Nexus ONE: An Architecture for Your Cisco Datacenter of Tomorrow

Max Ardica
Distinguished TME

CISCO Live !

BRKDCN-2602

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 13, 2025.

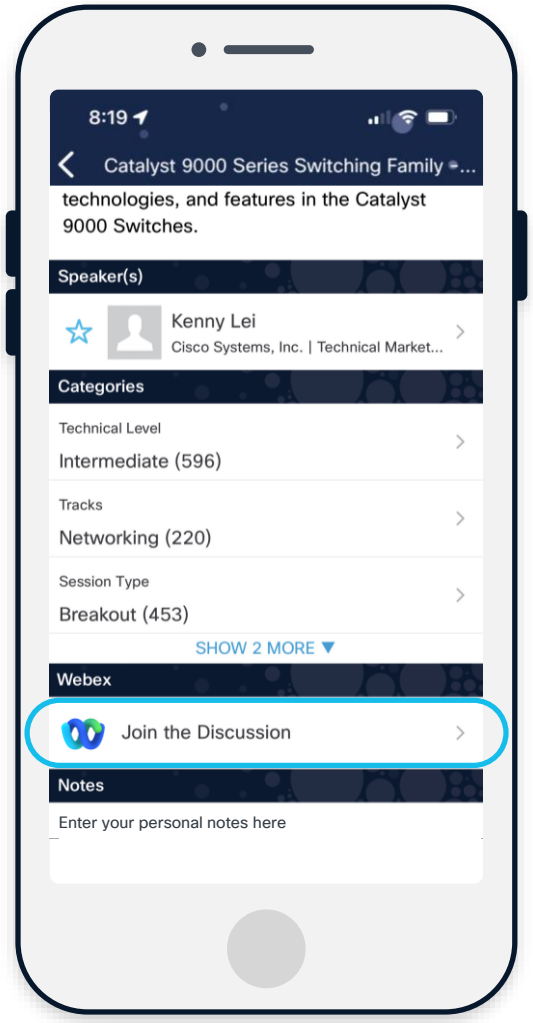

Agenda

01 DC Design Evolution

From a Single Fabric to a Distributed Architecture

02 Cisco Distributed DC Architectures

03 Nexus ONE Architecture

Group Policy Option GPO)

ACI BGWs

Unified Nexus Dashboard

DC Design Evolution

Leaf and Spine Topology

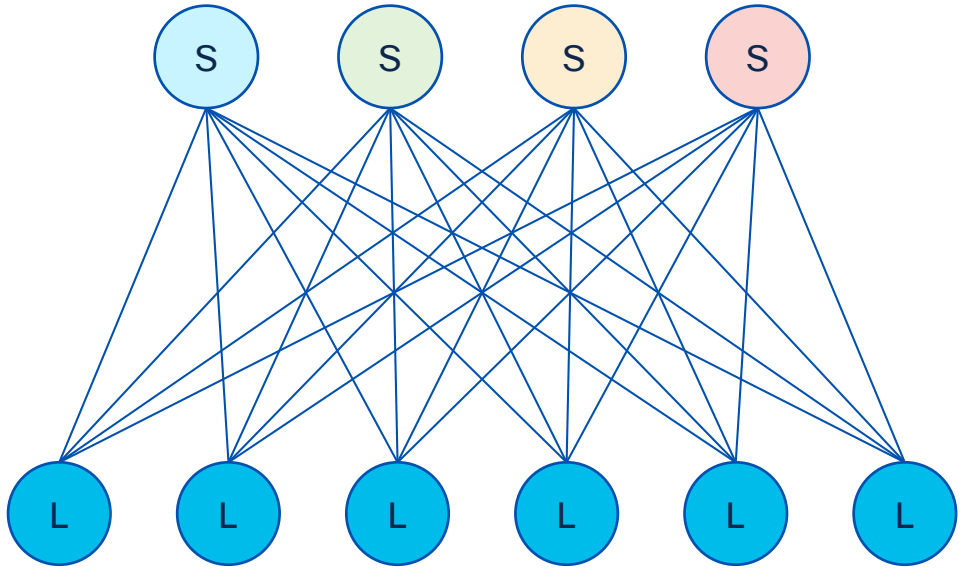

- A Leaf and Spine Topology
- Variations or Names of the same:
  - Fat Tree
  - Folded Clos
  - 3 Stage Clos
  - 2 Tier Network

Leaf and Spine Topology

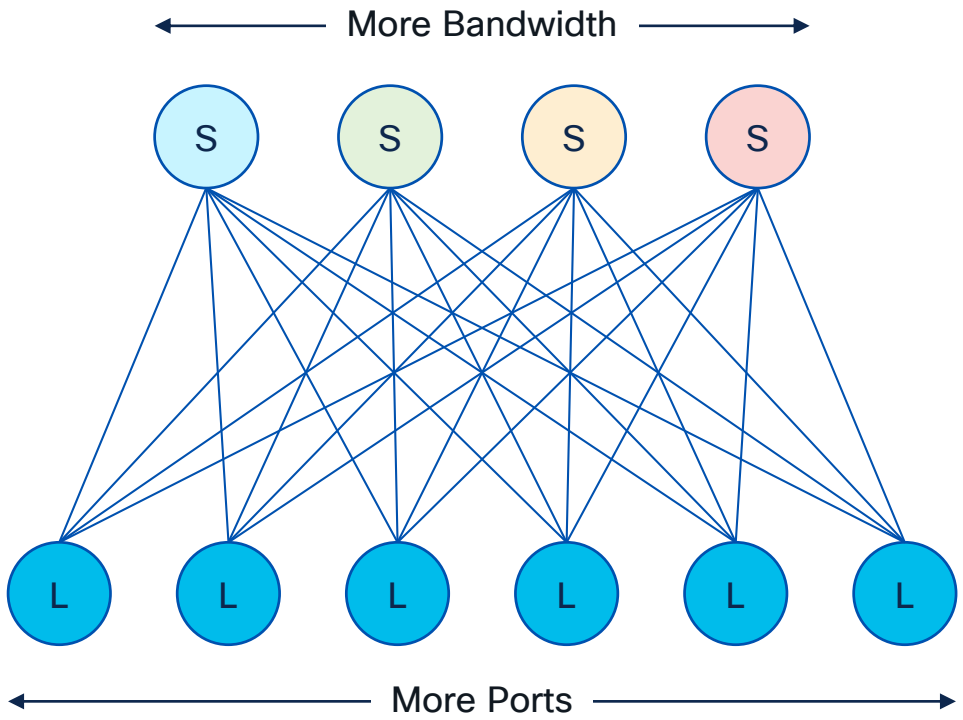

- A Scale Out Architecture
  - More Leaf = More Ports
  - More Spine = More Bandwidth

Leaf and Spine Topology

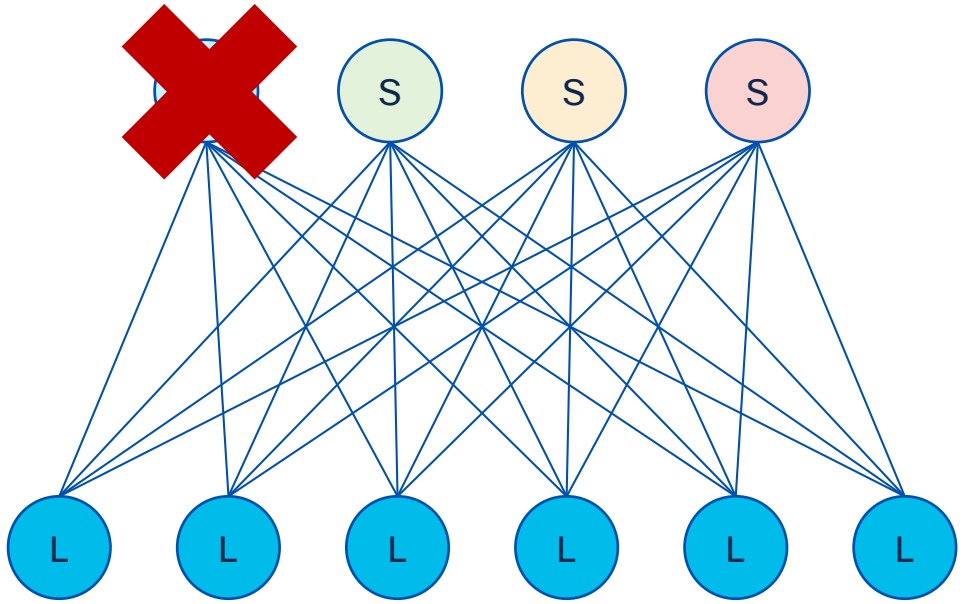

- N+1 Redundancy
- Redundancy increases by Building out the Topology
- On Spine failure
  - 4 Spine = 25% impact
  - 8 Spine = 12.5% impact

Leaf and Spine Topology

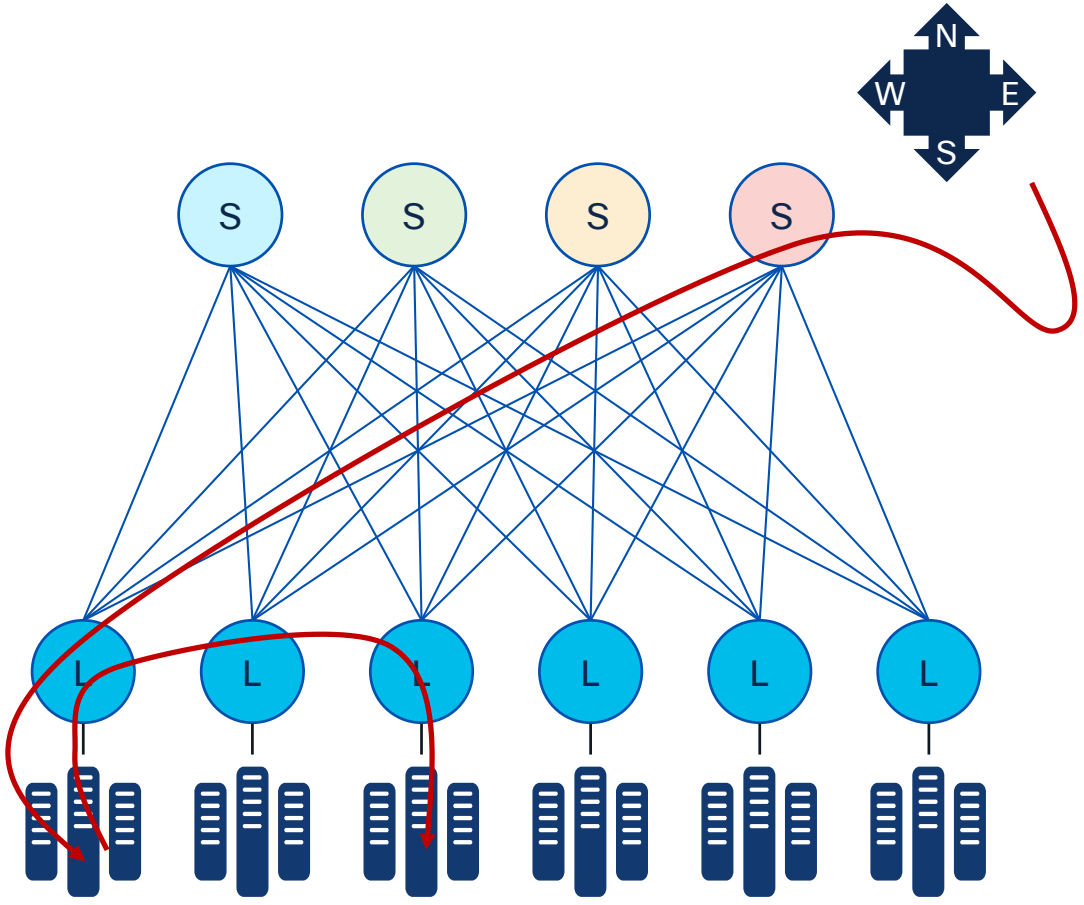

- Modern Application Needs
  - Every (1) North to South Connection, requires eight (8) East to West
  - User Access the Frontend (Web)
  - Frontend connects to App, DB, Storage etc.

Leaf and Spine Topology

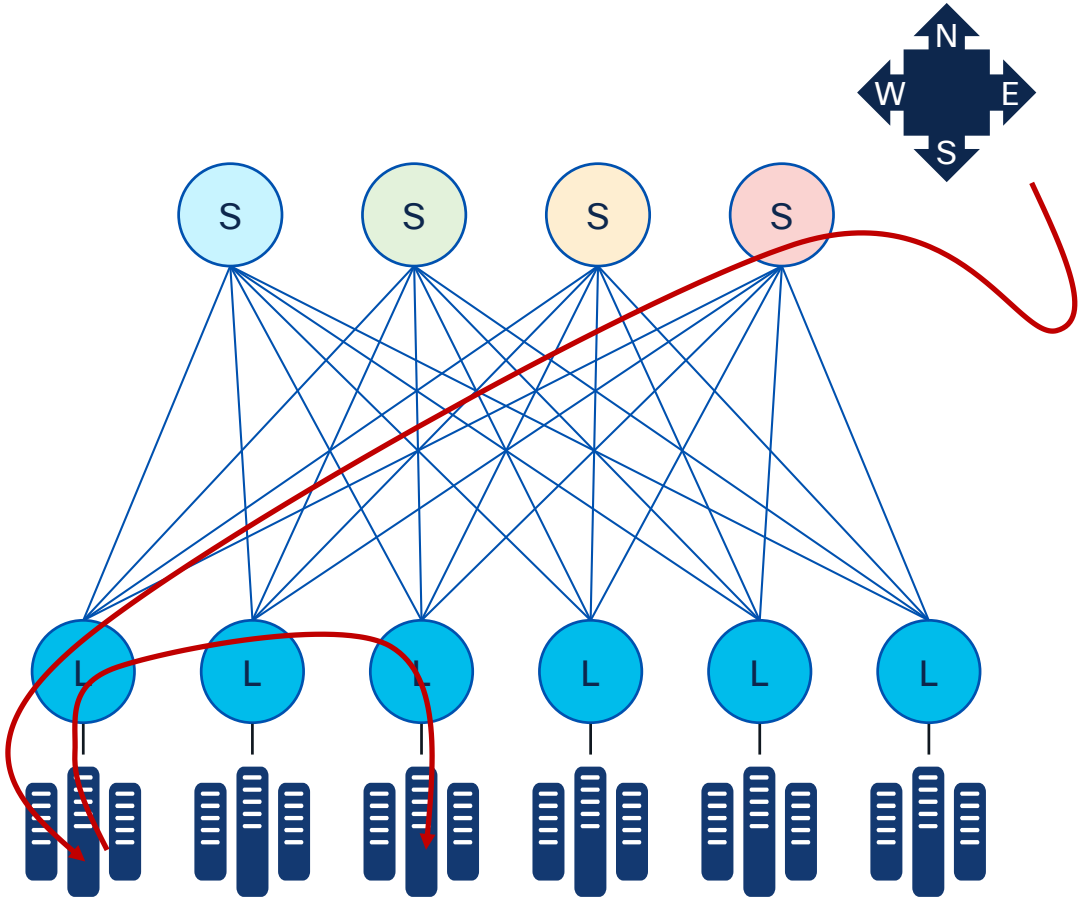

- Optimized for East to West
  - Consistent Latency from Leaf to Leaf
  - Wide ECMP
- Flexibility for North to South
  - External Connectivity at Leaf or Spine layer

Data Center Design Evolution

Various Considerations to Reduce the Fabric's Size

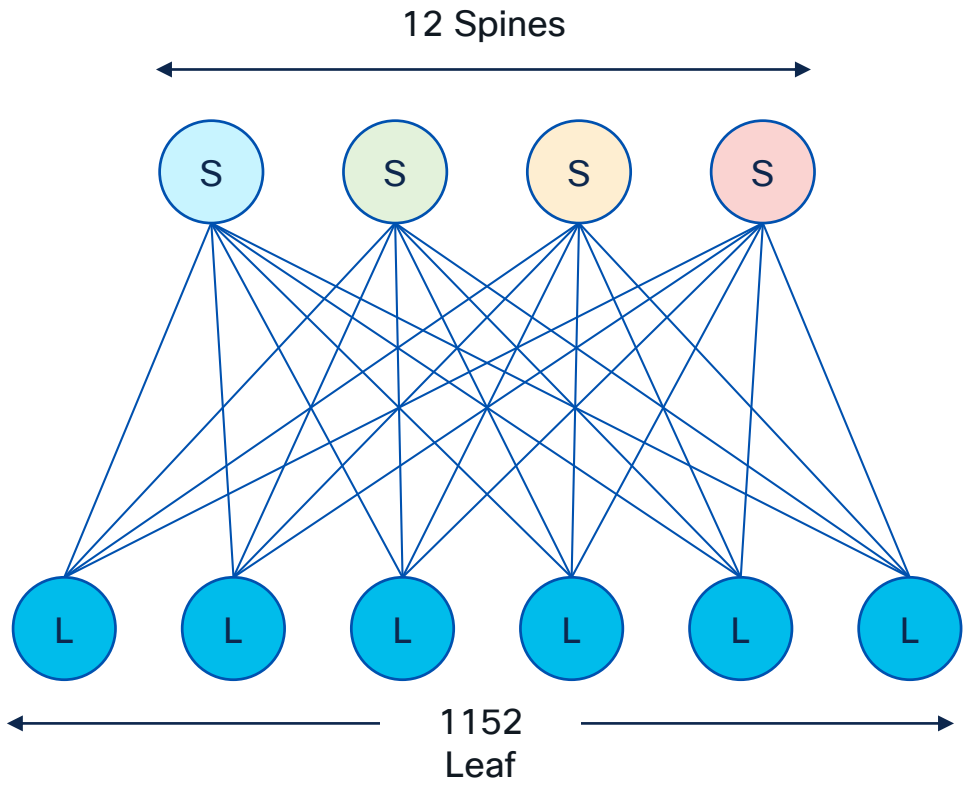

- What is my Failure Domain?
- What is my Change Domain?
- What is my Overall Scale?
- What is my Fabric Solution Scale?
- What is my Fabric SLA?
- What is my Maximum Downtime?

