



The bridge to possible

Deployment of VXLAN EVPN Gateways with Cisco ACI for the Interconnection of Heterogeneous Data Center Fabrics

Max Ardica, Distinguished TME
@maxardica
BRKDCN-2634



cisco *Live!*

#CiscoLive

Cisco Webex App

Questions?

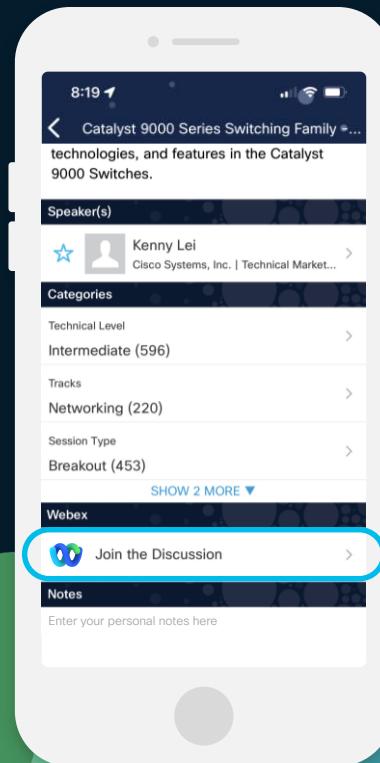
Use Cisco Webex App to chat with the speaker after the session

How

- 1 Find this session in the Cisco Live Mobile App
- 2 Click “Join the Discussion”
- 3 Install the Webex App or go directly to the Webex space
- 4 Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until June 7, 2024.

[https://ciscolive.ciscoevents.com/
ciscolivebot/#BRKDCN-2634](https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-2634)

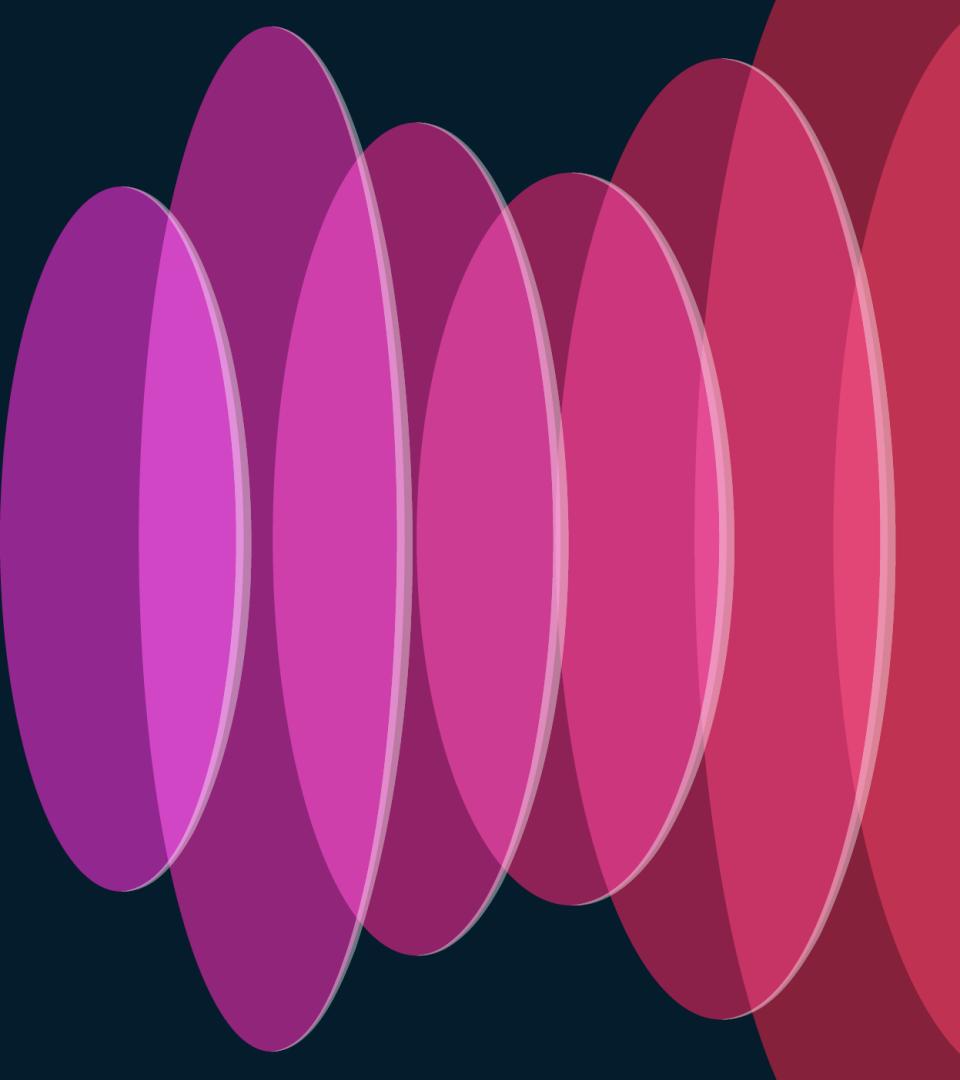




Agenda

- Introducing Cisco Nexus One Fabric Experience
- ACI Border Gateways (BGWs)
 - Introduction
 - Overview of Control-Plane and Data-Plane
 - Namespace Normalization
 - Workload Mobility across Domains
 - Policy Enforcement on ACI BGWs
- Secure Interconnection of Heterogeneous Fabrics

Introducing Cisco Nexus One Fabric Experience



What is Cisco Nexus One fabric experience?

Open networking Fabric Experience

Evolve multiple DCN fabrics into a single user experience to deliver consistent use cases

Nexus One Fabric Experience - Overview

3

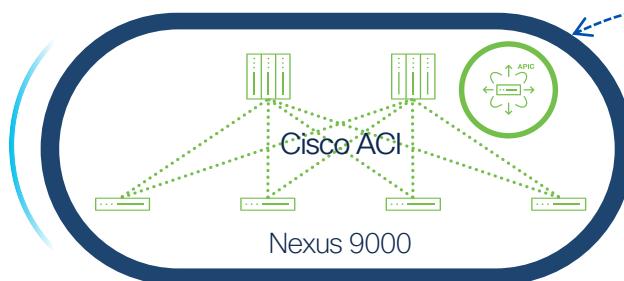
Cisco Nexus Dashboard as single point of control and operations

Cisco Nexus Dashboard



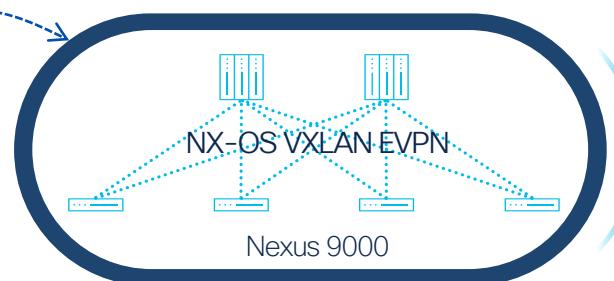
1

ACI VXLAN EVPN
Border Gateways



2

Policy in NX-OS
(Security Groups)

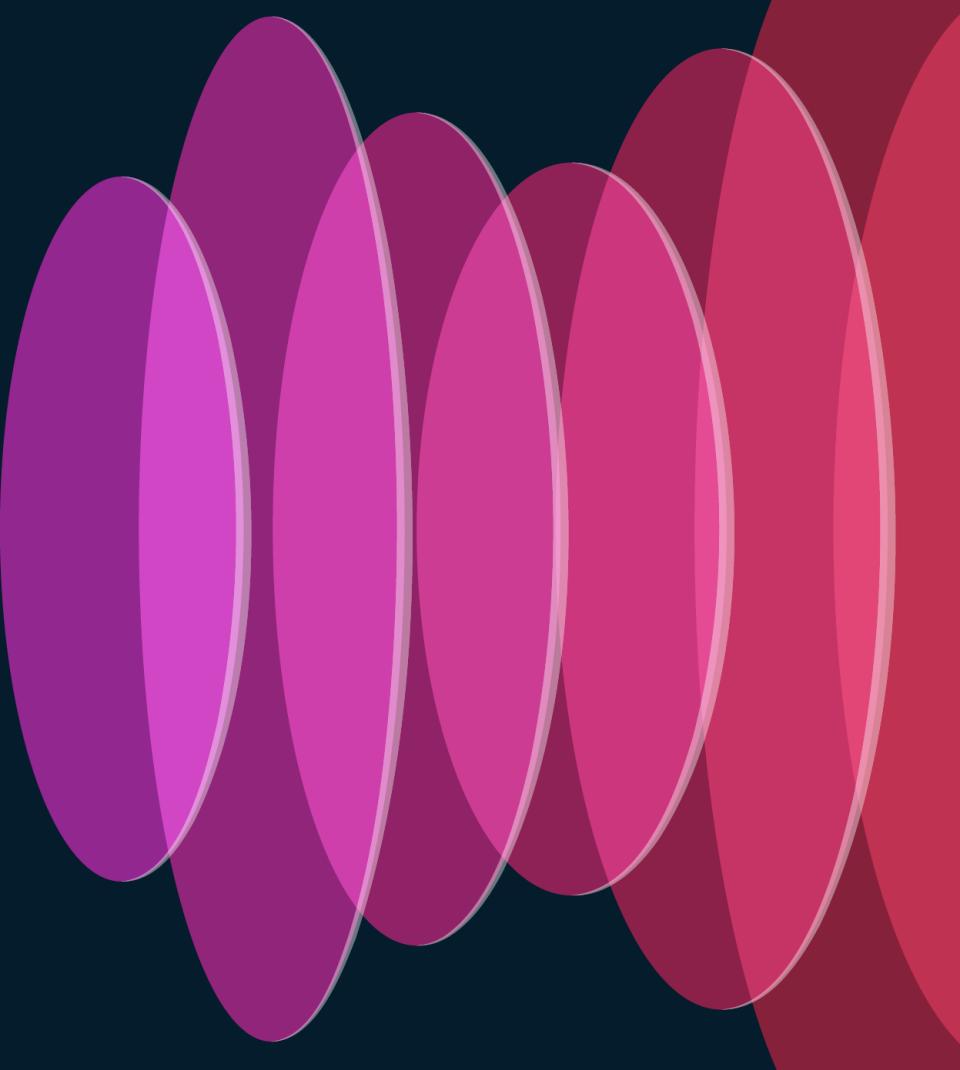


Different fabric architectures

Same outcome with common experience

ACI Border Gateways

Introduction



DC Design Evolution

From a Single Large Fabric to a Distributed Architecture

The slide is titled "Design Evolution" and subtitle "From a Single Large Fabric to a Distributed Architecture". It features two diagrams: "2 Tier Leaf Spine (5 Stages)" and "3 Tier Leaf-Fabric-Spine (5 Stages)".

2 Tier Leaf Spine (5 Stages): This diagram shows four Spine (S) nodes at the top, each connected to a group of Leaf (L) nodes below. The connections are dense, indicating a single large fabric design.

3 Tier Leaf-Fabric-Spine (5 Stages): This diagram shows a three-tier structure. The top tier has four Spine (S) nodes. The middle tier has four Fabric (F) nodes, each connected to a group of Leaf (L) nodes below. The connections are more sparse and hierarchical than the 2-tier model.

CISCO Live!

BRKDCN-2999

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

35

For more information on DC Multi-Tier Design Evolution please refer to BRKDCN-2099

DC Design Evolution

From a Single Large Fabric to a Distributed Architecture

The diagram is a screenshot of a presentation slide titled "Design Evolution: From a Single Large Fabric to a Distributed Architecture". It features two network architecture diagrams side-by-side, connected by a large orange arrow pointing from left to right.

Left Diagram: 2 Tier Leaf-Spine (5 Stages)

This diagram shows a single large fabric structure. At the top, there are four Spine (S) nodes, each connected to a group of Leaf (L) nodes below. The Leaf nodes are arranged in two rows: a top row of three and a bottom row of five. All connections are represented by grey lines.

Right Diagram: 3 Tier Leaf-Fabric-Spine (5 Stages)

This diagram illustrates a more distributed architecture. It consists of three tiers: a top tier of four Core (C) nodes, a middle tier of eight Spine (S) nodes, and a bottom tier of ten Leaf (L) nodes. Every node in one tier is connected to every node in the tier directly below it, with connections shown as blue lines.

Central Text: Let's Move Forward

Cisco Live!

BRKDCN-2999

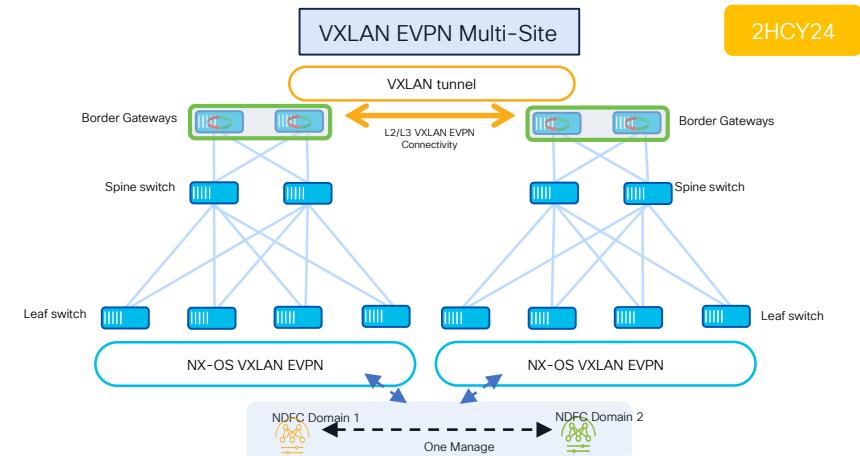
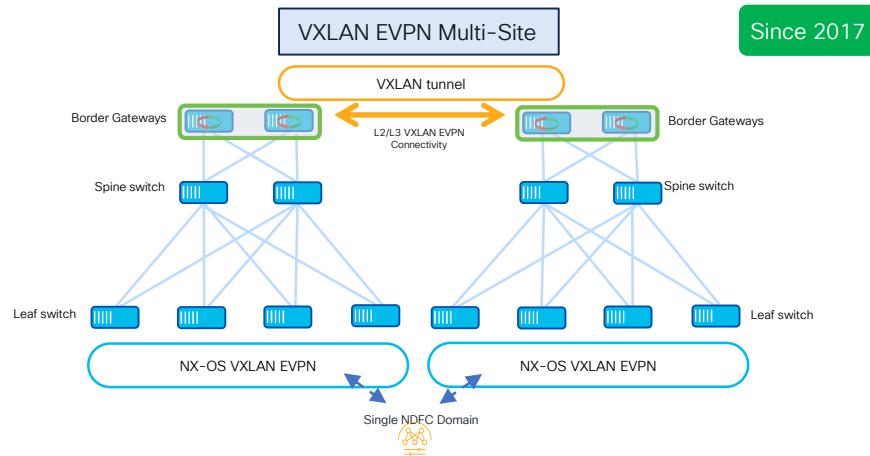
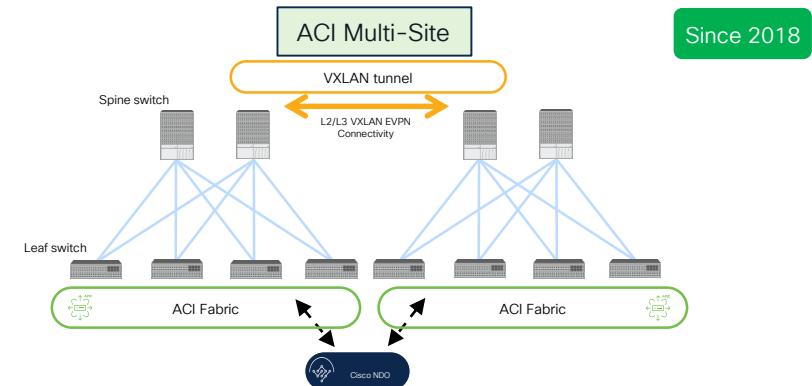
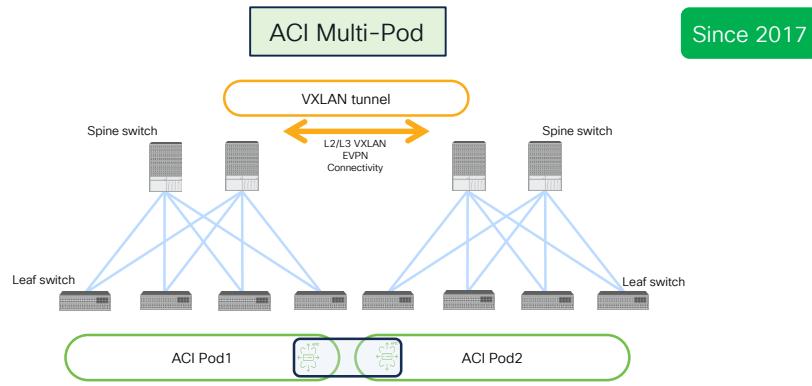
© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

36

For more information on DC Multi-Tier Design Evolution please refer to BRKDCN-2099

