The bridge to possible

Deployment of Micro-Segmentation in Cisco NX-OS VXLAN EVPN Fabrics with VXLAN Group Policy Option (GPO)

Max Ardica, Distinguished TME
@maxardica

BRKDCN-2633

cisco Live!

Cisco Webex App

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-2633

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.

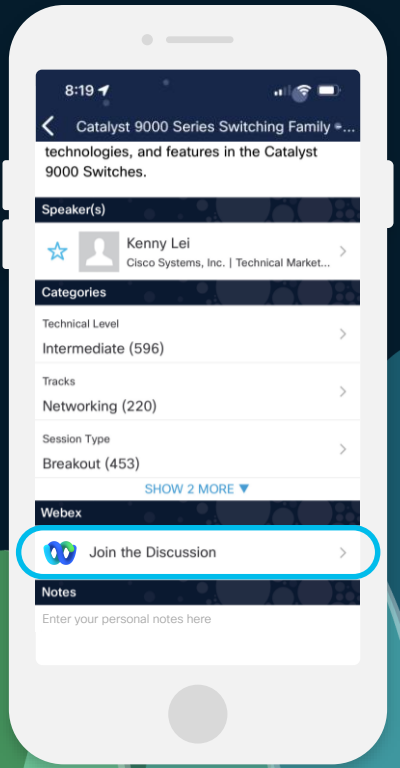

Agenda

- Introducing Cisco Nexus One Fabric Experience
- VXLAN GPO
  - VRF Modes of Operation
  - Classification and SGACLs
  - The Value of the Control Plane
  - VXLAN GPO and Multi-Site
- Secure Interconnection of Heterogeneous Fabrics

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

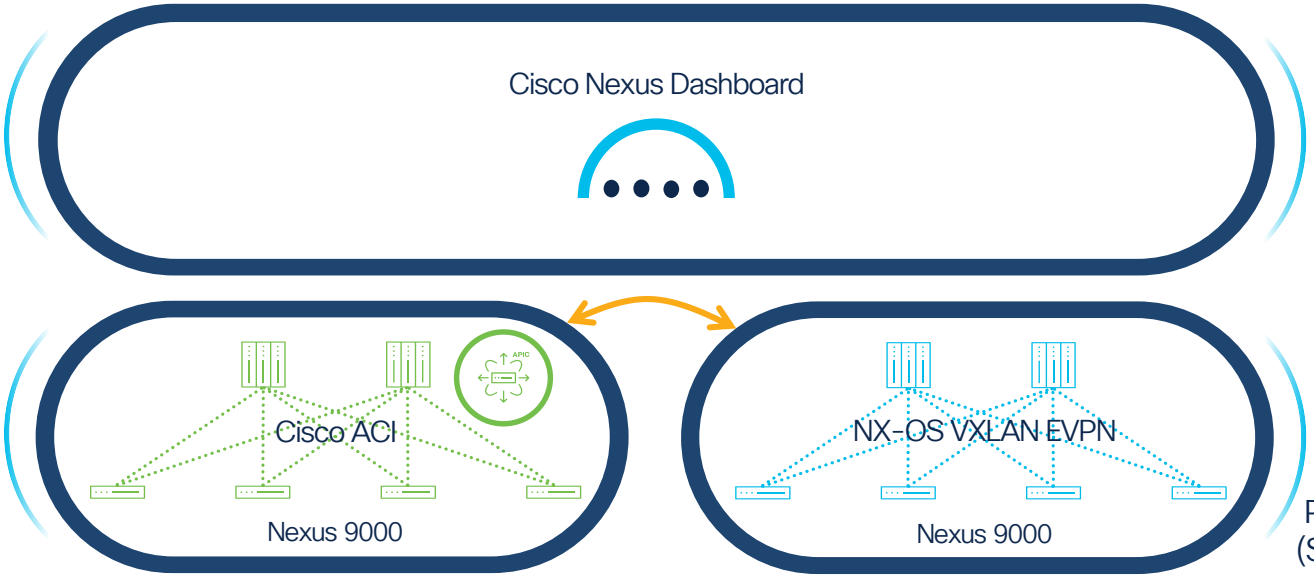

1 ACI VXLAN EVPN Border Gateways

2 Policy in NX-OS (Security Groups)

Different fabric architectures | Same outcome with common experience

VXLAN GPO
Introduction

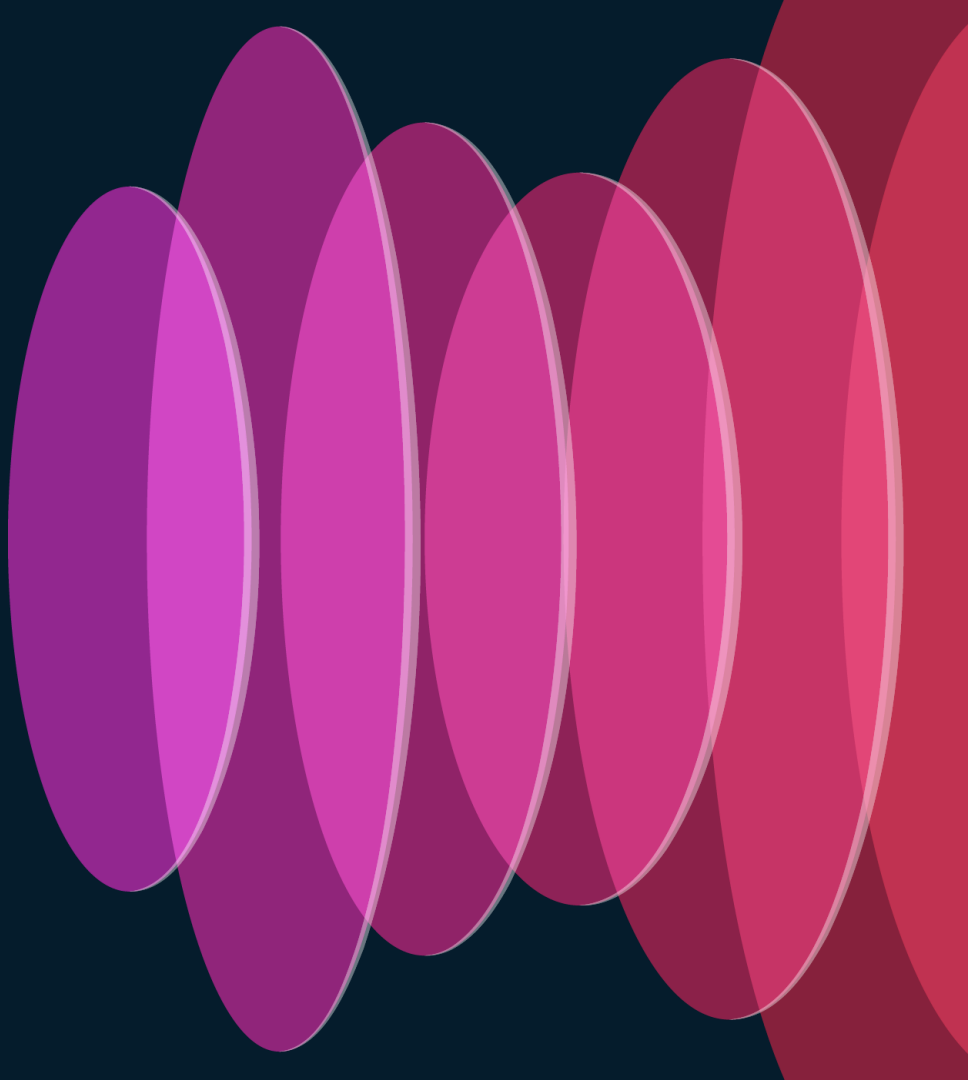

cisco Live!

VXLAN GPO with NX-OS

For More Information on VXLAN GPO with NDFC
BRKDCN-2629

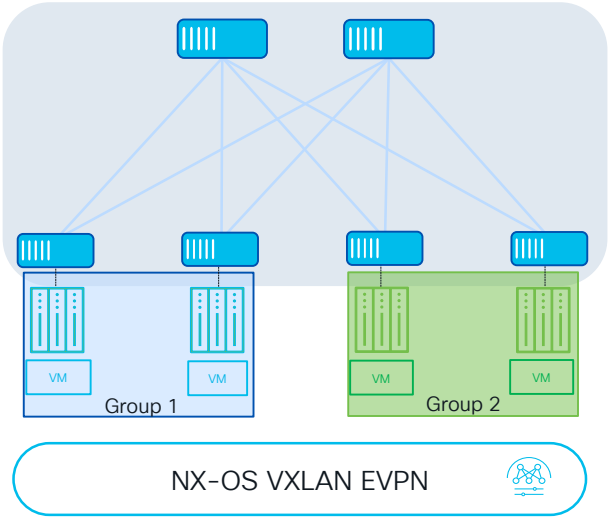

VXLAN GPO with NX-OS

- Group Policy Option 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

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)

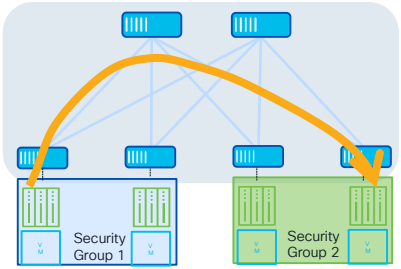

Service Chaining

- 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)

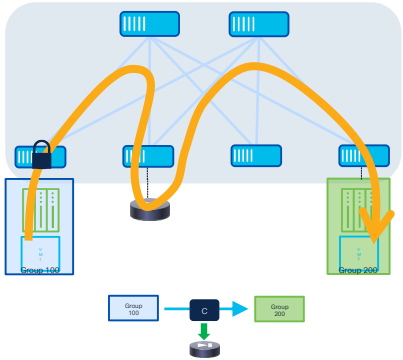

Hardware, Software and Licensing Support

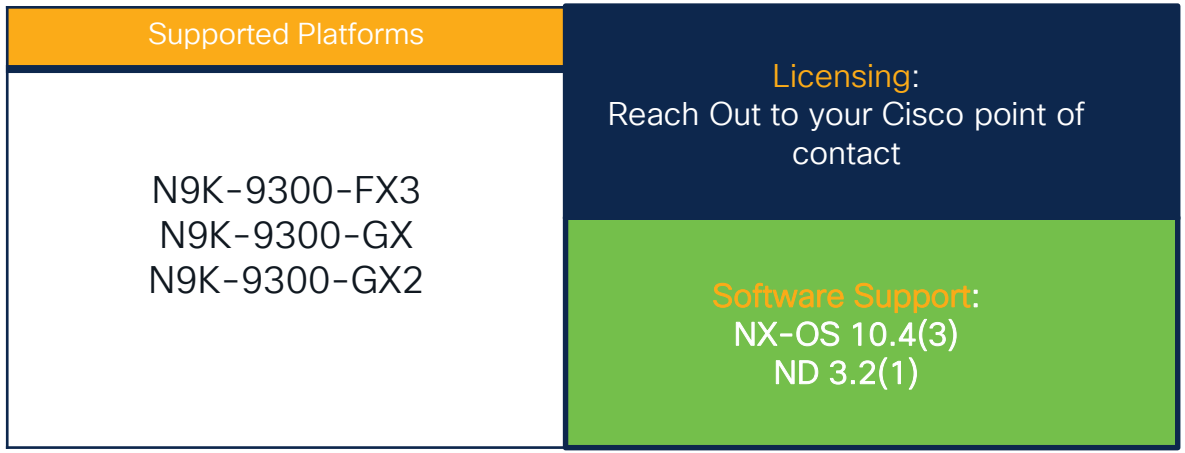

VXLAN GPO with NX-OS
Cisco GPO Data Plane and Control Plane Functionalities

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

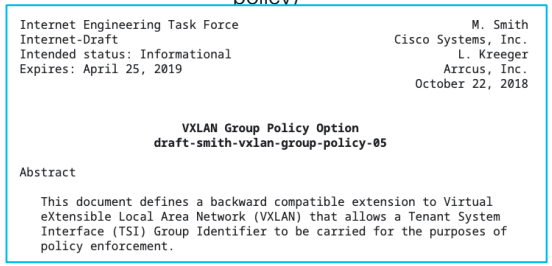

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: April 25, 2019

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)

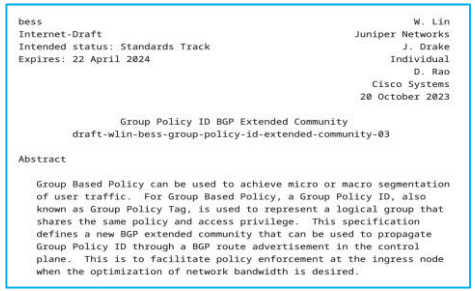

bess
Internet-Draft
Intended status: Standards Track
Expires: 22 April 2024

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)

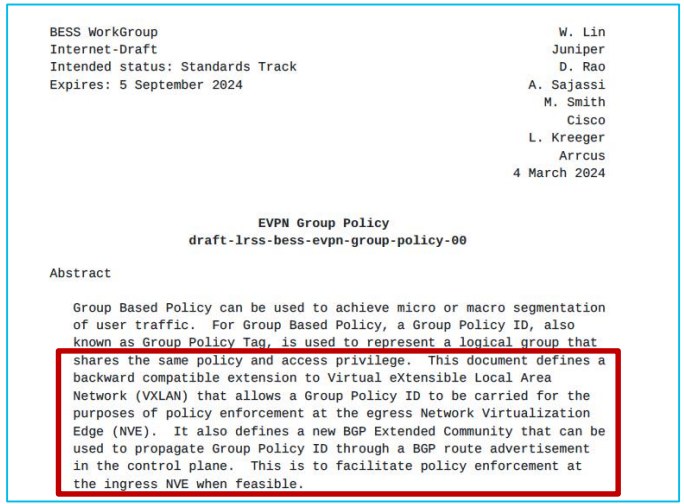

BESS WorkGroup
Internet-Draft
Intended status: Standards Track
Expires: 5 September 2024

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.

VXLAN GPO Header

VXLAN GPO (VXLAN Group Policy Option) as originally defined in draft-smith-vxlan-group-policy

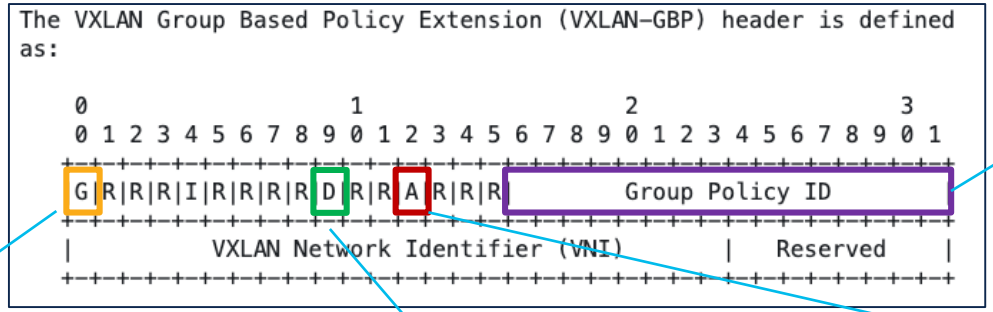

Group Policy ID

Security Group identifier

Group Based Policy Extension Bit

G = 1 indicates that the source Group membership is being carried within the Group Policy ID field

G = 0 indicates that the Group Policy ID is not being carried

Don't Learn Bit

D = 1 indicates that the egress VTEP MUST NOT learn the source address of the encapsulated frame

Policy Applied Bit (only relevant when G=1)

A = 1 indicates that the group policy has already been applied to this packet (the policy MUST NOT be applied by a device when the A bit is set)

A = 0 indicates that the group policy has not been applied to this packet (the policy MUST be applied by a device when the A bit is set to 0 and the destination Group can be determined)

VXLAN GPO
VRF Modes of Operation

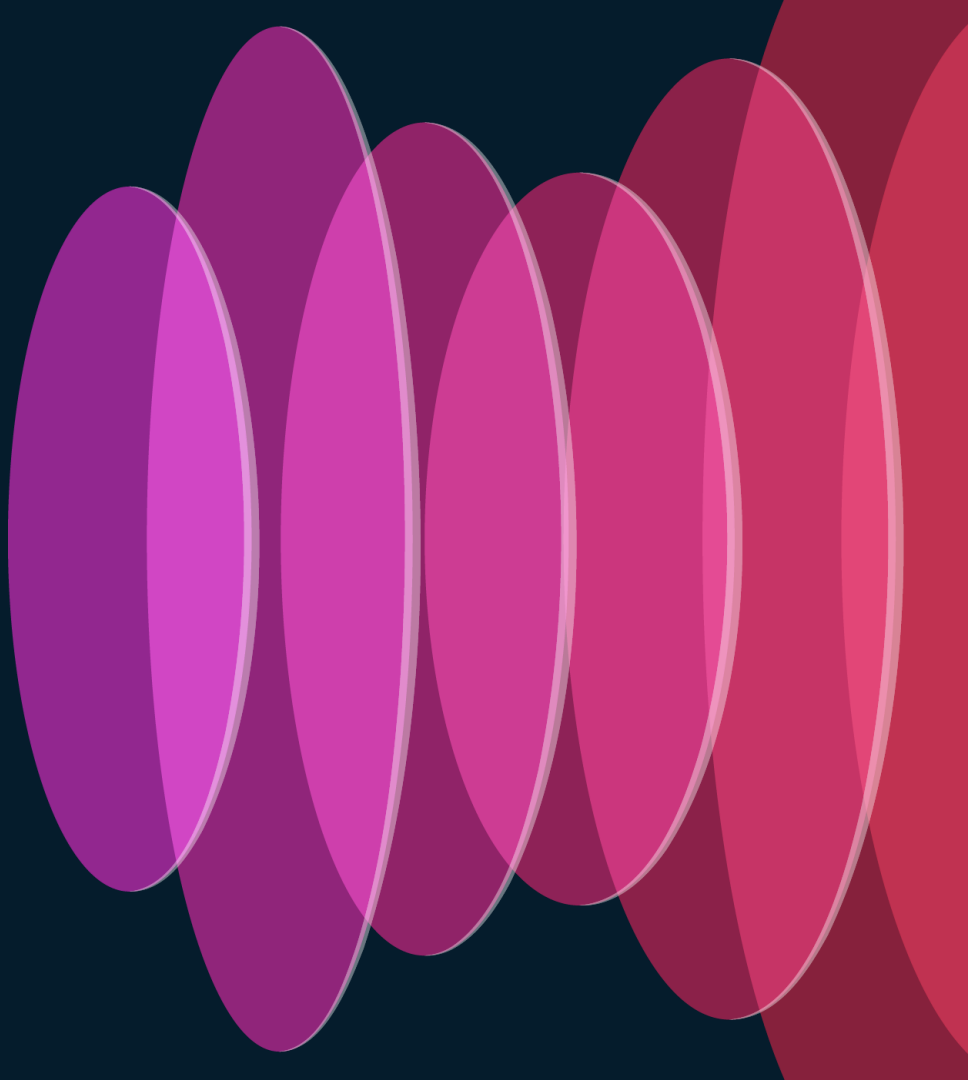

VRF Modes of Operation
VRF Unenforced

- Default VRF mode
- Can define Security Groups and associated rules to classify endpoints/prefixes in specific SGs
- SGACL contracts, even if configured, are not enforced in the VRF
- Not possible to verify if contracts applied between SGs are hit or not

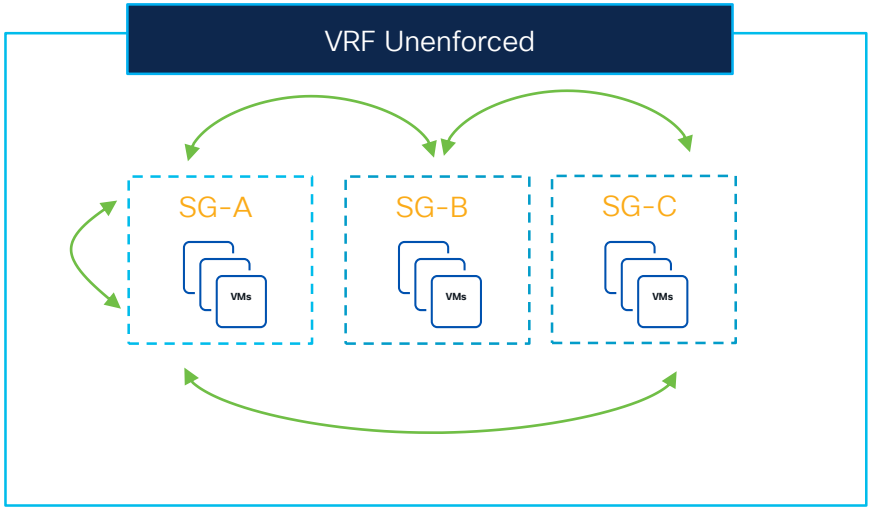

VRF Modes of Operation
VRF Enforced

Default-Permit Mode

- Open unicast communication between Security Groups (SGs) by default
- SGACLs must be applied to deny traffic between SGs

```
vrf context VRF1
  security enforce tag 17 default permit
```

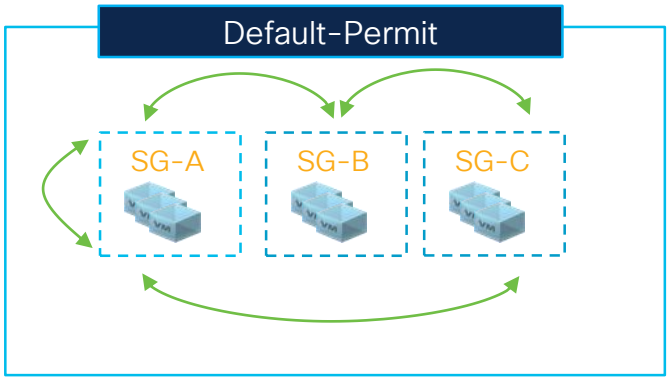

Default-Deny Mode

- No unicast communication between SGs by default
- SGACLs must be applied to allow traffic between SGs
- Zero Trust enforced

```
vrf context VRF1
  security enforce tag 17 default deny
```

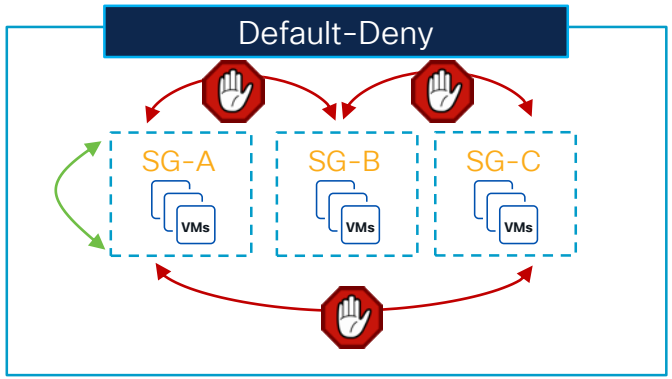

VXLAN GPO
Classification and SGACLs

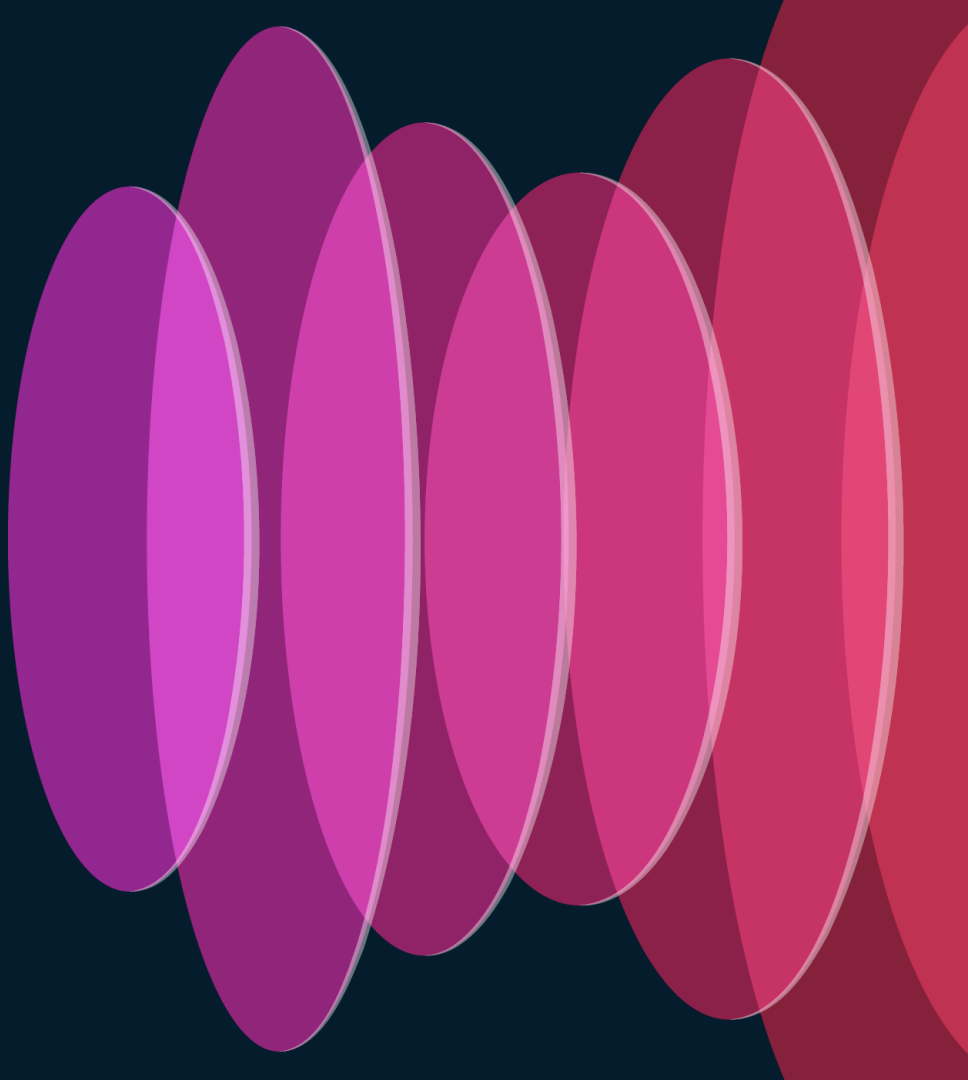

cisco Live!

