

CISCO *Live!*



#CiscoLive



The bridge to possible

VXLAN BGP EVPN Multi-Site

Architecture and Technology

Lukas Krattiger, Distinguished Engineer

@CCIE21921

BRKDCN-2913



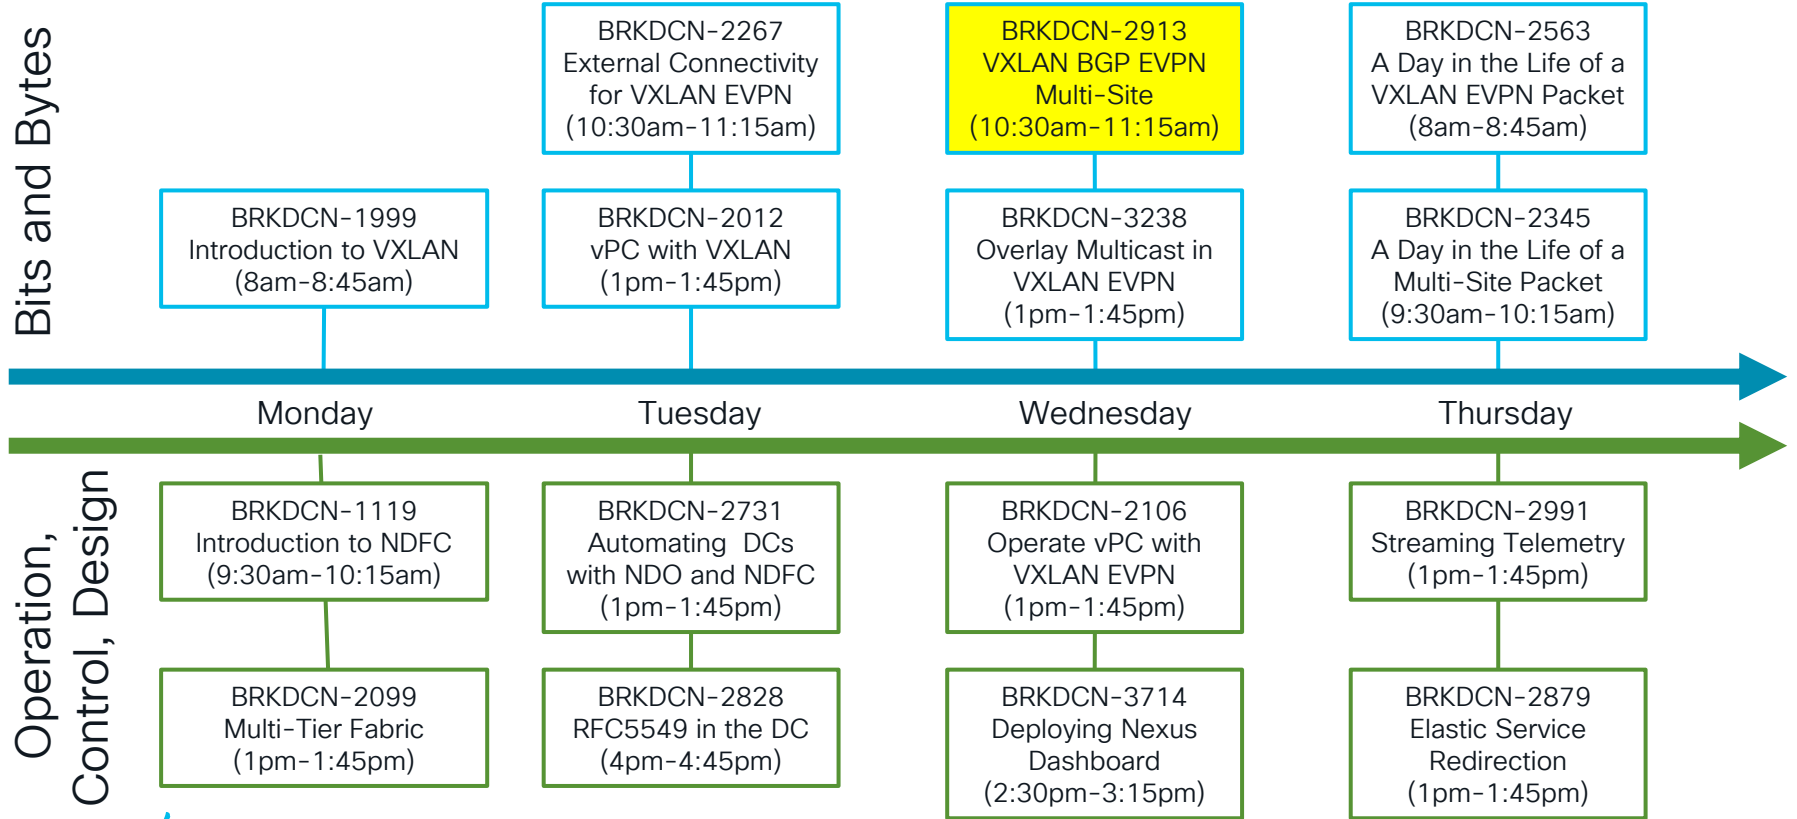
#CiscoLive

Abstract

VXLAN provides an excellent encapsulation for many applicability and EVPN provides extensive capabilities as a control-plane. With VXLAN and EVPN, we have excellent capabilities for Data Center fabric deployments with an integrated Layer-2 + Layer-3 approach. With the maturity of the control- and data-plane, new capabilities for interconnecting multiple fabrics is experiencing growing interest with VXLAN BGP EVPN. The goal of the session is describing the evolution of Cisco's approaches for efficiently connecting separate VXLAN BGP EVPN fabrics. We'll start by introducing the difference to Multi-Pod and Multi-Fabric, with the goal of understanding the evolution to the VXLAN Multi-Site approach that represents the optimal and recommended way to interconnect separate VXLAN EVPN islands.

Important Note: the session is exclusively focused on NX-OS standalone VXLAN EVPN and does not discuss the Multi-Pod and Multi-Site solutions offered with Cisco ACI.

Companion Sessions – Week at a Glance



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 17, 2022.



<https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-2913>



Agenda

- Introduction
- What is Multi-Site?
- Use Cases
- Multi-Site – A Deeper Look
- Conclusion

Introduction

Introduction

- A brief touchpoint of the work at the IETF (Internet Engineering Task Force) and what RFC (Request for Comment) are Standard and what Informational
- What is this Multi-Site about – why do we have it and what is it good for
- Use Cases and Deployment Scenarios
- The Border Gateway (BGW)

What is Multi-Site

What is RFC 9014?

By the Standards Body

[Search] [txt] [html] [xml] [pdf] [bibtex] [Tracker] [WG] [Email] [Diff1] [Diff2] [Nits]
From: draft-ietf-bess-dci-evpn-overlay-10 Proposed Standard
IPR declarations

Internet Engineering Task Force (IETF) J. Rabadan, Ed.
Request for Comments: 9014 S. Sathappan
Category: Standards Track W. Henderickx
ISSN: 2070-1721 Nokia
A. Sejassl
Cisco
J. Drake
Juniper
May 2021

Interconnect Solution for Ethernet VPN (EVPN) Overlay Networks

Abstract

This document describes how Network Virtualization Overlays (NVOs) can be connected to a Wide Area Network (WAN) in order to extend the Layer 2 connectivity required for some tenants. The solution analyzes the interaction between NVO networks running Ethernet Virtual Private Networks (EVPNs) and other Layer 2 VPN (L2VPN) technologies used in the WAN, such as Virtual Private LAN Services (VPLSs), VPLS extensions for Provider Backbone Bridging (PBB-VPLS), EVPN, or PBB-EVPN. It also describes how the existing technical specifications apply to the interconnection and extends the EVPN procedures needed in some cases. In particular, this document describes how EVPN routes are processed on Gateways (GWs) that interconnect EVPN-Overlay and EVPN-MPLS networks, as well as the Interconnect Ethernet Segment (I-ES), to provide multihoming. This document also describes the use of the Unknown MAC Route (UMR) to avoid issues of a Media Access Control (MAC) scale on Data Center Network Virtualization Edge (NVE) devices.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in [Section 2 of RFC 7944](#).

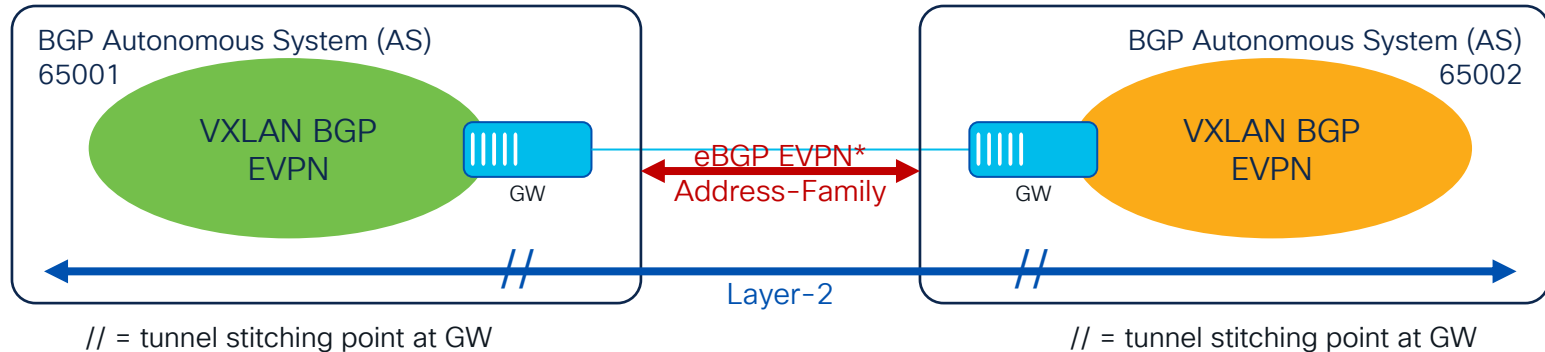
Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9014>.

- Internet Engineering Task Force (IETF) Request for Comment (RFC)
- Categorized for Standards Track
- Internet Standard since 2021
- Existing Industry Adoption
- Interconnect Solution for Ethernet VPN (EVPN) Overlay Networks
- Co-Authored by Cisco
- RFC 9014
 - <https://datatracker.ietf.org/doc/html/rfc9014>

RFC 9014 at a glance

By the Standards Body

- DCI EVPN Overlay (aka RFC 9014)
- Interconnect Solution for Ethernet VPN (EVPN) Overlay Networks
- From the Abstract “*extend the Layer 2 connectivity required for some tenants.*”




*RFC 9014 supports more than just EVPN for the Interconnect Network


RFC 9014 Gateway Model Side-by-Side

Decoupled and Integrated Gateway

Decoupled Gateway (Section 3)



Integrated Gateway (Section 4)




*RFC 9014 supports more than just EVPN for the Interconnect Network


RFC 9014 Gateway Model Side-by-Side

Decoupled and Integrated Gateway

Decoupled Gateway (Section 3)



Integrated Gateway (Section 4)



What about Layer-3?



Multi-Site Solution for Ethernet VPN (EVPN) Overlay

draft-sharma-bess-multi-site-evpn

What is Multi-Site?

By the Standards Body

[Search] [txt] [pdf] [bib] [Tracker] [Email] [Diff1] [Diff2] [Nits]
Versions: 00 01 02

INTERNET-DRAFT
Intended Status: Informational
Expires: November 13, 2022

L. Krattiger, Ed.
A. Banerjee, Ed.
A. Sajassi
R. Sharma
R. Sivaramu
Cisco Systems
May 12, 2022

Multi-Site Solution for Ethernet VPN (EVPN) Overlay draft-sharma-bess-multi-site-evpn-02

Abstract

This document describes the procedures for interconnecting two or more Network Virtualization Overlays (NVOs) via NVO over IP-only network. The solution interconnects Ethernet VPN network by using NVO with Ethernet VPN (EVPN) to facilitate the interconnect in a scalable fashion. The motivation is to support extension of Layer-2 and Layer-3, Unicast & Multicast, VPMs without having to rely on typical Data Center Interconnect (DCI) technologies like MPLS/VPLS. The requirements for the interconnect are similar to the ones specified in [RFC7289] "Requirements for Ethernet VPN (EVPN)". In particular, this document describes the difference of the Gateways (GWs) procedure and incremental functionality from [RFC9014] "Interconnect Solution for Ethernet VPN (EVPN) Overlay Networks", which this solution is interoperable to. This document updates and replaces all previous version of [SHARMA-MULTI-SITE].