What Have we Learned from the Cloud Titans?

Building Scalable Data Center Networks

#1

Simplicity is Key

Simple Design Principals

#2

Scale as you Go

Scale is Never Finite

#3

Fail but Fail Fast

Reduce Brown-Out Exposure

#4

Redundant and Repeatable

Risk is Never an Option

Design Evolution

Various Considerations to Reduce the Fabric's Size

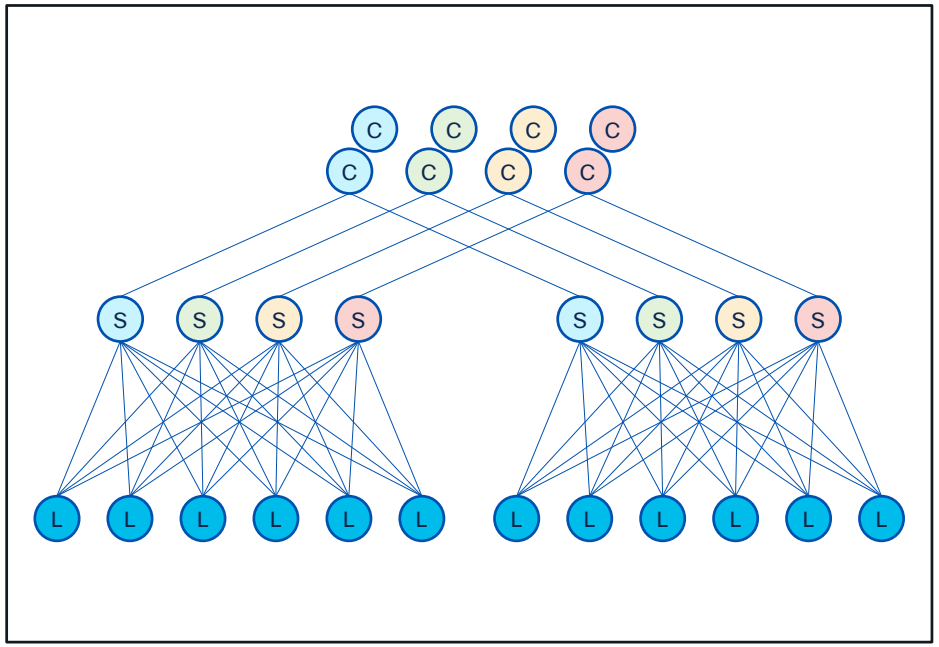

3 Tier Leaf-Spine-Core (or Super-Spine)

- Increasing Scale-Out in all Tiers
- Reducing Complexity
- Simple Design Principles
- Increasing the "Finite Scale"
- Scale as You Go
- Disaggregated Redundancy
- Flexible Link and Bandwidth Distribution
- Further Possibility for Cost Optimization

Design Evolution

From a Single Large Fabric to a Distributed Architecture

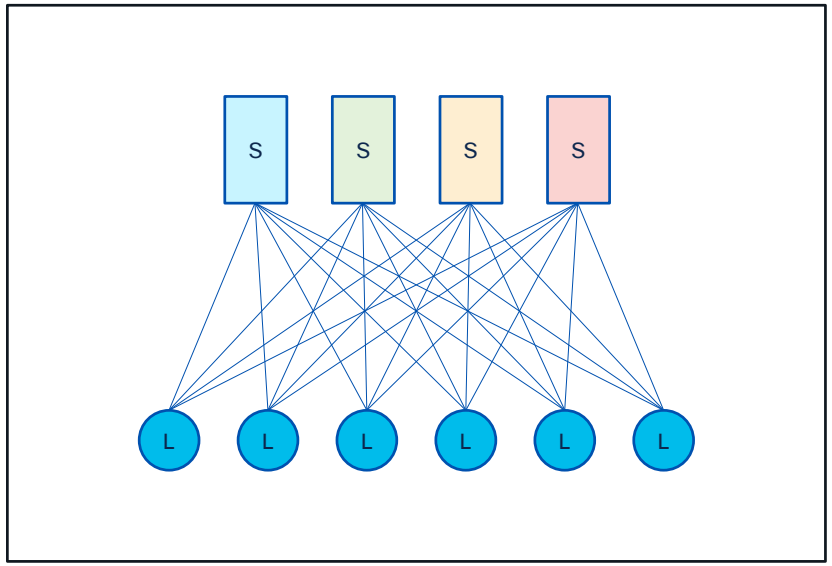

2 Tier Leaf Spine

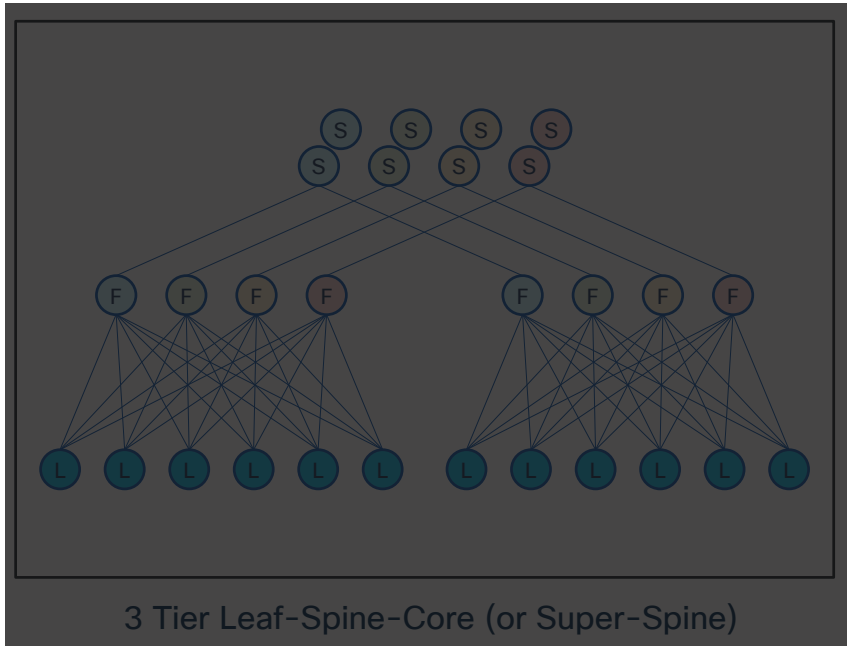

3 Tier Leaf-Spine-Core (or Super-Spine)

Design Evolution

From a Single Large Fabric to a Distributed Architecture

For More Information on Multi-Tier Fabric Deployments please refer to BRKDCN-2999

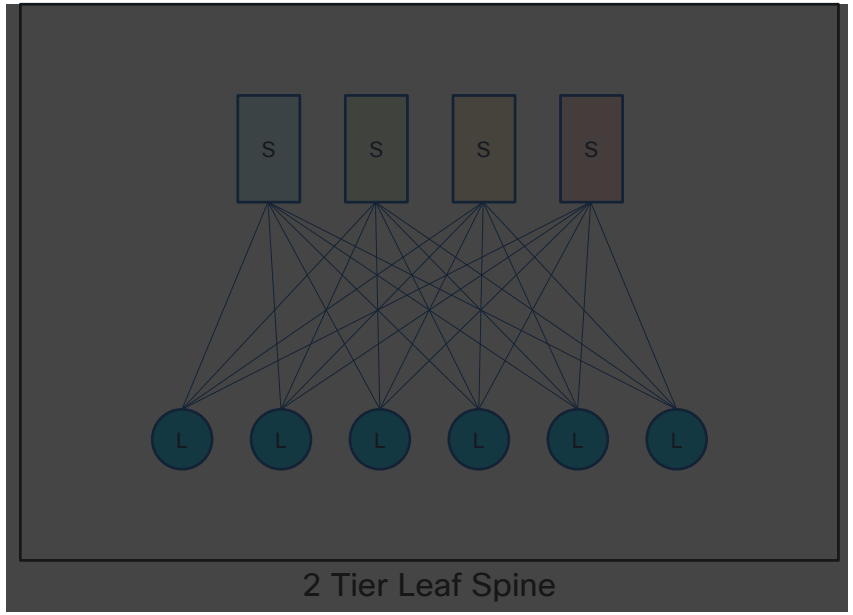

2 Tier Leaf Spine

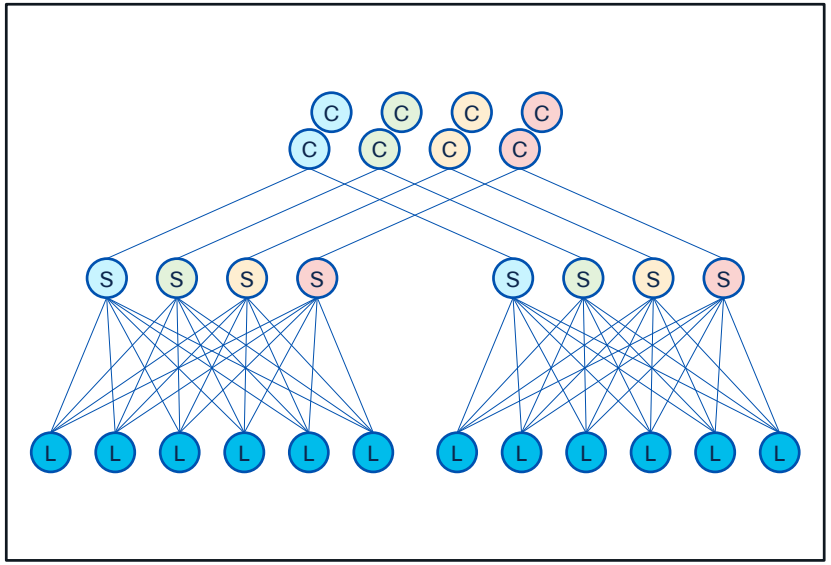

3 Tier Leaf-Spine-Core (or Super-Spine)

Let's Move Forward

Cisco Distributed DC Architectures

Building Distributed DC Architectures

Homogeneous Options

ACI Multi-Pod

Since 2017

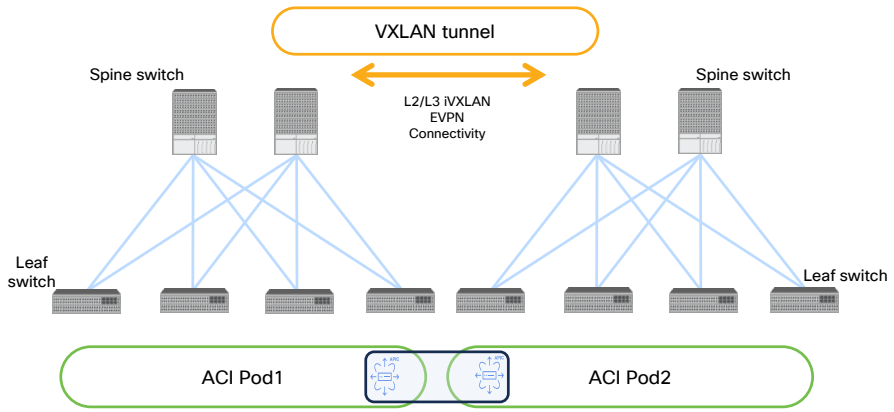

ACI Multi-Site

Since 2018

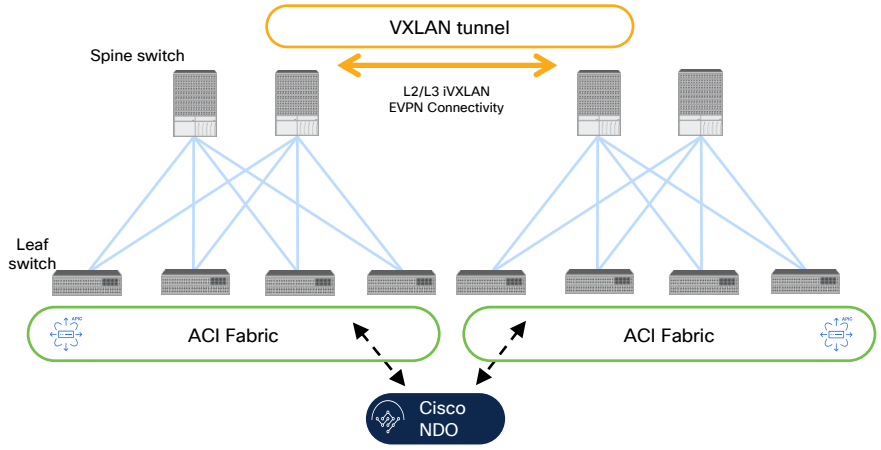

VXLAN EVPN Multi-Site

Since 2017

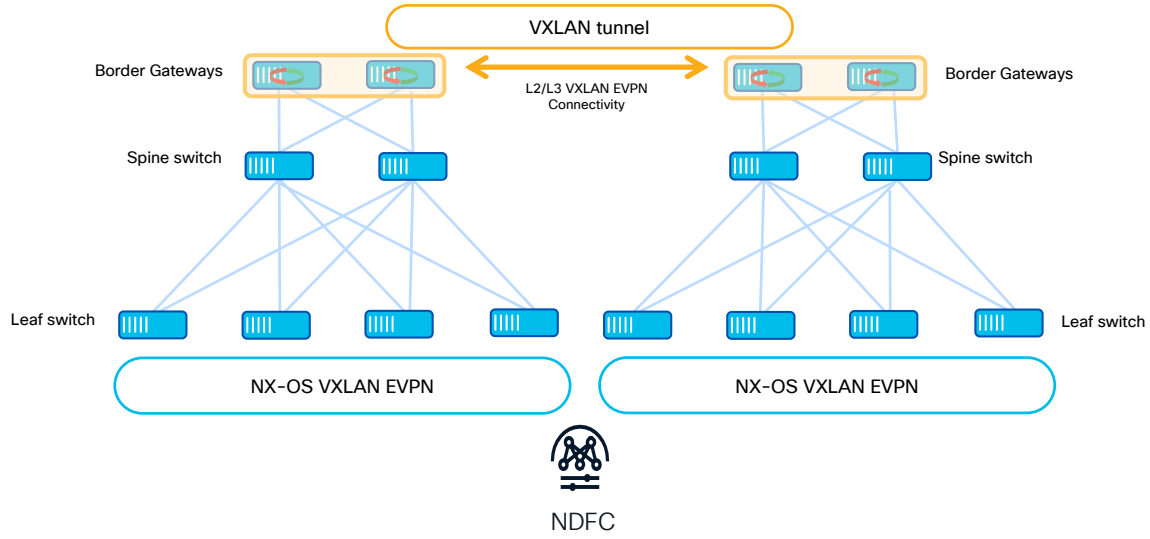

Use of BGWs for Distributed VXLAN Fabrics

For More Information on VXLAN Multi-Site please refer to BRKDCN-2913

Without BGWs (Single Logical Fabric)

- **Two Tiers**

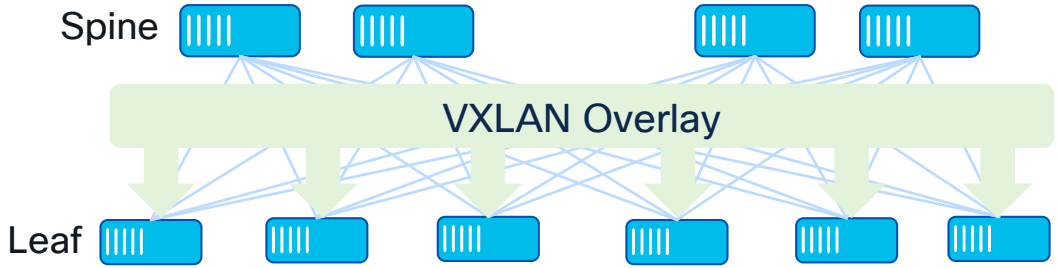

- **Three Tiers**

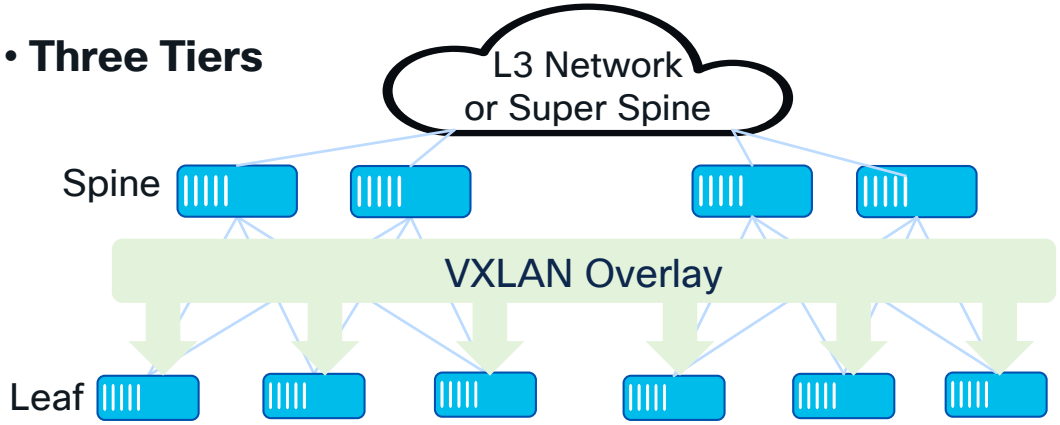

Scale Concerns
(VTEP Scale, BUM Flooding, etc.)

