



Shortest Path Bridging Multi-Area Design Guidelines

9037906-00 Rev. AA
October 2023



Copyright © 2023 Extreme Networks, Inc. All rights reserved.

Legal Notice

, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

and the logo are trademarks or registered trademarks of , Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on trademarks, see:

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses.

End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>




Table of Contents

Preface.....	4
Text Conventions.....	4
Documentation and Training.....	5
Help and Support.....	6
Subscribe to Product Announcements.....	6
Send Feedback.....	7
Overview.....	8
Fabric Connect (SPB) Multi-Area Overview.....	8
Multi-Area Architectural Concepts.....	10
Multi-Area Operational Considerations.....	12
Multi-Area Network Design Concepts.....	13
Introducing a Backbone Transit Area.....	14
Migrating to a Fabric Connect Edge while introducing Multi-Area.....	17
VSP Edge & Multi-Area.....	20
Fabric to the Edge with Multi-Area Blueprint.....	21
Appendix A: Acronyms.....	22
Index.....	25
Glossary.....	24



Preface

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.

Text Conventions

Unless otherwise noted, information in this document applies to all supported environments for the products in question. Exceptions, like command keywords associated with a specific software version, are identified in the text.

When a feature, function, or operation pertains to a specific hardware product, the product name is used. When features, functions, and operations are the same across an entire product family, such as ExtremeSwitching switches or SLX routers, the product is referred to as *the switch* or *the router*.

Table 1: Notes and warnings






Icon	Notice type	Alerts you to...
	Tip	Helpful tips and notices for using the product
	Note	Useful information or instructions
	Important	Important features or instructions
	Caution	Risk of personal injury, system damage, or loss of data
	Warning	Risk of severe personal injury

Table 2: Text

Convention	Description
screen displays	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
Key names	Key names are written in boldface, for example Ctrl or Esc . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press Ctrl+Alt+Del
<i>Words in italicized type</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
<i>italic text</i>	Italic text indicates variable content.
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member[member...]</i> .
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and software compatibility](#) for Extreme Networks products
[Extreme Optics Compatibility](#)
[Other resources](#) such as white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

Help and Support

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

Extreme Portal

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

The Hub

A forum for customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by employees, but is not intended to replace specific guidance from GTAC.

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2800. For the support phone number in your country, visit www.extremenetworks.com/support/contact.

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

- Your service contract number, or serial numbers for all involved products
- A description of the failure
- A description of any actions 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

Subscribe to Product Announcements

You can subscribe to email notifications for product and software release announcements, Field Notices, and Vulnerability Notices.

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

You can modify your product selections or unsubscribe at any time.

Send Feedback

The Information Development team at has made every effort to ensure that this document is accurate, complete, and easy to use. We strive to improve our documentation to help you in your work, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information.
- Broken links or usability issues.

To send feedback, do either of the following:

- Access the feedback form at .
- Email us at .

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.



Overview

This document provides design guidelines for SPB Multi-Area deployments. It outlines how to design greenfield networks using this functionality, and provides guidance for migrations to a Multi-Area solution from a brown field fabric deployment. It describes multi-area design elements, multi-area usage when migrating an existing (ERS) network to a VOSS Edge deployment, and overcoming nodal scale limitations in very large deployments using multi-area. Migration scenarios are discussed to show a smooth migration approach.

Fabric Connect (SPB) Multi-Area Overview

The Fabric Connect multi-area solution provides a number of benefits, including:

Fabric Node Scale

New Fabric-Connect capable edge switches at both the campus and the WAN edge can significantly increase the number of Fabric-enabled nodes. As a result, the number of IS-IS enabled nodes in a SPB area can get very large. The network can be segmented into areas of up to 500 fabric nodes each using multi-area to accommodate the large number of IS-IS enabled nodes on the SPB area.

Interconnecting those areas with multi-area SPB enables network deployments of tens of thousands of fabric-enabled switches with seamless end-to-end service automation.

Route Scaling

In the case of larger networks, in some instances, a best practice is to summarize IP routes. Prior to multi-area SPB, route summarization was not available and could only be applied at the IS-IS boundary to other protocols such as RIP, OSPF or BGP.

With Multi-Area SPB, route summarization is now supported between IS-IS area boundaries for the global routing table (VRFO) as well as all other VRFs/L3 VSNs.