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at


Sharma, et al. Expires November 13, 2022 [Page 1]

- Internet Engineering Task Force (IETF) Request for Comment (RFC)
- Categorized as Informational
- Internet Draft since 2016
 - Currently in Version 2
 - Overall, 7 versions
- Updated and Maintained by BESS version of draft
 - draft-sharma-bess-multi-site-evpn
- Shipping since 2017
- Multi-Site (BESS version)
 - <https://datatracker.ietf.org/doc/html/draft-sharma-bess-multi-site-evpn>
- Pre-Cursor Draft (replaced by BESS version)
 - <https://datatracker.ietf.org/doc/html/draft-sharma-multi-site-evpn>

Multi-Site

By the Standards Body

- Multi-Site Solution for Ethernet VPN (EVPN) Overlay (draft-sharma-bess-multi-site-evpn)
- Interconnect Solution for Ethernet VPN (EVPN) Overlay Networks
- From the Abstract “*support extension of Layer-2 and Layer-3, Unicast & Multicast, VPNs*”




RFC9014 and Multi-Site - Side by Side

	DCI-EVPN-Overlay (RFC 9014)	Multi-Site EVPN (draft-sharma-bess-multi-site-evpn)	
Interconnect	Integrated (1-Box), Decoupled (2-Box)	Integrated (1-Box)	
DCI Encap	VPLS, PBB-VPLS, EVPN-MPLS, PBB-EVPN, VXLAN	VXLAN	
Gateway Mode	Multipath PIP	Anycast VIP	Multipath PIP
ECMP	Underlay and Overlay	Underlay	Underlay and Overlay
EVPN RT-1	Consumed and Generated	None	Consumed and Generated
EVPN RT-2	Re-Originated with I-ESI	Re-Originated with ESI 0	Re-Originated with I-ESI
EVPN RT-3	Consumed and Generated	Consumed and Generated	Consumed and Generated
EVPN RT-4	Consumed and Generated	Consumed and Generated	Consumed and Generated
EVPN RT-5	- (not part of RFC)	Re-Originated	Re-Originated
Route Distinguisher (RD)	Separate RD for Intra and Inter DC	Separate RD for VIP and PIP	
Route-Target (RT)	Separate RT for Intra and Inter DC	Same RT for Intra and Inter DC	
VNI Allocation	Global and Downstream	Global and Downstream	
DF Election	Based on EVPN RT-4	Based on EVPN RT-4	
Identifier	I-ESI	I-ESI (= Site-ID)	
Split Horizon	Local Bias	Local Bias	
ESI-Type	Type 0 (Operator Managed)	Type 3 (MAC Based) or Type 5 (AS based)	
BUM Tree #	2, GW stitched (Intra and Inter DC)	2, GW stitched (Intra and Inter DC)	


RFC9014 and Multi-Site – Side by Side

In a Nutshell



RFC 9014

- Base Standard for Interconnecting EVPN
- Defines the Layer-2 Stitching
- Two Gateway Model
- Multiple Encapsulations
- Leverages Overlay and Underlay ECMP



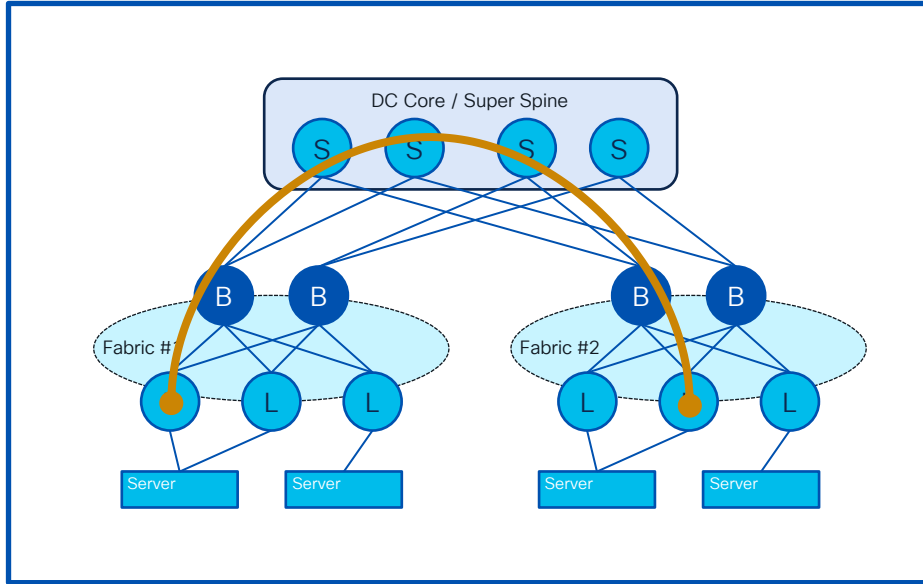
Multi-Site

- Extends RFC 9014 for Interconnecting EVPN
- Describes Layer-2 and Layer-3 Stitching
- Single Gateway Model (Two BGW* Model)
- Focus only on VXLAN Encapsulation
- Different ECMP model depending on BGW Model

*BGW – Border Gateway (BGW); Cisco's name for the VXLAN EVPN to VXLAN EVPN Gateway

Use Cases

Use Case #1 - Compartmentalization

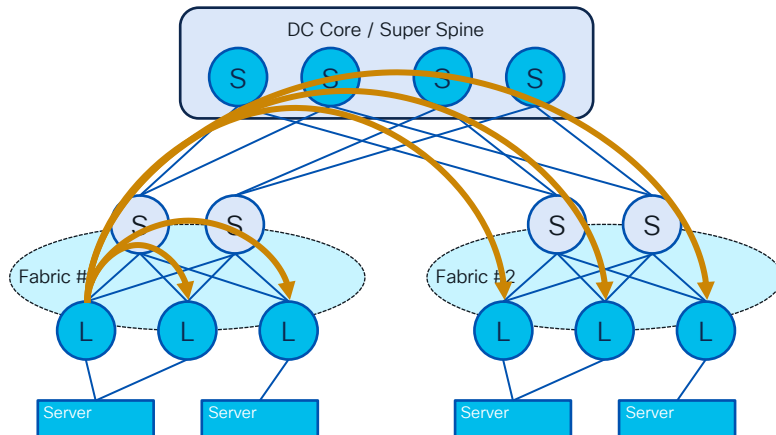


- Multiple Fabrics, single Data Center
 - Single or Multiple Data Halls
 - Within a Geographic Locations
- Control at BGW (Border Gateway)
 - Allows Extension of Layer-2
 - Allows Extension of Layer-3
 - Allows Extension of Layer-2 and Layer-3
 - Allows Traffic Control (BUM*)
 - Defines VNI allocation and stitching
 - Optimizes BUM* Replication

BUM Optimization

Use Case #1 – Compartmentalization


Single Fabric BUM with Ingress Replication



Single BUM Packet, 5x Replicated

3 Replication over DC Core / Super Spine (Between)
2 Replication for Fabric #1 (Local)

Multi-Site BUM with Ingress Replication

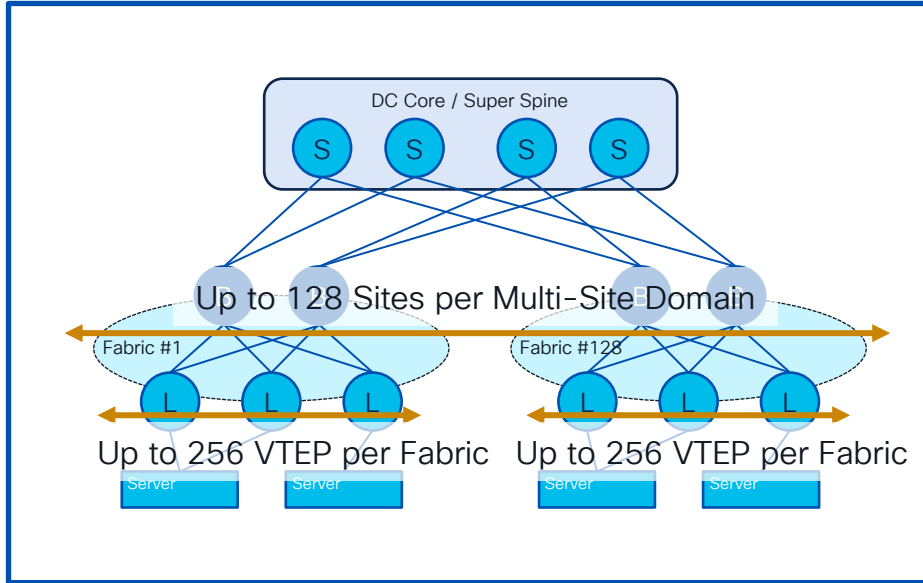


Single BUM Packet, 3x Replicated

1 Replication over DC Core / Super Spine (Between)
3 Replication for Fabric #1 (Local)
3 Replication for Fabric #2 (Local)

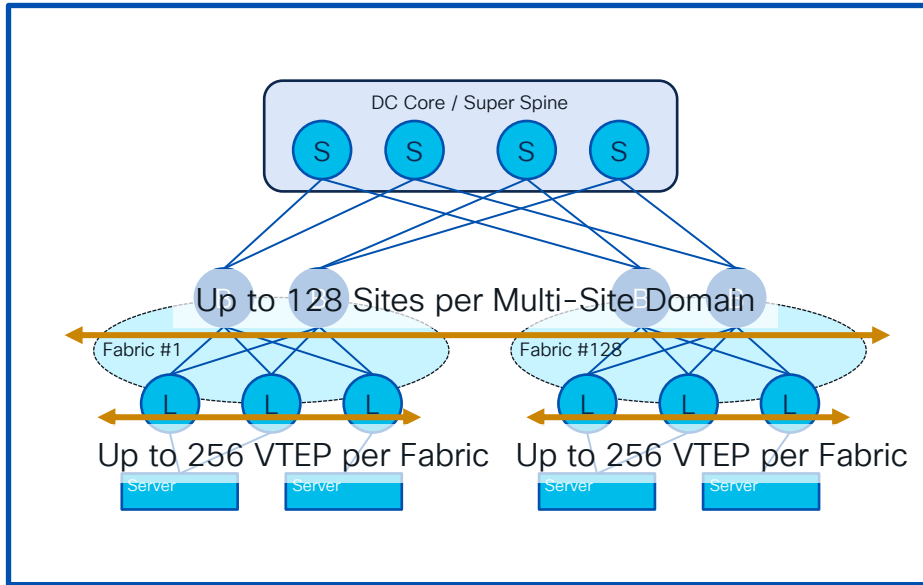
*BUM – Broadcast, Unknown Unicast, Multicast

Use Case #2 - Scale



- Multiple Fabrics , single or multiple Data Center
 - Single or Multiple Data Halls
 - Within or between Geographic Locations
- Control at BGW (Border Gateway)
 - Reduces Remote VTEP Count
 - Expands VTEP scale
- Scale through Hierarchy
 - Multiply VTEP with Sites

Use Case #2 - Scale



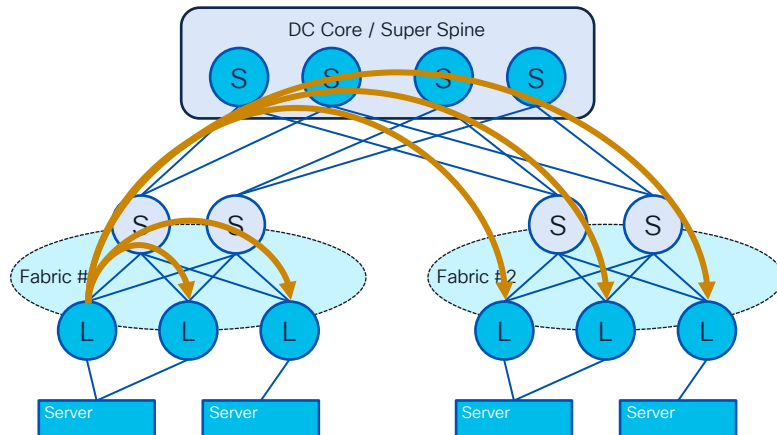
- Multiple Fabrics , single or multiple Data Center
 - Single or Multiple Data Halls
 - Within or between Geographic Locations
- Control at BGW (Border Gateway)
 - Reduces Remote VTEP Count
 - Expands VTEP scale
- Scale through Hierarchy
 - Multiply VTEP with Sites

32'768 VTEP to extend Layer-2 or/and Layer-3 segments to

VTEP Scale

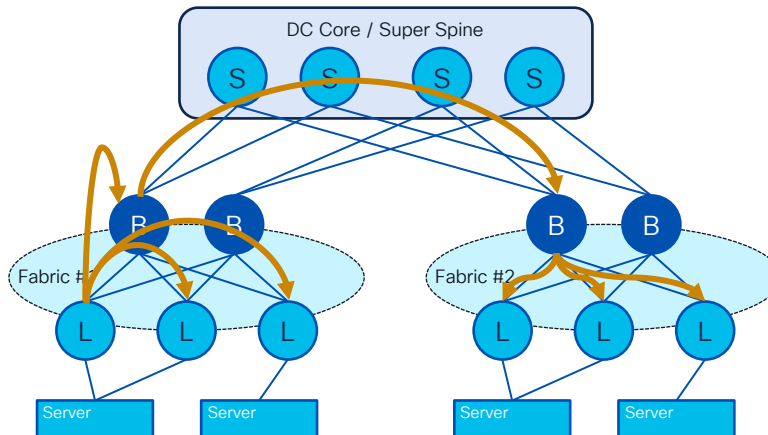
Use Case #2 - Scale

Single Fabric or Multi-POD



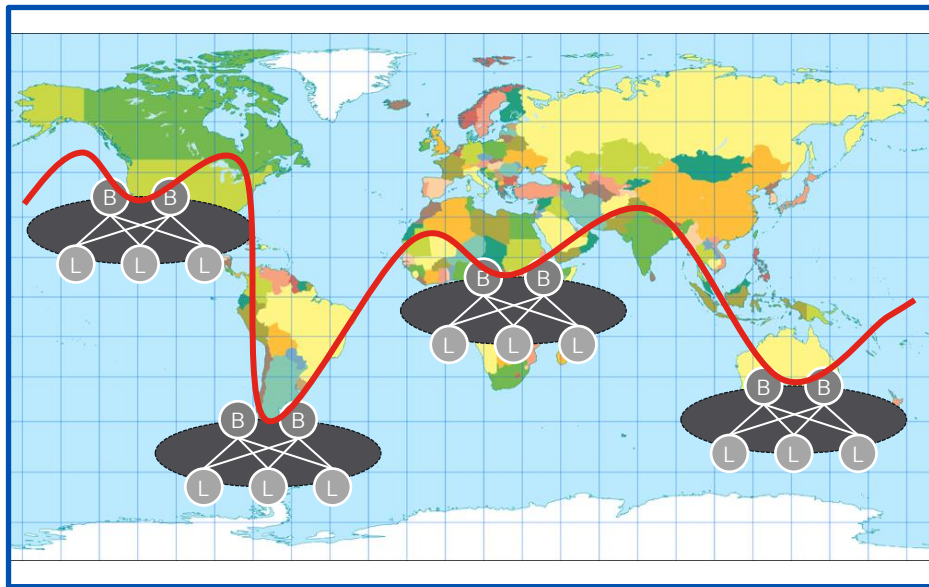
Leaf #1 sees every VTEP, 5 VTEP Peer
3 VTEP Peer for Fabric #2 (Between)
2 VTEP Peer for Fabric #1 (Local)

