basics and Layer 2 Services*.
- If an IS-IS internal metric-type is not an option, configure an IS-IS accept policy to change the preference of inter-VRF routes learned from the IST peer. The local inter-VRF route is preferred over the inter-VRF routes learned from the IST peer.

IS-IS redistribution policies

In the VRF, just like in the Global Router, the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP, within the context of a VRF. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

The VRF specific routes are transported in TLV 184 with the I-SID assigned to the VPNs. After extracting the IP VPN IP reachability information, the routes are installed in the route tables of the appropriate VRFs based on the I-SID association.

For each VRF, the next-hop for the installed VPN routes is the node from which the LSPs that carry the IP VPN routes with the same I-SID as the VRF are received. For the IP VPN, the next hop IP address is the internally generated IP address that corresponds to the nodal BMAC of the next hop that creates the virtual ARP for the node MAC address.

To make IS-IS retrieve the routes from the routing table of a specific VRF for which you enable IP VPN, and advertise the routes to IS-IS peers, use route redistribution and route policies. If you only need to advertise a subset of routes from a specific route type, use route policies, but under the specific VRF context.

The following example shows the configuration to export routes from directly connected interfaces into IS-IS from the SPBM cloud:

```
IP REDISTRIBUTION CONFIGURATION - VRF

router vrf blue
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
```

The following example shows the configuration to distribute IS-IS learned routes into BGP in a VRF context:

```
BGP CONFIGURATION - VRF

router vrf green
ip bgp
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf green
ip bgp redistribute isis
exit

ip bgp redistribute isis enable
```


Interconnection with OSPF or RIP networks

When you connect an SPBM core using Layer 3 VSNs to existing networks that run a routing protocol such as OSPF or RIP, a redundant configuration requires two switches:

- Both routers redistribute IP routes from Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) into IS-IS (IP) and redistribute IS-IS (IP) routes into RIP or OSPF. This can create a routing loop, special precaution need to be taken to prevent this.

The following figure illustrates this configuration.




Figure 10: Redundant OSPF or RIP network

In this scenario you must take extra care when redistributing through both switches. By default the preference value for IP routes generated by SPBM-IP (IS-IS) is 7. This is a higher preference than OSPF (20 for intra-area, 25 for inter-area, 120 for ext type1, 125 for ext type2) or RIP (100).

! Important:

The lower numerical value determines the higher preference.

In the preceding diagram both nodes (Switch-G and Switch-D) have an OSPF or a RIP route to 192.168.10.0/24 with the next-hop to Switch-A.

As soon as the Switch-G node redistributes that IP route into IS-IS, the Switch-D node learns the same route through IS-IS from Switch-G. (The Switch-G node already has the route through OSPF or RIP). Because IS-IS has a higher preference, Switch-D replaces its 192.168.10.0 OSPF route with an IS-IS one that points at Switch-G as the next-hop. The following figure illustrates this scenario.




Figure 11: Redistributing routes into IS-IS

This situation is undesirable and you must ensure that the two redistributing nodes (Switch-G and Switch-D) do not accept redistributed routes from each other. With IS-IS accept policies, you can associate an IS-IS accept policy on Switch-D to reject all redistributed IP routes received from Switch-G, and Switch-G to reject all redistribute IP routes from Switch-D.

ISIS Accept configuration used on Switch-G

```

router isis
  redistribute ospf
  redistribute ospf enable
exit
isis apply redistribute ospf

router ospf
  as-boundary-router enable
  redistribute isis
  redistribute isis enable
exit
ip ospf apply redistribute isis

route-map "reject" 1
  no permit
  enable
exit
router isis
  accept adv-rtr <SPB nickname of Switch-D>
  accept adv-rtr <SPB nickname of Switch-D> route-map "reject"
  accept adv-rtr <SPB nickname of Switch-D> enable
exit
isis apply accept

```

ISIS Accept configuration used on Switch-D

```

router isis

```

Layer 3 VSN configuration

```
    redistribute ospf
    redistribute ospf enable
exit
isis apply redistribute ospf

router ospf
  as-boundary-router enable
  redistribute isis
  redistribute isis enable
exit
ip ospf apply redistribute isis

route-map "reject" 1
  no permit
  enable
exit
router isis
  accept adv-rtr <SPB nickname of Switch-G>
  accept adv-rtr <SPB nickname of Switch-G> route-map "reject"
  accept adv-rtr <SPB nickname of Switch-G> enable
exit
isis apply accept
```

* Note:

Disable alternative routes by issuing the command **no ip alternative-route** to avoid routing loops on the SMLT Backbone Edge Bridges (BEBs).

In the preceding figure, if Switch-C advertises 25000 IS-IS routes to Switch-G and Switch-D, then both Switch-G and Switch-D install the 25000 routes as IS-IS routes. Since Switch-D and Switch-G have IS-IS to OSPF redistribution enabled, they also learn these 25000 routes as OSPF routes from each other. The OSPF route preference for external (Type1 or Type2) routes normally has a higher numerical value (120 or 125) than the default IS-IS route preference (7), so Switch-D and Switch-G keep the OSPF learned routes as alternative routes.

If Switch-C withdraws its 25000 IS-IS routes, Switch-G and Switch-D remove the IS-IS routes. While the IS-IS routes are removed the routing tables of Switch-G and Switch-D activate the alternative OSPF routes for the same prefix. Since Switch-G and Switch-D also have OSPF to IS-IS redistribution enabled, Switch-C will briefly learn these routes as IS-IS from both Switch-G and Switch-D and this causes a temporary, transient routing loop. This is because the alternative OSPF routes existed because they were redistributed from IS-IS in the first place, before the IS-IS route was withdrawn by Switch-B. To avoid these issues, it is better to simply disable alternative routes on redundant routers which are redistributing the same routes between two different routing protocols. To do this use the **no ip alternative-route** command to disable alternative routes on Switch-G and Switch-D to avoid routing loops.

```
no ip alternative-route
```

The following example demonstrates how to redistribute a default route, instead of all individual IS-IS routes, into an access OSPF or RIP network. In this example a RIP network example is used first then with OSPF. The following figure and sample configuration example illustrates this scenario.




Figure 12: Redistributing routes into RIP

Switch-G

```
enable
configure terminal

IP PREFIX LIST CONFIGURATION
ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION
route-map "inject-default" 1
permit
    set injectlist "default"
enable
exit

IP REDISTRIBUTION CONFIGURATION

router rip
    redistribute isis
    redistribute isis route-map "inject-default"
    redistribute isis enable
exit

RIP PORT CONFIGURATION

interface GigabitEthernet 1/12
ip rip default-supply enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip rip apply redistribute isis
```

Switch-A

```
RIP PORT CONFIGURATION

interface gigabitethernet 1/2
ip rip default-listen enable
exit

interface gigabitethernet 1/3
ip rip default-listen enable
exit
```

Switch-D

```
enable
configure terminal

IP PREFIX LIST CONFIGURATION

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION

route-map "inject-default" 1
permit
    set injectlist "default"
enable
exit

IP REDISTRIBUTION CONFIGURATION

router rip
redistribute isis
redistribute isis route-map "inject-default"
redistribute isis enable
exit

RIP PORT CONFIGURATION

interface GigabitEthernet 1/12
ip rip default-supply enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip rip apply redistribute isis
```

You can control the propagation of the default route on the RIP network so that both Switch-G and Switch-D supply the default route on their relevant interfaces, and not accept it on the same interfaces. Likewise, Switch-A will accept the default route on its interfaces to both Switch-G and Switch-D but it will not supply the default route back to them.

The preceding example where IS-IS IP routes are aggregated into a single default route when redistributed into the RIP network also applies to redistributing IS-IS IP routes into OSPF if that OSPF network is an access network to an SPBM core. In this case use the following redistribution policy configuration as an example for injecting IS-IS IP routes into OSPF:

```
enable
configure terminal

IP PREFIX LIST CONFIGURATION
```

```

ip prefix-list "default" 0.0.0.0/0 ge 0 le 32

IP ROUTE MAP CONFIGURATION

route-map "inject-default" 1
permit
set injectlist "default"
enable
exit

IP REDISTRIBUTION CONFIGURATION

router ospf
redistribute isis
redistribute isis route-map "inject-default"
redistribute isis enable
exit

OSPF CONFIGURATION

router ospf
ip ospf as-boundary-router enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

ip ospf apply redistribute isis

```

IS-IS accept policies

You can use Intermediate-System-to-Intermediate-System (IS-IS) accept policies to filter incoming IS-IS routes over the SPBM cloud and apply route policies to the incoming IS-IS routes. IS-IS accept policies enable the device to determine whether to add an incoming route to the routing table.

IS-IS accept policies and DvR

When you configure DvR in an SPB network, you can leverage IS-IS accept policies to control the DvR routes learned from the DvR backbone. The DvR backbone contains the master list of all the host routes learned from various DvR domains.

You can configure accept policies on a DvR Controller or a non-DvR BEB as a filter to determine which DvR host routes to accept into the routing table, from the DvR backbone. Accept policies apply to only those backbone (or inter-domain) host routes that are not part of the Controller's own DvR enabled subnets *and* do not have the same domain ID as that of the Controller.

For non-DvR BEBs, all the routes present in the backbone are learned, but you can still use the accept policies to filter specific routes.

For information on DvR, see *Configuring IPv4 Routing*.

IS-IS accept policy filters

You can filter traffic with IS-IS accept policies by:

- advertising BEB
- I-SID or I-SID list
- route-map

- backbone-route-map
- a combination of route-map and backbone-route-map

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

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

IS-IS accept policies can also apply route policies to determine what incoming traffic to accept into the routing table. With route policies the device can determine which routes to accept into the routing table based on the criteria you configure. You can match on the network or the route metric.

On DvR Controllers in a DvR domain, you can configure a backbone route policy to determine what host routes to accept from the DvR backbone, into the routing table. Also, just like on the route policy, you can configure match criteria, and set preferences on the backbone route policy.

To accept both IS-IS routes and host routes from the DvR backbone, you can configure both a route policy and a backbone route policy in the accept policy instance.

For more information on configuring route policies, see *Configuring IPv4 Routing*.

The following table describes IS-IS accept policy filters.

Filters into	Filter	Description
Global Routing Table (GRT)	accept route-map <i>WORD</i> <1-64>	By default, the device accepts all routes into the GRT and VRF routing table. This is the default accept policy.
	accept route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	This is the default accept policy with configuration to accept specific DvR host routes from the DvR backbone.
	accept adv-rtr <x.xx.xx> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB defined by the SPBM nickname. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept i-sid <1-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the I-SID, which represents a local or remote Layer 3 VSN. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept adv-rtr <x.xx.xx> i-sid <1-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN.

Table continues...

Filters into	Filter	Description
		The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the list of I-SIDs. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	accept adv-rtr <x.xx.xx> isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
Virtual Routing and Forwarding (VRF) routing table	isis accept adv-rtr <x.xx.xx> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB defined by the SPBM nickname. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept i-sid <0-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept adv-rtr <x.xx.xx> i-sid <0-16777215> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the specific advertising BEB and the I-SID, which represents a local or remote Layer 3 VSN. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept isid-list <i>WORD</i> <1-32> route-map <i>WORD</i> <1-64> backbone-route-map <i>WORD</i> <1-64>	The device filters based on the list of I-SIDs to which the IS-IS accept policy applies. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.

Table continues...

Filters into	Filter	Description
	isis accept adv-rtr <x.xx.xx> isid-list WORD<1-32> route-map WORD<1-64> backbone-route-map WORD<1-64>	The device filters based on the specific advertising BEB and the list of I-SIDs. The number 0 represents the Global Routing Table (GRT). The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.
	isis accept route-map WORD<1-64> route-map WORD<1-64> backbone-route-map WORD<1-64>	The device filters based on the route policy. The device, if DvR enabled, also filters the DvR host routes to accept from the DvR backbone. This is an optional filter.

IS-IS accept policies for the GRT and VRFs

You can create an IS-IS accept policy for incoming routes for the Global Routing Table (GRT), which accepts routes into the routing table, or for a Virtual Routing and Forwarding (VRF) instance, which accepts incoming routes to the routing table of the VRF.

If you create an IS-IS accept policy on the switch for either the GRT or a VRF that operates at a global default level, the accept policy applies to all routes for all BEBs in the GRT or VRF.

If you create an IS-IS accept policy on the switch for a specific advertising BEB for either the GRT or a VRF, the IS-IS accept policy instance applies for that specific advertising BEB. If you use a more specific filter, the system gives preference to the specific filter over the global default level.

IS-IS accept policies for inter-VRF route redistribution

You can also use the filter mechanism for IS-IS accept policies to redistribute routes between different VRFs, or between a VRF and the GRT. For inter-VRF route redistribution, you match the filter based on the I-SID, which represents the Layer 3 VSN context.

You can apply the filter at the global default level, where the IS-IS accept policy applies to all routes for that I-SID from all BEBs, or at a specific advertising BEB level, where the filter only applies to a specific advertising BEB. The device gives preference to a specific filter for a specific advertising BEB over the global default filter.

For inter-VRF route redistribution, an I-SID value of 0 represents the GRT. For inter-VRF route redistribution between VRFs, the I-SID is the source VRF (or remote VRF).

IS-IS accept policy considerations

Consider the following when you configure IS-IS accept policies:

- The switch does not support IS-IS accept policies for IPv6 addresses.
- If a VRF uses a different protocol to redistribute routes from another VRF, the IS-IS accept policy feature cannot be used. You can only use the IS-IS accept policy for inter-VSN route redistribution between VRFs.

Precedence rules in the same VSN

The following precedence rules apply for IS-IS accept policies used in the same VSN:

- You can only apply one configured IS-IS accept policy for each route.

- You can apply either a default filter for all advertising BEBs or a filter for a specific advertising BEB.
- If you disable the accept filter, the system ignores the filter and the filter with the next highest precedence applies.
- The device prefers the `accept adv-rtr` filter, which filters based on a specific advertising BEB, over the default filter for all advertising BEBs.
- The device accepts all routes within the same VSN by default. You can apply a route policy to filter or change the characteristics of the route by metric or preference.
- The `i-sid` or `isid-list` filters are not valid for routes within the same VSN.

