



Contents

Summary.....	1
Fabric Connect (Shortest Path Bridging) Multi-Area Overview.....	1
Multi-Area Architectural Concepts.....	3
Multi-Area Operational Considerations.....	5
Multi-Area Network Design Concepts.....	7
Migrating to a Fabric Connect Edge While Introducing Multi-Area.....	11
Fabric Edge and Multi-Area.....	15
Fabric to the Edge with Multi-Area Blueprint.....	16

Designing Multi-Area Fabric Connect/ Shortest Path Bridging Networks

This document provides design guidance for deploying multi-area Fabric Connect/SPB networks. It also discusses how multi-area can be used to scale the network when migrating an existing (ERS) network to a Fabric Edge deployment (using a VOSS-enabled switch). A couple migration scenarios are discussed to show a seamless approach.

Fabric Connect (Shortest Path Bridging) Multi-Area Overview

Extreme Networks' Fabric Connect (Shortest Path Bridging) technology simplifies the design, deployment and on-going management of enterprise networks. Long wait times required for network additions or changes and the rigid design constraints of legacy networks have made it challenging for network IT teams to move at the speed of the business. Fabric Connect eliminates these delays and enables our customers to deploy networks that are far more agile, enabling rapid application deployment and seamless services provisioning - while also improving the reliability, stability, and security of the network.

In the VOSS 8.4 release, Extreme Networks introduced multi-area capabilities to enhance Fabric Connect networks. Allowing them to scale to 10's of thousands of nodes, while also introducing capabilities to enhance the security of the network.

The specific benefits multi-area technology provides include:

Increased Fabric Node Scale

With the introduction of new Fabric Connect capable edge switches at both the campus and the WAN edge, the number of Fabric-enabled nodes can be significantly increased. As a result, the number of IS-IS enabled nodes in a SPB area can get very high. Multi-Area provides a solution to segment the network into areas of up to 500 fabric enabled nodes each.

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

Increased Route Scaling

In large networks, for certain use cases, it makes sense to summarize IP routes. Prior to multi-area Fabric Connect/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 (VRF0) as well as all other VRFs/L3 VSNs.

Increased Service Scaling

Within a Fabric Connect/SPB IS-IS area, 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, thus 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 Area Boundary Nodes to abstract or hide the services within the area behind them. Control plane as well as ASIC table usage can thus be reduced if required.

Enhanced Security with Service Zoning

In an SPB enabled network, services (ISIDs) can span anywhere. This enables seamless end-to-end service automation with end-point-only provisioning. However, this dynamic end-point provisioning is not always desired, especially when crossing network boundaries. 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 has to 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 and the shortest path calculations have to be done for a larger topology mesh.

Multi-Area Architectural Concepts

Multi-Area boundary node functionality for interconnecting areas is supported on the VSP7400 (VSP7400-32C and VSP7400-48Y) switches starting with the VOSS 8.4 release. The non-boundary nodes can operate with any existing VOSS release. There is no requirement 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. Extreme calls these 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.C.1)
- Remote area (49.DC.1)

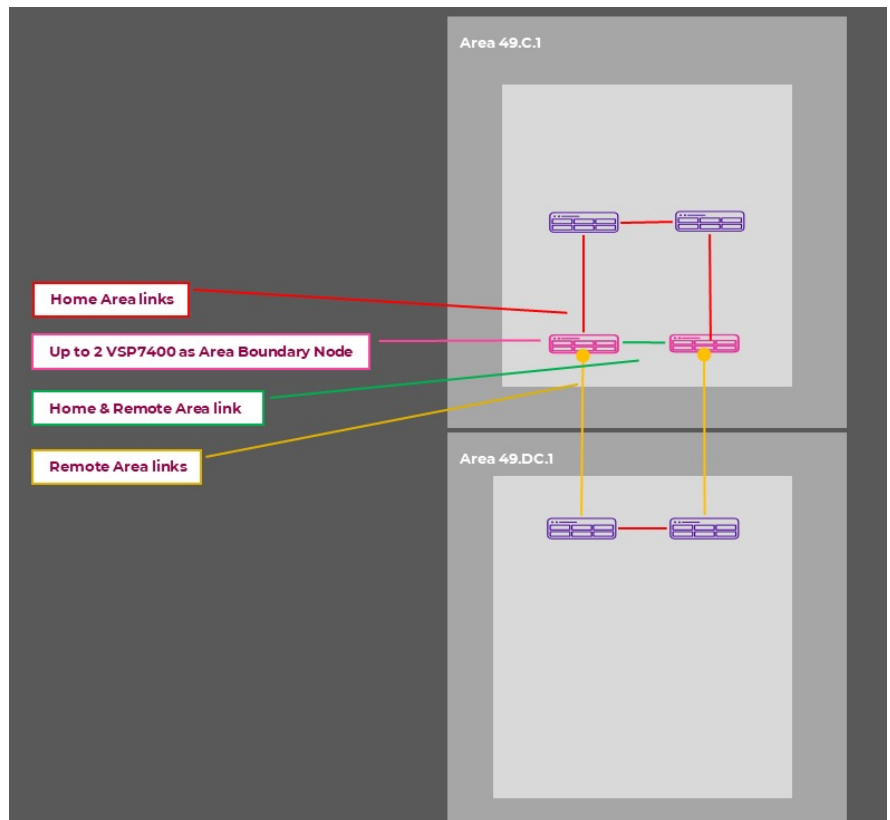


Figure 1: Multi-area architecture concepts (network view)

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

A link can either establish adjacencies in the home-area, the remote area or in home- and remote-area. The latter case is for boundary node interconnections to provide an elegant way to close the fabric topology for both areas using a (optional) single physical link.