Multiple Fabric with Multi-Site



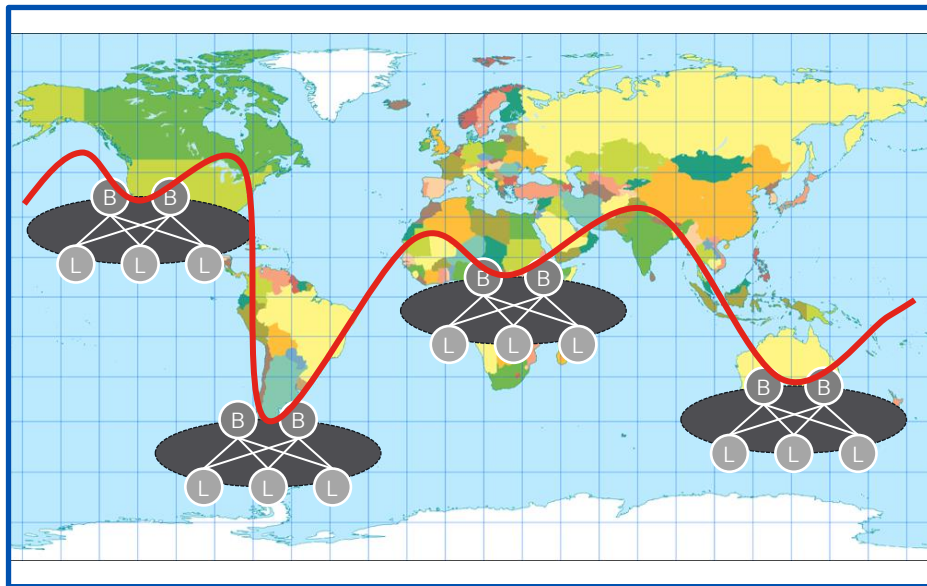
Leaf #1 sees only local VTEP, 3 VTEP Peer
1 VTEP Peer for Exit, BGW (Between)
2 VTEP Peer for Fabric #1 (Local)

Use Case #3 – Data Center Interconnect (DCI)



- Multiple Fabrics, Geographically Dispersed
- Classic DCI Use Case
 - Allows Extension of Layer-2
 - Allows Extension of Layer-3
 - Allows Extension of Layer-2 and Layer-3
 - Allows Traffic Control (BUM*)
 - Defines VNI allocation and stitching
 - Optimizes BUM* Replication

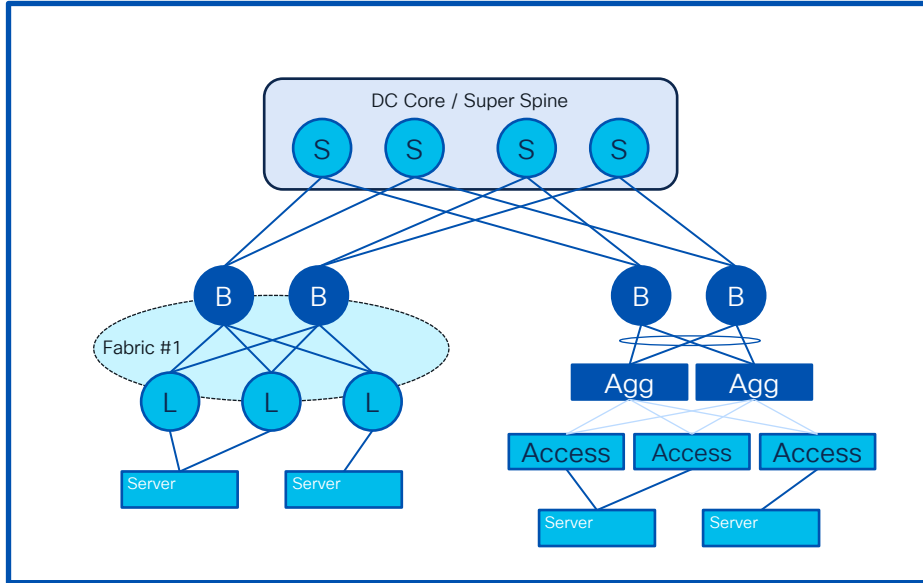
Use Case #3 – Data Center Interconnect (DCI)



- Multiple Fabrics, Geographically Dispersed
- Classic DCI Use Case
 - Allows Extension of Layer-2
 - Allows Extension of Layer-3
 - Allows Extension of Layer-2 and Layer-3
 - Allows Traffic Control (BUM*)
 - Defines VNI allocation and stitching
 - Optimizes BUM* Replication


Works Within a Geographic Location – Works Between Geographic Locations

Use Case #4 – Integration with Legacy Networks



- Integrating Fabrics with Legacy Networks
 - BGW Frontends Legacy Network
 - BGW Frontends New Network
- Host Mobility and Migration
 - Provides Distributed Default Gateway
 - Allows Layer-2 Extension where needed
- Benefits from all Multi-Site functions
 - Layer-2, Layer-3 Multicast and Unicast VPNs between different Networks for Migration or Co-Existence

Use Case #5 – Hybrid Cloud Connectivity



- Extending from On-Prem to Public Cloud
 - Using BGW to terminate VXLAN EVPN Tunnel from Cloud
 - Transported via Internet or Direct Connect / Express Route
 - Encryption Optional (IPSec)
- Multi-Tenancy to the Cloud
 - Extends Layer-3 with VRF awareness
 - Multiple On-Prem VRF mapped or integrated with Virtual Cloud Instances (ie VPC or VNET)
- Leveraging Cisco Cloud Router
 - C8kv (Catalyst 8000 Virtual) in the Cloud as termination point
 - Replacing CSR1000v

Multi-Site A Deeper Look

As we Talk about Scale

Hardware Support

Minimum Hardware and Software Requirements for BGW (Border Gateway)

Cisco Nexus Hardware	Cisco Nexus 9300 EX platform Cisco Nexus 9300 FX platform Cisco Nexus 9300 FX2 platform Cisco Nexus 9300 FX3 platform Cisco Nexus 9300 GX platform Cisco Nexus 9300 GX2 platform Cisco Nexus 9364C platform Cisco Nexus 9332C platform Cisco Nexus 9500 platform with X9700-EX line card Cisco Nexus 9500 platform with X9700-FX line card Cisco Nexus 9500 platform with X9700-GX line card
Cisco Nexus Software (NX-OS)	Cisco NX-OS Software Release 7.0(3)I7(1) or later*

*Check for Hardware Specific Support Releases

As we Talk about Scale

Scalability Values as of NX-OS 10.2(3)F

Multi-Site Scale

Number of Sites	128
Number of BGW per Site	6
Number of VTEP per Site (internal)	256


Border Gateway (BGW) Scale	EX	FX2	FX,FX3,GX,GX2	N9364C & N9332C
Number of Layer-2 VNI (VLAN)	3900			
Number of Layer-3 VNI (VRF)	2000			
MAC per BGW	92k			
IPv4 Host Routes per BGW*	450k	450k	1.1M	96k
IPv4 Network Routes per BGW*	450k	450k	1.1M	8k
IPv6 Host Routes per BGW*	24k	260k	620k	48k
IPv6 Network Routes per BGW*	200k	290k	620k	2k

*The values provided in these tables focus on the scalability of one particular Route scale at a time

Some Notes on BGW and VXLAN Tunnels

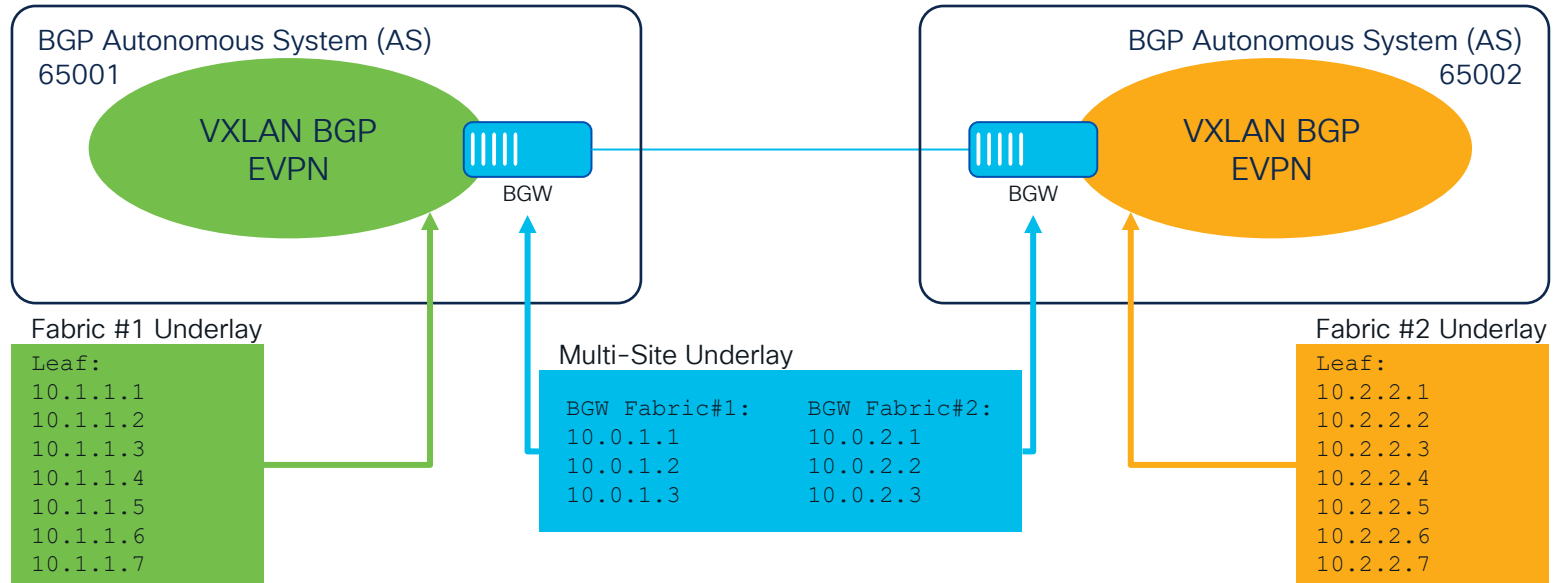
Multi-Site

- Tunnels are Stitched at the BGW (Border Gateway)
- Intra Fabric Tunnel goes from Leaf to Leaf or Leaf to BGW
- Inter Fabric Tunnel goes from BGW to BGW



Some Notes on the Interconnect and Underlay Multi-Site

- Fabric #1 Underlay (VTEP, Point-2-Point, Loopback etc) is not aware of Fabric #2
- Each Fabric maintains their Unique Network Topology, Protocols and IP Addressing
- Only BGW IP Addressing must be Unique and Aligned between Sites






Border Gateway Deployment Considerations


Multi-Site – A Deeper Look

Border Gateways Deployment Considerations

Anycast Border Gateway




vPC Border Gateway



- Border Gateways used for two main functions:
 - Interconnecting each site to the Inter-Site network (for East-West traffic flows)
 - Connecting each site to the external Layer 3 domain (for North-South traffic flows)
 - May also be used to connect endpoints and/or network service nodes (FWs, ADCs)
- Possible deployment models:
 - Anycast Border Gateways
 - vPC Border Gateways
- BGW function enablement in the VXLAN EVPN fabric:
 - BGWs on Leaf node (Border Gateway Leaf)
 - BGWs on Spine node (Border Gateway Spine)


Anycast Border Gateway



Anycast Border Gateway

- Up to 6 Border Gateways
- Border Gateway
 - Deploying as a Leaf node – 7.0(3)I7(1)
 - Deploying as a Spine node – 7.0(3)I7(2)
- Two Mode of Operation:
 - Can Operate as Multi-Site Anycast BGW with VIP
 - Focuses on Scale and Convergence
 - Using Virtual IP (VIP) for Tunnel Stitching
 - Uses Overlay ECMP
 - Can Operate in RFC 9014 BGW Mode with PIP
 - Focuses on 3rd Party Interop
 - Using Primary IP (PIP) for Tunnel Stitching
 - Uses Underlay and Overlay ECMP


vPC Border Gateway



vPC Border Gateway


- Up to 2 Border Gateways
- Border Gateway
 - Deploying as a Leaf node – 9.2(1)
- Common Use Case
 - Legacy Network Integration or Migration
 - Provides Multi-Chassis Link Aggregation
 - Integrates with Ethernet and FabricPath
 - Hosts the Distributed Anycast Gateway
 - Attachment of Network Services
 - Dual-Attachment of Firewalls and ADCs
 - Acts like a vPC when it comes to Routing