VXLAN GPO with NX-OS
Classification Criteria and SGACL Actions

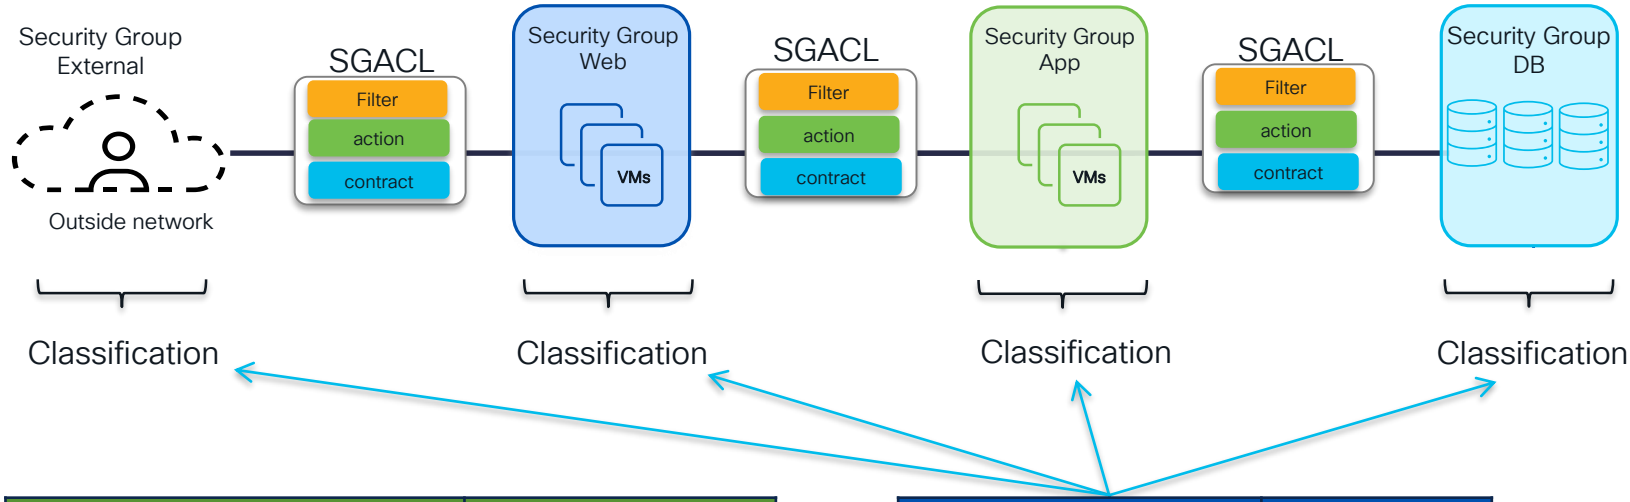

SGACL Action	NX-OS Version
Permit, Permit + Log	10.4(3)
Deny, Deny + Log	10.4(3)
Redirect (FW service only)	10.4(3)
Redirect (SLB, Service-Chain)	Roadmap

Security Group Attributes	NX-OS Version
IP Prefix	10.4(3)
VLAN	10.4(3)
Port + VLAN	Roadmap
VM Tags/Attributes	Roadmap

Classification is done assigning a fabric-wide unique tag (valid range: 16-65535)

VXLAN GPO with NX-OS
Connected Endpoints Classification

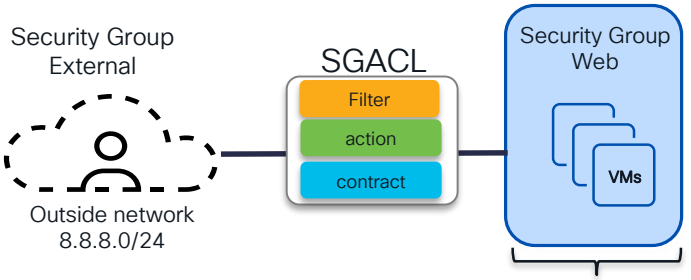


```
SW(config)# security-group 2000 name webserver
SW(config-security-group)# match connected-endpoints vrf vrf1 ipv4 10.1.1.141/32
SW(config-security-group)# match connected-endpoints vrf vrf1 ipv4 10.2.0.0/24
SW(config-security-group)# match vlan 10
```

- Endpoints internally connected to the fabric leaf nodes can be classified:
  - With host-level granularity (/32 or /128)
  - Using a less specific prefix, including a 0.0.0.0/0 ‘catch-all’ entry covering all the internal subnets in a VRF
  - Endpoints that are not classified by a specific prefix get assigned a configured global VRF tag
- The “match vlan” option ensures that all traffic received from/destined to hosts in that VLAN on a given switch is classified to the SG
  - In the future it will be supported to match a VLAN tag with per-port granularity

VXLAN GPO with NX-OS
External Subnets Classification

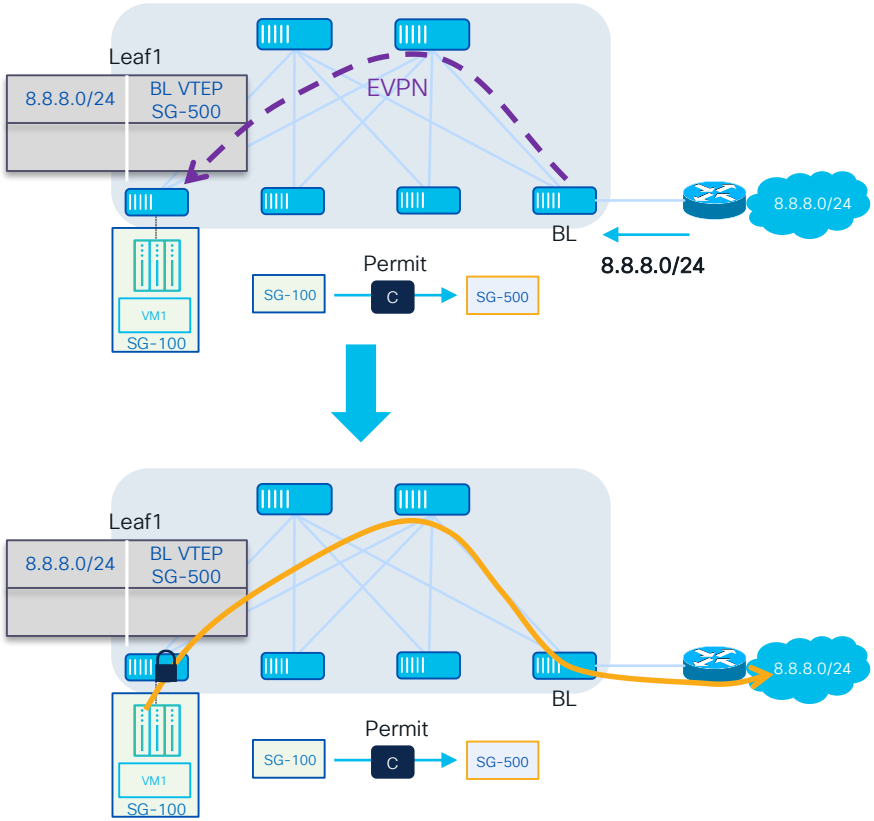

Classification on Border Leaf Node

```
SW(config)# security-group 500 name webclient
SW(config-security-group)# match external-subnets vrf vrf1 ipv4 8.8.8.0/24
```

- Leaf 1 receives the 8.8.8.0/24 prefix **matching** the classification subnet
- The 8.8.8.0/24 prefix is advertised into the fabric with the associated SG-500 tag
- Policy can be enforced between the endpoint in SG-100 and the external prefix based on the configured contract

VXLAN GPO with NX-OS
External Subnets Classification

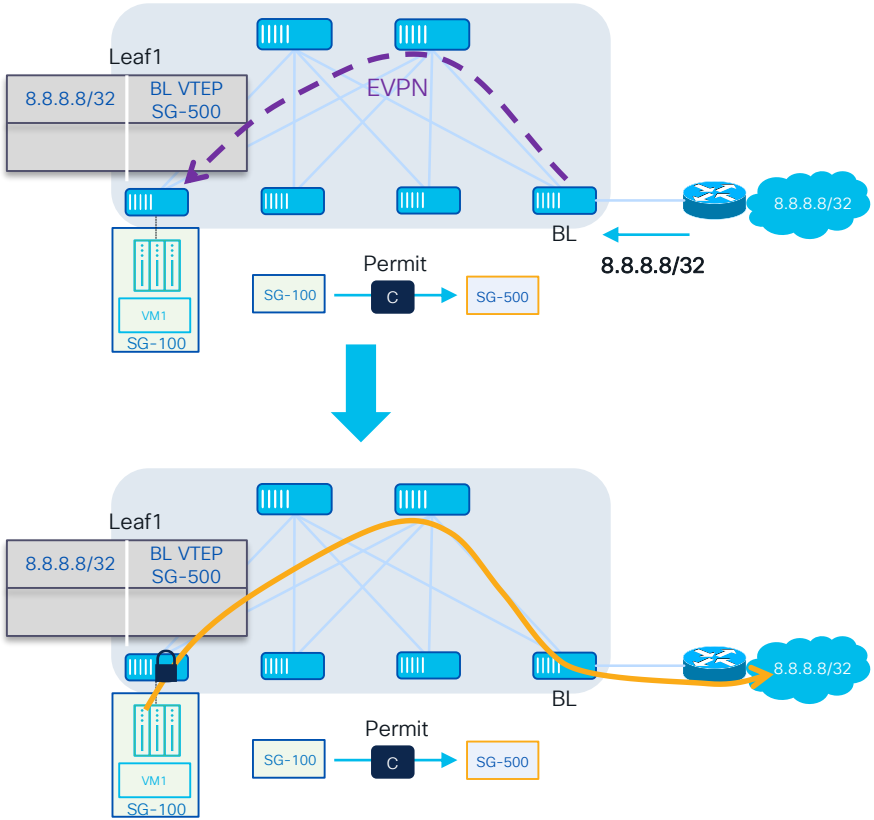

Classification on Border Leaf Node

```
SW(config)# security-group 500 name webclient
SW(config-security-group)# match external-subnets vrf vrf1 ipv4 8.8.8.0/24
```

- Leaf 1 receives the specific 8.8.8.8/32 prefix **covered** by the configured 8.8.8.0/24 classification subnet
- The 8.8.8.8/32 prefix is advertised into the fabric with the associated SG-500 tag
- Policy can be enforced between the endpoint in SG-100 and the external 8.8.8.8 destination based on the configured contract

VXLAN GPO with NX-OS
External Subnets Classification (2)

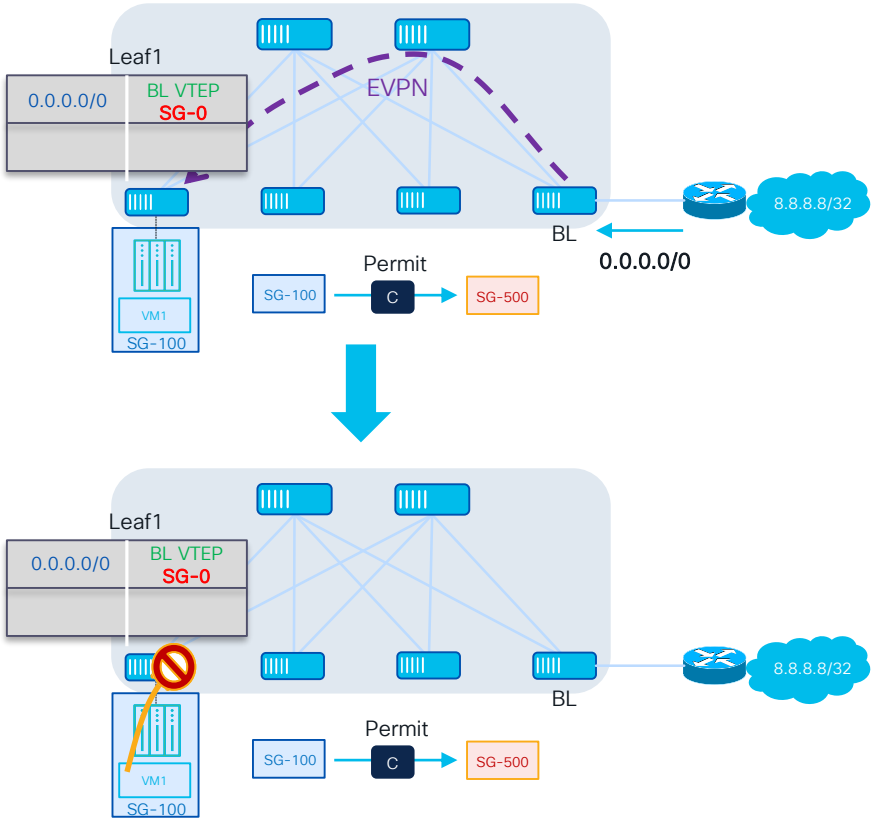

Classification on Border Leaf Node

```
SW(config)# security-group 500 name webclient
SW(config-security-group)# match external-subnets vrf vrf1 ipv4 8.8.8.0/24
```

OR

```
SW(config)# security-group 500 name webclient
SW(config-security-group)# match external-subnets vrf vrf1 ipv4 8.8.8.8/32
```

- Leaf 1 receives the generic 0.0.0.0/0 prefix and only more specific classification subnets are configured on the Border Leaf node
- The 0.0.0.0/0 prefix is advertised into the fabric without any associated SG tag
- Policy cannot be enforced between the endpoint in SG-100 and the external prefix 8.8.8.8 and traffic is dropped (assuming a ”default deny” VRF configuration)

VXLAN GPO with NX-OS
Creation of Bidirectional SGACLs (Contracts)

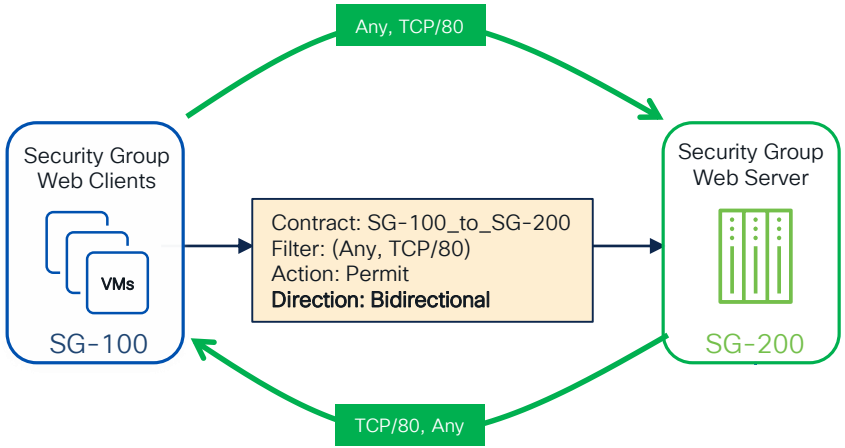

Return rule automatically created

- When the VRF is enforced in “default deny” mode, communication between different SGs is denied in absence of contracts
- A contract has a name and one (or more) rules with an associated action (permit, deny, redirect, etc.)
- Each rule should be defined with “**Bidirectional**” direction, to ensure that a two-way communication can be established (without requiring the definition of separate rules)

Rules

Direction	Action*	Protocol*	Match Summary
bidirectional	permit	http	IP TCP dport:80

VXLAN GPO
The Value of the Control Plane

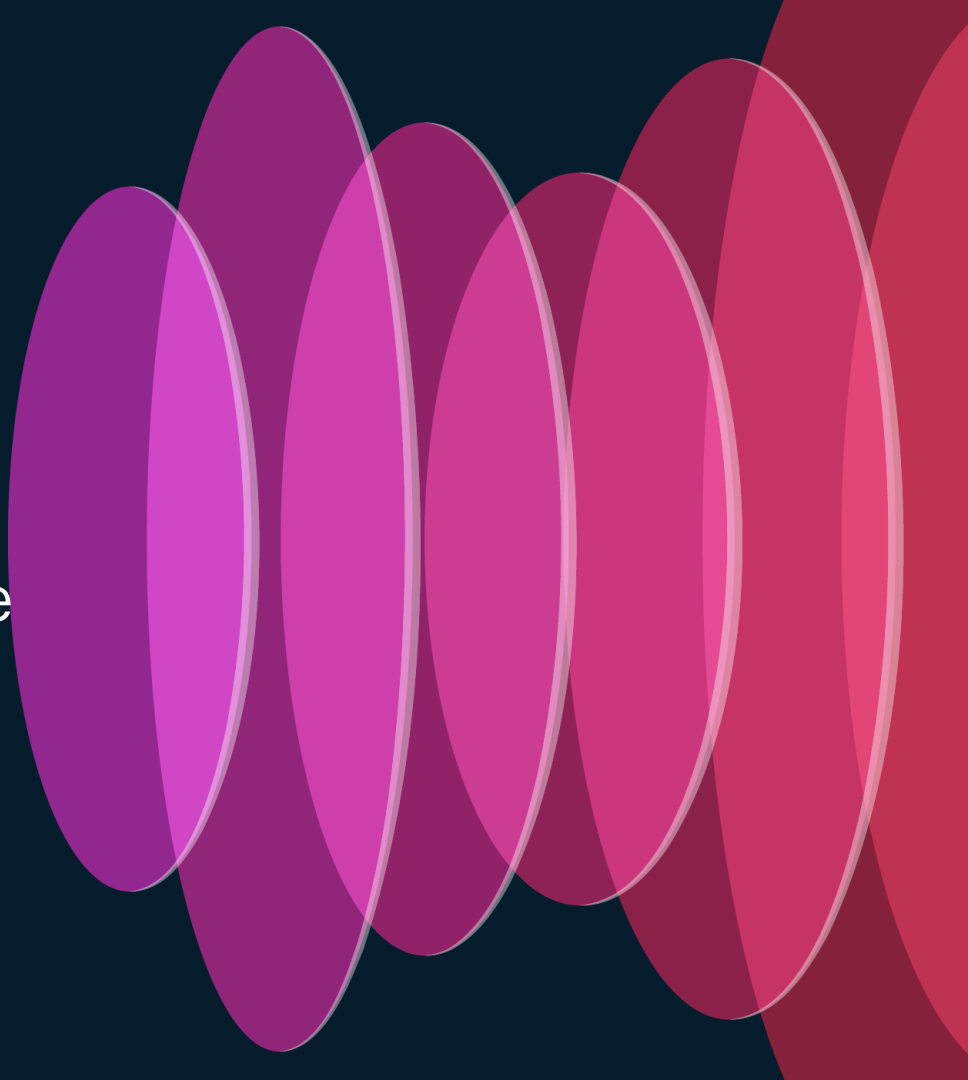

cisco Live!

VXLAN GPO with NX-OS
Egress Enforcement Only Possible without SG Info in Control Plane

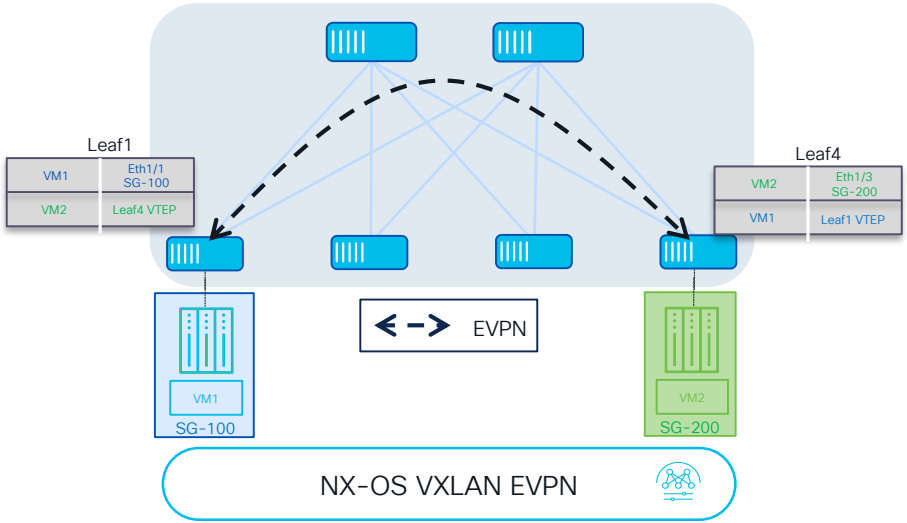

- Use of MP-BGP EVPN to propagate endpoints connectivity without policy information

VXLAN GPO with NX-OS
Egress Enforcement Only Possible without SG Info in Control Plane

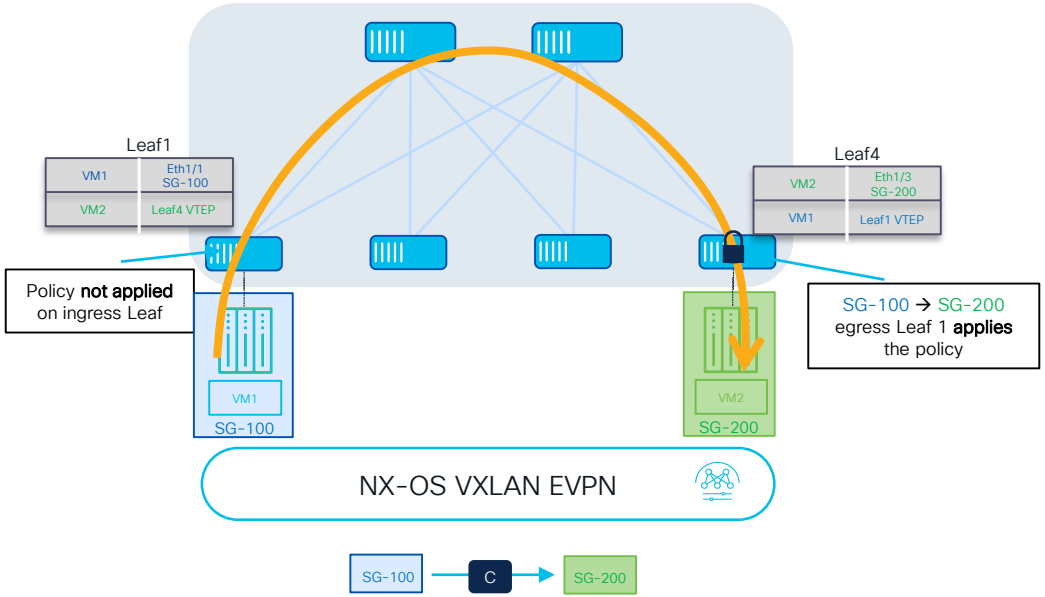

- Policy enforcement not possible on the ingress leaf node because missing info of the destination SG
- Egress leaf can apply the policy as the source SG is carried with the data packet

VXLAN GPO with NX-OS
Use of MP-BGP EVPN for Exchanging Security Tags

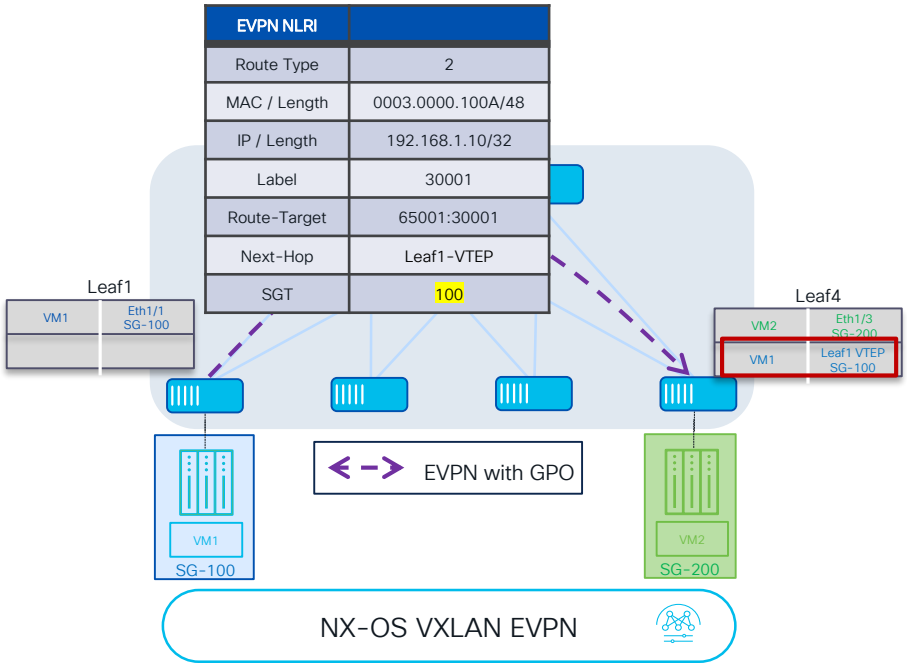

- Use of MP-BGP EVPN control plane to propagate endpoints connectivity and policy information inside the fabric
- The SGT information is propagated as a BGP extended community (Group Policy ID Extended Community)

VXLAN GPO with NX-OS
Use of MP-BGP EVPN for Exchanging Security Tags

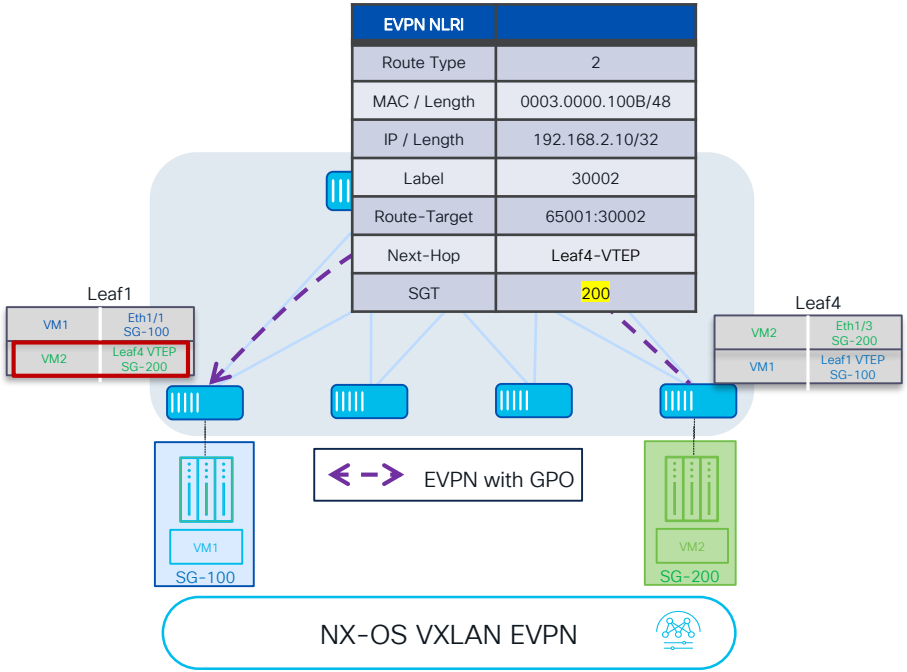

- Use of MP-BGP EVPN control plane to propagate endpoints connectivity and policy information inside the fabric
- The SGT information is propagated as a BGP extended community (Group Policy ID Extended Community)

VXLAN GPO with NX-OS
Use of MP-BGP EVPN for Exchanging Security Tags

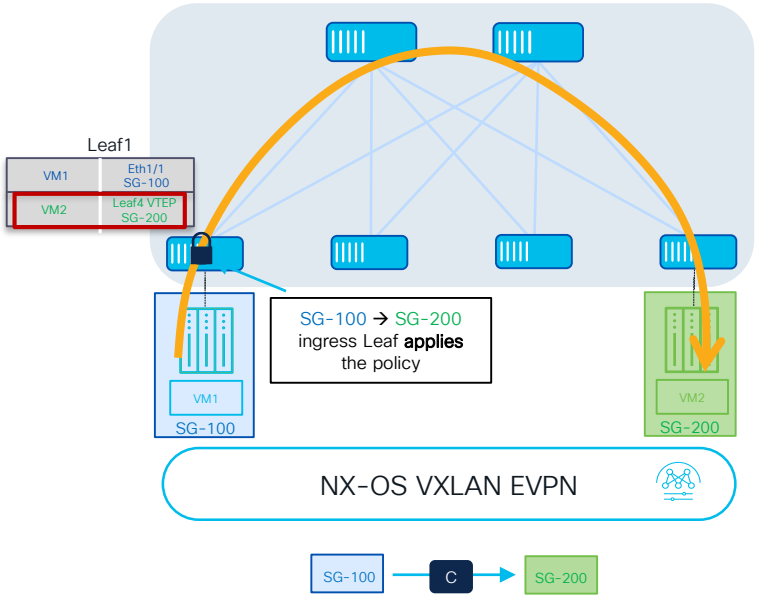

- Facilitate the enforcement of policy on the ingress leaf node (for both directions)
- Security Group Access Control Lists (SGACLs/contracts) enforced between groups

VXLAN GPO with NX-OS
Traffic Steering with Policy Based Redirection

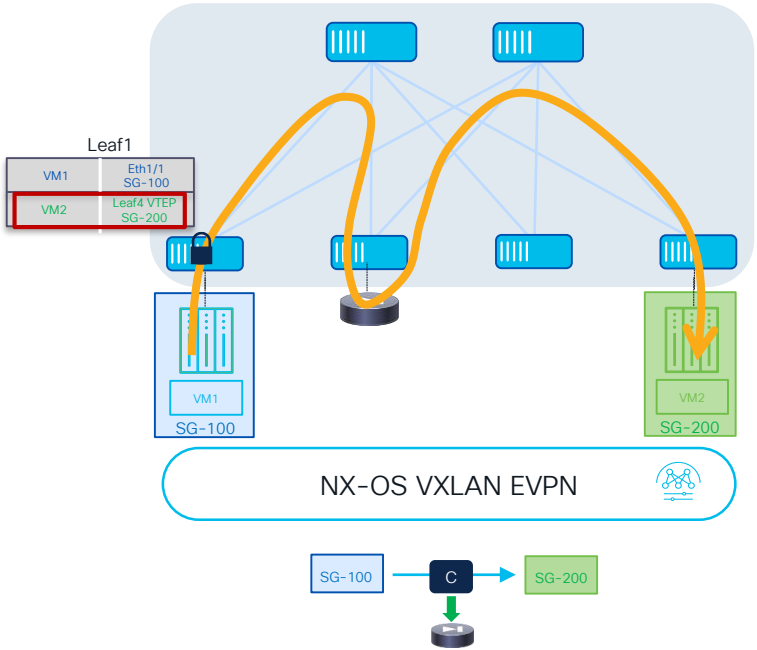

- Policy Based Redirection capabilities to steer through one or more service devices (firewall, load balancers, etc.) traffic flows between different security groups
- Redirection to a Firewall service function with NX-OS 10.4(3)F
- Other use cases, including traffic stitching through multiple services,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