Building Distributed DC Architectures

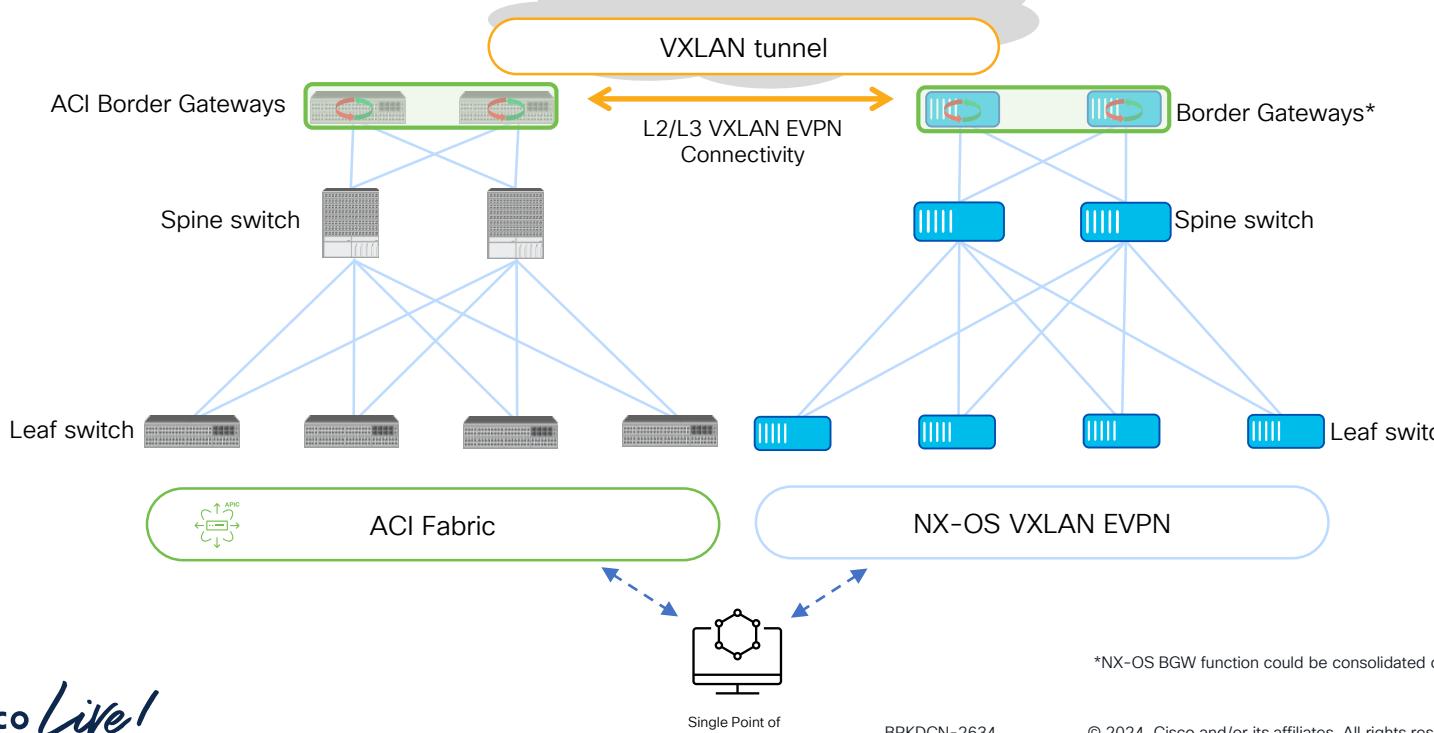
Homogeneous Options



Heterogeneous Fabrics

Introducing ACI Border Gateways

“Opening Up” L2/L3 Connectivity between ACI and VXLAN EVPN Fabrics



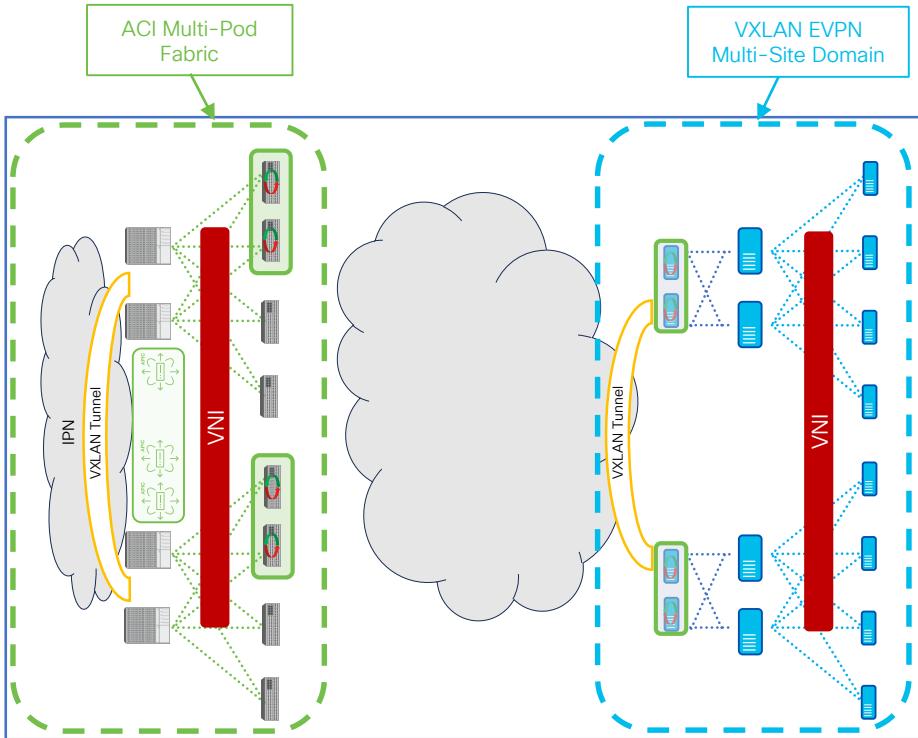
ACI Border Gateways

Deployment Considerations for 6.1(1) Release

- Hardware support for ACI BGWs: Nexus 9000 FX2 and above
- Dedicated leaf nodes for Border Gateway functionality
 - Coexistence with Border Leaf functions (L3Outs) planned for a future release
- IGMP snooping and L3 Multicast traffic not supported across domains
 - L2 Multicast traffic forwarded as BUM
- Symmetric namespace between ACI and VXLAN EVPN domains
 - VNI must be defined in the VXLAN EVPN domain to match the APIC assigned VNIDs
- “VRF unenforced” required on the ACI fabric for VRFs that need to be stretched
- Support for a single ACI fabric (can be Multi-Pod)

Heterogeneous Fabrics

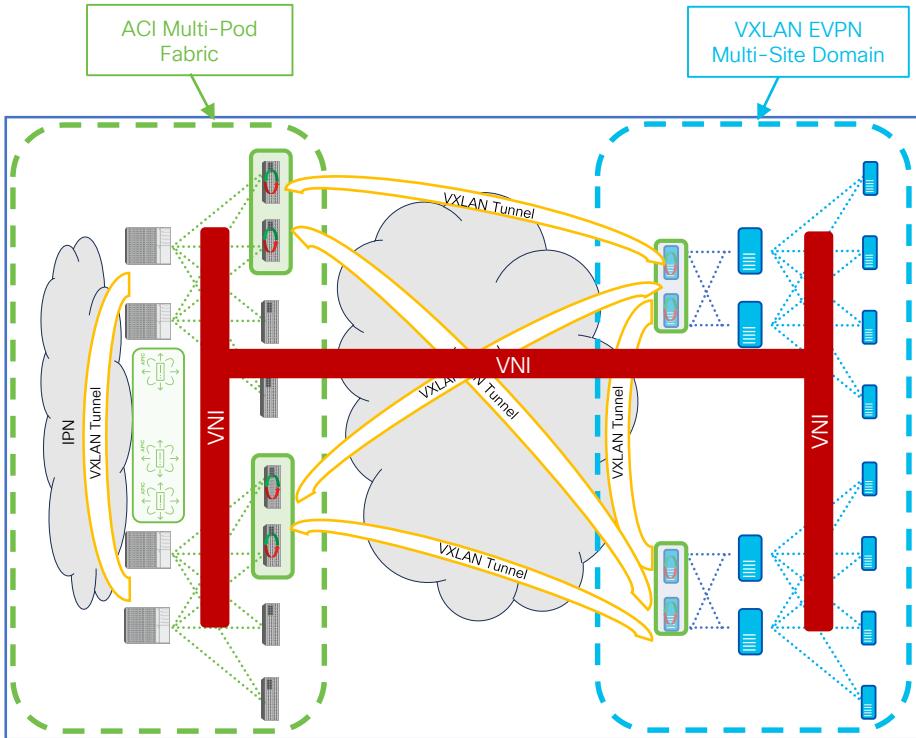
ACI Multi-Pod Fabric Support



- L2/L3 VXLAN connectivity between ACI Pods part of the same fabric achieved via the spine-to-spine data path (through the IPN)
 - No VXLAN EVPN connectivity between ACI BGWs of different ACI Pods
- Local instance of ACI BGWs mandatory in each Pod
- For each BD extended across domains, a specific ACI BGW is elected as DF (across all the BGWs in all the Pods)

Heterogeneous Fabrics

ACI Multi-Pod Fabric Support

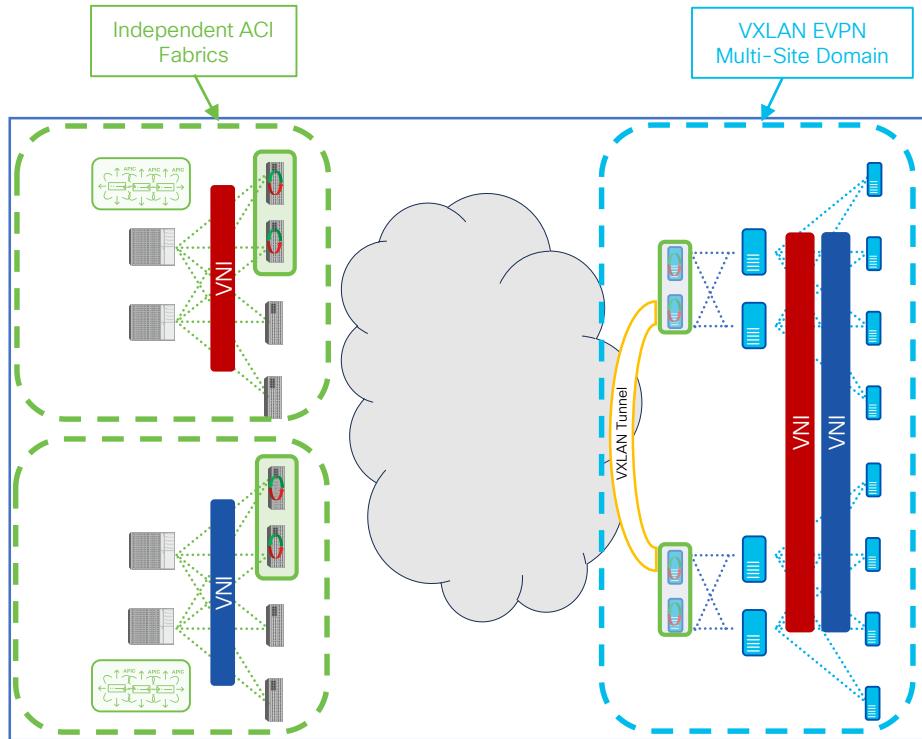


- L2/L3 VXLAN connectivity between ACI Pods part of the same fabric achieved via the spine-to-spine data path (through the IPN)
 - No VXLAN EVPN connectivity between ACI BGWs of different ACI Pods
- Local instance of ACI BGWs mandatory in each Pod
- For each BD extended across domains, a specific ACI BGW is elected as DF (across all the BGWs in all the Pods)

Heterogeneous Fabrics

Independent ACI Fabrics Support

Future

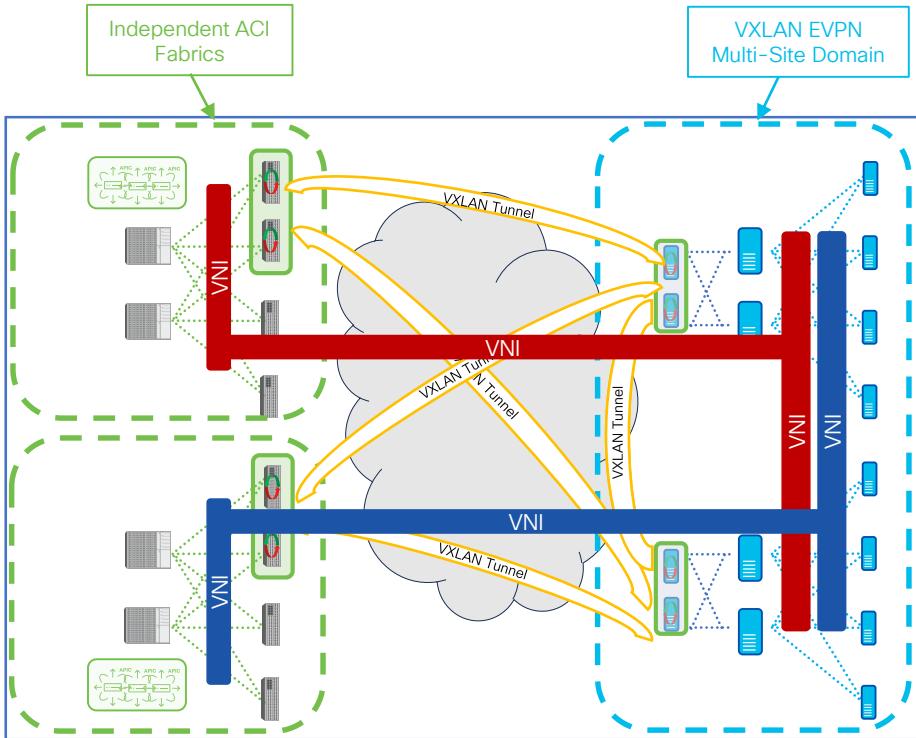


- Routed communications only via L3Out path possible between independent ACI fabrics
 - No VXLAN EVPN connectivity between ACI BGWs of different ACI Fabrics
- Different sets of VRFs/BDs can be extended between each ACI fabric and the VXLAN EVPN domain

Heterogeneous Fabrics

Independent ACI Fabrics Support

Future

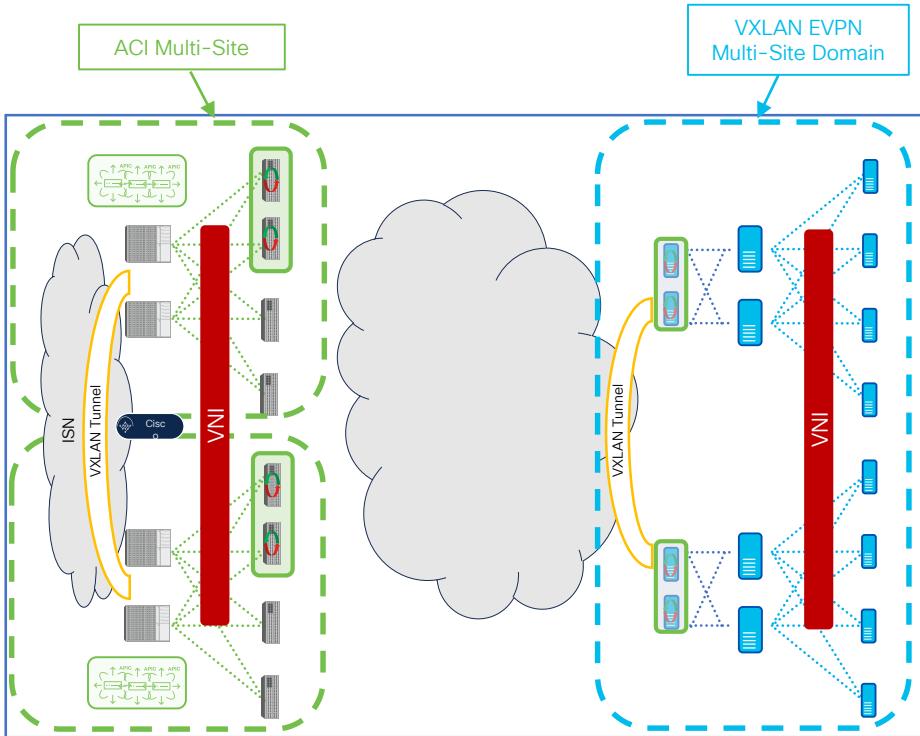


- Routed communications only via L3Out path possible between independent ACI fabrics
 - No VXLAN EVPN connectivity between ACI BGWs of different ACI Fabrics
- Different sets of VRFs/BDs can be extended between each ACI fabric and the VXLAN EVPN domain

Heterogeneous Fabrics

ACI Multi-Site Support

Future

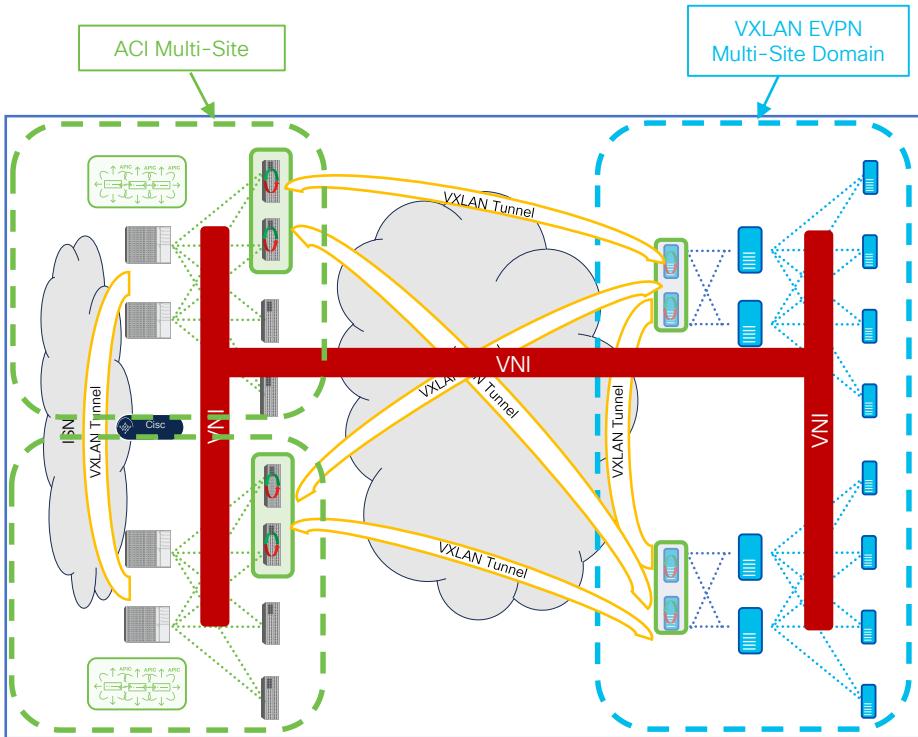


- L2/L3 VXLAN connectivity between ACI fabrics achieved via the spine-to-spine data path
 - No VXLAN EVPN connectivity between ACI BGWs of different ACI fabrics
- Each ACI fabric leverages a local instance of ACI BGWs to establish VXLAN EVPN connectivity with other domains
- NDO used for extending connectivity between ACI fabrics