When to use what BGW



Anycast Border Gateway

- Up to 6 BGW
 - Shared Nothing
 - Simple Failure Scenarios
- Any Deployments
 - No End-Point or Network Services Connectivity on BGW
- Greenfield Deployments




vPC Border Gateway

- 2 BGW with physical vPC Peer-Link
- Small Deployments
 - End-Point or Network Services Connectivity on BGW
- Migration Use-Cases (Brownfield)
 - Classic Ethernet/FabricPath to VXLAN EVPN


Details on the Different BGW

- Both Anycast and vPC Border Gateway needs to be configured with a common Multi-Site VIP address and an individual Primary IP (PIP) address
- vPC Border Gateways share a secondary IP address to be used as vPC virtual IP (vPC VIP)

Anycast Border Gateway



vPC Border Gateway






Control and Data Plane

Multi-Site – A Deeper Look


Control- and Data-Plane

- MP-eBGP EVPN for between Multiple Sites
- Leverages Natural BGP Next-hop behavior
 - VXLAN tunnel termination and re-origination
 - Loop protection (as-path attribute)
- Full mesh of MP-eBGP EVPN adjacencies between sites



Control- and Data-Plane

- Two main options for underlay and overlay Control-Plane (CP) deployment
 - I-E-I (Recommended)
 - Within Fabric: IGP (OSPF, IS-IS) as underlay CP, iBGP as overlay CP
 - E-E-E*
 - Within and Between Fabricis: eBGP for both underlay and overlay CPs

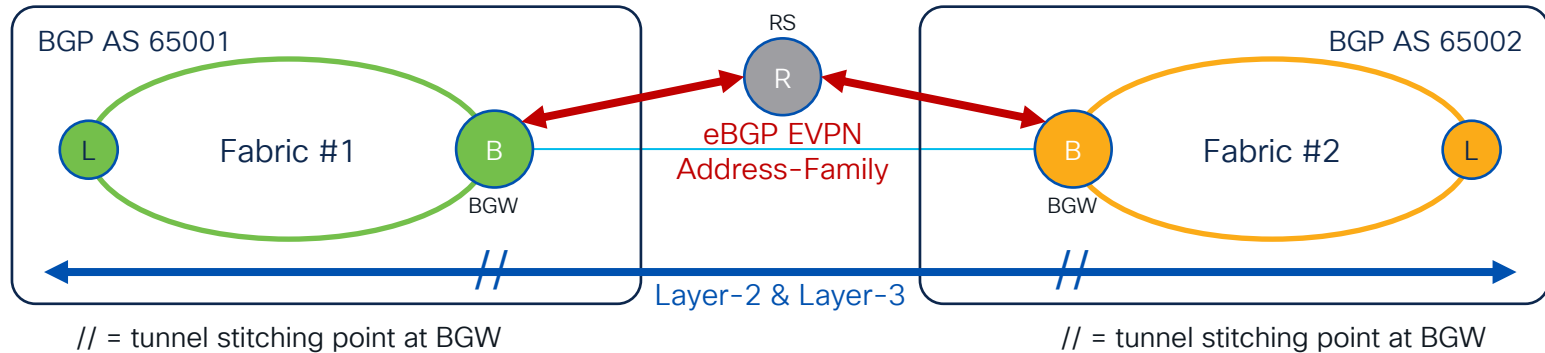


*For more information on why eBGP for both underlay and overlay CP is not a good idea:

https://learningnetwork.cisco.com/blogs/community_cafe/2017/10/17/the-magic-of-super-spines-and-rfc7938-with-overlays-guest-post


Control- and Data-Plane

- Full mesh of MP-eBGP EVPN adjacencies between sites
 - Recommended to deploy a couple of Route-Servers (RS) with 3 or more sites
- RS resides in a separate AS and only performs control plane functions
 - “eBGP Route-Reflectors”, IETF RFC 7947
- RS functions: EVPN routes reflection, next-hop-unchanged, route-target rewrite




It all starts with a Network Device

The Dating Network - When Control- meets Data-Plane




It all starts with a Network Device

The Dating Network - When Control- meets Data-Plane



Making the Network Device an NVE

The Dating Network - When Control- meets Data-Plane



Making the Network Device an NVE

The Dating Network - When Control- meets Data-Plane

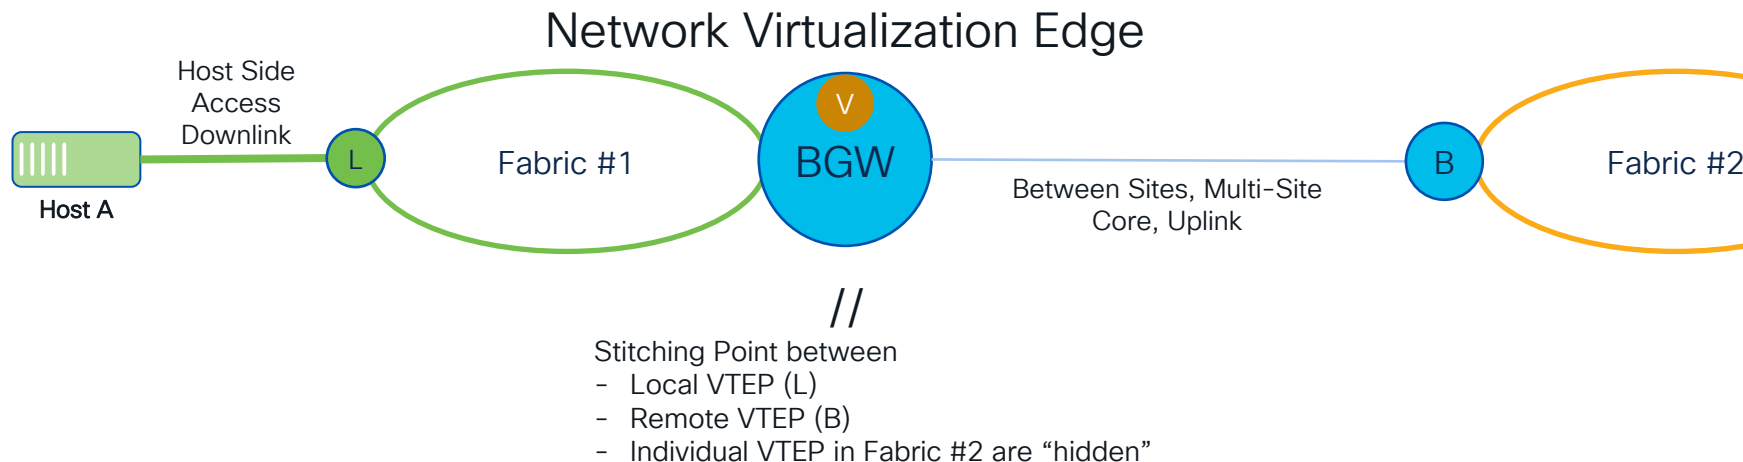


Still a ...