Precedence rules for inter-VSN route redistribution

The following precedence rules apply for IS-IS accept policies used for inter-VSN route redistribution:

- You can only apply one configured IS-IS accept policy for each route.
- You can apply filters at a global default level for all BEBs for a specific I-SID or I-SID list, or you can apply filters for a specific advertising BEB for a specific I-SID or I-SID list.
- If you disable the accept filter, the system ignores the filter and the filter with the next highest precedence applies.
- The device requires a specific filter to redistribute routes between VSNs through the use of the `i-sid` or `isid-list` filters.
- The `i-sid` filter takes precedence over the `isid-list` filter.
- The `adv-rtr` filter for a specific advertising BEB takes precedence over a filter with the same `i-sid` filter without the `adv-rtr` filter.
- The `i-sid` or `isid-list` filters only apply to routes for inter-VSN route redistribution.
- If multiple `isid-list` filters have the same I-SID within the list, the first on the list alphabetically has the higher precedence.

Route preference

The relative value of the route preference among different protocols determines which protocol the device prefers. If multiple protocols are in the routing table, the device prefers the route with the lower value. You can change the value at the protocol level, and you can also change the preference of incoming ISIS routes using the route-map with the ISIS Accept policy filter.

Route metric

Use route-map to change the metric of a route when you accept a remote IS-IS route with IS-IS accept policies.

You can use route-map to change the metric of a route when you redistribute the route from another protocol to IS-IS through the route redistribution mechanism.

You can also configure the route metric with the base `redistribute` command without the use of route-map.

For more information on the configuration of route-map, see *Configuring IPv4 Routing*.

Enable/disable ICMP Response on VRFs/Layer 3 VSNs

This feature supports VRFs/Layer 3 VSNs to operate in stealth mode by disabling ICMP responses on specific VRFs/Layer 3 VSNs.

If the ICMP response is disabled, the switch does not respond to any ICMP requests received on the VRFs/Layer 3 VSNs.

If the ICMP response is enabled, the switch responds to ICMP requests received on the VRF/Layer 3 VSNs.

Layer 3 VSN configuration using the CLI

This section provides a procedure to configure Layer 3 VSNs using the command line interface (CLI).

Configuring SPBM IPv4 Layer 3 VSN

After you have configured the SPBM infrastructure, you can enable SPBM Layer 3 VSN to advertise IPv4 routes across the SPBM network from one VRF to another using the following procedure.

SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF. In the VRF, just like in the Global Router (VRF 0), the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

Before you begin

- You must configure the required SPBM IS-IS infrastructure.
- You must configure a VRF on the switch. For more information, see *Configuring IPv4 Routing*.
- You must create the Customer VLANs and add slots/ports.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
router vrf WORD<1-16>
```

2. Create an IPv4 VPN instance on the VRF:

```
ipvpn
```

3. Configure SPBM Layer 3 VSN:

```
i-sid <0-16777215>
```

4. Enable IPv4 VPN on the VRF:

```
ipvpn enable
```

By default, a new IPv4 VPN instance is disabled.

5. Display all IPv4 VPNs:

```
show ip ipvpn [vrf WORD<1-16>] [vrfrids WORD<0-512>]
```

6. Identify routes on the local switch to be announced into the SPBM network:

```
isis redistribute {direct | bgp | ospf | rip | static}
```

7. Enable routes on the local switch to be announced into the SPBM network:

```
isis redistribute {direct | bgp | ospf | rip | static} enable
```

8. If you want to delete or disable the configuration, use the no option:

```
no isis redistribute {direct | bgp | ospf | rip | static}
```

```
no isis redistribute {direct | bgp | ospf | rip | static} enable
```

9. Identify other routing protocols to which to redistribute IS-IS routes:

```
ip {bgp | ospf | rip} redistribute isis
```

10. Enable IS-IS redistribution to other routing protocols::

```
ip {bgp | ospf | rip} redistribute isis enable
```

11. Exit Privileged EXEC mode:

```
exit
```

12. Apply the configured redistribution:

```
isis apply redistribute {direct | bgp | ospf | rip | static} vrf  
WORD<1-16>
```

```
ip bgp apply redistribute isis vrf WORD<1-16>
```

```
ip ospf apply redistribute isis vrf WORD<1-16>
```

```
ip rip apply redistribute isis vrf WORD<1-16>
```

13. Display the redistribution configuration:

```
show ip isis redistribute [vrf WORD<1-16>] [vrfrids WORD<0-512>]
```

Example**Create the IPv4 VPN instance:**

```
Switch:1>enable  
Switch:1#configure terminal  
Switch:1(config)#router vrf vrfred  
Switch:1(config)#ipvpn
```

Layer 3 VSN configuration

```
Switch:1(config)#i-sid 100
Switch:1(config)#ipvpn enable
Switch:1(config)#show ip ipvpn
      VRF Name      : vrfred
      Ipvpn-state   : enabled
      I-sid         : 100

1 out of 2 Total Num of VRF Entries displayed.
Switch:1(config)#isis redistribute ospf
Switch:1(config)#isis redistribute ospf enable
Switch:1(config)#isis redistribute ospf enable
Switch:1(config)#end
Switch:1(config)#isis apply redistribute ospf vrf vrfred
Switch:1(config)#show ip isis redistribute vrf vrfred
=====
                        ISIS Redistribute List - VRF vrfred
=====
SOURCE MET MTYPE      SUBNET  ENABLE LEVEL  RPOLICY
-----
LOC      1   internal  allow   FALSE  11
```

Variable definitions

Use the data in the following table to configure the **show ip ipvpn** command.

Variable	Value
vrf <i>WORD</i> <1-16>	Specifies the VRF name.
vrfids <i>WORD</i> <0-512>	Specifies the VRF ID.

Use the data in the following table to configure the **i-sid** command.

Variable	Value
<0-16777215>	Assigns an I-SID to the VRF being configured. Use the no or default option to remove the I-SID to VRF allocation for this VRF.

Use the data in the following table to configure the **isis redistribute** command.

Variable	Value
{ <i>direct</i> <i>bgp</i> <i>ospf</i> <i>rip</i> <i>static</i> }	Specifies the protocol.
enable	Enables the redistribution of the specified protocol into the SPBM network. The default is disabled. Use the no or default options to disable the redistribution.
metric <0-65535>	Configures the metric (cost) to apply to redistributed routes. The default is 1.
metric-type { <i>external</i> <i>internal</i> }	Configures the type of route to import into the protocol. The default is internal.
route-map <i>WORD</i> <0-64>	Configures the route policy to apply to redistributed routes. Specifies a name.

Table continues...

Variable	Value
subnets {allow suppress}	Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose suppress to advertise subnets aggregated to their classful subnet. Choose allow to advertise the subnets individually with the learned or configured mask of the subnet. The default is allow.

Use the data in the following table to configure the `isis apply redistribute` command.

Variable	Value
{direct bgp ospf rip static}	Specifies the protocol.
vrf WORD<1–16>	Applies IS-IS redistribute for a particular VRF. Specifies the VRF name.

Configuring IS-IS Accept Policies

Use the following procedure to create and enable IS-IS accept policies to apply to routes from all Backbone Edge Bridges (BEBs) or to all routes from a specific BEB.

Use IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud. Accept policies apply to incoming traffic and determine whether to add the route to the routing table.

If DvR is enabled on your switch, and the switch is either a DvR Controller or a non-DvR BEB within the domain, you can configure IS-IS accept policies to accept specific host routes from the DvR backbone. For information on DvR, see *Configuring IPv4 Routing*.

IS-IS accept policies are disabled by default.

* Note:

- The `isis apply accept [vrf WORD<1–16>]` command can disrupt traffic and cause temporary traffic loss. After you apply `isis apply accept [vrf <1–16>]`, the command reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. You should make all the relevant accept policy changes, and then apply `isis apply accept [vrf WORD<1–16>]` at the end.
- If the route policy changes, you must reapply the IS-IS accept policy, unless the IS-IS accept policy was the last sequence in the configuration.
- The `isis apply accept [vrf WORD<1–16>]` command is not saved in the configuration file. If you use a saved configuration file for IS-IS accept policy configuration, you must apply the `isis apply accept [vrf WORD<1–16>]` command at the end.
- The number of unique Layer 3 VSN I-SIDs used on a BEB is limited to the number of VRFs supported on the switch. This includes the I-SID values used for Layer 3 VSNs and the I-SID values specified for the ISIS accept policy filters, which can be configured using the `ip isid-list [ISID#], accept i-sid <value>`, or `accept adv-rtr <isis nn> i-sid <value>` commands.

The switch supports 24 VRFs by default, so, in a default configuration, you cannot create an `ip isid-list` or accept policy with more than 24 unique I-SID entries. However, the configured VRFs take up an entry, so the formula to calculate the limit is: $[24 \text{ VRF Limit} - (\text{currently configured VRFs})]$. This gives the number of unique I-SIDs that can be used directly in the IS-IS accept policy filters, which you implement with the `ip isid-list` or `accept policy` command. The I-SIDs used for Layer 3 VSNs can be reused in IS-IS accept policy filters without affecting the limit.

If you increase the VRF scaling, you can create more Layer 3 VSNs. For more information about how to increase the number of supported VRFs, see *Configuring IPv4 Routing*. The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see *Release Notes*.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.
- Ensure that DvR is enabled on the switch before you configure an IS-IS accept policy with a backbone route policy, to accept host routes from the DvR backbone.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. **(Optional)** If you want to accept routes from a variety of I-SIDs, create an I-SID list before you create an IS-IS accept policy for the I-SID list:

```
ip isid-list WORD<1-32> [<1-16777215>][list WORD<1-1024>]
```

3. **(Optional)** Delete an I-SID list:

```
no ip isid-list WORD<1-32> [<1-16777215>][list WORD<1-1024>]
```

4. Enter IS-IS Router Configuration mode:

```
enable
configure terminal
router isis
```

Configure IS-IS accept policies with a route policy or a backbone route policy or a combination of both, to determine which routes the IS-IS accept policy applies to.

Configure one of the following types of IS-IS accept policies.

- **An IS-IS accept policy with only the route policy:**

The IS-IS routes are selectively accepted based on the route policy. Since the backbone route policy is not configured, all host routes from the DvR backbone are *denied*.

If you do not configure a route policy, by default, all IS-IS routes are *accepted*.

- **An IS-IS accept policy with only the backbone route policy:**

The DvR host routes from the DvR backbone are selectively accepted based on the backbone route policy. Since the route policy is not configured, all IS-IS host routes are accepted.

If you do not configure a backbone route policy, all host routes from the DvR backbone are *denied*.

- **An IS-IS accept policy with both route policy and backbone route policy:**

IS-IS routes are selectively accepted based on the route policy and host routes from the DvR backbone are selectively accepted based on the backbone route policy.

5. Configure an IS-IS accept policy instance with a route policy.

Use one of the following options:

- a. Create an IS-IS accept policy instance to apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>]
```

- b. Create an IS-IS accept policy instance to apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> [i-sid <1-16777215>][isid-list WORD <1-32>]
```

- c. **(Optional)** Delete an IS-IS accept policy instance:

```
no accept [adv-rtr <x.xx.xx>][i-sid <1-16777215>][isid-list WORD <1-32>]
```

- d. Specify an IS-IS route policy to apply to routes from all BEBs:

```
accept route-map WORD<1-64>
```

- e. Specify an IS-IS route policy to apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx>[route-map WORD<1-64>]
```

- f. **(Optional)** Delete an IS-IS route policy:

```
no accept [adv-rtr <x.xx.xx>] [route-map]
```

- g. Enable an IS-IS route accept instance:

```
accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-list WORD<1-32>]
```

- h. **(Optional)** Disable an IS-IS route accept instance:

```
no accept [adv-rtr <x.xx.xx>][enable][i-sid <1-16777215>][i-sid-list WORD<1-32>]
```

6. Configure an IS-IS accept policy instance with a backbone route policy to accept host routes from the DvR backbone:

*** Note:**

IS-IS accept policies typically apply to all IS-IS routes. However, to accept DvR host routes from the DvR backbone, you *must* explicitly configure the IS-IS accept policy with a backbone route policy.

Use one of the following options:

- a. Create the default IS-IS accept policy instance to accept host routes from the DvR backbone:

```
accept backbone-route-map WORD <1-64>
```

- b. **(Optional)** Delete the default IS-IS accept policy instance with backbone route policy configuration:

```
no accept backbone-route-map
```

- c. Create an IS-IS accept policy instance to accept host routes from the DvR backbone, and apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>] backbone-  
route-map WORD<1-64>
```

- d. **(Optional)** Delete an IS-IS accept policy instance with backbone route policy configuration, which applies to all BEBs for a specific I-SID or I-SID list:

```
no accept [i-sid <1-16777215>][isid-list WORD <1-32>] backbone-  
route-map
```

- e. Create an IS-IS accept policy instance to accept host routes from the DvR backbone and apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> backbone-route-map WORD <1-64>
```

- f. **(Optional)** Delete an IS-IS accept policy instance with backbone route policy configuration, which applies to a specific advertising BEB

```
no accept adv-rtr <x.xx.xx> backbone-route-map
```

7. Configure an IS-IS accept policy with both route policy and backbone route policy, to selectively accept IS-IS routes as well as host routes from the DvR backbone.