Heterogeneous Fabrics

ACI Multi-Site Support

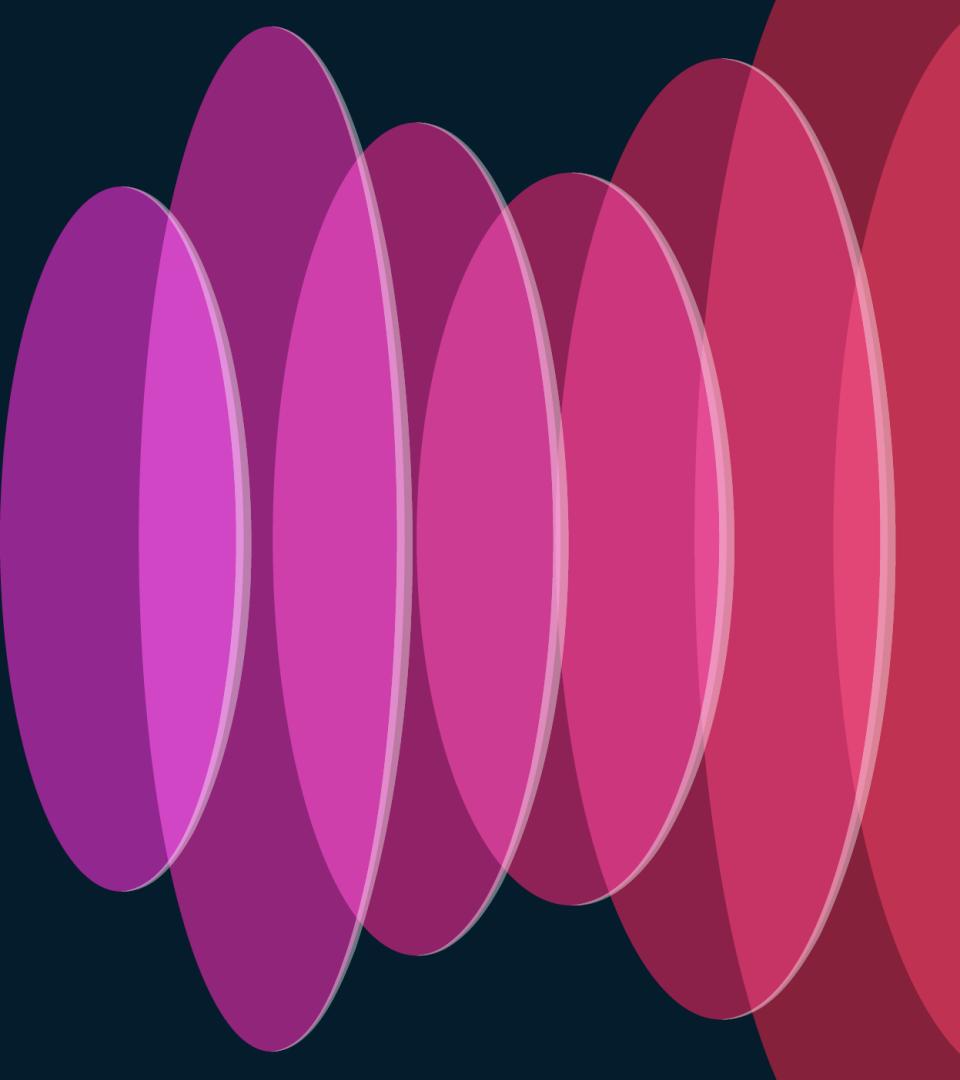
Future



- L2/L3 VXLAN connectivity between ACI fabrics achieved via the spine-to-spine data path
 - No VXLAN EVPN connectivity between ACI BGWs of different ACI fabrics
- Each ACI fabric leverages a local instance of ACI BGWs to establish VXLAN EVPN connectivity with other domains
- NDO used for extending connectivity between ACI fabrics