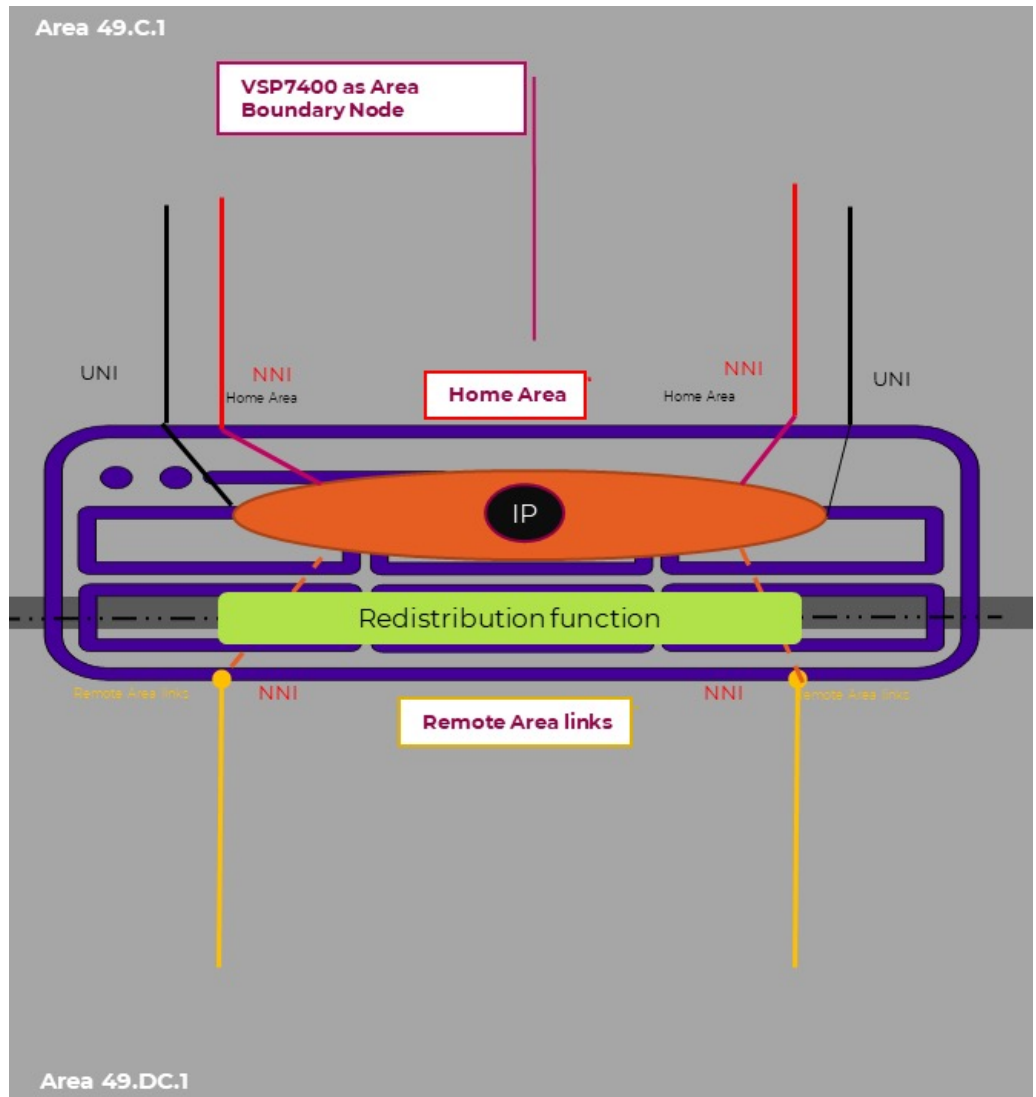


Figure 2: Multi-area architecture concepts (switch view)

The Home area is a concept on a boundary node where its local User interfaces (UNI) reside. This includes local IP interfaces and local services such as nick-name server, DVR controller or PIM-Gateway controller functionality. In this scenario, the boundary nodes would be inside the home area with external links to the 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 Extend) remote area links up to the total amount of supported NNI interfaces.

A Fabric Connect/Shortest Path Bridging network supports the following services:

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

By default, services are not redistributed between the home and remote area of a multi-area boundary node.

Redistribution policies can be configured to define which services are permitted to be stretched across multiple areas.