Service Scaling

Within an SPB IS-IS area the L2, L3 Unicast, and Multicast services (ISIDs) are exchanged through the IS-IS IGP protocol. The addressing space within an area is limited by the protocol itself. While it can exchange many thousands of services, IP routes, and IP Multicast streams, it has an upper limit. Additionally, smaller switches have smaller MAC, ARP and routing tables as well as smaller Control Plane CPUs. The overall network size can be restricted if the network consists only of one IS-IS area.

With multi-Area, SPB networks can be segmented into communication zones, and services between those zones can be selectively restricted and filtered, enabling

maximum overall network scalability by using the Area Boundary Nodes. Control plane as well as ASIC table usage can be reduced if needed.

Secure Service Zoning

In an SPB enabled network, services (ISIDs) can span anywhere within an area and across area boundaries, or both. The spanned services enable seamless end-to-end service automation with end-point-only provisioning, but such dynamic end-point provisioning is not always necessary, especially when crossing jurisdictions. A good example is that some of the Data Center subnets should not span into a Campus-edge or WAN area.

With Multi-Area SPB, a network operator can segment a network into secure service zones by using the fine-grain service redistribution capabilities for L2 and L3 IP Unicast and Multicast.

Faster Convergence

While SPB typically converges very fast when network failures occur, network convergence times can depend on area size, as topology change information must be propagated hop-by-hop throughout the network before network convergence can take place.

With Multi-Area SPB, smaller IGP (Interior Gateway Protocol) topologies sharing a Link-State-Database can be created. Convergence in smaller areas is faster compared to larger areas with many links and nodes where the topology change information has to be relayed to more nodes. Also, the shortest path calculations have to be done for a larger topology mesh.



Multi-Area Architectural Concepts

This topic discusses multi-area architectural concepts.

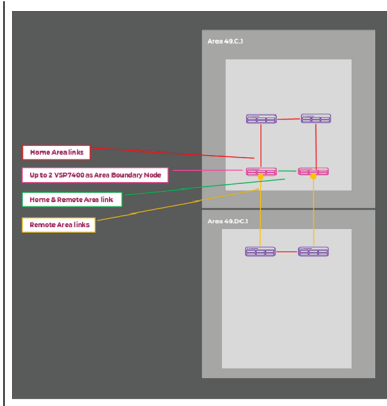
Multi-Area **boundary node** functionality for interconnecting areas is supported on the VSP7400, Universal switches 5520, 5720, 7520 and 7720, starting with the VOSS 8.10 release. The non-boundary nodes are supported with any existing VOSS release, so there is no need to upgrade the rest of the network to the release supporting the boundary node functionality.

A multi-area IS-IS boundary node supports two areas, which can be referred to as the *home* area and the *remote* area.

Both areas run their own independent IS-IS link-state-database and shortest path algorithms and both areas are identified by their individual IS-IS area ID:

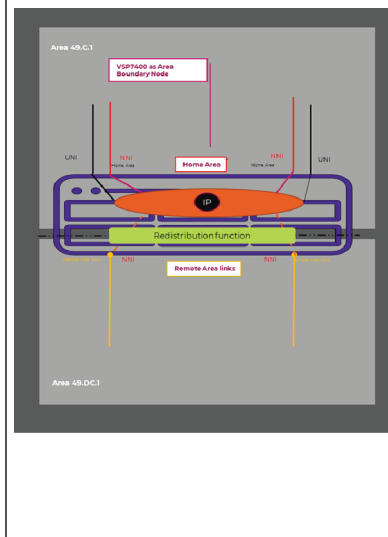
For example:

- Home area (49.C1)
- Remote area (49.DC1)



Areas are assigned to Network-to-Network (NNI) links, which can be physical or logical (Fabric-extended) links.

A link can either establish adjacencies in the home-area, the remote-area, or in home- and remote-areas. A link that establishes an adjacencies in home- and remote-areas creates boundary node interconnections that provide an elegant way to close the fabric topology for both areas using an optional single physical link.



The home-area is a concept on a boundary node which defines the IS-IS area where its local User interfaces (UNI) reside. This also includes local IP interfaces and local services such as nick-name server, DVR controller or PIM-Gateway controller functionality.


Boundary nodes are typically depicted in their home area with links out to their remote area.