- a. Create the default IS-IS accept policy instance with a route policy to accept IS-IS routes and a backbone route policy to accept host routes from the DvR backbone:

```
accept route-map WORD<1-32> backbone-route-map WORD <1-64>
```

- b. **(Optional)** Delete the default IS-IS accept policy with route policy and backbone route policy configuration:

```
no accept route-map backbone-route-map
```

- c. Create an accept policy instance to selectively accept IS-IS routes and host routes from the DvR backbone, and apply to all BEBs for a specific I-SID or I-SID list:

```
accept [i-sid <1-16777215>][isid-list WORD <1-32>] route-map  
WORD<1-32> backbone-route-map WORD<1-64>
```


- d. **(Optional)** Delete an accept policy instance with route policy and backbone route policy configuration, which applies to all BEBs for a specific I-SID or I-SID list:

```
no accept [i-sid <1-16777215>][isid-list WORD <1-32>] route-map
backbone-route-map
```

- e. Create an IS-IS accept policy instance to selectively accept IS-IS routes and host routes from the DvR backbone, and apply to a specific advertising BEB:

```
accept adv-rtr <x.xx.xx> route-map WORD<1-32> backbone-route-map
WORD <1-64>
```

- f. **(Optional)** Delete an IS-IS accept policy instance with route policy and backbone route policy configuration, which applies to a specific advertising BEB:

```
no accept adv-rtr <x.xx.xx> route-map backbone-route-map
```

8. Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

```
isis apply accept [vrf WORD <1-16>]
```

9. Exit IS-IS Router Configuration mode:

```
exit
```

You are in Global Configuration mode.

Example

Configure an I-SID based IS-IS accept policy with the route policy `test`:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#route-map test 1
Switch:1(route-map)#enable
Switch:1(route-map)#exit

Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 101
Switch:1(config-isis)#accept i-sid 101 route-map test
Switch:1(config-isis)#accept i-sid 101 enable
Switch:1#exit
Switch:1(config)#isis apply accept
```

Configuration of IS-IS accept policy to accept host routes from the DvR backbone

Example 1:

To accept host routes from the DvR backbone, you must configure a backbone route policy and apply it to the IS-IS accept policy.

1. Configure a route policy for DvR:

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#route-map dvmap1 1
Switch:1(route-map)#enable
```

2. Configure an IS-IS accept policy for I-SID 10, and apply the route policy as a backbone route policy:

```
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 backbone-route-map dvrmap1
Switch:1(config-isis)#accept i-sid 10 enable
Switch:1(config-isis)#exit
```

OR

Configure the default accept policy for IS-IS and DvR, and apply the route policy as a backbone route policy:

```
Switch:1(config)#route-map isismap1 1
Switch:1(route-map)#enable
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept route-map isismap1 backbone-route-map dvrmap1
```

3. Apply the IS-IS accept policy:

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
Switch:1(config)#exit
```

4. Verify the configuration:

```
Switch:1#show ip isis accept
=====
Isis Accept - GlobalRouter
=====
ADV_RTR  I-SID    ISID-LIST          ENABLE POLICY      BACKBONE
POLICY
-----
-         10      -                  TRUE               dvrmap1
-         -       -                  isismap1          dvrmap1
2 out of 2 Total Num of Isis Accept Policies displayed
```

Example 2:

Configure an IS-IS accept policy for I-SID 10 that accepts DvR host routes in a subnet, for example, subnet 126.1.1.0/24.

1. Configure an IP prefix list:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list listPrefix 126.1.1.0/24
```

2. Create the route policy dvrmap2 to match the IP prefix list:

```
Switch:1(config)#route-map dvrmap2 1
Switch:1(route-map)#match network listPrefix
Switch:1(route-map)#enable
```

3. Create an IS-IS accept policy with I-SID 10 and apply the route policy as a backbone route policy:

```
Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 backbone-route-map dvrmap2
Switch:1(config-isis)#accept i-sid 10 enable
```

4. Apply the IS-IS accept policy:

```
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
```

The above command causes IS-IS to accept all routes with I-SID 10. To deny IS-IS routes and accept only DvR host routes, you can configure an additional IS-IS route policy as follows:

```
Switch:1(config)#route-map isismap2 1
Switch:1(route-map)#no permit
Switch:1(route-map)#enable

Switch:1(route-map)#exit
Switch:1(config)#router isis
Switch:1(config-isis)#accept i-sid 10 route-map isismap2 backbone-route-map dvrmap2
Switch:1(config-isis)#accept i-sid 10 enable
Switch:1(config-isis)#exit
Switch:1(config)#isis apply accept
```

5. Verify the configuration:

```
Switch:1(config)#exit
Switch:1#show ip isis accept
```

```
=====
                        Isis Accept - GlobalRouter
=====
ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
-----  -
-         10      -                        TRUE  isismap2                dvrmap2

1 out of 1 Total Num of Isis Accept Policies displayed
```

Configuration of IS-IS accept policies for a specific VRF instance**Example 1:**

Configure IS-IS accept policies to accept host routes from the DvR backbone, for a specific VRF instance.

1. In the VRF green context, configure the route policy `dvrmap3` for DvR:

```
Switch:1(config)#router vrf green
Switch:1(router-vrf)#route-map dvrmap3 1
Switch:1(router-vrf-routemap)#enable
```

2. Use one of the following options to configure an IS-IS accept policy, and apply the route policy as a backbone route policy:

Configure an IS-IS accept policy for a specific advertising BEB with nickname `1.11.11`:

```
Switch:1(router-vrf-routemap)#isis accept adv-rtr 1.11.11 backbone-route-map
dvrmap3
Switch:1(router-vrf-routemap)#exit
Switch:1(router-vrf)#isis accept adv-rtr 1.11.11 enable
```

```
Switch:1(router-vrf)#show ip isis accept vrf green
```

```
=====
                        Isis Accept - VRF green
=====
ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
```

Layer 3 VSN configuration

```
-----POLICY-----
1.11.11 - - TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
Switch:1(config)#show ip isis accept vrfids 2
=====
Isis Accept - VRF green
=====
ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY
-----
1.11.11 - - TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure an accept policy for I-SID 10:

```
Switch:1(router-vrf)#isis accept i-sid 10 backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
=====
Isis Accept - VRF green
=====
ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY
-----
- 10 - TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure an accept policy for the I-SID list listisids:

```
Switch:1(router-vrf)#isis accept isid-list listisids backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green
=====
Isis Accept - VRF green
=====
ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
POLICY
-----
- 10 listisids TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed
```

Configure the default accept policy for IS-IS and DvR:

```
Switch:1(router-vrf)#route-map isismap3 1
Switch:1(router-vrf-routemap)#
Switch:1(router-vrf-routemap)#enable
Switch:1(router-vrf-routemap)#
Switch:1(router-vrf-routemap)#isis accept route-map isismap3 backbone-route-map
dvrmap3
Switch:1(router-vrf)#
Switch:1(router-vrf)#show ip isis accept vrf green
=====
Isis Accept - VRF green
=====
ADV_RTR I-SID ISID-LIST ENABLE POLICY BACKBONE
```

```

-----POLICY
- - - TRUE isismap3 dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed

Configure the default accept policy for DvR:

Switch:1(router-vrf)#isis accept backbone-route-map dvrmap3
Switch:1(router-vrf)#show ip isis accept vrf green

=====
Isis Accept - VRF green
=====

ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
-----POLICY
- - - TRUE dvrmap3
1 out of 1 Total Num of Isis Accept Policies displayed

```

Example 2:

Configure an accept policy for I-SID 10 that accepts DvR host routes in a subnet, for example, subnet 126.1.1.0/24.

1. Configure an IP prefix list:

```

Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#ip prefix-list listPrefix 126.1.1.0/24

```

2. For a specific VRF instance, create a route policy to match the IP prefix list:

```

Switch:1(config)#router vrf green
Switch:1(router-vrf)#route-map dvrmap4 1
Switch:1(router-vrf-routemap)#match network listPrefix
Switch:1(router-vrf-routemap)#enable
Switch:1(router-vrf-routemap)#exit
Switch:1(router-vrf)#

```

3. Create an IS-IS accept policy with I-SID 10, and apply the route policy as the backbone route policy:

```

Switch:1(router-vrf)#accept i-sid 10 backbone-route-map dvrmap4
Switch:1(router-vrf)#accept i-sid 10 enable

```

4. Apply the IS-IS accept policy:

```

Switch:1(router-vrf)#exit
Switch:1(config)#isis apply accept

```

5. Verify the configuration:

```

Switch:1(config)#exit
Switch:1(router-vrf)#show ip isis accept vrf green

=====
Isis Accept - VRF green
=====

ADV_RTR  I-SID    ISID-LIST                ENABLE POLICY            BACKBONE
-----POLICY
- - - TRUE dvrmap4

```

1 out of 1 Total Num of Isis Accept Policies displayed

Variable definitions

Use the data in the following table to use the `ip isid-list` command.

Variable	Value
<code>WORD<1-32></code>	Creates a name for your I-SID list.
<code><1-16777215></code>	Specifies an I-SID number.
<code>list WORD<1-1024></code>	Specifies a list of I-SID values. For example, in the format 1,3,5,8-10.

Use the data in the following table to use the `accept` command.

Variable	Value
<code>adv-rtr <x.xx.xx></code>	Specifies the SPBM nickname for each advertising BEB to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter.
<code>backbone-route-map WORD<1-64></code>	Specifies the DvR backbone route map.
<code>enable</code>	Enables an IS-IS accept policy.
<code>i-sid <1-16777215></code>	Specifies an I-SID number to represent a local or remote Layer 3 VSN to which the IS-IS accept policy applies. Use the parameter to apply a filter for routes from a specific I-SID that represents the remote VSN. Based on the routing policy the system applies, the system can redistribute the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).
<code>isid-list WORD<1-32></code>	Specifies the I-SID list name that represents the local or remote Layer 3 VSNs to which the IS-IS accept policy applies. Use the parameter to apply a default filter for all routes from a specific I-SID that represents the remote VSN. Based on the routing policy the system applies, the system redistributes the remote VSN to the VSN where you applied the filter. An I-SID value of 0 represents the global routing table (GRT).
<code>route-map WORD<1-64></code>	Specifies a route policy by name.

Table continues...

Variable	Value
	You must configure the route policy earlier in a separate procedure.

Use the data in the following table to use the `isis apply accept` command.

Variable	Value
vrf <i>WORD</i> <1-16>	Specifies a specific VRF instance.

Configuring inter-VRF accept policies on VRFs

Configure IS-IS accept policies on a VRF to use inter-VRF accept policies in the SPB cloud. You can use IS-IS accept policies to redistribute routes between different VRFs, including the global routing table (GRT). First you apply the filter, and then you match the filter based on the I-SID, which represents the Layer 3 VSN context.

* Note:

- The `isis apply accept [vrf WORD<1-16>]` command can disrupt traffic and cause temporary traffic loss. After you apply `isis apply accept [vrf<1-16>]`, the command reapplies the accept policies, which deletes all of the IS-IS routes and adds the IS-IS routes again. You should make all the relevant accept policy changes, and then apply `isis apply accept [vrf WORD<1-16>]` at the end.
- If you use the `accept` command for inter-VRF routes based on the remote I-SID, the device only accepts routes coming from remote BEBs. For instance, if a local Layer 3 VSN exists with the same I-SID, the device does not add the local routes. The assumption is that the device uses existent methods, either through use of another protocol or static configuration, to obtain those routes.
- If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure that a route policy exists.
- Ensure that the VRFs exist.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
router vrf WORD<1-16>
```

2. **(Optional)** If you want to accept routes from a variety of I-SIDs, create an I-SID list before you create an IS-IS accept policy for the I-SID list:

```
ip isid-list WORD<1-32> [<0-16777215>][list WORD<1-1024>]
```

3. Create an IS-IS accept policy instance to apply to routes from all Backbone Edge Bridges (BEBs):

```
isis accept [i-sid <0-16777215>][isid-list WORD<1-32>]
```

4. Create an IS-IS accept policy instance to apply to routes for a specific BEB:

```
isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list WORD<1-32>]
```

5. **(Optional)** Delete an IS-IS accept policy instance:

```
no isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list WORD<1-32>]
```

6. Specify an IS-IS route policy to apply to routes from all BEBs:

```
isis accept route-map WORD<1-64>
```

7. Specify an IS-IS route policy to apply for a specific BEB:

```
isis accept adv-rtr <x.xx.xx> route-map WORD<1-64>
```

8. **(Optional)** Delete an IS-IS route policy:

```
no isis accept [adv-rtr <x.xx.xx>] [route-map]
```

9. Enable a configured IS-IS accept policy instance:

```
isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list WORD<1-32>] [enable]
```

10. **(Optional)** Disable a configured IS-IS accept policy instance:

```
no isis accept [adv-rtr <x.xx.xx>][i-sid <0-16777215>][isid-list WORD<1-32>] [enable]
```

11. Exit VRF Router Configuration mode:

```
exit
```

You are in Global Configuration mode.

12. Apply the IS-IS accept policy changes, which removes and re-adds all routes with updated filters:

```
isis apply accept [vrf WORD<1-16>]
```

Example

Configure Inter-VRF accept policies on a VRF:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#router vrf green
Switch:1(router-vrf)#isis accept i-sid 100
```



```
Switch:1(router-vrf)#isis accept i-sid 100 enable
Switch:1(router-vrf)#exit
Switch:1(config)#isis apply accept vrf green
```

Variable definitions

Use the data in the following table to use the `ip isid-list` command.

Variable	Value
<code>WORD<1-32></code>	Creates a name for your I-SID list.
<code><0-16777215></code>	Specifies an I-SID value.
<code>list WORD<1-1024></code>	Specifies a list of I-SID values. For example, in the format 1,3,5,8-10.

Use the data in the following table to use the `isis accept` command.

Variable	Value
<code>adv-rtr <x.xx.xx></code>	Specifies a specific advertising BEB in which to apply the IS-IS accept policy to routes for a specific advertising BEB. <code>x.xx.xx</code> specifies an SPBM nickname. The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter. The system requires an explicit filter to redistribute routes from a particular VSN. If the default global filter or the filter for a specific advertising BEB does not exist, the system does not redistribute the routes from the remote VSN.
<code>enable</code>	Enables the IS-IS accept policy.
<code>i-sid <0-16777215></code>	Configures the I-SID to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).
<code>isid-list WORD<1-32></code>	Configures a list of I-SIDs to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).
<code>route-map WORD <1-64></code>	Specifies a route policy. You must configure a route policy earlier in a separate procedure.