With BGWs (Multi-Site)

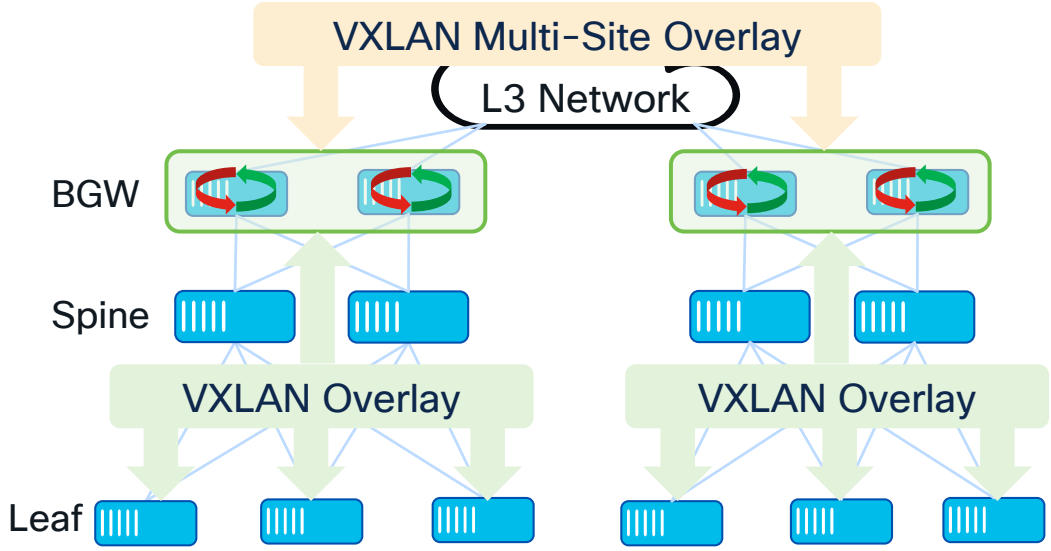

- ✓ Scalable
- ✓ Easy to extend and stitch
- ✓ Flood control across fabrics

(BGW and spine can be the same devices)