The remote-area of a boundary node is the area where its services to and from the home-area are redistributed. A switch can have any number of physical or logical (Fabric-extended) remote area links up to the total amount of supported NNI interfaces.

An SPB fabric supports the following network services:

- L2 VSNs & ETree: Any-to-any L2 stretched VLAN and point to multi-point VLAN
- GRT & L3 VSNs: Network wide IP VPNs
- Bridged & Routed Multicast: Bridged or routed IP Multicast
- DVR Backbone: Anycast Gateway solution
- ISID Mirroring: Ability to span mirrored traffic across fabric

Services are not redistributed between the home- and remote-area of a multi-area boundary node by default. Redistribution policies can be configured to define which services are permitted to be stretched across multiple areas.



```
L2:
multi-area l2 isid-list "l2-isid-redirect" 1-80
multi-area l2 redistribute i-sid deny-all except-isid-list "l2-isid-redirect"
```

```
L3 Unicast:
route-map "area-20-to-30" 10
permit
enable
match network "area-20"
route-map "area-20-to-30" 65535
no permit
enable

router isis
multi-area ip redistribute unicast home-to-remote
multi-area ip redistribute unicast home-to-remote route-map "area-20-to-30"
```


```
L3 Multicast:
ip prefix-list "area-20-mc-group-all" 224.0.1.0/24 id 4 ge 24 le 32

route-map "inter-area-mc" 30
permit
enable
match multicast-group "area-20-mc-group-all"
route-map "inter-area-mc" 65535
no permit
enable

router isis
multi-area ip redistribute routed-multicast home-to-remote
multi-area ip redistribute routed-multicast home-to-remote route-map "inter-area-mc"
multi-area ip redistribute routed-multicast home-to-remote enable
exit
```

```
DVR:
Router isis
multi-area dvr redistribute
```

Multi-Area Operational Considerations:



Multi-area uses two equal IS-IS L1 areas. ISIDs are redistributed between areas using redistribution policies. It is important to ensure that redistribution policies on both boundary nodes are the same.

Connected areas appear as virtual nodes in an area, such as VG and VB. Virtual nodes represent an area and all redistributed services from the connected area. For example, virtual node VG represents area 49.0030 and its redistributed services in area 49.0050. Virtual node VB represents area 49.0050 and all its redistributed services in area 49.0030.

Up to two boundary nodes can share a virtual node with the initial release.

The node count in an area includes the physical boundary nodes and the virtual nodes when deploying multi-area. For example, the node count in area 49.0030 is four (G1, BG1, BG2, VB).


Multi-area boundary node functionality increases nodal count by the boundary nodes as well as the virtual nodes per area. Connected areas appear as a virtual-node and count as 1 additional node each (+1 to 2 boundary nodes)

When designing multi-area networks, a best practice is to limit the number of IS-IS nodes per area to 500.

All traffic to a connected area is destined to the virtual node representing the connected area.

Area topologies need to be loop free. A built-in area loop detection mechanism among all area boundary nodes protects against area loops and disables the area boundary node, closing an area-loop.