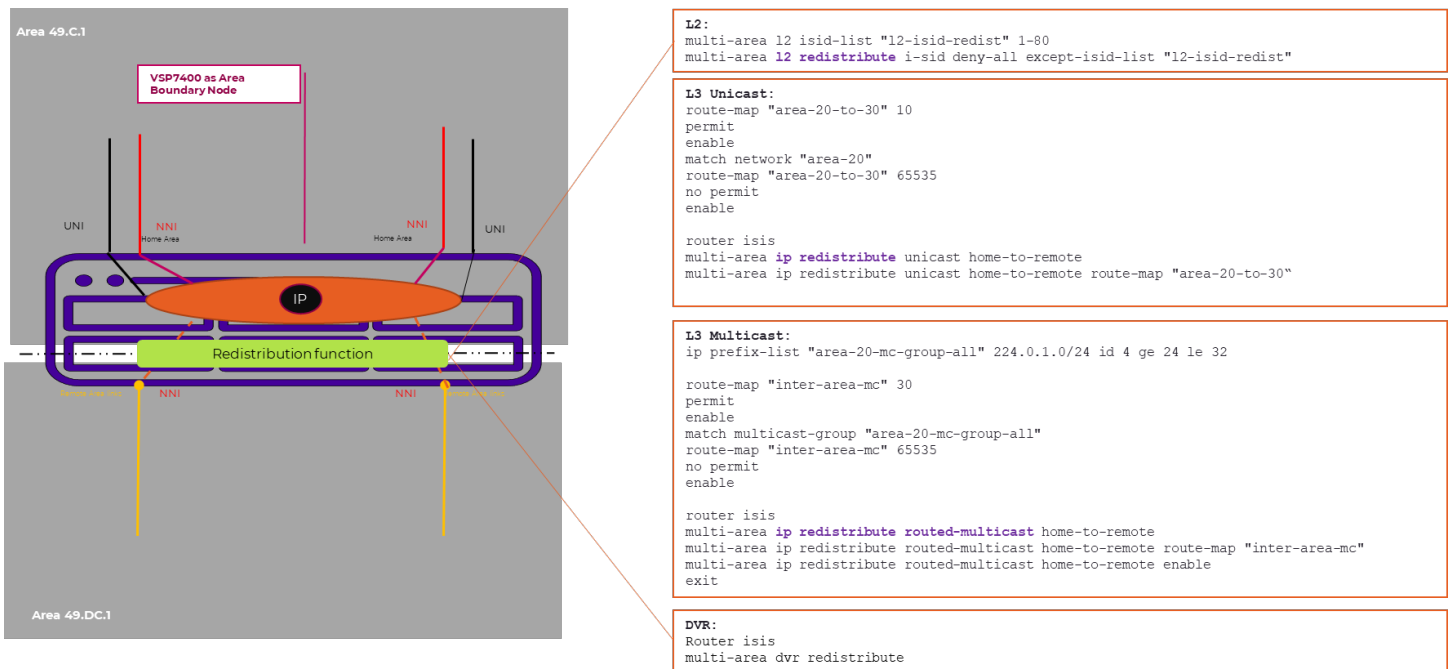


Figure 3: Redistribution policies

Multi-Area Operational Considerations

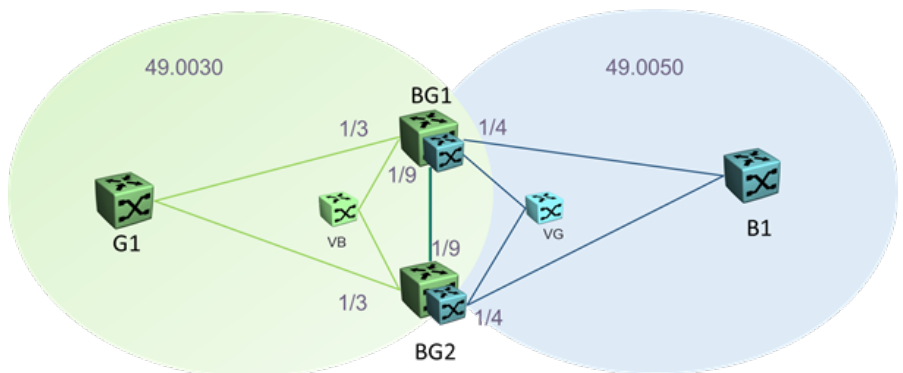


Figure 4: Multi-area operational concepts

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 representing all redistributed services from the connected area. For example, virtual node green (represented by VG on the diagram) represents area 49.0030 and its redistributed services in area 49.0050 and similarly, virtual node blue (represented by VB on the diagram) represents area 49.0050 and its redistributed services in area 49.0030.

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

When deploying multi-area, node count in an area includes the physical boundary nodes and the virtual nodes. 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 will appear as virtual-node and count as 1 additional node each (+1 to 2 boundary nodes)

When designing multi-area networks, it is recommended that you 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 will protect against area loops and will disable the area boundary node closing an area-loop.

You can chain up to 5 areas. This will still allow scaling to very large multi-area networks in a hierarchical design.

Multi-Area Network Design Concepts

In this section, multi-area design concepts are discussed. There are multiple ways to build scalable multi-area networks.

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

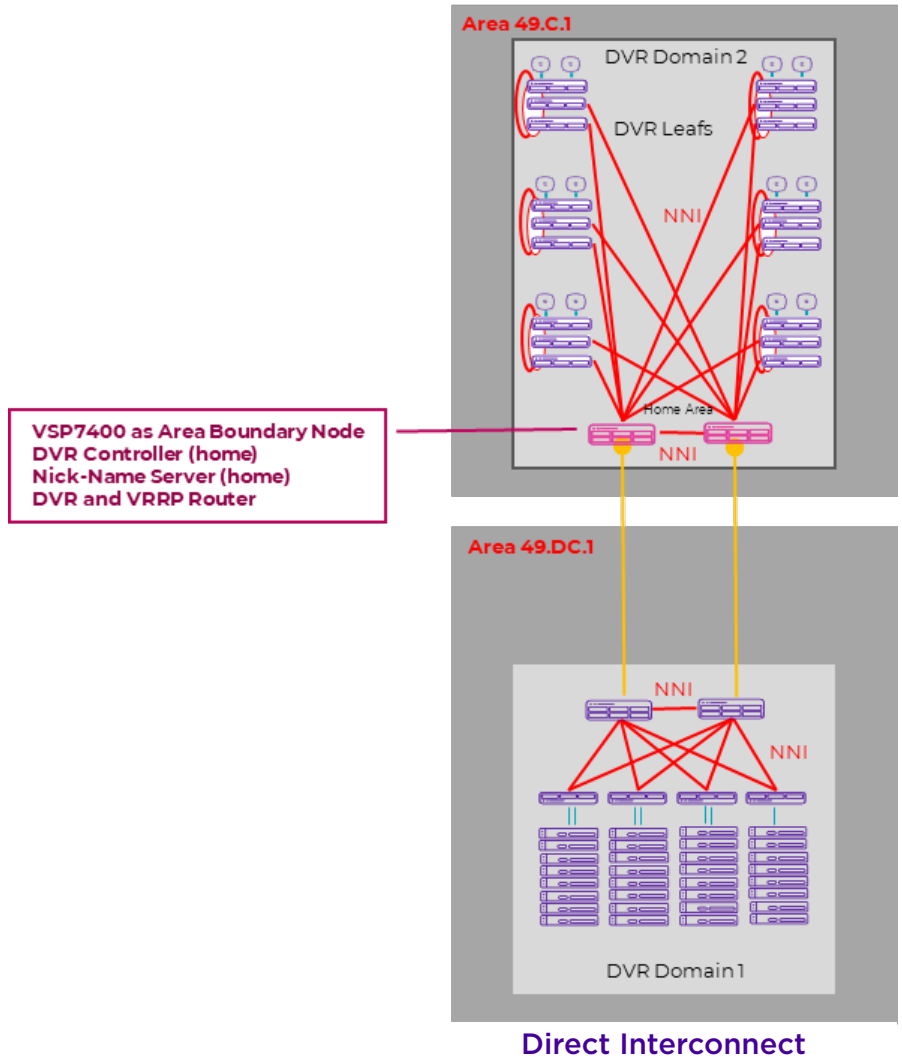


Figure 5: Use of boundary nodes to create multiple areas

Two VSP7400 are used in their Campus home area (49.C.1) to interconnect to a Data Center which is using IS-IS area id: 49.DC.1.