Use the data in the following table to use the `isis apply accept` command.

Variable	Value
<code>vrf WORD<1-16></code>	Specifies a specific VRF instance.

Configuring SPBM IPv6 Layer 3 VSN

About this task

After you have configured the SPBM infrastructure, you can enable SPBM Layer 3 VSN to advertise IPv6 routes across the SPBM network using the following procedure.

Before you begin

- You must enable IPv6 Shortcuts.
- You must configure the required SPBM IS-IS infrastructure.
- You must configure a VRF instance on the switch. For more information, see *Configuring IPv6 Routing*.
- Not all parameters are available on non-default VRFs.

Procedure

1. Enter VRF Router Configuration mode for a specific VRF context:

```
enable
configure terminal
router vrf WORD<1-16>
```

2. Create an IPv6 VPN instance on the VRF:

```
ipv6 ipvpn
```

3. Configure SPBM Layer 3 VSN:

```
i-sid <0-16777215>
```

4. Enable IPv6 VPN on the VRF:

```
ipv6 ipvpn enable
```

5. Display all IPv6 VPNs:

```
show ipv6 ipvpn [vrf WORD<1-16> | vrfids WORD<0-512>]
```

6. Enable routes on the local switch to be announced into the SPBM network:

```
ipv6 isis redistribute {direct | static} enable
```

7. Display the redistribution configuration:

```
show ipv6 isis redistribute [vrf WORD<1-16> | vrfids WORD<0-512>]
```

8. Verify IPv6 IS-IS routes:

```
show ipv6 route vrf WORD<1-16>
```

Example

Create the IP VPN instance:

```
Switch:1>enable
Switch:1#configure terminal
```

```

Switch:1(config)#router vrf vrfred
Switch:1(router-vrf)#ipv6 ipvpn
Switch:1(router-vrf)#i-sid 100
Switch:1(router-vrf)#ipv6 ipvpn enable
Switch:1(router-vrf)#show ipv6 ipvpn
      VRF Name           : vrfred
      Ipv6 Ipvpn-state    : enabled
      Ipv4 Ipvpn-state    : disabled
      I-sid               : 100

Total active Ipv6 L3 VSN : 1

1 out of 3 Total Num of VRF Entries displayed.
Switch:1(router-vrf)#ipv6 isis redistribute direct enable
Switch:1(router-vrf)#show ipv6 route vrfred
=====
IPv6 Routing Table Information - VRF vrfred
=====
Destination Address/PrefixLen      NEXT HOP          VID/BID/TID      PROTO    COST   AGE
TYPE      PREF
-----
55:0:0:0:0:0:0:0/64              Switch           V-2              ISIS     10
0      B      7
-----
1 out of 1 Total Num of Route Entries displayed.
-----
TYPE Legend:
A=Alternative Route, B=Best Route, E=Ecmp Route

```

Variable definitions

Use the data in the following table to configure the **ipv6 ipvpn** command.

Variable	Value
enable	Enables IPv6 IPVPN. The default is disabled.

Use the data in the following table to configure the **show ipv6 ipvpn** command.

Variable	Value
vrf <i>WORD</i> <1-16>	Specifies the VRF name.
vrfids <i>WORD</i> <0-512>	Specifies the VRF ID.

Use the data in the following table to configure the **i-sid** command.

Variable	Value
<0-16777215>	Assigns an I-SID to the VRF being configured.

Use the data in the following table to configure the **isis redistribute** command.

Variable	Value
{ <i>direct</i> <i>static</i> }	Specifies the protocol.

Table continues...

Variable	Value
enable	Enables the redistribution of the specified protocol into the SPBM network. The default is disabled.

Viewing IS-IS accept policy information

Use the following procedure to view IS-IS accept policy information on the switch.

Procedure

1. Display IS-IS accept policy information:

```
show ip isis accept [vrf WORD<1-16>][vrfids WORD<0-512>]
```

2. Display I-SID list information:

```
show ip isid-list [vrf WORD<1-16>][vrfids WORD<0-512>][WORD<1-32>]
```

3. Display route information:

```
show ip route [vrf WORD<1-16>]
```

The NH VRF/ISID column displays the I-SID for inter-Virtual Services Network (VSN) routes redistributed with IS-IS accept policies, only if the I-SID redistributed does not have an IP VSN associated with it. If an IP VSN exists for that I-SID, the VRF name displays. If the I-SID is 0, the column represents and displays as the GlobalRouter.

The existing IS-IS routes for Layer 3 VSNs continue to display as the VRF name of the IP VSN.

4. Display the SPBM IP unicast Forwarding Information Base (FIB):

```
show isis spbm ip-unicast-fib [all] [id <1-16777215>][spbm-nh-as-mac]
```

Example

View IS-IS accept policy information:

```
Switch:1#show ip route vrf test
=====
                        IP Route - VRF test
=====
DST          MASK          NEXT          NH          INTER
VRF/ISID    COST  FACE  PROT  AGE  TYPE  PRF
-----
1.1.1.5      255.255.255.255  1.1.1.5      GlobalRouter  0    0    ISIS  0    IB    200
1.1.1.13     255.255.255.255  Switch13     GRT          10   1000 ISIS  0    IBSV  7
1.1.1.200    255.255.255.255  Switch200    GRT          10   1000 ISIS  0    IBSV  7
5.7.1.0      255.255.255.0    5.7.1.1      -            1    7    LOC   0    DB    0
13.7.1.0     255.255.255.0    Switch13     GlobalRouter  10   1000 ISIS  0    IBSV  7
100.0.0.0    255.255.255.0    100.0.0.1    GlobalRouter  0    100  ISIS  0    IB    200
111.1.1.0    255.255.255.0    111.1.1.1    hub          0    111  ISIS  0    IB    200

Switch:1(config)#show isis spbm ip-unicast-fib
```

```

=====
                        SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF      VRF      DEST      Destination  NH BEB  VLAN  OUTGOING  SPBM  PREFIX  IP ROUTE
ISID     ISID                                     INTERFACE  COST  COST    PREFERENCE
-----
GRT      -         101       1.1.1.13/32  Switch13 1000  1/7      10    44      7
GRT      -         101       1.1.1.13/32  Switch13 1001  1/7      10    44      7
-----

Total number of SPBM IP-UNICAST FIB entries 2
-----
Switch:1(config)#show ip isid-list test
=====
                        IP ISID LIST
=====
List Name      I-SID      VRF
-----
test           1          GlobalRouter
              3          GlobalRouter
              4          GlobalRouter
              5          GlobalRouter
              10         GlobalRouter
              22         GlobalRouter

All 6 out of 6 Total Num of Isid Lists displayed
Switch:1(router-vrf)#show ip isid-list vrf red
=====
                        IP ISID LIST red
=====
List Name      I-SID      VRF
-----
test1          11         1
              12         1
              13         1
              14         1
              15         1

```

Variable definitions

Use the data in the following table to use the **show ip isis accept** command.

Variable	Value
vrf <i>WORD</i> <1-16>	Displays I-SID list information for a particular VRF by name.
vrfids <i>WORD</i> <0-512>	Displays I-SID list information for a particular VRF ID.

Use the data in the following table to use the **show ip isid-list** command.

Variable	Value
vrf <i>WORD</i> <1-16>	Displays I-SID list information for a particular VRF by name.

Table continues...

Variable	Value
vrfids <i>WORD</i> <0-512>	Displays I-SID list information for a particular VRF ID.
<i>WORD</i> <1-32>	Displays I-SID list information for a particular I-SID list name.

Use the data in the following table to use the `show ip route` command.

Variable	Value
vrf <i>WORD</i> <1-16>	Displays I-SID list information for a particular VRF by name.

Use the data in the following table to use the `show isis spbm ip-unicast-fib` command.

Variable	Value
all	Displays all IS-IS SPBM IP unicast Forwarding Information Base (FIB) information.
id <1-16777215>	Displays IS-IS SPBM IP unicast FIB information by I-SID ID.
spbm-nh-as-mac	Displays the next hop B-MAC of the IP unicast FIB entry.

Displaying SPBM IPv6 unicast Forwarding Information Base

About this task

Perform this procedure to display SPBM IPv6 unicast Forwarding Information Base (FIB).

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Display SPBM IPv6 unicast FIB:

```
show isis spbm ipv6-unicast-fib [all] [id <1-16777215>] [spbm-nh-as-mac]
```

Example

```
Switch:1>enable
Switch:1#show isis spbm ipv6-unicast-fib all
```

```

=====
=====
                                SPBM IPV6-UNICAST FIB ENTRY INFO
=====
=====
PREFIX                               OUTGOING          SPBM
VRF      ISID  Destination                NH BEB          VLAN  INTERFACE      COST
COST
-----
=====

```

```

GRT      -      1111:0:0:0:0:0:1111/128      VSP-8284XSQ      2      1/1
10      1
GRT      -      1111:0:0:0:0:0:1111/128      VSP-8284XSQ      3      1/1
10      1
test     5555    5555:0:0:0:0:0:5555/128      VSP-8284XSQ      2      1/1
10      1
test     5555    5555:0:0:0:0:0:5555/128      VSP-8284XSQ      3      1/1
10      1
-----
-----

```

Variable definitions

Use the data in the following table to use the `show isis spbm ipv6-unicast-fib` command.

Variable	Value
all	Displays all IS-IS SPBM IPv6 unicast Forwarding Information Base (FIB) information for all VRFs.
id <1-16777215>	Displays IS-IS SPBM IPv6 unicast FIB information by I-SID ID.
spbm-nh-as-mac	Displays the next hop B-MAC of the IPv6 unicast FIB entry.

Layer 3 VSN configuration using EDM

This section provides procedures to configure Layer 3 Virtual Services Networks (VSNs) using Enterprise Device Manager (EDM).

Configuring SPBM IPv4 Layer 3 VSN

After you have configured the SPBM infrastructure, you can enable SPBM Layer 3 Virtual Services Network (VSN) to advertise IPv4 routes across the SPBM network from one VRF to another using the following procedure.

SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF. In the VRF, just like in the Global Router (VRF 0), the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

Before you begin

- You must configure the required SPBM IS-IS infrastructure.
- You must configure a VRF and IP VPN instance on the switch. For more information, see *Configuring IPv4 Routing*.
- You must create the Customer VLANs and add slots/ports.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IP-VPN**.
3. Click the **VPN** tab.
4. To create an IP VPN instance, click **Insert**.
5. Click the ellipsis button (...), select a VRF to associate with the IP VPN, and click **Ok**.
6. Click **Insert**.
7. In the **Enable** column, select **enable** to enable the IP VPN on the VRF.
8. In the **IsidNumber** column, specify an I-SID to associate with the VPN.
9. Click **Apply**.
10. In the navigation pane, expand the **Configuration > IP** folders.
11. Click **Policy**.
12. To identify routes on the local switch to be announced into the SPBM network, click the **Route Redistribution** tab.
13. Click **Insert**.
14. In the **DstVrflid** box, click the ellipsis button (...), select the destination VRF ID and click **Ok**.
15. In the **Protocol** box, click **isis** as the route destination.
16. In the **SrcVrflid** box, click (...) button, select the source VRF ID and click **Ok**.
17. In the **RouteSource** box, click the source protocol.
18. In the **Enable** box, click **enable**.
19. In the **RoutePolicy** box, click the ellipsis (...) button, choose the route policy to apply to the redistributed routes and click **Ok**.
20. Configure the other parameters as required.
21. Click **Insert**.
22. To apply the redistribution configuration, click the **Applying Policy** tab.
23. Select **RedistributeApply**, and then click **Apply**.

Configuring IS-IS redistribution

Use this procedure to configure IS-IS redistribution. In the Virtual Routing and Forwarding (VRF), just like in the Global Router, the routes are not redistributed into IS-IS automatically. To advertise the VRF routes, you must explicitly redistribute one of the following protocols into IS-IS: direct, static, RIP, OSPF, or BGP, within the context of a VRF. Routing between VRFs is also possible by using redistribution policies and injecting routes from the other protocols.

The VRF specific routes are transported in TLV 184 with the I-SID assigned to the VPNs. After extracting the IP VPN IP reachability information, the routes are installed in the route tables of the appropriate VRFs based on the I-SID association.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Redistribute** tab.
4. Click **Insert**.
5. Complete the fields as required.
6. Click **Insert**.

IS-IS Redistribute field descriptions

Use the data in the following table to configure the **IS-IS Redistribute** tab.

Name	Description
DstVrflid	Specifies the destination Virtual Routing and Forwarding (VRF) ID used in the redistribution.
Protocol	Specifies the protocols that receive the redistributed routes.
SrcVrflid	Specifies the source VRF ID used in the redistribution. For IS-IS, the source VRF ID must be the same as the destination VRF ID.
RouteSource	Specifies the source protocol for the route redistribution entry.
Enable	Enables or disables a redistribution entry. The default is disable.
RoutePolicy	Specifies the route policy to be used for the detailed redistribution of external routes from a specified source into the IS-IS domain.
Metric	Specifies the metric for the redistributed route. The value can be a range between 0 to 65535. The default value is 0. Use a value that is consistent with the destination protocol.
MetricType	Specifies the metric type. Specifies a type1 or a type2 metric. For metric type1, the cost of the external routes is equal to the sum of all internal costs and the external cost. For metric type2, the cost of the external routes is equal to the external cost alone. The default is type2.
Subnets	Indicates whether the subnets are advertised individually or aggregated to their classful subnet. Choose suppress to advertise subnets aggregated to their classful subnet. Choose allow to advertise the subnets individually with the learned or configured mask of the subnet. The default is allow.

Applying IS-IS accept policies globally

Apply IS-IS accept policies globally. Use IS-IS accept policies to filter incoming IS-IS routes the device receives over the SPBM cloud. Accept policies apply to incoming traffic and determine whether to add the route to the routing table.

After you apply the IS-IS accept filters, the device removes and re-adds all routes with updated filters.