Introducing Cisco Nexus ONE Architecture

Cisco Nexus ONE Architecture

Cisco ACI vs Cisco NX-OS Fabrics

Customers can choose one or more Cisco technologies for their DC network

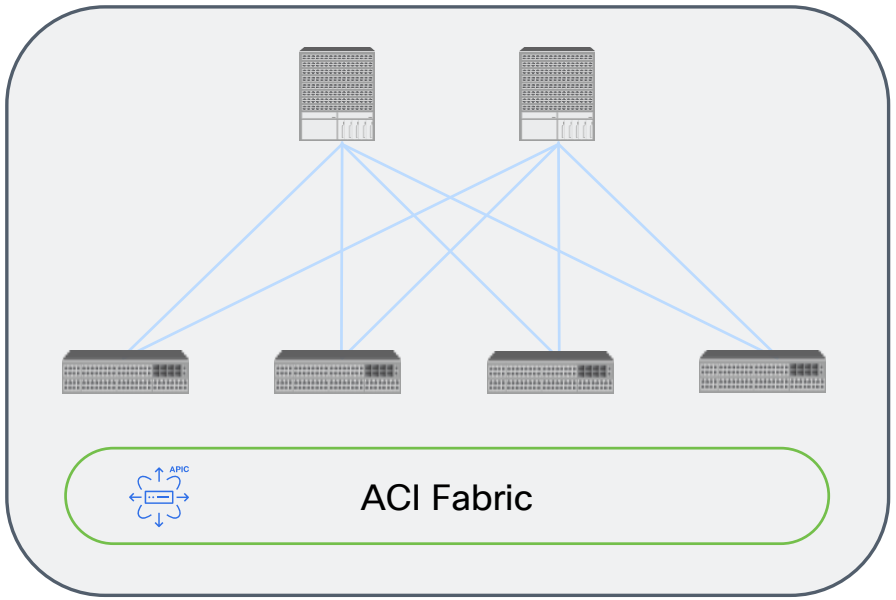

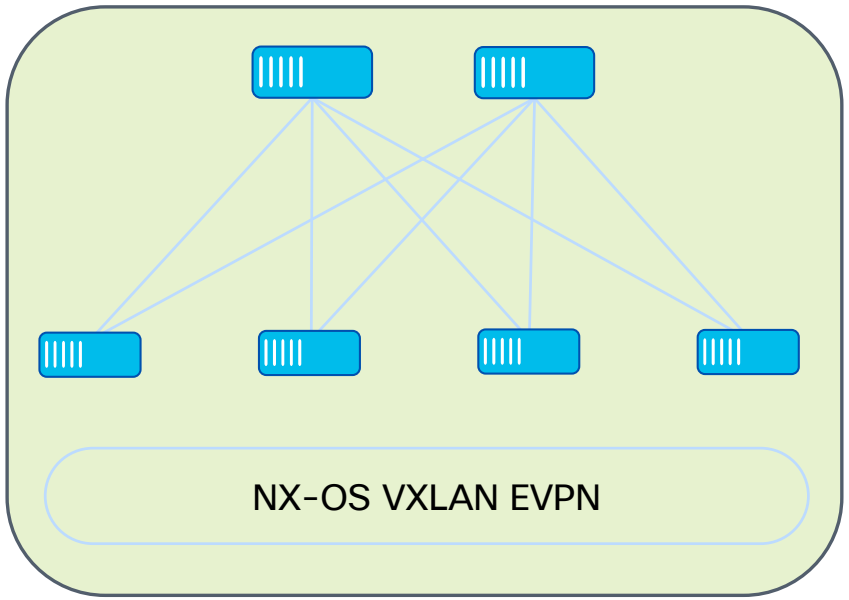

Cisco Nexus ONE Architecture

The Challenges with Interconnecting Heterogeneous Fabrics

Customers can choose one or more Cisco technologies for their DC network

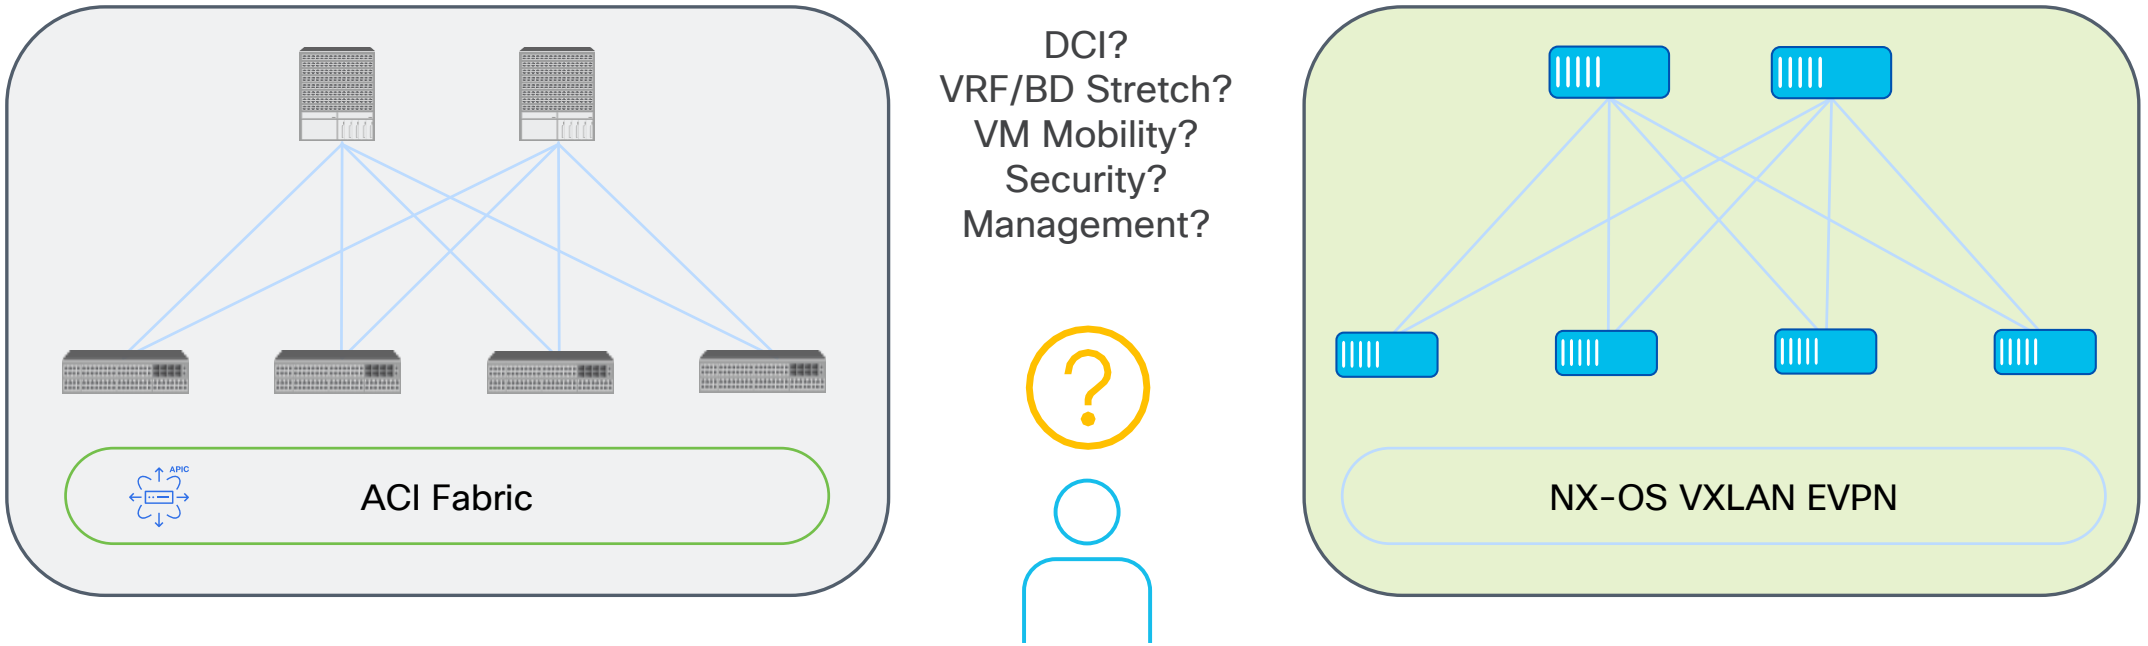

Open NEtworking Fabric Experience

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

Nexus ONE Architecture

Functional Components

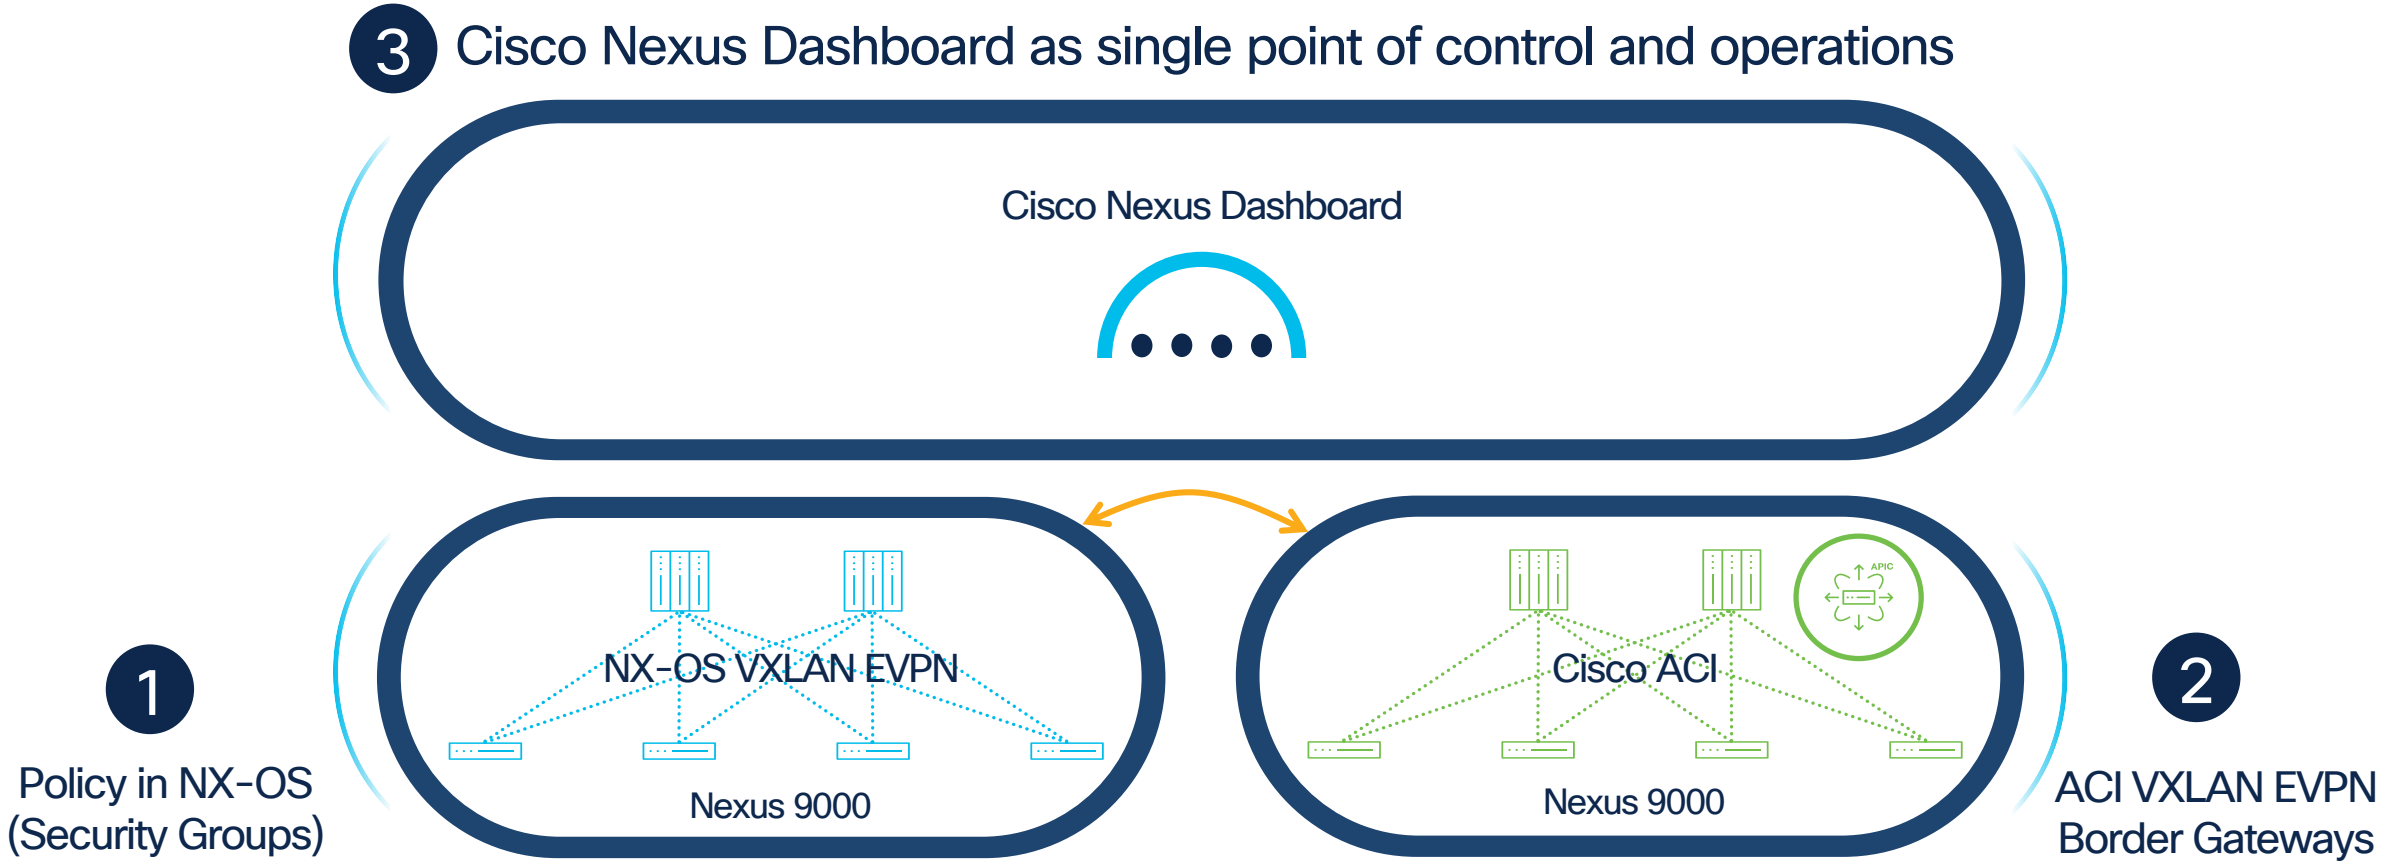

Different fabric architectures | Same outcome with common experience

Nexus ONE Future Evolution

Strategic Vision

For More Information on Cisco Nexus Hyperfabric please refer to BRKDCN-2944

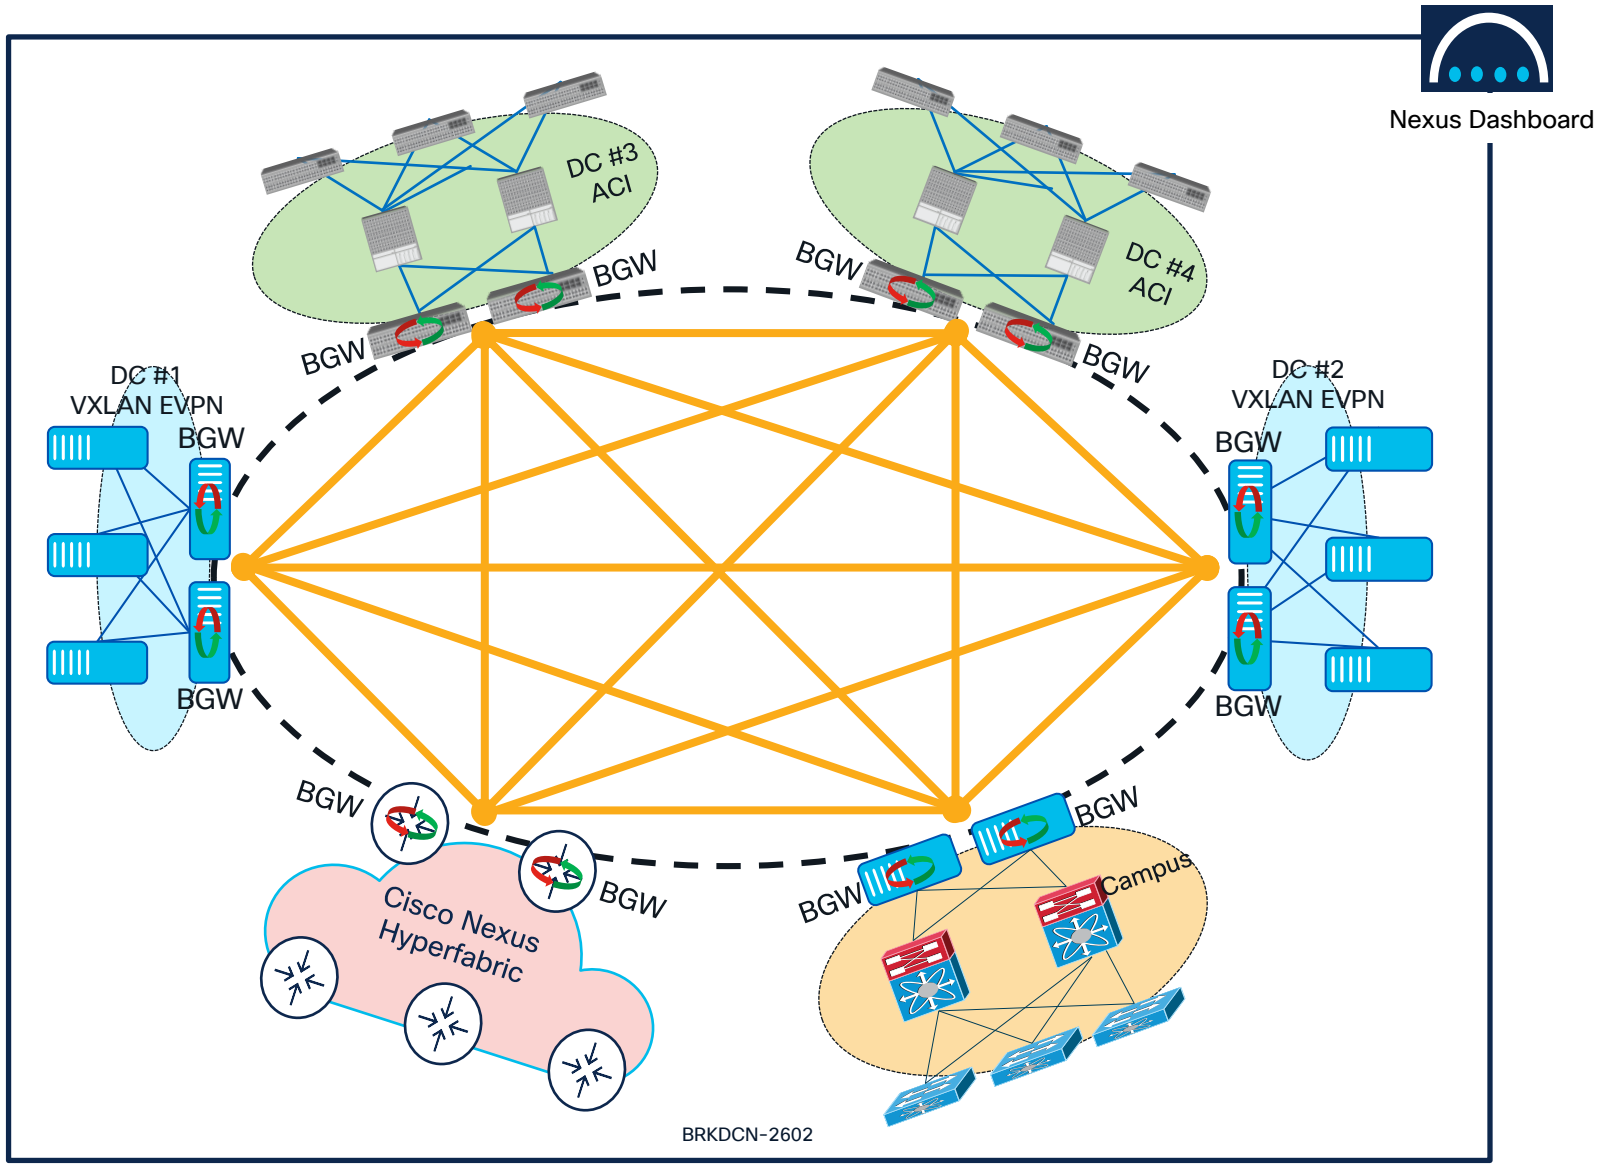

For More Information on VXLAN GPO please refer to BRKDCN-2633

VXLAN GPO

Why Group Policy Option (GPO)?

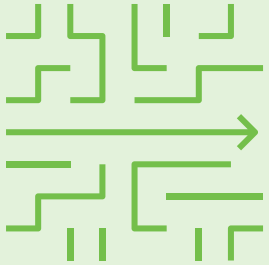

Ability to segment east-west traffic

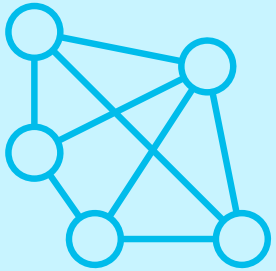

Smaller attack surface and better security

Auditing, compliance and conformance

VXLAN GPO with NX-OS

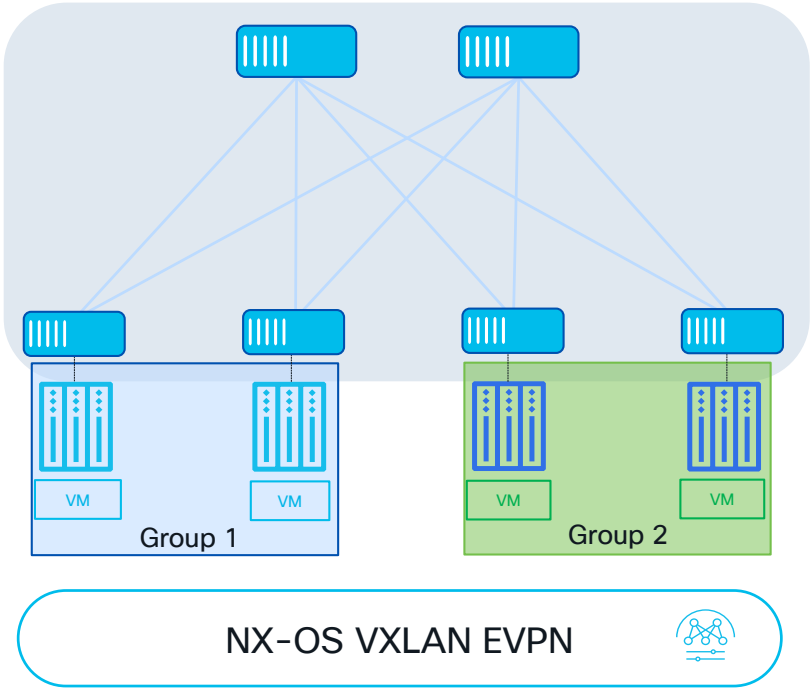

VXLAN GPO with NX-OS

- Group Policy Object carried in standard VXLAN header
- BGP Extended Community to advertise policy information on EVPN control plane advertisements
- 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 Endpoint Security Groups (ESGs*)
- Possible actions: permit, deny, redirect (service chaining)

Automation

- Automate using Nexus Dashboard (ND) 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

Creation of Security Zones

- VXLAN GPO allows to define policies for enforcing security policies (SGACLs) between endpoint security groups (ESGs)
- 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)

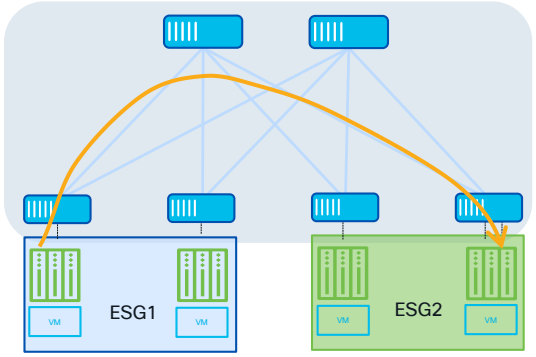

Service Chaining

- 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)

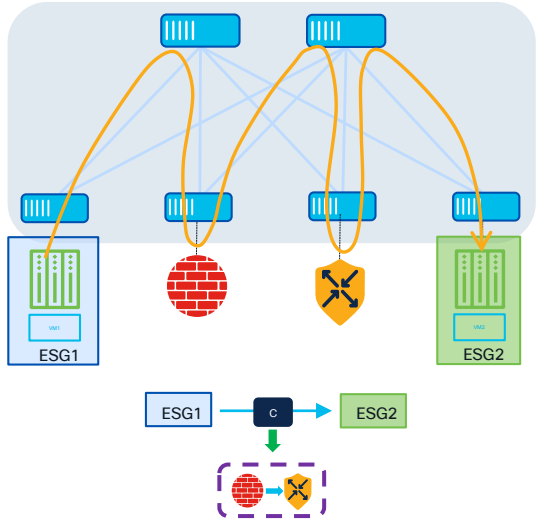

VXLAN GPO with NX-OS

Cisco GPO Control Plane Functionalities

Data Plane
(draft-smith-vxlan-group-policy)

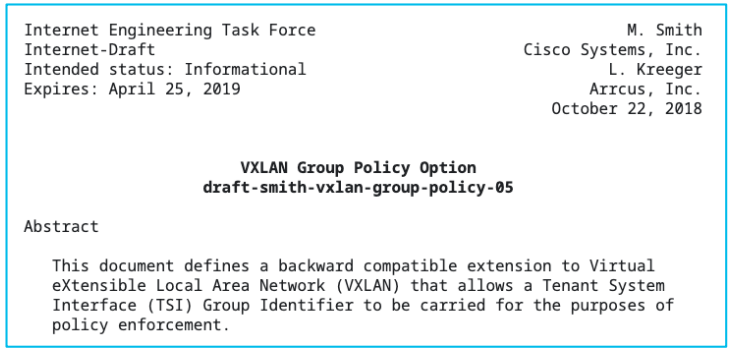

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: April 25, 2019

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)

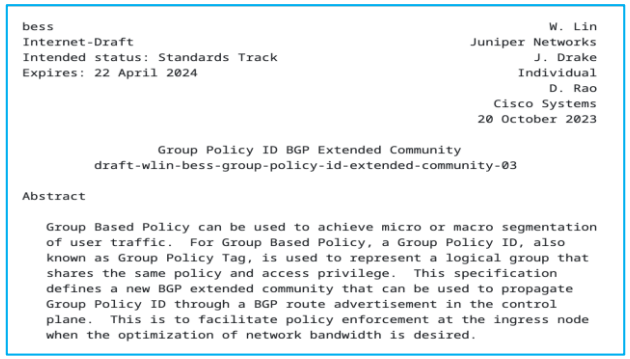

bess
Internet-Draft
Intended status: Standards Track
Expires: 22 April 2024

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.