VXLAN GPO

VXLAN GPO and Multi-Site

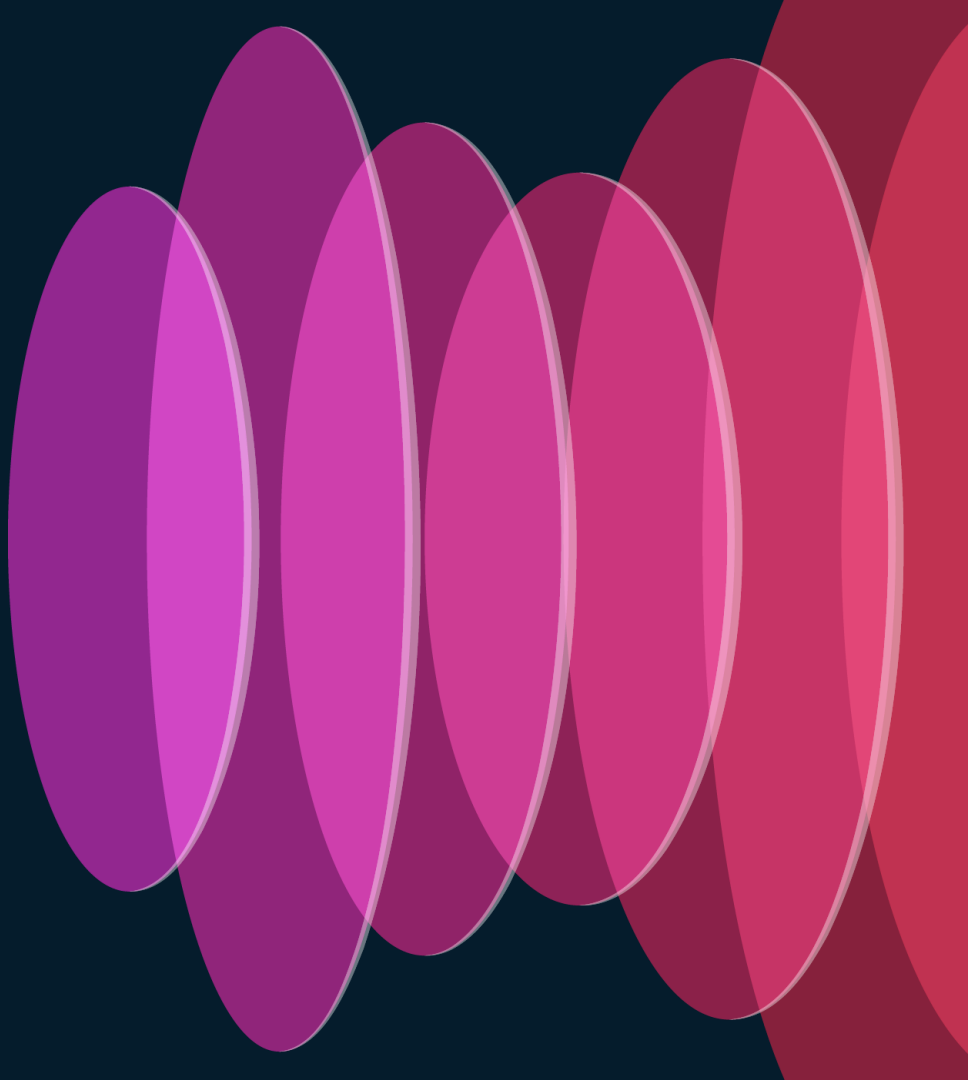

cisco Live!

VXLAN EVPN Multi-Site
Functional Components

For More Information on
VXLAN Multi-Site
BRKDCN-2913

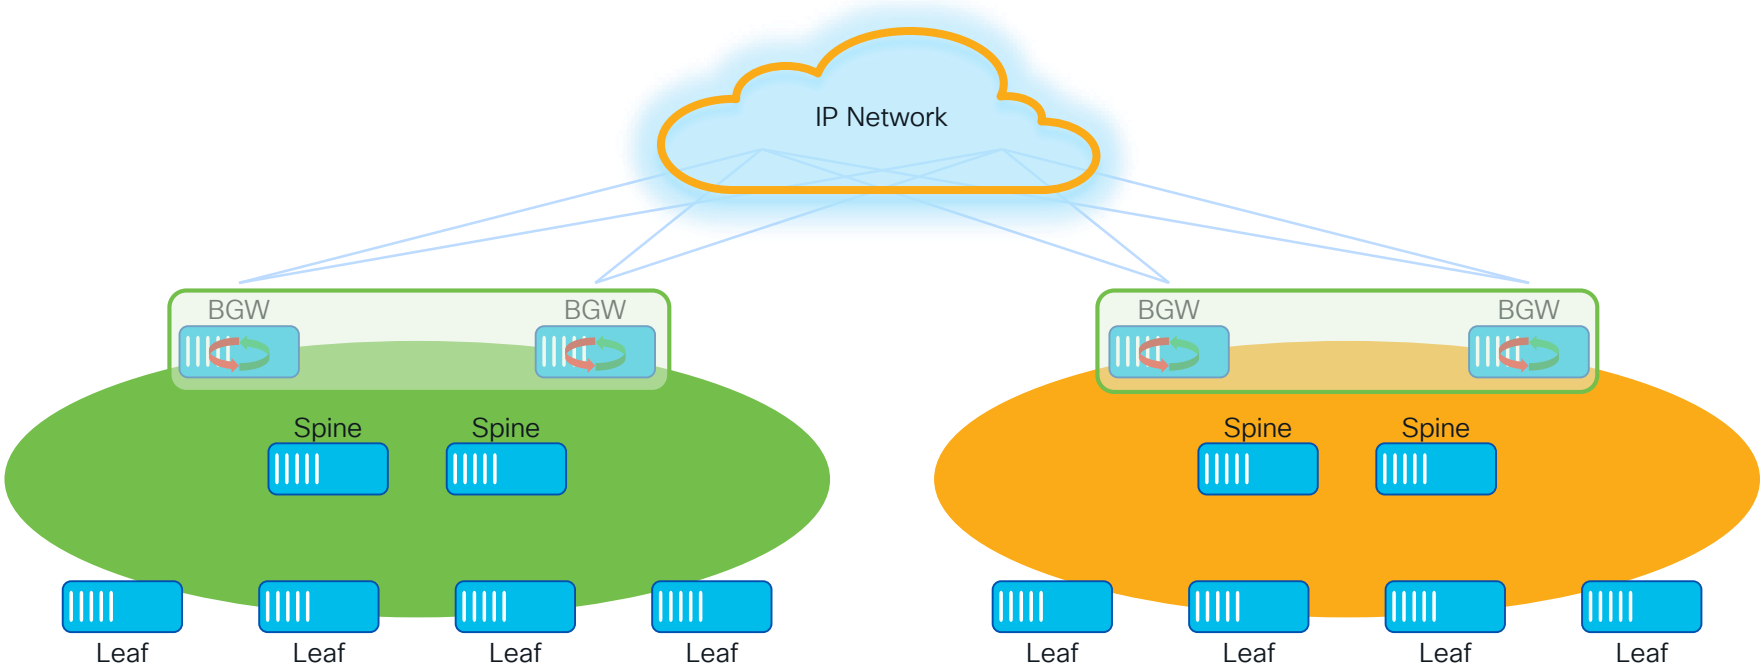

draft-sharma-bess-multi-site-evpn

VXLAN GPO with Multi-Site
Deployment Considerations

NX-OS 10.4(3)F

- Anycast BGWs and vPC BGWs are both supported with the Security Group feature
- All Security Group tags belong to a global namespace valid across multiple sites
  - Only symmetric namespace across fabrics is supported with NX-OS 10.4(3) release
- Support for connectivity and policy extension between policy aware sites but also with policy unaware fabrics
- All policy unaware sites endpoints or external prefix routes are mapped to a single global tag (“vxlan-evpn-sg” tag = 15)
  - The “vxlan-evpn-sg” tag is allowed in contract’s CLI to apply security policies to traffic originated from (or destined to) policy unaware fabrics
- Service redirection with Multi-Site planned for a future NX-OS release

VXLAN GPO with Multi-Site

Policy Aware to Policy Aware Fabrics

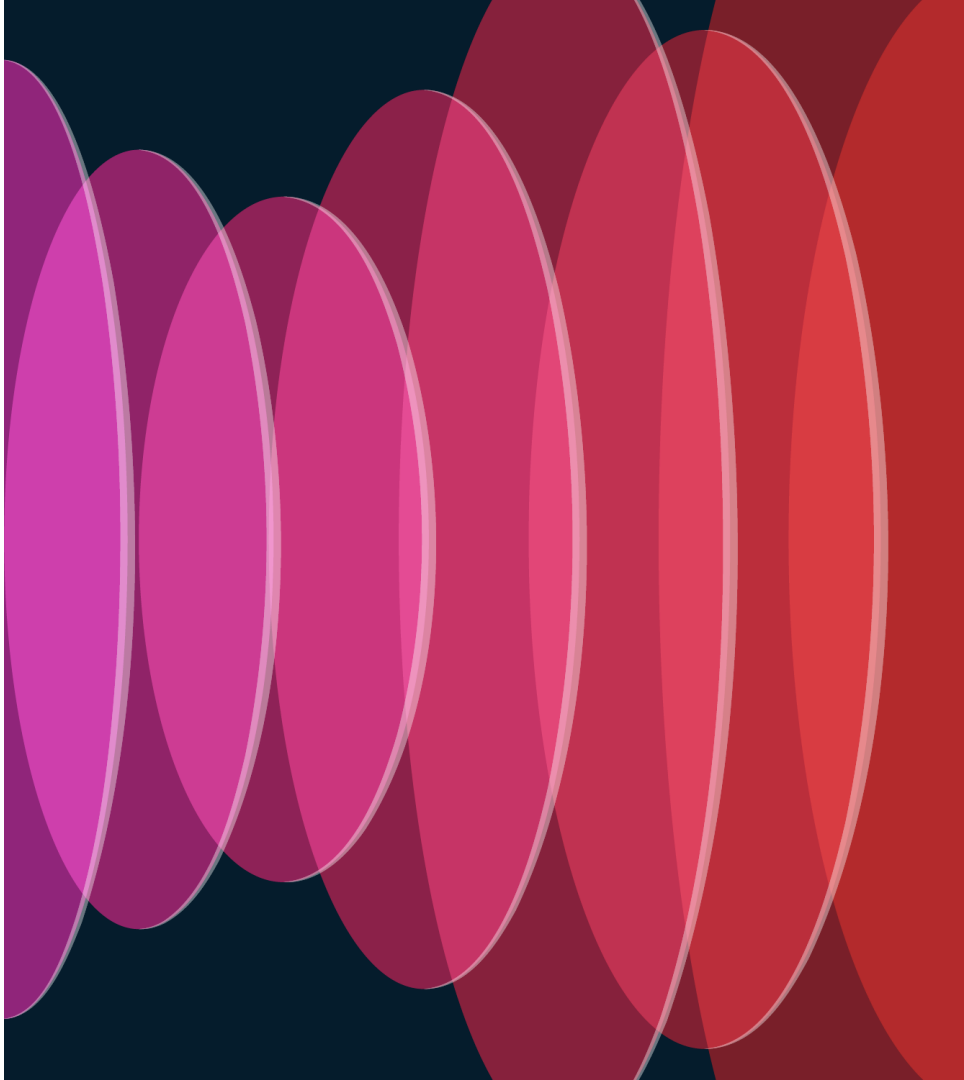

cisco Live!