IS-IS accept policies are disabled by default.

*** Note:**

- After you apply IS-IS accept policies globally the application can disrupt traffic and cause temporary traffic loss. After you configure the IS-IS accept policies value to **Apply**, the device reapplies the accept policies, which deletes all of the IS-IS routes, and adds the IS-IS routes again. You should make all the relevant accept policy changes, and then apply IS-IS accept policies globally at the end.
- If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- Ensure the IP IS-IS filter exists.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Global** tab.
4. Select a name from the list or enter name in the **DefaultPolicyName** field to specify the route policy name for the default filter.
5. Select **Apply** to apply the default policy.

Accept Global field descriptions

Use the data in the following table to configure the **Accept Global** tab.

Name	Description
DefaultPolicyName	Specifies the route policy name for the default filter.
DefaultBackbonePolicyName	Specifies the backbone host route policy name for the default filter.

Table continues...

Name	Description
Apply	Applies the default policy when you configure the field to apply. The device only activates the default policy if the route map (the default policy name) has a value. If you do not select apply, the device takes no action. The GRT always returns no action.

Configuring an IS-IS accept policy for a specific advertising BEB

Configure an IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB). Specify the SPBM nickname and the IS-IS accept policy name to allow you to apply the IS-IS accept policy.

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.

* Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Nick Name** tab.
4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.
8. Click **Insert**.

Accept Nick Name field descriptions

Use the data in the following table to configure the **Accept Nick Name** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter for a specific advertising BEB is present the device applies the specific filter. The value is 2.5 bytes in the format <x.xx.xx>.
Enable	Enables or disables the SPBM nickname advertising router entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies a route policy. You must configure a policy earlier in a separate procedure.
BackbonePolicyName	Specifies the route policy for the backbone routes. You must configure a policy earlier in a separate procedure.

Configuring an IS-IS accept policy for a specific advertising BEB and I-SID

Configures a specific advertising Backbone Edge Bridge (BEB) with a specific I-SID to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB.

* Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Nick-Name Isid** tab.
4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Isid** field, specify an I-SID number.

7. Select enable in the **Enable** check box to enable the filter.
8. In the **PolicyName** field, specify the route-map name.
9. Click **Insert**.

Accept Nick-Name Isid descriptions

Use the data in the following table to configure the **Accept Nick-Name Isid** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The value is 2.5 bytes in the format <x.xx.xx>.
Isid	Specifies an I-SID used to filter. The value 0 is used for the Global Router.
Enable	Enables or disables the I-SID entry. The default is disabled.
PolicyName	Specifies the route policy name. You must configure a policy earlier in a separate procedure.
BackBonePolicyName	Specifies the backbone route policy name. You must configure a policy earlier in a separate procedure.

Configuring an I-SID list for an IS-IS accept policy

Configures a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IS-IS accept policy applies. After you create the list of I-SID numbers, you must then create, configure, and enable the IS-IS accept policy.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Isid-List** tab.
4. Click **Insert**.
5. In the **Name** field, specify a name for the I-SID list.
6. Select **Isid** or **Isid-List**.
7. Specify an I-SID number or a list of I-SID numbers.
8. Click **Insert**.

Isid-List field descriptions

Use the data in the following table to configure the **Isid-List** tab.

Name	Description
Name	Specifies the name of the I-SID list.
Isid or Isid-List	Specifies that you either want to add a particular I-SID or a list of I-SID numbers.
Isid	Specifies a particular I-SID number or a list of I-SID numbers that represent local or remote Layer 3 VSNs to which the IS-IS accept policy applies. An I-SID value of 0 represents the global routing table (GRT).

Configuring an IS-IS accept policy for a specific I-SID list

Configure an IS-IS accept policy for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.

* Note:

If the route policy changes, you must re-apply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Isid-List** tab.
4. Click **Insert**.
5. In the **Name** field, specify the I-SID list name.
6. Select enable in the **Enable** check box to enable the filter.
7. In the **PolicyName** field, specify the route-map name.
8. Click **Insert**.

Accept Isid-List field descriptions

Use the data in the following table to configure **Accept Isid-List** tab.

Name	Description
Name	Specifies the name of I-SID list.
Enable	Enables or disables the I-SID list entry. The value must be enabled to filter. The default is disabled.
PolicyName	Specifies the route policy name.
BackBonePolicyName	Specifies the backbone route policy name.

Configuring an IS-IS accept policy for a specific advertising BEB and I-SID-list

Configure an IS-IS accept policy to apply to a specific advertising Backbone Edge Bridge (BEB) for a specific I-SID list to represent local or remote Layer 3 VSNs, which allows the system to redistribute the remote VSNs to the VSN where you applied the filter.

* Note:

If the route policy changes, you must reapply the IS-IS accept policy, unless it was the last sequence in the configuration.

Before you begin

- Enable IS-IS globally.
- Ensure the manual area exists.
- You must configure a route-map to apply. For more information, see *Configuring IPv4 Routing*.

About this task

The system uses the default global filter unless a filter for a specific advertising BEB exists, in which case the system applies a more specific filter.

Procedure

1. In the navigation pane, expand the **Configuration > IP** folders.
2. Click **IS-IS**.
3. Click the **Accept Nick-Name Isid-List** tab.
4. Click **Insert**.
5. In the **AdvertisingRtr** field, specify the SPBM nickname.
6. In the **Name** field, specify an I-SID list name.
7. Select enable in the **Enable** check box to enable the filter.
8. In the **PolicyName** field, specify the route-map name.
9. Click **Insert**.

Accept Nick–Name Isid-List field descriptions

Use the data in the following table to configure the **Accept Nick-Name Isid-List** tab.

Name	Description
AdvertisingRtr	Specifies the SPBM nickname to allow you to apply the IS-IS accept policy to routes for a specific advertising BEB. The system first uses the default filter, but if a more specific filter is present the device applies the specific filter. The value is 2.5 bytes in the format <x.xx.xx>.
Name	Specifies the name of the I-SID list used to filter.
Enable	Enables or disables the SPBM nickname advertising router entry. You must enable the value to filter. The default is disabled.
PolicyName	Specifies a route policy name.
BackBonePolicyName	Specifies a backbone route policy name.

Configuring SPBM IPv6 Layer 3 VSN

About this task

After you have configured the SPBM infrastructure, you can enable SPBM Layer 3 Virtual Services Network (VSN) to advertise IPv6 routes across the SPBM network from one VRF to another using the following procedure.

SPBM Layer 3 VSN uses IS-IS to exchange the routing information for each VRF.

Before you begin

- You must enable IPv6 Shortcuts.
- You must configure the required SPBM IS-IS infrastructure.
- You must configure a VRF and IPv6 VPN instance on the switch. For more information, see *Configuring IPv6 Routing*.
- Not all parameters are configurable on the non-default VRFs.

Procedure

1. In the navigation pane, expand the **Configuration > IPv6** folders.
2. Click **IPv6-VPN**.
3. Click the **VPN** tab.
4. Click **Insert**.
5. Click the ellipsis [...], and select a VRF.
6. Click **Ok**.
7. Click **Insert**.
8. In the **IsidNumber** column, double-click the **0** value, and then enter the service instance identifier (I-SID) to assign to the IPv6-VPN.

9. Click **Apply**.
10. In the **Enable** column, select **true** or **false**.
11. Click **Apply**.
12. In the navigation pane, expand the following folders: **Configuration > VRF Context View**.
13. Click **Set VRF Context View**.
14. Click the **VRF** tab.
15. Select a context to view.
16. Click **Launch VRF Context view**.
A new browser tab opens containing the selected VRF view
17. In the navigation pane, expand the **Configuration > IPv6** folders.
18. Click **IS-IS**.
19. Click the **Redistribute** tab.
20. For the type of route source, double-click the cell in the **Enable** column to change the value.
21. Select **enable**.
22. Click **Apply**.

Layer 3 VSN configuration example

The following figure shows a sample Layer 3 VSN deployment.




Figure 13: Layer 3 VSN

The following sections show the steps required to configure the Layer 3 VSN parameters in this example.

Note that IP IS-IS redistribution needs to be configured to inject the VRF routes into IS-IS.

You must first configure basic SPBM and IS-IS infrastructure.

VRF green configuration

The following figure shows the green VRF in this Layer 3 VSN example.




Figure 14: Layer 3 VSN — VRF green

The following sections show the steps required to configure the green VRF parameters in this example.

VRF green – Switch-C

```
VRF CONFIGURATION
ip vrf green vrfid 1

VLAN CONFIGURATION
vlan create 101 type port-mstprstp 0
vlan mlt 101 1
vlan members 101 1/2 portmember
interface Vlan 101
vrf green
ip address 10.1.101.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION
router vrf green
ipvpn
i-sid 13990001
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF
router vrf green
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS
isis apply redistribute direct vrf green
```

VRF green – Switch-D

```

VRF CONFIGURATION

ip vrf green vrfid 1

VLAN CONFIGURATION

vlan create 102 type port-mstprstp 0
vlan mlt 102 1
vlan members add 102 1/2 portmember
interface vlan 102
vrf green
ip address 10.1.102.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf green
ipvpn
i-sid 13990001
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf green
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct vrf green

```

VRF red configuration

The following figure shows the red VRF in this Layer 3 VSN example.




Figure 15: Layer 3 VSN — VRF red

The following sections show the steps required to configure the red VRF parameters in this example.

VRF red – Switch-C

```

VRF CONFIGURATION

```

Layer 3 VSN configuration

```
ip vrf red vrfid 2

VLAN CONFIGURATION

vlan create 201 type port-mstprstp 0
vlan mlt 201 1
vlan members 201 1/2 portmember
interface Vlan 201
vrf red
ip address 10.2.201.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf red
ipvpn
i-sid 13990002
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf red
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS

isis apply redistribute direct vrf red
```

VRF red – Switch-D

```
VRF CONFIGURATION

ip vrf red vrfid 2

VLAN CONFIGURATION

vlan create 202 type port-mstprstp 0
vlan mlt 101 1
vlan members 202 1/2 portmember
interface Vlan 202
vrf red
ip address 10.3.202.1 255.255.255.0 1
exit

ISIS PLSB IPVPN CONFIGURATION

router vrf red
ipvpn
i-sid 13990002
ipvpn enable
exit

IP REDISTRIBUTION CONFIGURATION - VRF

router vrf red
isis redistribute direct
isis redistribute direct metric 1
isis redistribute direct enable
exit

IP REDISTRIBUTE APPLY CONFIGURATIONS
```

```
isis apply redistribute direct vrf red
```

Verifying Layer 3 VSN operation

The following sections show the steps required to verify the Layer 3 VSN configuration in this example.

Switch-C

```
Switch-C:1# show isis spbm ip-unicast-fib
```

```
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   VRF   DEST          OUTGOING  SPBM   PREFIX  IP ROUTE
VRF  ISID  ISID  Destination NH BEB   VLAN  INTERFACE COST   COST   PREFERENCE
-----
GRT   -    -    10.0.0.2/32  Switch-D 4000  1/3    20     1     7
GRT   -    -    10.0.14.0/24 Switch-D 4000  1/3    20     1     7
-----
Total number of SPBM IP-UNICAST FIB entries 2
=====
```

```
Switch-C:1# show isis spbm ip-unicast-fib id 13990001
```

```
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   VRF   DEST          OUTGOING  SPBM   PREFIX  IP ROUTE
VRF  ISID  ISID  Destination NH BEB   VLAN  INTERFACE COST   COST   PREFERENCE
-----
green -    13990001 10.1.101.0/24 Switch-D 4000  1/2    20     1     7
-----
Total number of SPBM IP-UNICAST FIB entries 1
=====
```

```
Switch-C:1# show isis spbm ip-unicast-fib id 13990002
```

```
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   VRF   DEST          OUTGOING  SPBM   PREFIX  IP ROUTE
VRF  ISID  ISID  Destination NH BEB   VLAN  INTERFACE COST   COST   PREFERENCE
-----
red   -    13990002 10.2.202.0/24 Switch-D 4000  1/3    20     1     7
-----
Total number of SPBM IP-UNICAST FIB entries 1
=====
```

```
Switch-C:1# show isis spbm ip-unicast-fib id all
```

```
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   VRF   DEST          OUTGOING  SPBM   PREFIX  IP ROUTE
VRF  ISID  ISID  Destination NH BEB   VLAN  INTERFACE COST   COST   PREFERENCE
-----
GRT   -    -    10.0.0.2/32  Switch-D 4000  1/3    20     1     7
GRT   -    -    10.0.14.0/24 Switch-D 4000  1/3    20     1     7
green -    13990001 10.1.102.0/24 Switch-D 4000  1/3    20     1     7
red   -    13990002 10.2.202.0/24 Switch-D 4000  1/3    20     1     7
-----
```