VXLAN GPO with NX-OS

Use of MP-BGP EVPN for Exchanging Security Tags

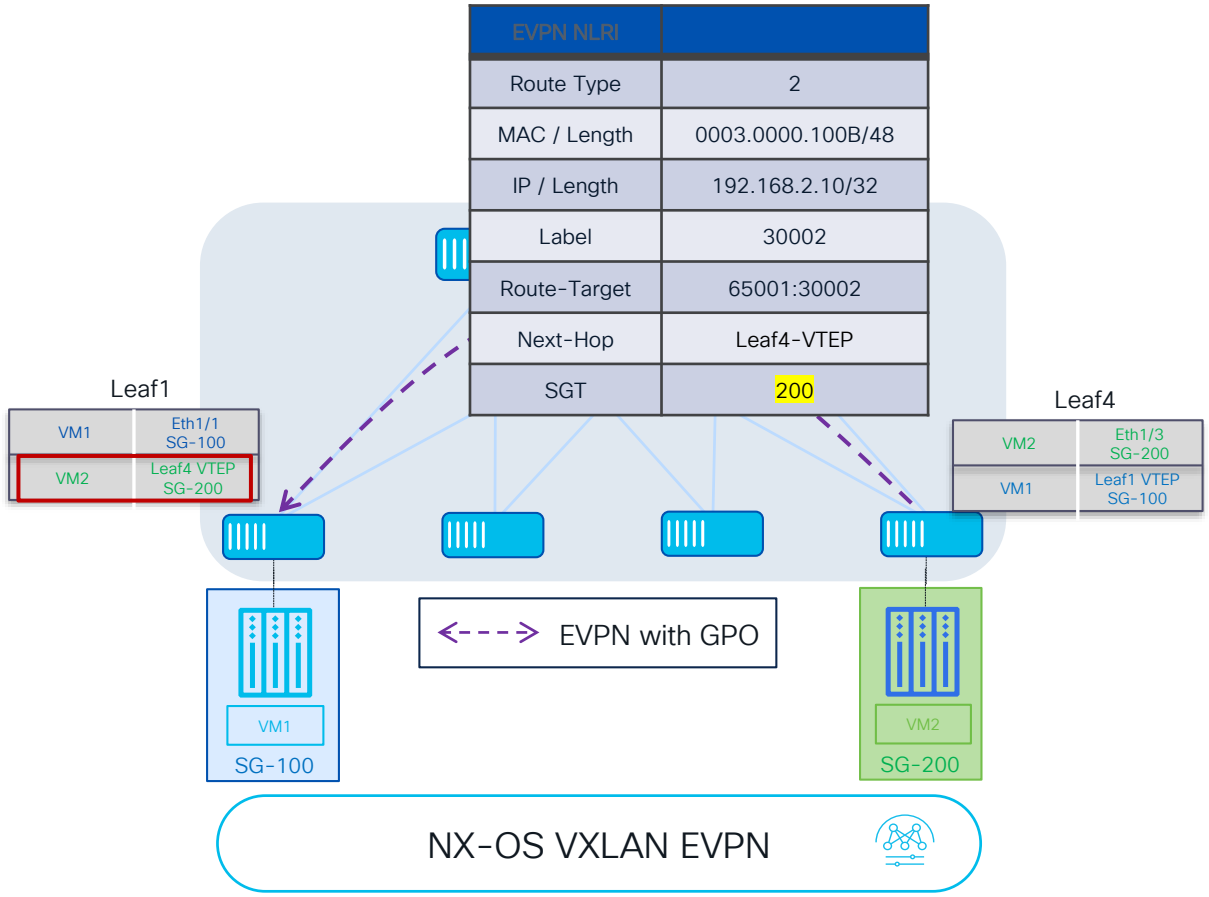

- Use of MP-BGP EVPN control plane to propagate endpoints connectivity and policy information inside the fabric
- The SGT information is propagated as a BGP extended community (Group Policy ID Extended Community)

VXLAN GPO with NX-OS

Enforcing Policy on the Ingress Leaf

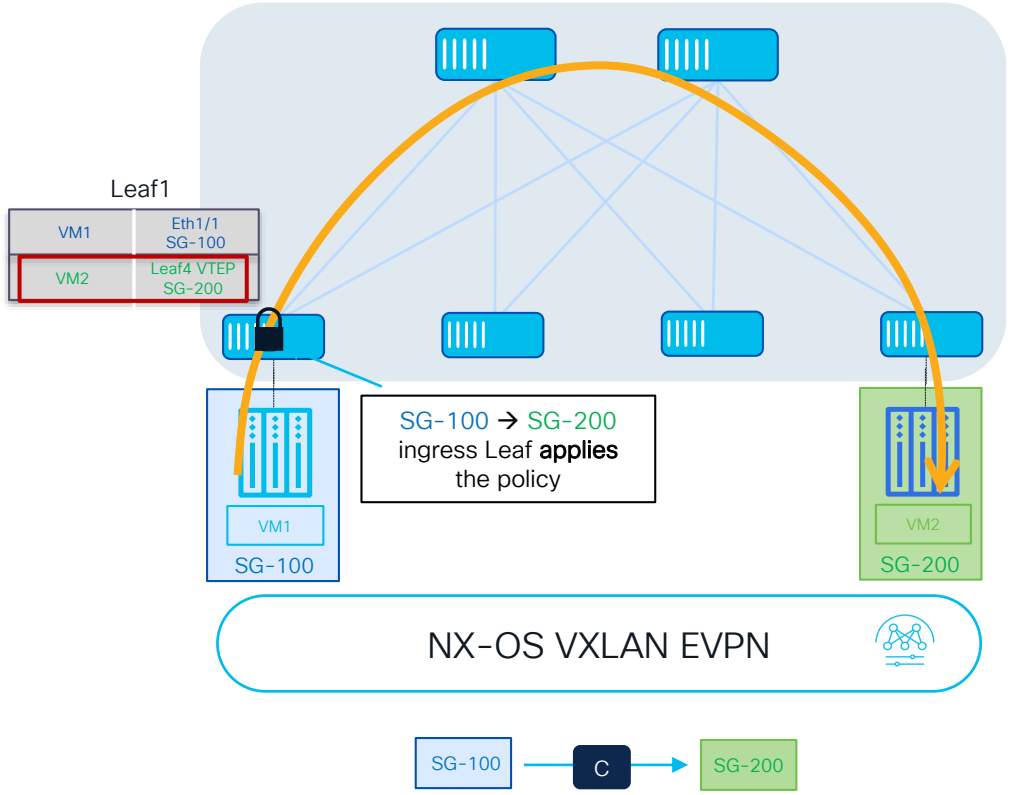

- Facilitate the enforcement of policy on the ingress leaf node (for both directions)
- Security Group Access Control Lists (SGACLs/contracts) enforced between groups

VXLAN GPO with NX-OS

Traffic Steering with Policy Based Redirection

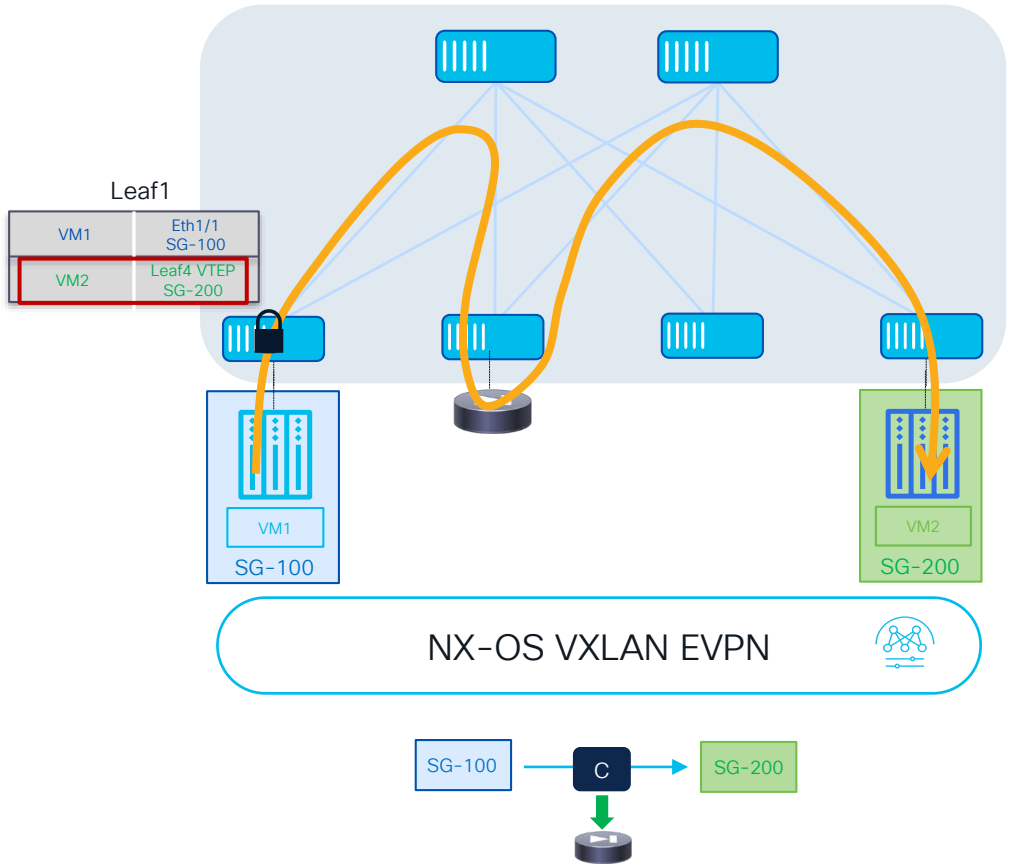

- Policy Based Redirection capabilities to steer through one or more service devices (firewall, load balancers, etc.) traffic flows between different security groups
- Redirection to a Firewall service function with NX-OS 10.4(3)F
- Other use cases, including traffic stitching through multiple services, planned for NX-OS 10.5(x) release train

VXLAN GPO with NX-OS

Multi-Site Support

Policy-Aware to Policy-Aware

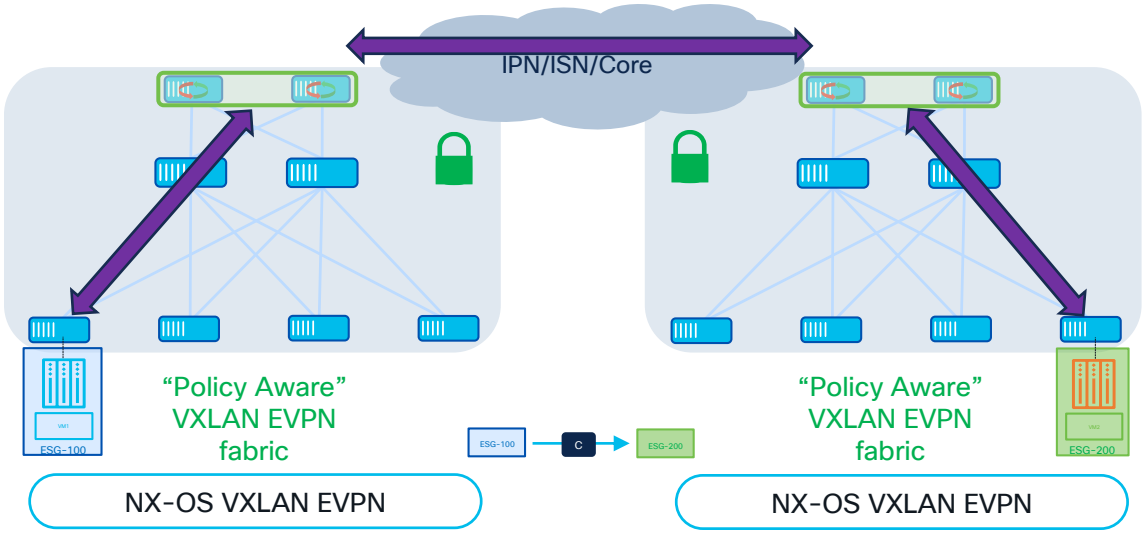

Policy-Aware to Policy-Unaware

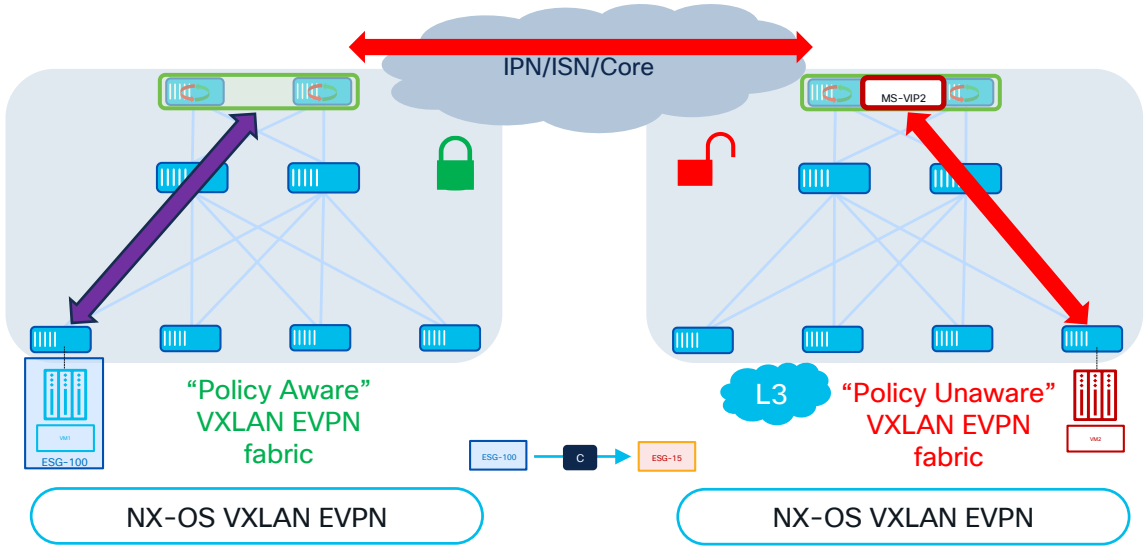

For More Information on ACI BGWs please refer to BRKDCN-2634

ACI Border Gateways (BGWs)

Heterogeneous Fabrics

ACI 6.1(x)