In this diagram the VSP 7400 switches are not only doing the area boundary function, but are also aggregating a series of VSP edge switches. For zero-touch deployments, they are also providing the seed node functionality by providing the nick-name server within the home area.

Additionally, this example shows a Campus DVR setup where the aggregation nodes are DVR controllers and the wiring closet (IDF) switches are DVR leafs. All uplinks are configured as SPB NNIs in the home area 49.C.1.

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

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.C.1 and another pair for the Data Center IS-IS area 49.DC.1, 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:

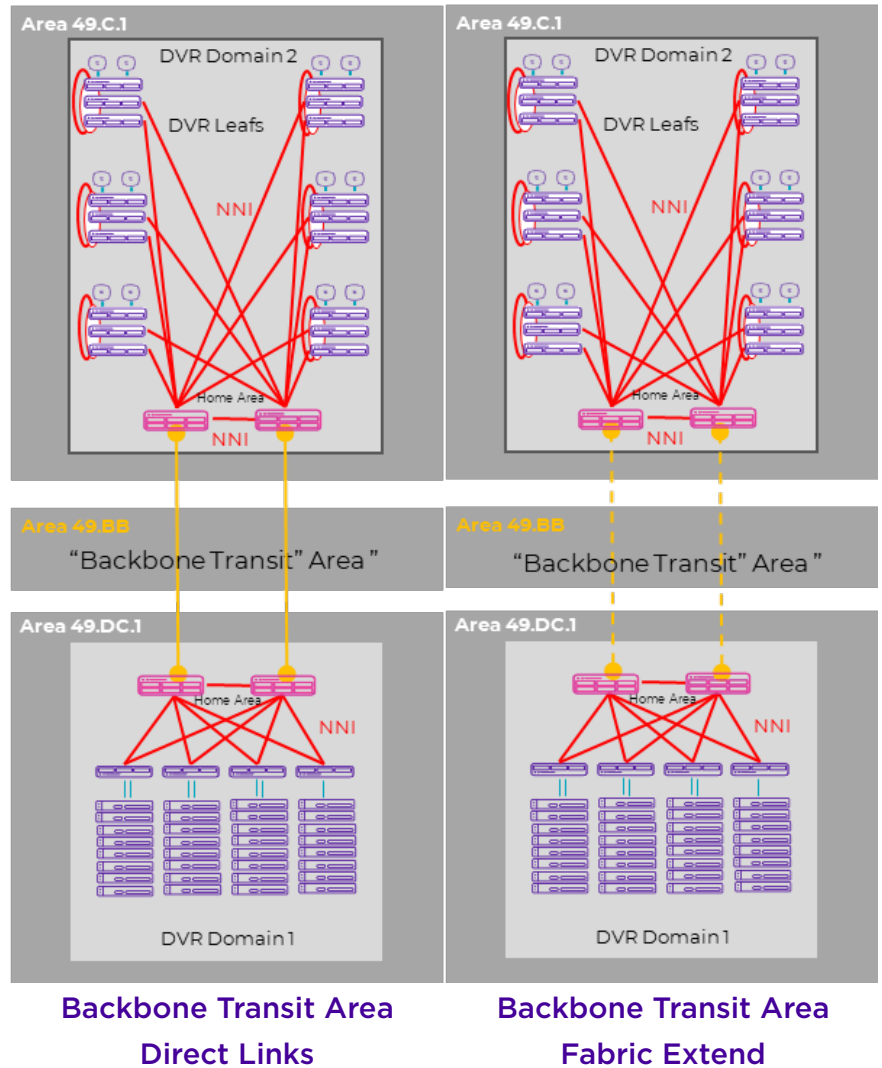


Figure 6: Use of Backbone Transit Areas to create multiple areas

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 is sufficient to create this BBT 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.