Switch
Router
Network Device
or Edge Device

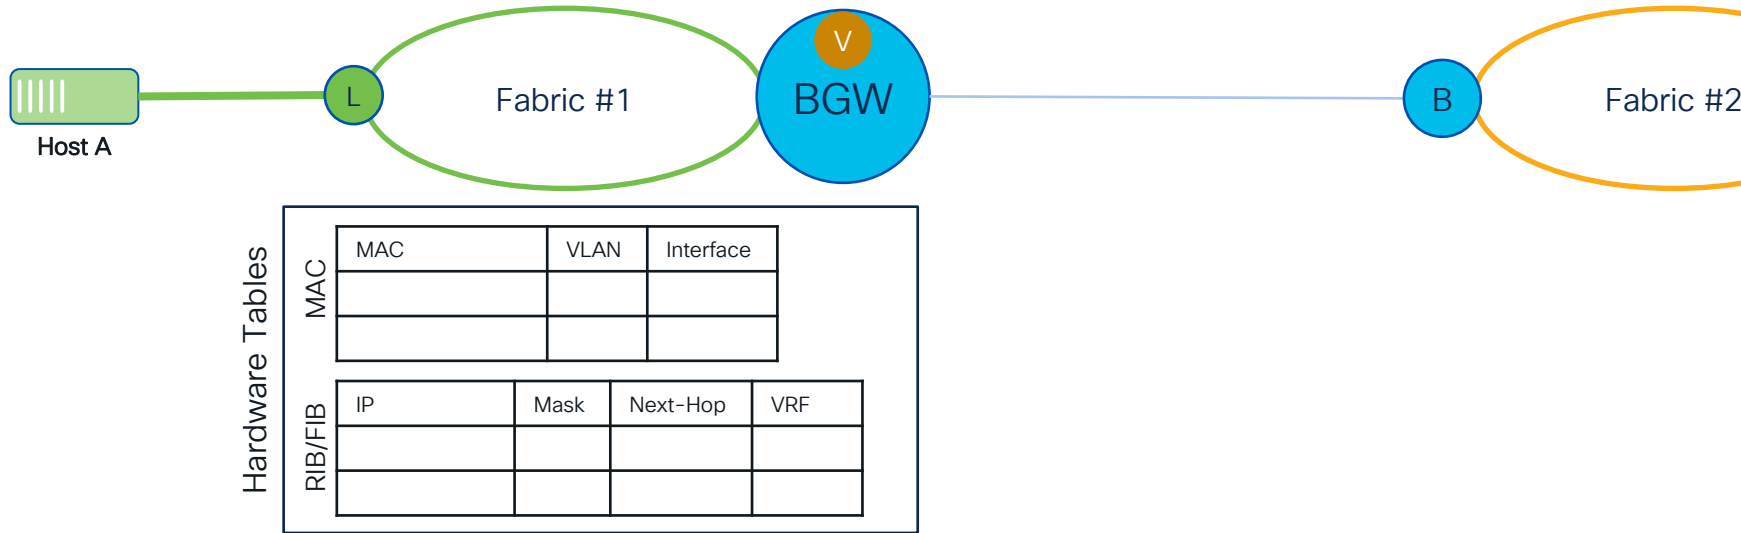
Expanding to the BGW – A “special” NVE

The Dating Network – When Control- meets Data-Plane



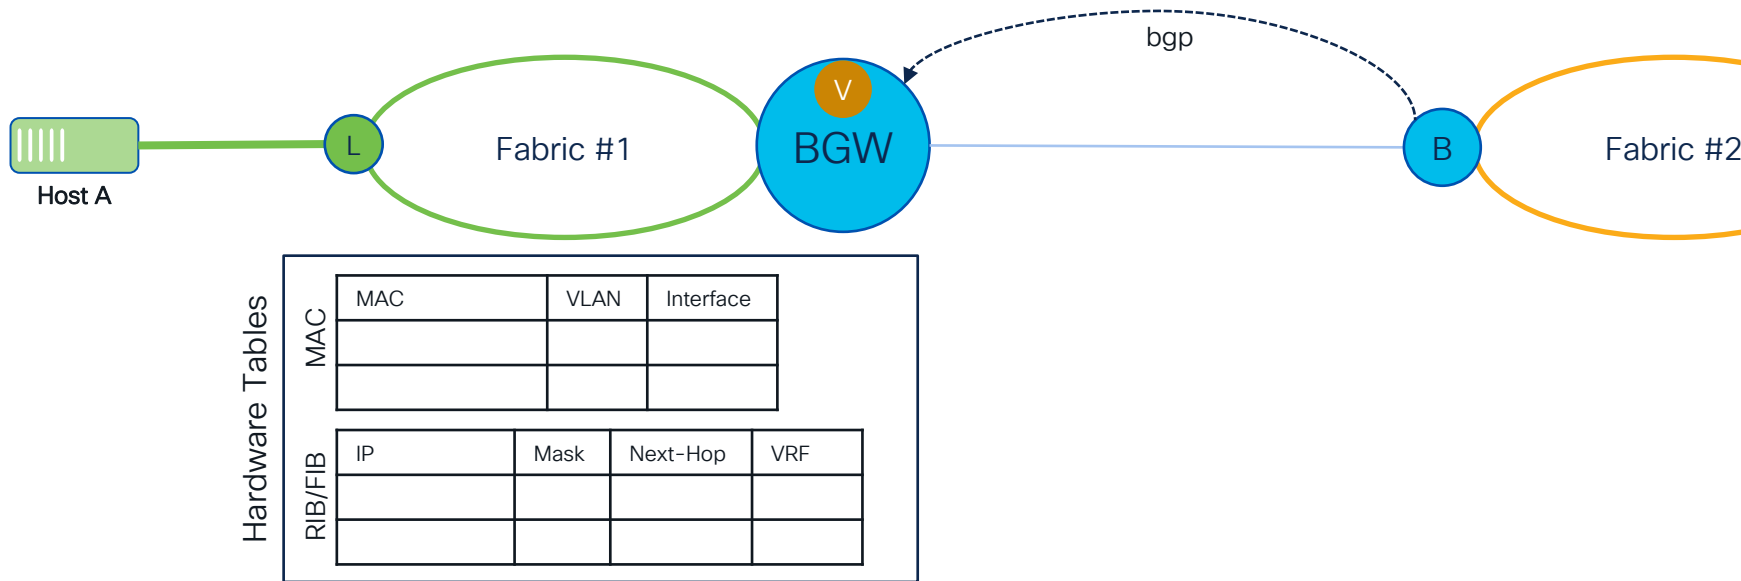
The BGW and Some Important Table

The Dating Network - When Control- meets Data-Plane



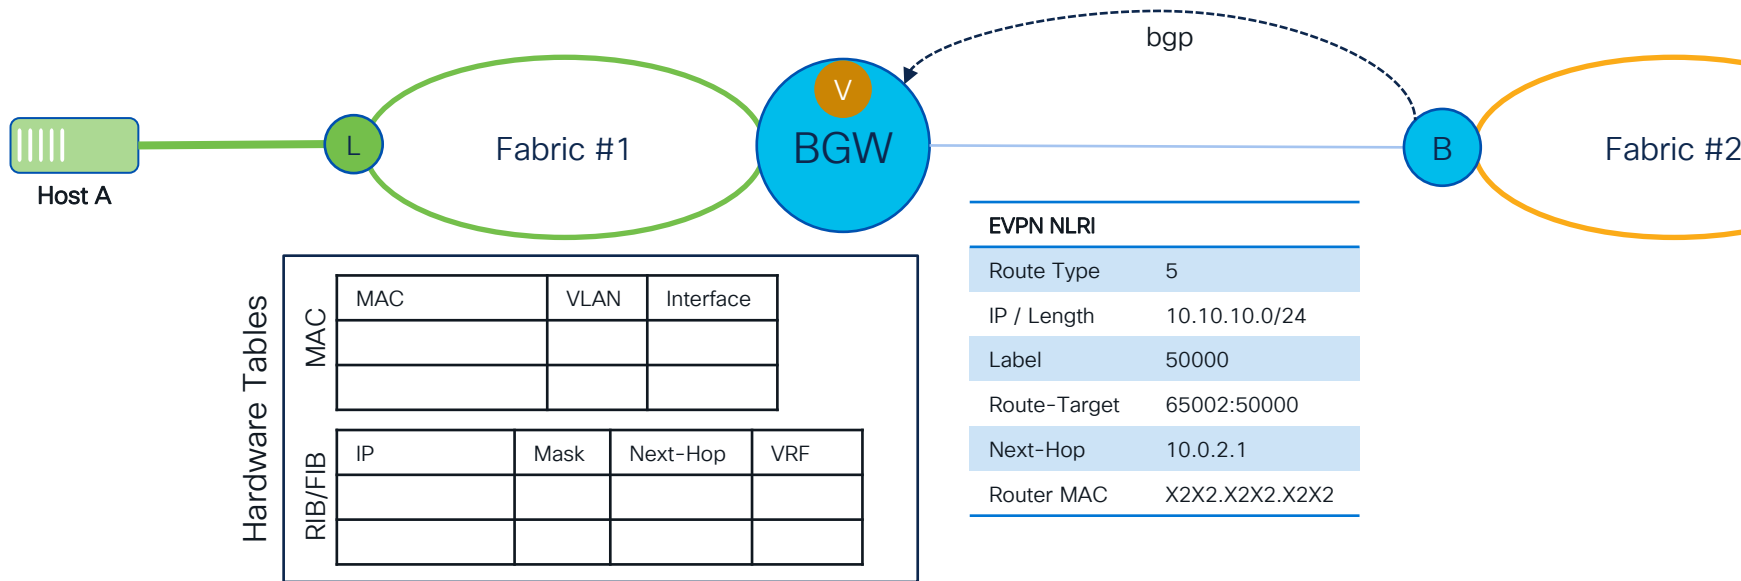
Always Remote Learning on a BGW

The Dating Network - When Control- meets Data-Plane



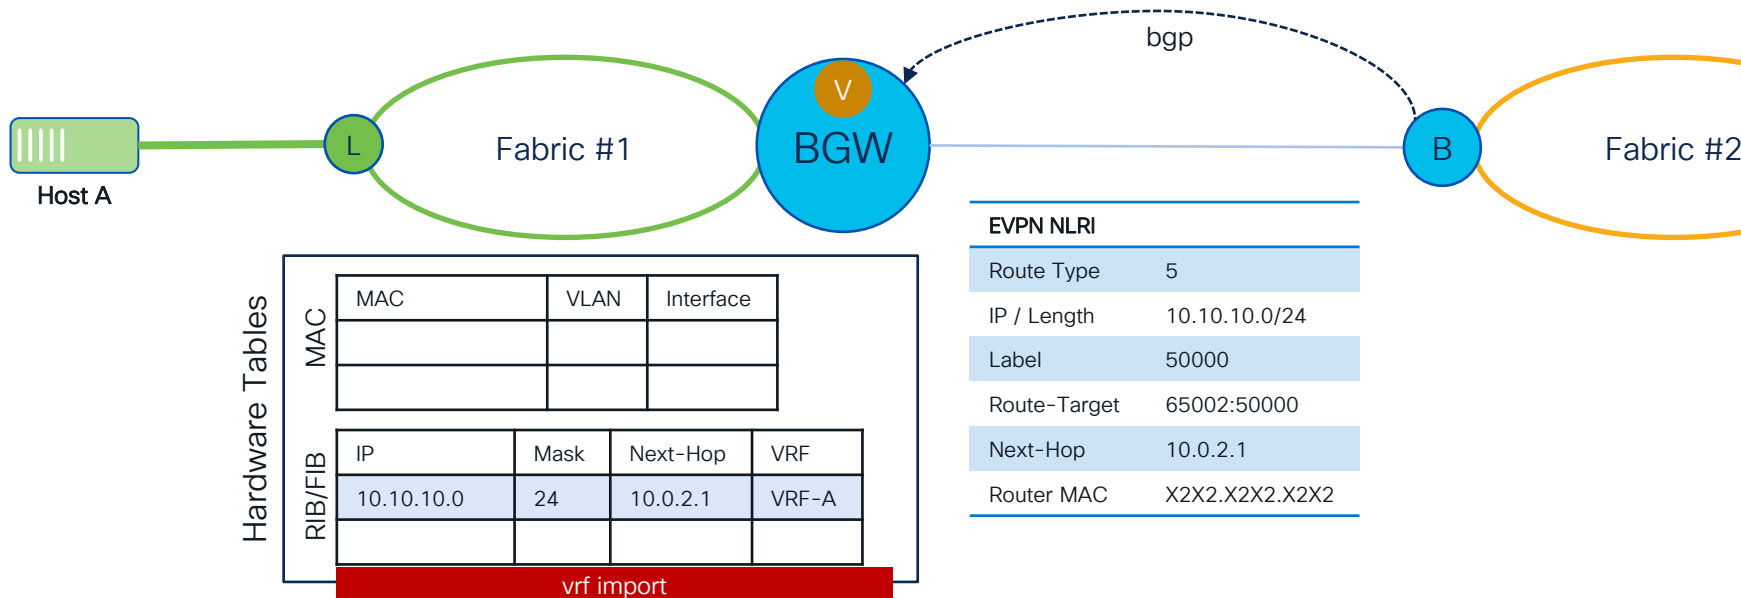
Remote Learning from Fabric External

The Dating Network - When Control- meets Data-Plane



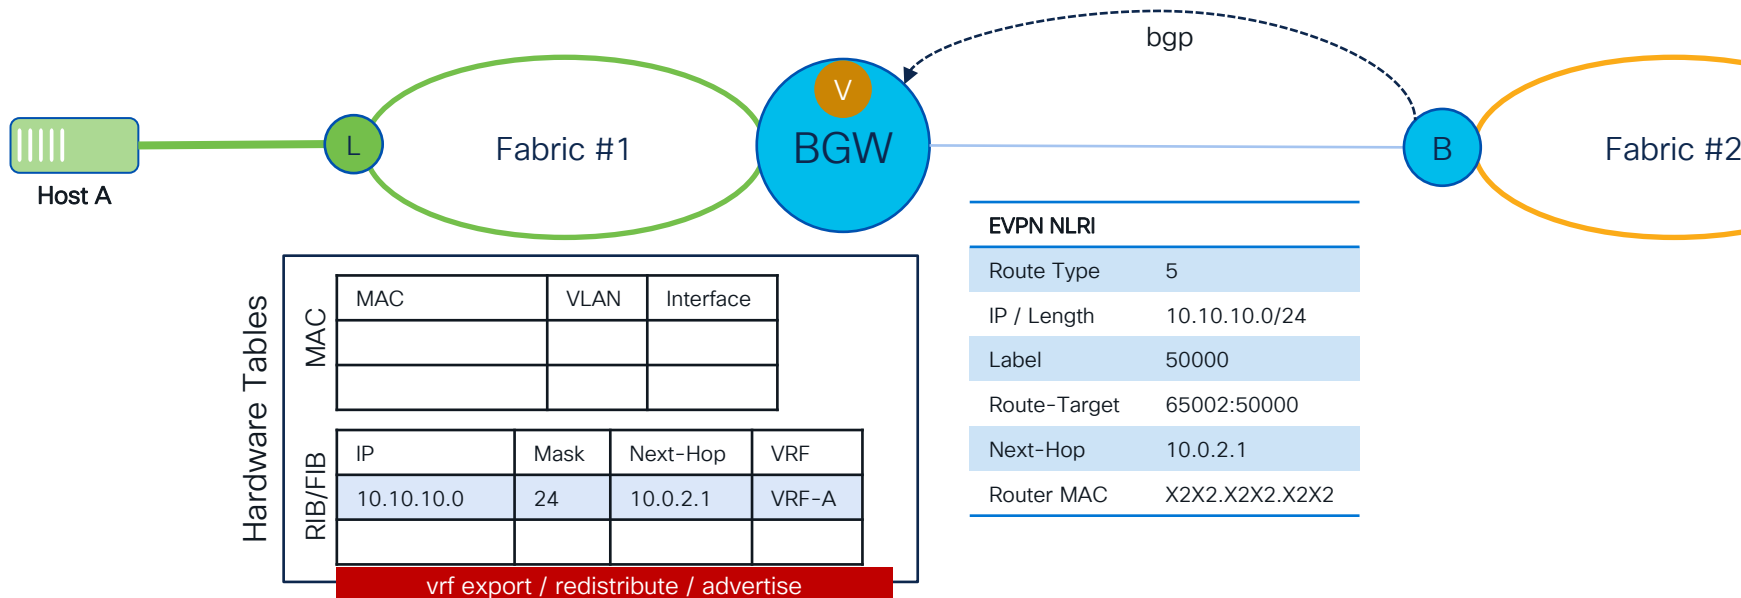
Remote Learning from Fabric External

The Dating Network - When Control- meets Data-Plane