Introducing ACI Border Gateways

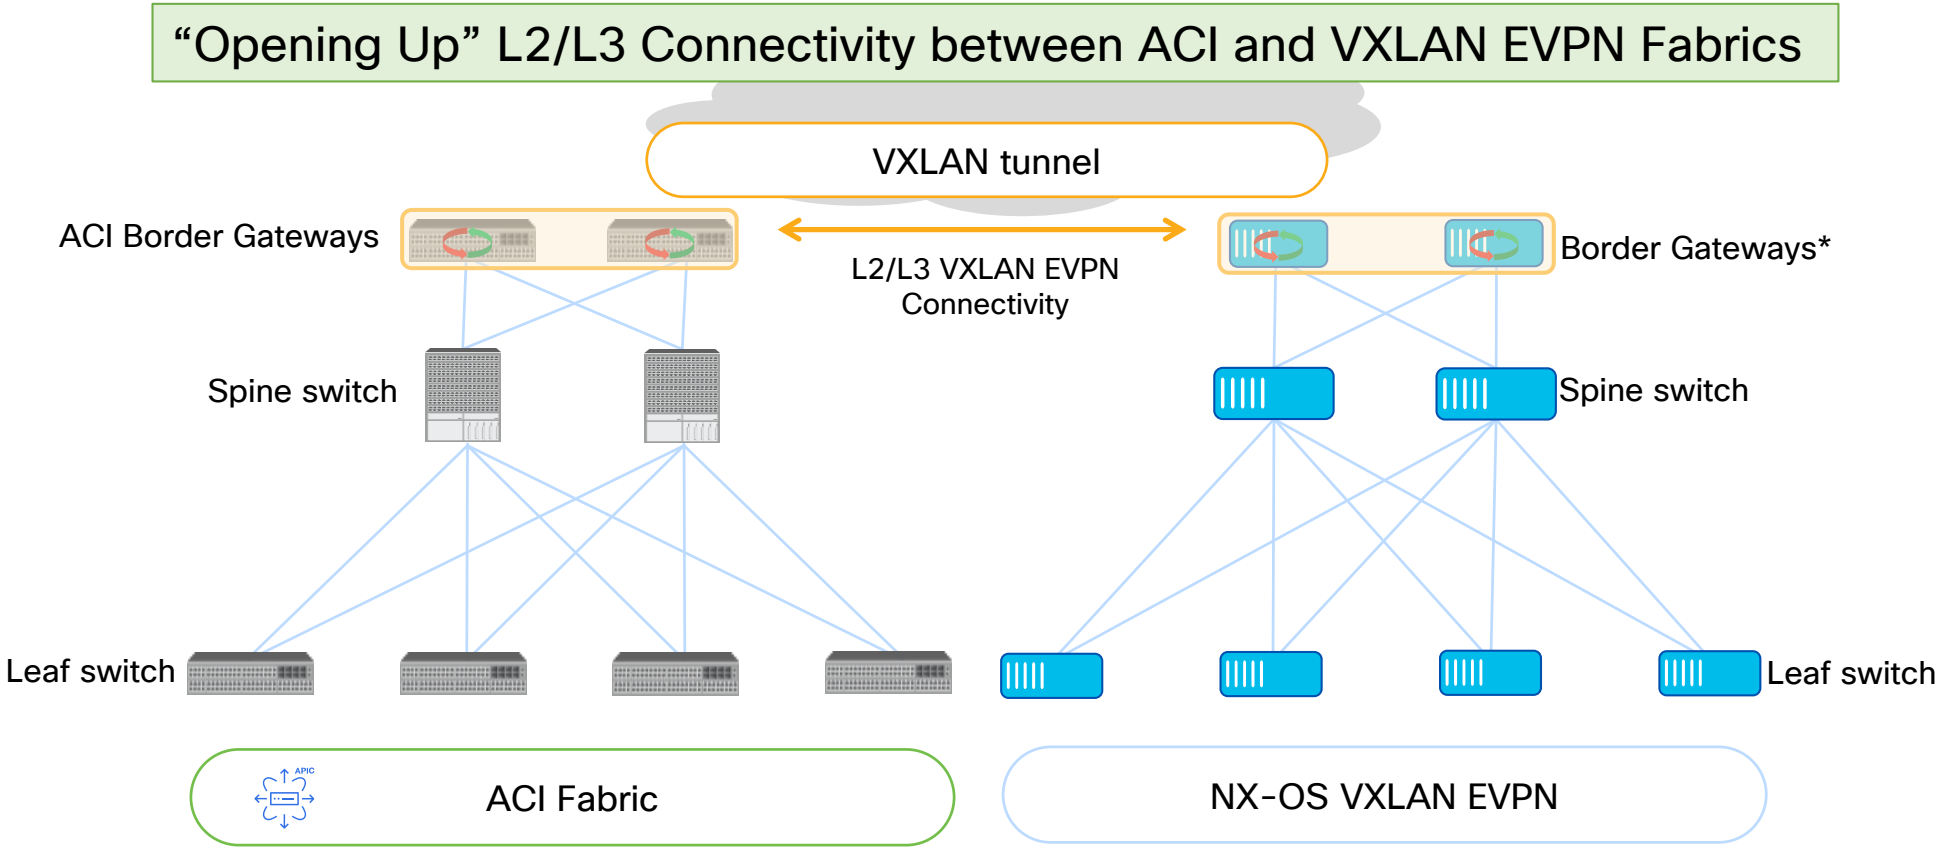

ACI Border Gateways

Overlay EVPN Connectivity across Domains

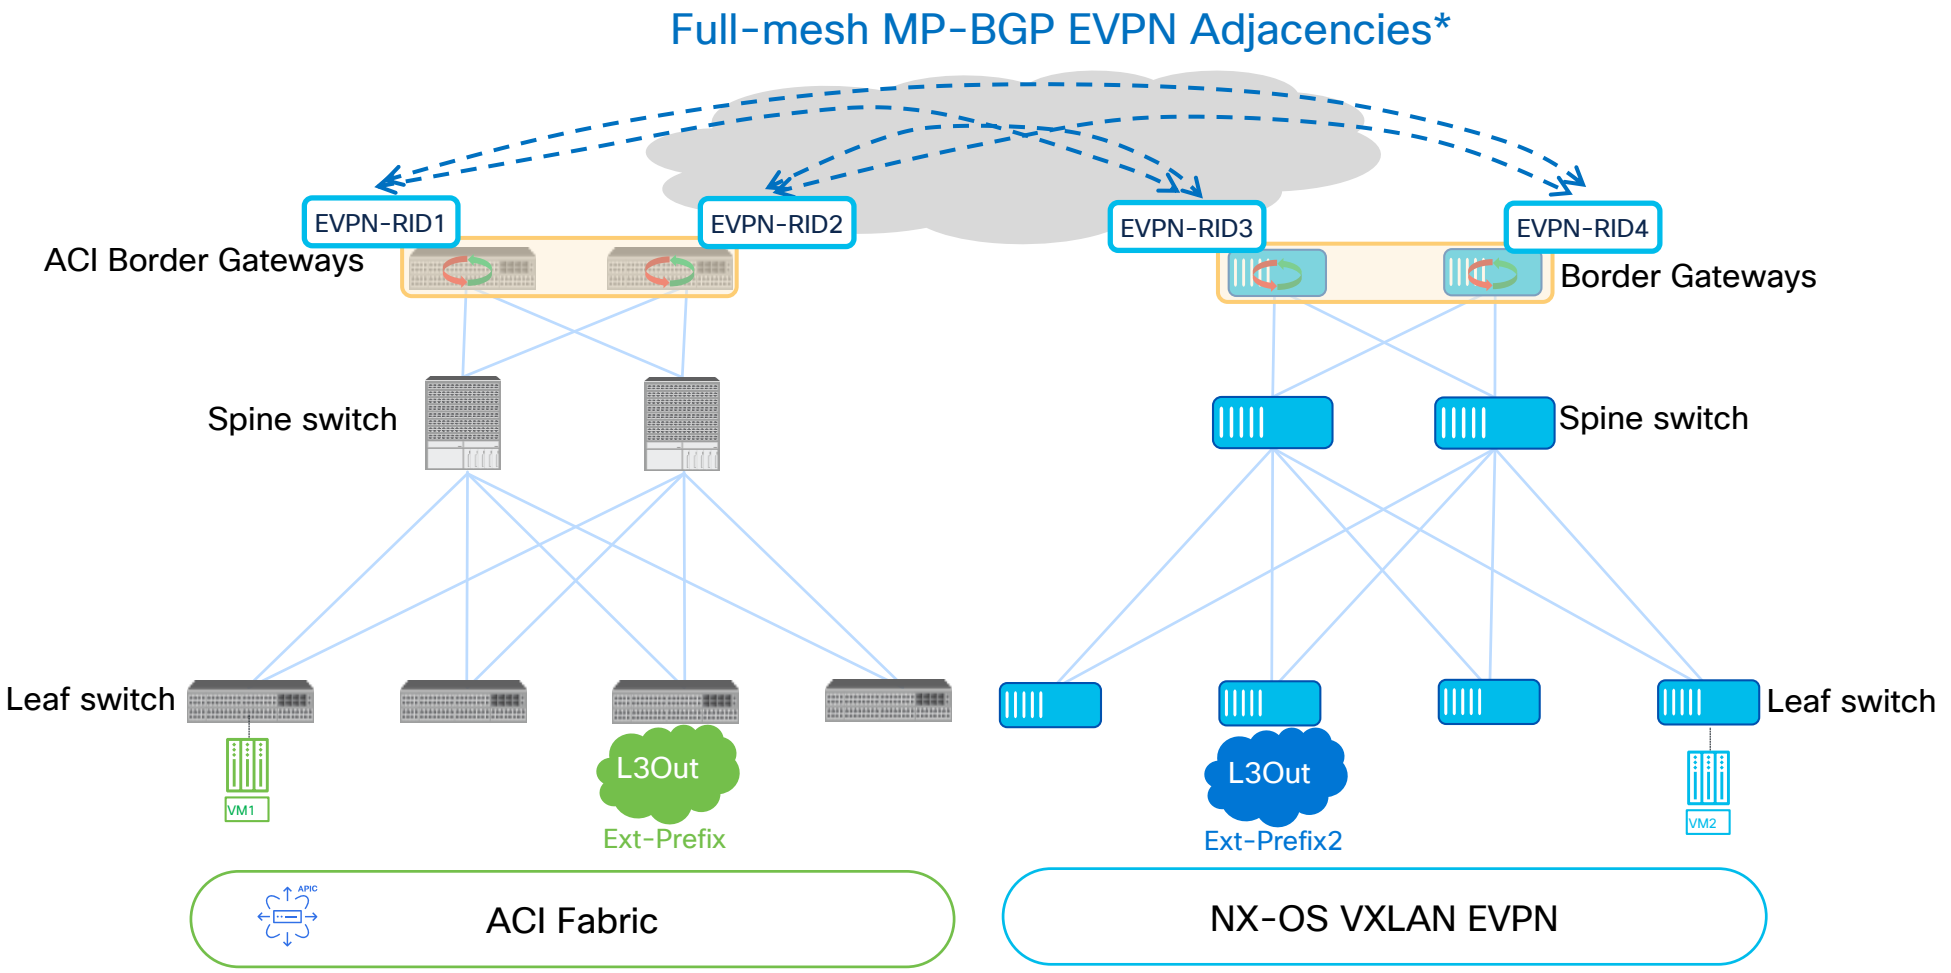

*No current Route-Server support

ACI Border Gateways

Data-Plane Overview

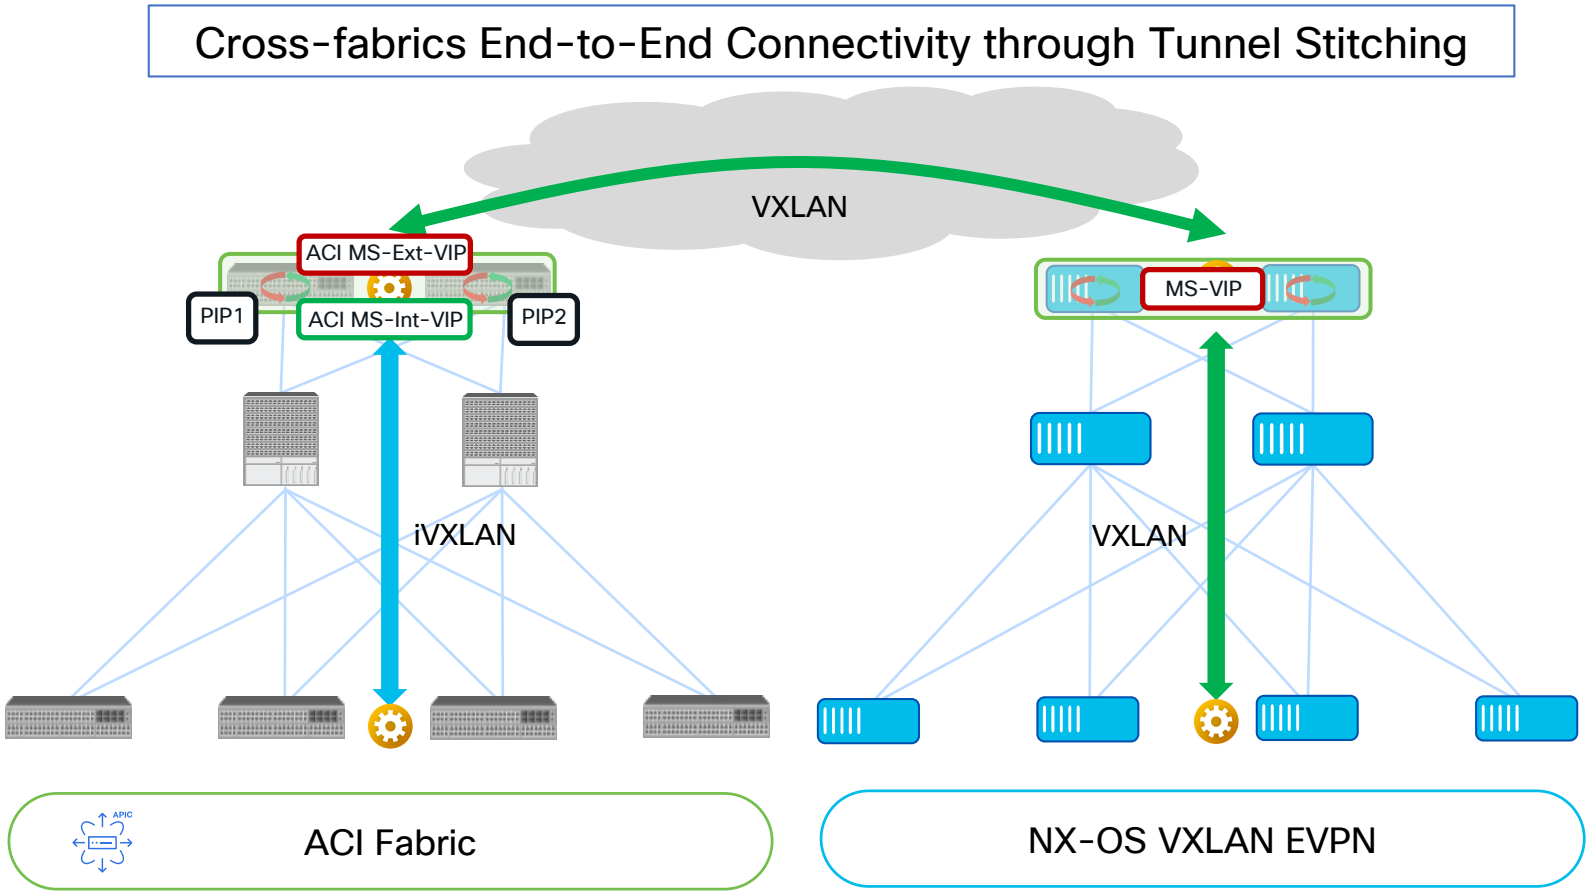

NX-OS GPO + ACI BGWs
End-to-End Policy Domain between Heterogeneous Fabrics

Use Case 1

Cisco ACI to Policy Unaware VXLAN EVPN

Cisco ACI to Policy Unaware VXLAN EVPN

Classification and Enforcement only on Cisco ACI

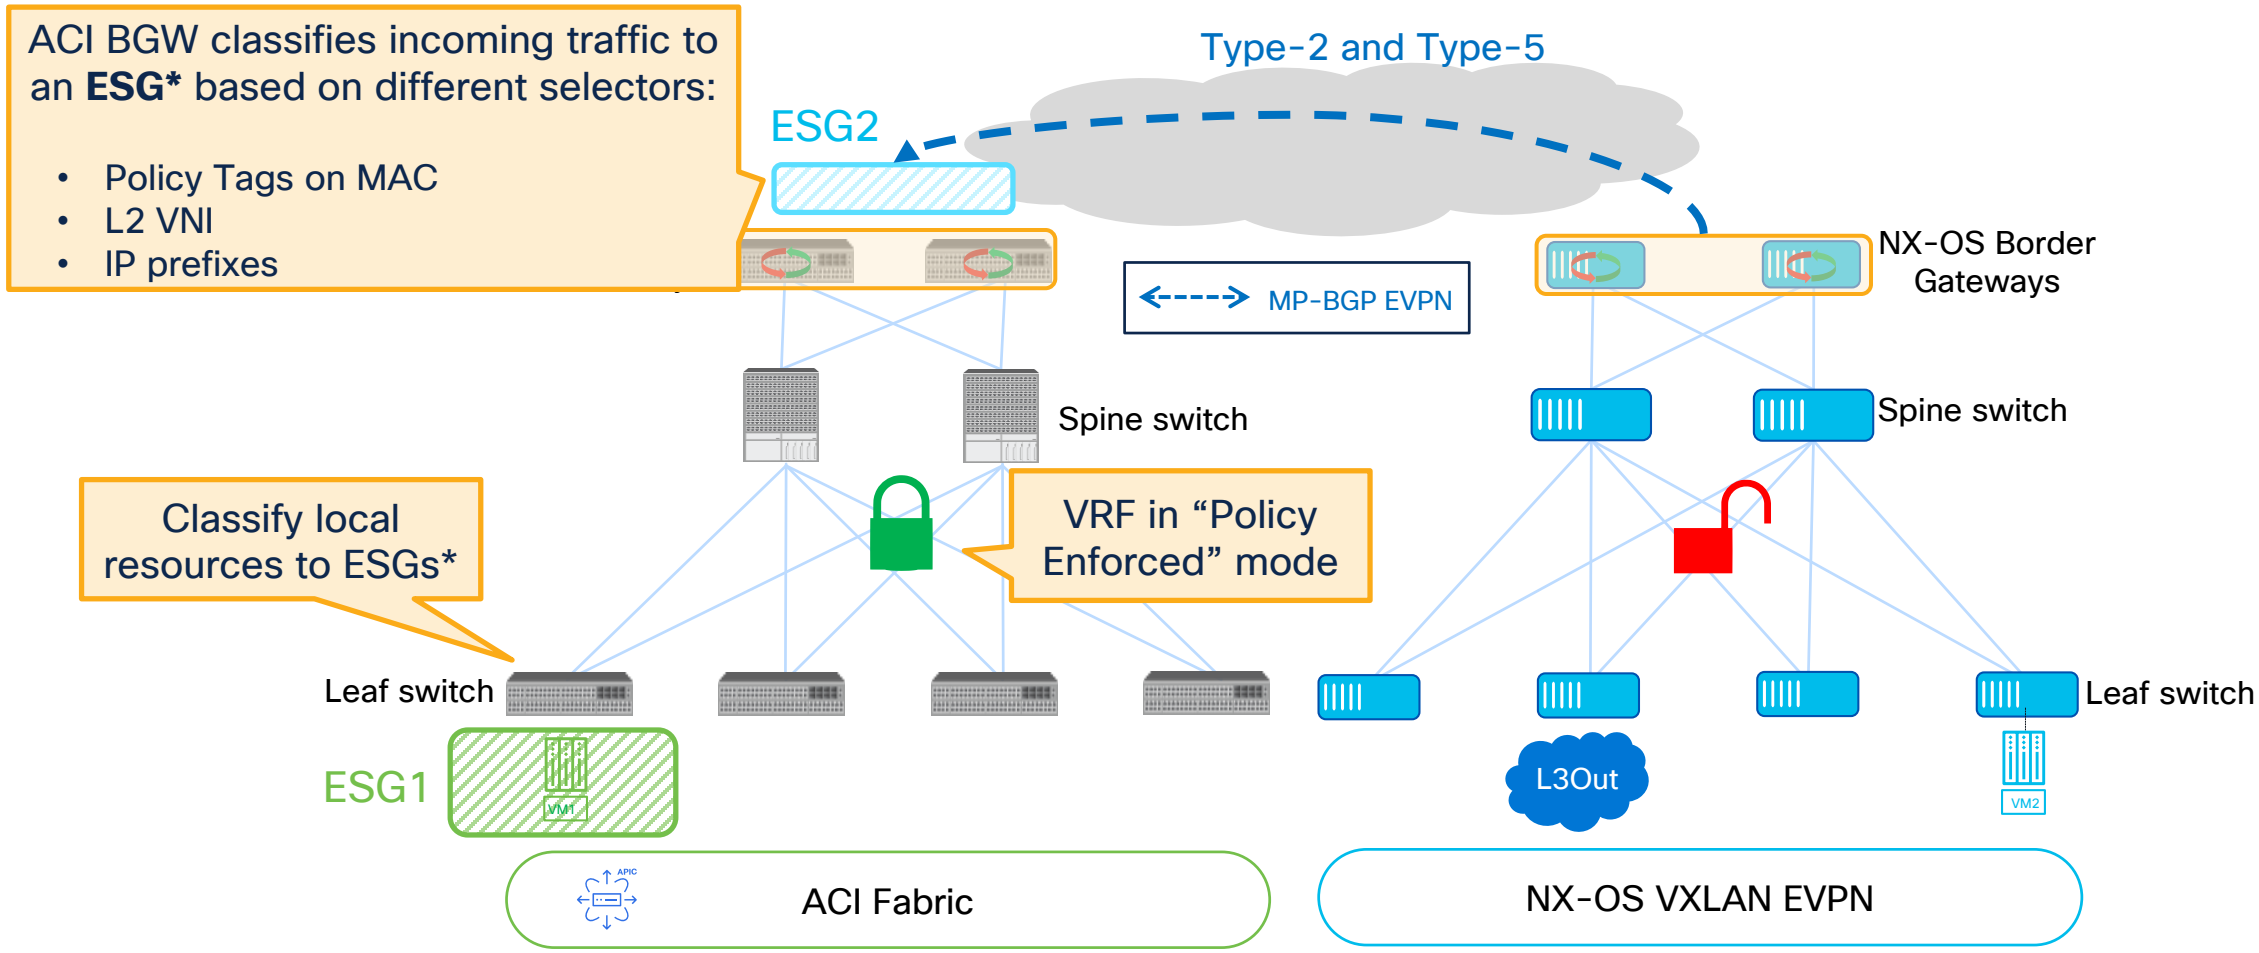

*Nexus ONE mandates the use of Endpoint Security Groups (ESGs) in ACI for classification and enforcement. EPG to ESG migration procedure ("uplifting") should be performed upfront.