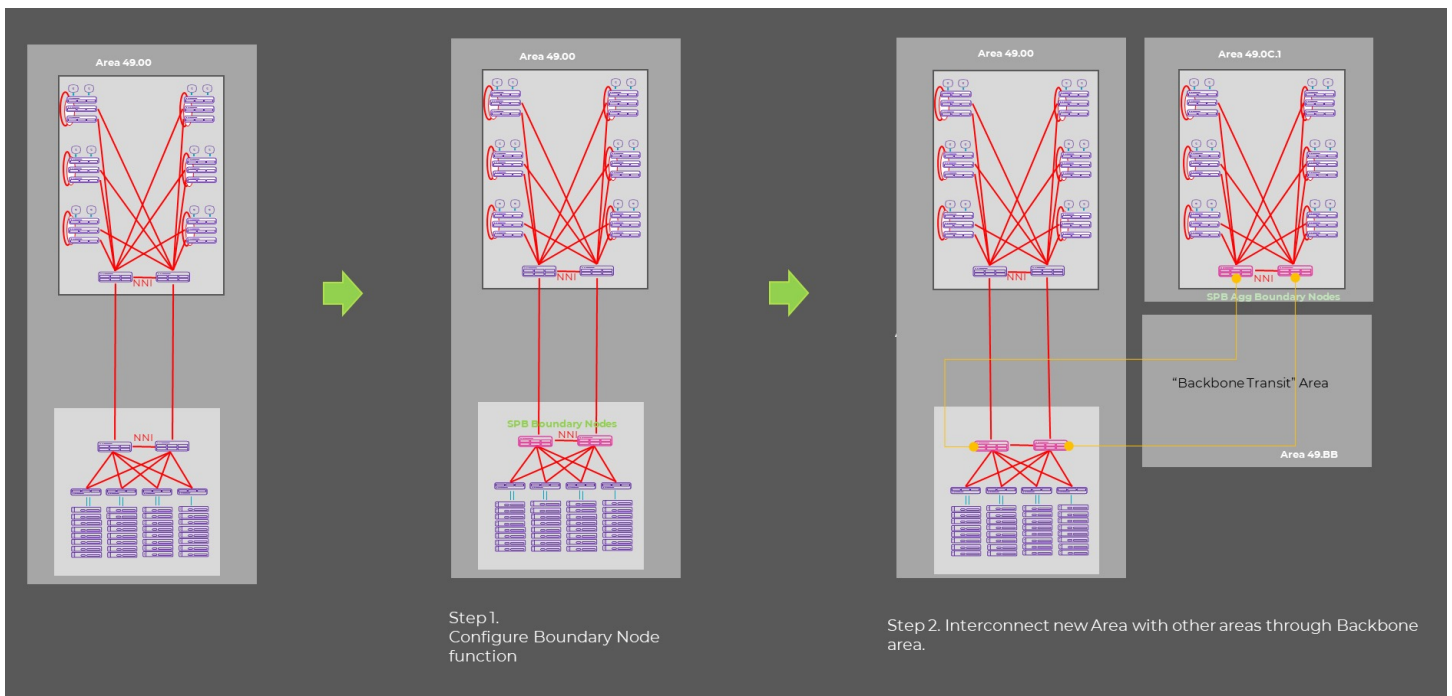


Figure 7: Migrating to a Backbone Transit Area

There are several BBT deployment options:

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

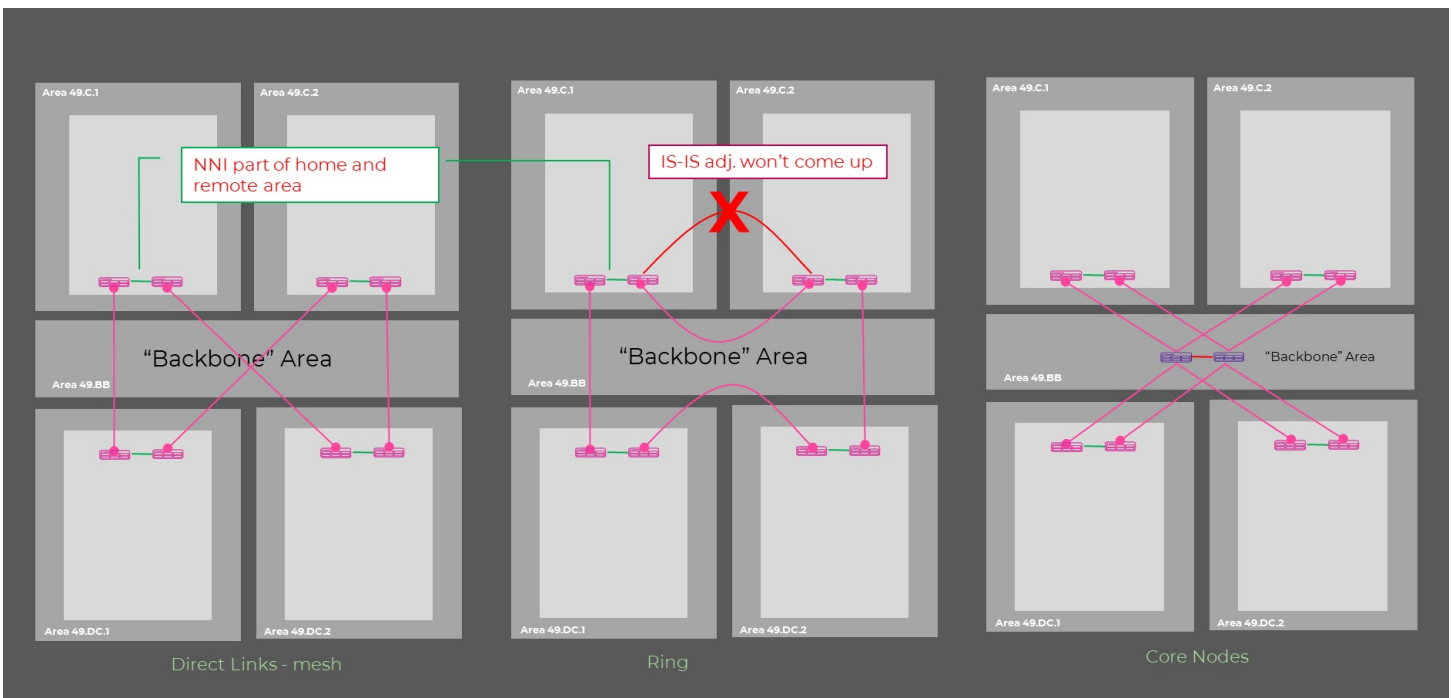


Figure 8: Design options for Backbone Transit Areas

These different options to design a Backbone Transit Area, illustrate the flexibility of a scale-out multi-area design. The main benefit being 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 allow for minimal disruption when expanding the network with additional areas since adding new areas does not impact existing area topologies. This type of design lends itself nicely to large campuses with 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 it is recommended to ensure that both areas of a boundary node (home and remote) are redundantly connected. Typically, a direct NNI link between the boundary nodes is used for this. This direct NNI link (or emulated link) should have both areas configured on it (home and remote), 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 to the design must ensure that both areas (home and remote) have a redundant connection between the boundary nodes. These 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. This is illustrated in Figure 8 above.