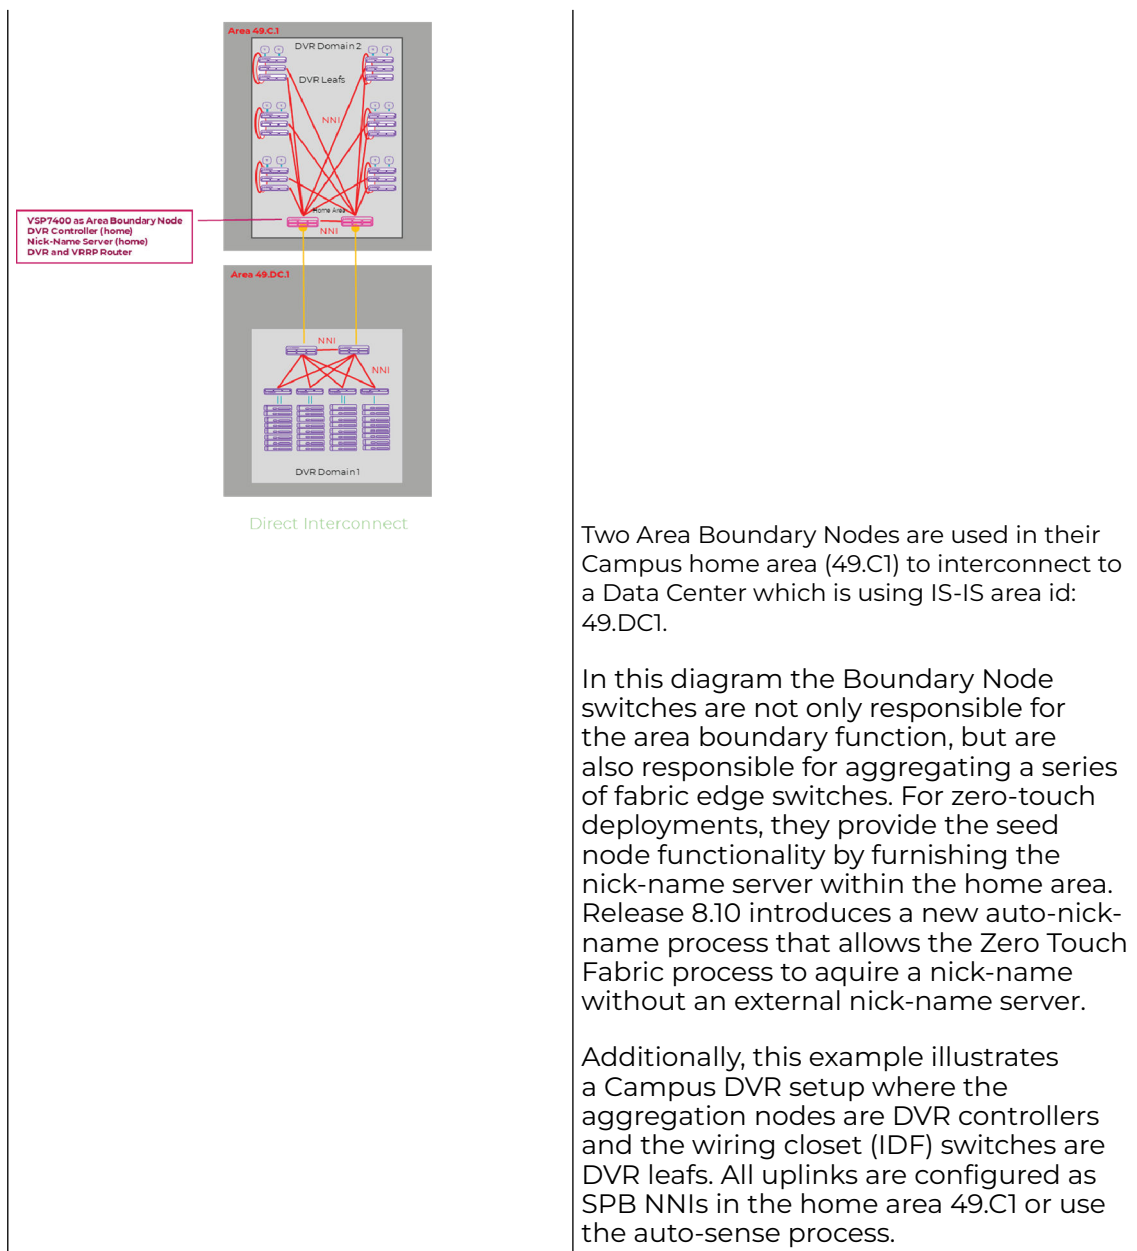
You can chain up to 5 areas, and still be able to scale to very large multi-area networks in a hierarchical design.



Multi-Area Network Design Concepts

This topic discusses multi-area design concepts. There are multiple ways to build scalable multi-area networks.

1. Interconnecting two areas with a pair of SPB Multi-area boundary nodes:



This direct interconnect of two areas is typically used for smaller deployments. In larger multi-area deployments, the introduction of a “backbone transit” area might be necessary.



Note

For Fabric Edge, DVR is a deployment option, but is not required for a successful and feature rich fabric edge solution. With Network Access Control, a non DVR deployment can be fully automated through Radius-based authentication and service attribute provisioning.

2. Backbone Transit Area design for larger multi-area deployments:

By electing two pairs of boundary nodes, one pair for the Campus IS-IS area 49.C1 and another pair for the Data Center IS-IS area 49.DC1, a Backbone Transit Area can be created. This can be implemented by interconnecting the boundary nodes directly using the “yellow” remote area links.