Layer 3 VSN configuration

```
Total number of SPBM IP-UNICAST FIB entries 4
```

Switch-D

```
Switch-D:1# show isis spbm ip-unicast-fib
```

```
=====
VRF   VRF   DEST
VRF  ISID  ISID  Destination  NH BEB  VLAN  OUTGOING SPBM  PREFIX IP ROUTE
      ISID  ISID  Destination  NH BEB  VLAN  INTERFACE COST  COST  PREFERENCE
-----
GRT   -    -    10.0.0.1/32  Switch-C 4000  1/2    20    1    7
GRT   -    -    10.0.13.0/24 Switch-C 4000  1/2    20    1    7
-----
Total number of SPBM IP-UNICAST FIB entries 2
=====
```

```
Switch-D:1# show isis spbm ip-unicast-fib id 13990001
```

```
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   VRF   DEST
VRF  ISID  ISID  Destination  NH BEB  VLAN  OUTGOING SPBM  PREFIX IP ROUTE
      ISID  ISID  Destination  NH BEB  VLAN  INTERFACE COST  COST  PREFERENCE
-----
green -    13990001 10.1.101.0/24 Switch-C 4000  1/2    20    1    7
-----
Total number of SPBM IP-UNICAST FIB entries 1
=====
```

```
Switch-D:1# show isis spbm ip-unicast-fib id 13990002
```

```
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   VRF   DEST
VRF  ISID  ISID  Destination  NH BEB  VLAN  OUTGOING SPBM  PREFIX IP ROUTE
      ISID  ISID  Destination  NH BEB  VLAN  INTERFACE COST  COST  PREFERENCE
-----
red   -    13990002 10.2.201.0/24 Switch-C 4000  1/2    20    1    7
-----
Total number of SPBM IP-UNICAST FIB entries 1
=====
```

```
Switch-D:1# show isis spbm ip-unicast-fib id all
```

```
=====
SPBM IP-UNICAST FIB ENTRY INFO
=====
VRF   VRF   DEST
VRF  ISID  ISID  Destination  NH BEB  VLAN  OUTGOING SPBM  PREFIX IP ROUTE
      ISID  ISID  Destination  NH BEB  VLAN  INTERFACE COST  COST  PREFERENCE
-----
GRT   -    -    10.0.0.1/32  Switch-C 4000  1/2    20    1    7
GRT   -    -    10.0.13.0/24 Switch-C 4000  1/2    20    1    7
green -    13990001 10.1.101.0/24 Switch-C 4000  1/2    20    1    7
red   -    13990002 10.2.201.0/24 Switch-C 4000  1/2    20    1    7
-----
Total number of SPBM IP-UNICAST FIB entries 4
=====
```

VRF green—Switch-C

```
Switch-C:1# show ip route vrf green
```

```
=====
IP Route - VRF green
=====
DST          MASK          NEXT          NH          INTER
VRF/ISID    COST  FACE  PROT  AGE  TYPE  PRF
-----
10.1.101.0   255.255.255.0 10.1.101.1   -           1    101  LOC  0   DB  0
10.1.102.0   255.255.255.0 Switch-D     vrf green   20   4000 ISIS 0   IBSV 7
=====
```

```
2 out of 2 Total Num of Route Entries, 0 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
```

VRF green—Switch-D

```
Switch-D:1# show ip route vrf green
```

```
=====
IP Route - VRF green
=====
DST                MASK                NEXT                NH                INTER
VRF/ISID          COST FACE PROT AGE TYPE PRF
-----
10.1.101.0         255.255.255.0      Switch-C           vrf green         20  4000 ISIS 0  IBSV 7
10.1.102.0         255.255.255.0      10.1.102.1        -                  1   102  LOC  0  DB   0
```

```
2 out of 2 Total Num of Route Entries, 0 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
```

VRF red—Switch-C

```
Switch-C:1# show ip route vrf red
```

```
=====
IP Route - VRF red
=====
DST                MASK                NEXT                NH                INTER
VRF/ISID          COST FACE PROT AGE TYPE PRF
-----
10.2.201.0         255.255.255.0      10.2.201.1        -                  1  201  LOC  0  DB   0
10.2.202.0         255.255.255.0      Switch-D           vrf red           20  4000 ISIS 0  IBSV 7
```

```
2 out of 2 Total Num of Route Entries, 0 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
```

VRF red—Switch-D

```
Switch-D:1# show ip route vrf red
```

```
=====
IP Route - VRF red
=====
DST                MASK                NEXT                NH                INTER
VRF/ISID          COST FACE PROT AGE TYPE PRF
-----
10.2.201.0         255.255.255.0      Switch-C           vrf red           20  4000 ISIS 0  IBSV 7
10.2.202.0         255.255.255.0      10.2.202.1        -                  1   202  LOC  0  DB   0
```

```
2 out of 2 Total Num of Route Entries, 0 Total Num of Dest Networks displayed.
```

```
-----
TYPE Legend:
```

```
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp Route,
```

Layer 3 VSN configuration

```
e,  
U=Unresolved Route, N=Not in HW, F=Replaced by FTN, V=IPVPN Route, S=SPBM Route  
PROTOCOL Legend:  
v=Inter-VRF route redistributed
```


Chapter 6: Layer 3 video surveillance basic configuration

The following sections describe the Layer 3 Video Surveillance install script, that you can use to perform a basic configuration to quickly deploy a video surveillance network.

Layer 3 Video Surveillance install script

*** Note:**

This feature is supported only on the Virtual Services Platform 4000 Series platform. For more information about feature support, see *Release Notes*.

*** Note:**

The Layer 3 Video Surveillance install script performs the same function as the `run vms endura` script. However, the switch continues to support the `run vms endura` script for backward compatibility.

The `run vms layer-3 switch` command runs the Layer 3 Video Surveillance install script that pre-configures basic and common configuration parameters to deploy a video surveillance network. You can use this script to quickly and easily deploy a video surveillance network in accordance with best practices, using networking equipment.

The install script is specifically targeted for use with the VSP 4450GSX-PWR+ and VSP 4850GTS/VSP 4850GTS-PWR+ platforms. This script enables customers to use a single command on a VSP 4450 Series switch to set up the core switch configuration where the VMS/video surveillance management and operation systems reside. Similarly, using the same command, customers can set up each edge VSP 4850 Series switch where the IP cameras are connected.

The install script requires the VSP 4000 Series switch to be in a factory default state, to ensure correct operation of the configuration. In a VMS solution, use a VSP 4450GSX-PWR+ switch as the core switch and a VSP 4850GTS or a VSP 4850GTS-PWR+ as the edge/access switch. Use an edge/access switch such as the VSP 4850GTS-PWR+ PoE switch to connect and power the IP video surveillance cameras.

Essentially, the Layer 3 Video Surveillance install script performs the following tasks:

- Creates an SPB network core solution with IP Shortcuts to connect IP subnet zones between the core and edge IP subnets.

- Configures all network edge IP subnet areas containing IP cameras with an IP gateway address, that is redistributed over the SPB fabric. This enables the fabric core to act as a single IP routing entity for the solution.
- Relays DHCP services between each IP subnet area and the central server, for IP camera address allocation.
- Enables IP multicast over Fabric Connect virtualization, to support and allow efficient IP multicast communication over the fabric core from IP cameras to central VMS servers, for viewing and recording video streams.

CLI command switch value

You must specify a value for the switch in the install script command, and the value must be between 5 and 99. Use the value 5 for a core switch where the VMS core systems are connected. It is recommended that you use a VSP 4450GSX-PWR+ as the core switch.

Use the range 6–99 for switch values when you run the script on edge/access layer switches. Ensure that the switch value is unique for each additional switch that is part of the solution.

For example, the first edge/access switch with the IP Cameras connected would use a value of switch 6. For additional edge/access switches, use switch 7, switch 8 and so on, for each IP subnet and IP camera zone. You can connect up to 48 IP cameras to a switch within an IP subnet zone.

Switch parameters configured by the script

The following major parameters are configured by the `run vms layer-3` command and utilizes the switch “#” value in the command to set up parameters:

- SNMP-Server switch hostname
- SPB parameters such as: System ID, Nickname, SPB Area ID, Backbone VLAN ID's (4051 and 4052), Multicast virtualization and CFM
- IP loopback interface addresses
- IP redistribution over IS-IS (IP Shortcuts)
- All SFP ports as SPB NNI ports
- All copper RJ-45 ports as end device ports with Spanning Tree enabled
- Spanning Tree mstprstp mode
- VLAN port memberships
- VLAN IP address (Gateway IP for VLAN)
- DHCP Relay

Note:

DHCP Relay parameters are configured only when the script is run on VSP 4850GTS and VSP 4850GTS-PWR+ switches.

Configuration file

Upon successful completion of the Layer 3 Video Surveillance install script, the switch saves the configuration with a filename based on the switch value used when the script is run. The switch primary boot config file flags are updated with the new filename.

For example, running the command `run vms layer-3 switch 5` results in a switch configuration filename of `spb-switch-5.cfg`.

Running the Layer 3 Video Surveillance install script

Use the following procedure to run the Layer 3 Video Surveillance install script.

This procedure applies only to the VSP 4000 Series platform. For more information about feature support, see *Release Notes*.

*** Note:**

The install script requires the switch to be in a factory default state and prompts you to confirm this.

*** Note:**

The `run vms layer-3 switch` command performs the same function as the `run vms endura switch` command. The switch supports the `run vms endura switch` command only for backward compatibility.

About this task

Use `switch 5` on a VSP 4450GSX-PWR+, which is suited for the network core where the video management systems exist.

Use `switch 6` on a VSP 4850GTS or VSP 4850GTS-PWR+, which is suited for the network edge where the IP cameras will connect. Up to 48 IP cameras can connect to a VSP 4850GTS-PWR+ switch within an IP subnet zone.

For each additional area and switch, increment the switch number by one. For example, use `switch 7` for the second edge switch. The configuration is customized based on that number for the IP subnet, loopback addresses, and SPB information.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Run the Layer 3 Video Surveillance install script:

```
run vms Layer-3 switch <switch #>
```

Where, `<switch #>` is a numerical value from 5 to 99 used to seed unique values in the configuration script.

*** Note:**

If the script causes a configuration conflict or cannot run a command, an error message displays and the script stops.

Example

The following example shows the configuration of a switch in the VMS core, and shows the configuration file created by the script.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch:1(config)#run vms layer-3 switch 5
Do you want to execute the run vms script? Device needs to be in factory default state.
(y/n) ? y
CP1 [05/05/17 07:48:33.760:IST] 0x000045e3 00000000 GlobalRouter SNMP INFO Save config
successful.
CP1 [05/05/17 07:48:37.951:IST] 0x000045e3 00000000 GlobalRouter SNMP INFO Save config
successful.

**Previous configurations stored in pre_vms_install.cfg**
**New VMS configurations stored in new primary config file spb-switch-5.cfg**

*** VMS script execution complete ***
Switch:1(config)#exit
Switch:1#
```

The following example shows the configuration of a switch at the edge, and shows the configuration file created by the script.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.


switch:1(config)#run vms layer-3 switch 6
Do you want to execute the run vms script? Device needs to be in factory default state.
(y/n) ? y
CP1 [05/05/17 07:54:04.046:IST] 0x000045e3 00000000 GlobalRouter SNMP INFO Save config
successful.
CP1 [05/05/17 07:54:05.760:IST] 0x000045e3 00000000 GlobalRouter SNMP INFO Save config
successful.

**Previous configurations stored in pre_vms_install.cfg**
**New VMS configurations stored in new primary config file spb-switch-6.cfg**

*** VMS script execution complete ***
Switch:1(config)#exit
Switch:1#
```

Variable definitions

Use the data in the following table to use the `run vms Layer-3 switch` command.

Variable	Value
<5-99>	<p>The numeric switch value used as a common element to configure switch parameters such as name, VLAN ID, SPB and IP parameters.</p> <p> Note: Use a value of 5 with a VSP 4450GSX-PWR+ for the network core where the VMS servers are connected. Use a</p>

Variable	Value
	value of 6 onwards (until and including 99) for all VSP 4850GTS/ VSP 4850GTS-PWR+ switches used for connecting IP Cameras at the network edge/access layer.

Appendix A: SPBM reference architectures

Reference architectures

SPBM has a straightforward architecture that simply forwards encapsulated C-MACs across the backbone. Because the B-MAC header stays the same across the network, there is no need to swap a label or perform a route lookup at each node. This architecture allows the frame to follow the most efficient forwarding path from end to end.

The following reference architectures illustrate SPBM with multiple switches in a network.

For information about solution-specific architectures like Video Surveillance or Data Center implementation using the VSP switch, see [Solution-specific reference architectures](#) on page 135.

The following figure shows the MAC-in-MAC SPBM domain with BEBs on the boundary and BCBs in the core.




Figure 16: SPBM basic architecture

Provisioning an SPBM core is as simple as enabling SPBM and IS-IS globally on all the nodes and on the core facing links. To migrate an existing edge configuration into an SPBM network is just as simple.

The boundary between the MAC-in-MAC SPBM domain and the 802.1Q domain is handled by the BEBs. At the BEBs, VLANs or VRFs are mapped into I-SIDs based on the local service provisioning.

Services (whether Layer 2 or Layer 3 VSNs) only need to be configured at the edge of the SPBM backbone (on the BEBs). There is no provisioning needed on the core SPBM nodes.

The following figure illustrates an existing edge that connects to an SPBM core.




Figure 17: Access to the SPBM Core

For Layer 2 virtualized bridging (Layer 2 VSN), identify all the VLANs that you want to migrate into SPBM and assign them to an I-SID on the BEB.

For Layer 3 virtualized routing (Layer 3 VSN), map IPv4-enabled VLANs to VRFs, create an IP VPN instance on the VRF, assign an I-SID to the VRF, and then configure the desired IP redistribution of IP routes into IS-IS.

All BEBs that have the same I-SID configured can participate in the same VSN. That completes the configuration part of the migration and all the traffic flows return to normal operation.

Campus architecture

For migration purposes, you can add SPBM to an existing network that has SMLT configured. In fact, if there are other protocols already running in the network, such as Open Shortest Path First (OSPF), you can leave them in place too. SPBM uses IS-IS, and operates independently from other protocols. However, it is recommended that you eventually eliminate SMLT in the core and eliminate other unnecessary protocols. This reduces the complexity of the network and makes it much simpler to maintain and troubleshoot.

Whether you configure SMLT in the core, the main point to remember is that SPBM separates services from the infrastructure. For example, in a large campus, a user may need access to other

sites or data centers. With SPBM you can grant that access by associating the user to a specific I-SID. With this mechanism, the user can work without getting access to confidential information of another department.

The following figure depicts a topology where the BEBs in the edge and data center distribution nodes are configured in SMLT clusters. Prior to implementing SPBM, the core nodes would also have been configured as SMLT clusters. When migrating SPBM onto this network design, it is important to note that you can deploy SPBM over the existing SMLT topology without network interruption. After the SPBM infrastructure is in place, you can create VSN services over SPBM or migrate them from the previous end-to-end SMLT-based design.