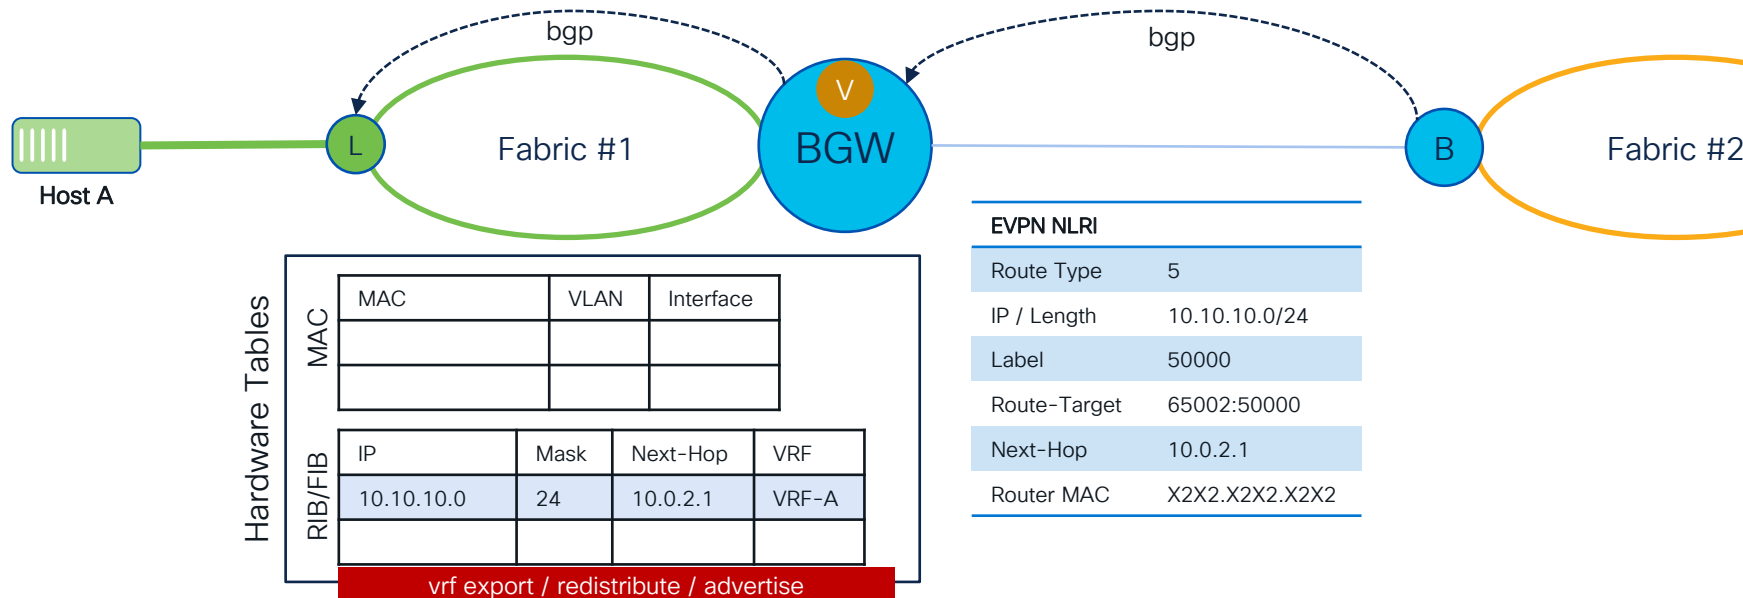
Remote Learning from Fabric External

The Dating Network - When Control- meets Data-Plane



Remote Learning from Fabric External

The Dating Network - When Control- meets Data-Plane

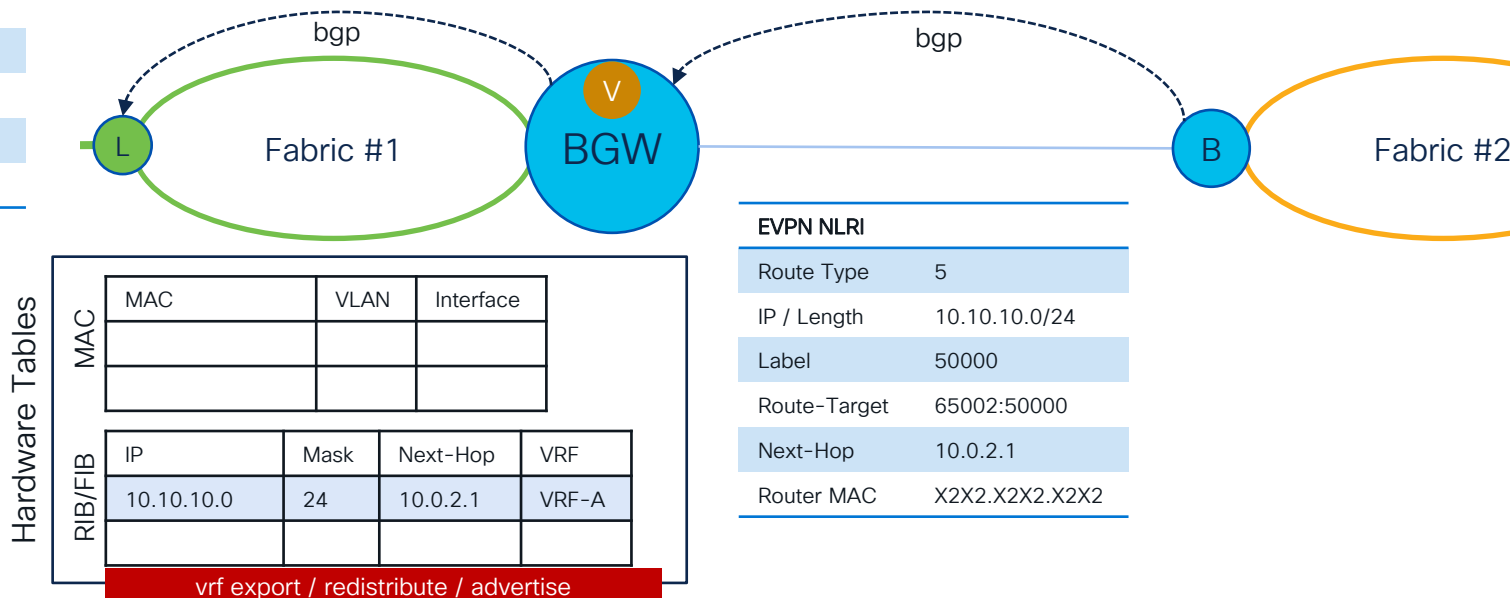


Remote Learning from Fabric External

The Dating Network - When Control- meets Data-Plane

EVPN NLRI

Route Type	5
IP / Length	10.10.10.0/24
Label	50000
Route-Target	65001:50000
Next-Hop	10.0.1.1
Router MAC	X1X1.X1X1.X1X1

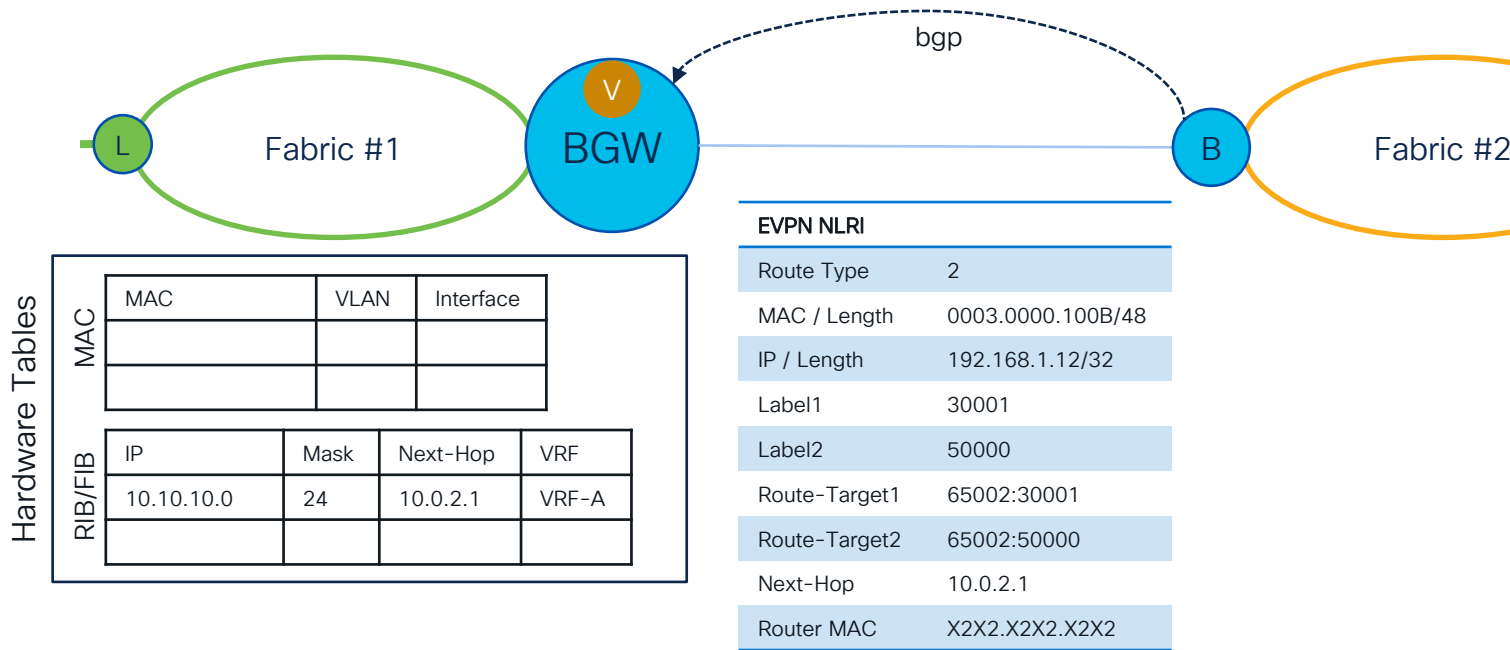


EVPN NLRI

Route Type	5
IP / Length	10.10.10.0/24
Label	50000
Route-Target	65002:50000
Next-Hop	10.0.2.1
Router MAC	X2X2.X2X2.X2X2