Cisco ACI to Policy Unaware VXLAN EVPN

Classification and Enforcement only on Cisco ACI

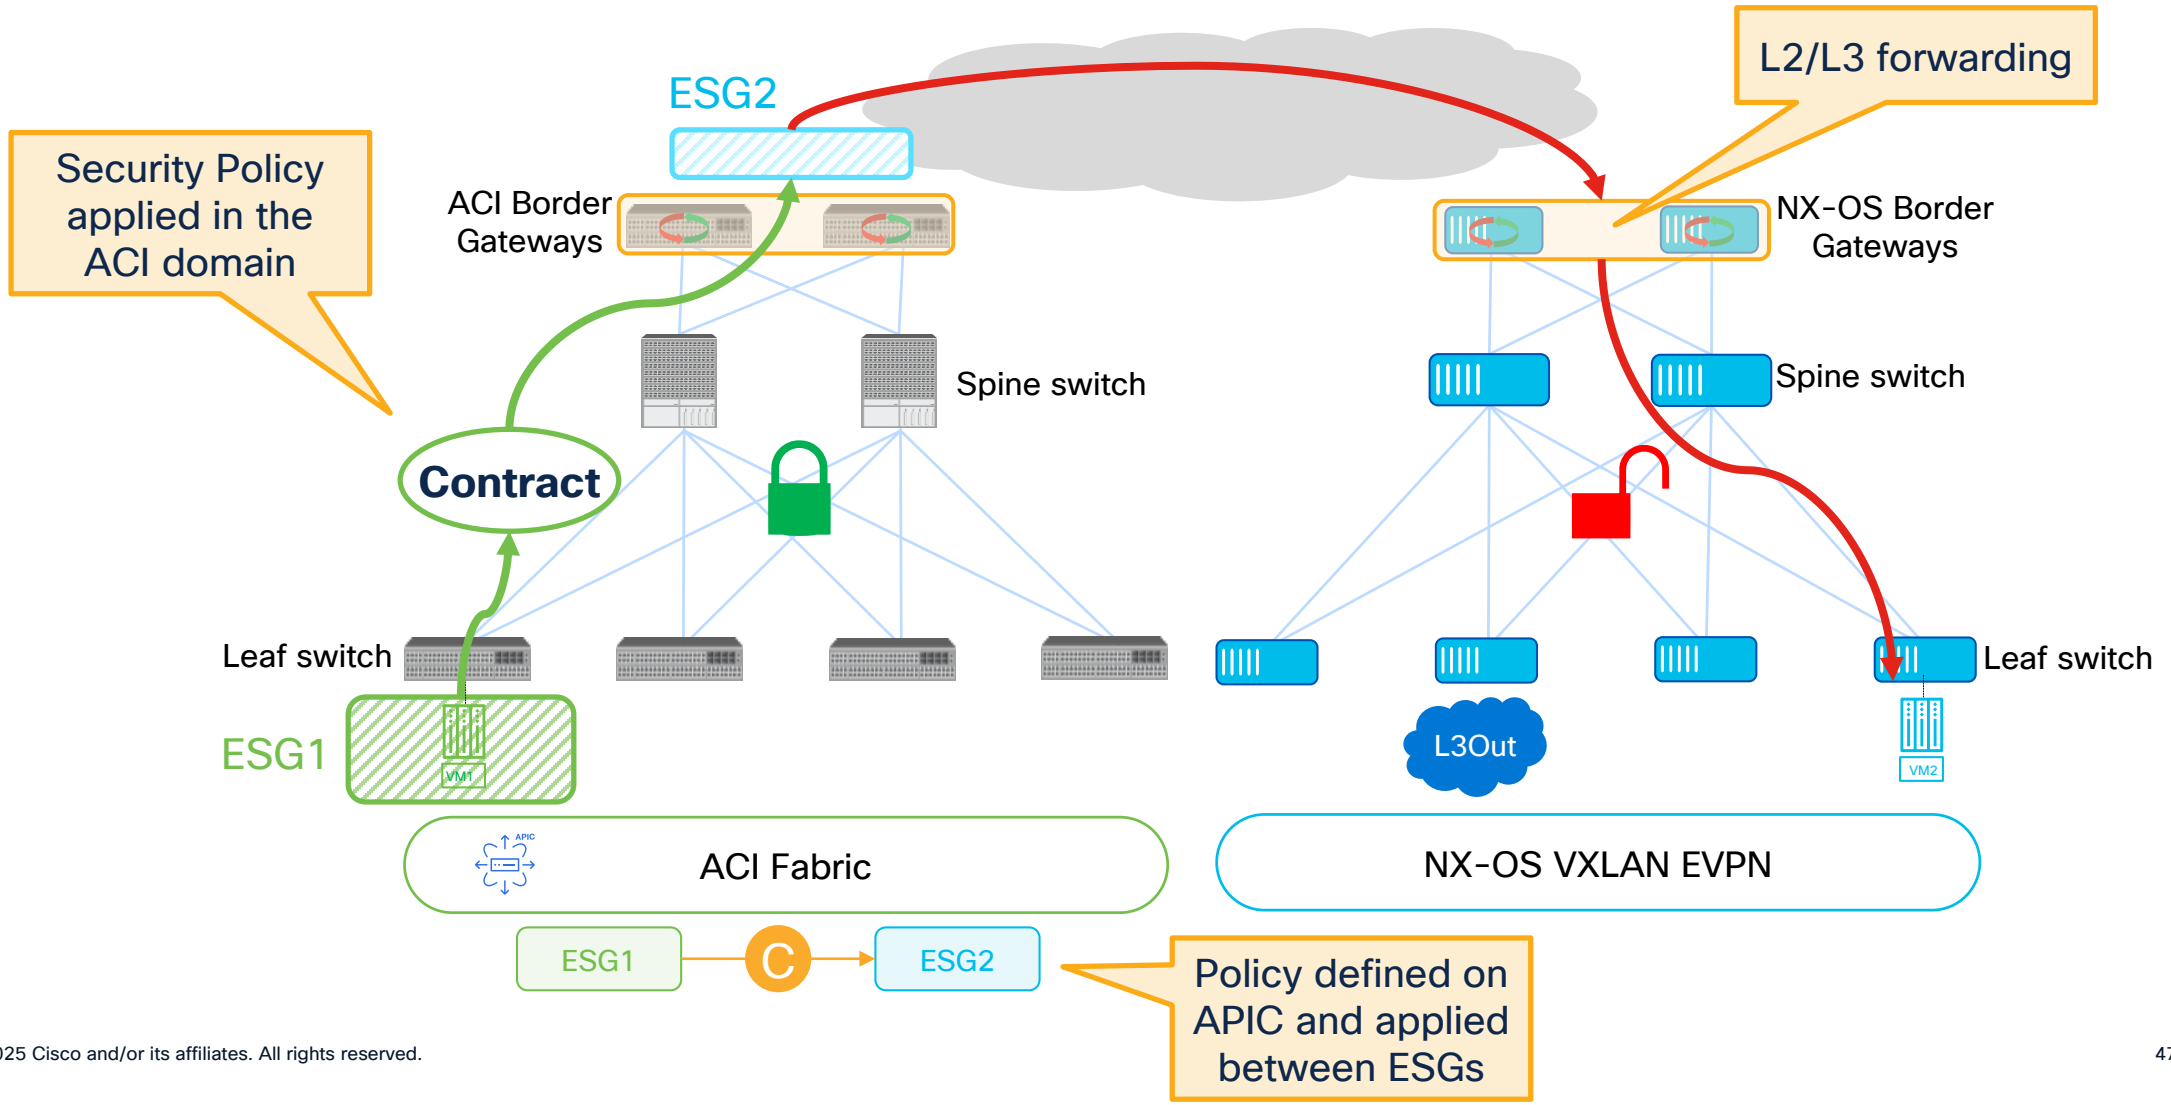

Use Case 2

Cisco ACI to Policy Aware VXLAN EVPN

Cisco ACI to Policy Aware VXLAN EVPN

Classification and Enforcement in both Domains

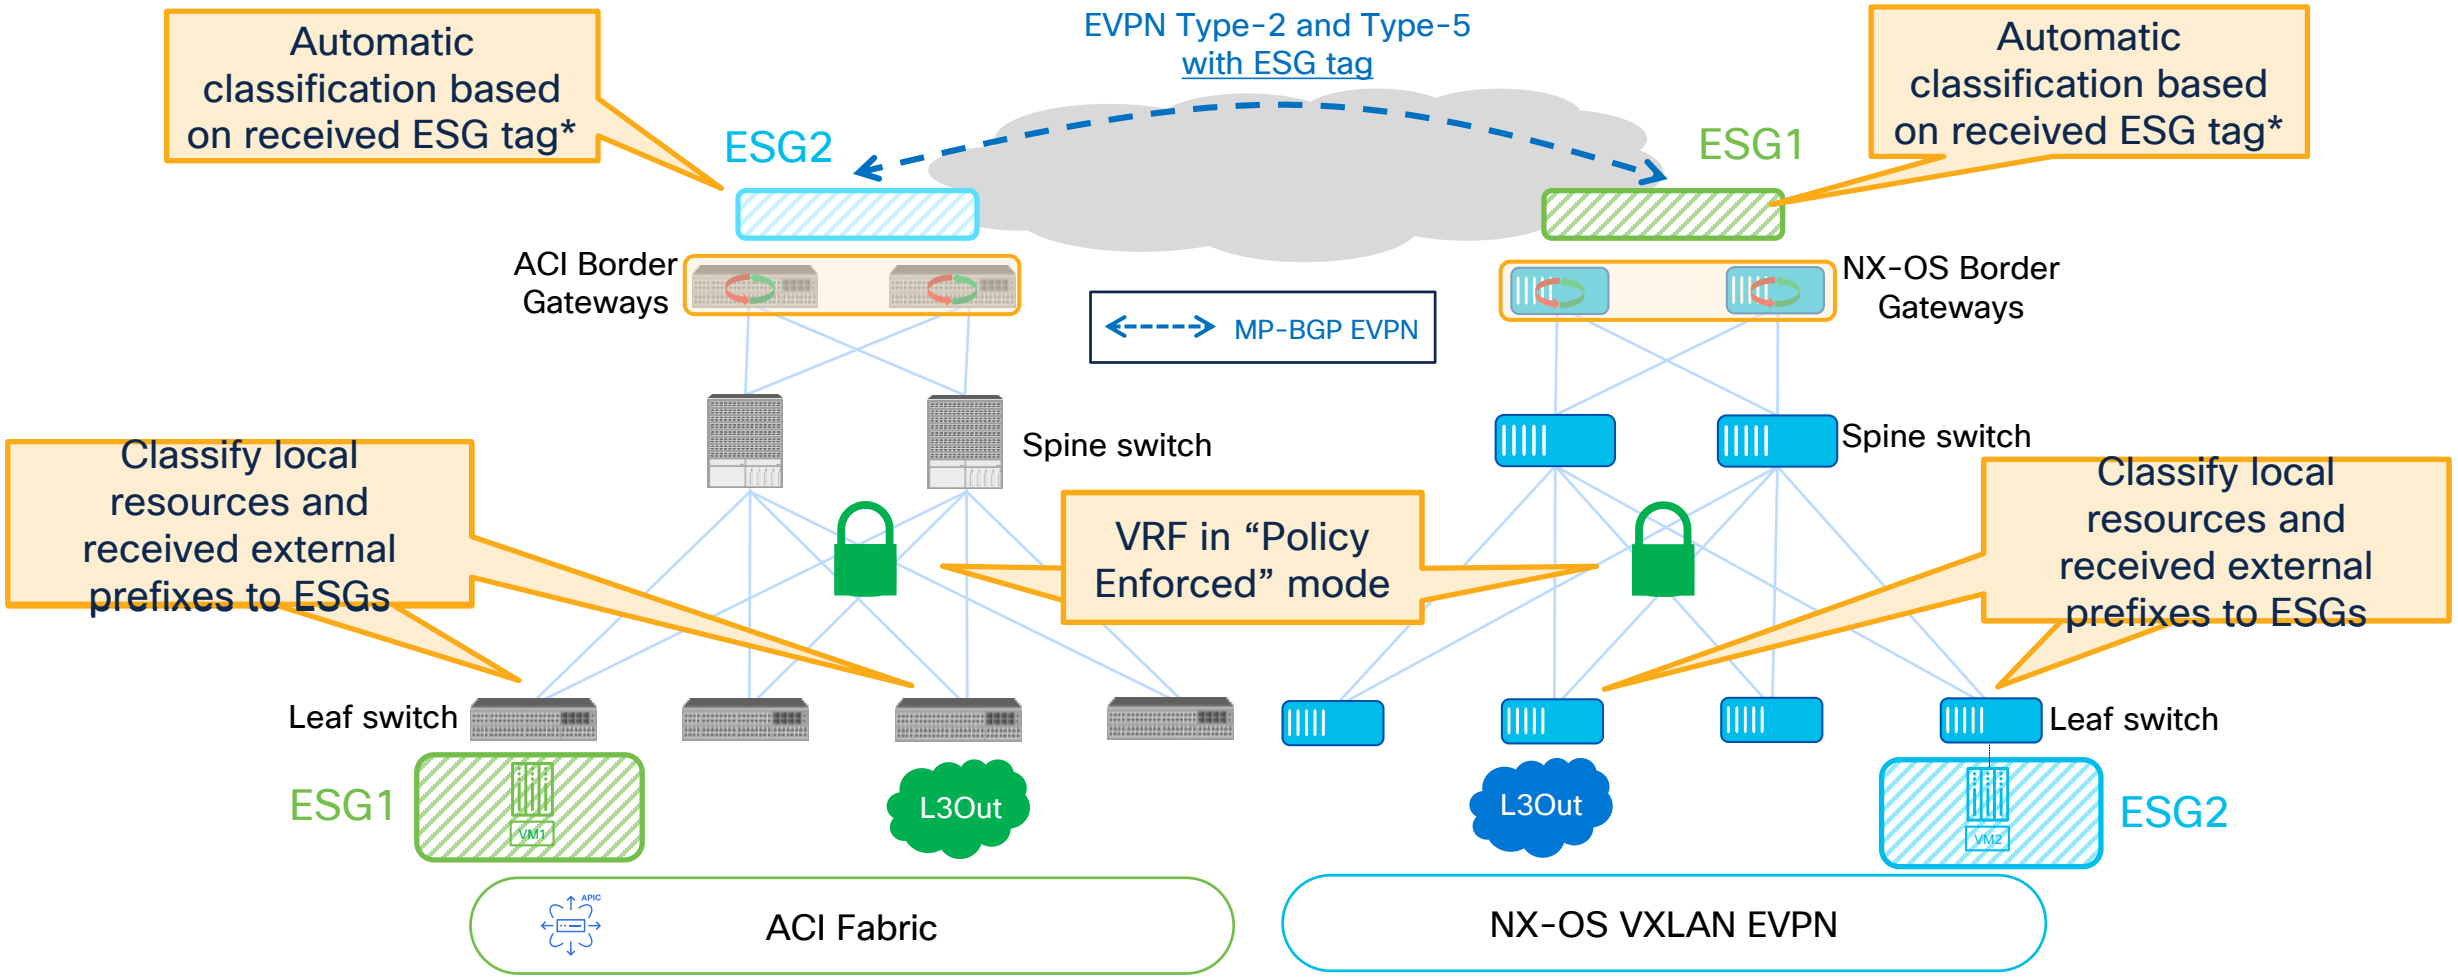

*Tag normalization performed on the ACI BGWs

Cisco ACI to Policy Aware VXLAN EVPN

Classification and Enforcement in both Domains

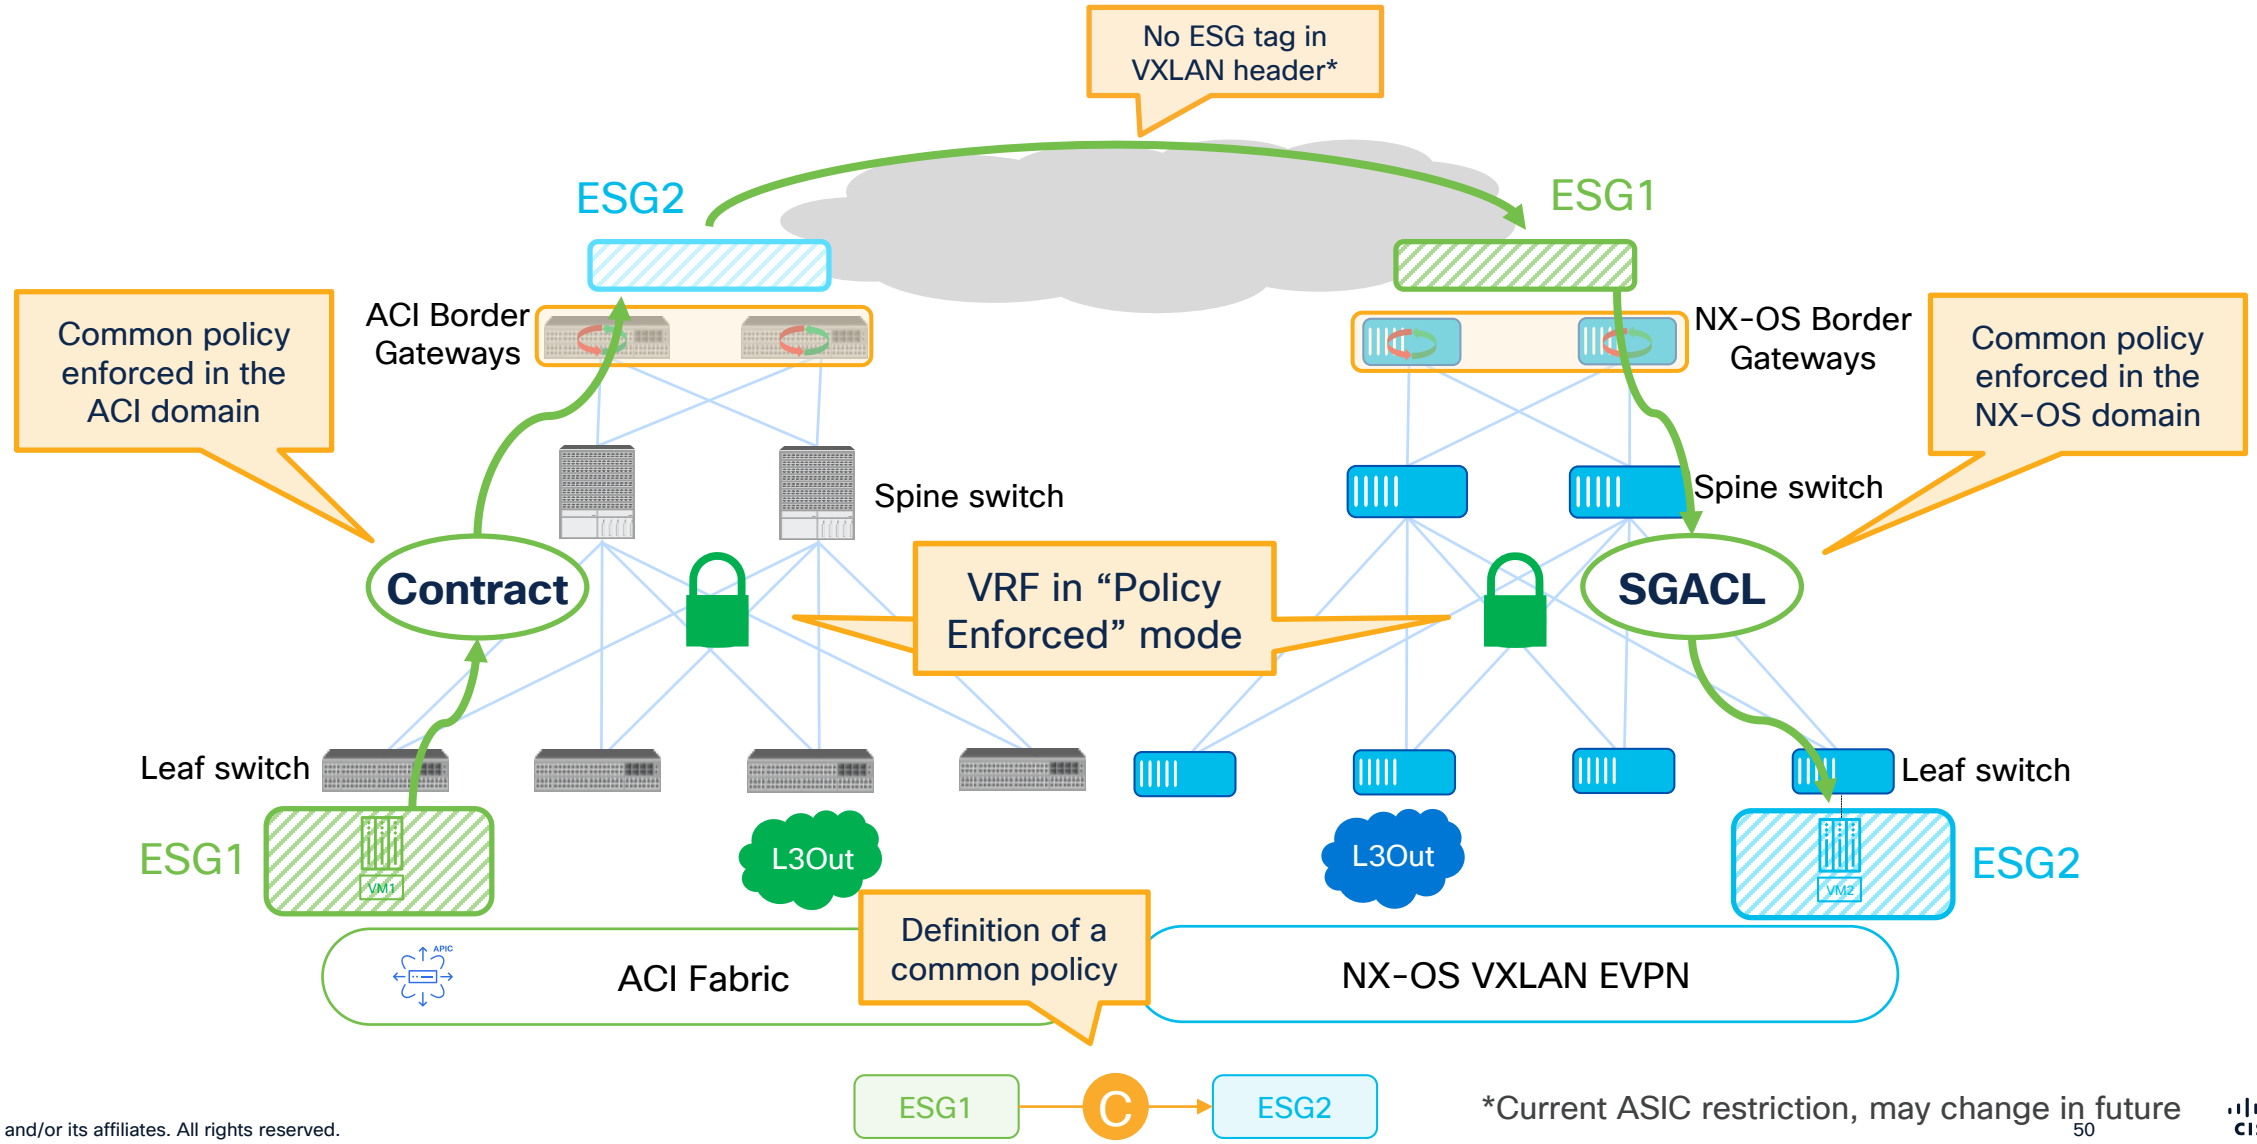

*Current ASIC restriction, may change in future

For More Information on Unified Nexus Dashboard please refer to PSODCN-1009

Unified Nexus Dashboard

We have embarked on a journey

Towards simplification and unification

Products have evolved through the years with microservices-based architectures & other enhancements

2004 2008 2014 2017 2018 2019 2021 2025

Network Assurance Engine

Network Insights Advisor

Network Insights Resources

Insights

Unified Nexus Dashboard

DCNM SAN

DCNM LAN

Multi-Site Orchestrator

Orchestrator

2-way integration

Fabric Controller

ACI

Which leads us here...

Unified Nexus Dashboard

Provisioning

Operations

Access to Nexus Dashboard services/features is based on the Cisco Nexus 9000 switch license tier

Unified Nexus Dashboard

ND 4.2

Single Point of Management and Operation

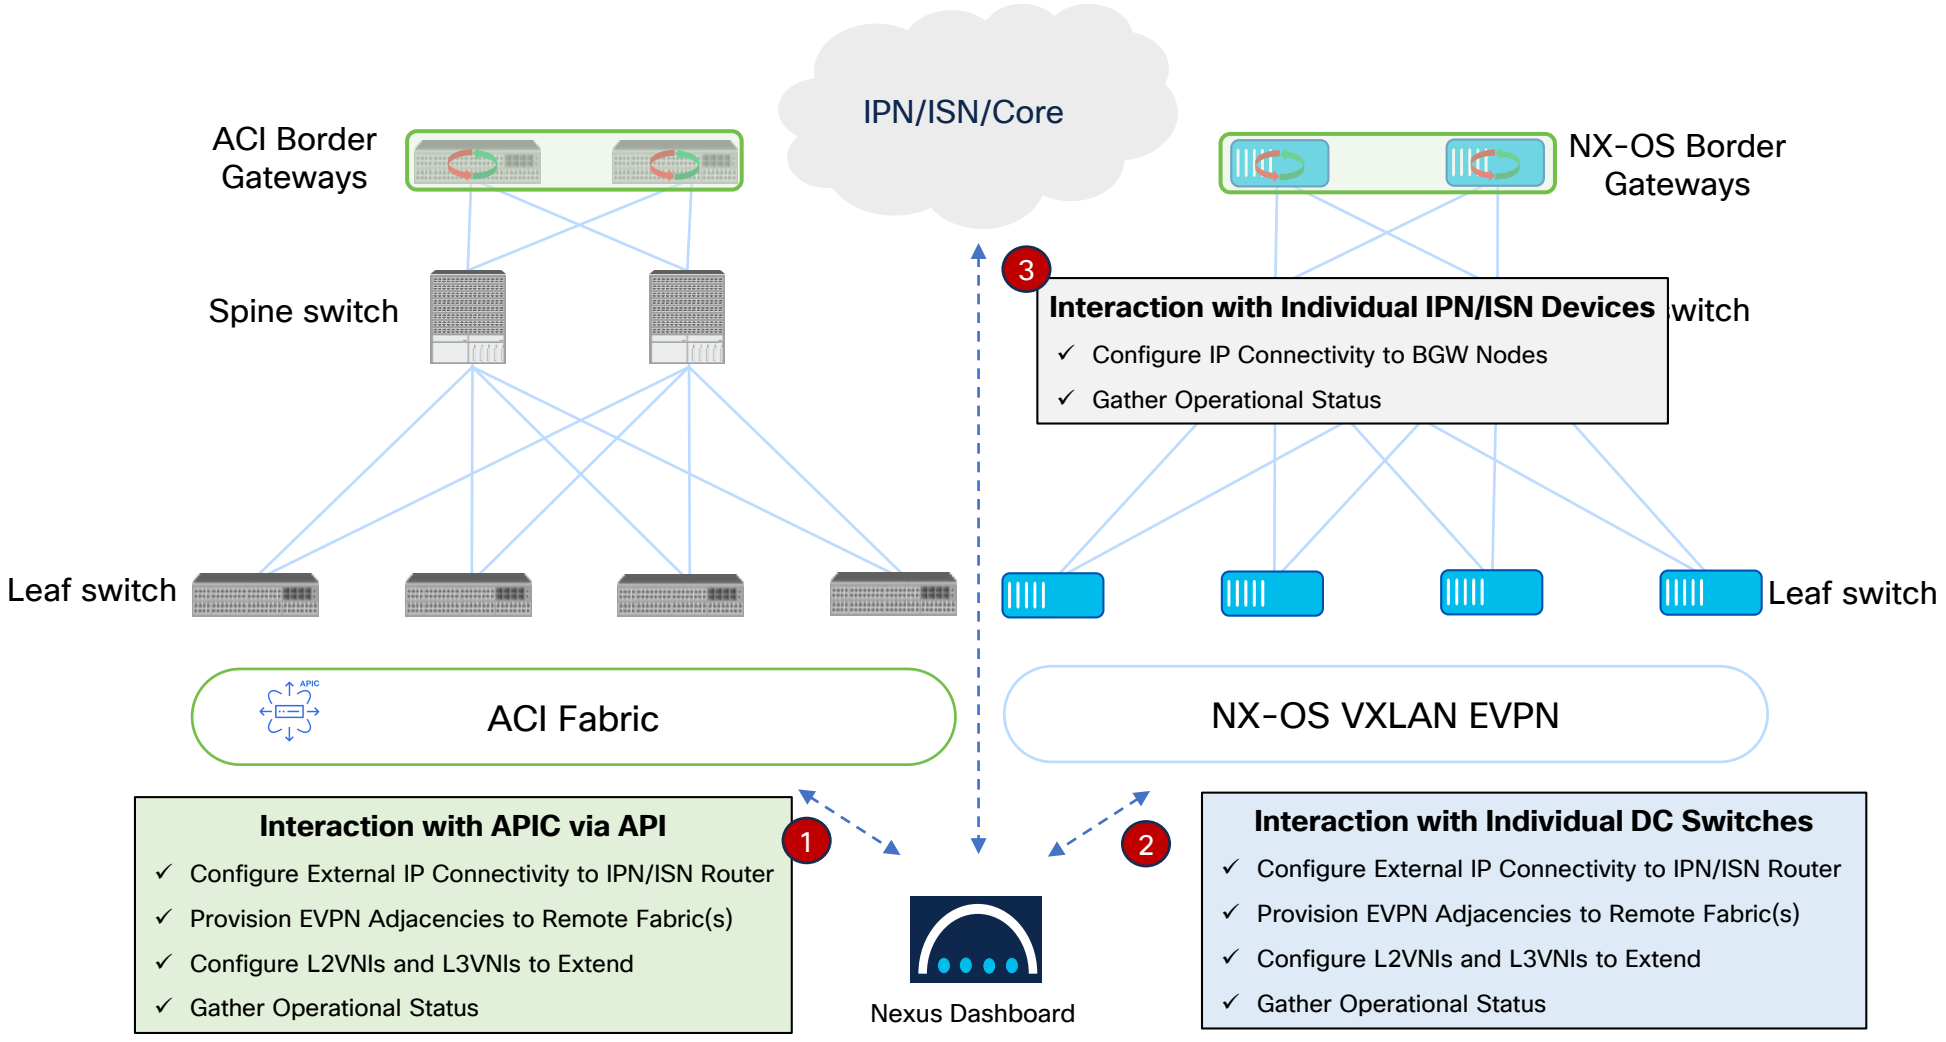