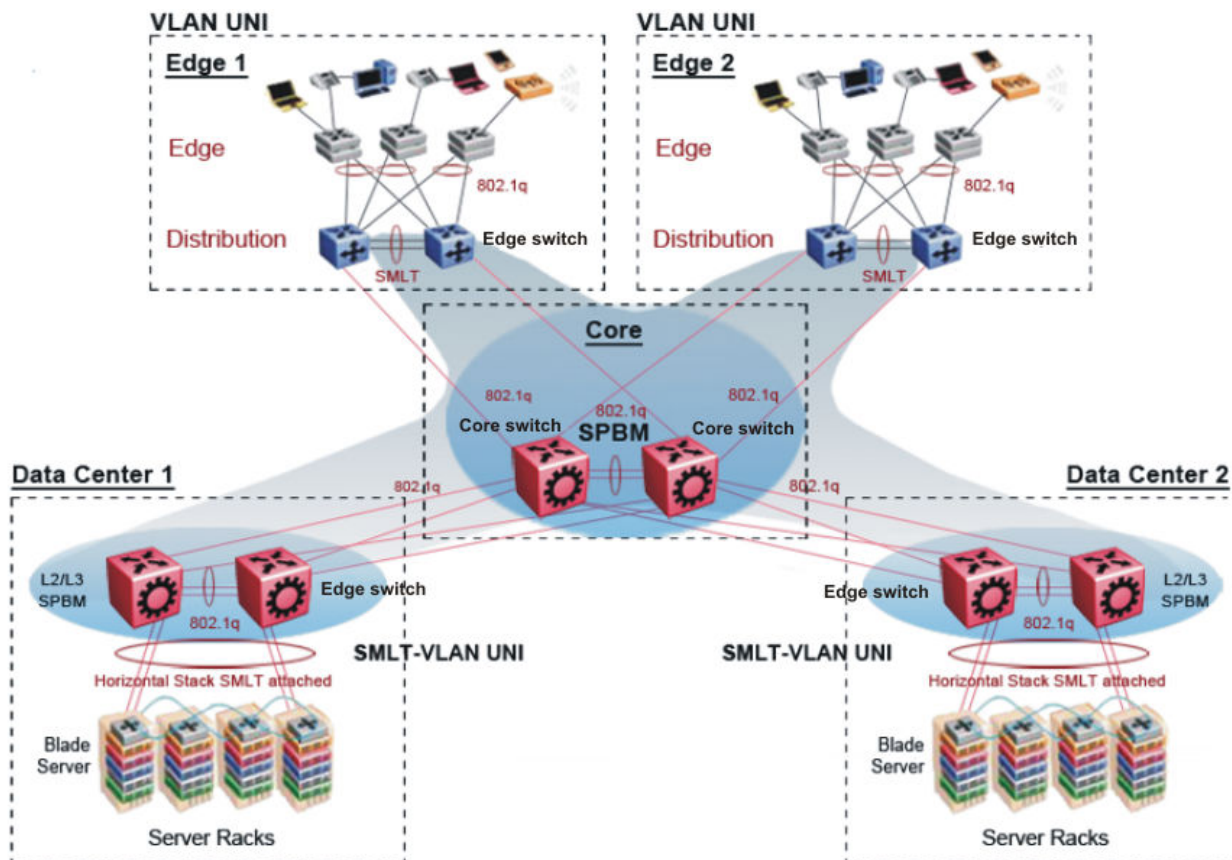


Figure 18: SPBM campus without SMLT

After you migrate all services to SPBM, the customer VLANs (C-VLANs) will exist only on the BEB SMLT clusters at the edge of the SPBM network. The C-VLANs will be assigned to an I-SID instance and then associated with either a VLAN in an Layer 2 VSN or terminated into a VRF in an Layer 3 VSN. You can also terminate the C-VLAN into the default router, which uses IP shortcuts to IP route over the SPBM core.

In an SPBM network design, the only nodes where it makes sense to have an SMLT cluster configuration is on the BEB nodes where VSN services terminate. These are the SPBM nodes where C-VLANs exist and these C-VLANs need to be redundantly extended to non-SPBM devices such as Layer 2 edge stackable switches. On the BCB core nodes where no VSNs are terminated and no Layer 2 edge stackables are connected, there is no longer any use for the SMLT clustering functionality. Therefore, in the depicted SPBM design, the SMLT/vIST configuration can be removed

from the core nodes because they now act as pure BCBs that simply transport VSN traffic and the only control plane protocol they need to run is IS-IS.

Because SMLT BEB nodes exist in this design (the edge BEBs) and it is desirable to use equal cost paths to load balance VSN traffic across the SPBM core, all SPBM nodes in the network are configured with the same two B-VIDs.

Where the above figure shows the physical topology, the following two figures illustrate a logical rendition of the same topology. In both of the following figures, you can see that the core is almost identical. Because the SPBM core just serves as a transport mechanism that transmits traffic to the destination BEB, all the provisioning is performed at the edge.

In the data center, VLANs are attached to Inter-VSNs that transmit the traffic across the SPBM core between the data center on the left and the data center on the right. A common application of this service is VMotion moving VMs from one data center to another.

The following figure uses IP shortcuts that route VLANs. There is no I-SID configuration and no Layer 3 virtualization between the edge distribution and the core. This is normal IP forwarding to the BEB.

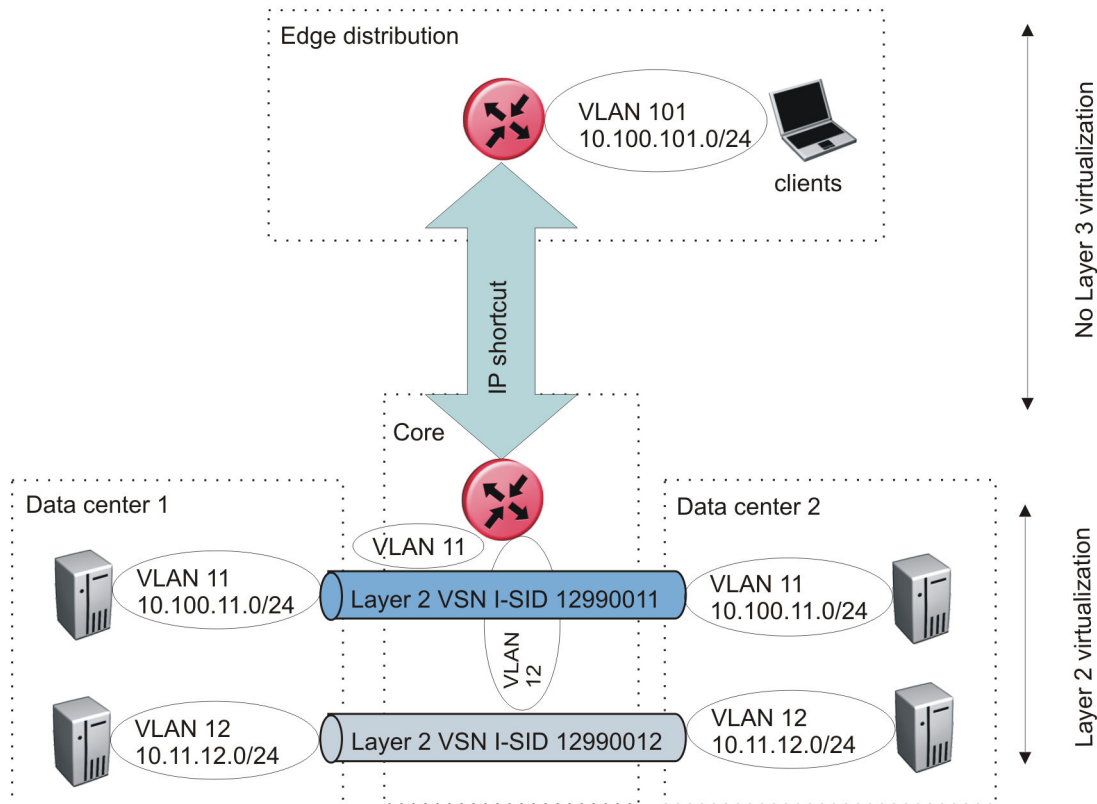


Figure 19: IP shortcut scenario to move traffic between data centers

The following figure uses Layer 3 VSNs to route VRFs between the edge distribution and the core. The VRFs are attached to I-SIDs and use Layer 3 virtualization.




Figure 20: VRF scenario to move traffic between data centers

Large data center architecture

SPBM supports data centers with IP shortcuts, Layer 2 VSNs, or Layer 3 VSNs. If you use vMotion, you must use Layer 2 between data centers (Layer 2 VSN). With Layer 2 VSNs, you can add IP addresses to the VLAN on both data centers and run Virtual Router Redundancy Protocol (VRRP) between them to allow the ESX server to route to the rest of the network.

The following figure shows an SPBM topology of a large data center. This figure represents a full-mesh data center fabric using SPBM for storage over Ethernet. This topology is optimized for storage transport because traffic never travels more than two hops.

*** Note:**

It is recommended that you use a two-tier, full-mesh topology for large data centers.




Figure 21: SPBM data center—full mesh

Traditional data center routing of VMs:

In a traditional data center configuration, the traffic flows into the network to a VM and out of the network in almost a direct path.

The following figure shows an example of a traditional data center with VRRP configured. Because end stations are often configured with a static default gateway IP address, a loss of the default gateway router causes a loss of connectivity to the remote networks. VRRP eliminates the single point of failure that can occur when the single static default gateway router for an end station is lost.




Figure 22: Traditional routing before moving VMs

A VM is a virtual server. When you move a VM, the virtual server is moved as is. This action means that the IP addresses of that server remain the same after the server is moved from one data center to the other. This in turn dictates that the same IP subnet (and hence VLAN) exist in both data centers.

In the following figure, the VM moved from the data center on the left to the data center on the right. To ensure a seamless transition that is transparent to the user, the VM retains its network connections through the default gateway. This method works, but it adds more hops to all traffic. As you can see in the figure, one VM move results in a complicated traffic path. Multiply this with many moves and soon the network look like a tangled mess that is very inefficient, difficult to maintain, and almost impossible to troubleshoot.




Figure 23: Traditional routing after moving VMs

Optimized data center routing of VMs:

Two features make a data center optimized:

- VLAN routers in the Layer 2 domain (green icons)
- VRRP BackupMaster

The VLAN routers use lookup tables to determine the best path to route incoming traffic (red dots) to the destination VM.

VRRP BackupMaster solves the problem of traffic congestion on the vIST. Because there can be only one VRRP Master, all other interfaces are in backup mode. In this case, all traffic is forwarded over the vIST link towards the primary VRRP switch. All traffic that arrives at the VRRP backup interface is forwarded, so there is not enough bandwidth on the vIST link to carry all the aggregated riser traffic. VRRP BackupMaster overcomes this issue by ensuring that the vIST trunk is not used in such a case for primary data forwarding. The VRRP BackupMaster acts as an IP router for packets destined for the logical VRRP IP address. All traffic is directly routed to the destined subnetwork and not through Layer 2 switches to the VRRP Master. This avoids potential limitation in the available vIST bandwidth.

The following figure shows a solution that optimizes your network for bidirectional traffic flows. However, this solution turns two SPBM BCB nodes into BEBs where MAC and ARP learning will be enabled on the Inter-VSN routing interfaces. If you do not care about top-down traffic flows, you can omit the Inter-VSN routing interfaces on the SPBM BCB nodes. This makes the IP routed paths top-down less optimal, but the BCBs remain pure BCBs, thus simplifying core switch configurations.




Figure 24: Optimized routing before moving VMs

In the traditional data center, chaos resulted after many VMs were moved. In an optimized data center as shown in the following figure, the incoming traffic enters the Layer 2 domain where an edge switch uses Inter-VSN routing to attach an I-SID to a VLAN. The I-SID bridges traffic directly to the destination. With VRRP BackupMaster, the traffic no longer goes through the default gateway; it takes the most direct route in and out of the network.




Figure 25: Optimized routing after moving VMs

Solution-specific reference architectures

The following sections describe solution-specific reference architectures, like for example for Video Surveillance or Data Center implementation, using the VSP 4000.

Multi-tenant — fabric connect

This fabric connect-based solution leverages the fabric capabilities of the VSP platforms: a VSP 7000 core and a VSP 4000 edge. This solution provides the ability to run, by default, up to 24 VRFs for each wiring closet and is well suited for multi-tenant applications. The zero-touch core is enabled by the fabric connect endpoint provisioning capabilities.

* Note:

You can increase VRF scaling to run more than 24 VRFs. The maximum number of supported VRFs and Layer 3 VSNs differs depending on the hardware platform. For more information about maximum scaling numbers, see *Release Notes*.

If this solution must support IPv6, then a central router-pair routes all IPv6 traffic. The IPv6 traffic is tunneled from each wiring closet to the IPv6 routers by extending Layer 2 VSNs to the q-tagged router interfaces.