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

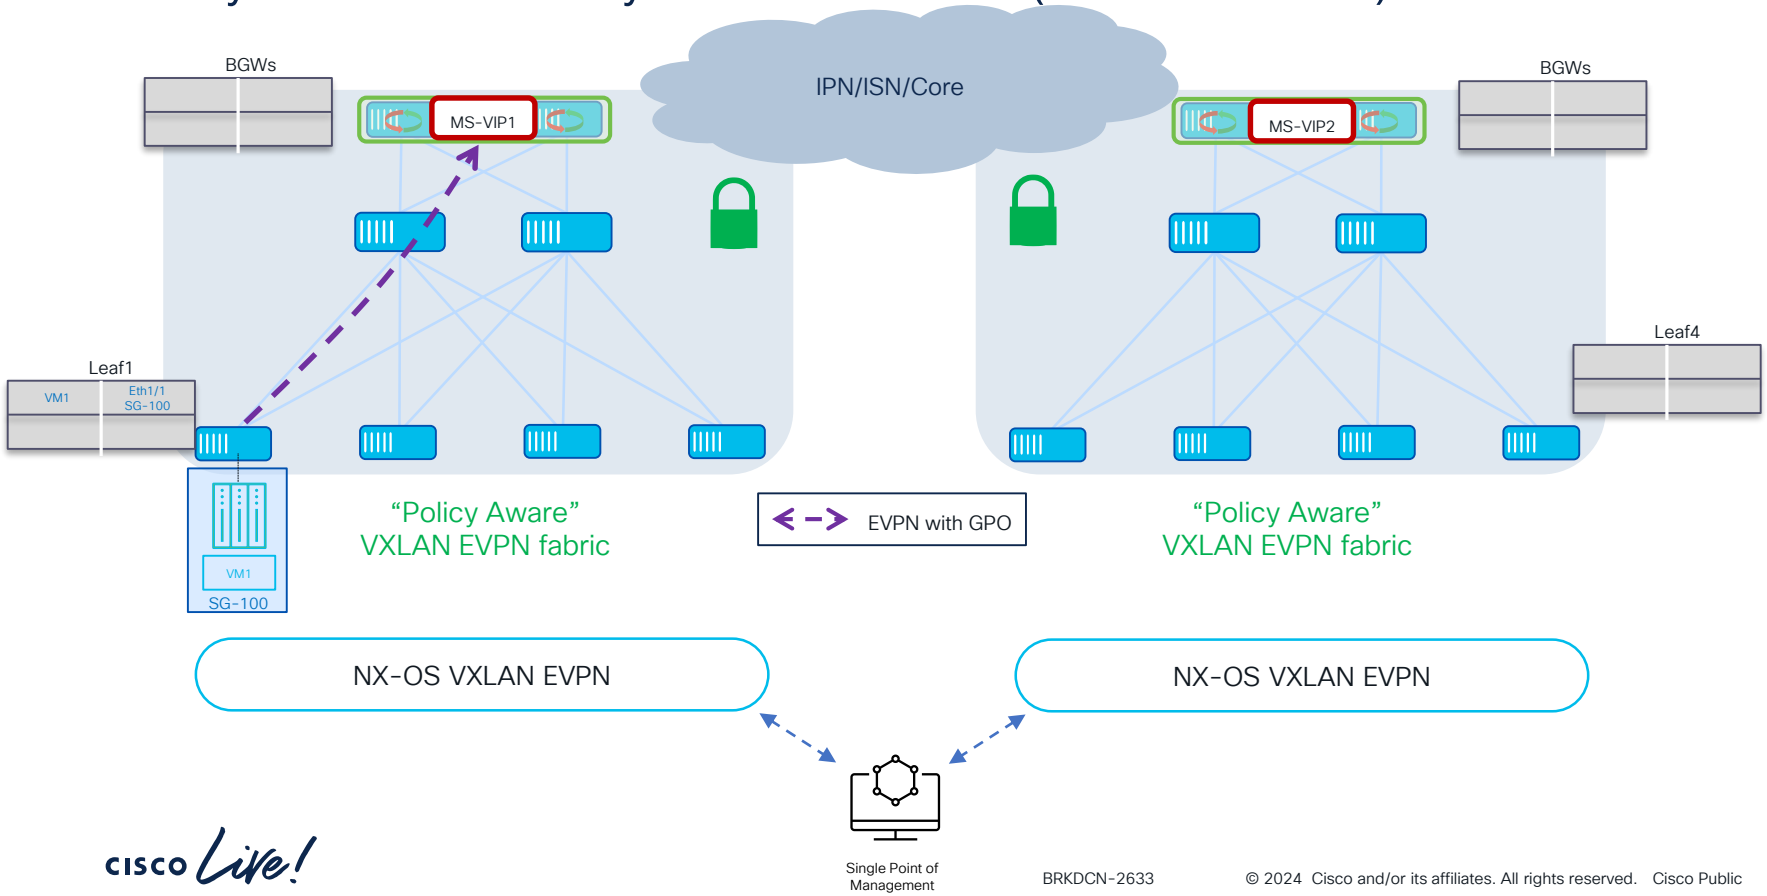

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

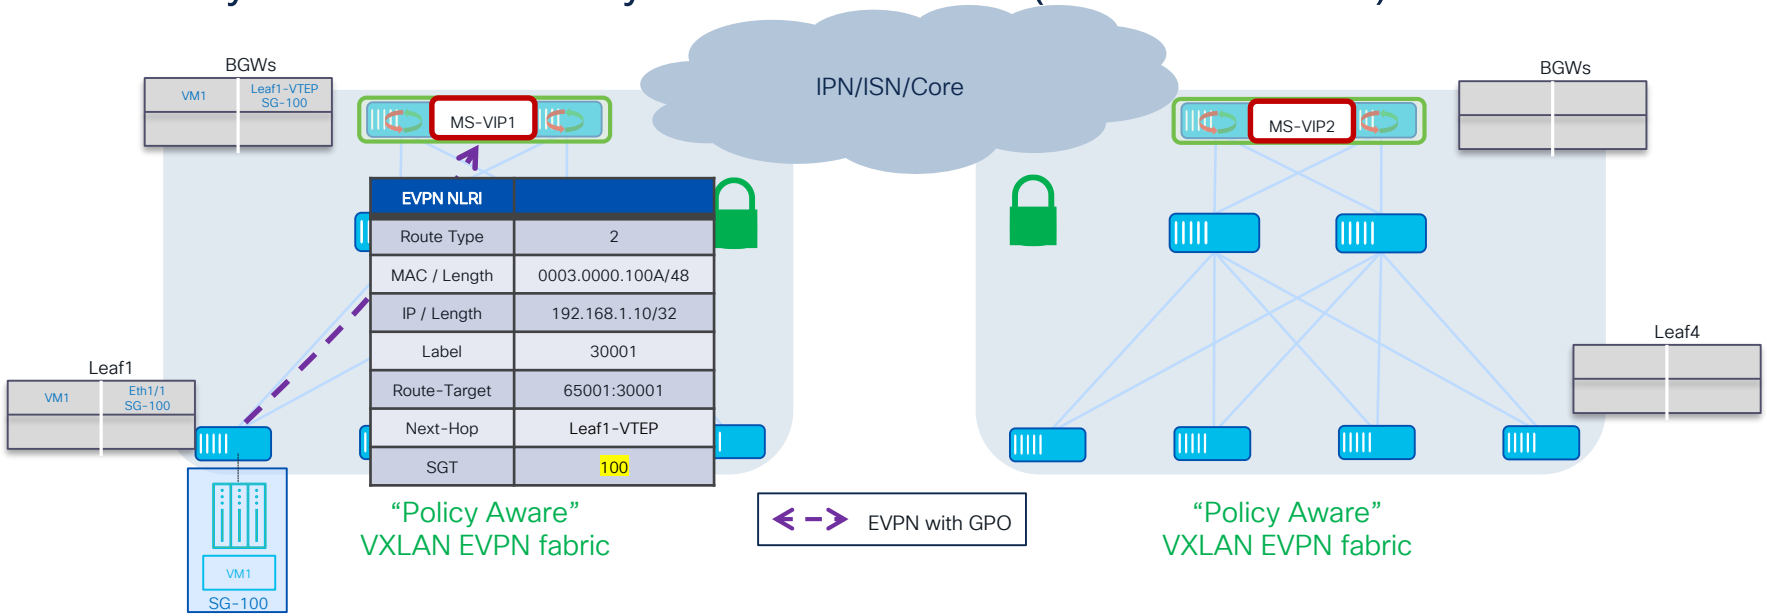

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

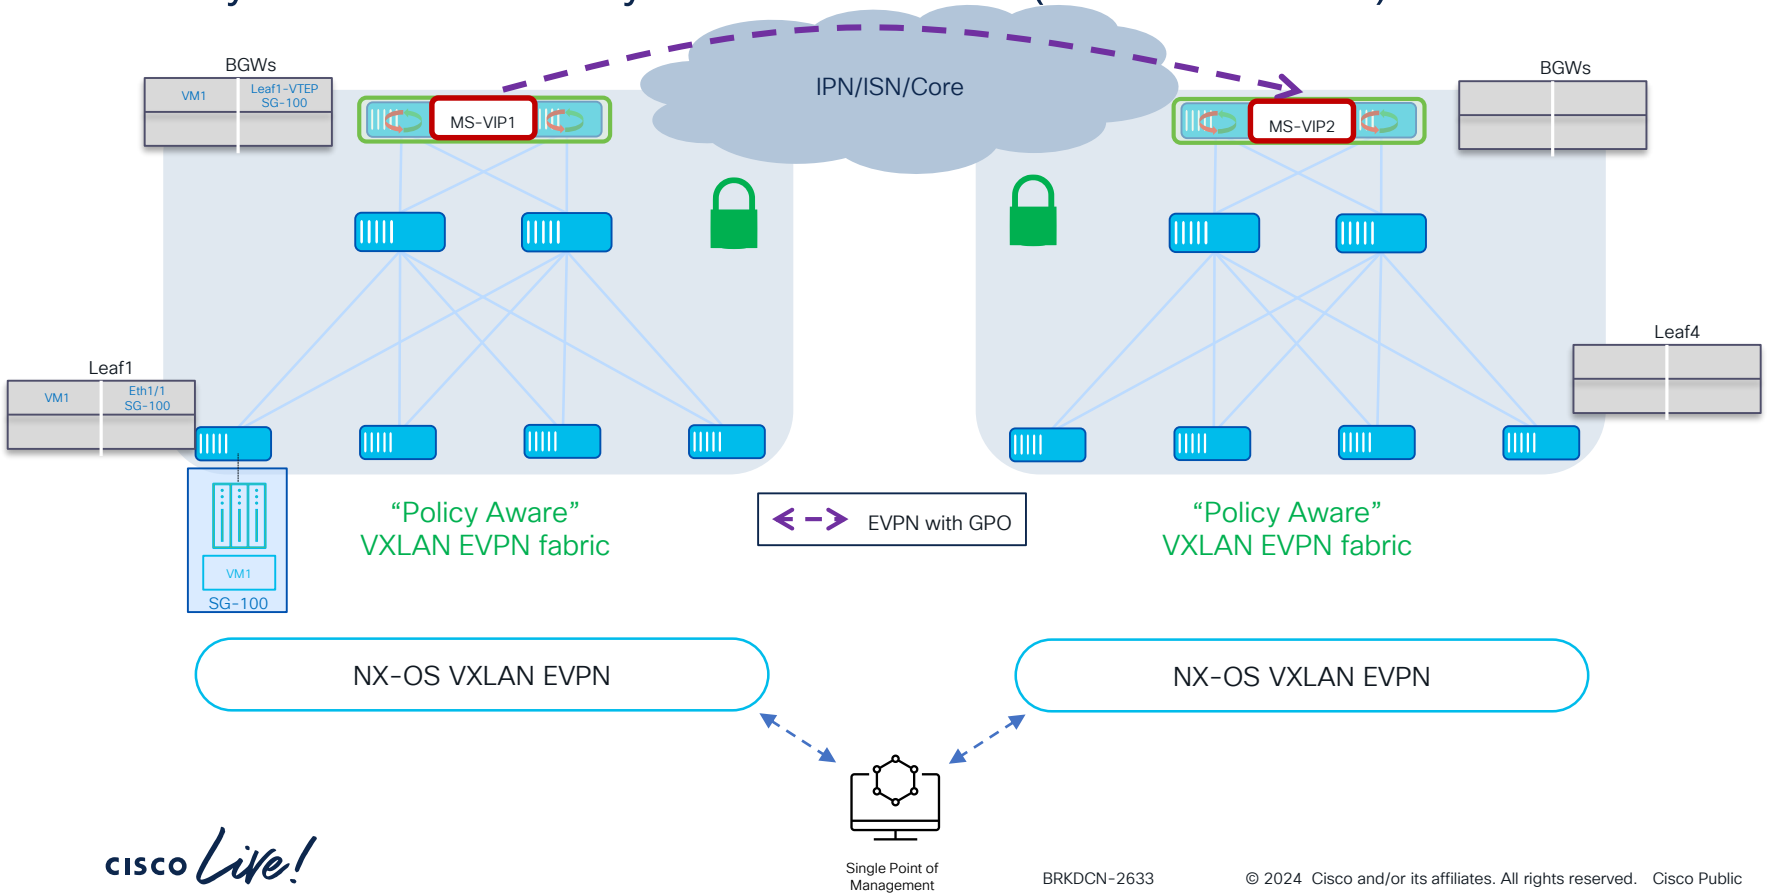

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

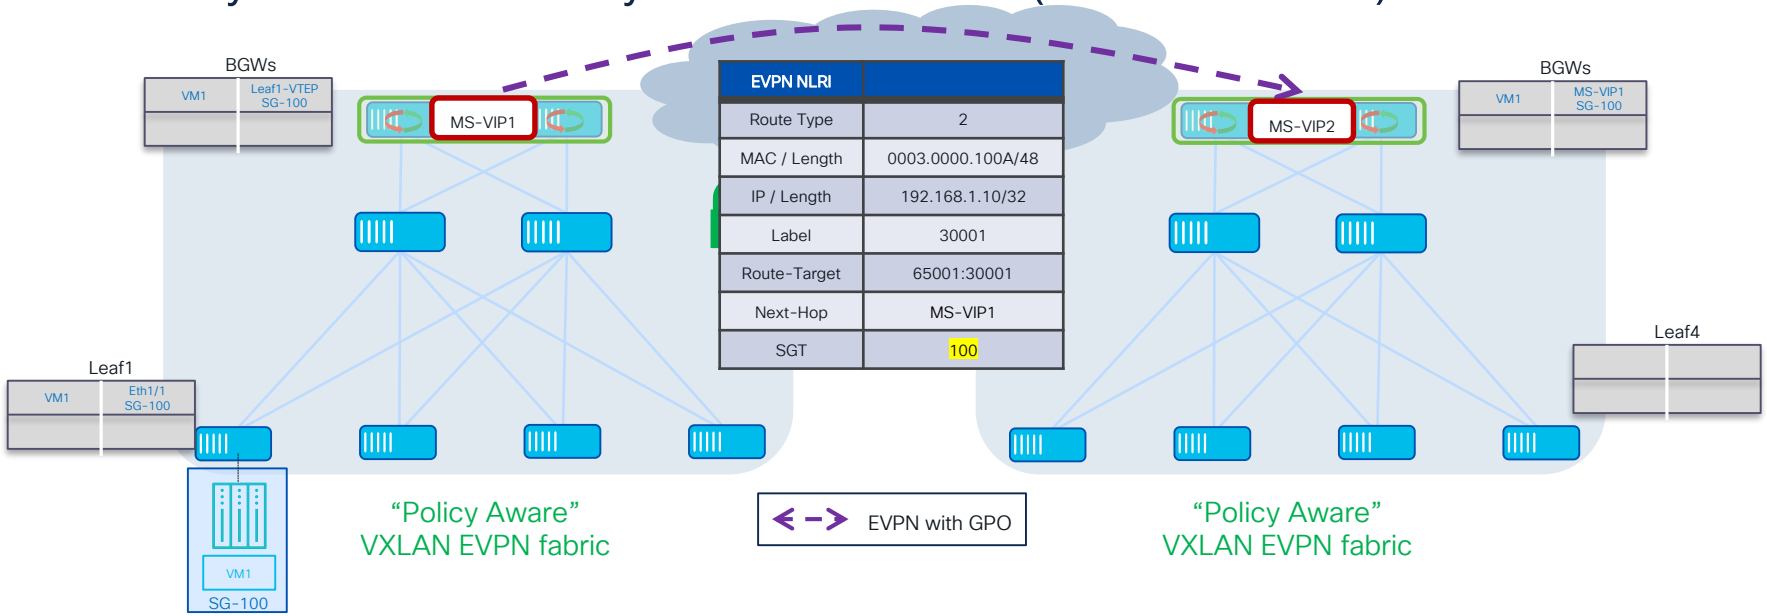

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

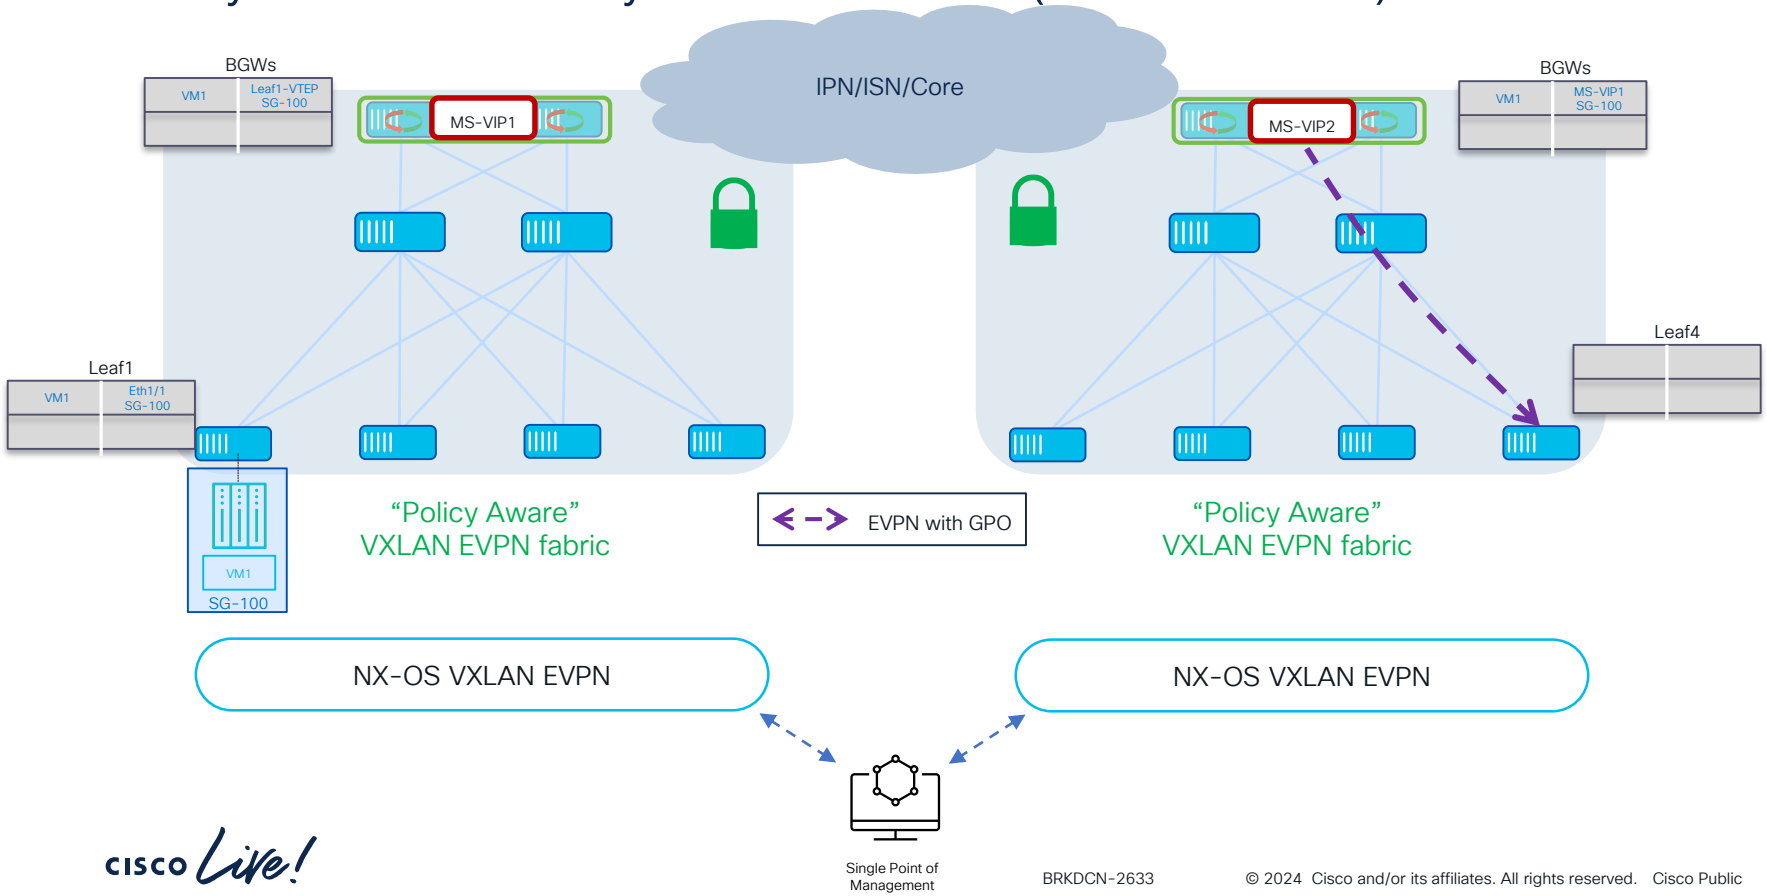

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

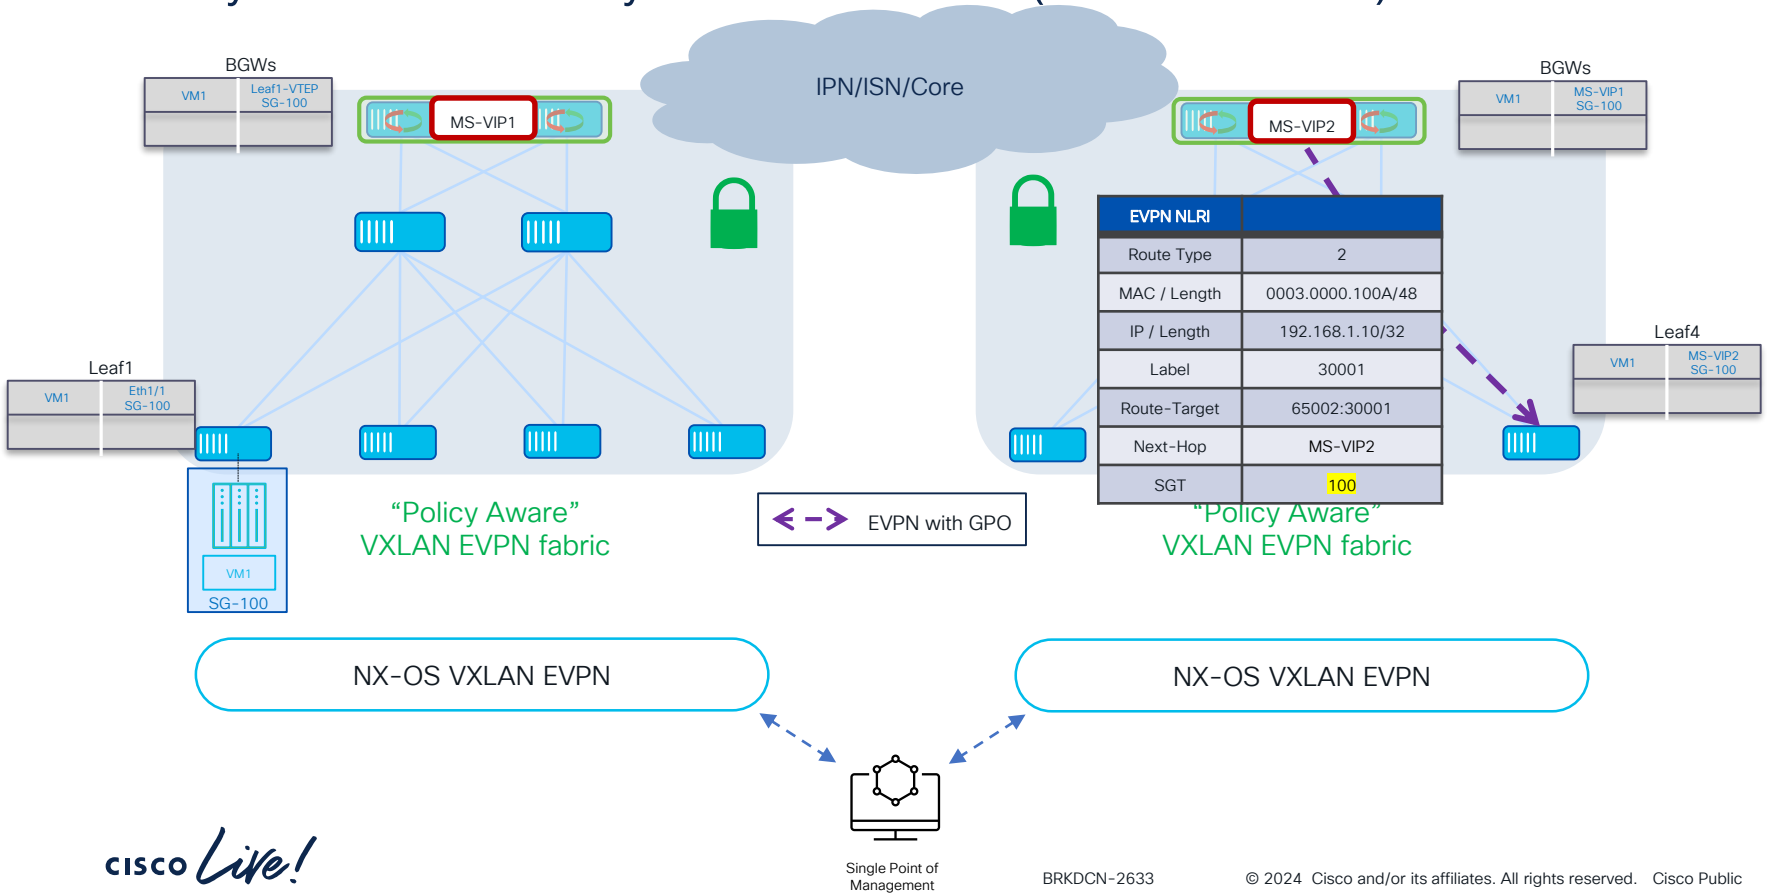

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

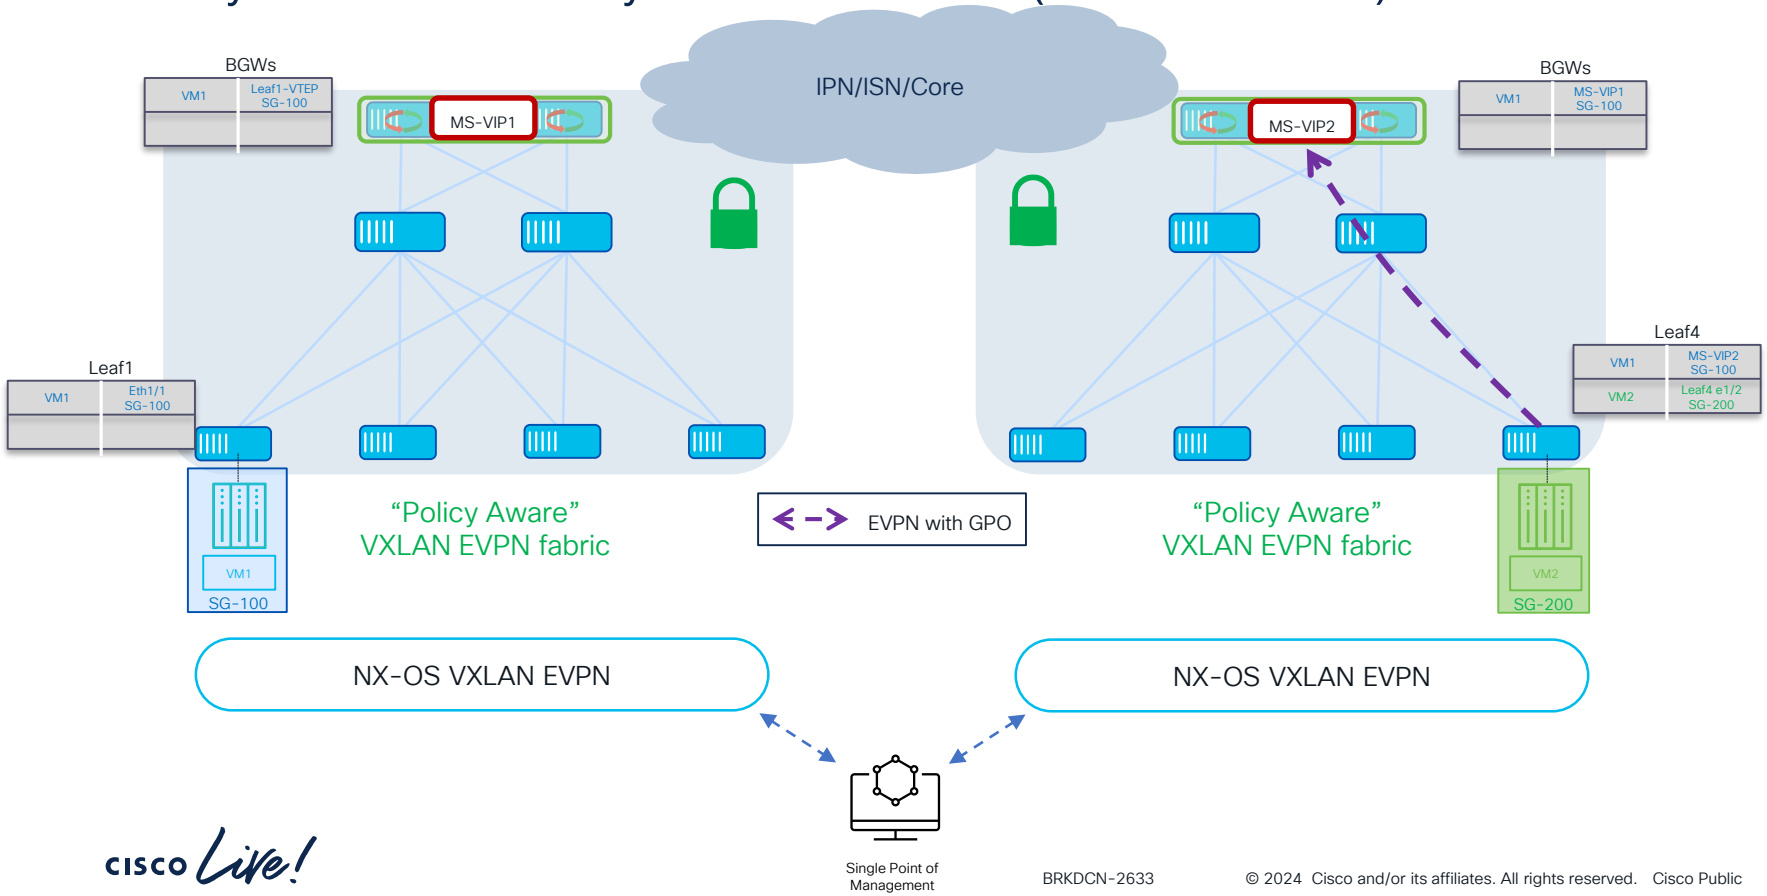

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

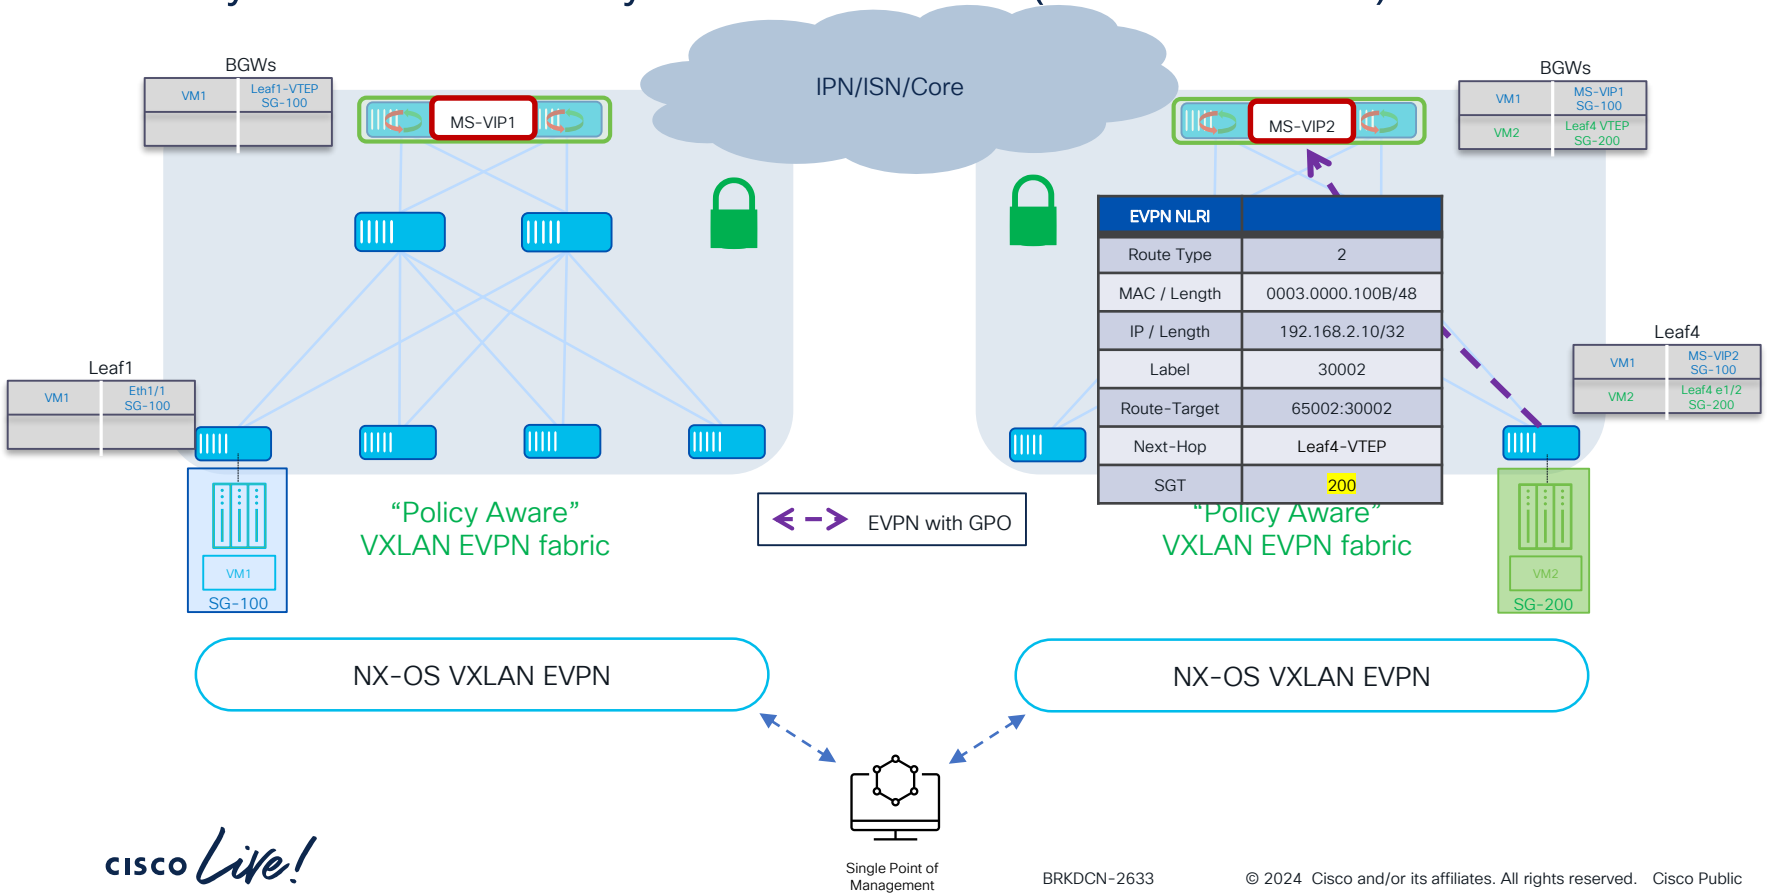