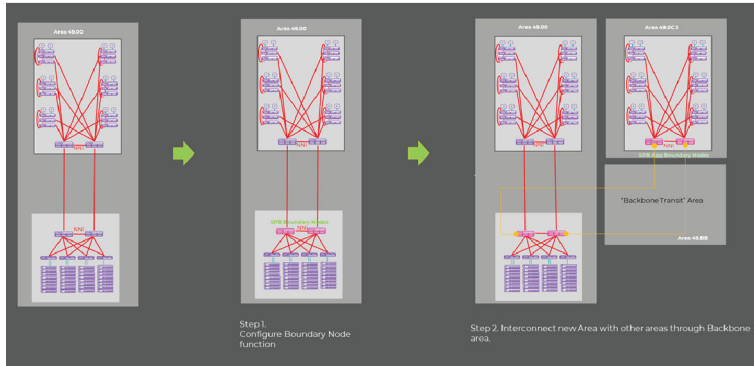
The NNI links residing in the backbone transit area “49.BB” can be direct physical links, or logical Fabric Extend links:



Introducing a Backbone Transit Area

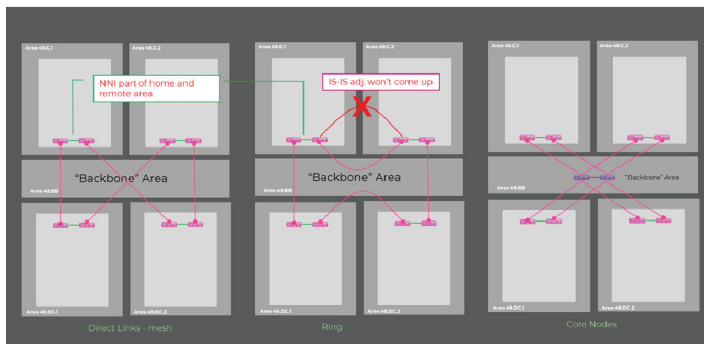
A Backbone Transit area (BBT) can be used to expand an existing single area SPB network to a multi-area network. This BBT area does not need to consist of any nodes. It can be created by configuring the links between the area boundary nodes with the BBT area ID.

The BBT is an arbitrary area that is “elected” to be the backbone transit area where other areas are connected to. There is no special configuration required.



There are several BBT deployment options:

- BBT with direct-link mesh
- BBT Ring
- BBT with Core Backbone Nodes



These different options to design a Backbone Transit Area illustrate the flexibility of a scale-out multi-area design. The main benefit is that additional areas can be added to the network without redesigning the IS-IS area constellation. Choosing a design with dedicated backbone core nodes acting as area aggregator nodes within the Backbone Transit Area can result in minimal disruption when expanding the network with additional areas since adding new areas does not impact existing area topologies. This type of design works nicely with large campuses that have available fiber, especially where larger buildings running Fabric Connect right to the edge wiring closet are part of the solution.

Ring designs are typically useful where fiber is scarce.

A couple of key design considerations are:

- To provide redundancy in a multi-area fabric design, a best practice is to ensure that both areas of a boundary node (home and remote) are redundantly connected. A direct NNI link between the boundary nodes is typically used for this. This direct NNI

link (or emulated link) should have both home- and remote-areas configured on it, ensuring that both areas are multi-homed to the boundary nodes for redundancy.

**Note**

If it is not possible to connect two boundary nodes directly with each other, then the design must ensure that both areas (home and remote) have a redundant connection between the boundary nodes. Redundant connections can be established by connecting with multiple links to either area.


- The backbone transit area is established by connecting the backbone area interfaces with each other over an NNI connection. If two areas are mistakenly interconnected with NNIs that don't have the same area, the IS-IS adjacency won't establish.



Migrating to a Fabric Connect Edge while introducing Multi-Area

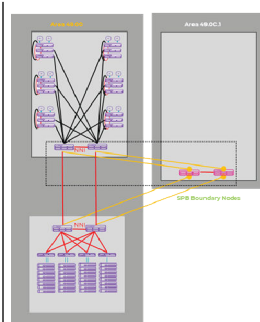
This topic describes the migration from a traditional Fabric Attach /MLAG /SMLT edge to a Fabric edge network with the introduction of a new area at the same time.

Any existing SPB fabric can be expanded using this step by step approach, minimizing the risk of disrupting network operation during the migration.