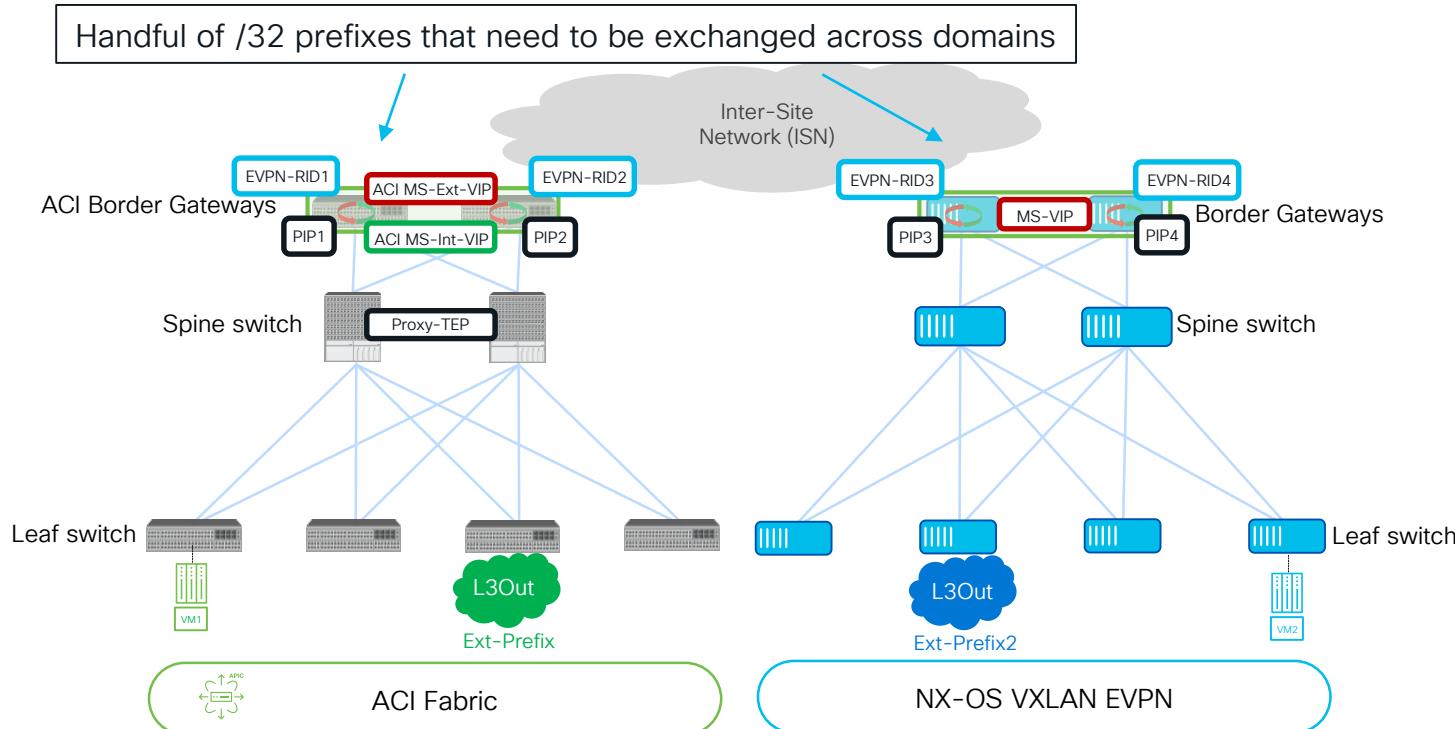
ACI Border Gateways

Overview of Control-Plane and Data-Plane



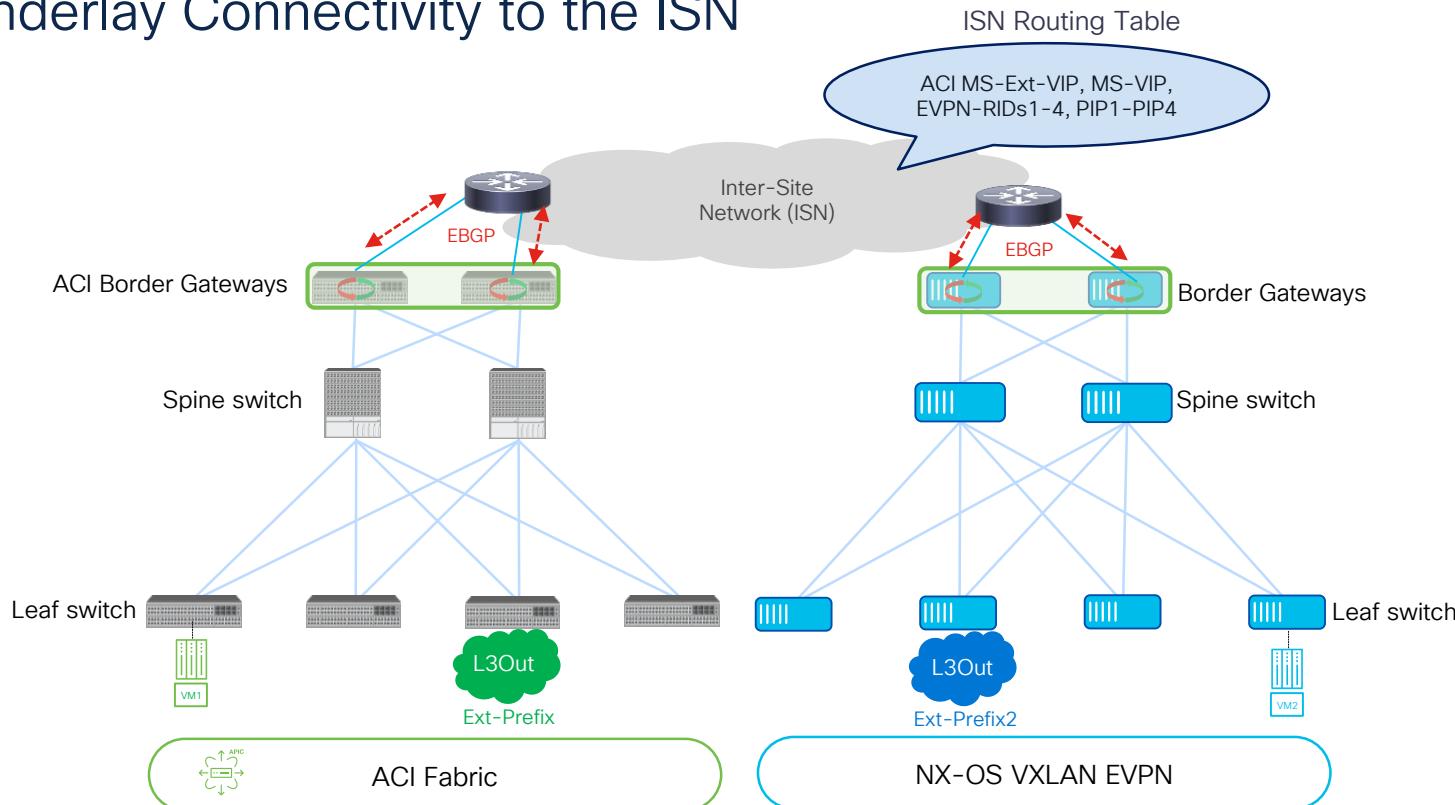
ACI Border Gateways

External and Internal Multi-Site VIP Addresses



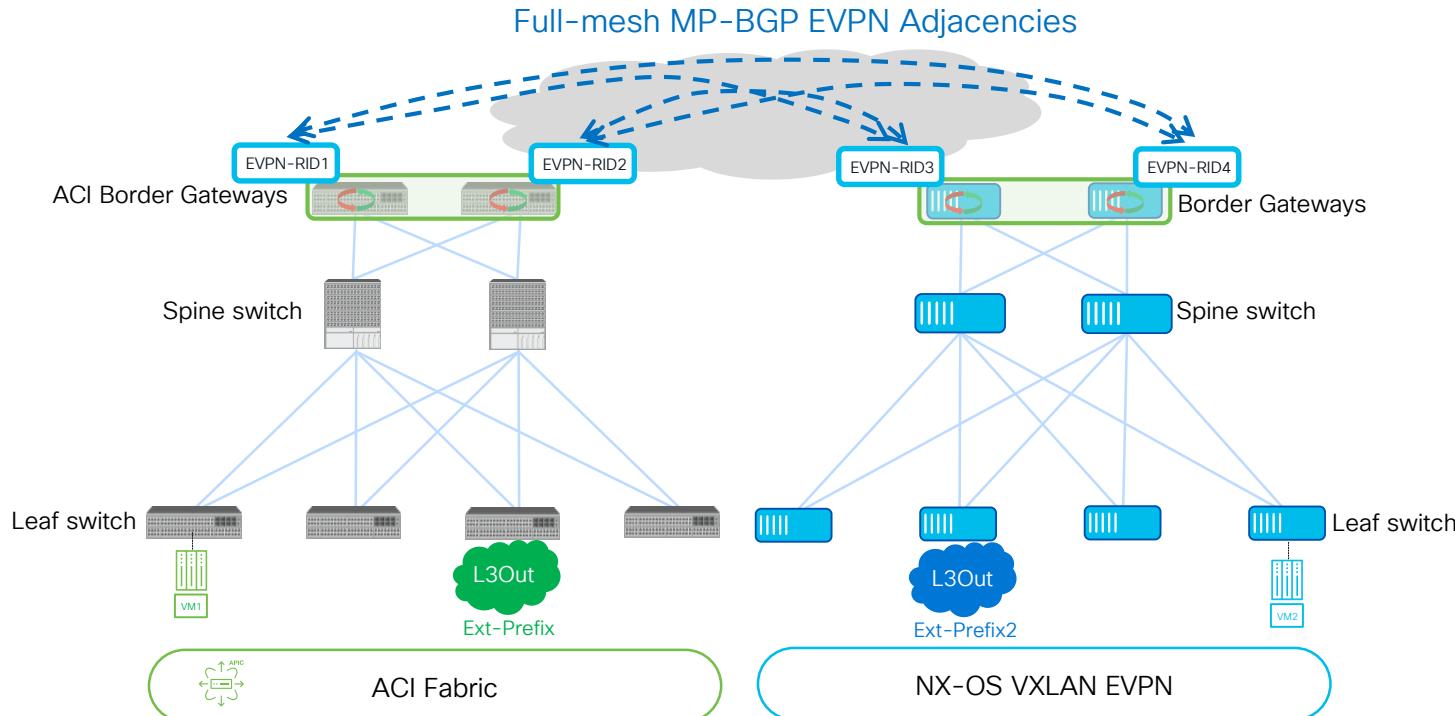
ACI Border Gateways

Underlay Connectivity to the ISN



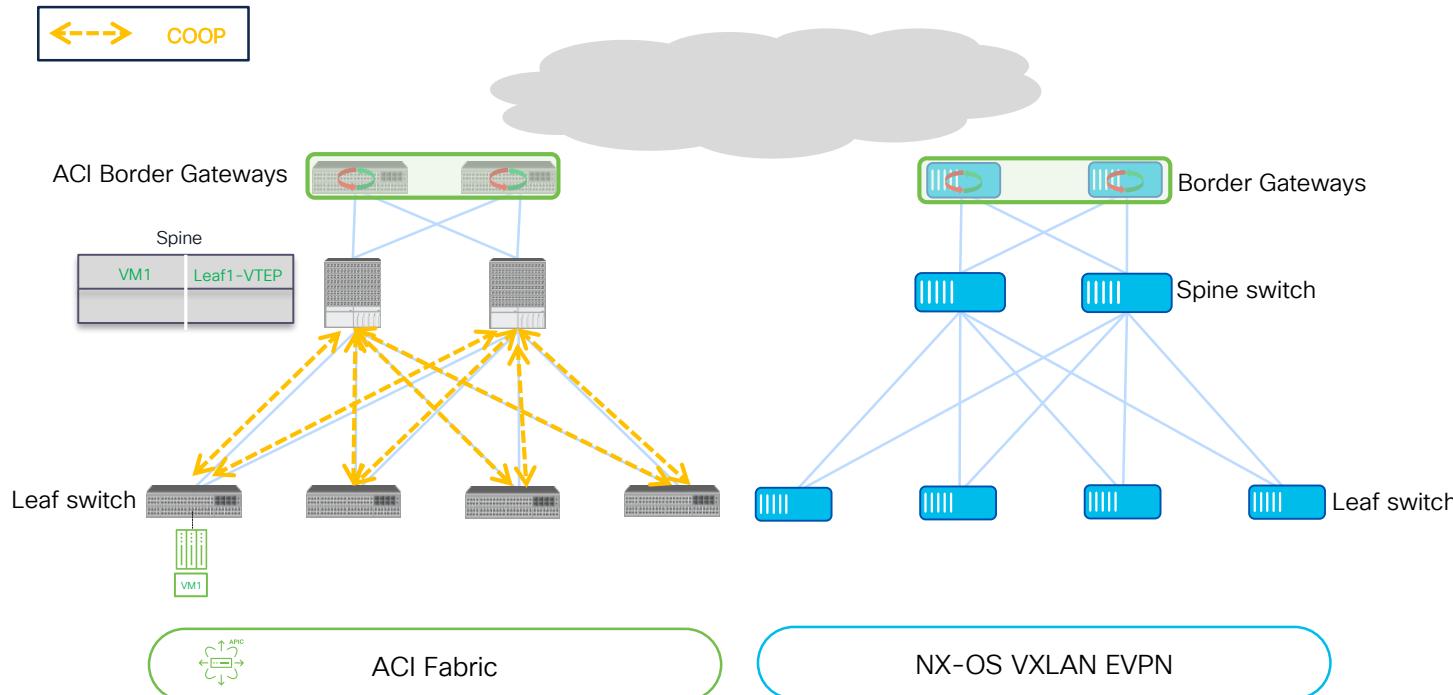
ACI Border Gateways

Overlay EVPN Connectivity between BGWs



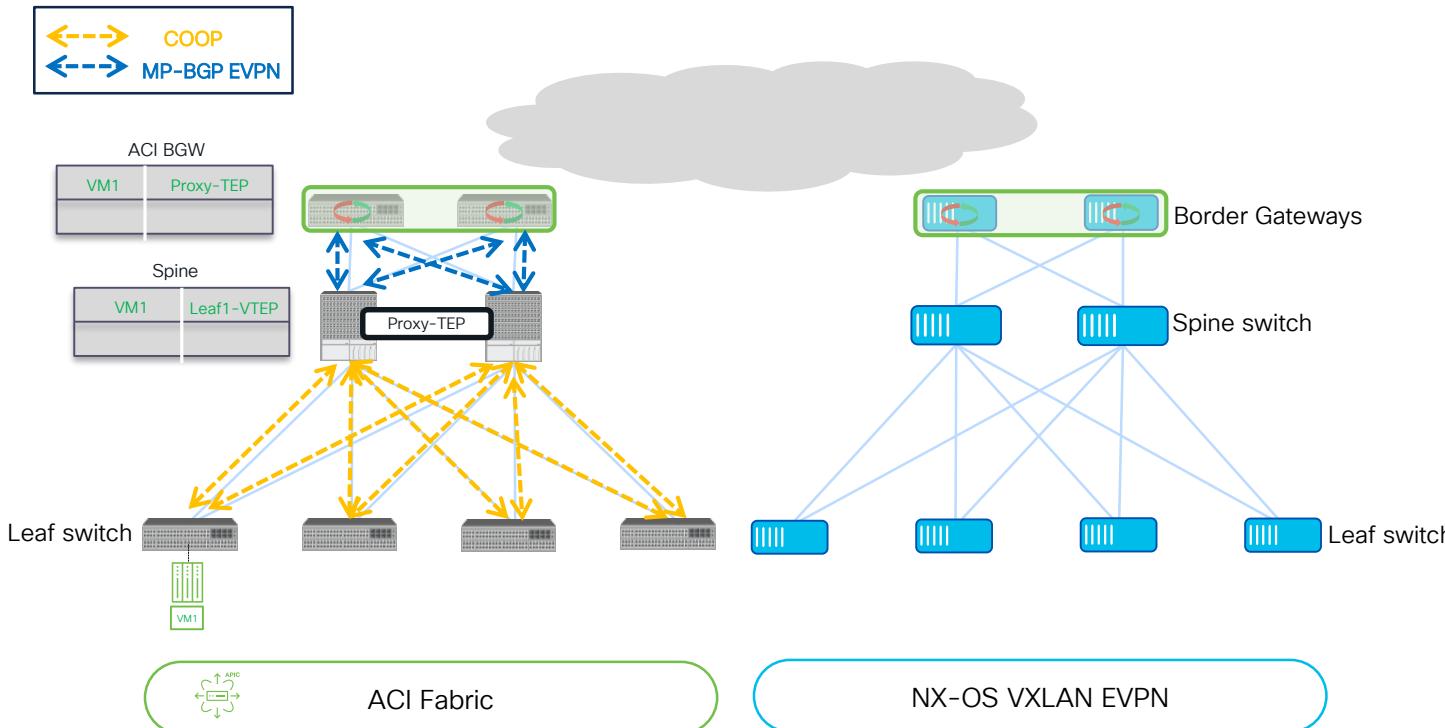
ACI Border Gateways

Control-Plane Overview



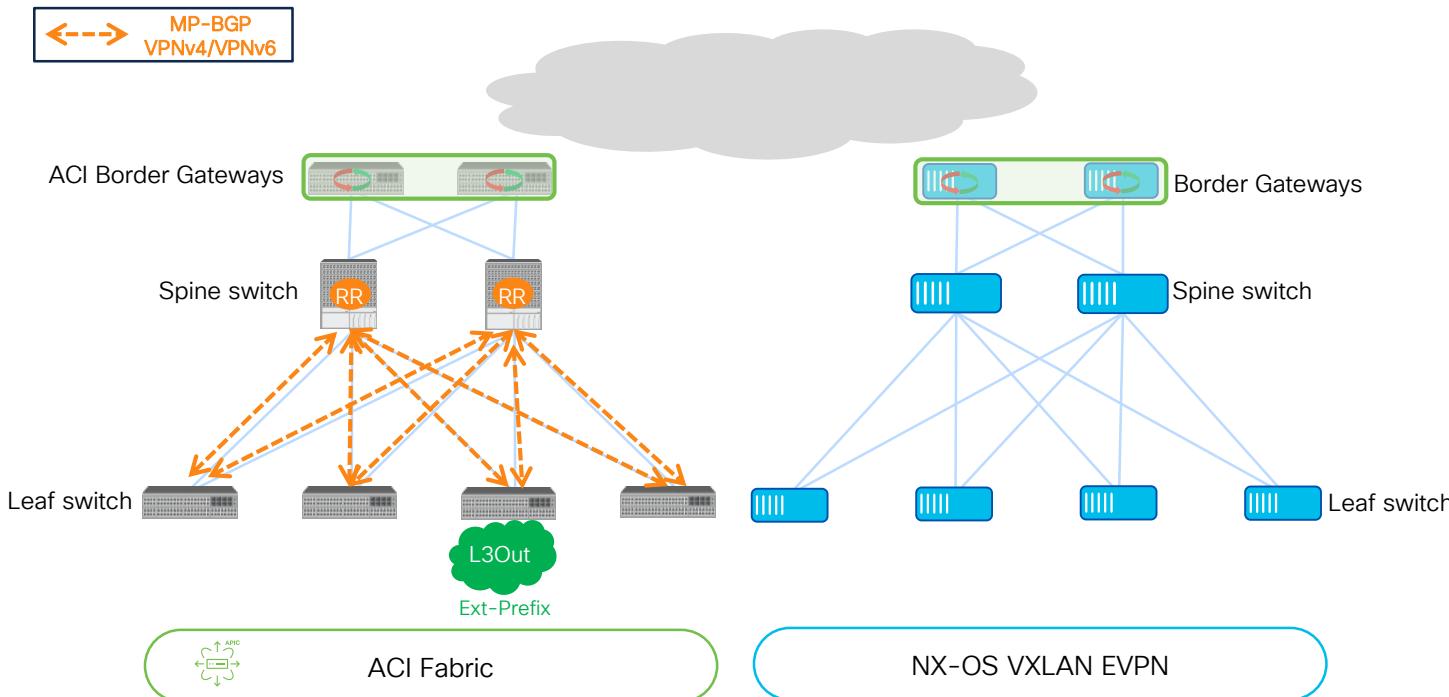
ACI Border Gateways