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

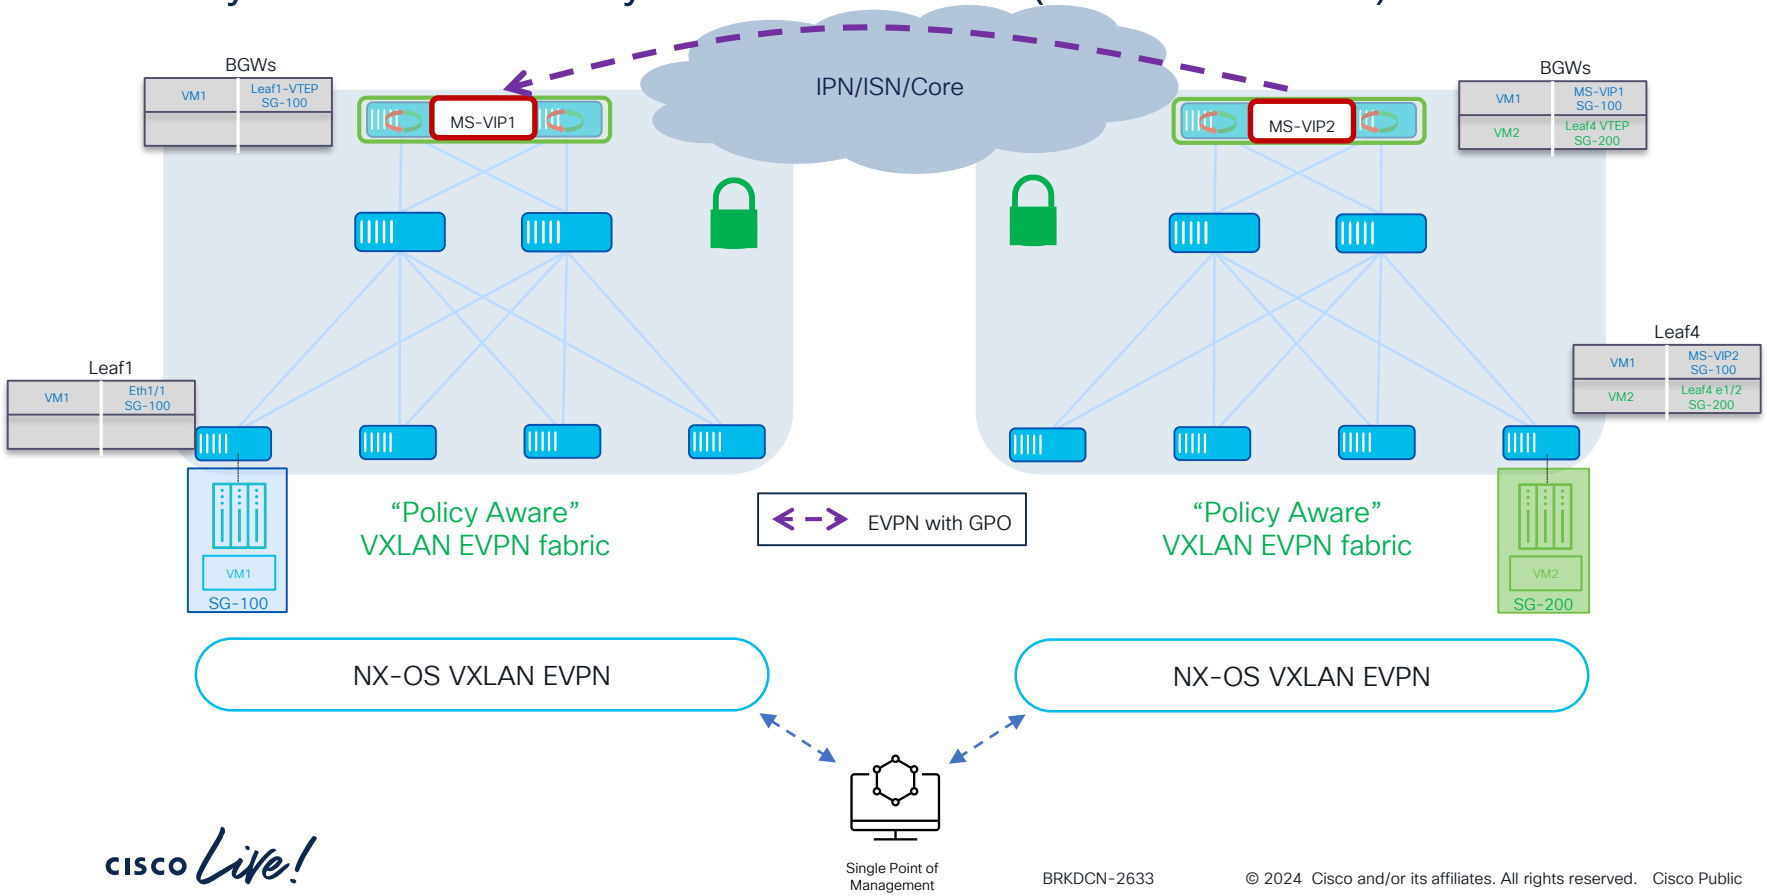

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

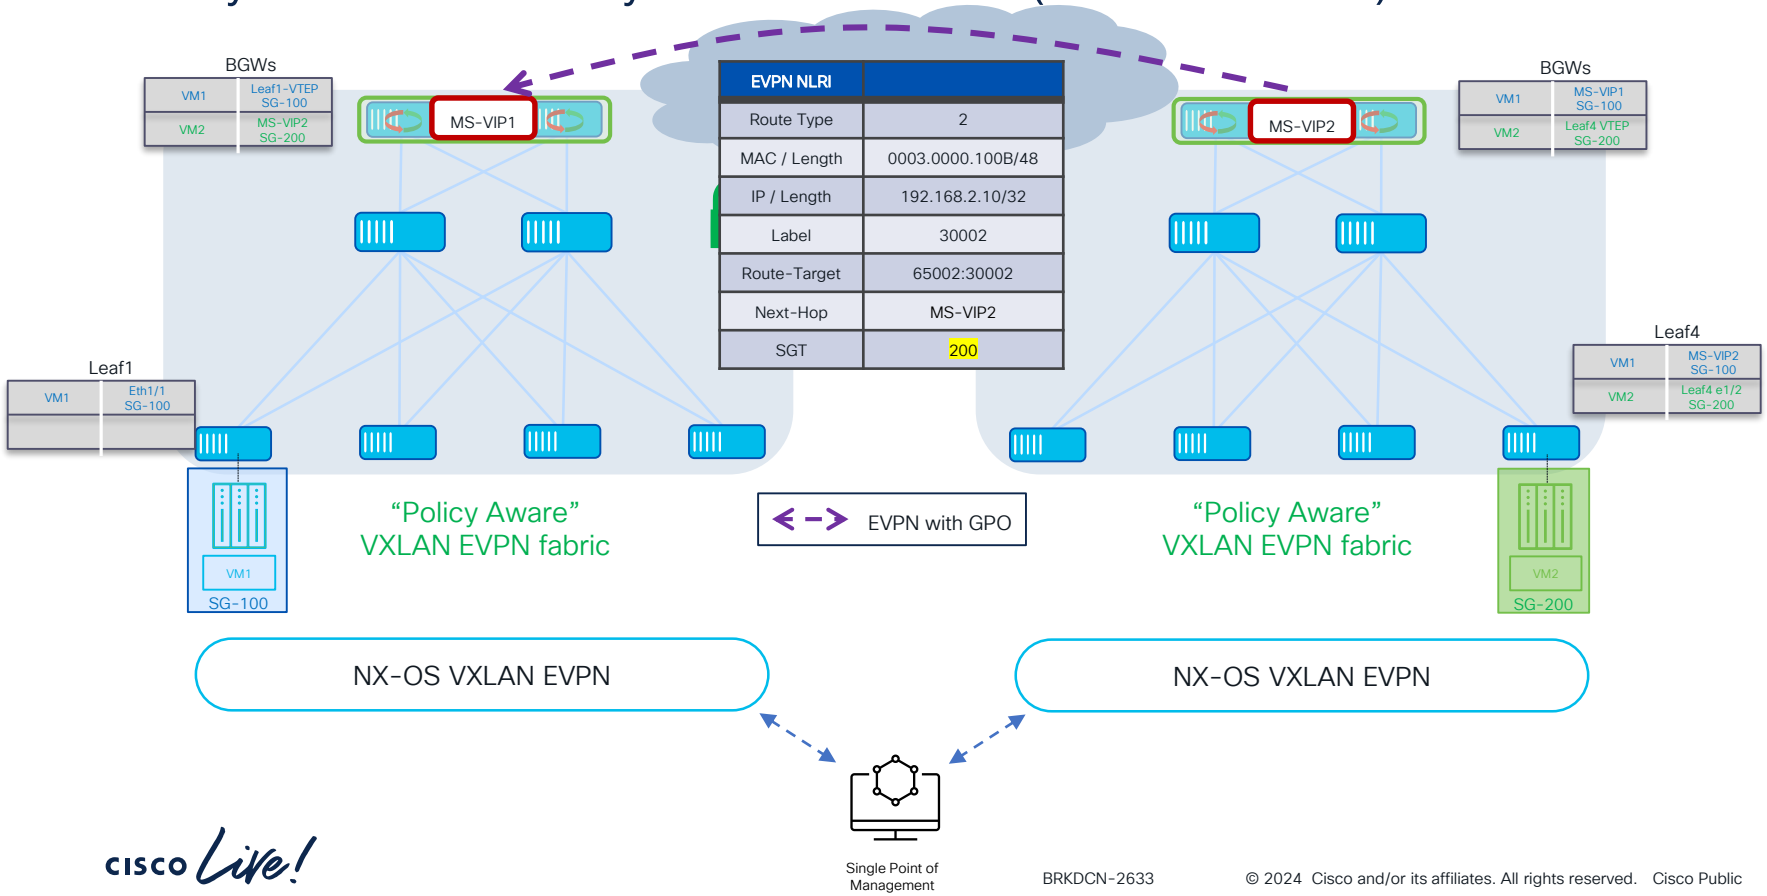

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

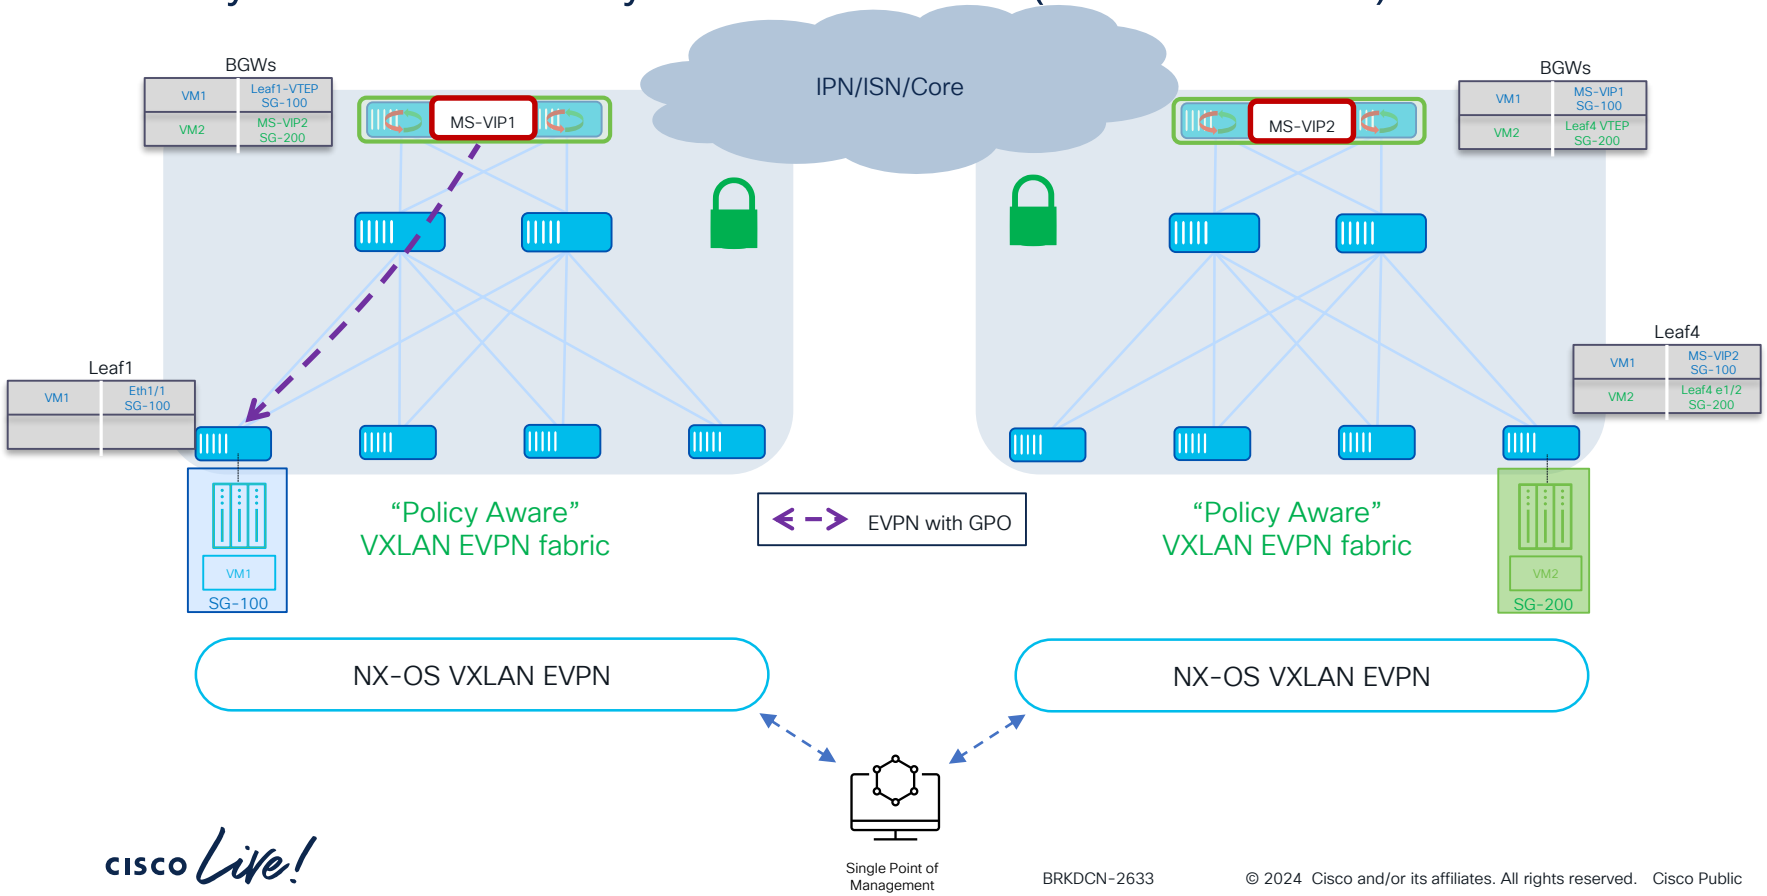

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Control Plane)

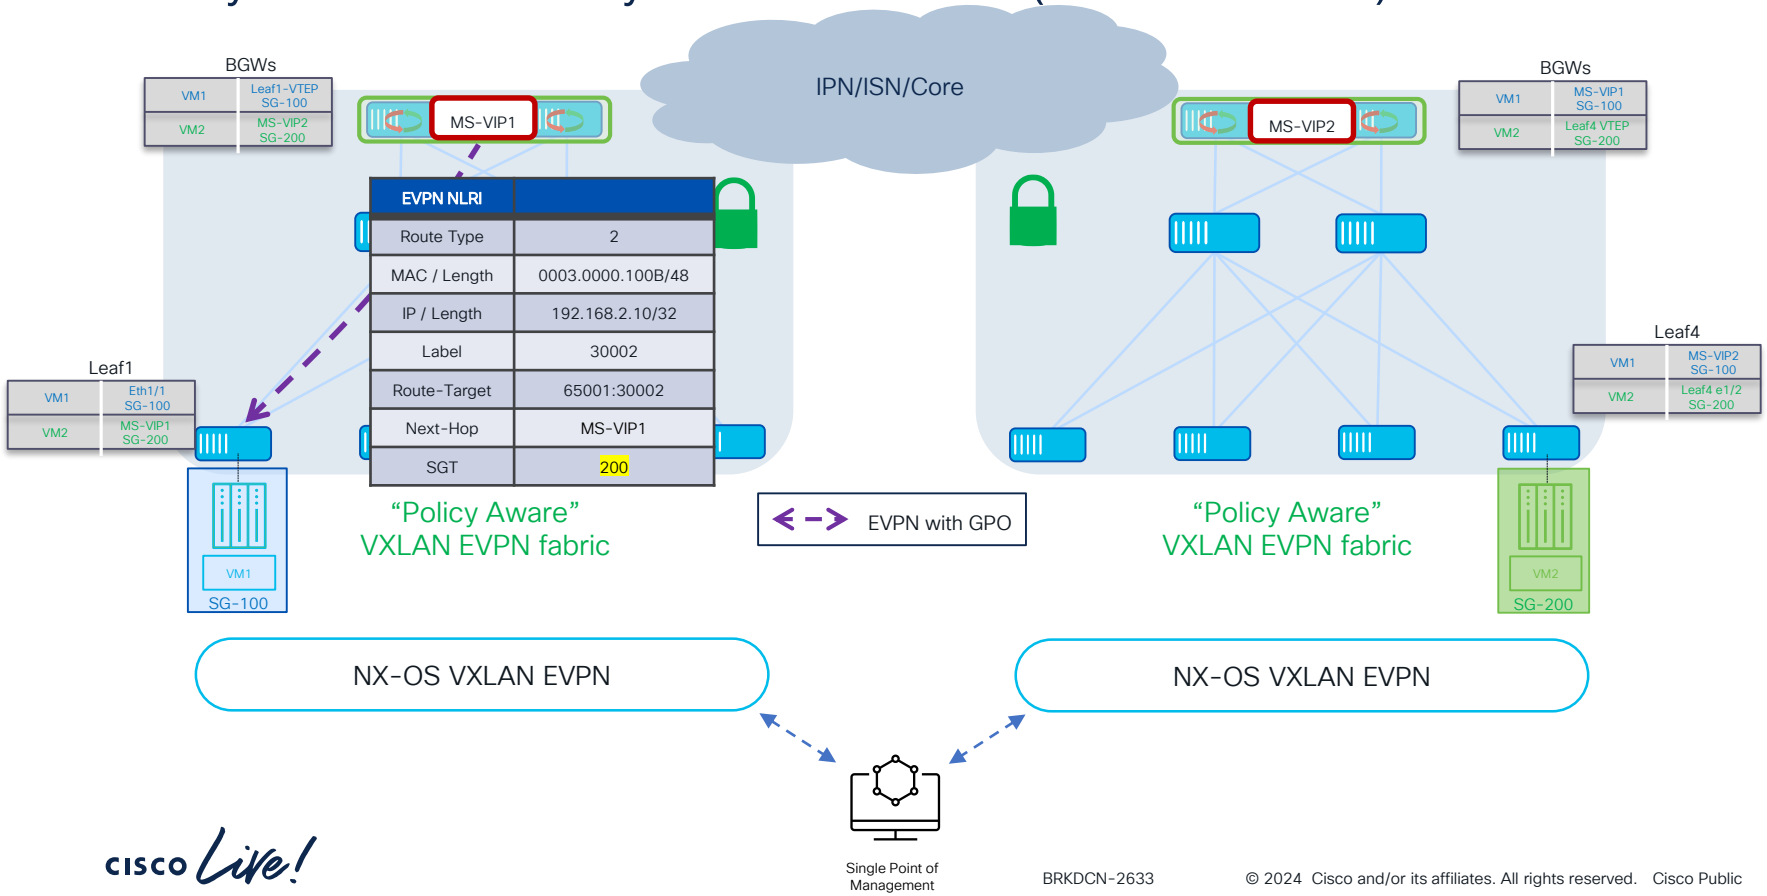

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Data Plane)

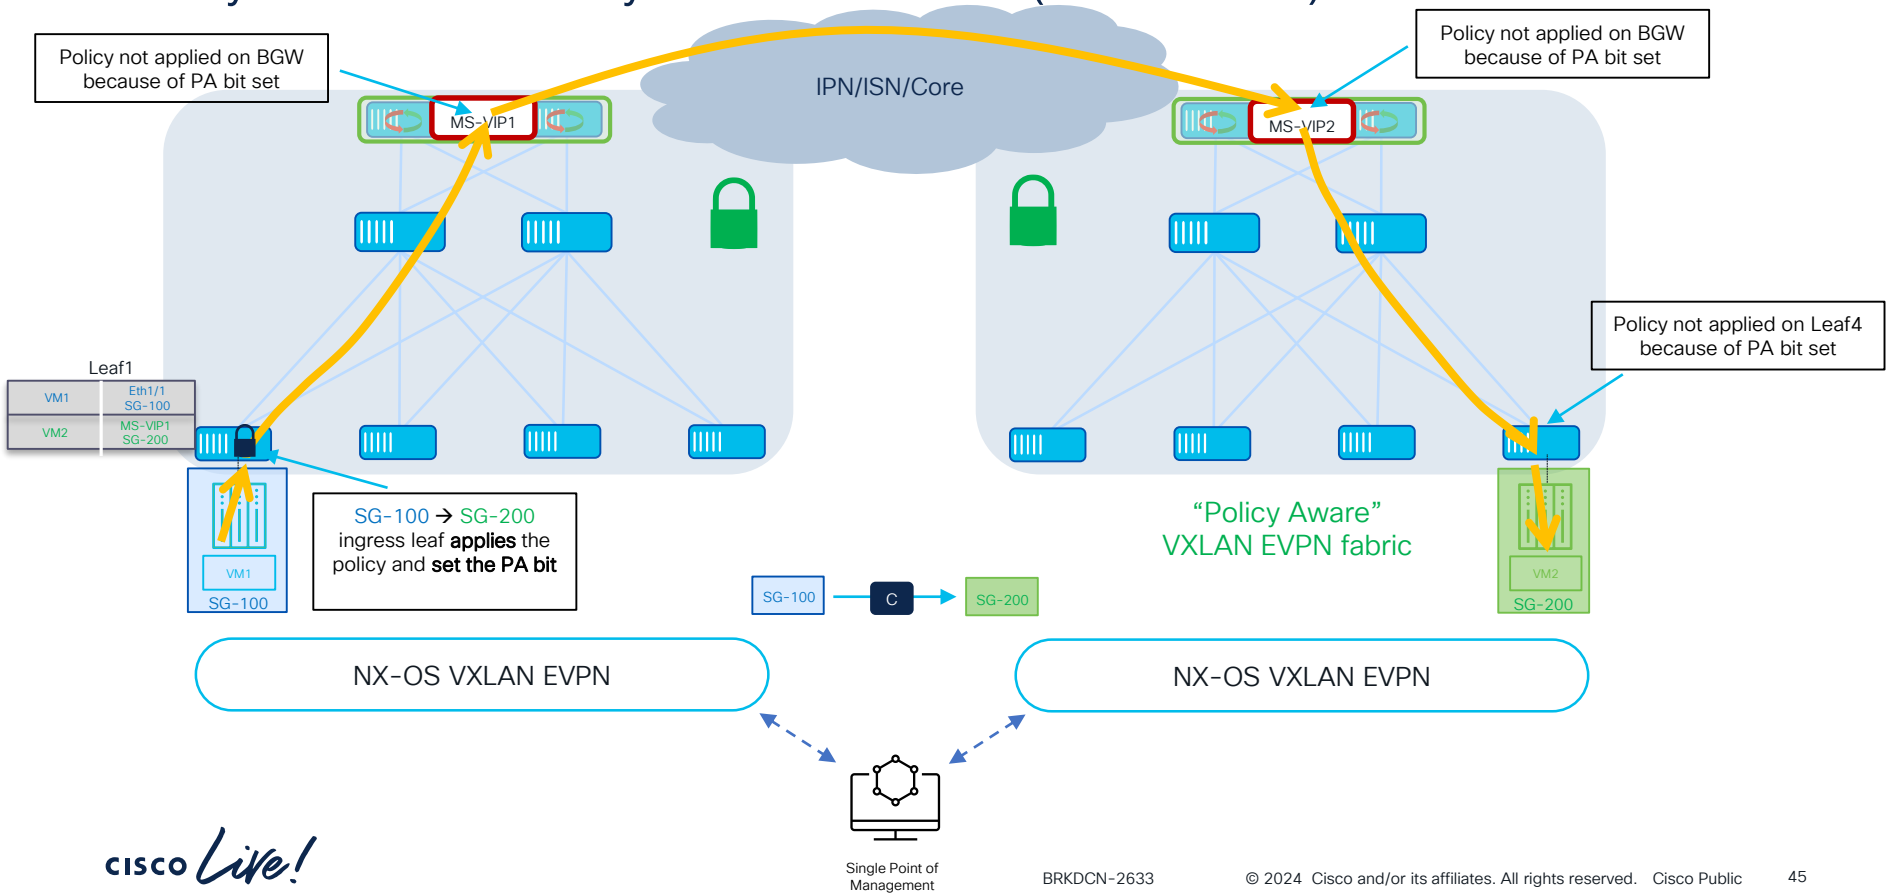

VXLAN GPO with Multi-Site
Policy Aware to Policy Aware Fabrics (Data Plane)

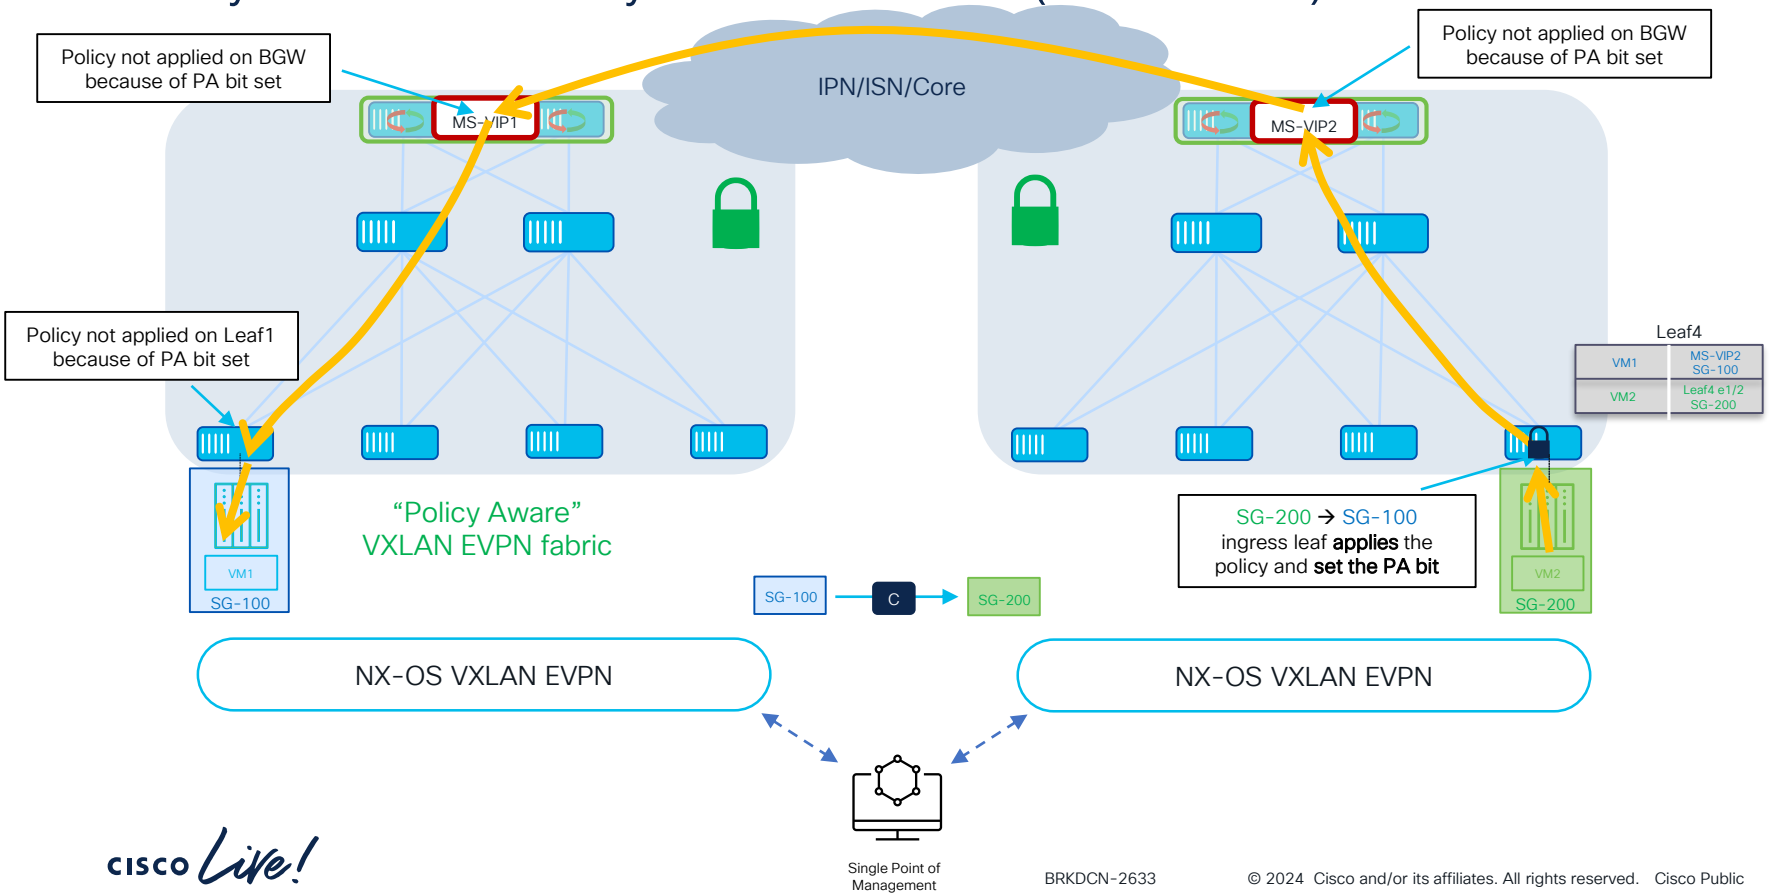

VXLAN GPO with Multi-Site

Policy Aware to Policy Unaware Fabrics

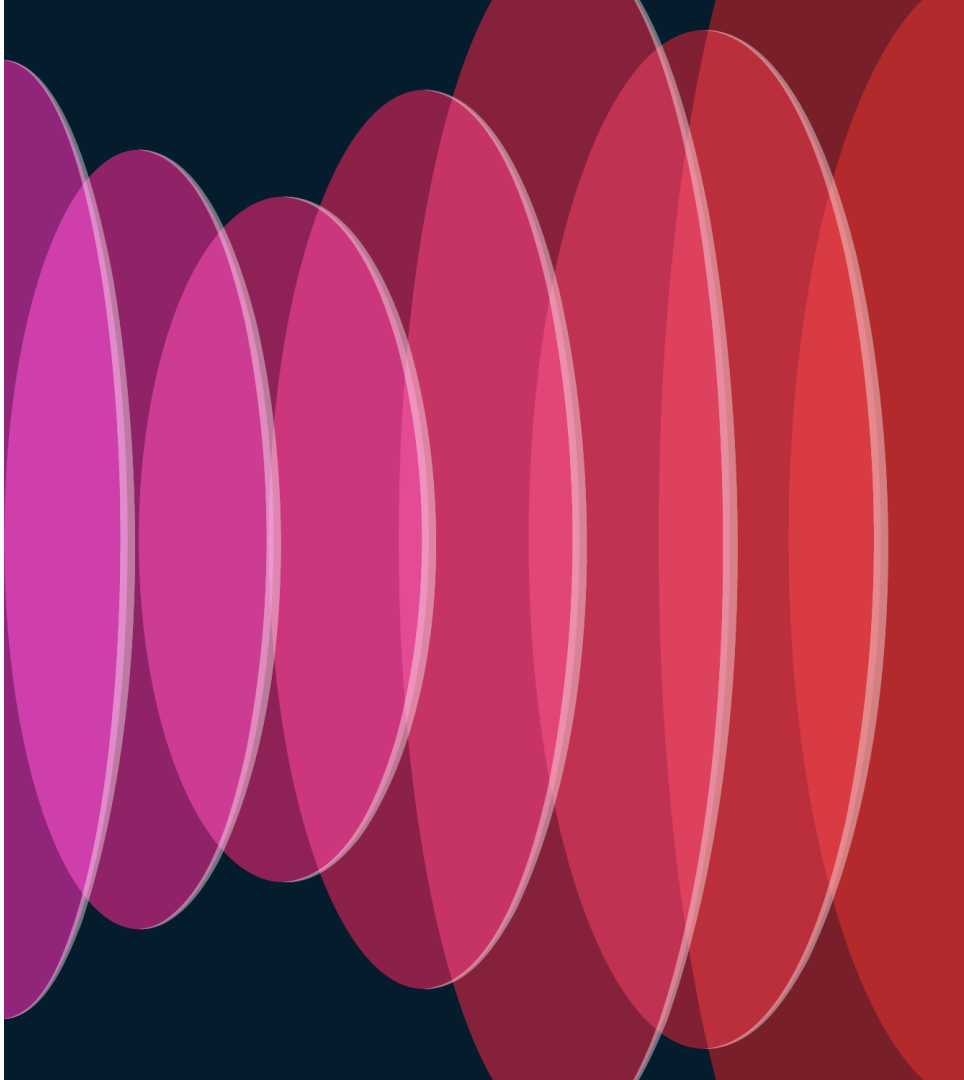

cisco Live!