Migrating to a Fabric Connect Edge While Introducing Multi-Area

This section describes the migration from a traditional Fabric Attach/MLAG/SMLT edge to a Fabric Connect enabled edge, while at the same time, implementing multi-area.

Any existing Fabric Connect network can be scaled using this step-by-step approach to minimize the risk of disrupting network operation during the migration process.

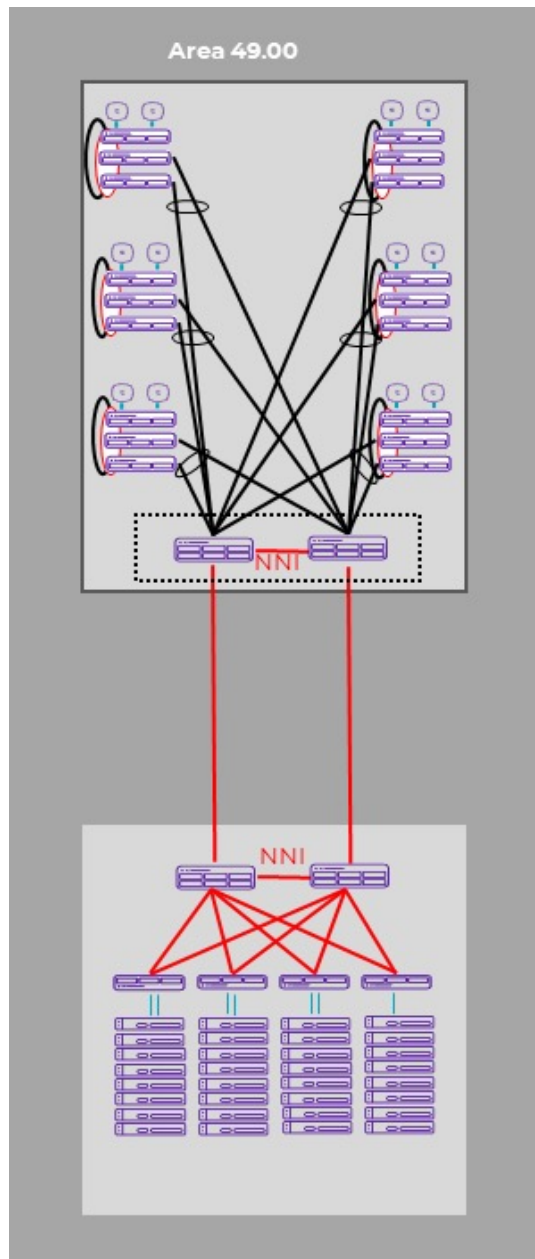


Figure 9: Edge Migration: PMO

In this example, a network with Campus and Data Center components are depicted. 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 typically also provide the default gateway function for the campus building.




Figure 10: Edge Migration: Step 1 - introduction of boundary nodes

To provide a fabric to edge solution, while introducing a new IS-IS area, it is recommended to introduce a new set of aggregation layer nodes that can provide the multi-area capability.

In this example, the existing MDF nodes and the new area boundary nodes are co-located. This is illustrated by the dotted square. The interconnections to the existing IS-IS area will be 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, inter-area redistribution policies must be configured.

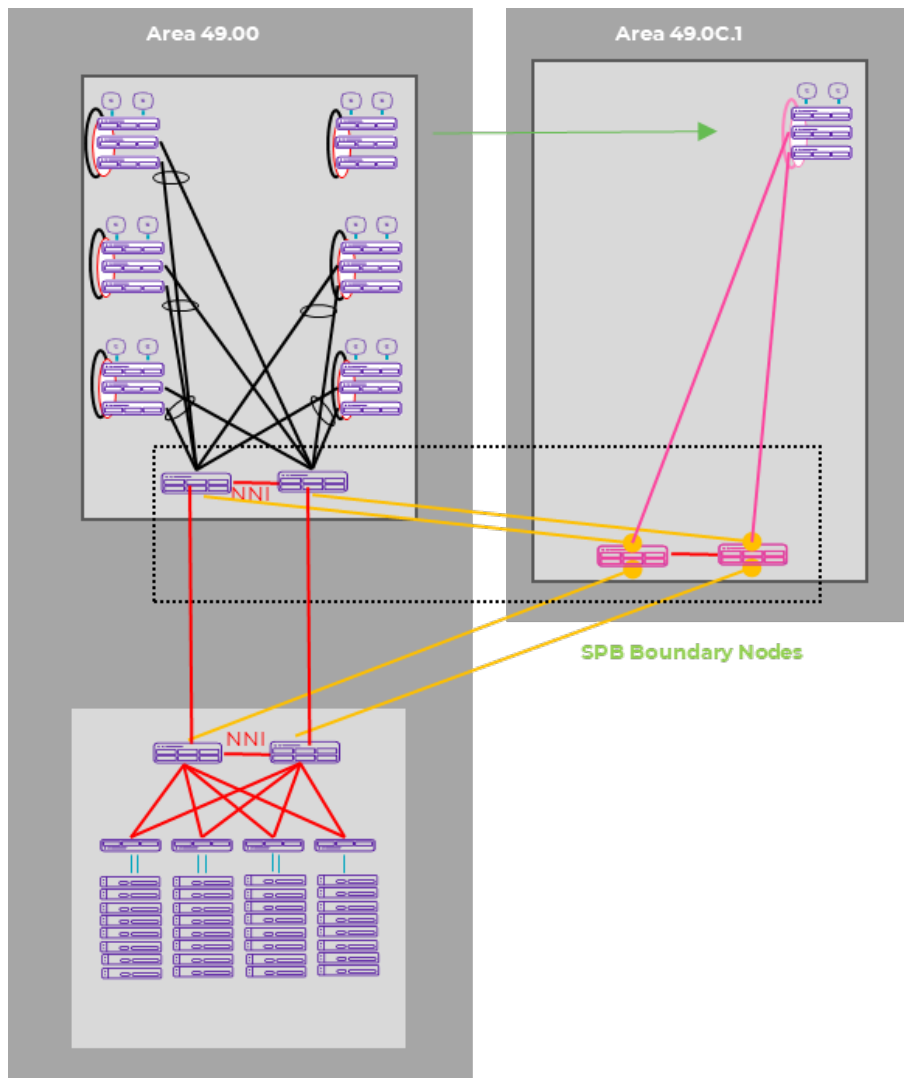


Figure 11: Edge Migration: Step 2 – introduction of new area for new wiring closets

Once the MDF switches are in place, the new fabric to the edge nodes can be introduced to the new area. Users can be migrated over to the new IDF switches at your own pace.