This example contains a network with Campus and Data Center components. The Campus wiring closets (IDF) are stacked switches using a link-trunking method and are dual-homed to two aggregation layer switches.

The aggregation layer (MDF) nodes can be fabric attach server switches and have fabric connect (SPB) connections to the data center and possibly other buildings that are fabric enabled. These aggregation layer nodes also typically provide the default gateway function for the campus building.



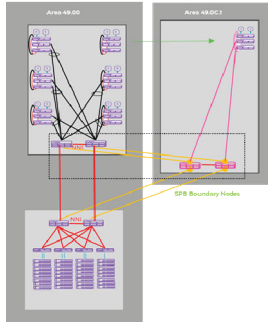
To provide a fabric to edge solution while introducing a new IS-IS area, a best practice is to introduce a new set of aggregation layer nodes that are able to provide the multi-area capability.

In this example the existing MDF nodes and the new area boundary nodes are co-located illustrated with the dotted square. The interconnections to the existing IS-IS area are remote area links.

These new switches can be introduced without any disruption to the existing infrastructure.

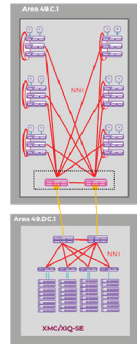
Typically, L3 default gateway routing would remain on the existing MDF switches until the IDF migration has been completed.

To ensure all services are available in the new area as well as the existing area,

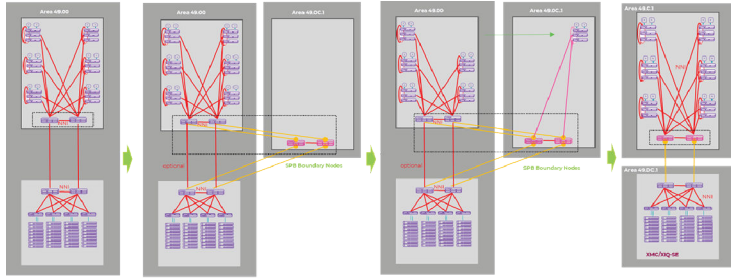
	<p>inter-area redistribution policies must be configured.</p> <p>After the MDF switches are in place, the new fabric to the edge nodes can be introduced to the new area and users can be migrated over to the new IDF switches when ready.</p> <p>The IDF switches can either be deployed using the automated XMC/XIQ-SE onboarding (ZTP+) and using Zero Touch Fabric (ZTF) deployment option where the new edge fabric nodes are deployed without any configuration, or they can be pre-staged.</p> <p>In the ZTP+/ZTF deployment option, the switches join the fabric automatically, get their management IP address through DHCP, and then onboard to XMC/XIQ-SE automatically, where they then get their initial configuration.</p> <p>To facilitate this deployment option, the onboarding ISID 159999999 needs to provide reachability to the management segment, including the DHCP server.</p>
---	---

For ZTF, a nick-name server needs to be enabled in the new Campus area (49.C1) on the MDF multi-area boundary nodes.

Finally, the L3 Default Gateway routing function can be migrated to the new aggregation (MDF) nodes.

	<p>When all users have been migrated off the old IDF switches, all the old switches can be decommissioned or redeployed for other uses.</p> <p>The result of this migration is that now the network consists of two SPB IS-IS areas with a Campus area that is operating as a fabric to the edge network.</p>
---	---

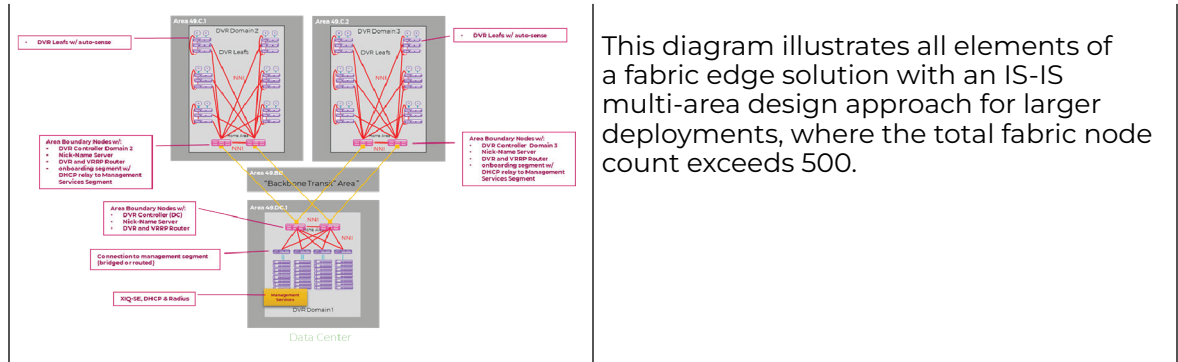
The same approach can be used for fabric to the edge solutions with ERS switches:





VSP Edge & Multi-Area

This topic describes a fabric edge solution with a multi-area design.