Control-Plane Overview



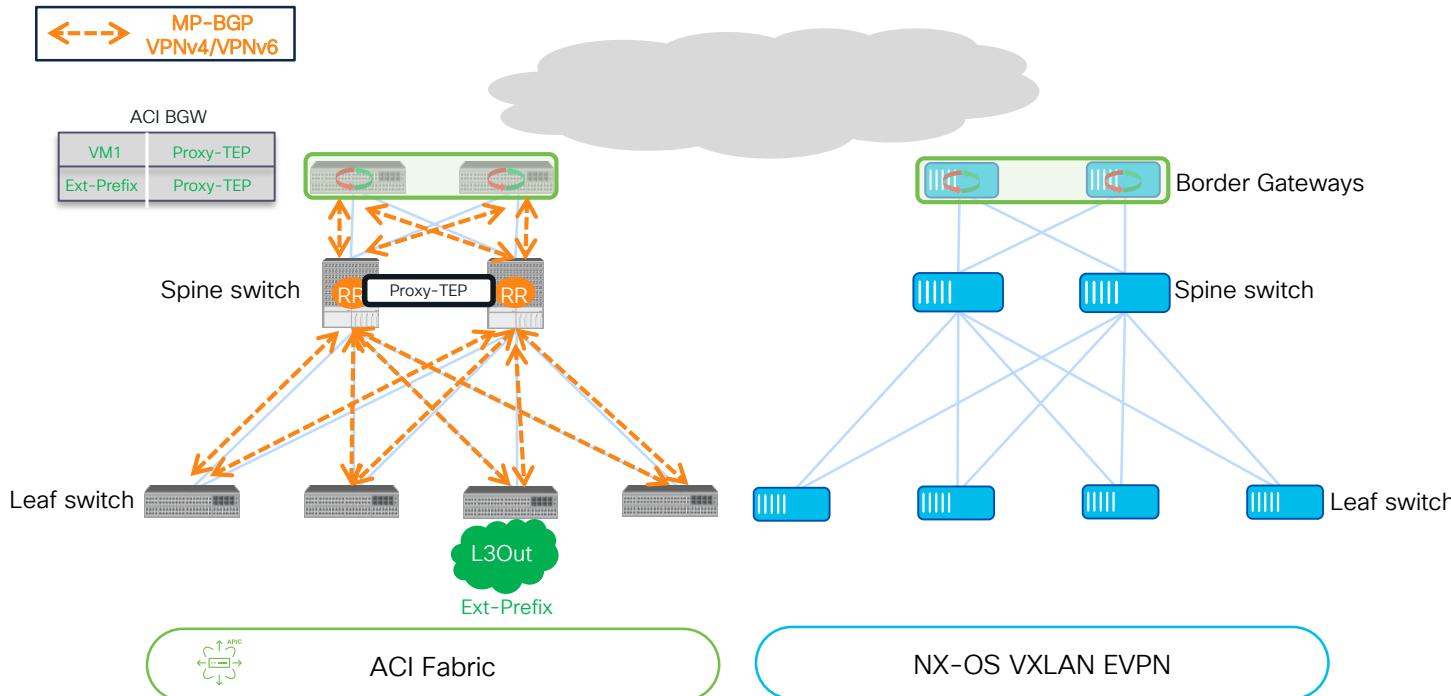
ACI Border Gateways

Control-Plane Overview



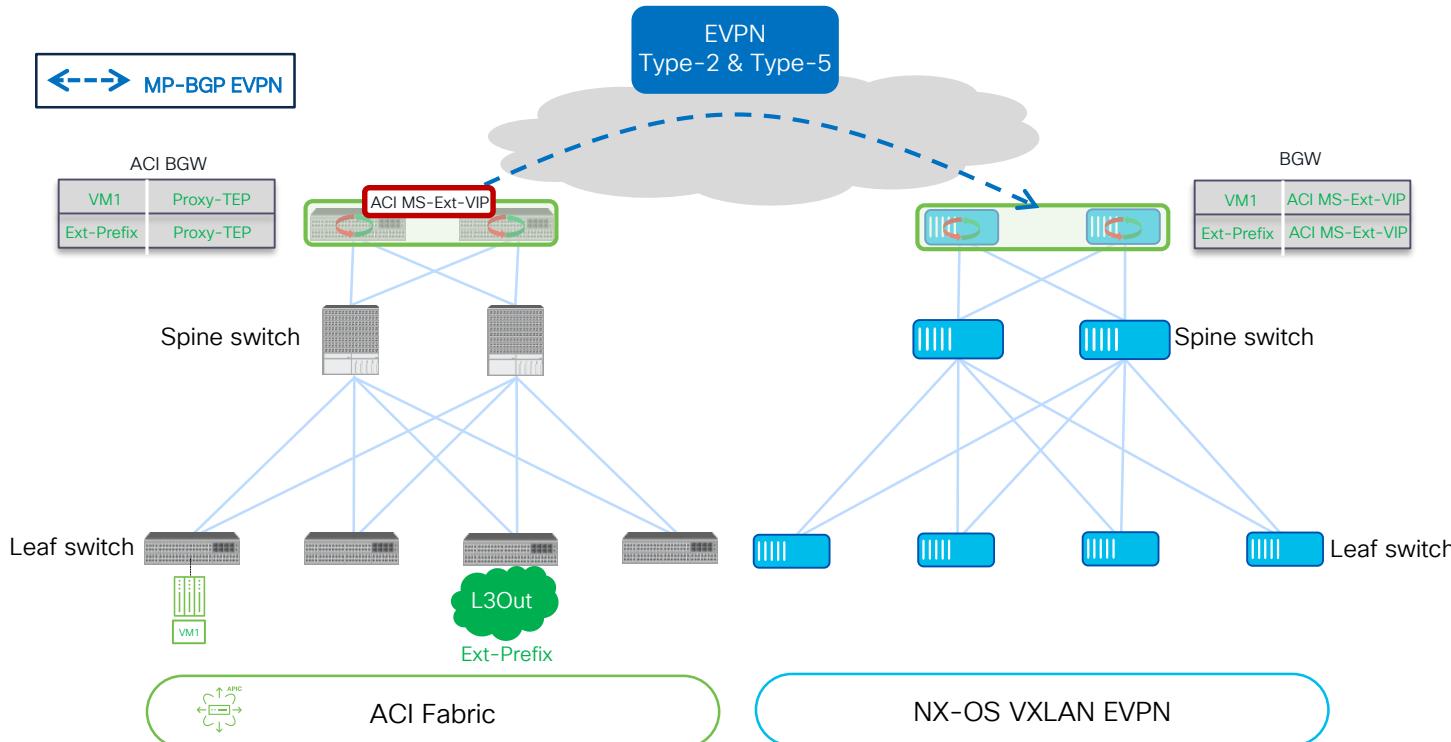
ACI Border Gateways

Control-Plane Overview



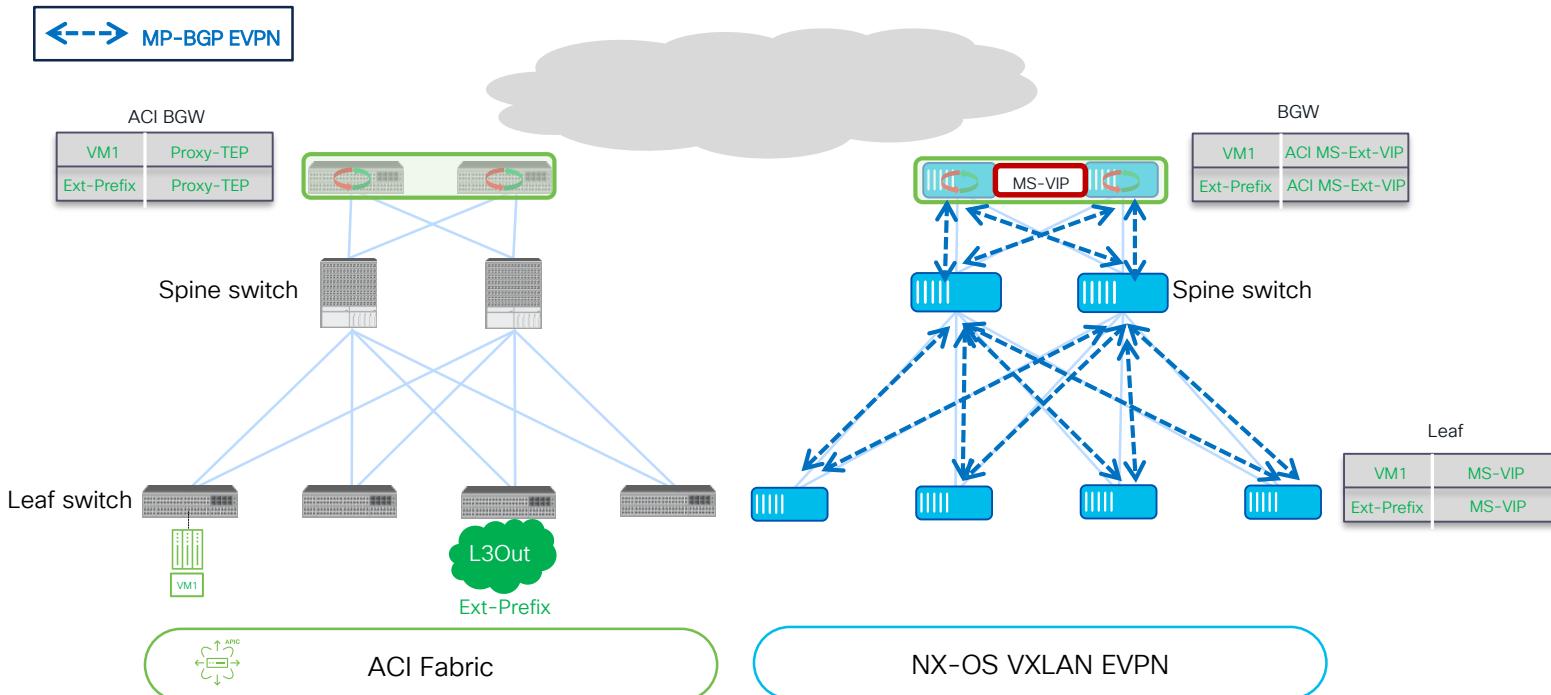
ACI Border Gateways

Control-Plane Overview



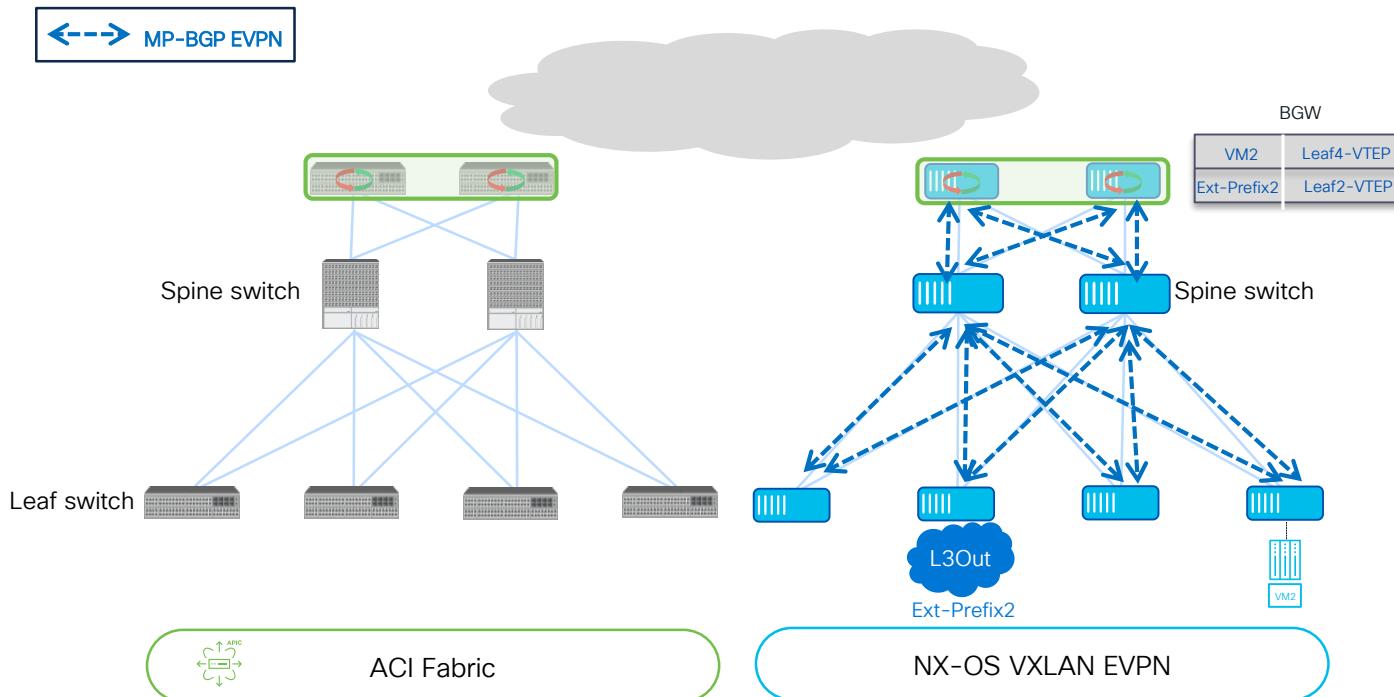
ACI Border Gateways

Control-Plane Overview



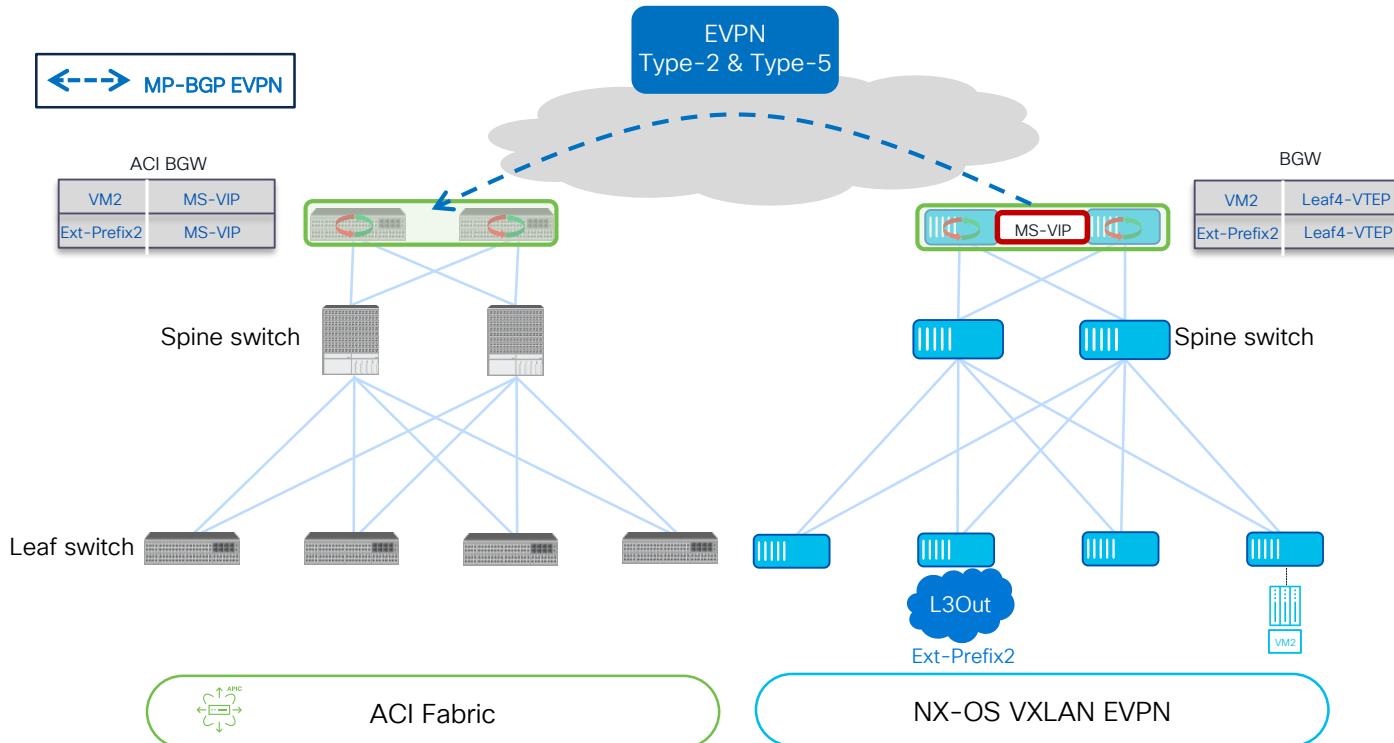
ACI Border Gateways

Control-Plane Overview



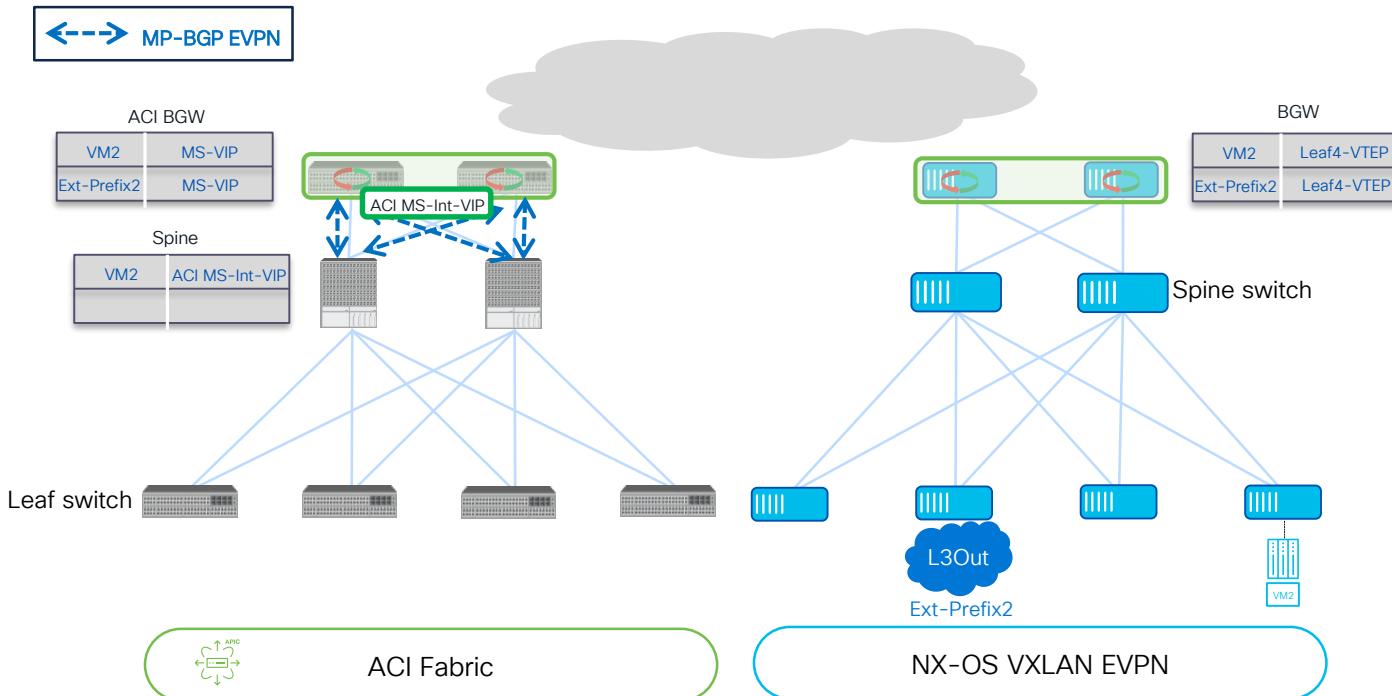
ACI Border Gateways

Control-Plane Overview