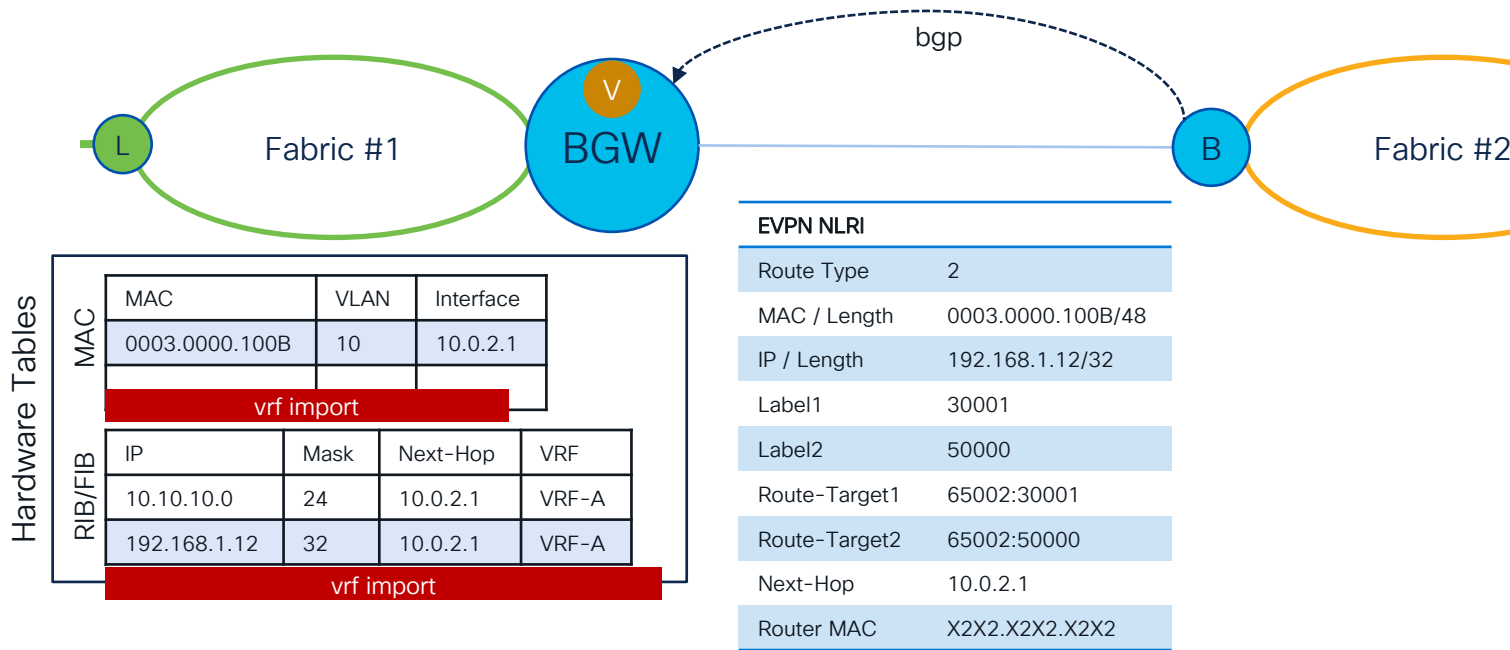
Remote Learning from Fabric External

The Dating Network - When Control- meets Data-Plane



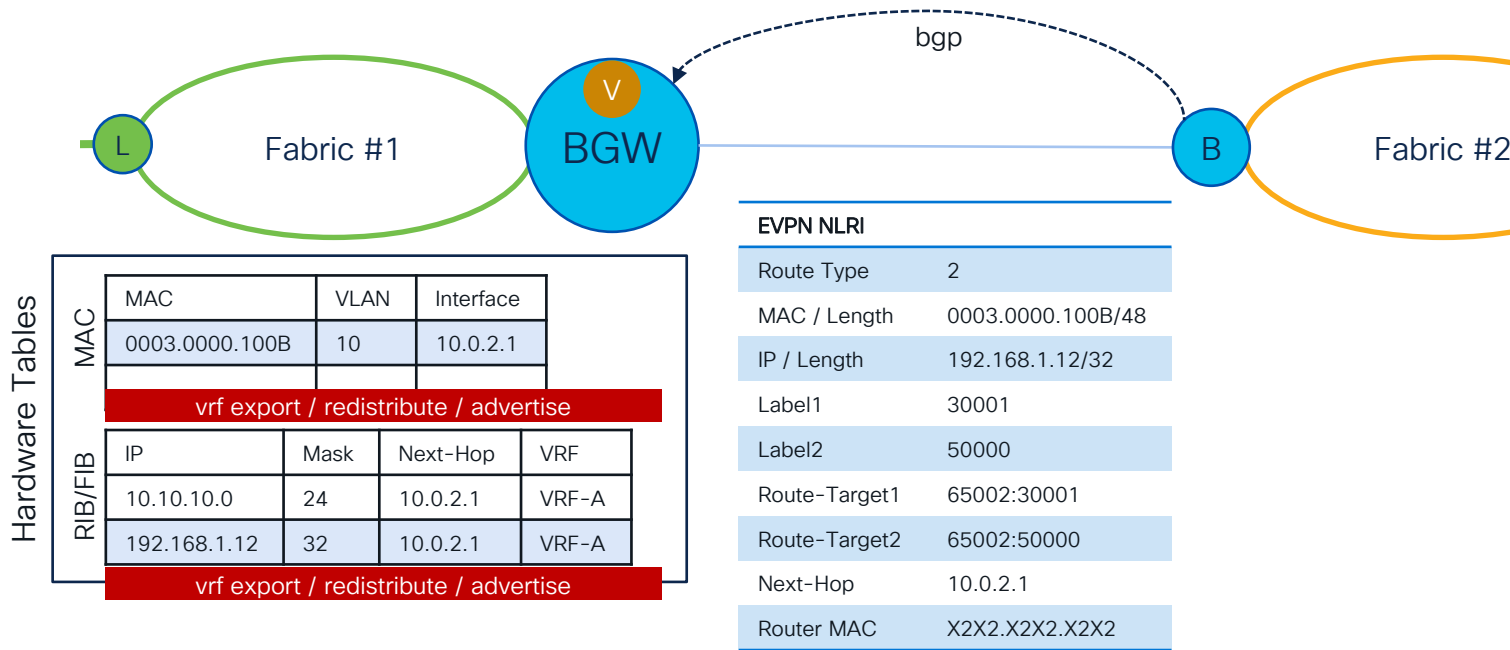
Remote Learning from Fabric External

The Dating Network - When Control- meets Data-Plane



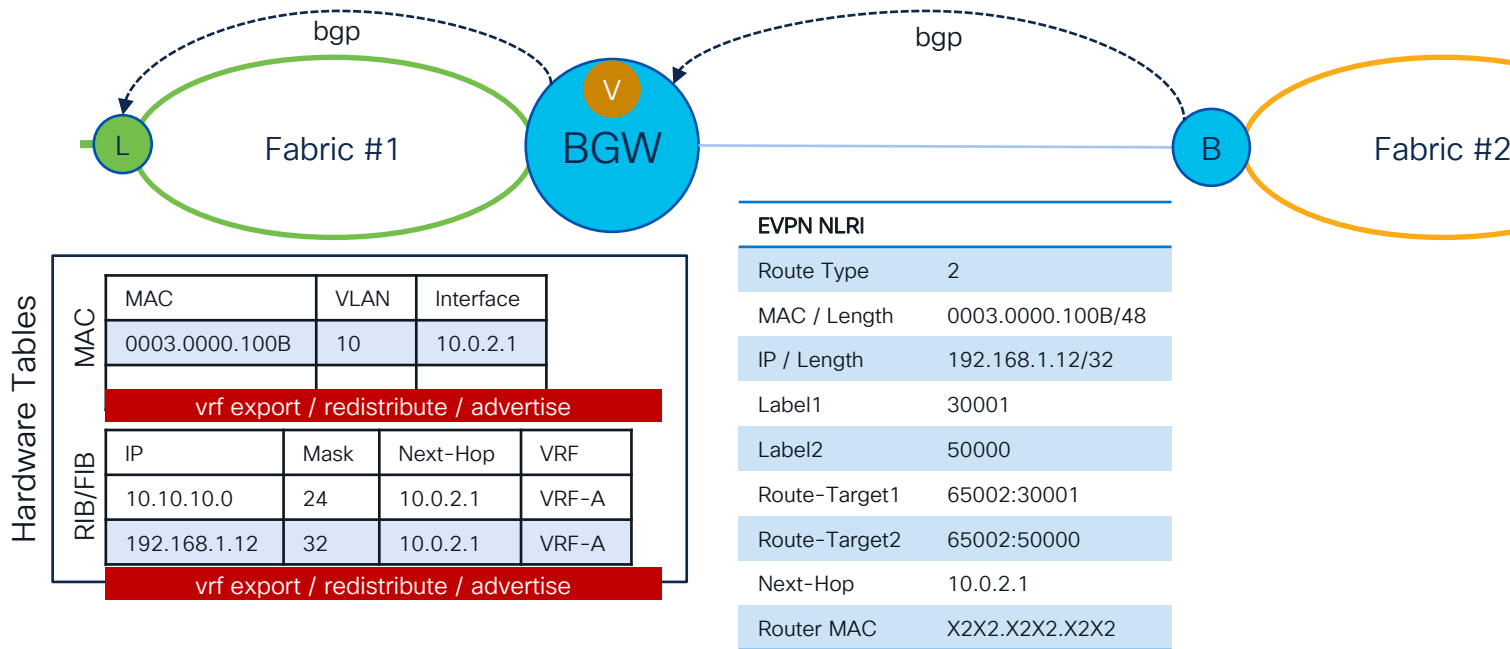
Remote Learning from Fabric External

The Dating Network - When Control- meets Data-Plane



Remote Learning from Fabric External

The Dating Network - When Control- meets Data-Plane



Remote Learning from Fabric External

The Dating Network - When Control- meets Data-Plane

EVPN NLRI	
Route Type	2
MAC / Length	0003.0000.100B/48
IP / Length	192.168.1.12/32
Label1	30001
Label2	50000
Route-Target1	65001:30001
Route-Target2	65001:50000
Next-Hop	10.0.1.1
Router MAC	X1X1.X1X1.X1X1

Hardware Tables

MAC


MAC	VLAN	Interface
0003.0000.100B	10	10.0.2.1

vrf export / redistribute / advertise

RIB/FIB

IP	Mask	Next-Hop	VRF
10.10.10.0	24	10.0.2.1	VRF-A
192.168.1.12	32	10.0.2.1	VRF-A

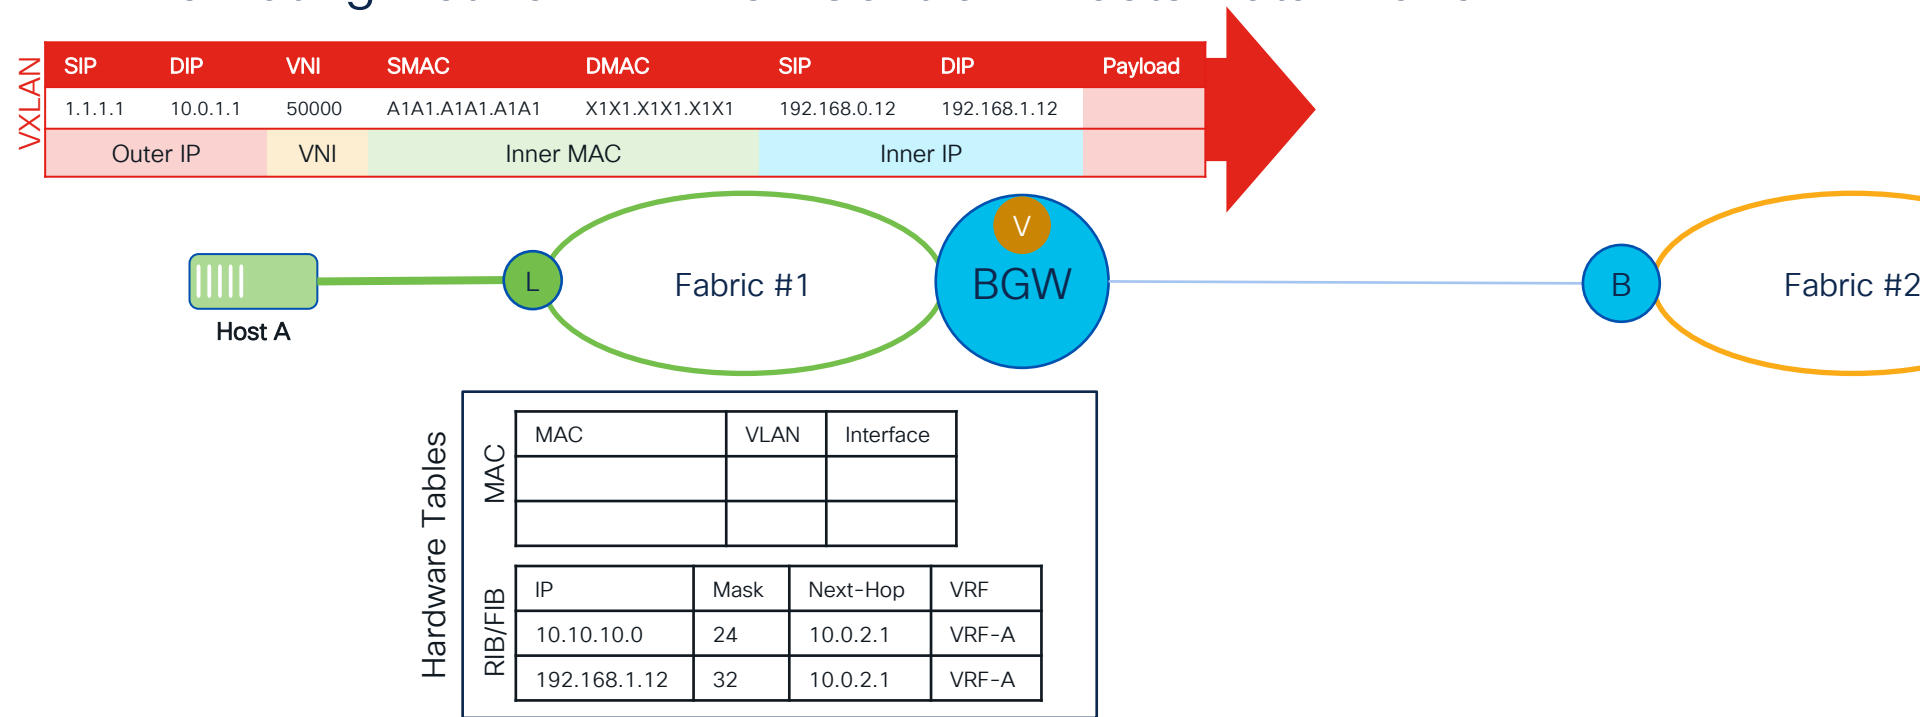
vrf export / redistribute / advertise



EVPN NLRI	
Route Type	2
MAC / Length	0003.0000.100B/48
IP / Length	192.168.1.12/32
Label1	30001
Label2	50000
Route-Target1	65002:30001
Route-Target2	65002:50000
Next-Hop	10.0.2.1
Router MAC	X2X2.X2X2.X2X2