Conclusions

Conclusions

- Building distributed infrastructures is key to the deployment of resilient and scalable designs
- Cisco Nexus ONE Architecture aims to seamlessly interconnect and operate a mix of heterogeneous fabrics (ACI and VXLAN EVPN)
- The three main functional components of Cisco Nexus ONE 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

Recommended Sessions

Nexus ONE Architecture

- BRKDCN-2633: Deployment of Micro-Segmentation in Cisco NX-OS VXLAN EVPN Fabrics with VXLAN Group Policy Option (GPO)
- BRKDCN-2634: Deployment of VXLAN EVPN Gateways with Cisco ACI for the Interconnection of Heterogeneous Data Center Fabrics
- PSODCN-1009: Experience the Unified Nexus Dashboard

DC Design Evolution

- BRKDCN-2999: Multi-Tier Fabrics: Network Designs for the Modern Data Center

NX-OS EVPN VXLAN Multi-Site

- BRKDCN-2913: VXLAN BGP EVPN Multi-Site

ACI Multi-Pod/Multi-Site

- BRKDCN-2949: Cisco ACI Multi-Pod Design and Deployment
- BRKDCN-2980: ACI Multi-Site Architecture and Deployment

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 2026.

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

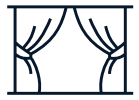

Visit the Cisco Showcase for related demos

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

Contact me at: Insert preferred comms method

Thank you

cisco Live !


CISCO