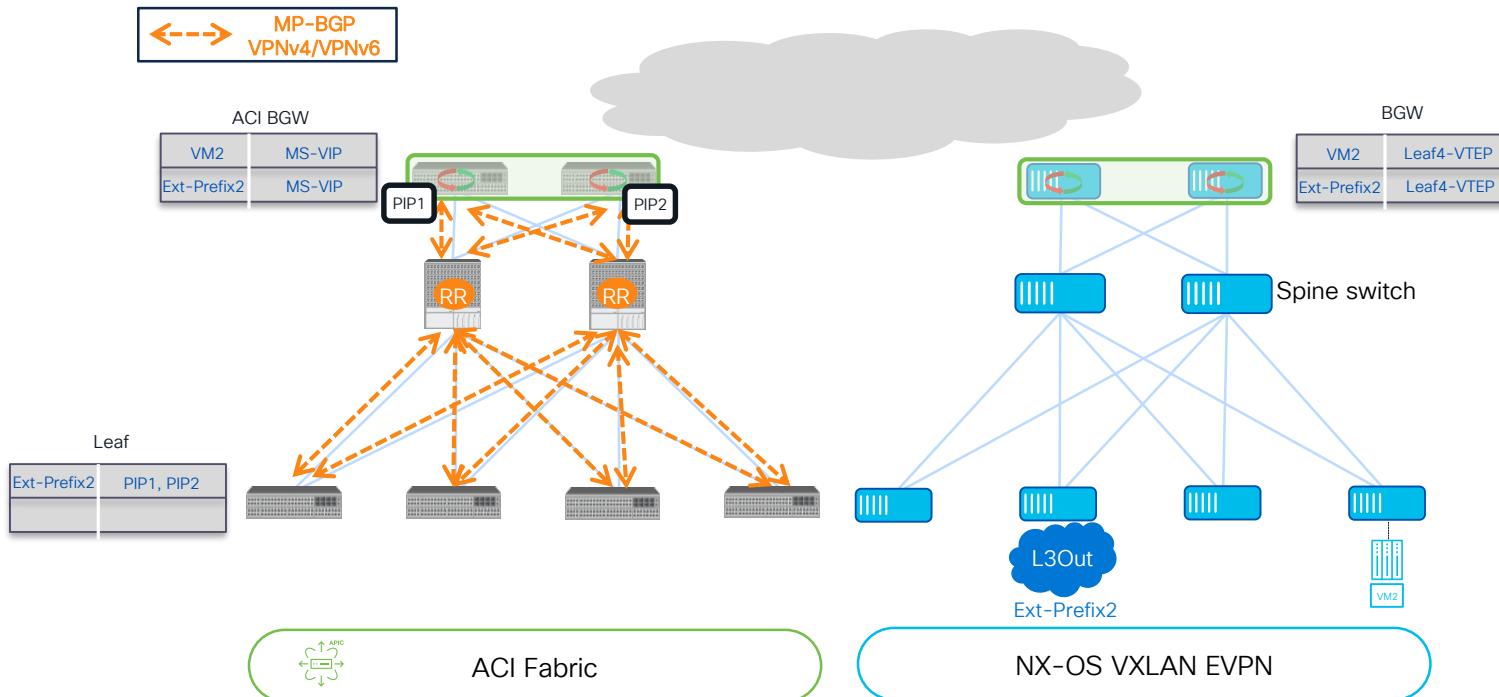
ACI Border Gateways

Control-Plane Overview



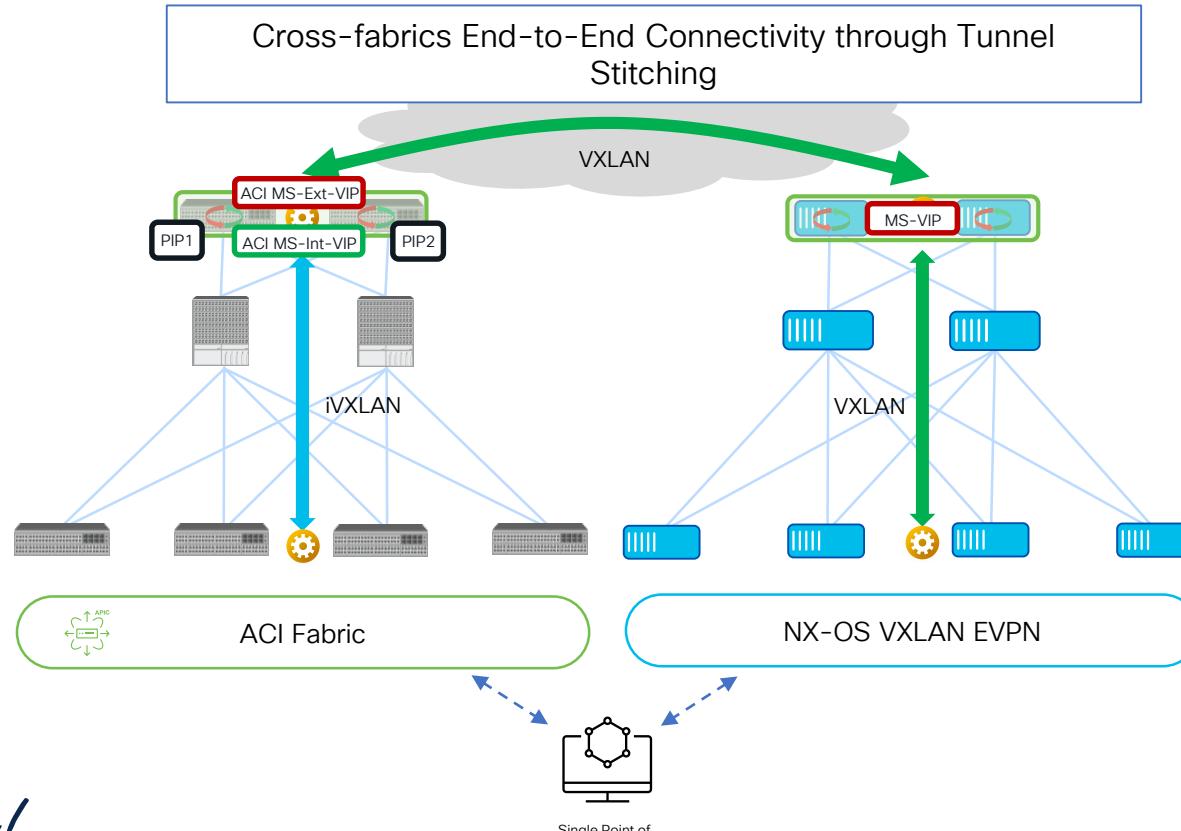
ACI Border Gateways

Control-Plane Overview



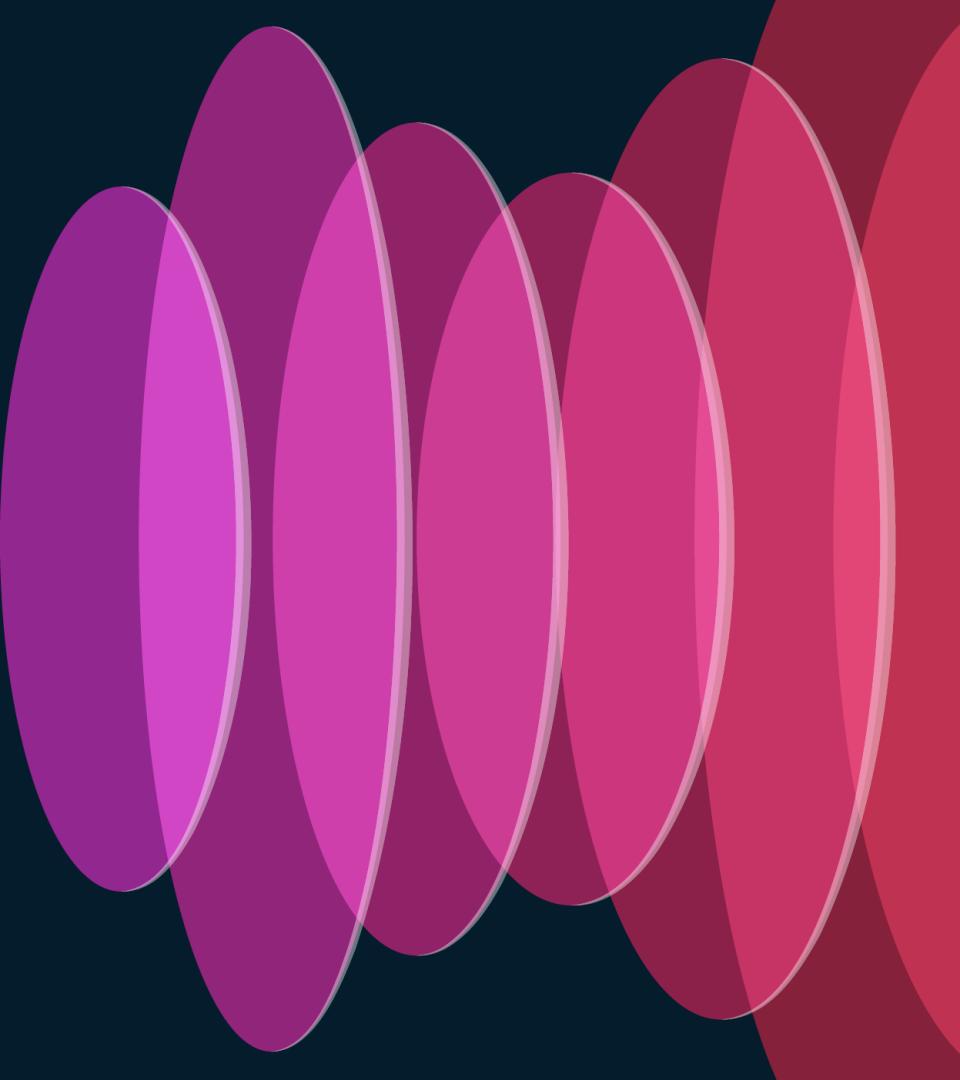
ACI Border Gateways

Data-Plane Overview



ACI Border Gateways

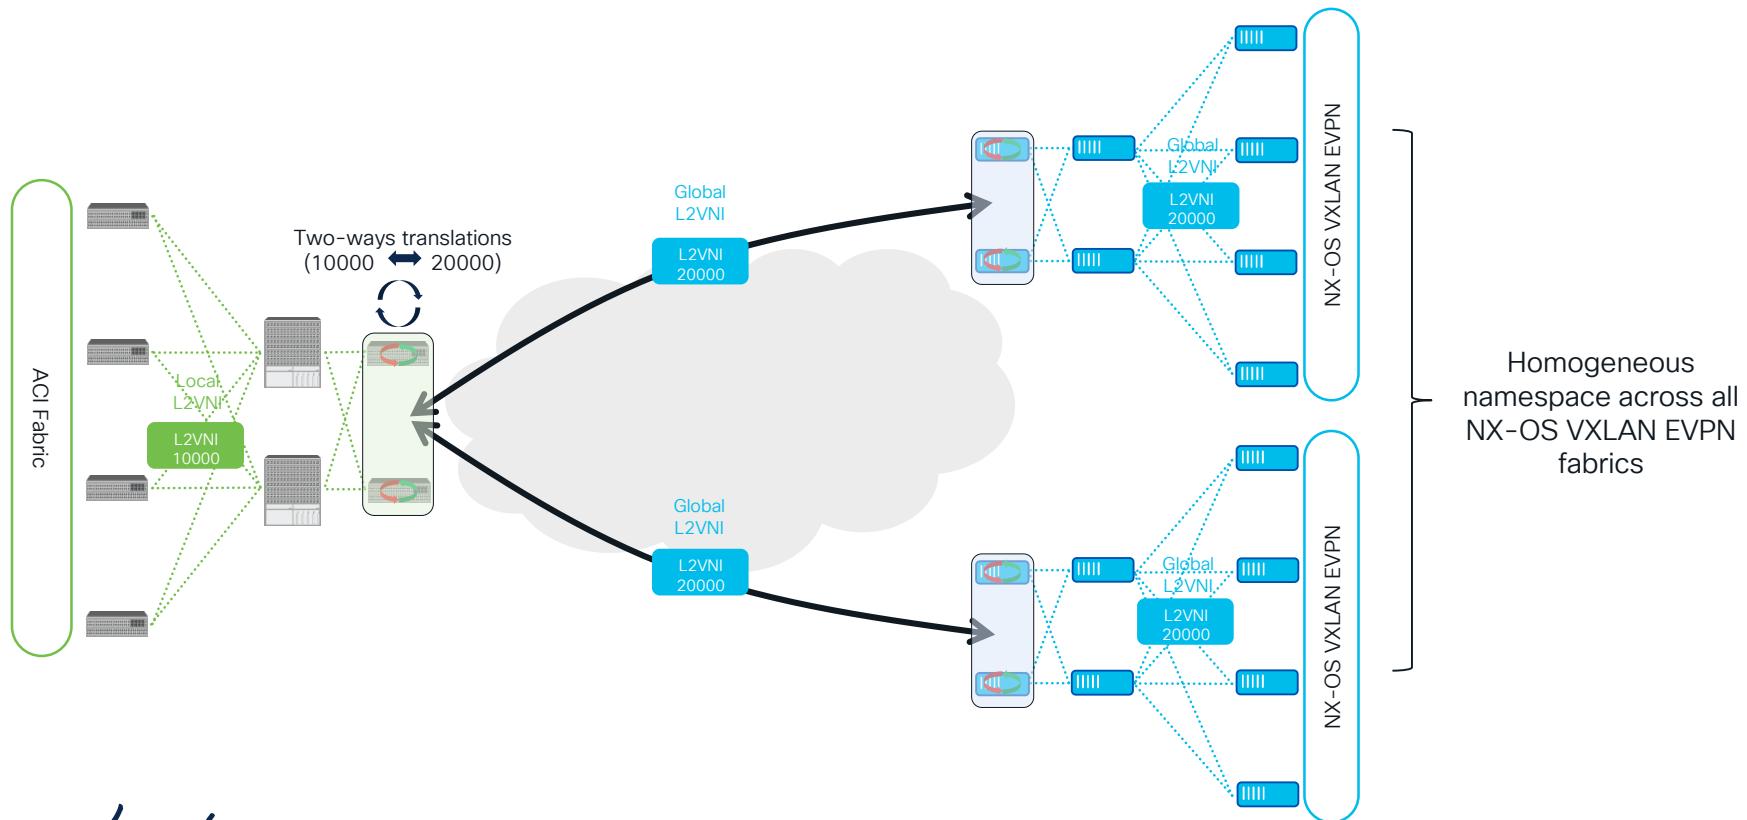
Namespace Normalization



ACI Border Gateways

Namespace Normalization for Stretched BDs

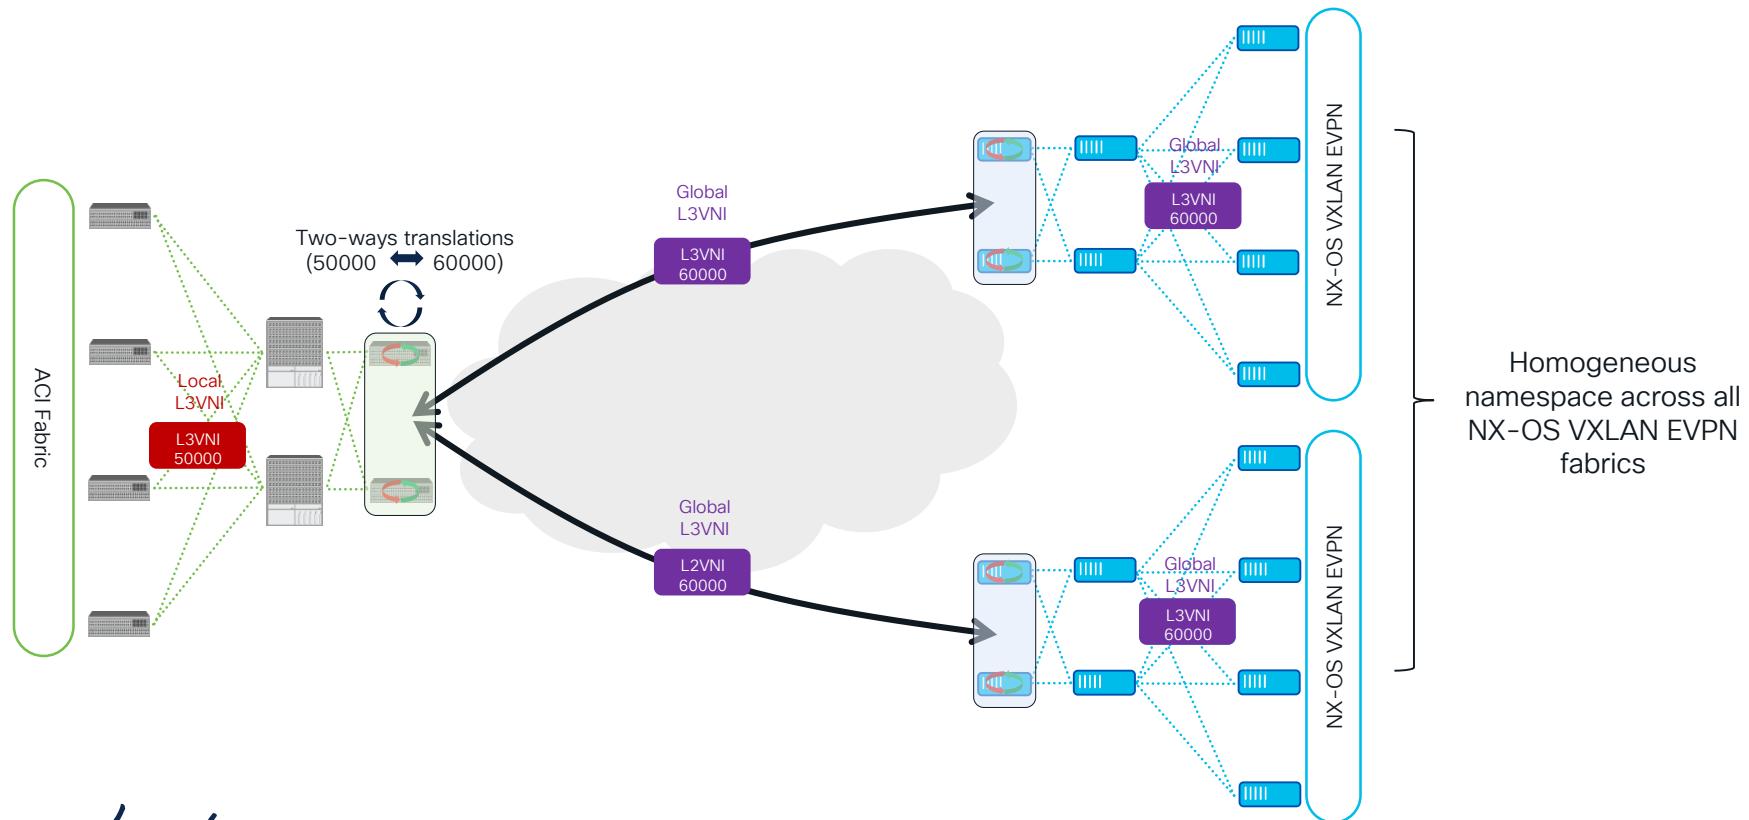
ACI 6.1(x)



ACI Border Gateways

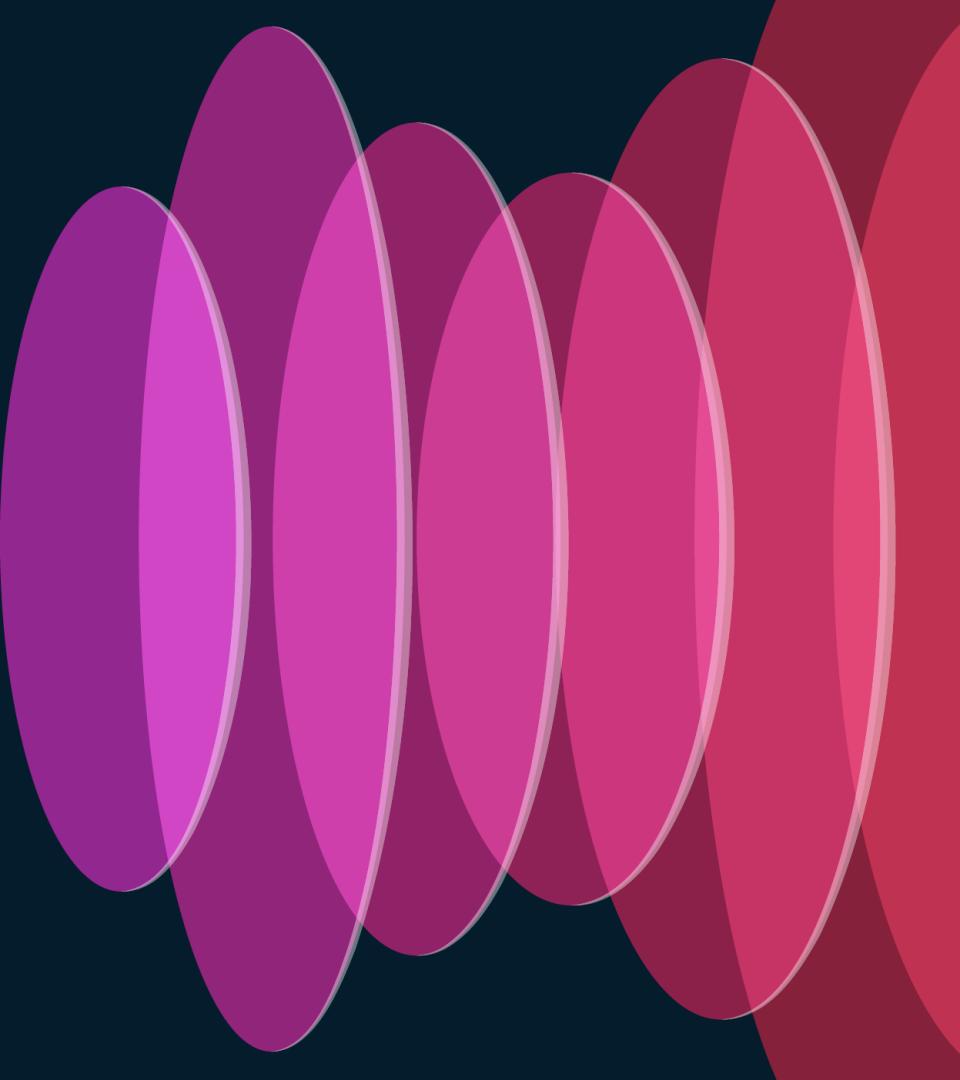
Namespace Normalization for Stretched VRFs

ACI 6.1(x)