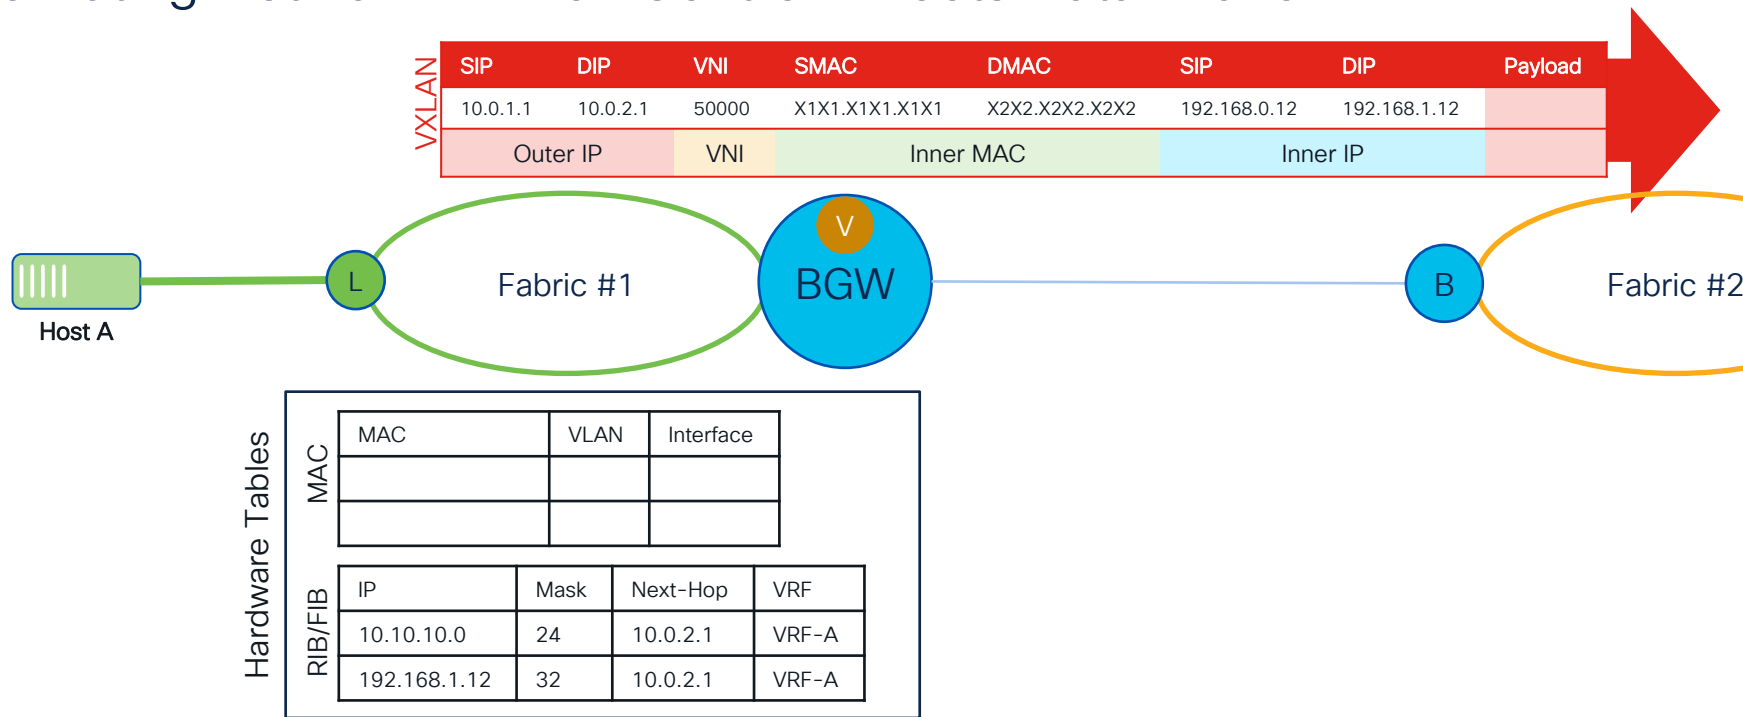
Routing from Local VTEP to Remote BGW

The Dating Network - When Control- meets Data-Plane



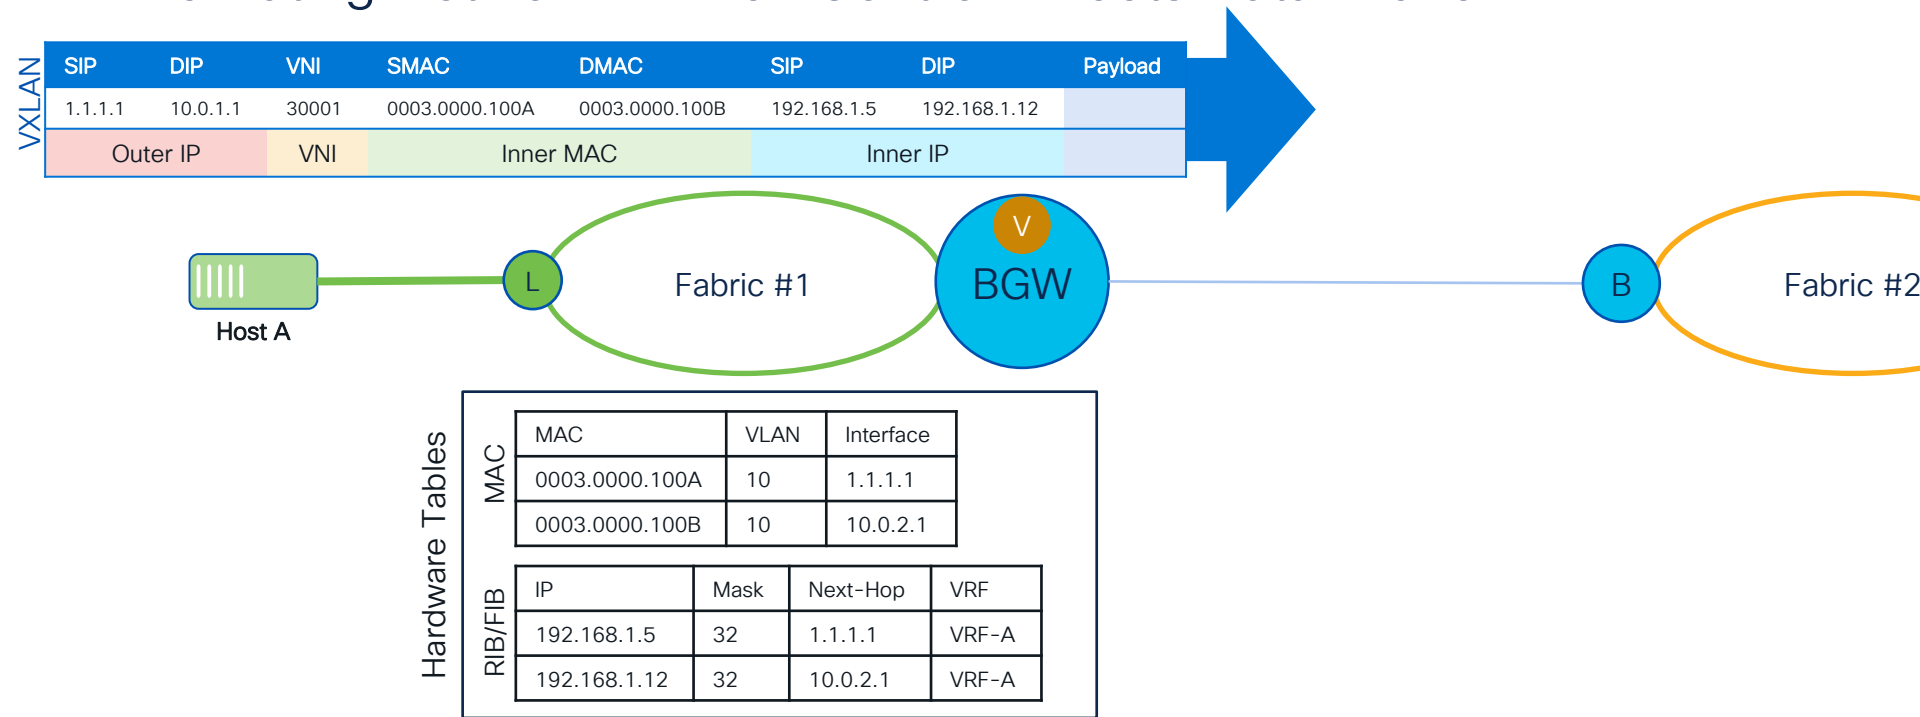
Routing from Local VTEP to Remote BGW

The Dating Network - When Control- meets Data-Plane



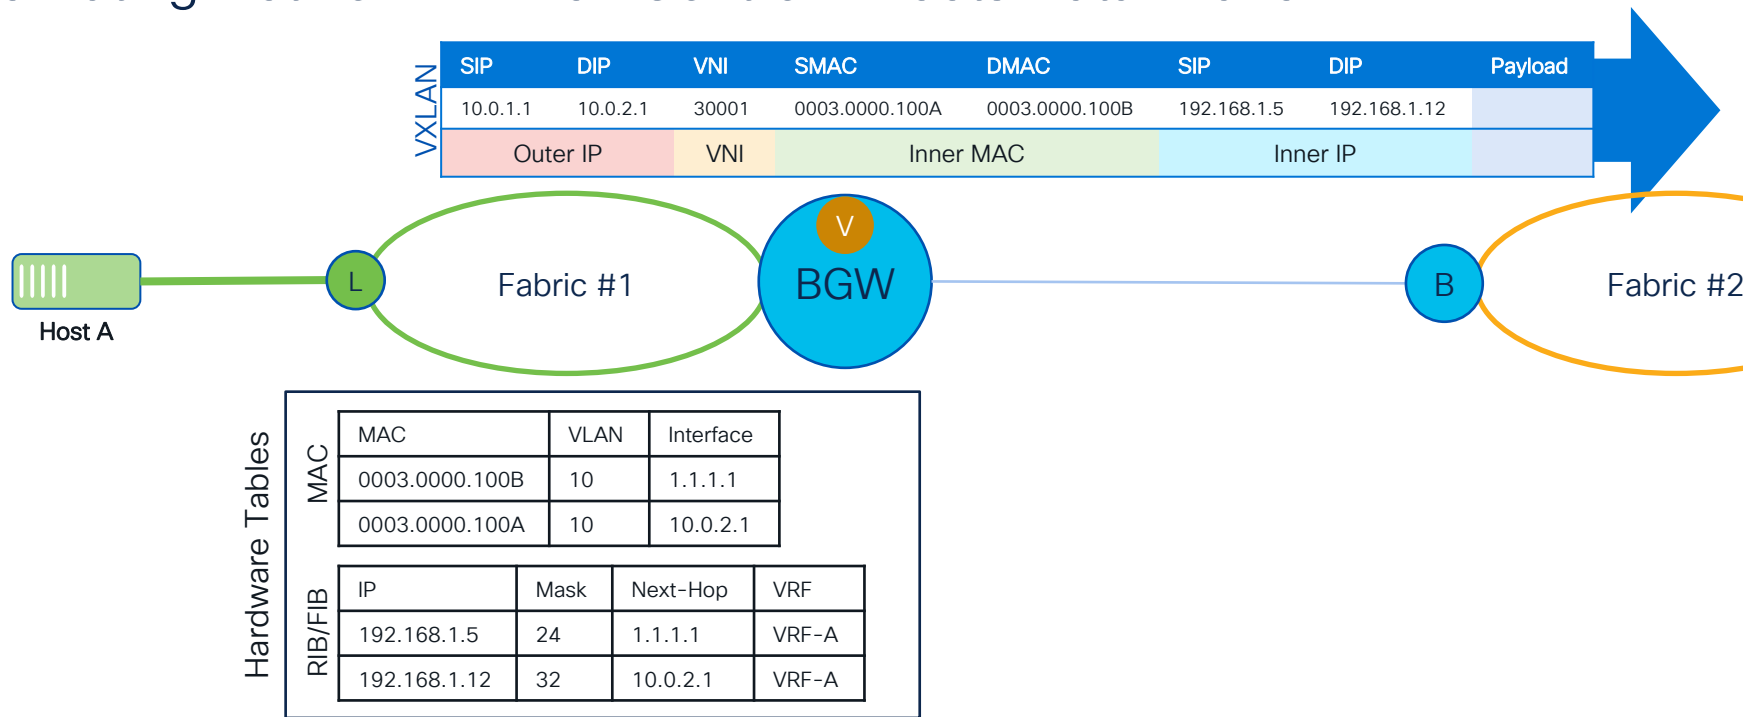
Bridging from Local VTEP to Remote BGW

The Dating Network - When Control- meets Data-Plane



Bridging from Local VTEP to Remote BGW

The Dating Network - When Control- meets Data-Plane



Conclusion

Conclusion

#1

Border Gateway (BGW)

- A Gateway (GW) to stitch multiple VXLAN BGP EVPN domains
- Provides Control- and Data-Plane separation
- Extends Layer-2 and Layer-3 with Control
- Allows to Scale beyond any Fabric Scale
- Facilitates Multi-DC and Multi-Pod Use-Cases
- More than just a Data Center Interconnect (DCI)

#2

VXLAN BGP EVPN Multi-Site

- A Simple add or drop-in
- First introduced in September 2017 – proven and deployed
- A Solution beyond EVPN DCI Overlay (RFC9014)
- Provides Layer-2 and Layer-3 extension
- Wide Hardware Support
- Flexible Deployment Option – Not just for VXLAN Fabrics

Resources – Cisco NX-OS

- VXLAN EVPN Multi-Site Design and Deployment White Paper
 - <https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/white-paper-c11-739942.html>
 - Supported since NX-OS 7.0(3)I7(1) – EX, FX, FX2, FX3, GX, GX2
- NextGen DCI with VXLAN EVPN Multi-Site Using vPC Border Gateways White Paper
 - <https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/whitepaper-c11-742114.html>
 - Supported since NX-OS 9,.2(1) – EX, FX, FX2, FX3, GX, GX2

Resources – IETF

- draft-sharma-bess-multi-site-evpn – Multi-Site Solution for Ethernet VPN (EVPN) Overlay (Layer-2 and Layer-3)
 - <https://datatracker.ietf.org/doc/html/draft-sharma-bess-multi-site-evpn>
- RFC 9014 – Interconnect Solution for Ethernet VPN (EVPN) Overlay Networks (Layer-2 only)
 - <https://datatracker.ietf.org/doc/html/rfc9014>
- draft-ietf-bess-evpn-ipvpn-interworking – EVPN Interworking with IPVPN (Layer-3 only)
 - <https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-ipvpn-interworking>

Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.



Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. www.cisco.com/go/certs

Pay for Learning with Cisco Learning Credits

(CLCs) are prepaid training vouchers redeemed directly with Cisco.



Learn

Cisco U.

IT learning hub that guides teams and learners toward their goals

Cisco Digital Learning

Subscription-based product, technology, and certification training

Cisco Modeling Labs

Network simulation platform for design, testing, and troubleshooting

Cisco Learning Network

Resource community portal for certifications and learning



Train

Cisco Training Bootcamps

Intensive team & individual automation and technology training programs

Cisco Learning Partner Program

Authorized training partners supporting Cisco technology and career certifications

Cisco Instructor-led and Virtual Instructor-led training

Accelerated curriculum of product, technology, and certification courses



Certify

Cisco Certifications and Specialist Certifications

Award-winning certification program empowers students and IT Professionals to advance their technical careers

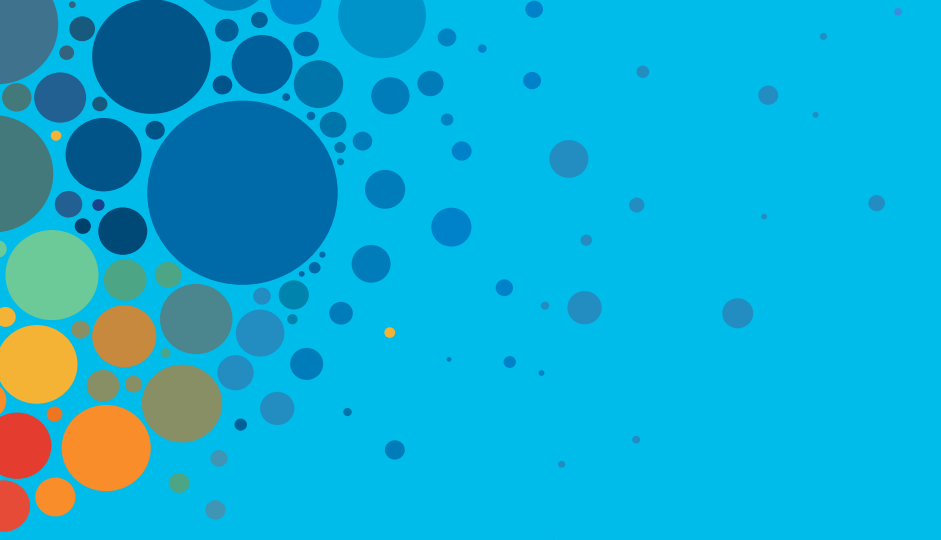
Cisco Guided Study Groups

180-day certification prep program with learning and support

Cisco Continuing Education Program

Recertification training options for Cisco certified individuals

Here at the event? Visit us at **The Learning and Certifications lounge at the World of Solutions**



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



The bridge to possible

Thank you

CISCO *Live!*



#CiscoLive