The IDF switches can either be deployed using the automated Extreme Management Center/ExtremeCloud IQ-Site Engine onboarding (ZTP+) and leveraging the Zero Touch Fabric (ZTF) deployment option, where new edge fabric nodes are deployed without any manual configuration.

In the ZTP+/ZTF deployment option, the switches join the fabric automatically, get their management IP address through DHCP and then onboard to the centralized management system where they then get their initial configuration.

To facilitate this deployment option, the onboarding ISID 15999999 will need to provide reachability to the management segment as well as the DHCP server.

For ZTF, a Nick-name server needs to be enabled in the new Campus area (49.0C1) on the MDF multi-area boundary nodes.

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

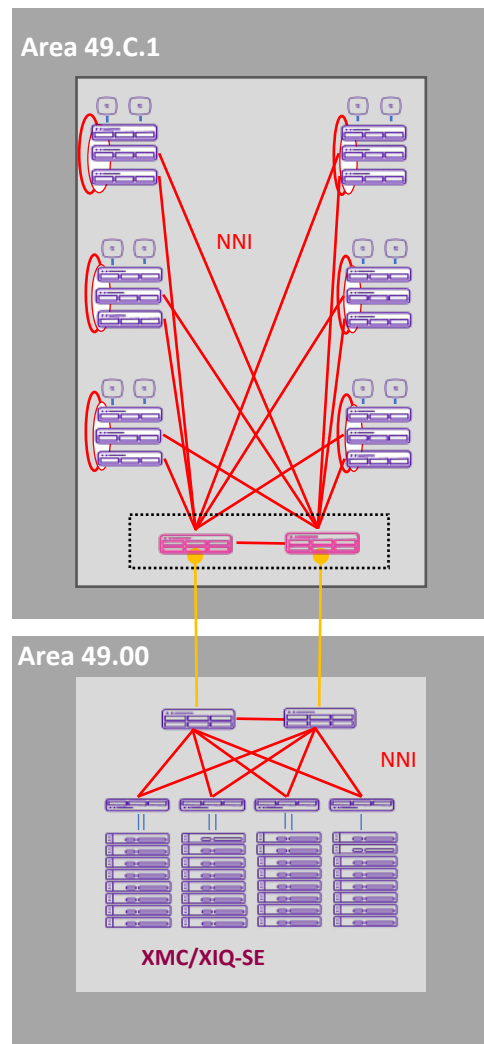


Figure 12: Edge migration - remove old aggregation nodes after migration of services

Once all users have been migrated from the old IDF switches, all old switches can be decommissioned or redeployed for other use cases.

The result of this migration is that a now the network consists of two SPB IS-IS areas and the campus area is enhanced with more automation as Fabric Connect is extended right to the network edge.

The same approach can be used for fabric to the edge solution with ERS switches.