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

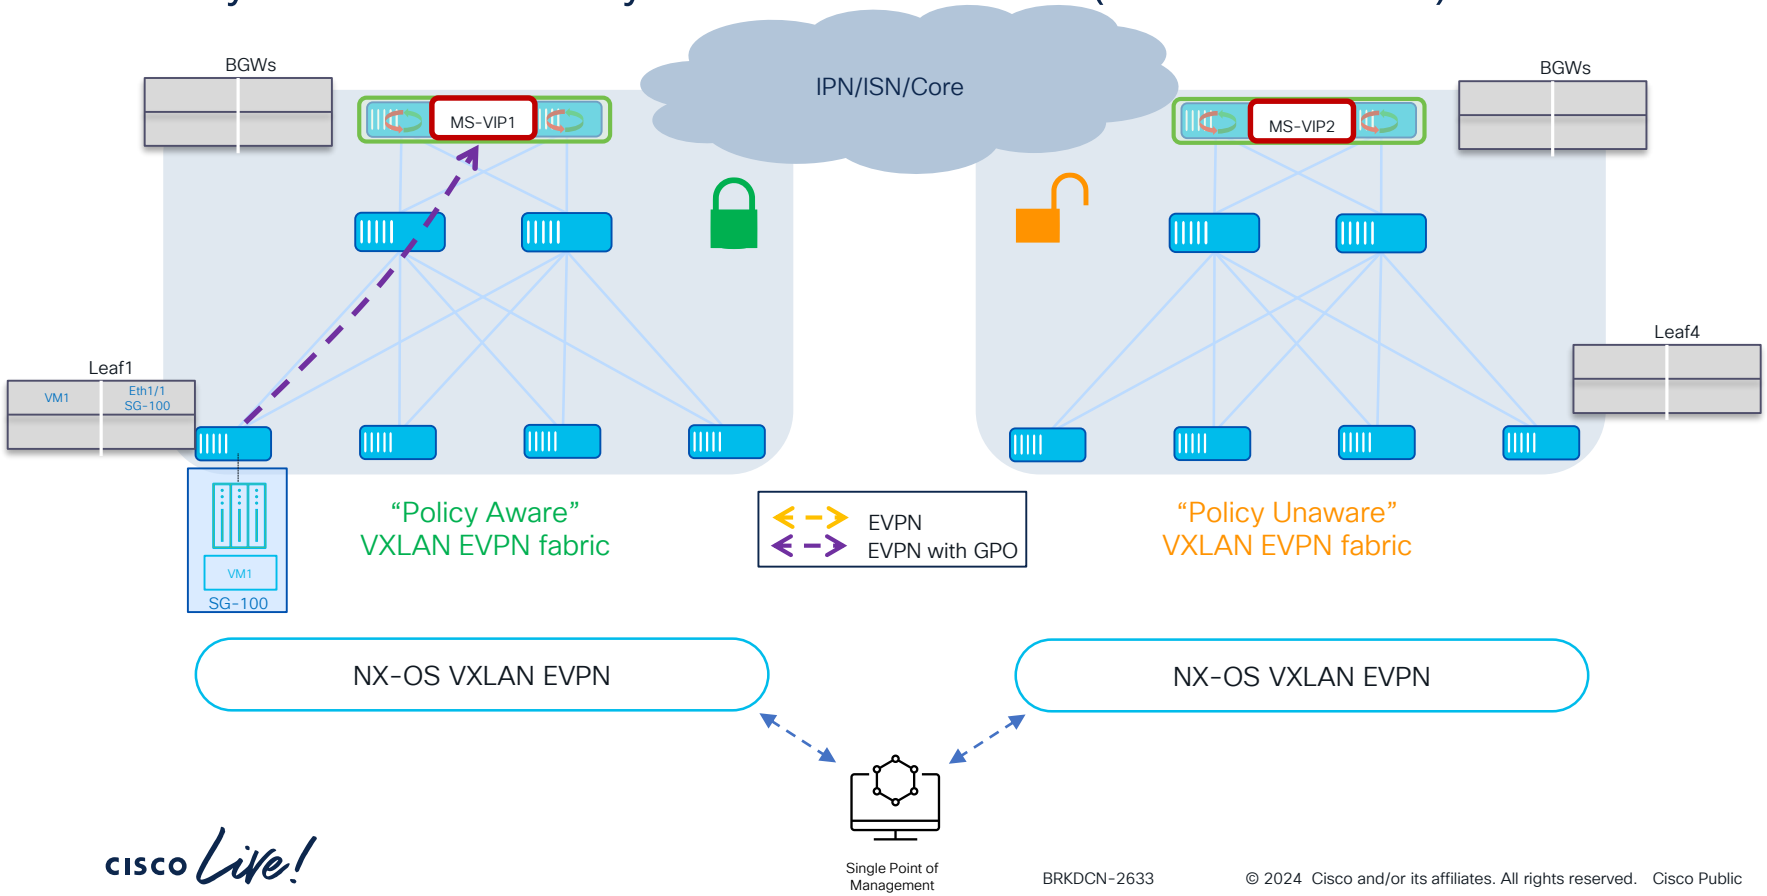

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

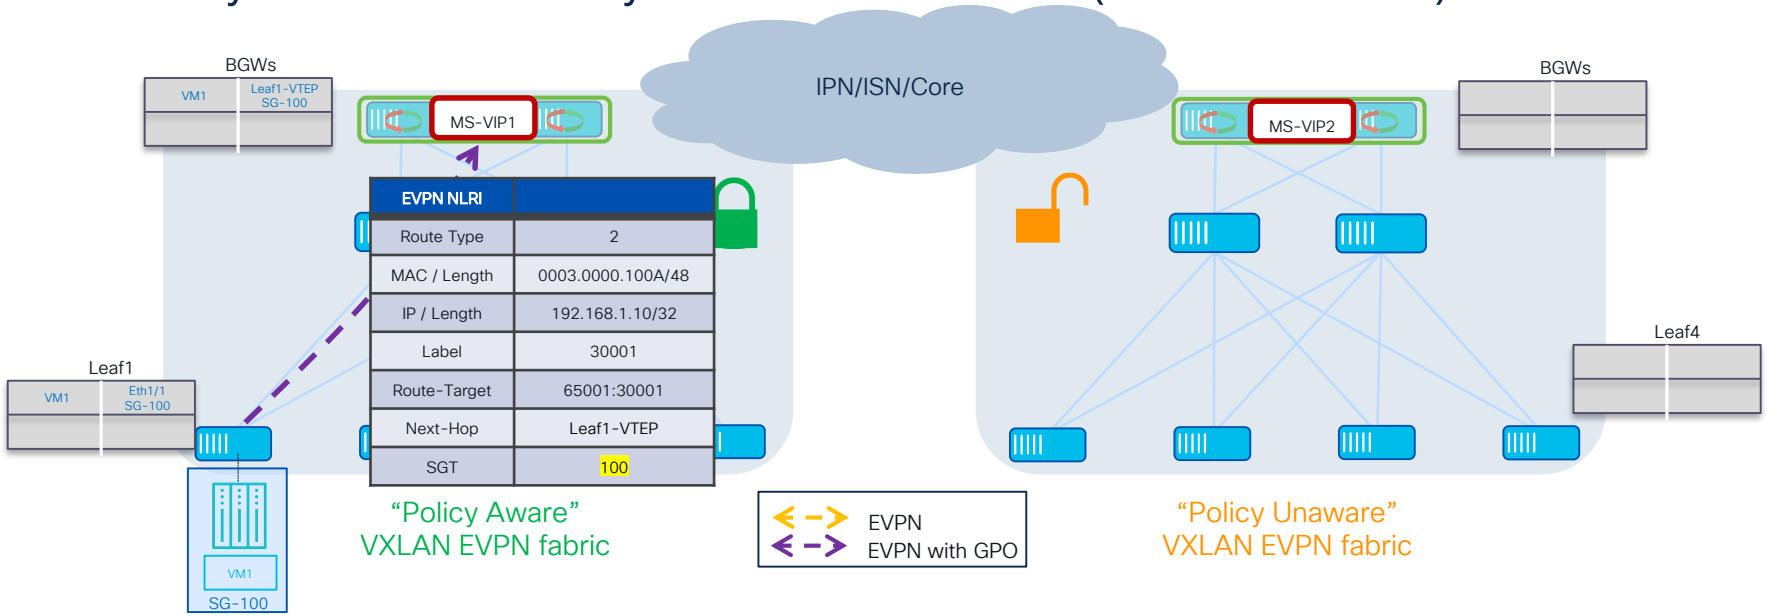

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

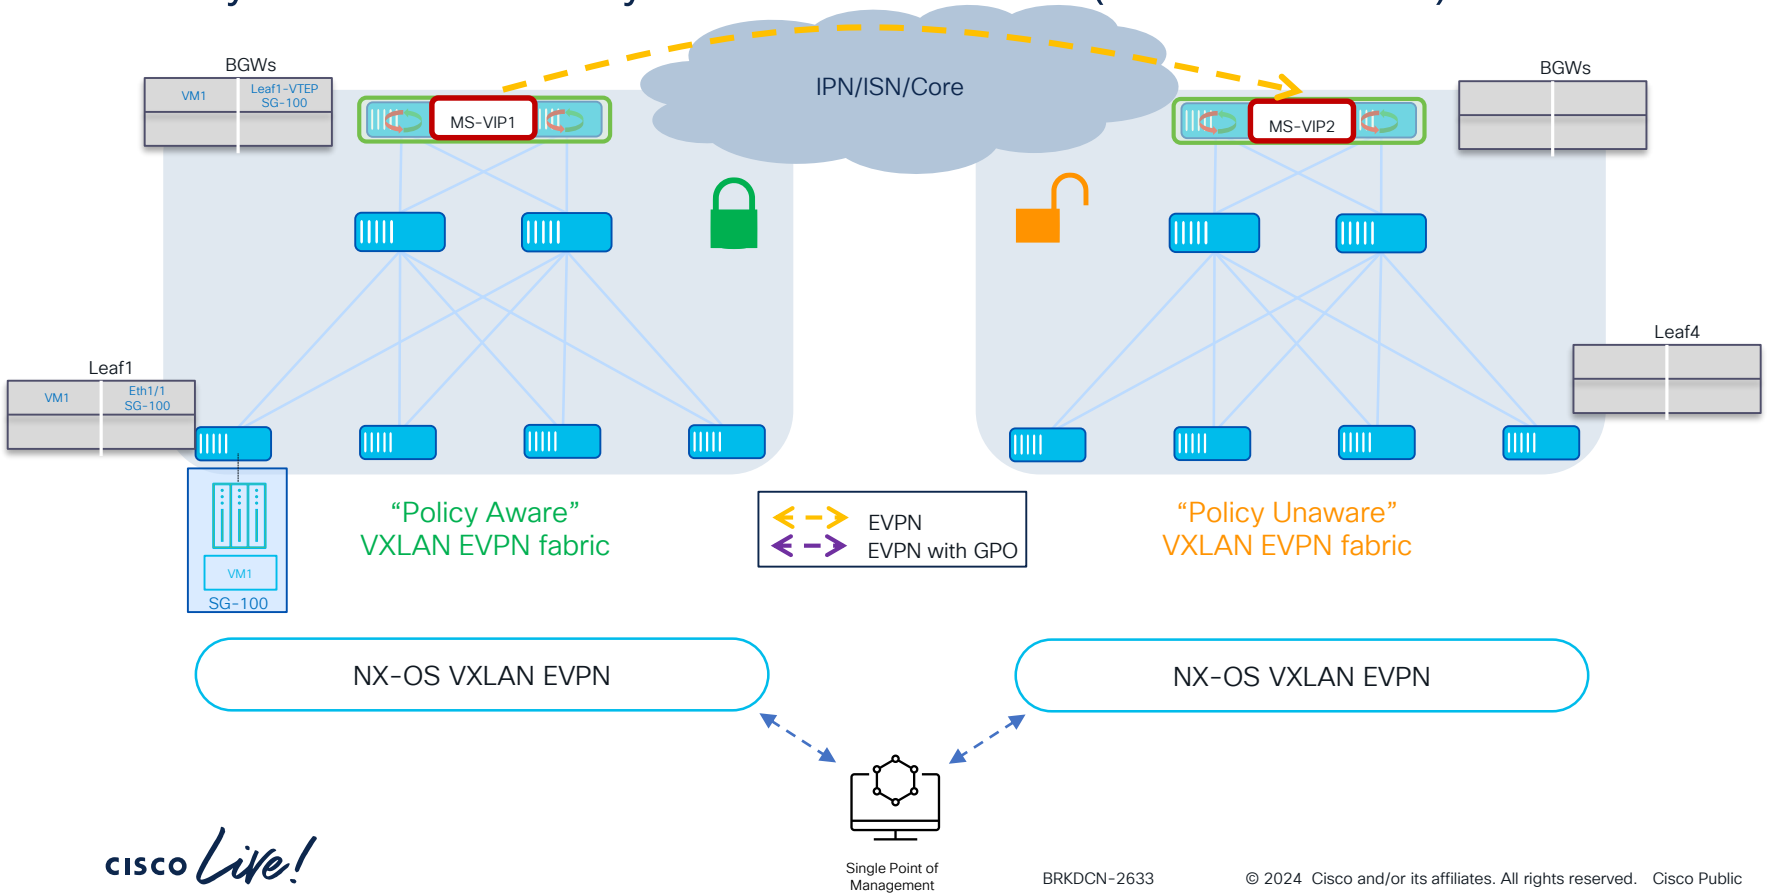

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

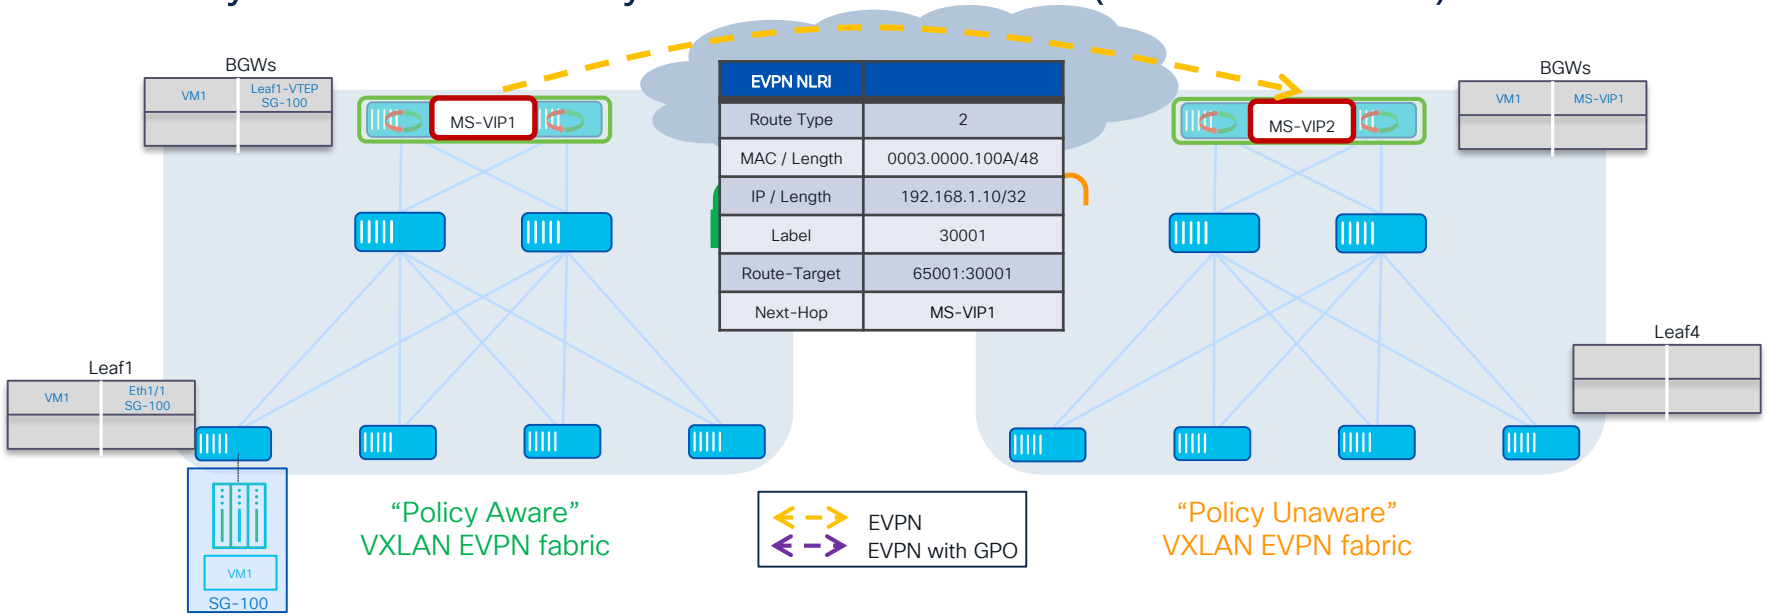

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

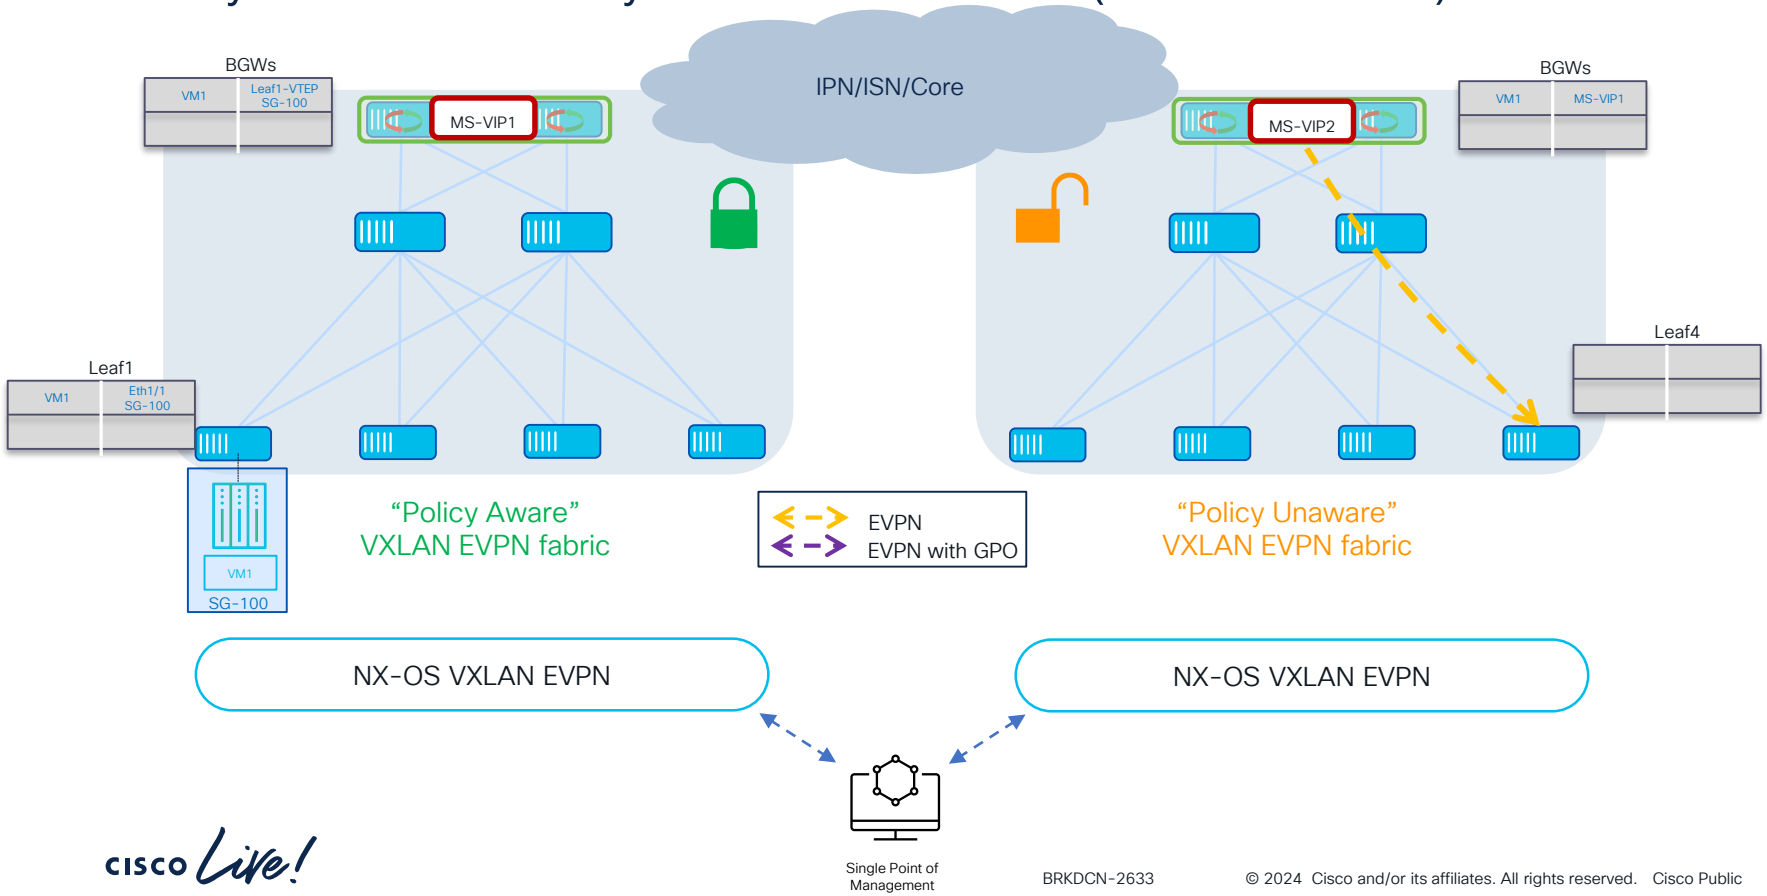

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

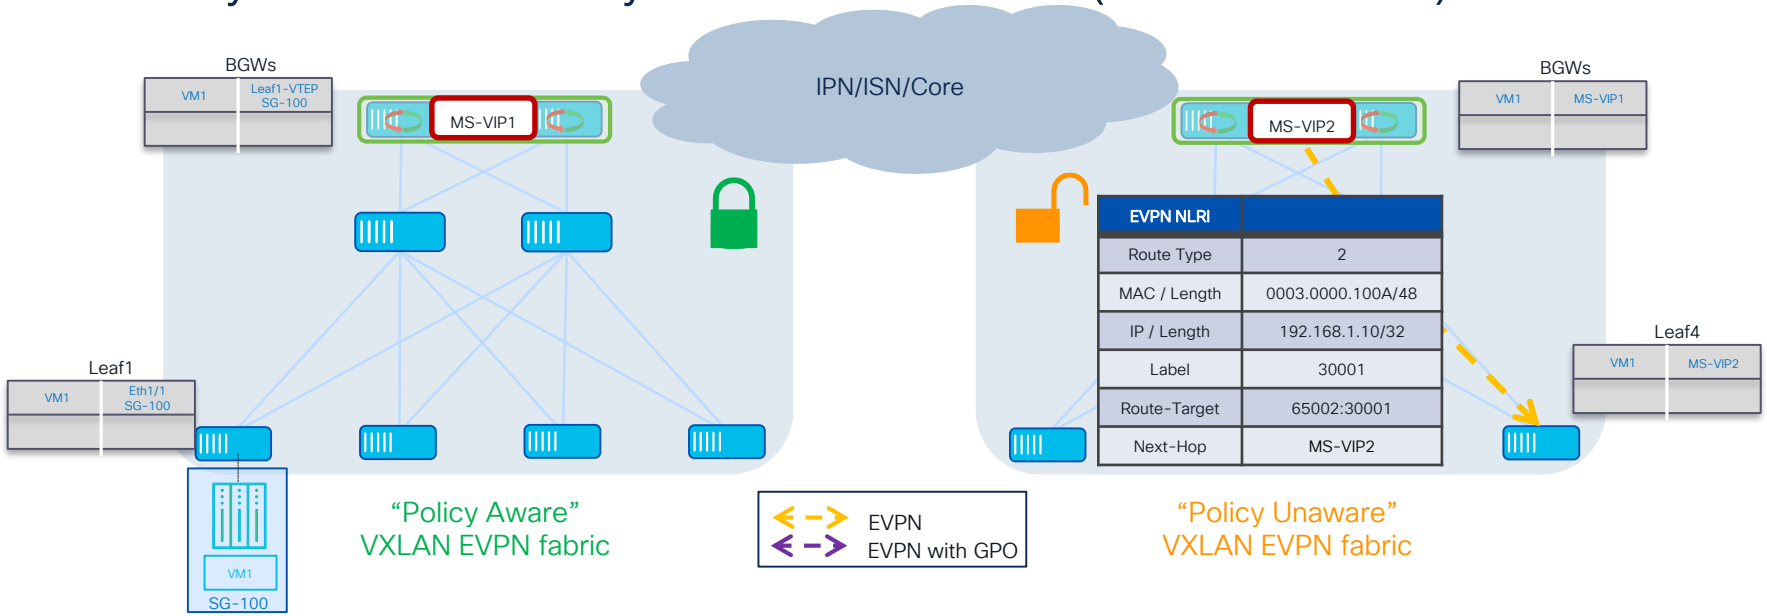

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