This diagram illustrates all elements of a fabric edge solution with an IS-IS multi-area design approach for larger deployments, where the total fabric node count exceeds 500.

The Campus network consists of MDF nodes using the following features:


- Multi-Area boundary nodes, for large buildings where a lot of Fabric edge switches are deployed. The home-area is the area internal to the building.
- Nick-Name server to enable Zero Touch Fabric (ZTF) deployment for the IDF switches.



Note

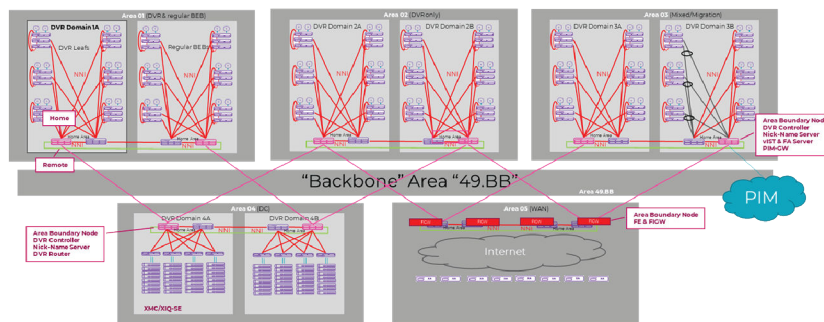
Boundary nodes must be configured manually and cannot operate in ZTF deployment mode. Also, auto-nickname can only be used in one area (typically the edge area) separated by area boundary nodes. The other area must use either static nick-name assignments or the nick-name server function.

- Preferred: L3 routing and redundancy using VRRP in home-area on aggregation nodes, which can also be the area boundary nodes.
- Optionally: DVR controller enabling distributed L3 routing capability for stretched subnets, as well as enabling per ISID multicast, removing the need to configure anything multicast related on the IDF nodes.
- Reachability to the management segment including DHCP server for IP address acquisition using the onboarding ISID 15999999.
- Data Center networks are typically deployed as DVR domains in dedicated Data Center areas.




Fabric to the Edge with Multi-Area Blueprint

This topic describes a scale-out Multi-Area blueprint with Backbone Transit Area and several fabric and non-fabric enabled IDF closets. The solution consists of three campus areas, one data center area, and a WAN area.



Each area is connected with two multi-area boundary nodes to the backbone transit area.

By adding another layer such as a WAN backbone area, individual metros can be interconnected, building an even further scaled solution and enabling networks with tens of thousands of fabric nodes.





Appendix A: Acronyms

DVR

Distributed Virtual Routing

GRT

Global Routing Table

IDF

Intermediate Distribution Frame

IGP

Interior Gateway Protocol

ISID

I-component Service Identifier per IEEE 802.1ah.

IS-IS

Intermediate Systems to Intermediate Systems Protocol established at IETF in 1990 as rfc1142.

MDF

Main Distribution Frame

Multi-Area

Interconnection of multiple IS-IS areas through multi-area boundary nodes.

Nick-name Server

Function to automatically distribute SPB Nick-names to other devices

NNI

Network to Network Interface.

SPB

Shortest Path Bridging: IEEE 802.1aq/ IEEE 802.1Q-2018.

UNI

User to Network Interface.

VRF

Virtual Route Forwarder

VSN


Virtual Services Network

ZTF

Zero Touch Fabric for SPB deployments

ZTP+

Zero Touch Provisioning with automated XMC/XIQ-SE onboarding



Glossary



Index

A

announcements 6

C

conventions

notice icons 4

text 4

D

documentation

feedback 7

location 5

F

feedback 7

N

notices 4

P

product announcements 6

S

support, *see* technical support

T

technical support

contacting 6

W

warnings 4