Fabric Edge and Multi-Area

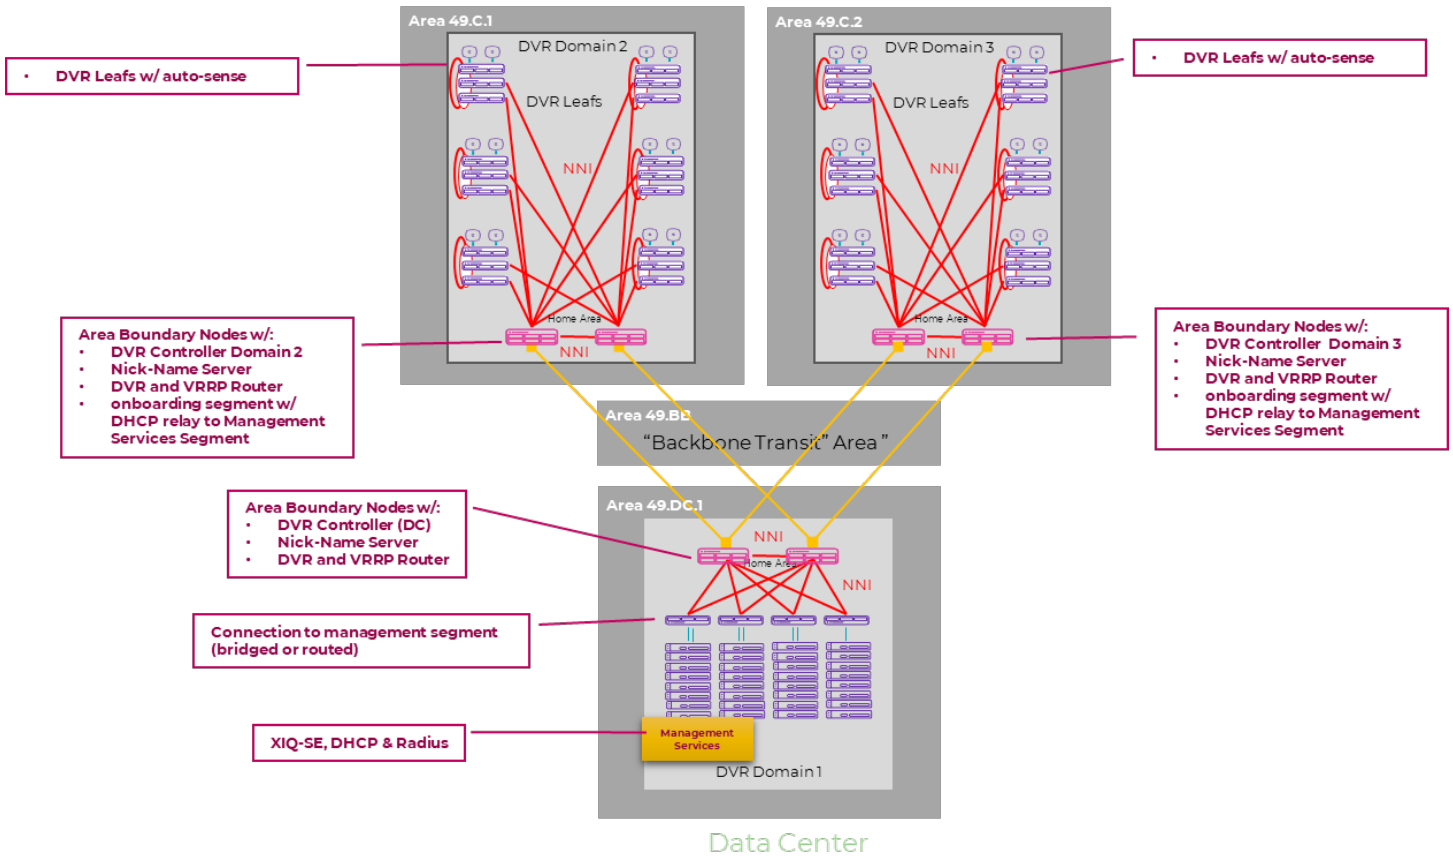


Figure 13: Fabric Edge and Multi-area constructs

This diagram shows all elements of a Fabric edge solution (based on VOSS-enabled switches) with an IS-IS multi-area design approach for larger deployments, where the total fabric node count exceeds 500.

The network does not require the use all of the features that enable automation such as ZTF, ZTP+ and DVR; however, these features are recommended to reap all the benefits of simplified operations.

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

- Multi-area boundary nodes (for large buildings where many Fabric Connect-enabled edge switches are being deployed) with the home area being the switches inside the building.
- Nick-Name server to enable Zero Touch Fabric Deployment for the IDF switches.
- DVR controller enabling distributed L3 routing capability for stretched subnets, as well as enabling per ISID multicast. This removes 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 .

The Data Center Spine nodes shall provide similar configuration to the IDF nodes if a multi-area design is deployed.

Fabric to the Edge with Multi-Area Blueprint

Scale-out multi-area blueprint with Backbone Transit Area and several fabric and non-fabric enabled IDF closets. Solution consists of 3 campus areas, 1 data center area and a WAN area.

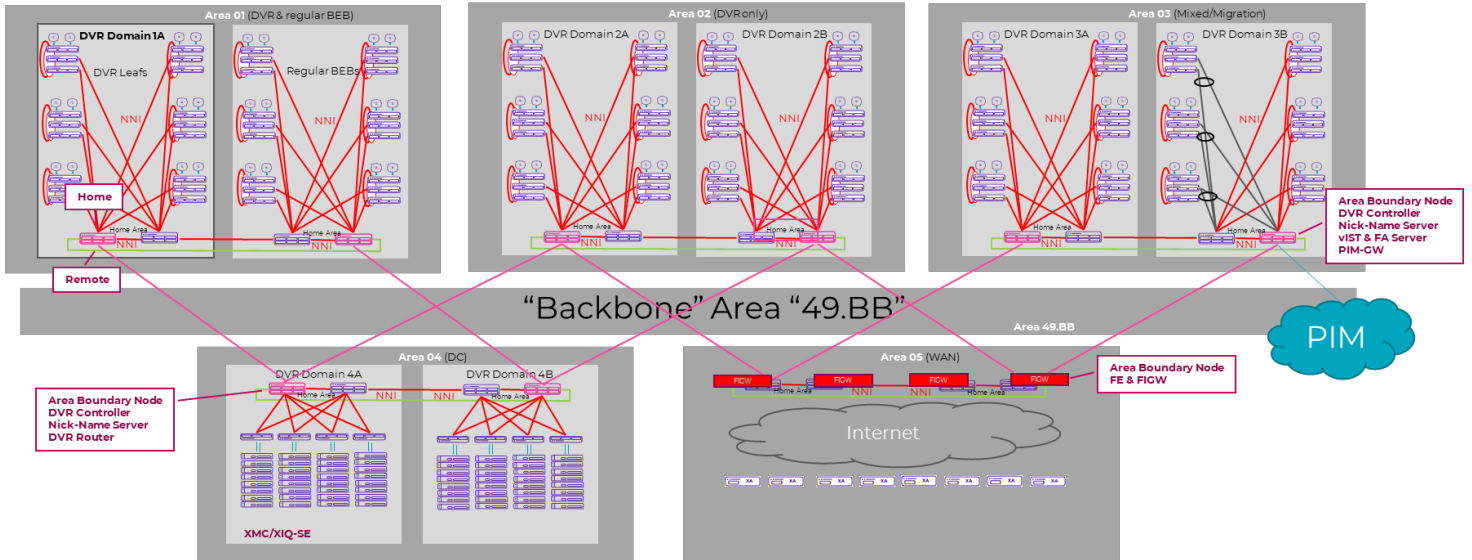


Figure 14: Large multi-area deployment with Backbone Transit Area

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 out solution, enabling networks to scale to tens of thousands of fabric nodes.




Figure 15: Designing for massive scalability

For more information:

For more information on Fabric Connect/Shortest Path Bridging multi-area, please check out the following video or contact your Extreme sales representative.

- [Multi-area: Ultimate scaling for the Infinite Enterprise](#)



<http://www.extremenetworks.com/contact>

©2021 Extreme Networks, Inc. All rights reserved. Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries. All other names are the property of their respective owners. For additional information on Extreme Networks Trademarks please see <http://www.extremenetworks.com/company/legal/trademarks>. Specifications and product availability are subject to change without notice. 39940-0821-23