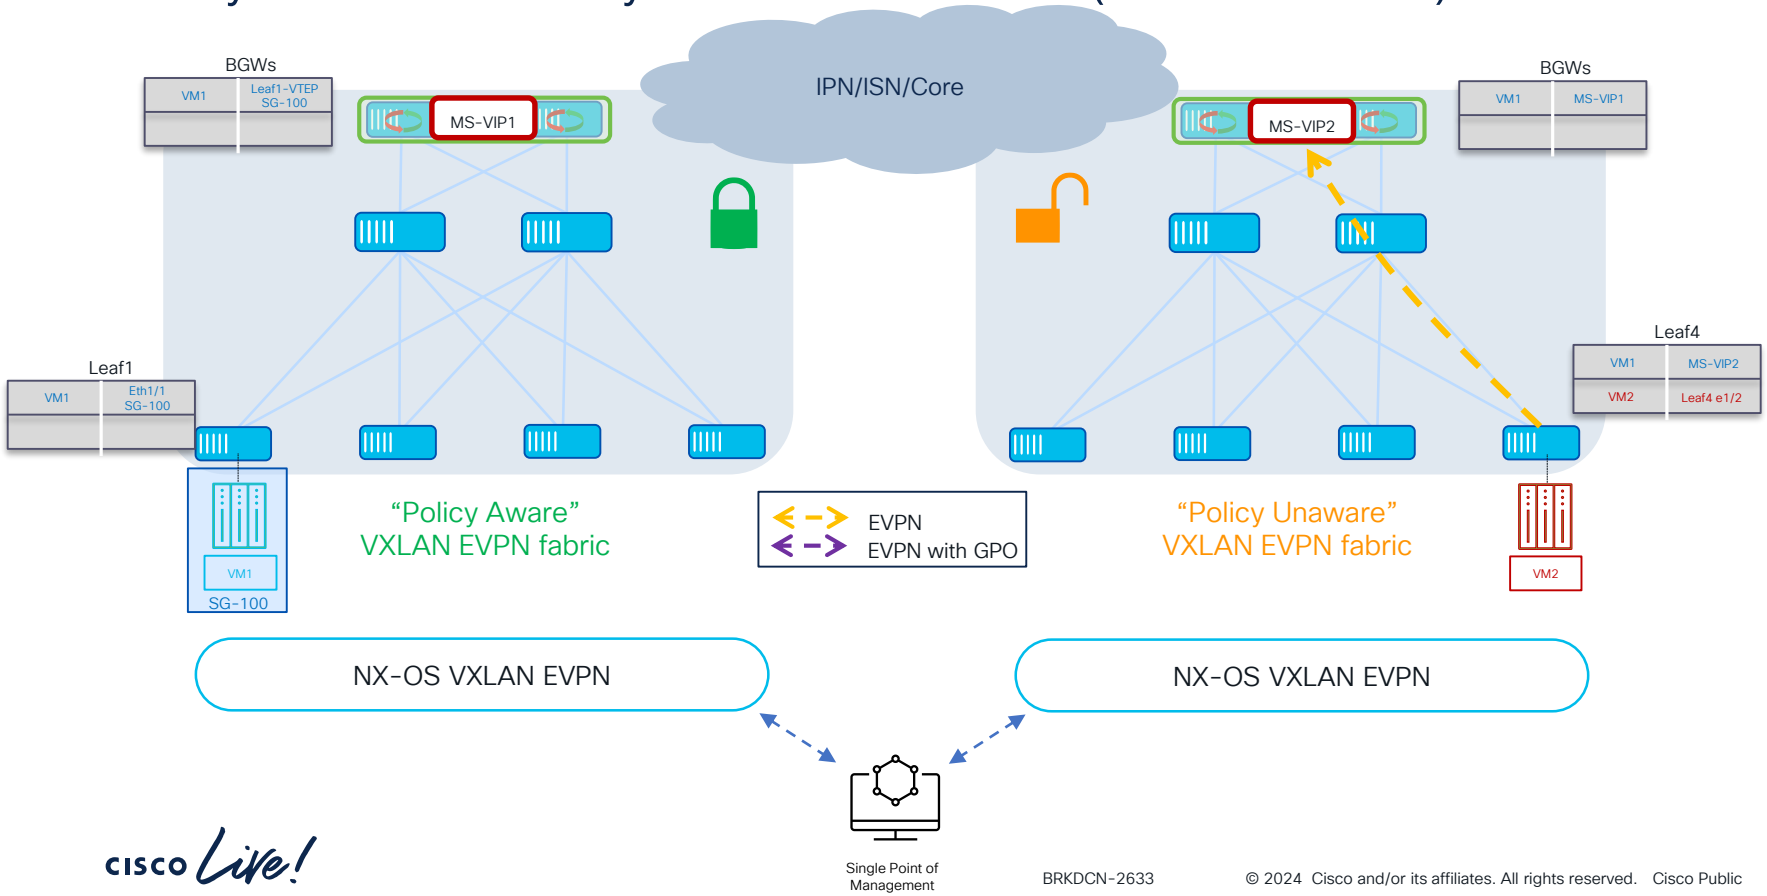

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

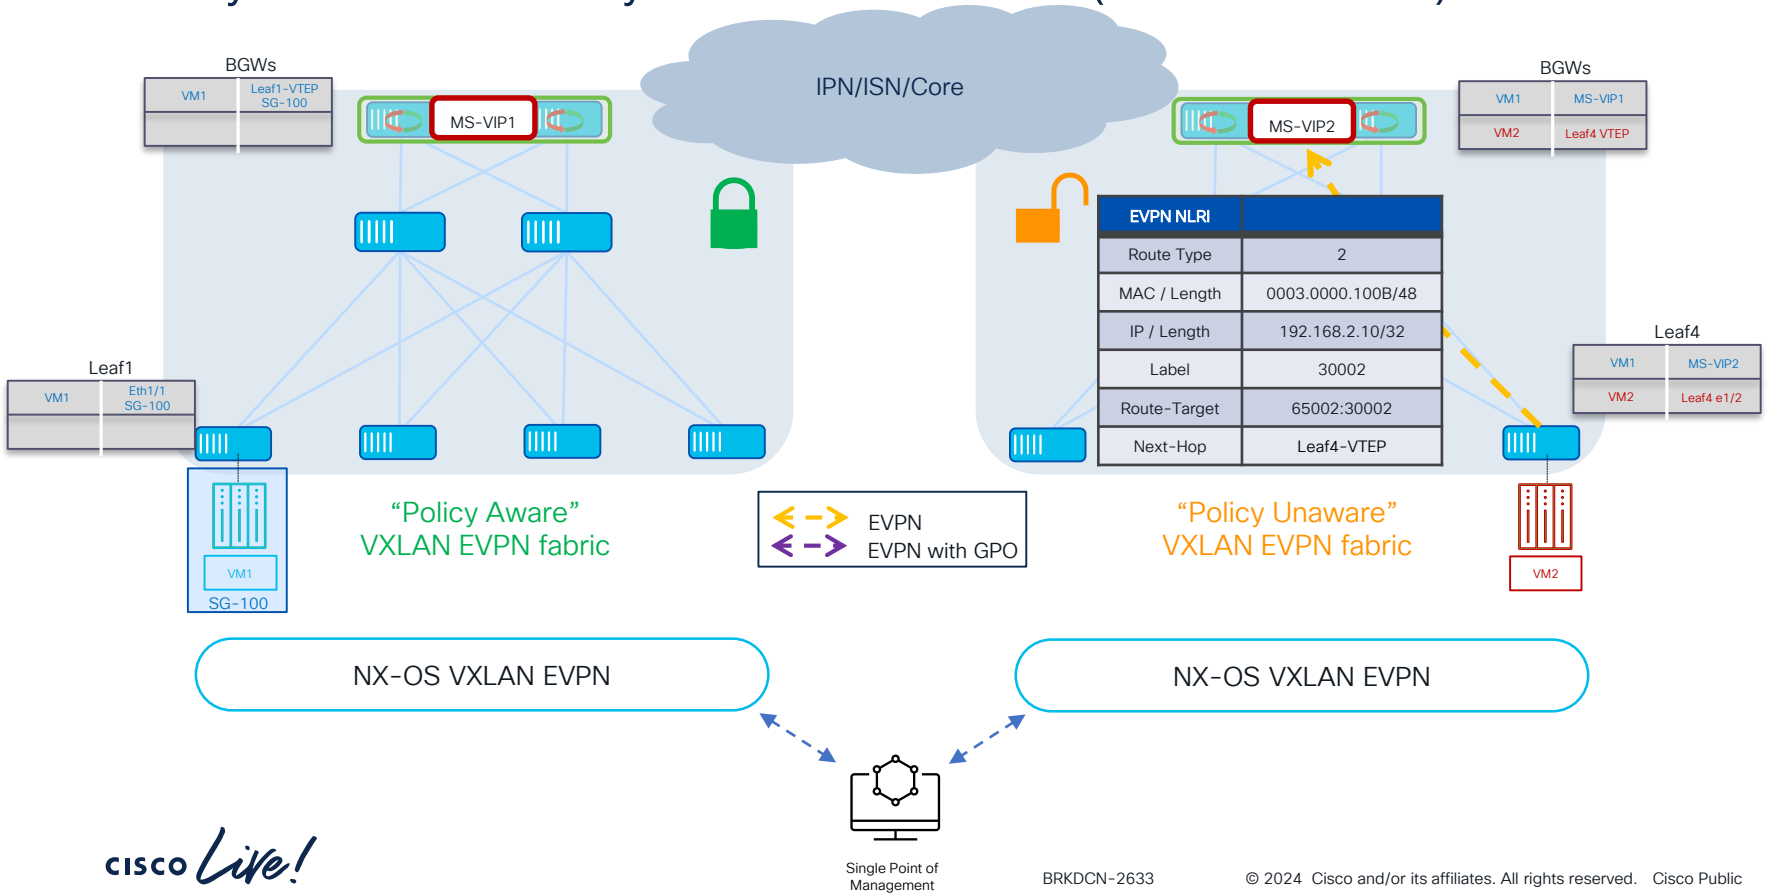

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

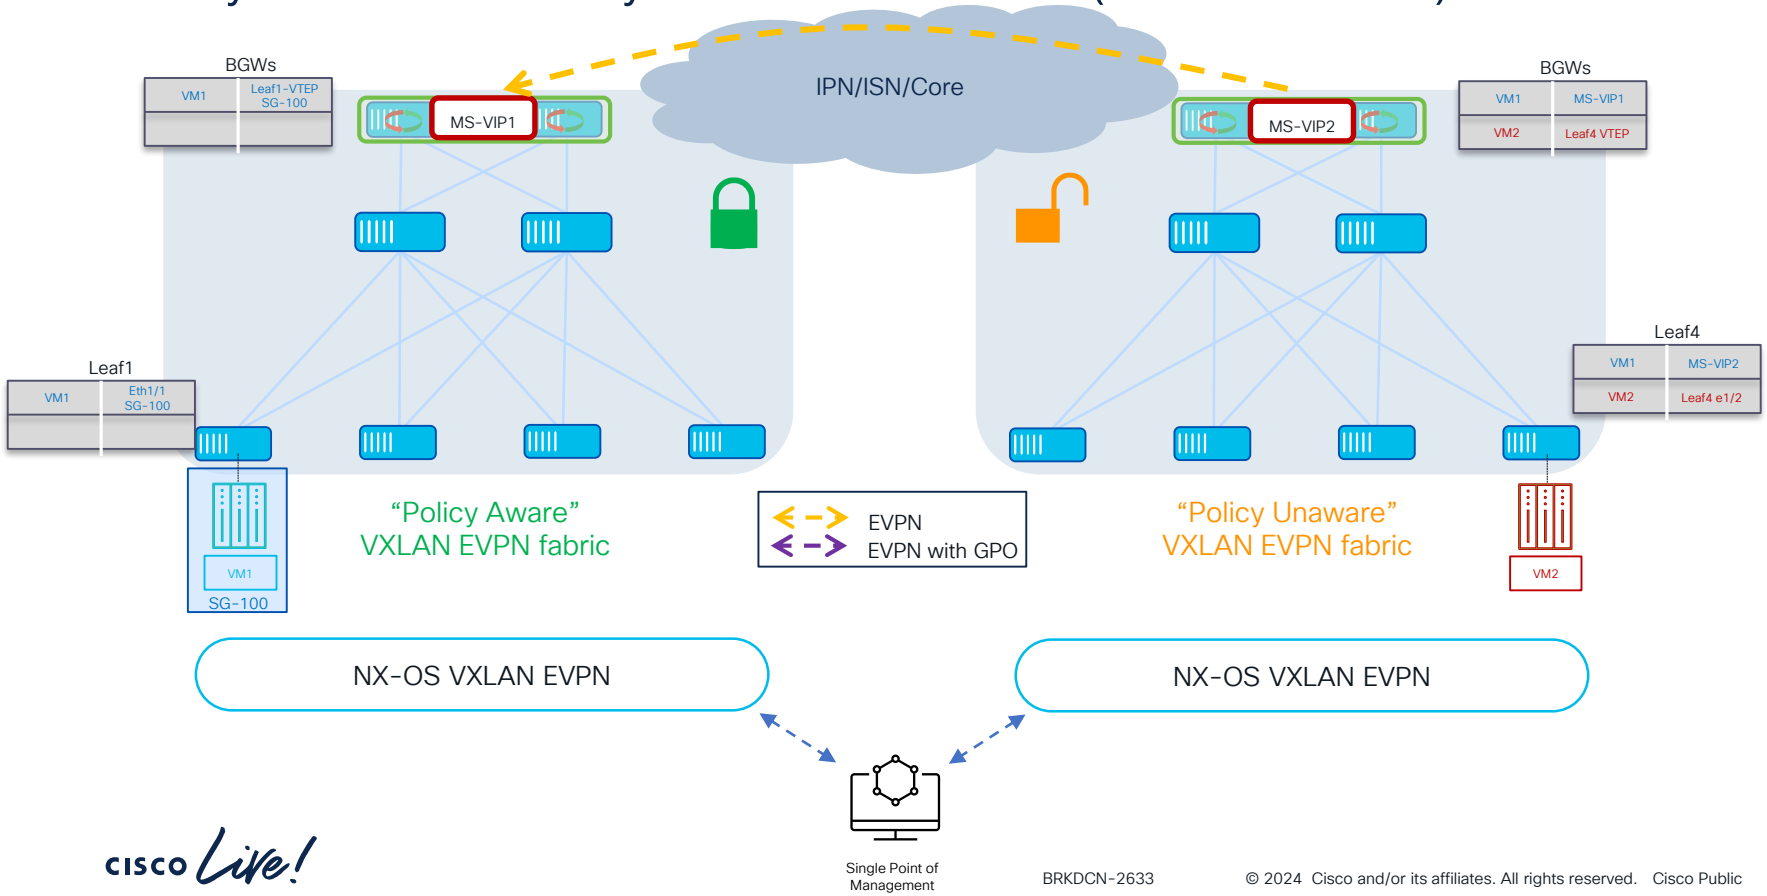

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

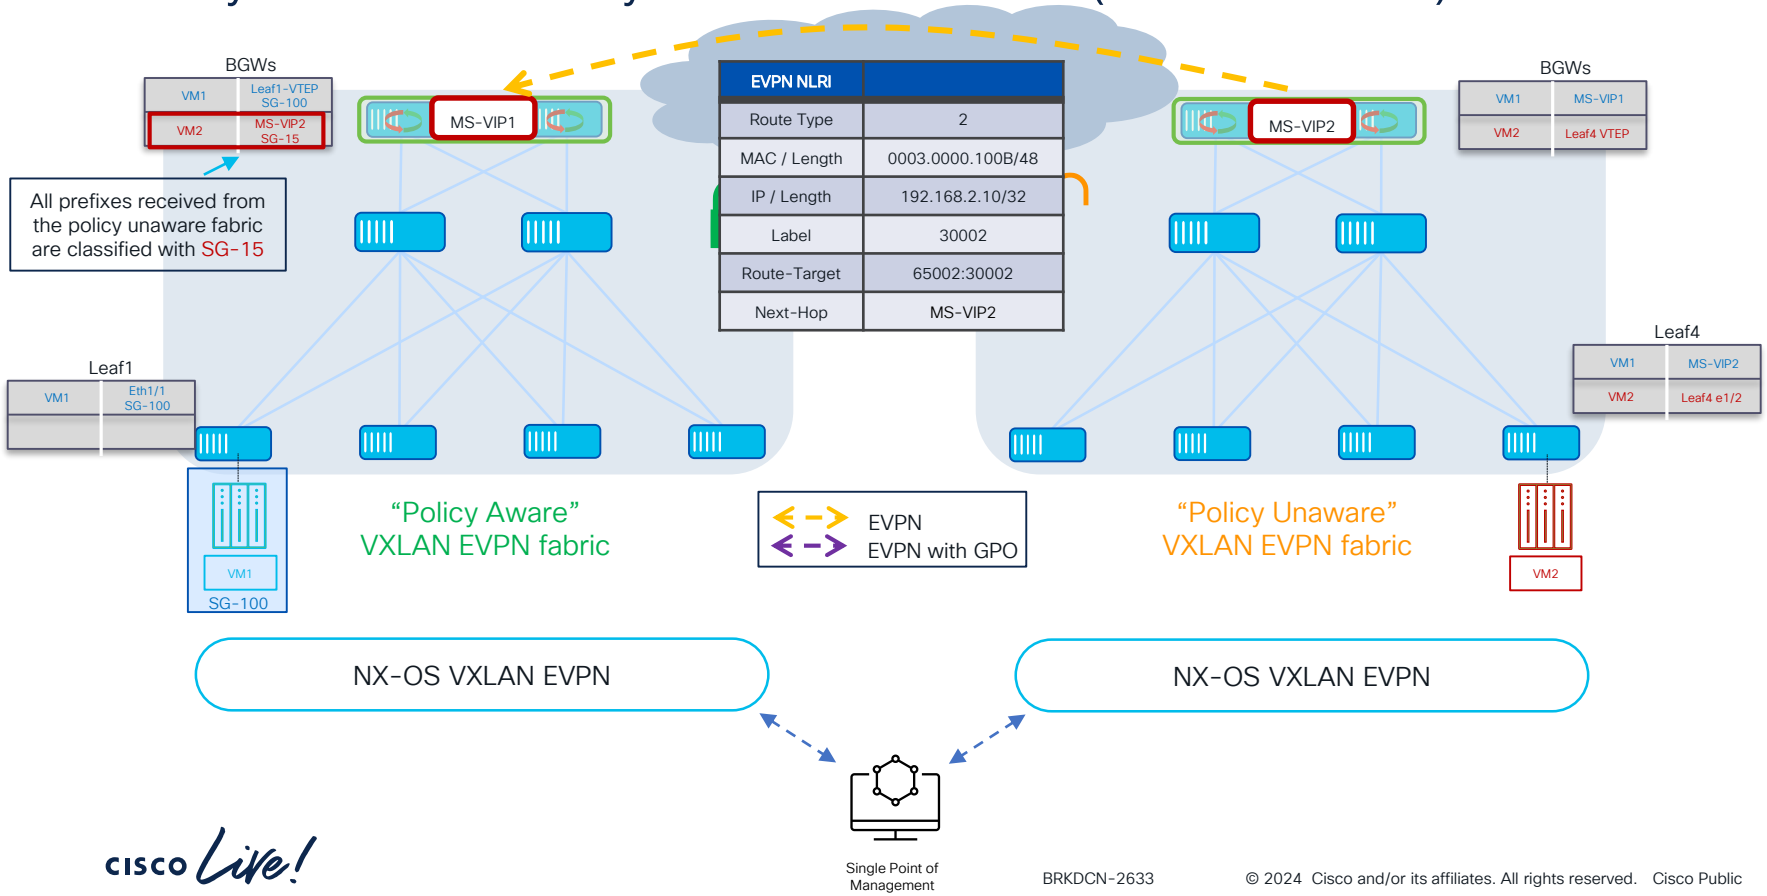

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

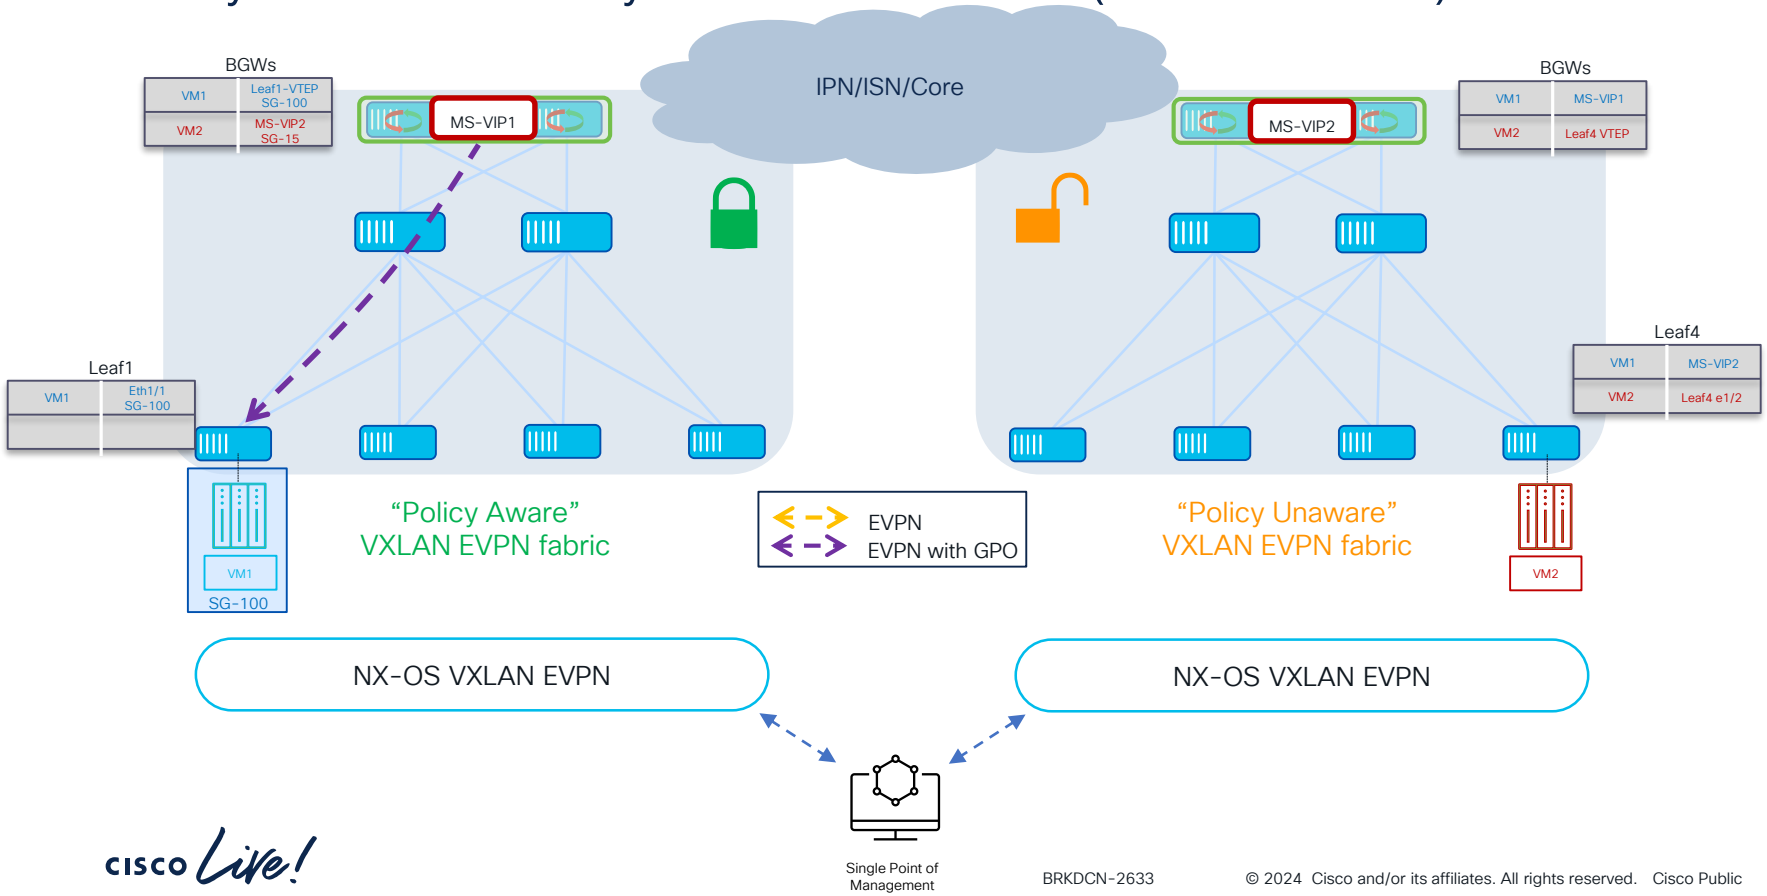

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Control Plane)

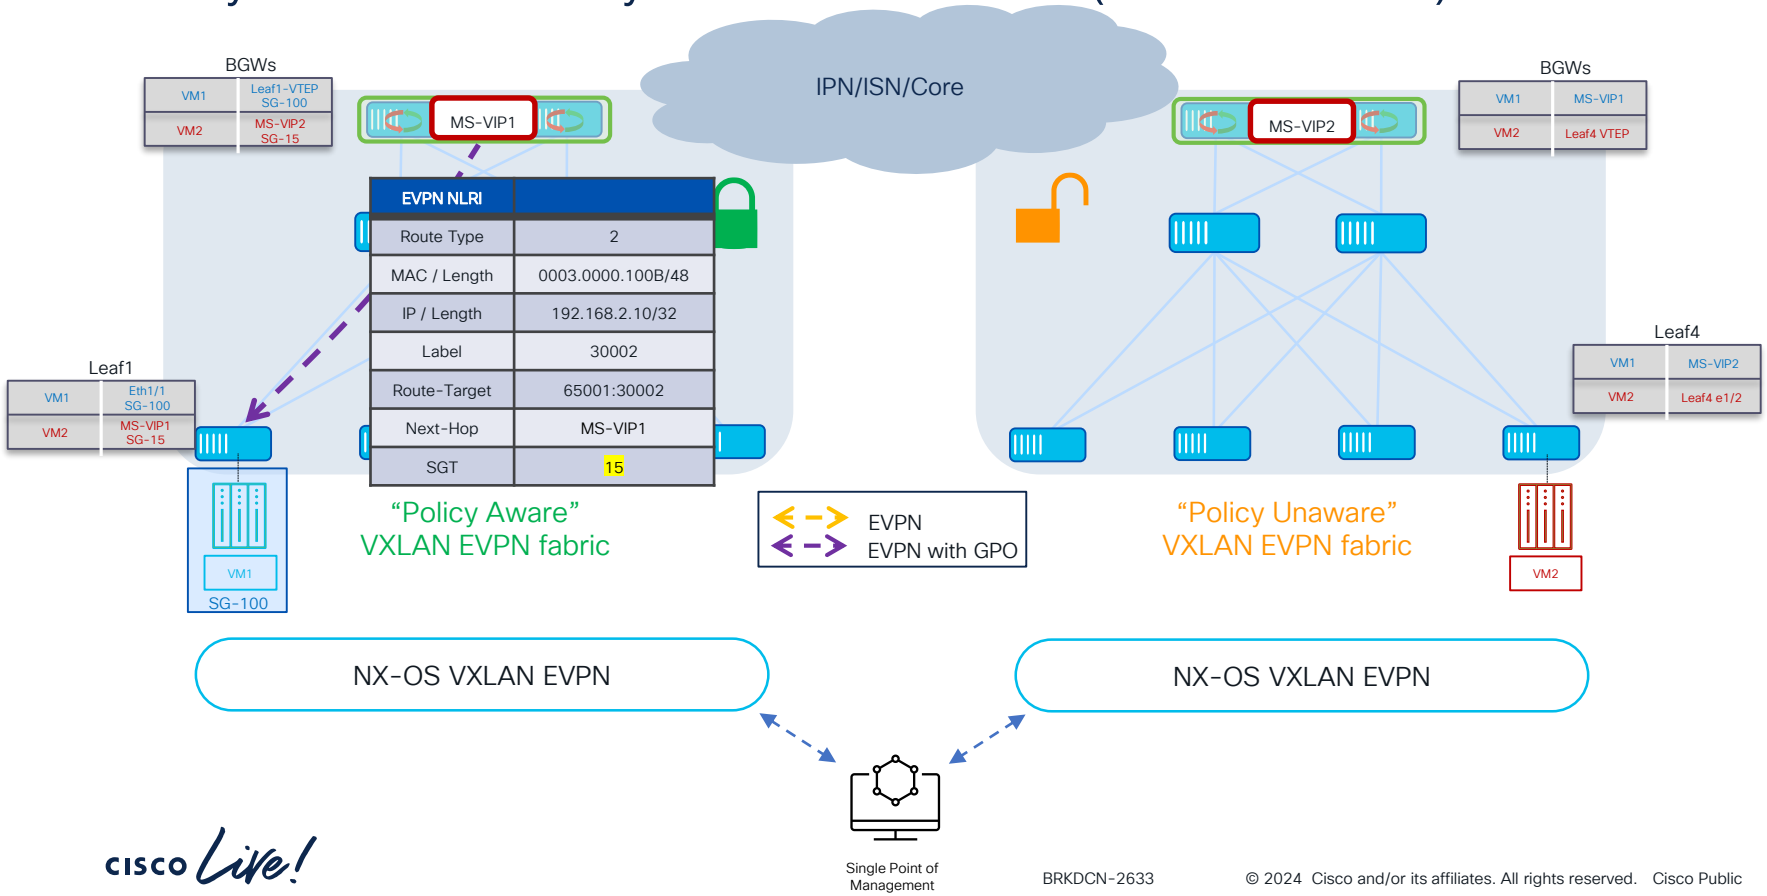

VXLAN GPO with Multi-Site
Policy Aware to Policy Unaware Fabrics (Data Plane)