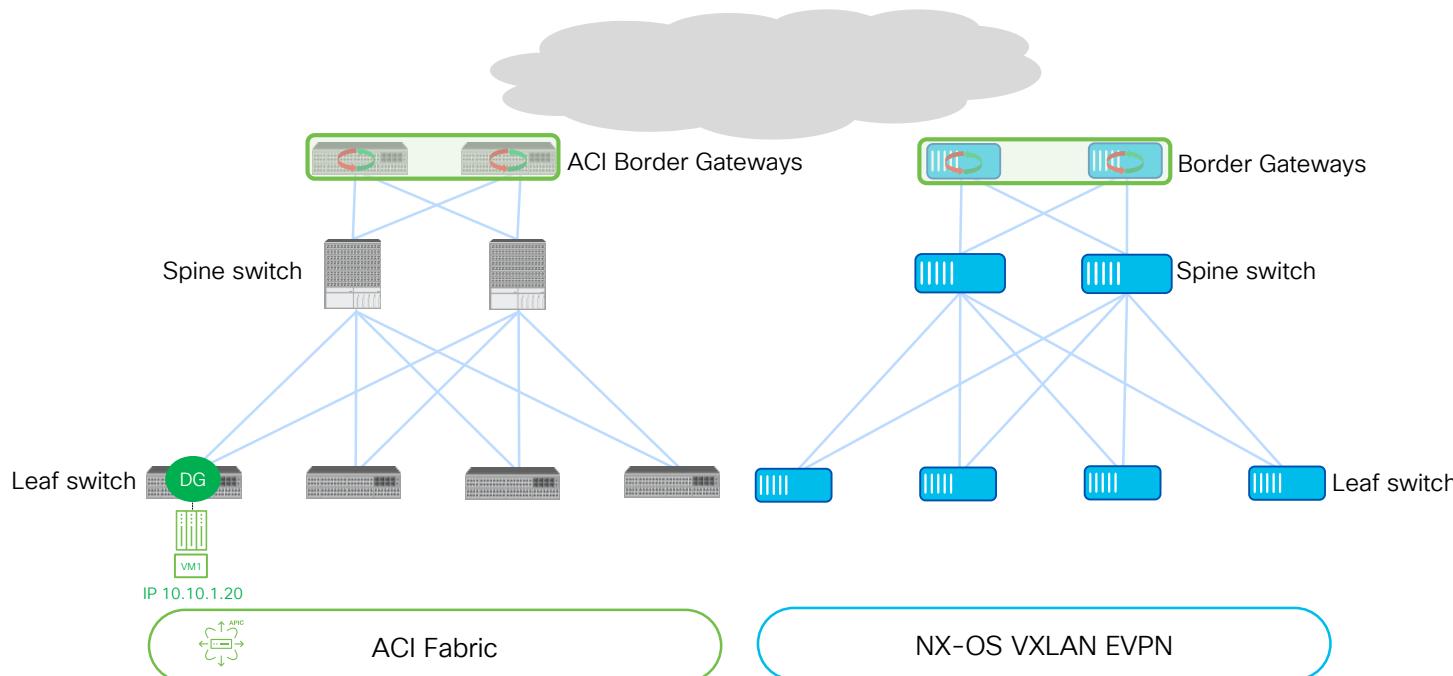
ACI Border Gateways

Workload Mobility across Domains



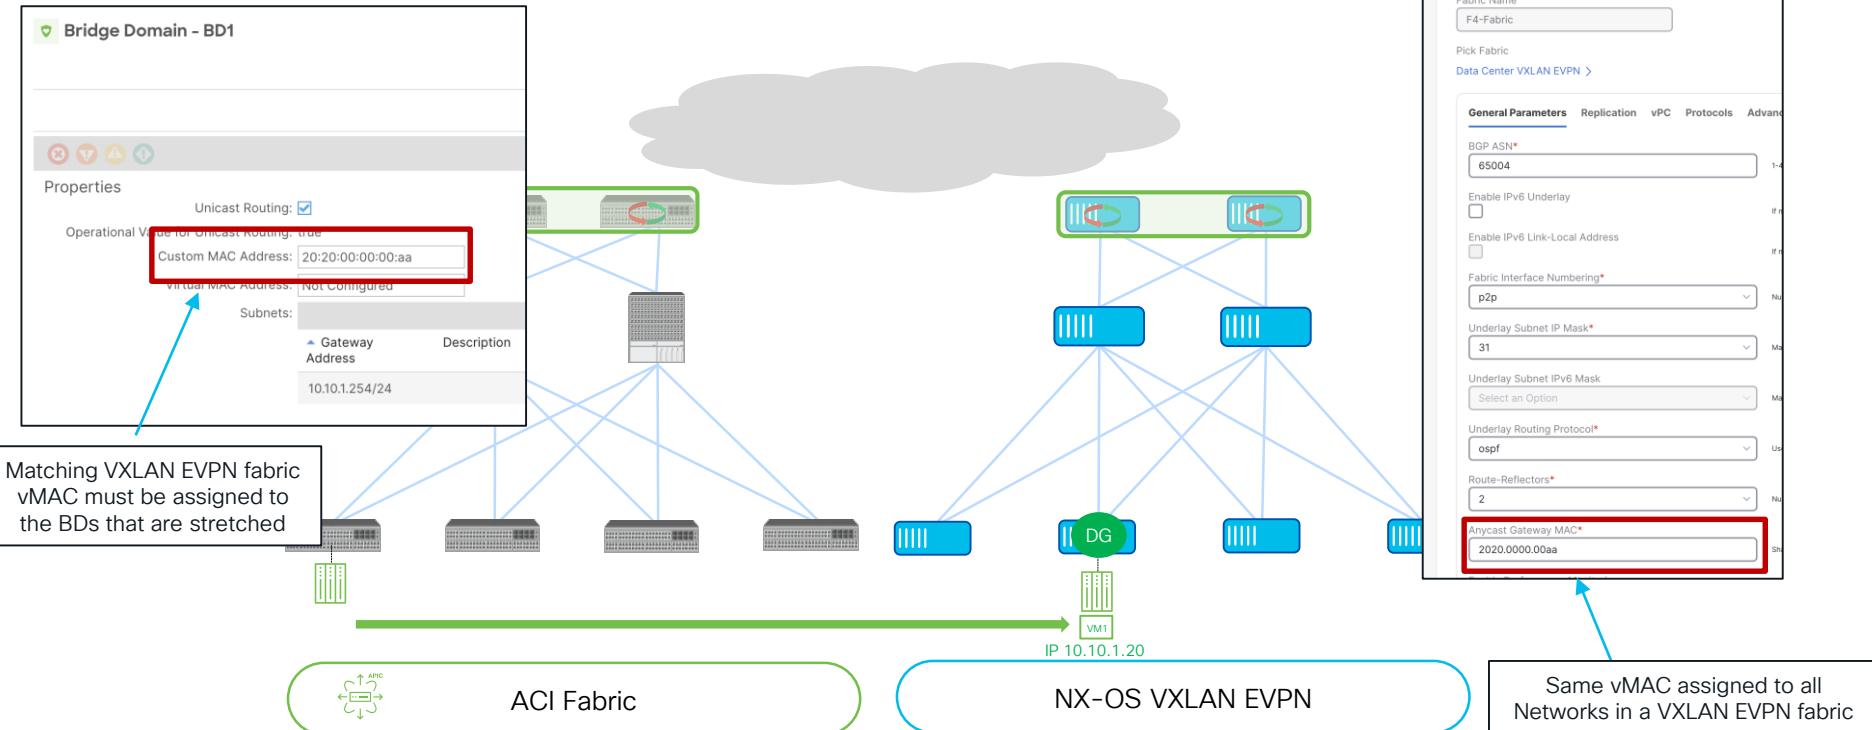
Workload Mobility

Configure a Consistent vMAC/VIP



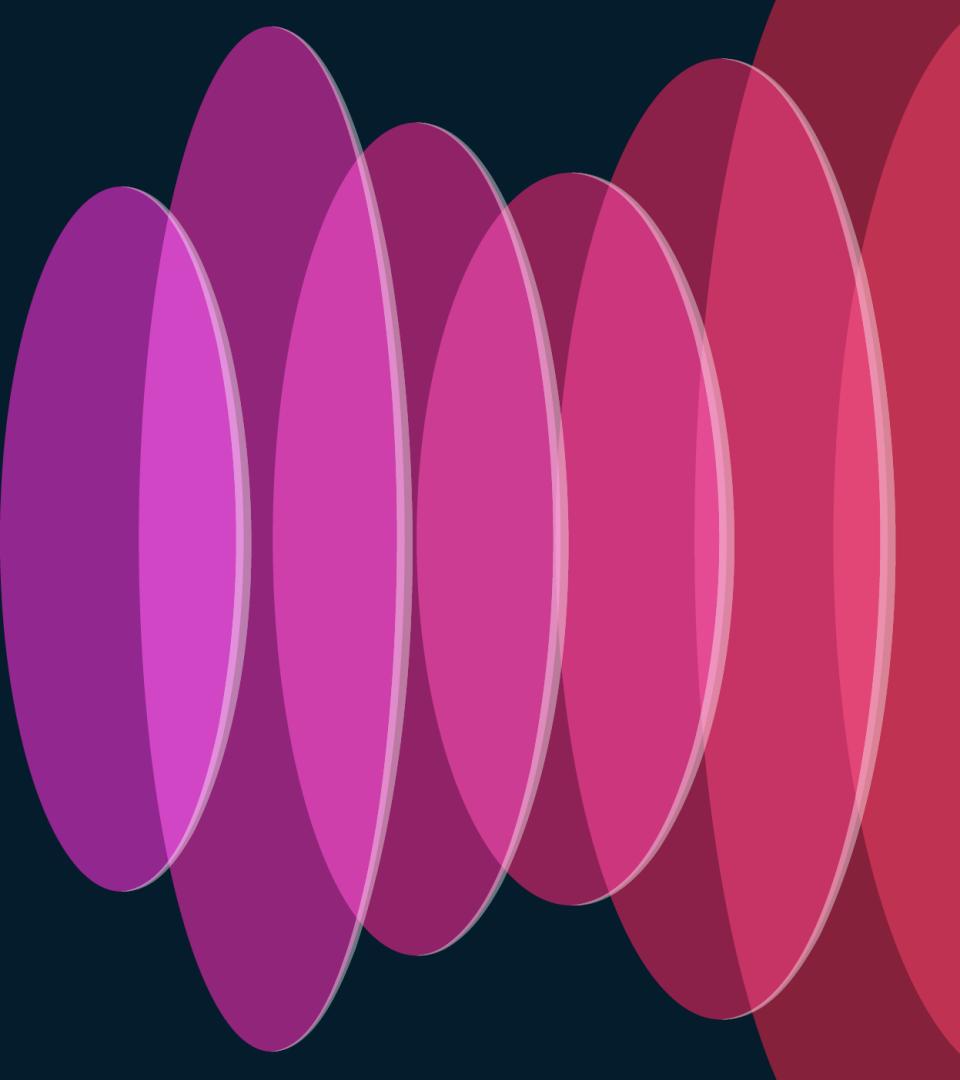
Workload Mobility

Configure a Consistent vMAC/VIP



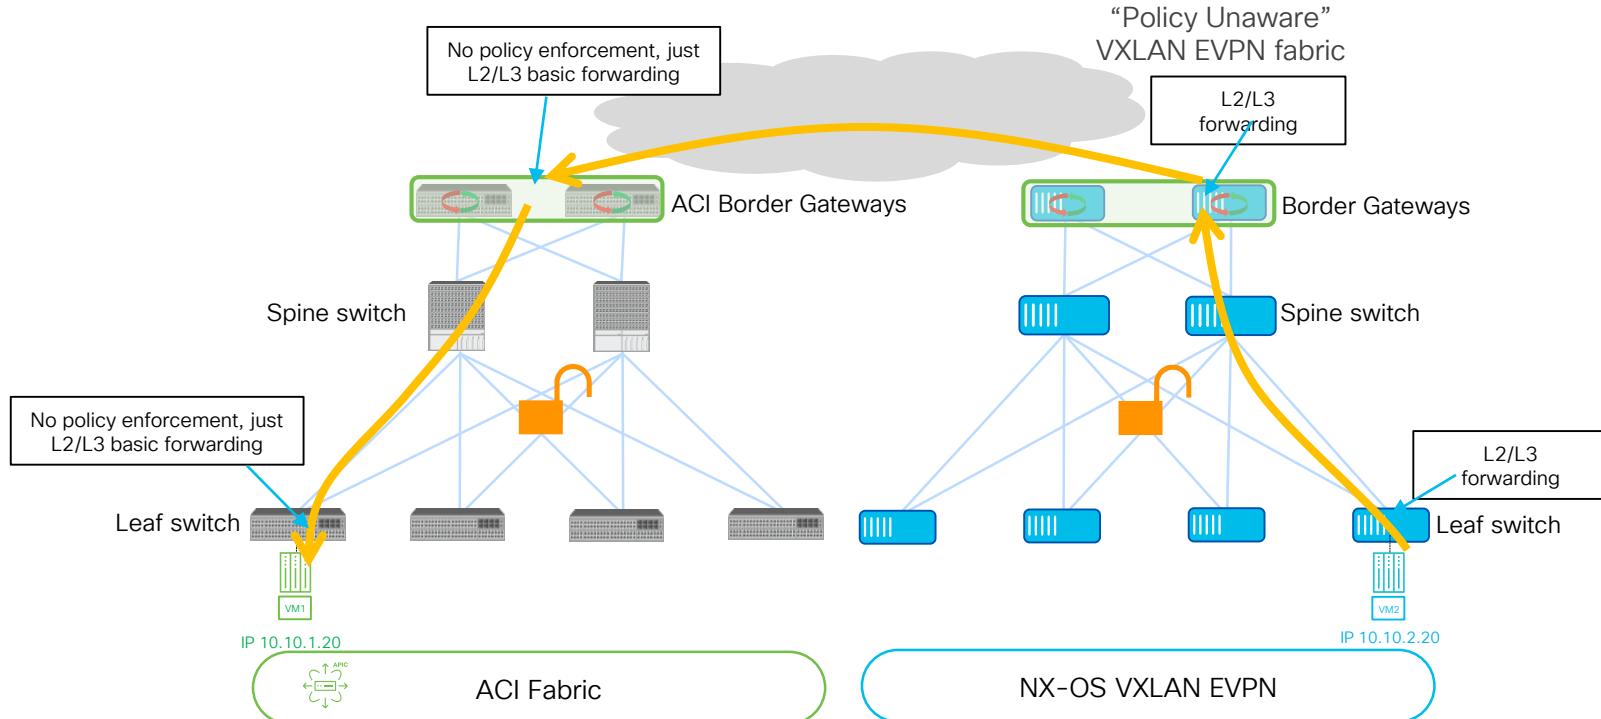
ACI Border Gateways

Policy Enforcement on ACI BGWs



Heterogeneous Fabrics

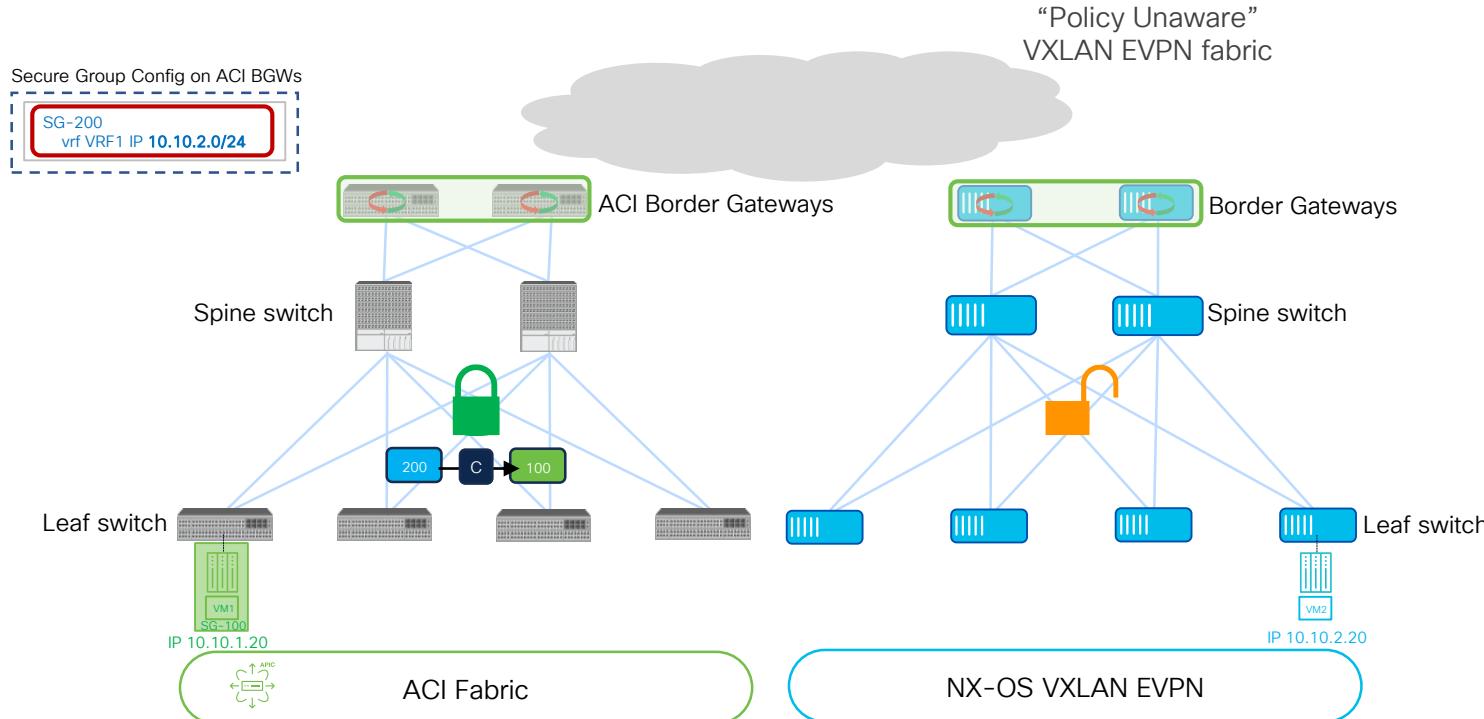
VRF Unenforced in ACI 6.1(1) Release



Heterogeneous Fabrics

Policy Enforcement on ACI BGWs

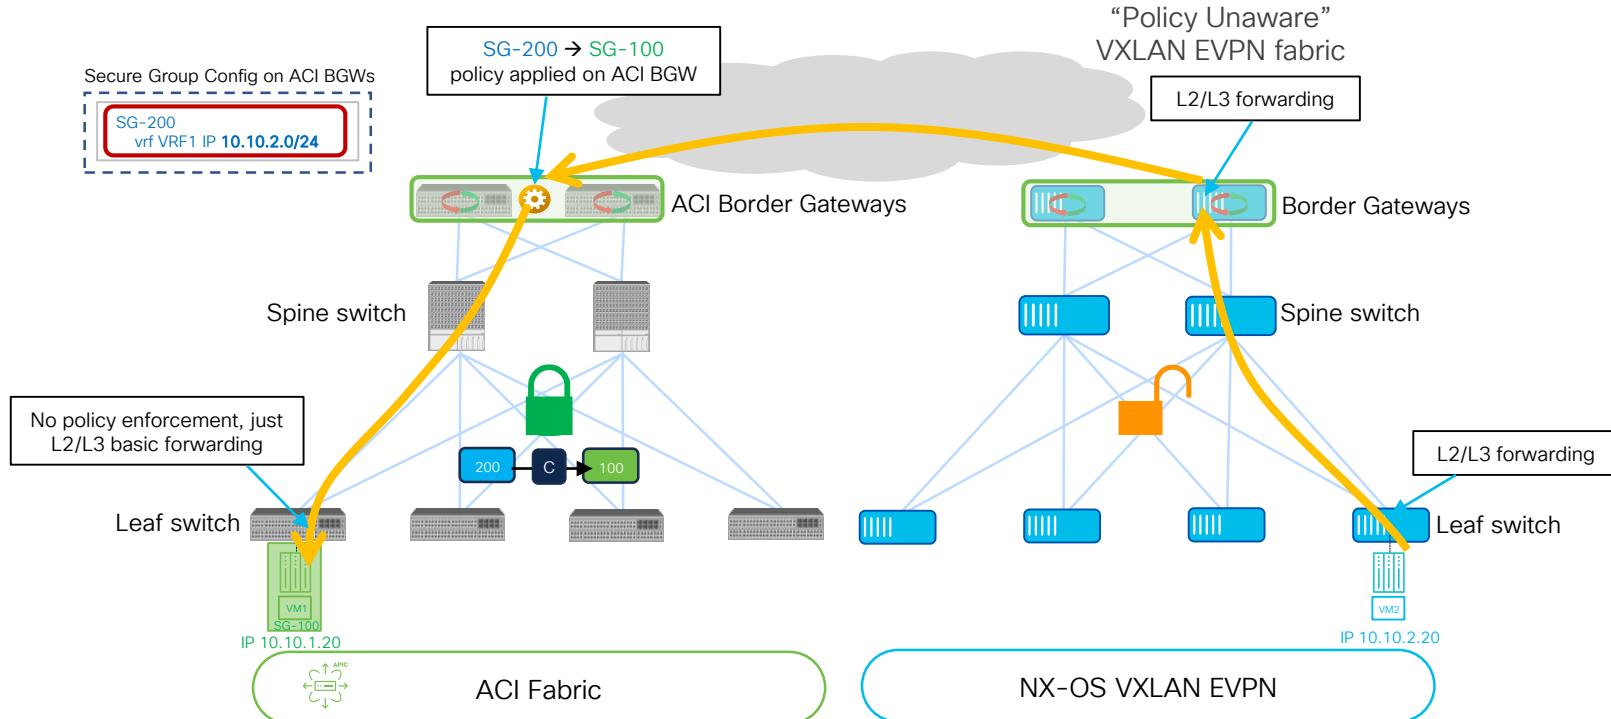
ACI 6.1(x)