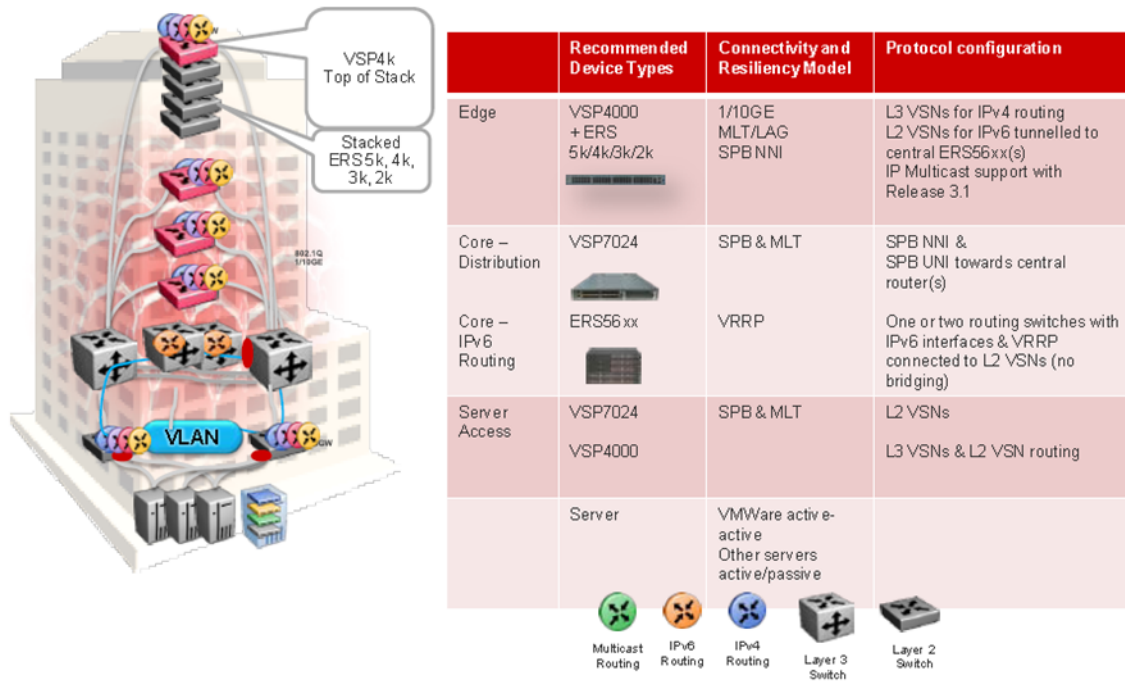


Figure 26: Small core — multi-tenant

The following list outlines the benefits of the fabric connect-based solution:

- Endpoint provisioning
- Fast failover
- Simple to configure
- Layer 2 and Layer 3 virtualized

Hosted data center management solution — E-Tree

In some hosted data center solutions, the hosting center operating company takes responsibility for managing customer servers. For this shared management, shown in the following figure, servers that control the operating system level of the production servers, such as the patch level, are deployed. Because customer production servers do not communicate with each other, a distributed private VLAN solution based on fabric connect is deployed to manage all production servers. This solution builds a distributed set of E-Trees for each management domain.

The VSP switches as access, provide an elegant network-wide E-Tree solution. Spokes, or managed servers, cannot communicate to each other over this network, but the shared management servers on the hub ports can access all spokes. Because of the Layer 2 – E-Tree nature of this setup, the managed servers do not require any route entries, and only require one IP interface in this management private VLAN. This solution supports tagged and untagged physical and virtual (VM) servers.

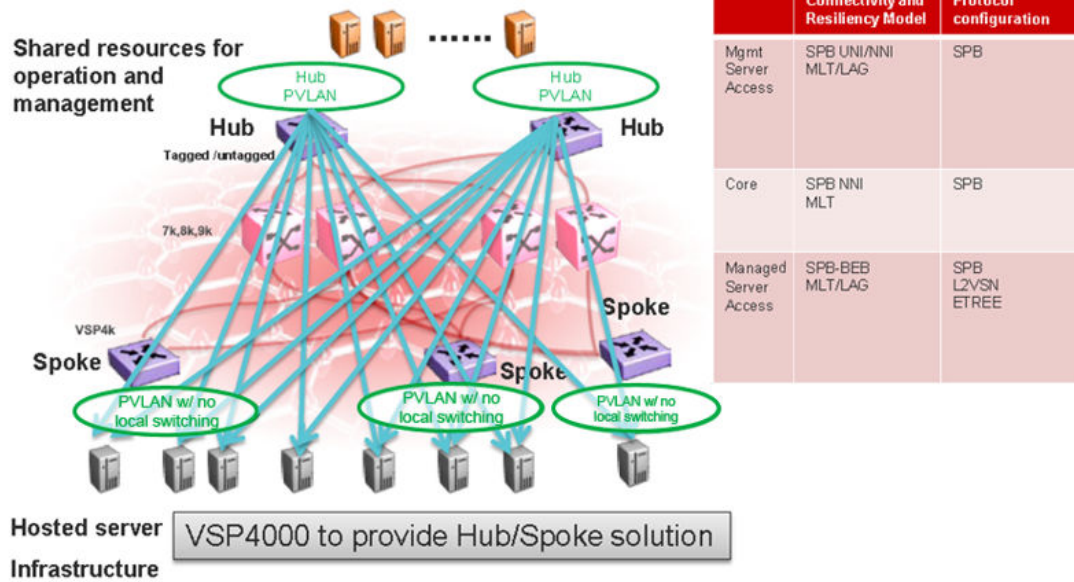


Figure 27: Data center hosting private VLAN

The following list outlines the benefits of the hosted data center management solution:

- Easy endpoint provisioning
- Optimal resiliency
- Secure tenant separation

Video surveillance — bridged

In a video surveillance solution, optimal traffic forwarding is a key requirement to ensure proper operation of the camera and recorder solutions. However, signaling is also important to ensure quick channel switching. This is achieved by deploying a fabric connect based IP multicast infrastructure that is optimized for multicast transport, so that the cameras can be selected quickly, and so that there is no unnecessary traffic sent across the backbone.

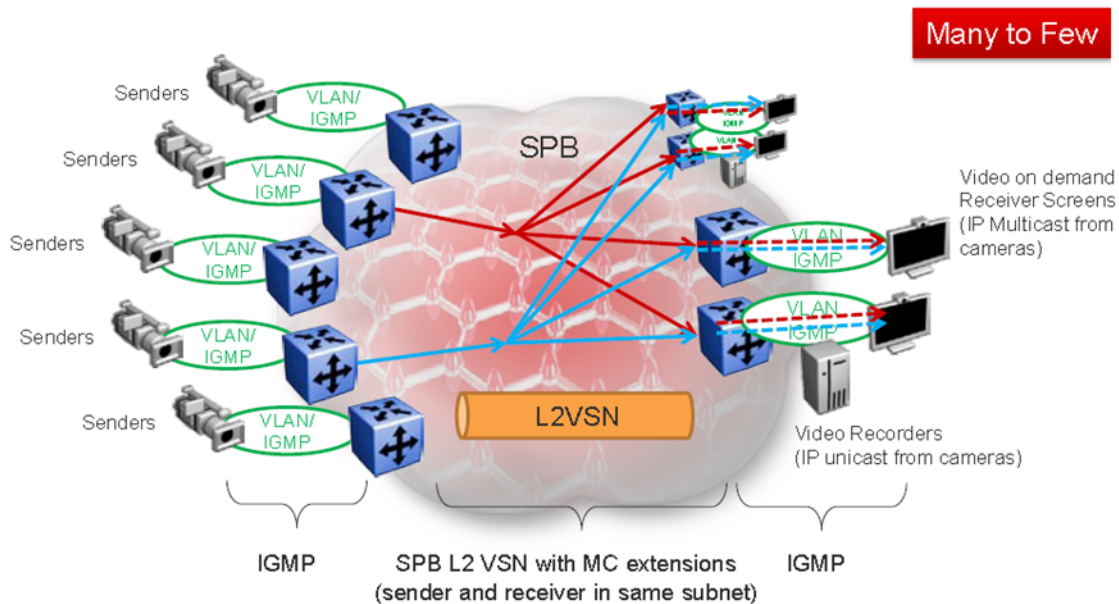


Figure 28: Deployment scenario — bridged video surveillance and IP camera deployment for transportation, airports, and government

The following list outlines the benefits of the bridged video surveillance solution:

- Easy end-point provisioning
- sub second resiliency and mc forwarding
- secure tenant separation
- quick camera switching

Video surveillance — routed

In a video surveillance solution, optimal traffic forwarding is a key requirement to ensure proper operation of the camera and recorder solutions. However, signaling is also important to ensure quick channel switching. This is achieved by deploying an IP multicast infrastructure that is optimized for multicast transport, so that the cameras can be selected quickly, and so that there is no unnecessary traffic sent across the backbone. In the topology shown in the following figure, each camera is attached to its own IP subnet. In a larger topology, this can reduce network overhead. To increase network scalability, you can attach a set of cameras to a Layer 2 switch that has IGMP, and then connect the cameras to the fabric edge (BEB) which has a routing instance.

In many customer scenarios, surveillance must be separated from the rest of the infrastructure. This can be achieved by deploying a Layer 3 VSN for the surveillance traffic to keep the surveillance traffic isolated from any other tenant. For more information, see *Configuring Fabric Layer 3 Services*.




Figure 29: Deployment scenario — Routed video surveillance and IP camera deployment for transportation, airports, and government

The following list outlines the benefits of the routed video surveillance solution:

- Easy endpoint provisioning
- Optimal resiliency and mc forwarding
- Secure tenant separation
- Rapid channel/camera switching

Metro-Ethernet Provider solution

VSP switches provide an end-to-end Metro-Ethernet Provider solution. Leveraging fabric connect throughout the infrastructure enables a scalable and flexible wholesale provider infrastructure.

This use case extends the Transparent Port UNI functionality to transparently forward any customer VLAN across the services.




Figure 30: Metro ring access solution

The following list outlines the benefits of the Metro-Ethernet Provider solution:

- Easy endpoint provisioning
- Optimal resiliency
- Secure tenant separation

Glossary

Backbone Core Bridge (BCB)	Backbone Core Bridges (BCBs) form the core of the SPBM network. The BCBs are SPBM nodes that do not terminate the VSN services. BCBs forward encapsulated VSN traffic based on the Backbone MAC Destination Address (B-MAC-DA). A BCB can access information to send that traffic to any Backbone Edge Bridges (BEBs) in the SPBM backbone.
Backbone Edge Bridge (BEB)	Backbone Edge Bridges (BEBs) are SPBM nodes where Virtual Services Networks (VSNs) terminate. BEBs handle the boundary between the core MAC-in-MAC Shortest Path Bridging MAC (SPBM) domain and the edge customer 802.1Q domain. A BEB node performs 802.1ah MAC-in-MAC encapsulation and decapsulation for the Virtual Services Network (VSN).
Backbone MAC (B-MAC)	Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation encapsulates customer MAC addresses in Backbone MAC (B-MAC) addresses. MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header that SPBM uses for delivery from end to end. As the MAC header stays the same across the network, no need exists to swap a label or perform a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end. In Shortest Path Bridging MAC (SPBM), each node has a System ID, which is used in the topology announcement. This same System ID also serves as the switch Backbone MAC address (B-MAC), which is used as the source and destination MAC address in the SPBM network.
Customer MAC (C-MAC)	For customer MAC (C-MAC) addresses, which is customer traffic, to forward across the service provider back, SPBM uses IEEE 802.1ah Provider Backbone Bridging MAC-in-MAC encapsulation. The system encapsulates C-MAC addresses within a backbone MAC (B-MAC) address pair made up of a BMAC destination address (BMAC-DA) and a BMAC source address (BMAC-SA).
Customer VLAN (C-VLAN)	A traditional VLAN with MAC learning and flooding, where user devices connect to the network. In SPBM, C-VLANs are mapped to a Service Instance Identifier (I-SID) at the Backbone Edge Bridges (BEBs).
Fabric Connect	Fabric Connect is a single network-wide protocol that enables virtualized network segmentation across the network infrastructure.

Global Routing Table (GRT)	The Global Routing Table (GRT) is a table that maintains the information needed to forward an IP packet along the best route.
Layer 2 Virtual Services Network	The Layer 2 Virtual Services Network (L2 VSN) feature provides IP connectivity over SPBM for VLANs. Backbone Edge Bridges (BEBs) handle Layer 2 virtualization. At the BEBs you map the end-user VLAN to a Service Instance Identifier (I-SID). BEBs that have the same I-SID configured can participate in the same Layer 2 Virtual Services Network (VSN).
Layer 3 Virtual Services Network	The Layer 3 Virtual Services Network (L3 VSN) feature provides IP connectivity over SPBM for VRFs. Backbone Edge Bridges (BEBs) handle Layer 3 virtualized. At the BEBs through local provisioning, you map the end-user IP enabled VLAN or VLANs to a Virtualized Routing and Forwarding (VRF) instance. Then you map the VRF to a Service Instance Identifier (I-SID). VRFs that have the same I-SID configured can participate in the same Layer 3 Virtual Service Network (VSN).
Protocol Independent Multicast, Source Specific (PIM-SSM)	PIM-SSM is a multicast routing protocol for IP networks. PIM-SSM uses only shortest-path trees to provide multicast services based on subscription to a particular (source, group) channel. PIM-SSM eliminates the need for starting with a shared tree by immediately joining a source through the shortest path tree. This method enables PIM-SSM to avoid using a rendezvous point (RP) and RP-based shared tree, which can be a potential bottleneck.
Protocol Independent Multicast, Sparse Mode (PIM-SM)	PIM-SM is a multicast routing protocol for IP networks. PIM-SM provides multicast routing for multicast groups that can span wide-area and inter-domain networks, where receivers are not densely populated. PIM-SM sends multicast traffic only to those routers that belong to a specific multicast group and that choose to receive the traffic. PIM-SM adds a Rendezvous Point router to avoid multicast-data flooding. Use PIM-SM when receivers for multicast data are sparsely distributed throughout the network.
rendezvous point (RP)	The root of the shared tree. One RP exists for each multicast group. The RP gathers information about available multicast services through the reception of control messages and the distribution of multicast group information. Protocol Independent Multicast (PIM) uses RPs.
Shortest Path Bridging (SPB)	Shortest Path Bridging is a control Link State Protocol that provides a loop-free Ethernet topology. There are two versions of Shortest Path Bridge: Shortest Path Bridging VLAN and Shortest Path Bridging MAC. Shortest Path Bridging VLAN uses the Q-in-Q frame format and encapsulates the source bridge ID into the VLAN header. Shortest Path Bridging MAC uses the 802.1 ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header.

**Shortest Path
Bridging MAC
(SPBM)**

Shortest Path Bridging MAC (SPBM) uses the Intermediate-System-to-Intermediate-System (IS-IS) link-state routing protocol to provide a loop-free Ethernet topology that creates a shortest-path topology from every node to every other node in the network based on node MAC addresses. SPBM uses the 802.1ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header. SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based link-state protocol, which can provide virtualization services, both layer 2 and layer 3, using a pure Ethernet technology base.