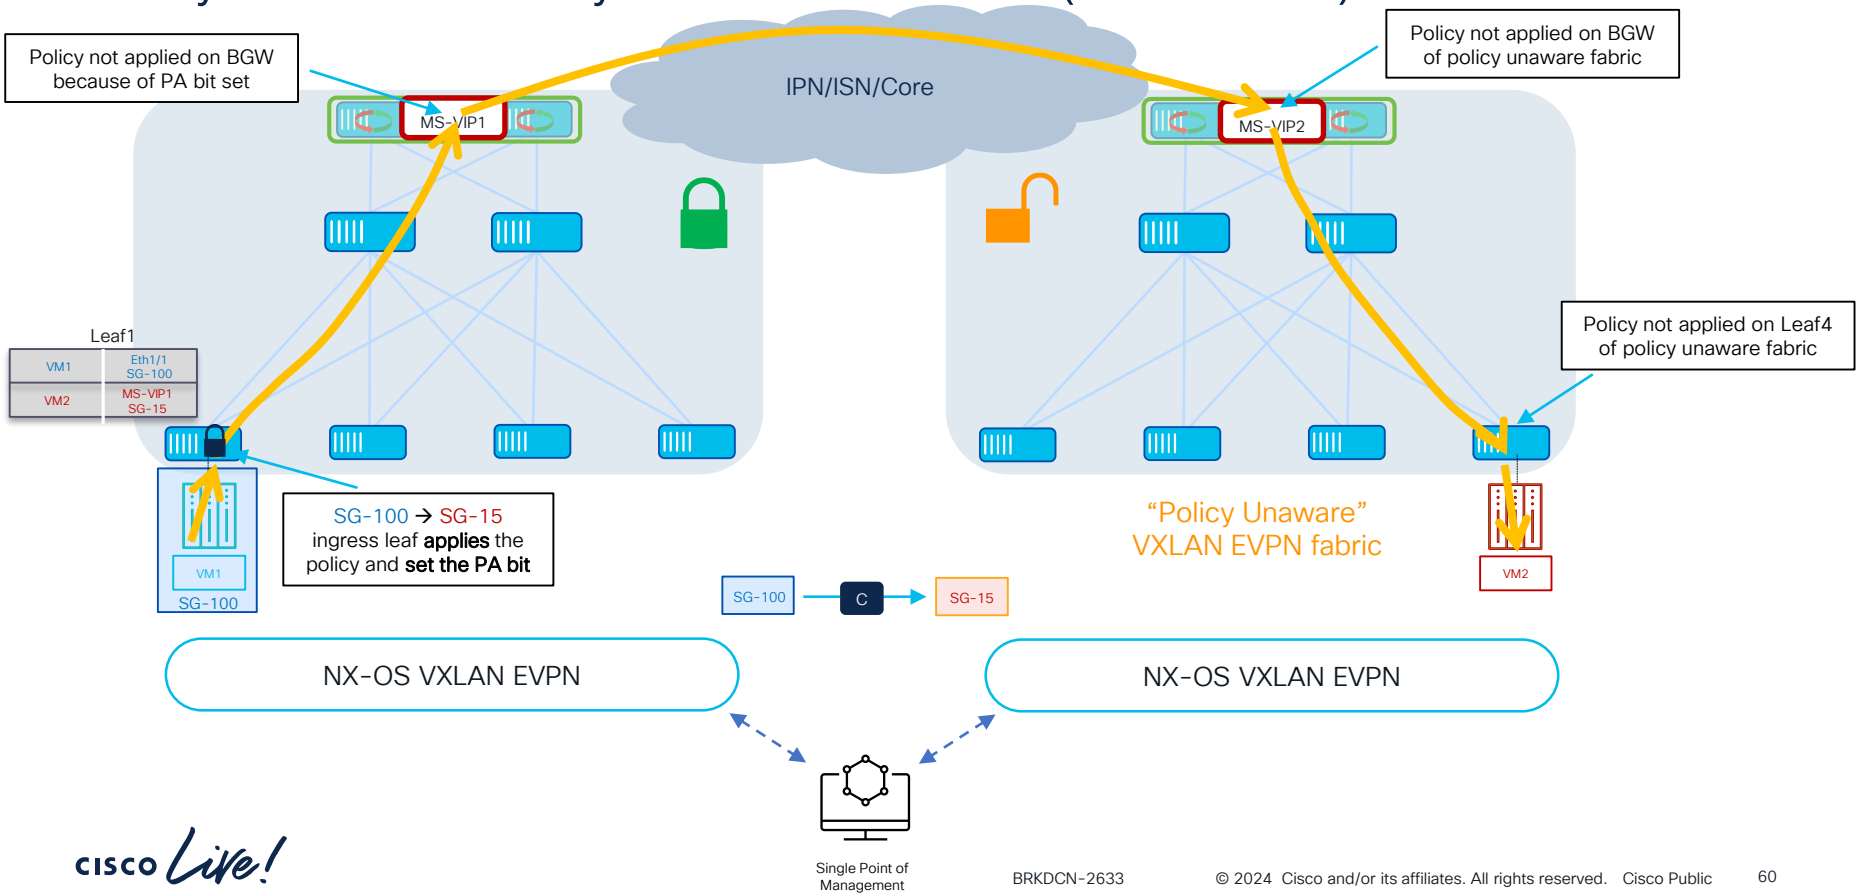

VXLAN GPO with Multi-Site

Policy Aware to Policy Unaware Fabrics (Data Plane)

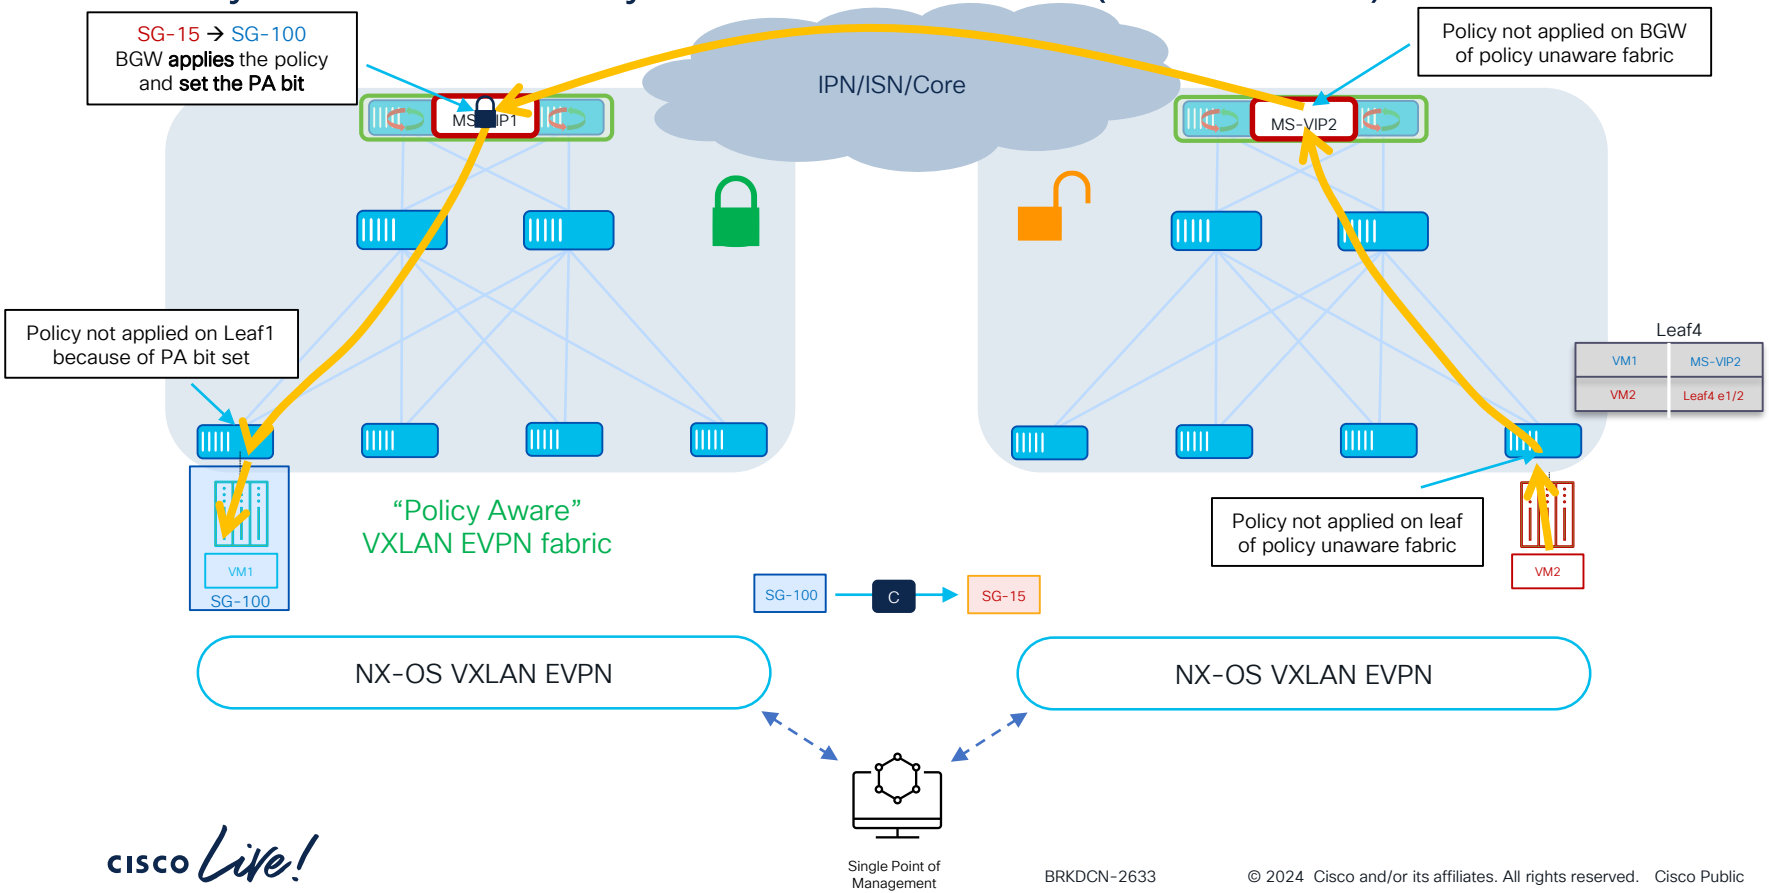

Secure Interconnection of Heterogeneous Fabrics

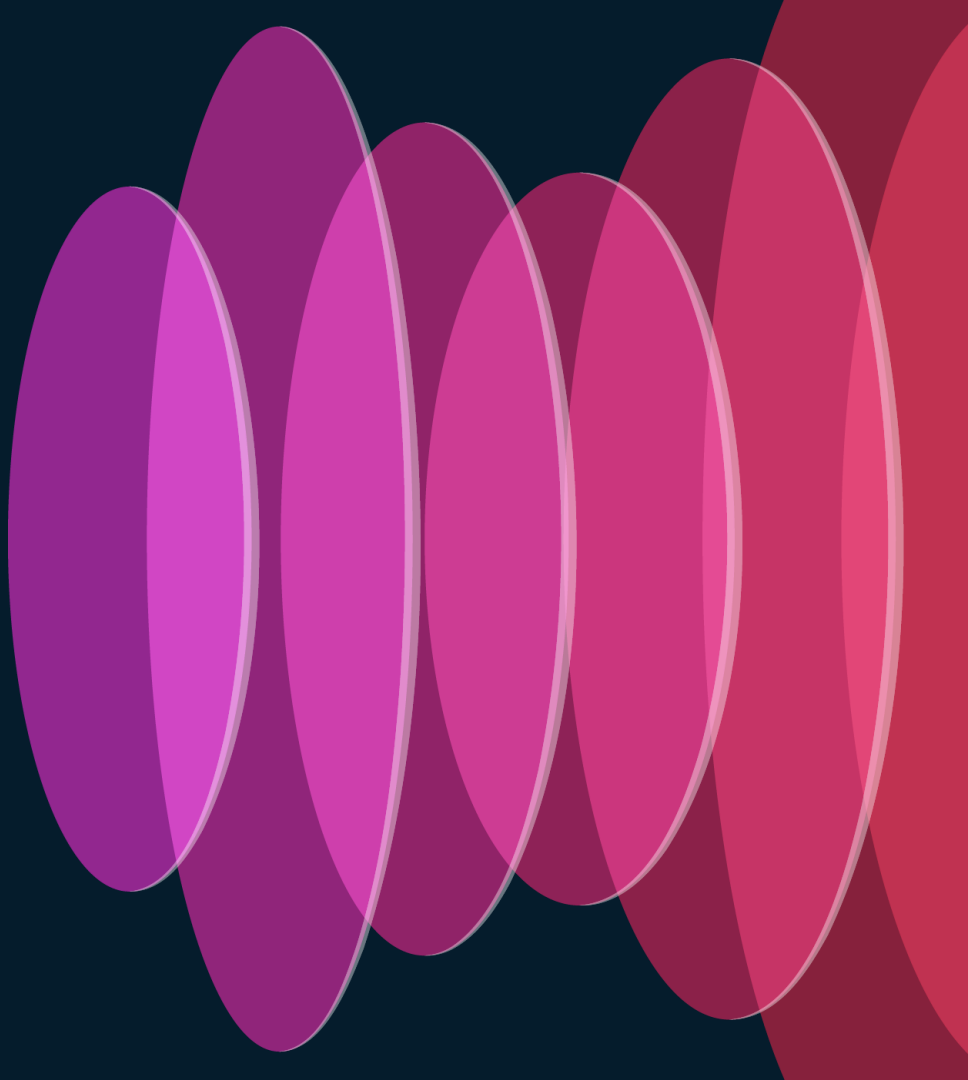

Secure Interconnection of Fabrics
Homogeneous Options

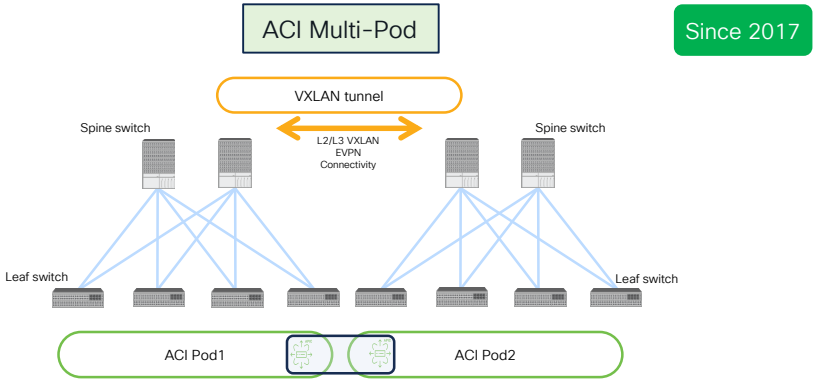

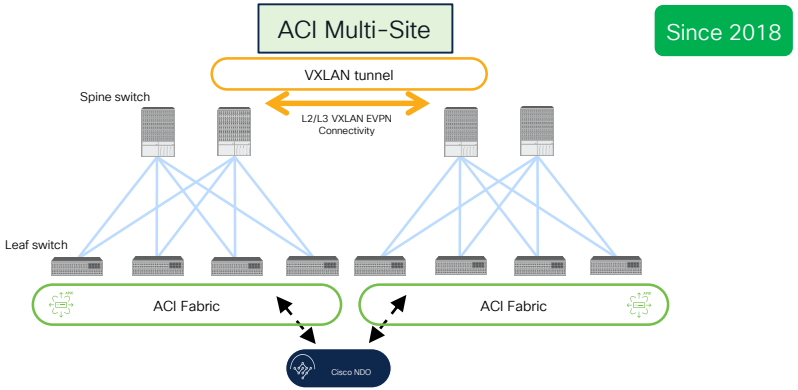

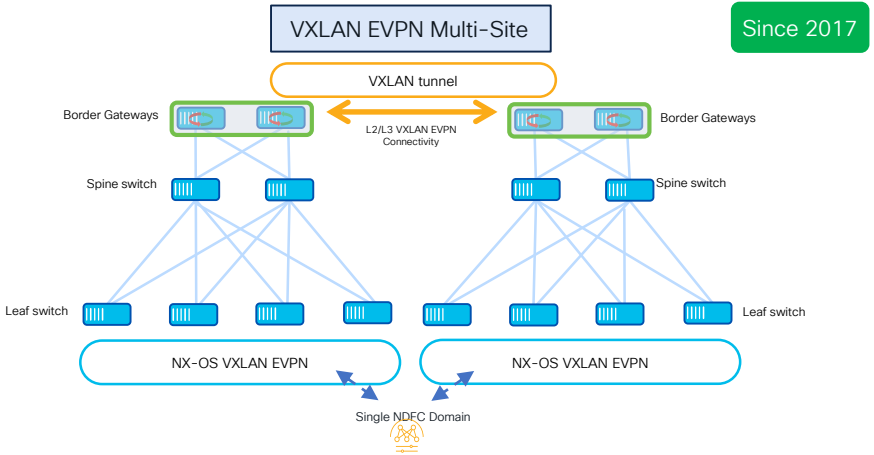

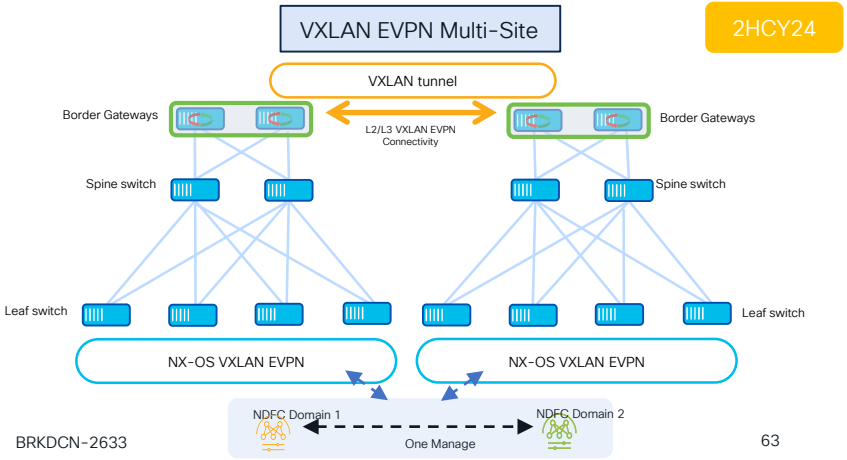

Heterogeneous Fabrics
Introducing ACI Border Gateways

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

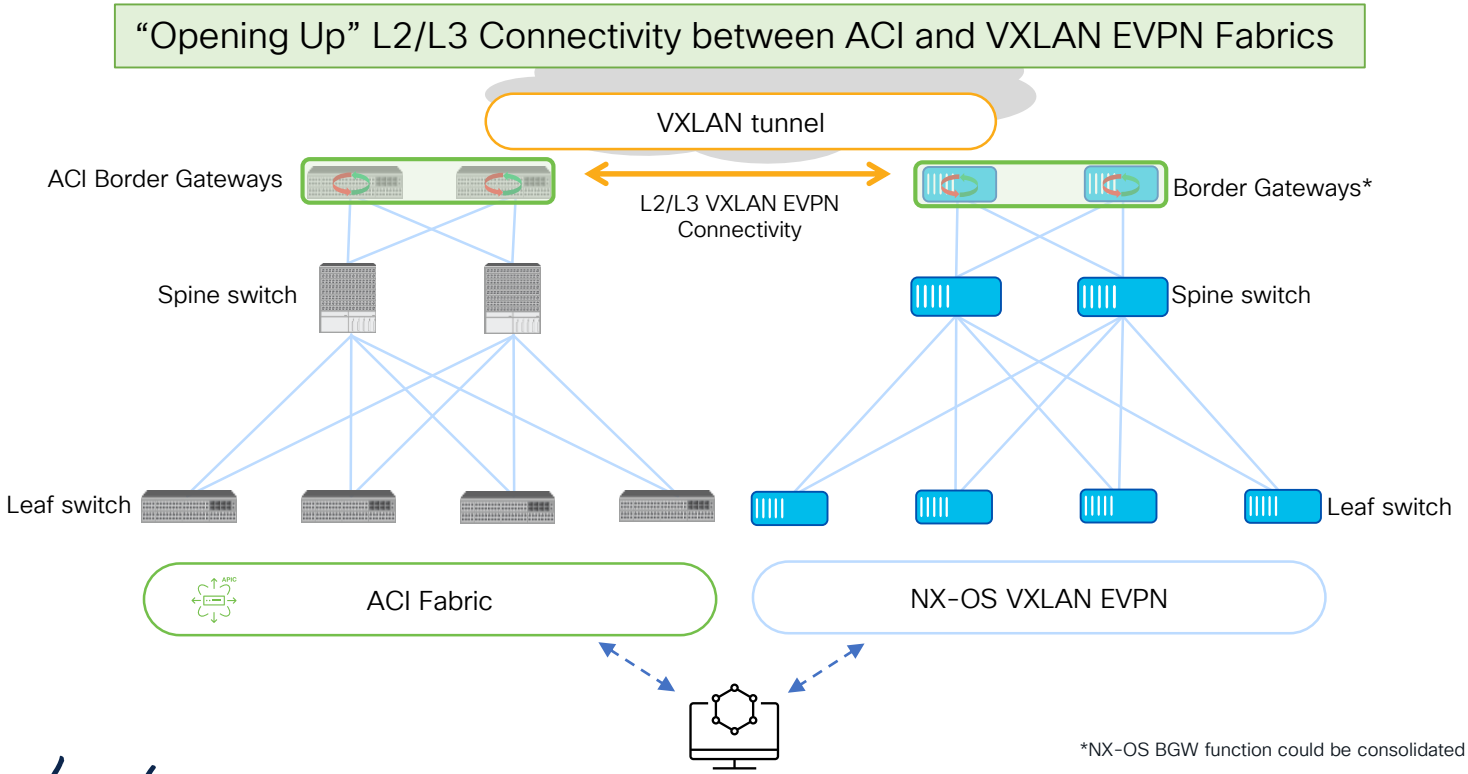

*NX-OS BGW function could be consolidated on the spines if desired

Heterogeneous Fabrics
Policy Enforcement End-to-End

Future

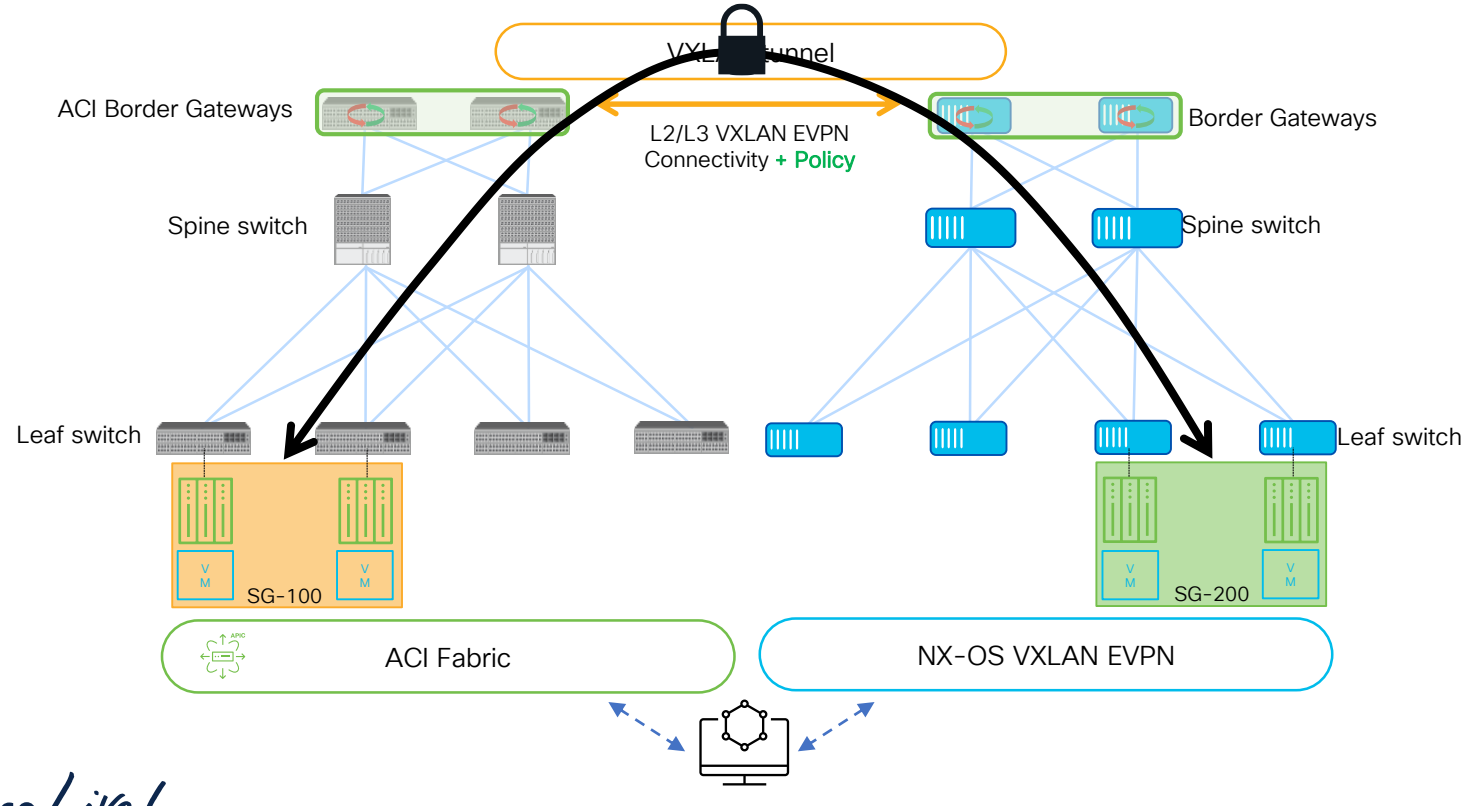

Conclusions

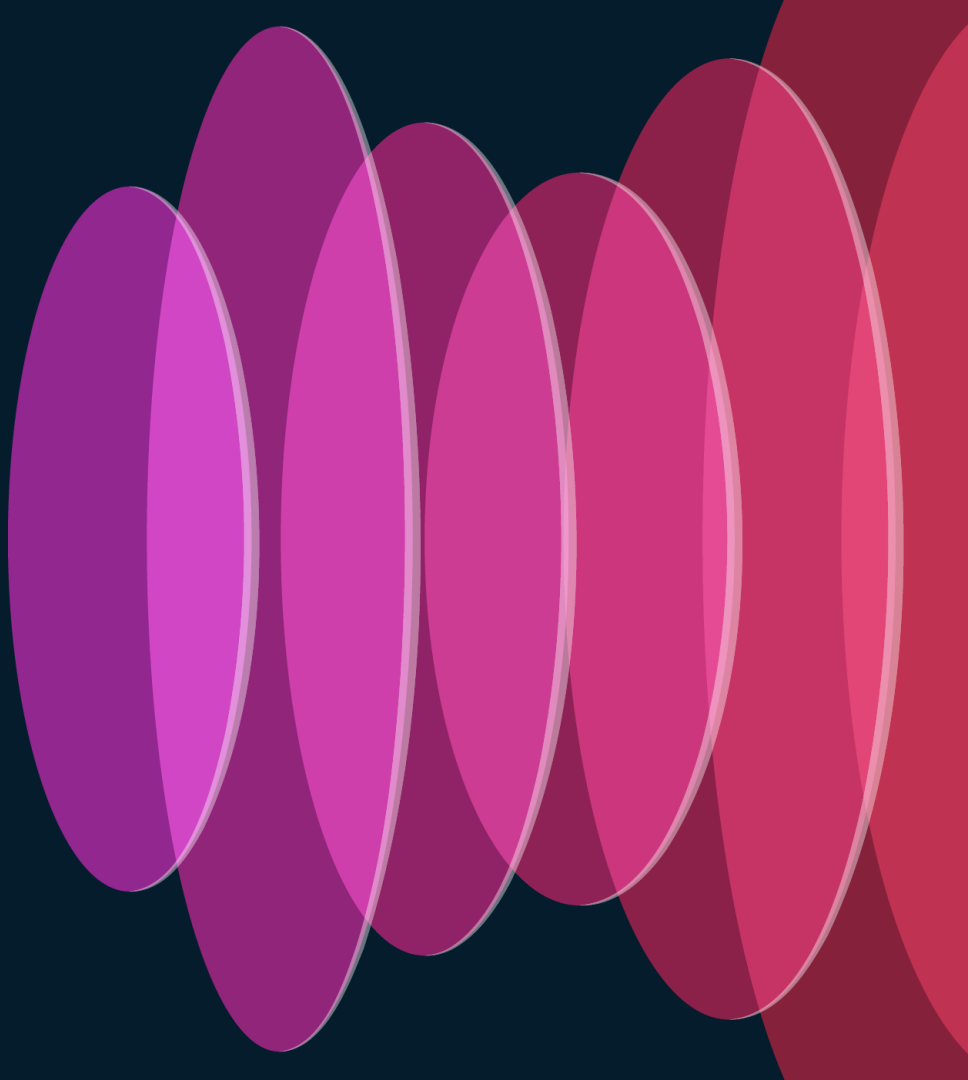

Conclusions

- The introduction of GPO in VXLAN EVPN fabrics provides policy enforcement and redirection capabilities between different secure groups
- Cisco One Fabric Experience aims to seamlessly and securely 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

- 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

Thank you