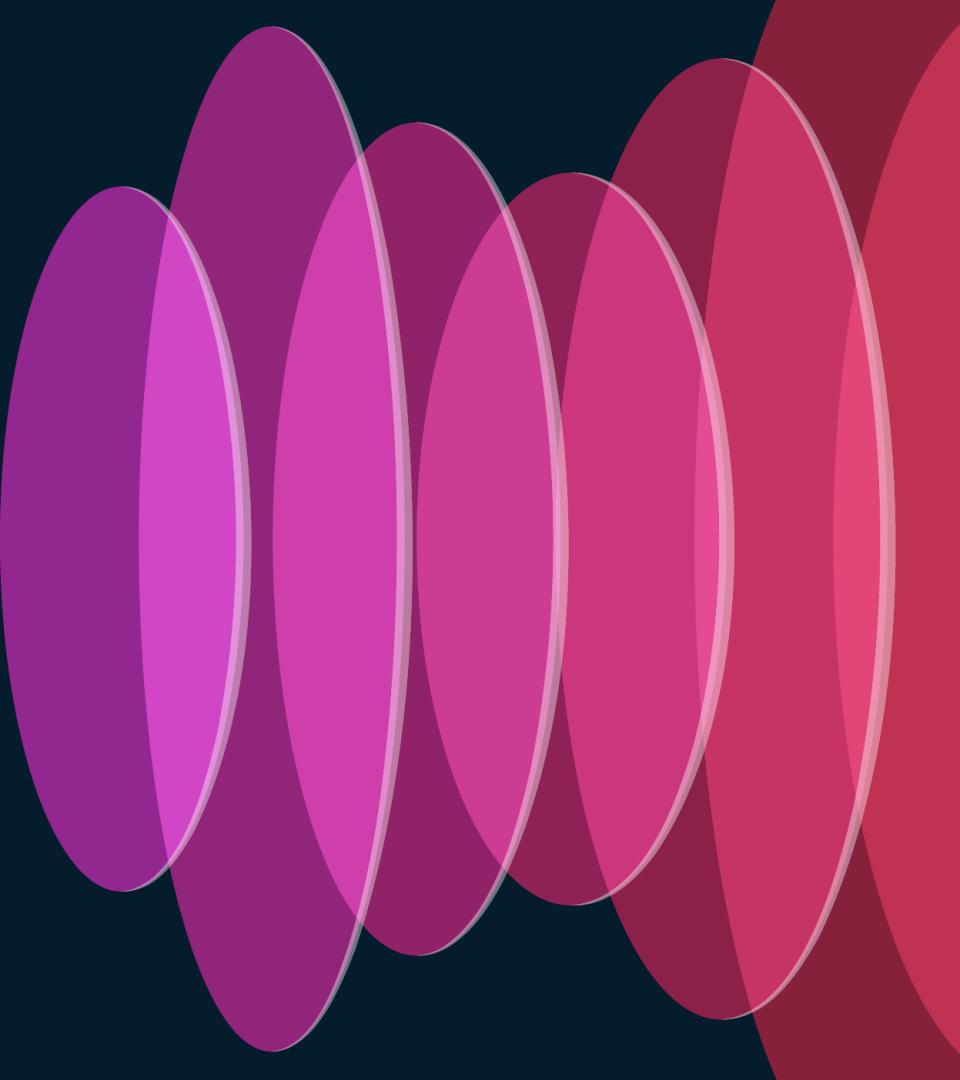
Heterogeneous Fabrics

Policy Enforcement on ACI BGWs

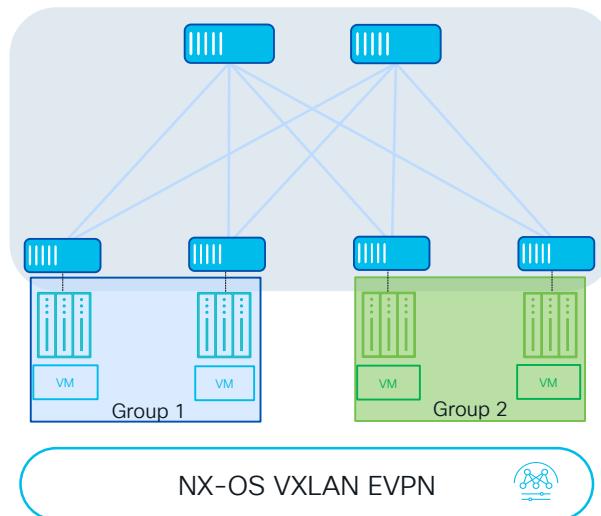
ACI 6.1(x)



Secure Interconnection of Heterogeneous Fabrics



VXLAN GPO with NX-OS



VXLAN GPO with NX-OS

- Group Policy Object carried in standard VXLAN header
- Decoupling network connectivity and security

Grouping

- Classify endpoints to create security groups
- Based on IP, VLAN, VM attributes, etc. across VRFs

Policy enforcement

- Create contracts/SGACLs between security groups
- Possible actions: permit, deny, redirect (service chaining)

Automation

- Automate using [NDFC](#) or [Open APIs](#)

Benefits

Segment East-West traffic

Flexible security isolation

Reduce attack surface

Automate your way

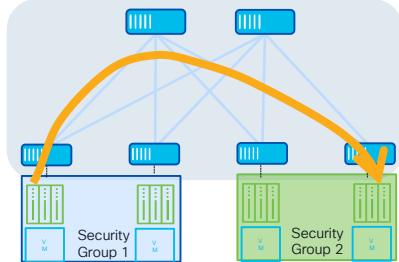
VXLAN GPO with NX-OS

Main Use Cases

For More Information on VXLAN GPO with NX-OS
BRKDCN-2629 & BRKDCN-2633

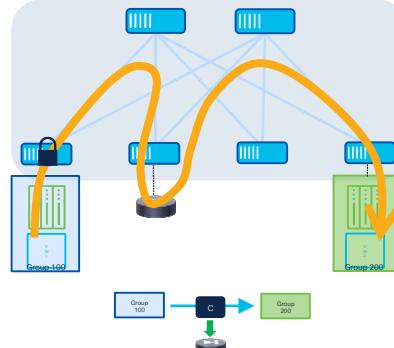
Creation of Security Zones

- VXLAN GPO allows to define policies for enforcing security policies (SGACLS) between security groups (SGs)
- SGACLS are a simpler, more flexible and more scalable policy enforcement mechanism compared to traditional ACLs
- Provides better control over the flow of network traffic (both east-west and north-south)



Service Chaining

- VXLAN GPO can be used to insert network services into a packet flow based on specific policy criteria
- Service chaining steers flows through the appropriate network services functions (such as firewalls, load balancers, or intrusion detection systems)



VXLAN GPO with NX-OS

Cisco GPO Data Plane and Control Plane Functionalities

Data Plane

(draft-smith-vxlan-group-policy)

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: April 25, 2019

M. Smith
Cisco Systems, Inc.
L. Kreeger
Arrcus, Inc.
October 22, 2018

VXLAN Group Policy Option
draft-smith-vxlan-group-policy-05

Abstract

This document defines a backward compatible extension to Virtual extensible Local Area Network (VXLAN) that allows a Tenant System Interface (TSI) Group Identifier to be carried for the purposes of policy enforcement.



Control Plane

(draft-wlin-bess-group-policy-id-extended-community)

bess
Internet-Draft
Intended status: Standards Track
Expires: 22 April 2024

W. Lin
Juniper Networks
J. Drake
Individual
D. Rao
Cisco Systems
20 October 2023

Group Policy ID BGP Extended Community
draft-wlin-bess-group-policy-id-extended-community-03

Abstract

Group Based Policy can be used to achieve micro or macro segmentation of user traffic. For Group Based Policy, a Group Policy ID, also known as Group Policy Tag, is used to represent a logical group that shares the same policy and access privilege. This specification defines a new BGP extended community that can be used to propagate Group Policy ID through a BGP route advertisement in the control plane. This is to facilitate policy enforcement at the ingress node when the optimization of network bandwidth is desired.

Data Plane and Control Plane

(draft-lrss-bess-evpn-group-policy)

BESS WorkGroup
Internet-Draft
Intended status: Standards Track
Expires: 5 September 2024

W. Lin
Juniper
D. Rao
A. Sajassi
M. Smith
Cisco
L. Kreeger
Arrcus
4 March 2024

EVPN Group Policy
draft-lrss-bess-evpn-group-policy-00

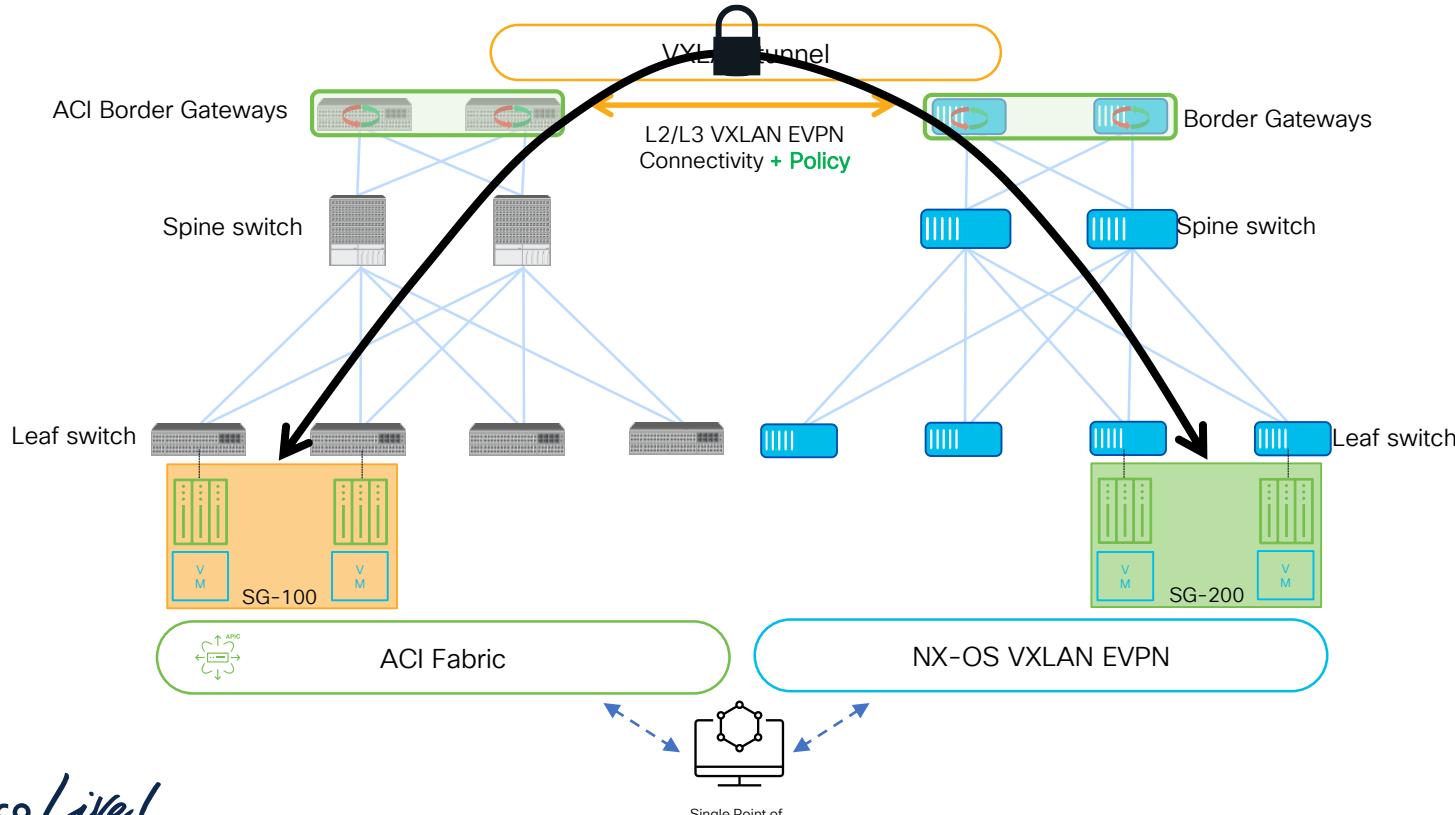
Abstract

Group Based Policy can be used to achieve micro or macro segmentation of user traffic. For Group Based Policy, a Group Policy ID, also known as Group Policy Tag, is used to represent a logical group that shares the same policy and access privilege. This document defines a backward compatible extension to Virtual extensible Local Area Network (VXLAN) that allows a Group Policy ID to be carried for the purposes of policy enforcement at the egress Network Virtualization Edge (NVE). It also defines a new BGP Extended Community that can be used to propagate Group Policy ID through a BGP route advertisement in the control plane. This is to facilitate policy enforcement at the ingress NVE when feasible.

Heterogeneous Fabrics

Policy Enforcement End-to-End

Future



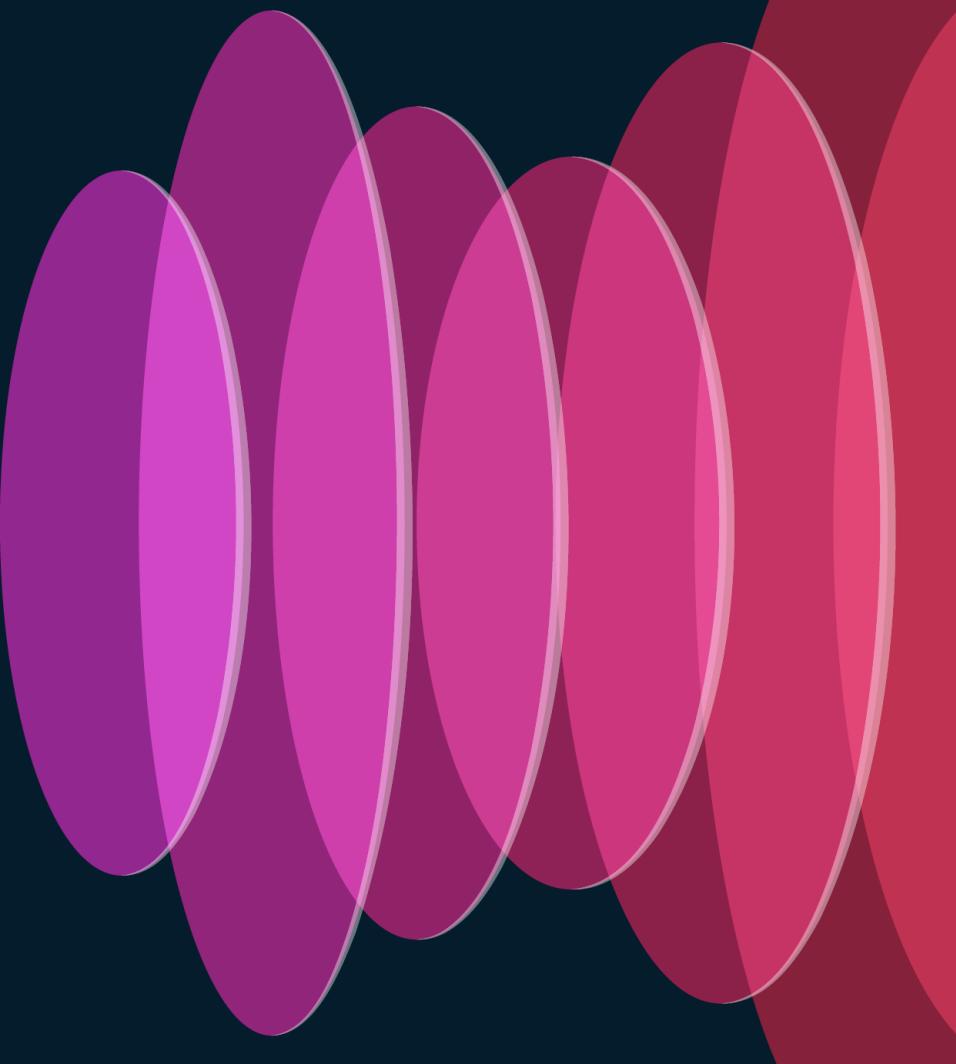
CISCO Live!

BRKDCN-2634

© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Public

48

Conclusions



Conclusions

- Building distributed infrastructures is key to the deployment of resilient and scalable designs
- Cisco One Fabric Experience aims to seamlessly interconnect and operate a mix of heterogeneous fabrics (ACI and VXLAN EVPN)
- The three main pillars to realize the One Fabric Experience vision are:
 1. BGW function for ACI fabrics
 2. Security policies in VXLAN EVPN fabrics (GPO)
 3. Introduction of centralized management and operation platforms for heterogeneous fabric on Nexus Dashboard

Complete Your Session Evaluations



Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to **win 1 of 5 full conference passes** to Cisco Live 2025.



Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.



Level up and earn **exclusive prizes!**



Complete your surveys in the **Cisco Live mobile app**.



Continue your education

CISCO *Live!*

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand



The bridge to possible

Thank you

CISCO *Live!*

#CiscoLive